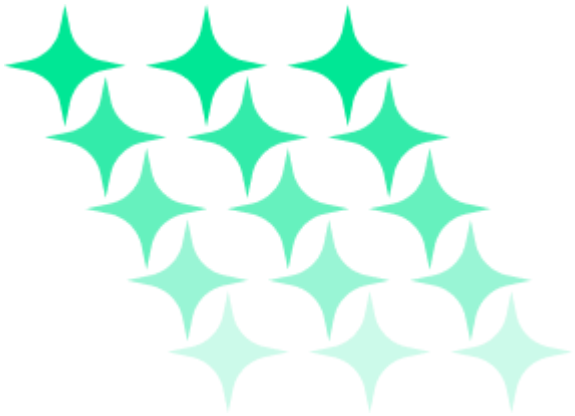
GPT for Work Documentation

User Guide for Microsoft Excel & Word

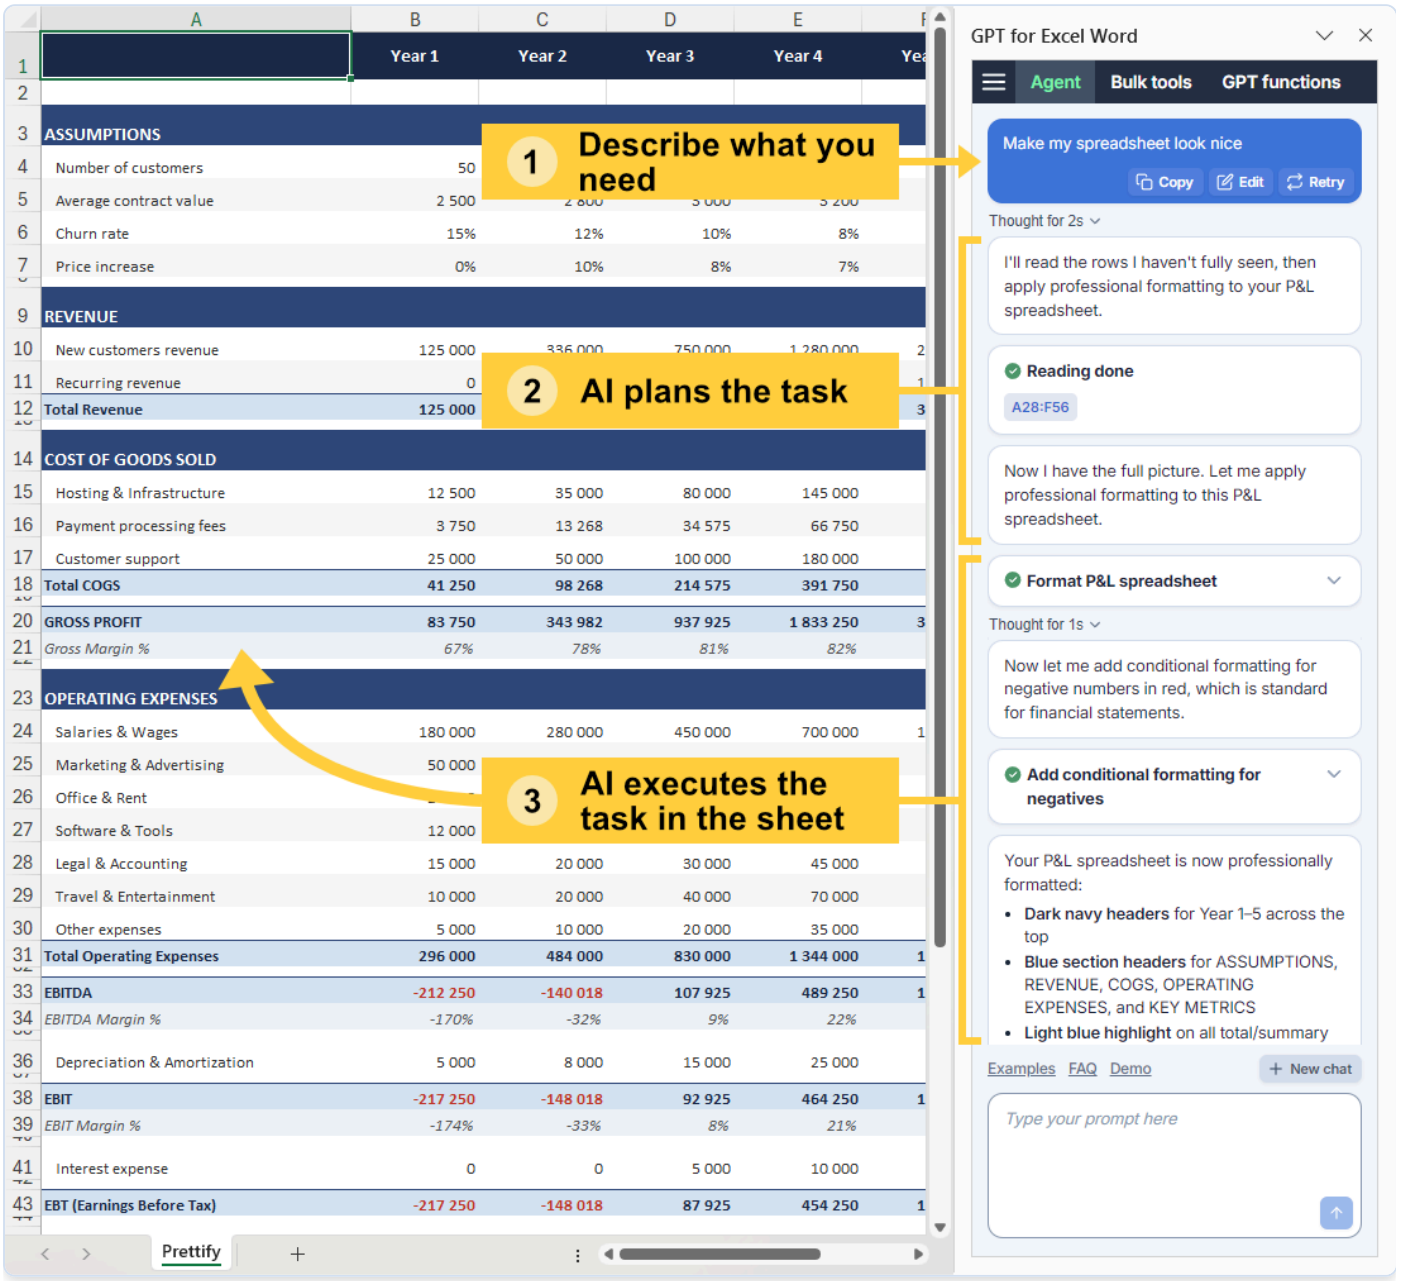
2026-10-19



GPT for Excel user guide

Install for Microsoft

GPT for Excel is a Microsoft Excel add-in that integrates an **AI agent** into the Excel user interface. Use plain-language instructions to automate everyday spreadsheet tasks — such as writing formulas, fixing errors, formatting cells, or creating charts and pivot tables — or process large datasets row by row with operations such as translation, categorization, enrichment, and scoring.



Use cases

Here are common use cases that you can easily handle with GPT for Excel.

Everyday spreadsheet tasks

- **Write** spreadsheet formulas like a pro.
- **Format** cells, ranges, or entire sheets to match your preferred style or conventions.
- **Standardize** inconsistent date formats, trailing whitespace, or mixed capitalization across your sheet.
- **Create** charts and pivot tables from your existing data to visualize trends and summarize results.
- **Find and fix** errors in formulas, references, or inconsistent data.

Row-by-row bulk processing

- **Generate** SEO-optimized product descriptions from product specifications.
- **Translate** product descriptions to multiple target languages while observing appropriate localization conventions.
- **Enrich** company or contact records with data from AI-powered web lookups.
- **Categorize** support tickets, survey responses, or product feedback into groups that match your taxonomy.
- **Clean up** customer relationship management (CRM) data by reformatting fields, recategorizing records, and removing duplicates.
- **Score and qualify** sales leads in a consolidated prospecting spreadsheet.
- **Analyze** customer sentiment by rating reviews and support requests based on criteria relevant to you.

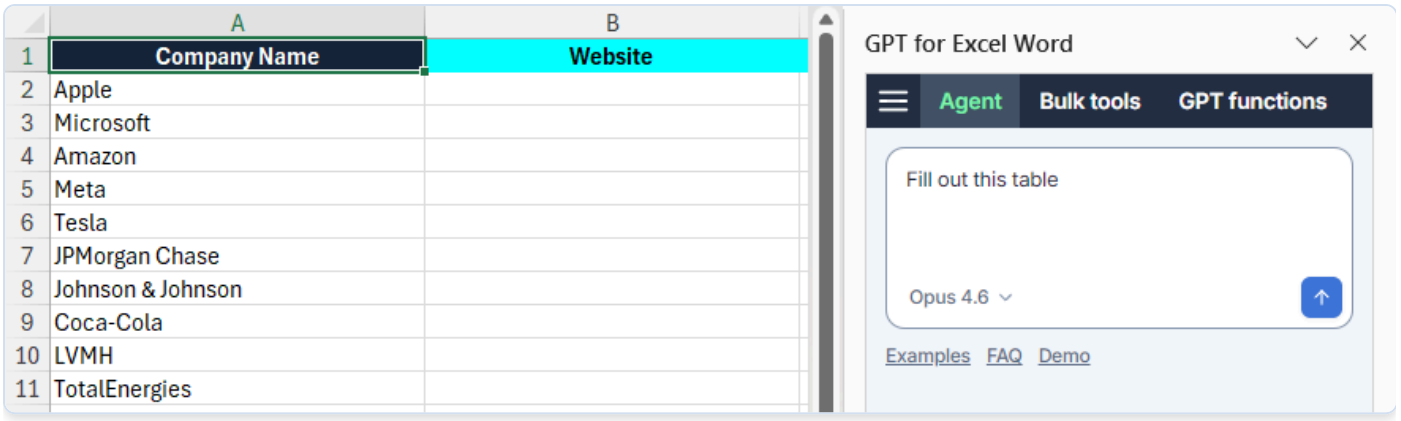
Features

GPT for Excel provides the following main features for using AI in spreadsheets:

- Agent
- Bulk tools
- GPT functions

Agent

The **Agent** is the chat interface for spreadsheet work in GPT for Excel. Describe what you need in plain language — from everyday spreadsheet tasks to row-by-row bulk processing across large datasets — and the Agent reads your data, determines the required steps, and executes them directly in your spreadsheet.

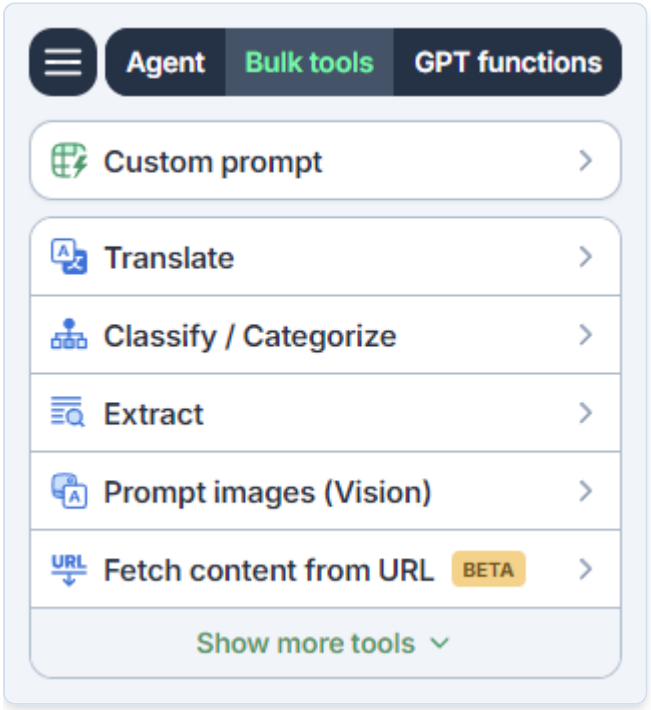


The following video shows you how to use the Agent to fill in missing company information in a market research spreadsheet. Based on a simple prompt, the Agent fetches the missing information from the web and fills the empty columns.

Bulk tools

Bulk tools allow you to run prompts on specific spreadsheet columns at once without writing any formulas. You configure and run bulk tools from the add-in sidebar. You can also easily reuse past prompts and configurations with the **History** feature.

GPT for Excel provides dedicated bulk tools for selected use cases, such as translation, classification/categorization, and extraction, as well as a generic bulk tool for fully custom prompts.

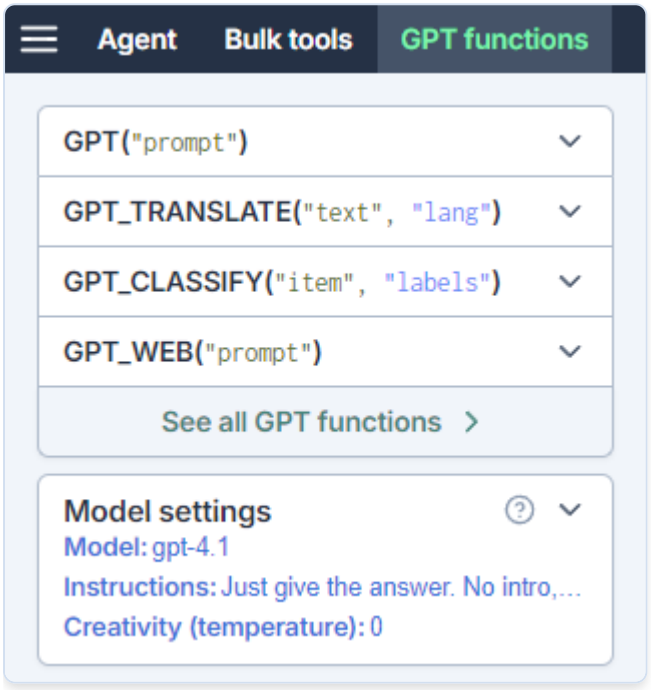


The following video shows you how to generate some text with the **Custom prompt** bulk tool. You simply tell the tool what you want to generate and select the column where you want to place the generated text. You do not need to define any other settings if the default setup meets your needs.

GPT functions

GPT functions are custom spreadsheet functions that allow you to prompt AI from inside spreadsheet cells. GPT functions work exactly like native functions in that you can use them on their own or combine them with other functions when creating formulas.

GPT for Excel provides dedicated GPT functions for selected use cases, such as translation, classification, and web search, as well as a generic GPT function for all other use cases. You can define function-specific parameters to further refine how the AI generates responses. You can find a list of all available functions with usage examples in the add-in sidebar.



The following video shows you how to generate some text in bulk with the **GPT** function.

Which feature to use when

The following table shows how easy it is to use the different features and describes the overall experience of using each. As a rule, start with the Agent.

Agent	Bulk tools	GPT functions
Easy	Moderate	Advanced

Agent	Bulk tools	GPT functions
• Tell the AI what you want in plain language. • Automate spreadsheet tasks or workflows across multiple columns. • Let the AI automatically figure out the steps.	• Process a column row by row. • Precisely control the prompt, model, and settings. • Easily reuse prompts and configurations.	• Use AI directly inside spreadsheet formulas. • Drag formulas across rows or columns. • Get results that update automatically when inputs change.

Benefits

- **Prompt AIs from inside your spreadsheets.** No more copy-pasting between Excel and AI chatbots like ChatGPT.
- **Chat with your spreadsheets in plain language.** Simply tell the Agent in your own words what you want to do, and it will automatically plan and execute the task. No prompt engineering required.
- **Run prompts at scale and at lightning speed.** Execute up to 1,000 prompts per minute, reliably process up to 1,000,000 rows in one go, and generate up to 10,000 results per hour.
- **Use the best AI model for your bulk tasks.** Choose from a large selection of models from OpenAI, Anthropic, Google, xAI, and Perplexity.
- **Bring an API key or set up your own API endpoint to get more features.** You do not need your own API keys or endpoints to use Excel, but they give you access to more models and afford you more control and privacy.

What's next

- Install GPT for Excel if you have not already.
- Select an AI model that best meets your needs.
- Get going with the Agent.
- Get going with bulk tools.
- Get going with GPT functions.

Get started with GPT for Excel

GPT for Excel is an **AI agent** for spreadsheet work in **Microsoft Excel**. It integrates directly into your spreadsheet and lets you automate spreadsheet tasks through plain-language instructions.

This guide gives you a quick walkthrough of the add-in from installation to trying out the Agent.

Install the add-in

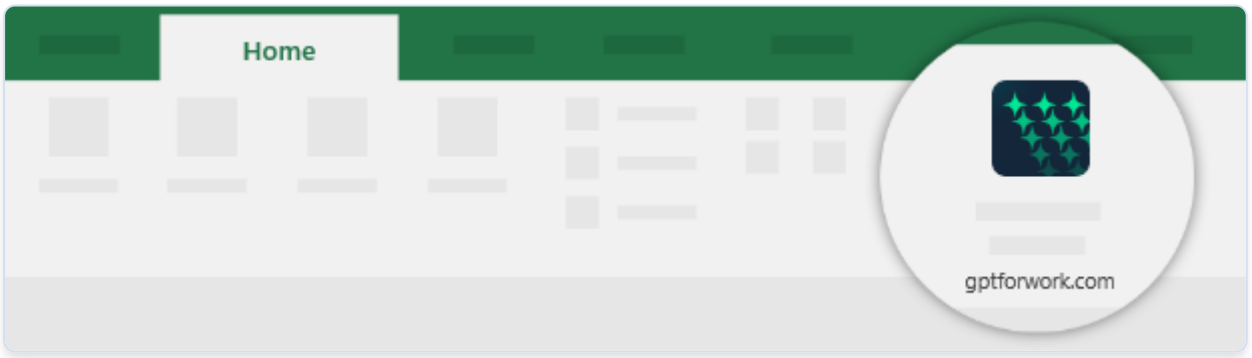
To use GPT for Excel, install the add-in:

Install for Microsoft

For detailed installation instructions, see Install GPT for Excel.

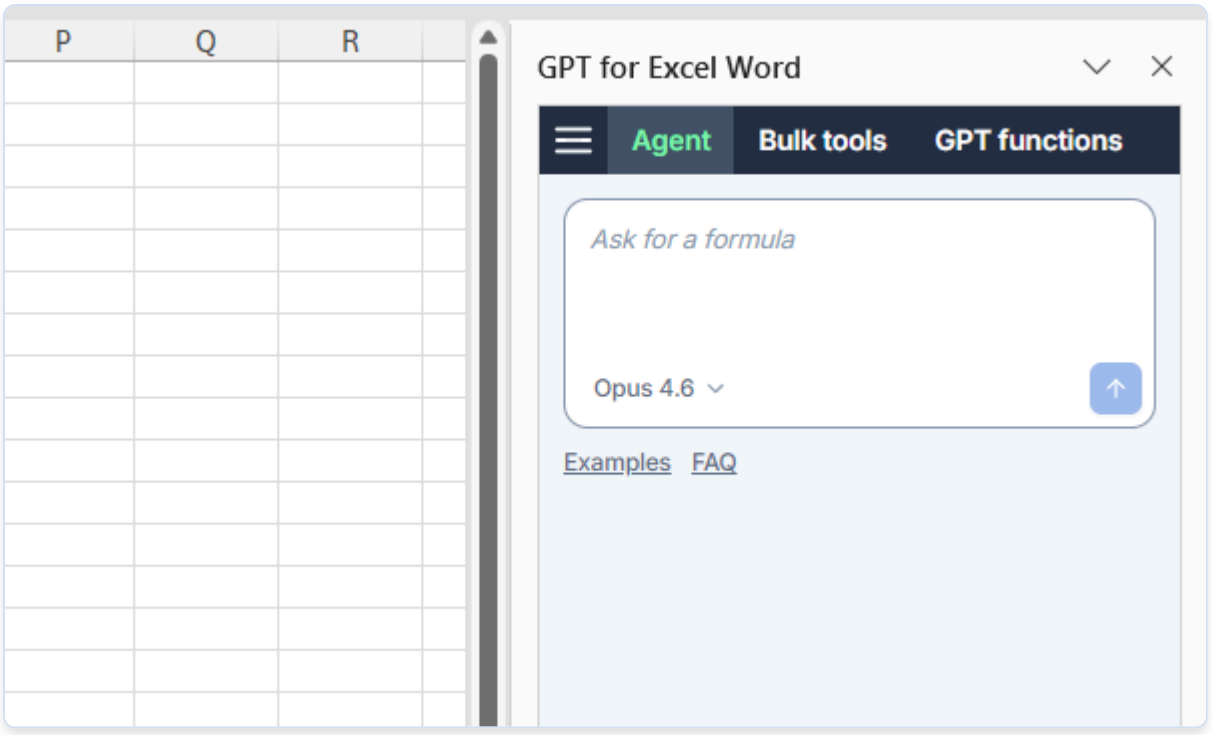
Open the add-in

1. Create a new blank workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

The GPT for Excel sidebar opens.



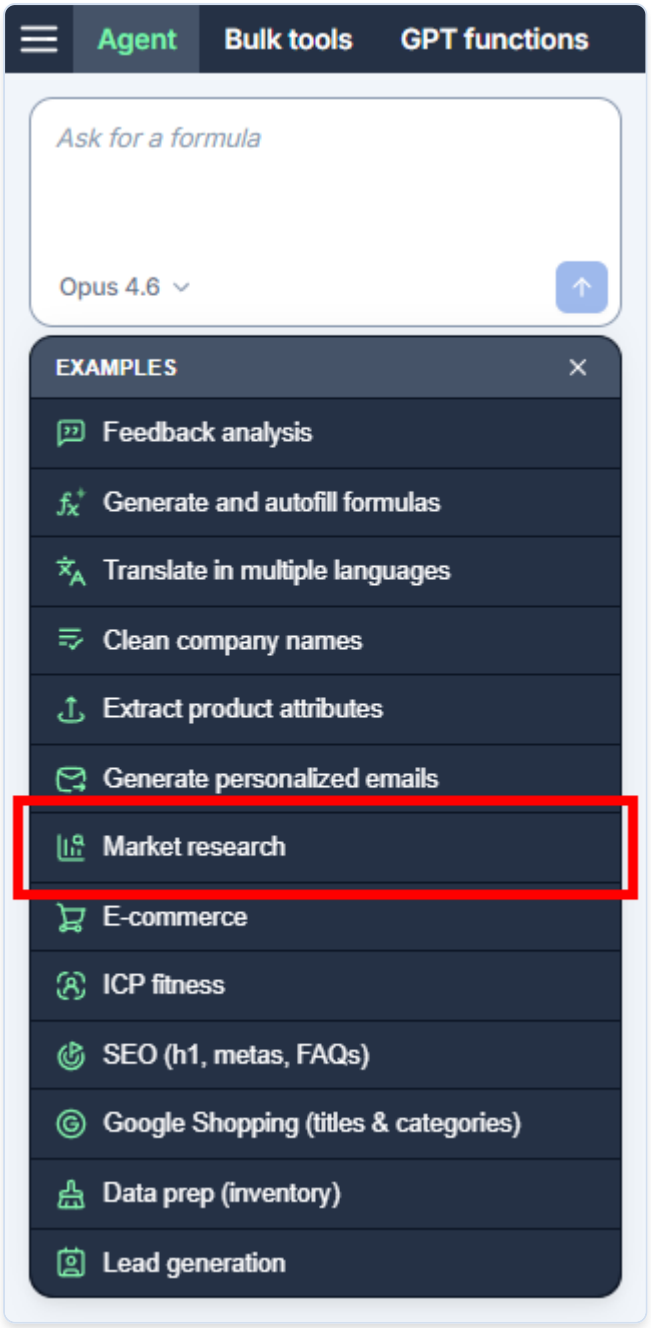
Use the Agent

The **Agent** is the chat interface for spreadsheet work in GPT for Excel. Describe what you need in plain language — from everyday spreadsheet tasks to row-by-row bulk processing across large datasets — and the Agent reads your data, determines the required steps, and executes them directly in your spreadsheet.

Let's have the Agent do some quick market research for us. We'll use one of the built-in examples for convenience.

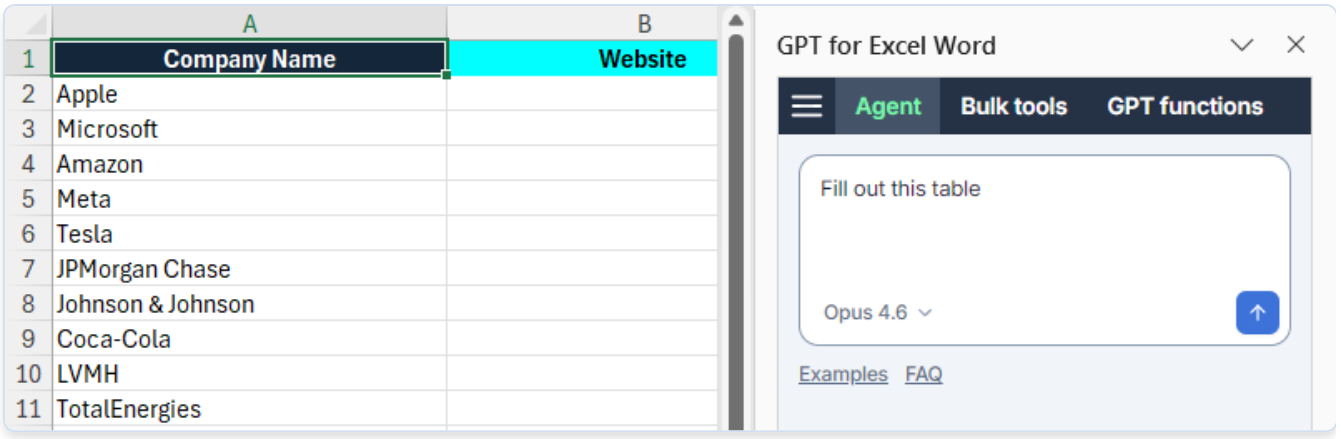
1. In the sidebar, select **Agent**.

2. Click **Examples** and select **Market research**.



GPT for Excel:

1. Creates a new sheet prepopulated with example data. The data is a list of company names and a series of empty columns for company information that we want the Agent to fill in.
2. Prefills the Agent chat field with a prompt that will accomplish the example task: `Fill out this table`. This simple prompt is enough for the Agent to fetch the missing company information from the web and fill the empty columns.



3. Click the send button.

   The Agent:

   1. Analyzes the prompt and drafts a plan for accomplishing the requested task.
   2. Executes the plan. In this example, the Agent fetches the missing company information from the web and fills the empty columns B-E with what it found, one by one. Note that for column F, which does not require web search, the Agent uses gpt-4.1 (no web search).

   > **TIP**
   >
   > You can follow the Agent's thinking and workflow live in the sidebar.

The following video shows the entire process.

What's next

- Learn more about GPT for Excel.
- Learn more about the Agent.
- Learn about the other features available in GPT for Excel:
  - Bulk tools
  - GPT functions
- Try out different AI models.
- Try out example Agent use cases for different business applications and product capabilities.

Set up GPT for Excel

Install GPT for Excel, open the add-in, and sign in. If needed, configure API keys and endpoints to access more models. If you no longer use GPT for Excel, remove its access to your Microsoft or Google account or uninstall the add-in.

Install

Install the GPT for Excel Word add-in to use AI in your Excel workbooks.

Open

Open the GPT for Excel Word add-in in Microsoft Excel.

Sign in

Use a Microsoft or Google account to sign in to GPT for Excel, switch between accounts, and sign out.

Manage models

Connect AI models to Microsoft Excel with GPT for Excel.

Remove account access

Remove Microsoft or Google account access from the GPT for Excel Word add-in.

Uninstall

Remove the GPT for Excel Word add-in from your Microsoft account.

Install GPT for Excel

GPT for Excel is part of the **GPT for Excel Word** add-in. One installation covers Microsoft Excel and Microsoft Word, on both desktop and the web; this guide focuses on installing and using the add-in for Microsoft Excel.

Prerequisites

The following table lists the **minimum versions of Microsoft Excel** compatible with the GPT for Excel Word add-in on different Microsoft Office platforms.

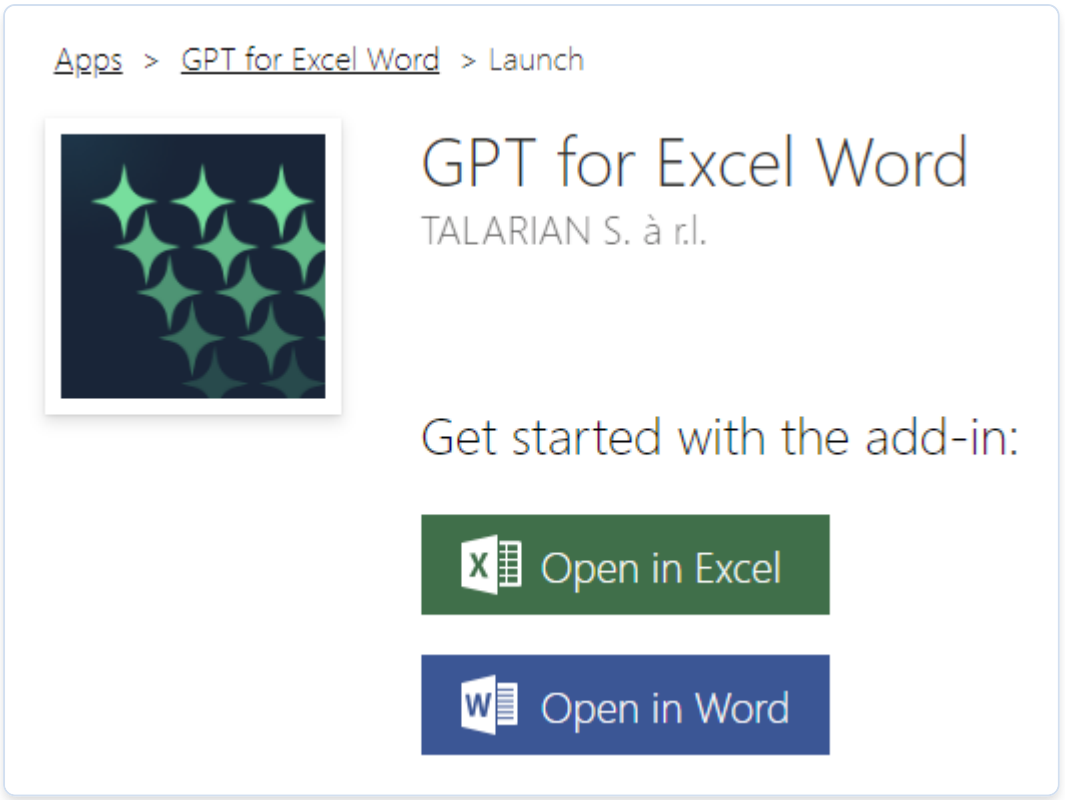
	Windows	macOS	Web
	**Microsoft 365 subscription:** Version 2002 (Build 12527.20092) **Retail perpetual:** Version 2002 (Build 12527.20092) **Volume-licensed perpetual:** Excel 2021: Version 2108 (Build 12527.20092)	Version 16.35 (Build 20030802)	Supported on: • Google Chrome (latest version) • Microsoft Edge (latest version)

Which version of Excel am I using?

Installation

1. Go to the installation page.
2. Click **Open in Excel**.

> **INFO**
> The add-in will be installed on your Microsoft account for both Excel and Word.

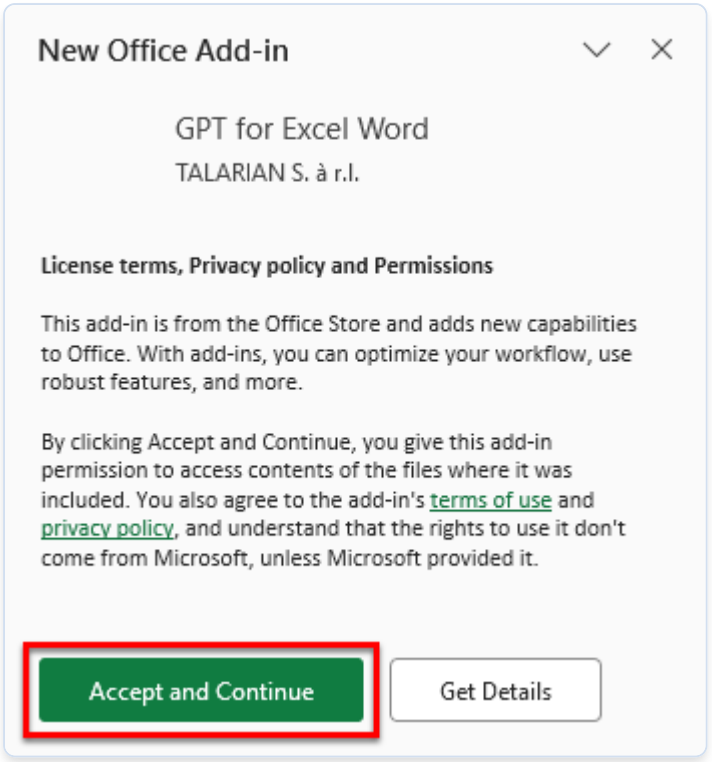


A pop-up appears at the top of the window, requesting your permission to proceed.

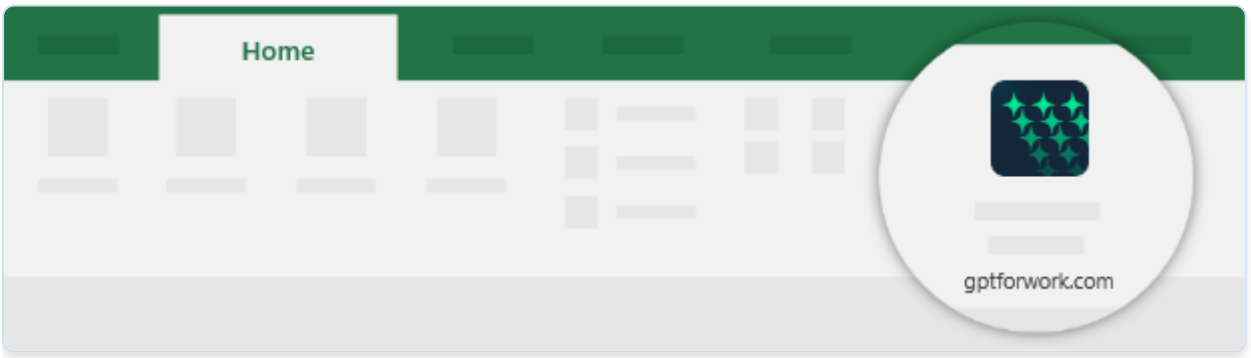
3. Confirm to open the Excel desktop app.

> **INFO**
>
> If this method doesn't work, install the add-in from the app:
>
> 1. Open Excel.
> 2. On the ribbon, select **Home > Add-ins**.
> 3. Search for "GPT for Excel Word".
> 4. In the search results, click **Add** for **GPT for Excel Word**.

4 Click **Accept and Continue** to accept the permissions.



You have installed GPT for Excel Word on your Microsoft account. The add-in is now available from the **Home** tab in the desktop and web versions of Excel. If you don't see it, go to **Home > Add-ins > My add-ins > GPT for Excel Word**.



FOR IT ADMINISTRATORS

If your organization uses a firewall or proxy that restricts network access, you must whitelist the following domains for the add-in to work:

```
*.autosheet.com
*.gptforwork.com
*.microsoft.com
*.microsoftonline.com
*.live.com
*.office.com
```

Having trouble installing?

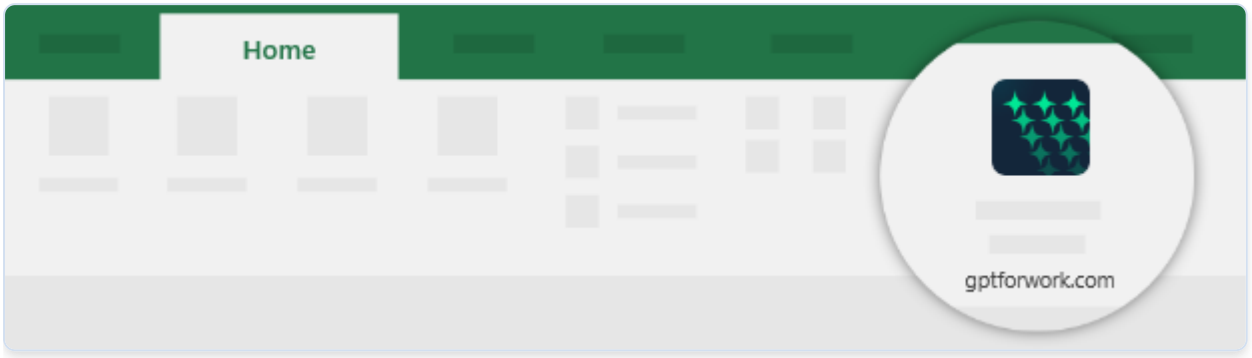
If you run into any issues while installing the add-in, see the troubleshooting guide.

What's next

- Open GPT for Excel.
- Learn about GPT for Excel.
- Select a model:
  - Select the models for the Agent.
  - Select a model for bulk tools.
  - Select a model for GPT functions.
- Get started with the Agent, bulk tools, and GPT functions.

Open GPT for Excel

On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

What's next

- Sign in with your Microsoft or Google account.
- If you run into issues, see the troubleshooting guide.

Sign in to GPT for Excel

This section explains how to sign in to GPT for Excel with your Microsoft or Google account, how to check which account you're using, and how to switch between accounts or sign out.

Sign in with Microsoft or Google

Sign in to GPT for Excel with your Microsoft or Google account.

Check which account you're using

How to check which account GPT for Excel is using.

Switch to another account

Use a different Microsoft or Google account with GPT for Excel.

Sign out

How to sign out of GPT for Excel.

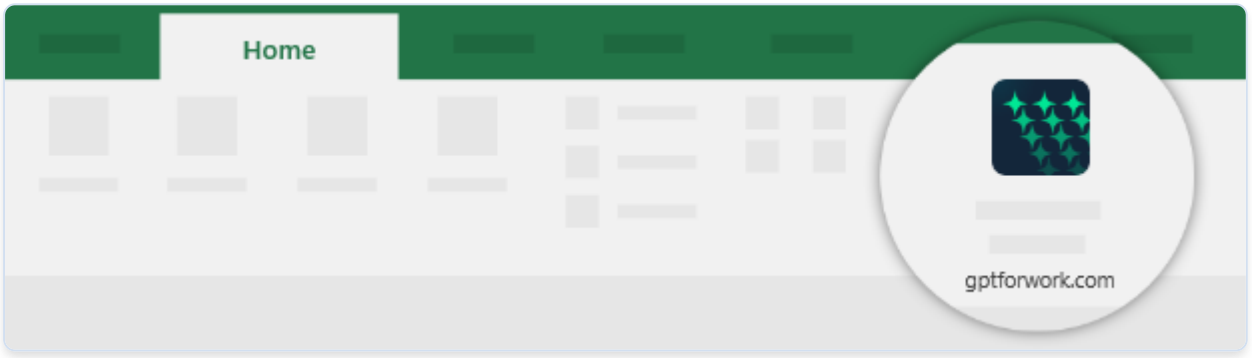
Sign in with Microsoft or Google

You can use GPT for Excel with a Microsoft account or a Google account.

If you have multiple accounts, such as a personal account and a work account, you can check which account you're using and switch to another account to ensure you are using the correct account in GPT for Excel.

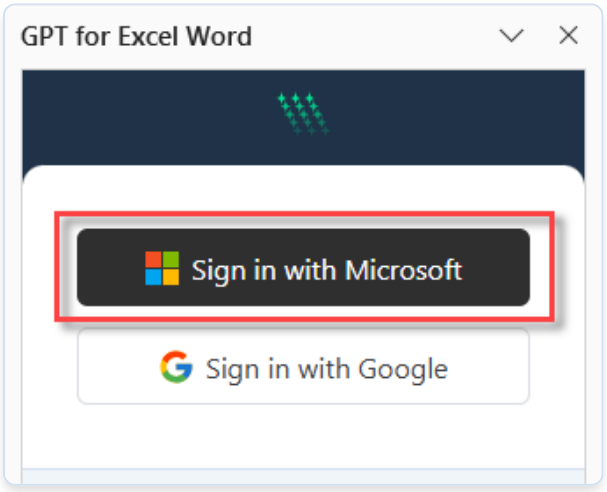
Microsoft

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. On the sign-in screen in the sidebar, select **Sign in with Microsoft**.



4. In the dialog that opens, select your Microsoft account and complete the sign-in process.

You are signed in to GPT for Excel. The add-in keeps you signed in the next time you open it, until you sign out.

> **INFO**
>
> Signing in to GPT for Excel also signs you in to GPT for Word, since they use the same account session.

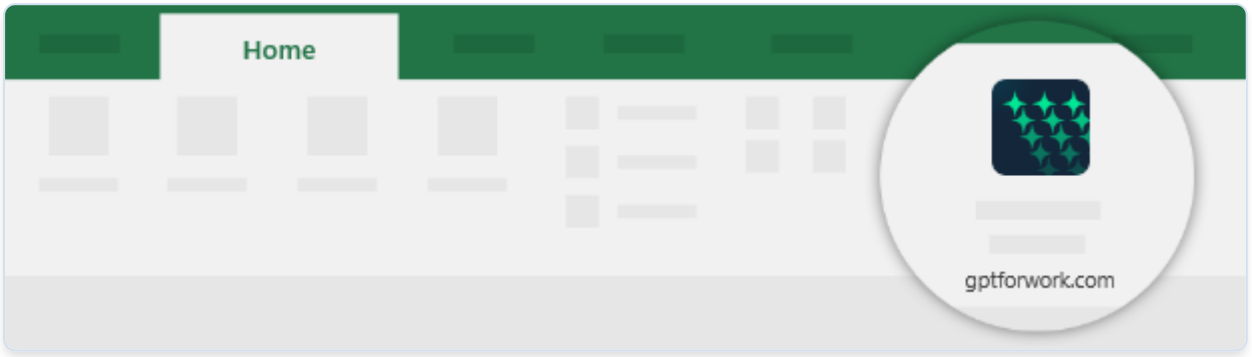
What's next

- Check which account is used by GPT for Excel.
- If needed, switch to another account.

Check which account you're using

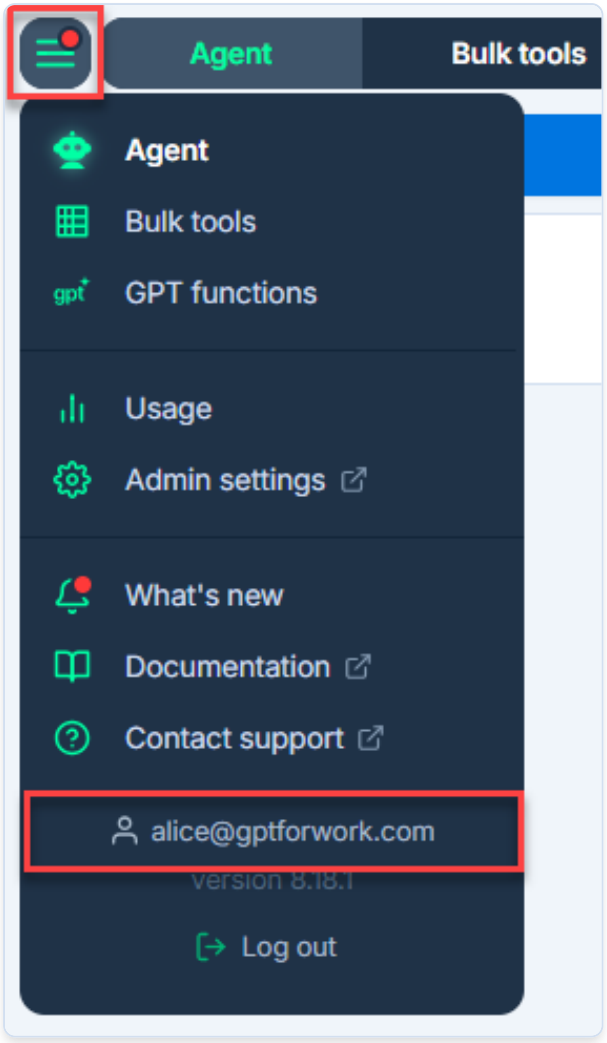
To check which account you are using in GPT for Excel:

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3 In the sidebar, open the main menu. Your account is displayed at the bottom of the menu.



What's next

- If needed, switch to another account.

Switch to another account

GPT for Excel authenticates independently of Microsoft Excel: switching accounts in Excel does not change the account used by the add-in.

If the account currently used by GPT for Excel is not the one you want to use, switch to another account:

1. Sign out of GPT for Excel.
2. Sign in again with the Microsoft or Google account you want to use.

> **INFO**
>
> Switching accounts in GPT for Excel also switches the account used by GPT for Word, since they use the same account session.

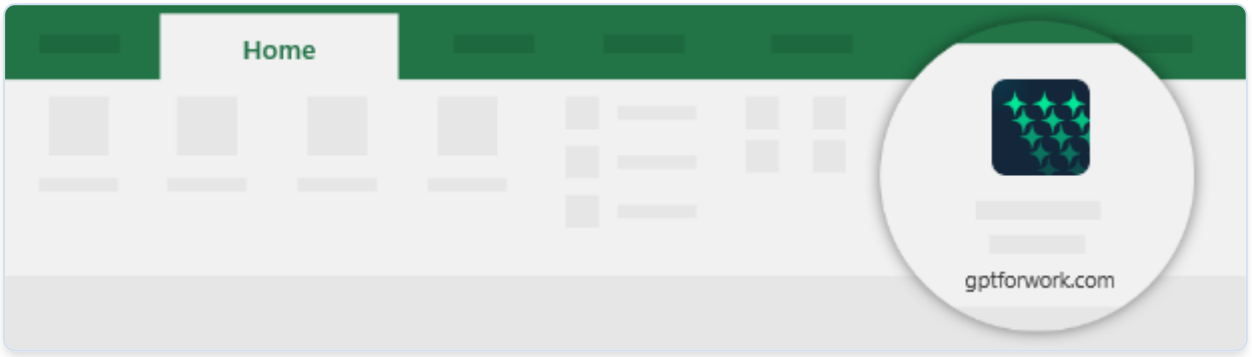
You have switched to another account. You can now use GPT for Excel with the new account.

> **TIP**
>
> If you signed in with the wrong account during installation, contact support to request an account transfer.

Sign out of GPT for Excel

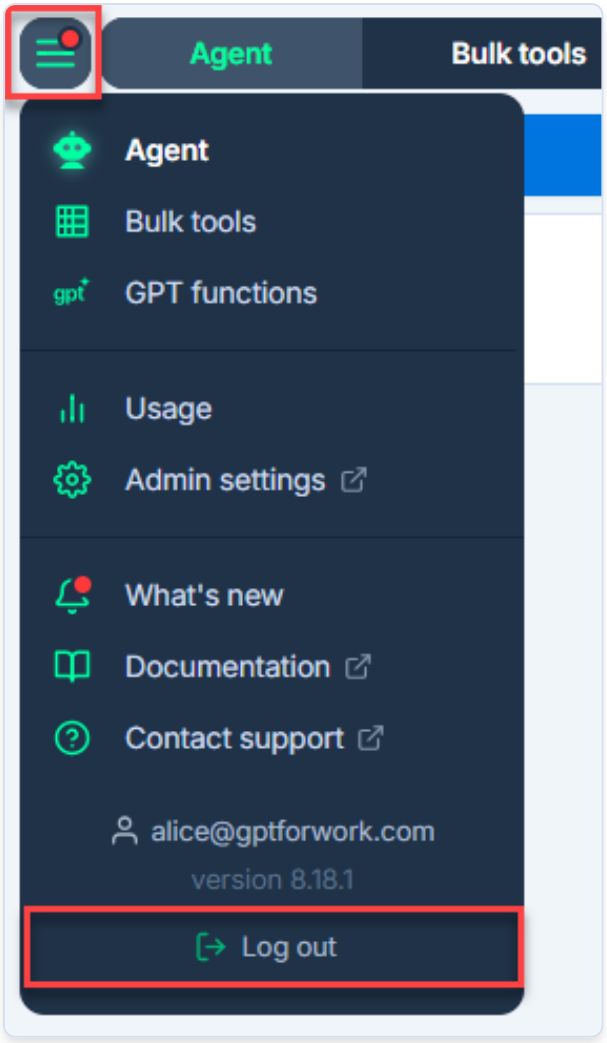
To sign out of GPT for Excel:

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3 In the sidebar, open the main menu. Then, select **Log out**.



INFO

Signing out of GPT for Excel also signs you out of GPT for Word, since they use the same account session.

You have signed out of GPT for Excel. You'll need to sign in again when you next launch the add-in.

What's next

- Sign in with Microsoft or Google.

Manage models for GPT for Excel

Models are the engines that power GPT for Excel. GPT for Excel supports models from OpenAI, Perplexity, Google, Anthropic, Parallel, DeepSeek, Mistral, and xAI. GPT for Excel also supports Azure models and open-source models through any OpenAI-compatible API endpoint.

> **INFO**
>
> As a space owner or admin, you can control which models are available to users in your space through the AI configuration section in the GPT for Work dashboard.

Which model is the right one depends on your specific needs in terms of response quality, speed, privacy, and cost. The following table summarizes the kinds of models available in GPT for Excel, from zero setup to more advanced options:

Use	Setup	When
Built-in models	No setup. Use the **default model** or select a different model: • Agent • Bulk tools • GPT functions	You want to get started quickly without setting up API keys. ✅ Available models
Provider-connected models 🔑	Set an API key for an AI provider.	• You want to benefit from the terms of service you have with an AI provider, such as cost, privacy, and rate limits. • You want to use a model that is not available as a built-in model. • You want more control. ✅ Available models

Use	Setup	When
Endpoint-connected models from cloud platforms	**Azure OpenAI models** Set up an Azure OpenAI API endpoint.	• Your company is already using Azure with a consumption commitment. Using Azure OpenAI counts towards that commitment. • You want to use models that are available through Azure OpenAI. ✅ Available models
	OpenAI-compatible models Set up an OpenAI-compatible platform endpoint.	• You want to use models available on a cloud-based AI platform such as Anyscale or Fireworks AI. • You need a fast and powerful inference infrastructure that scales to your performance needs.
Endpoint-connected models from local servers	**OpenAI-compatible models** Set up an OpenAI-compatible server endpoint.	• You want to use models available on a local LLM server such as LM Studio or Ollama. • You want to run models on a local machine. • You do not want to send or receive any prompt or response data over the internet.

What's next

- Set and manage your API keys.
- Use Azure OpenAI models.
- Use OpenAI-compatible models through API endpoints.
- Use Ollama models.

Manage API keys for GPT for Excel

Using API keys offers several benefits:

- **More models.** Access a wider range of models from a bigger pool of AI providers than what is available without an API key.
- **More control.** Monitor and manage AI usage and costs directly on the AI provider's platform.
- **Improved privacy.** We do not log inputs and outputs when you use a connected model. For more information, see our security and privacy FAQ.
- **Reduced costs.** If you're on a pay-as-you-go space, using an API key reduces costs. You pay the model usage fees directly to the AI provider, while paying a reduced service fee for GPT for Work.

Create API keys

> **INFO**
>
> Make sure your keys belong to an organization or project with the following rate limits:
>
> - RPM (requests per minute) ≥ 1,000
> - TPM (tokens per minute) ≥ 1,000,000
>
> The actual limits required depend on your workload requirements.

To create an API key, go to the AI provider's platform and follow their instructions:

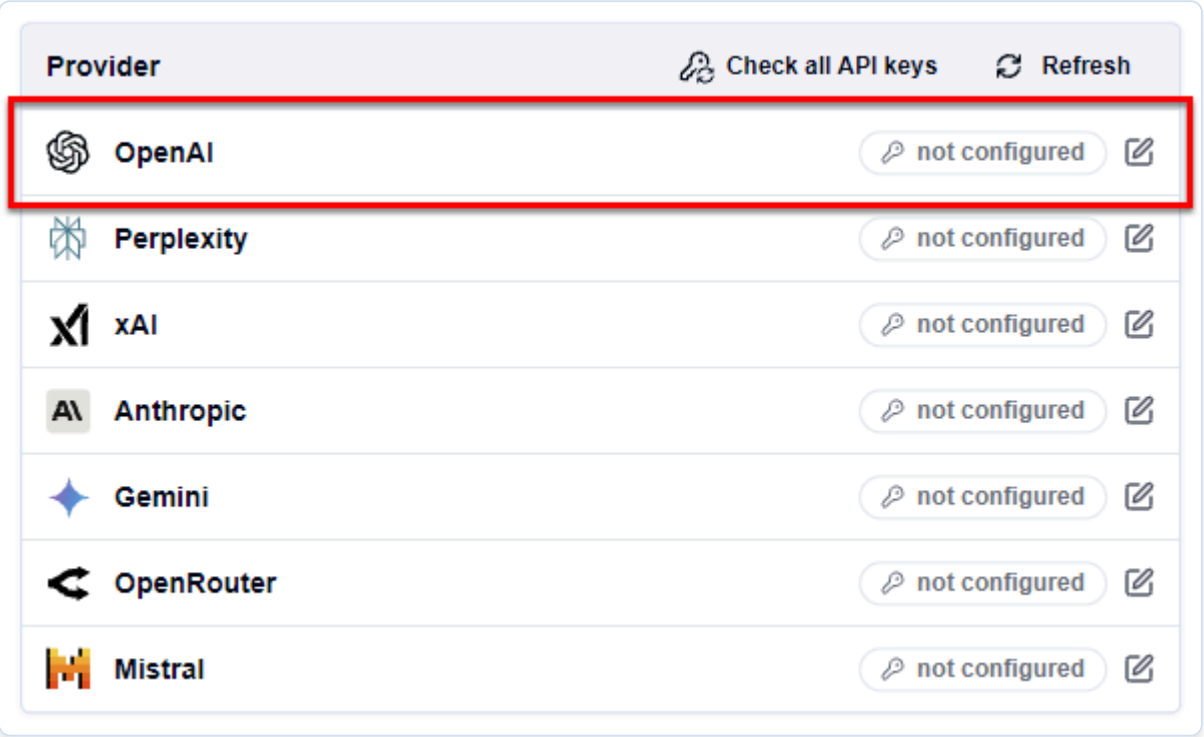
- OpenAI Platform for GPT models and fine-tuned models
- Anthropic Console for Claude models
- DeepSeek Platform for DeepSeek models
- Google AI Studio for Gemini models
- Mistral AI for Mistral models
- Perplexity for Perplexity models
- xAI Cloud Console for xAI models

Set API keys

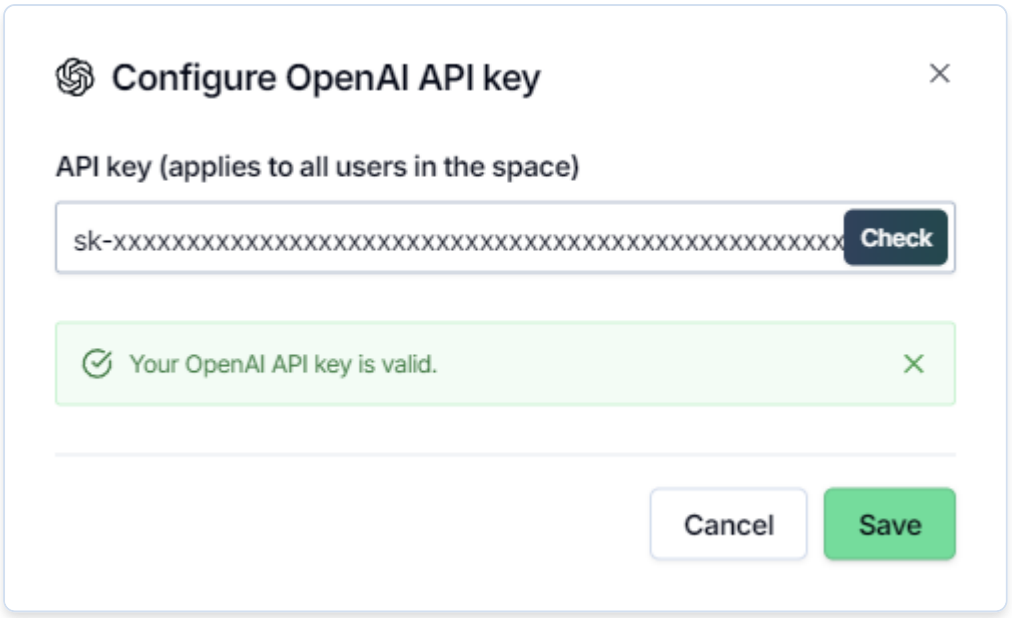
You set API keys for your space in the GPT for Work dashboard. The keys you set for a GPT for Excel and Word space apply to GPT for Excel.

> **INFO**
>
> The API keys you set work for all users in your space. Users automatically have access to models associated with the keys, but only the space owner (you) and admins can set and manage the keys. Billing for a key reflects its usage by all users.

1. Create an API key or request an API key from your AI provider platform admin. Learn more.
2. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
3. In the sidebar, select **AI provider keys**.
4. In the AI providers list, click the provider whose API key you want to set.



5. Set the API key:
   1. In the text field, enter the key.
   2. Click **Check** to verify that the key is valid.
   3. Click **Save**.



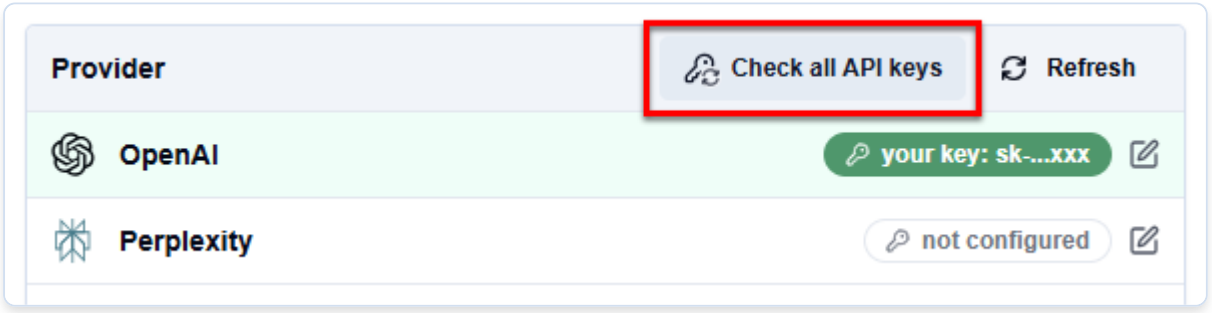
You have set the API key. You can now use any model available with the key in GPT for Excel.

Check API keys

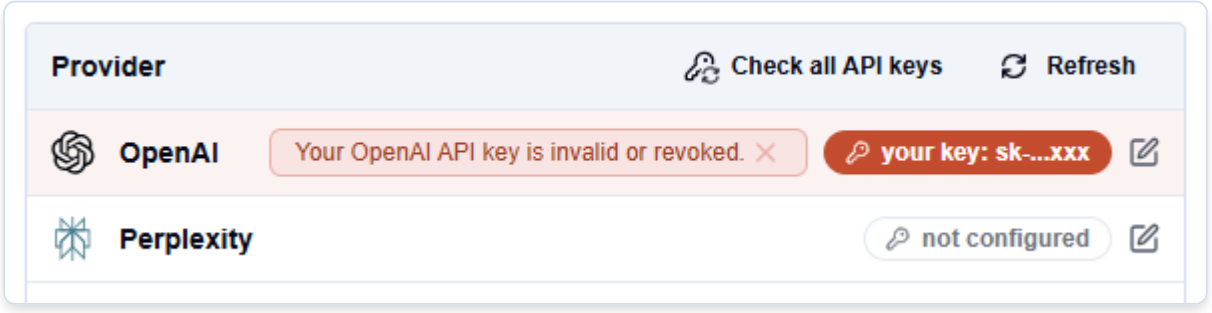
You can verify whether the API keys you have set are valid.

Check all API keys at once

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **AI provider keys**.
3. Click **Check all API keys**.



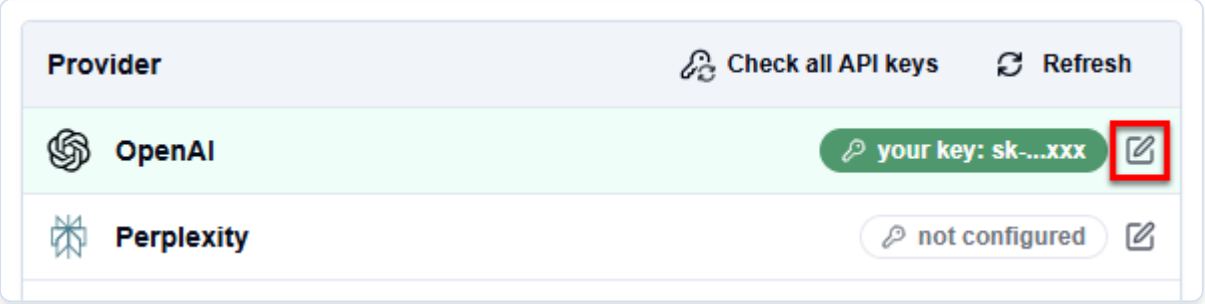
The page displays a warning for each key that is invalid.



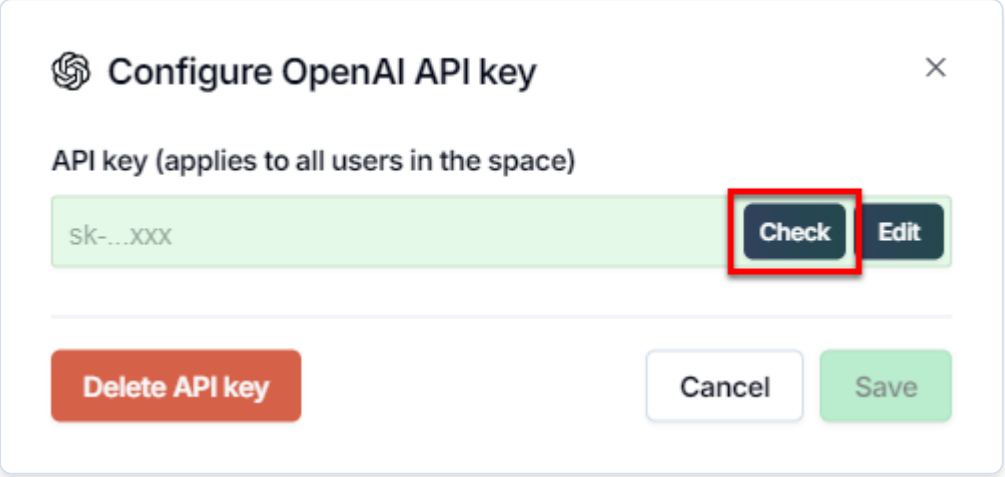
If a key is invalid, create a new key and replace the invalid value with the new one.

Check a specific API key

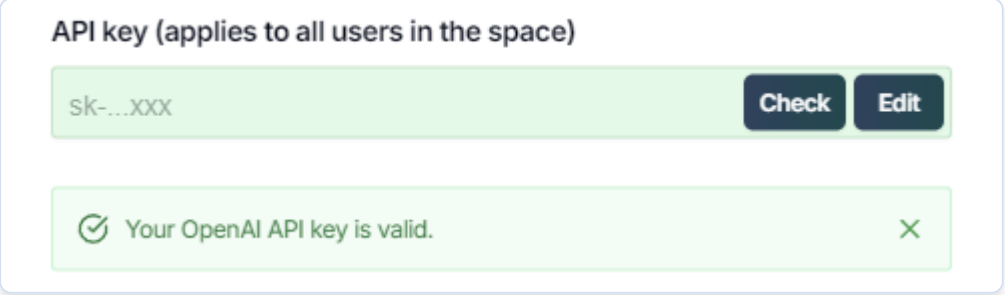
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **AI provider keys**.
3. Click the edit icon for the key you want to check.

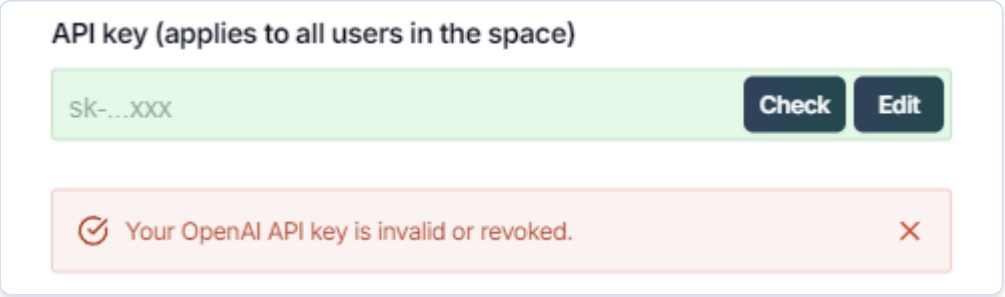


4. Click **Check**.



The UI displays a message indicating whether the key is valid or invalid.





If the key is invalid, create a new key and replace the invalid value with the new one.

Troubleshoot OpenAI API keys

If your API key is not working, check on the OpenAI Platform that the key's permissions are not restricted. The following permissions are required:

Permission	Description
List models	List all available models.
Model capabilities	Use model capabilities, such as chat completion.
Fine-tuning	Connect OpenAI fine-tuned models to Microsoft Excel.

What's next

- Check which provider-connected models you can use.
- Select a provider-connected model:
  - Select the Agent models.
  - Select a model for bulk tools.
  - Select a model for GPT functions.
- Need to use your own AI service endpoint for models? Connect to an OpenAI-compatible endpoint.

Connect to Azure OpenAI from Excel

GPT for Excel can connect to Azure AI to use OpenAI models. In Azure AI, each model has a dedicated API endpoint that you use to access the model.

Prerequisites

- Your space is on the Business or Enterprise plan or pay-as-you-go pricing.
- You have an Azure account with Azure AI service access.
- You have deployed a model in a project.

To use an Azure OpenAI model:

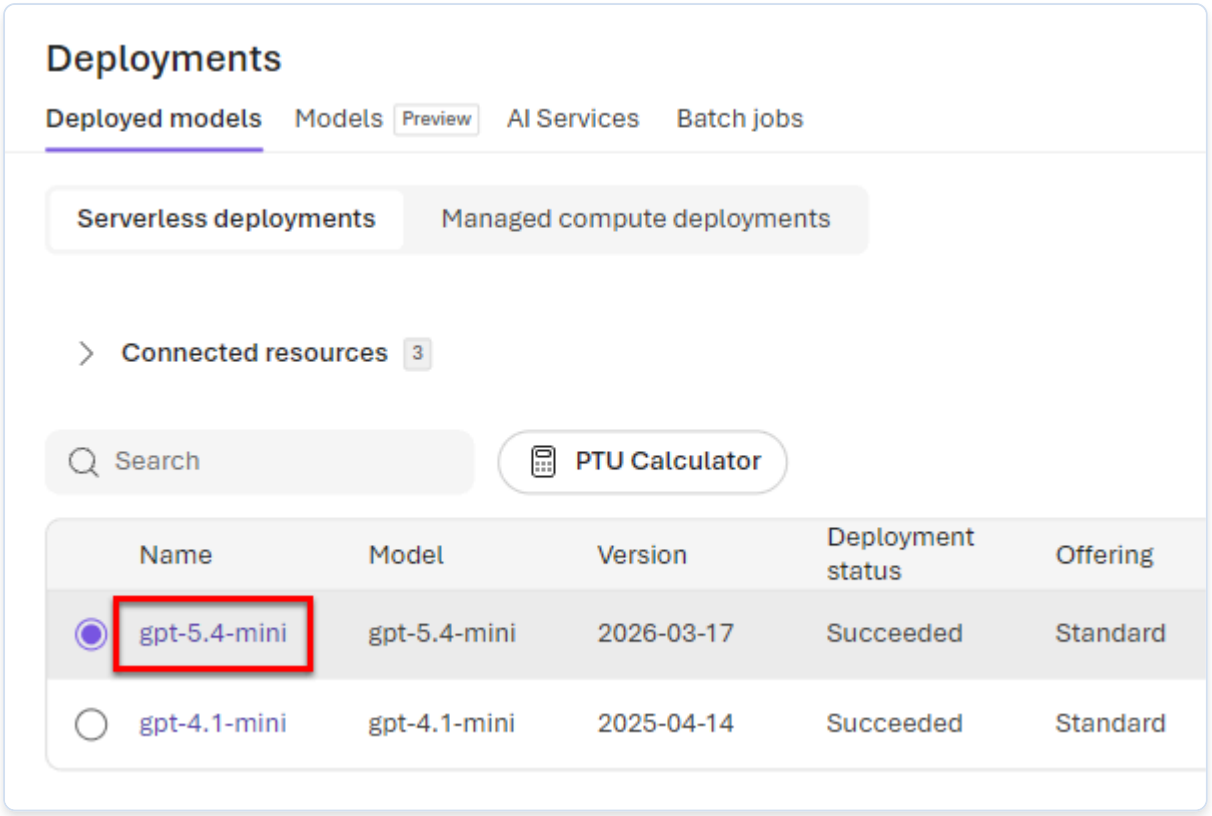
1. Get the Azure OpenAI API endpoint details for the model.
2. Set up the Azure OpenAI API endpoint.
3. Select the Azure OpenAI model.

Get the Azure OpenAI API endpoint details

To connect to Azure AI from GPT for Excel, you need the endpoint URL, API key, deployment name, and model name for the Azure OpenAI model you want to use.

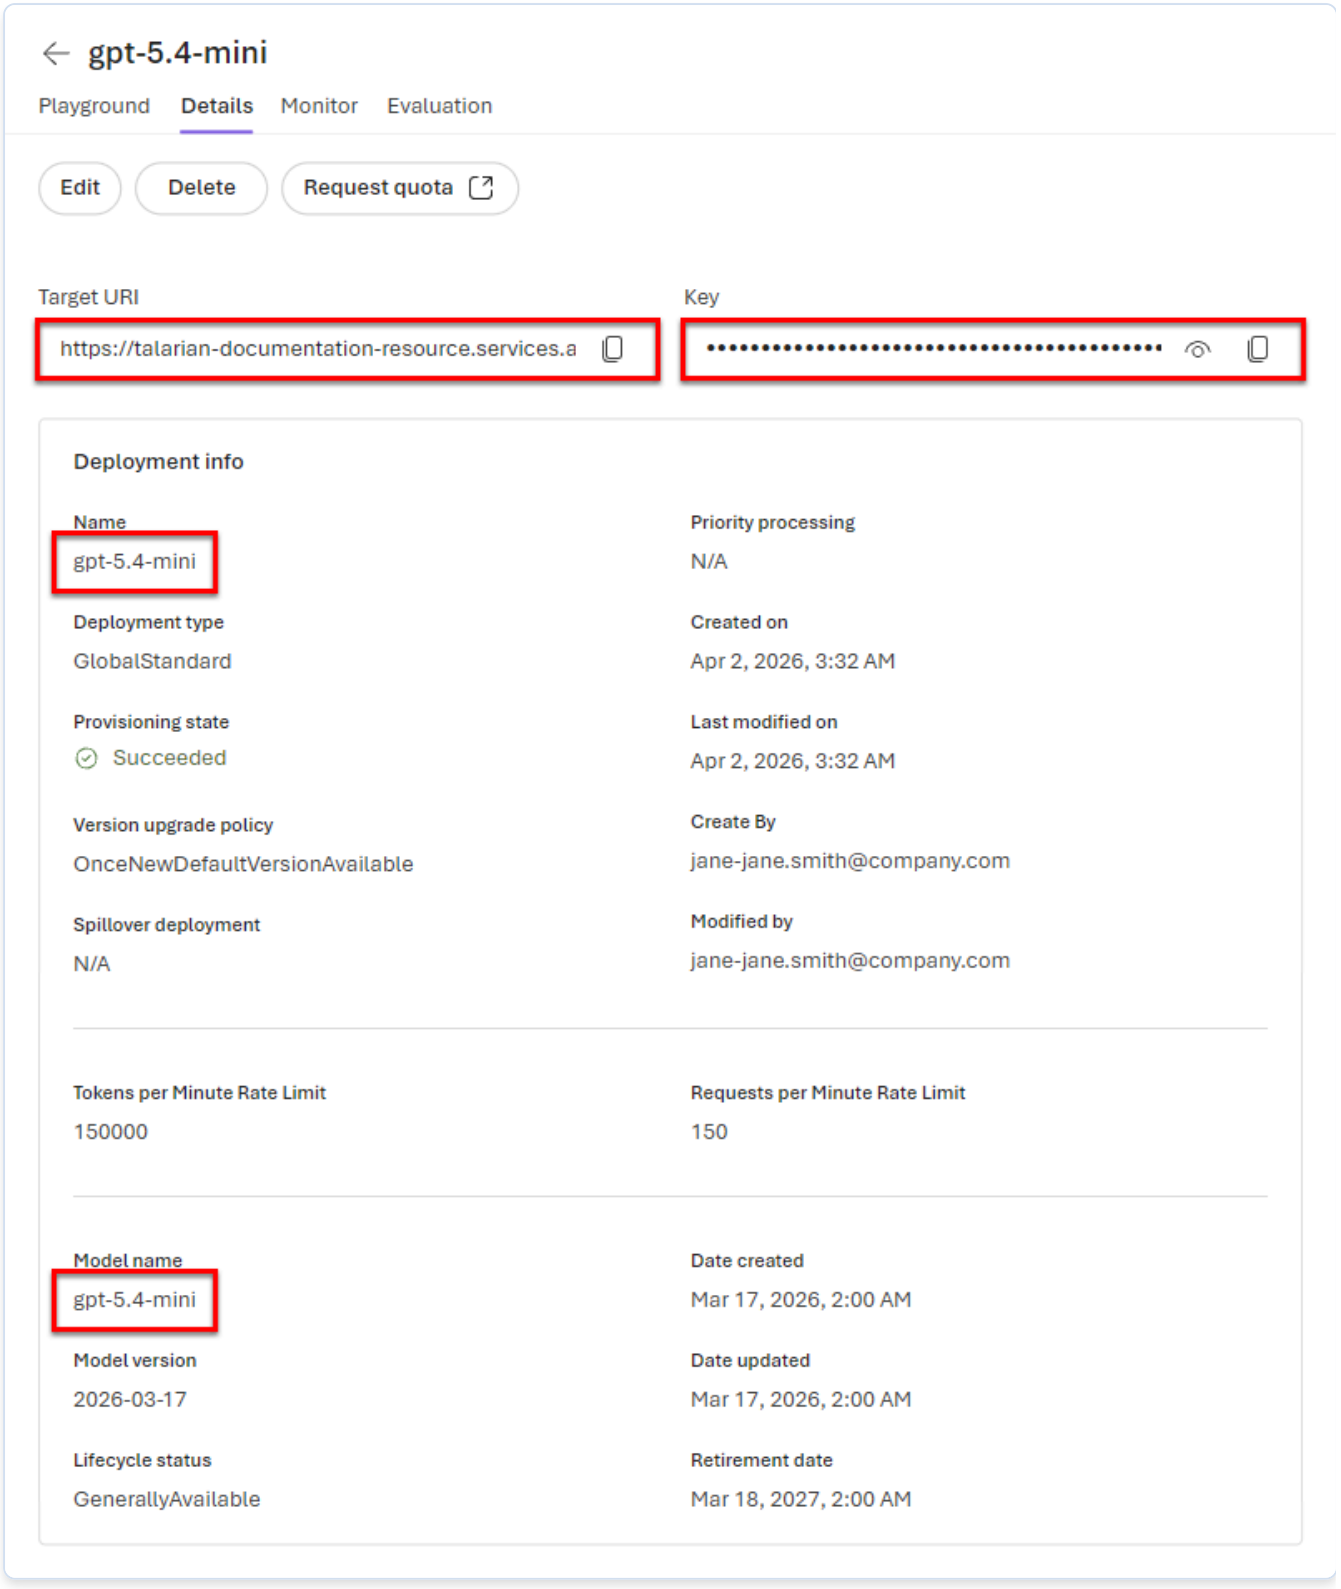
1. Sign in to Microsoft Foundry.
2. Select the project that has the model you want to use.
3. Click **View deployments**.

4 In **Deployed models**, in the models list, click the name of the model you want to use.



5 Select **Details**.

6 In the model details, copy the **Target URI**, **Key**, **Deployment name**, and **Model name**.



Set up the Azure OpenAI API endpoint

1 Sign in to the GPT for Work dashboard with your Microsoft or Google account.

2 In the sidebar, select **Custom API endpoints**.

3 Select the **Bulk models** tab.

4 Click **Azure OpenAI endpoints (Chat Completions API)**.

> **TIP**
>
> If you already have Azure OpenAI API endpoints set up, click **Add an Azure OpenAI API endpoint**.

5 Define the endpoint settings:

- **Endpoint URL**: Enter the endpoint URL for the model using the Azure OpenAI Chat Completions API endpoint format:

  ```
  https://<resource-name>.<domain>/openai/deployments/<deployment-name>/chat/completions?api-version=<api-version>
  ```

  If the URL you copied does not use the above format, rewrite it. For example, let's assume that your deployment name is `gpt-5.4-mini`, the API version is `2025-04-01-preview`, and you copied the following URL:

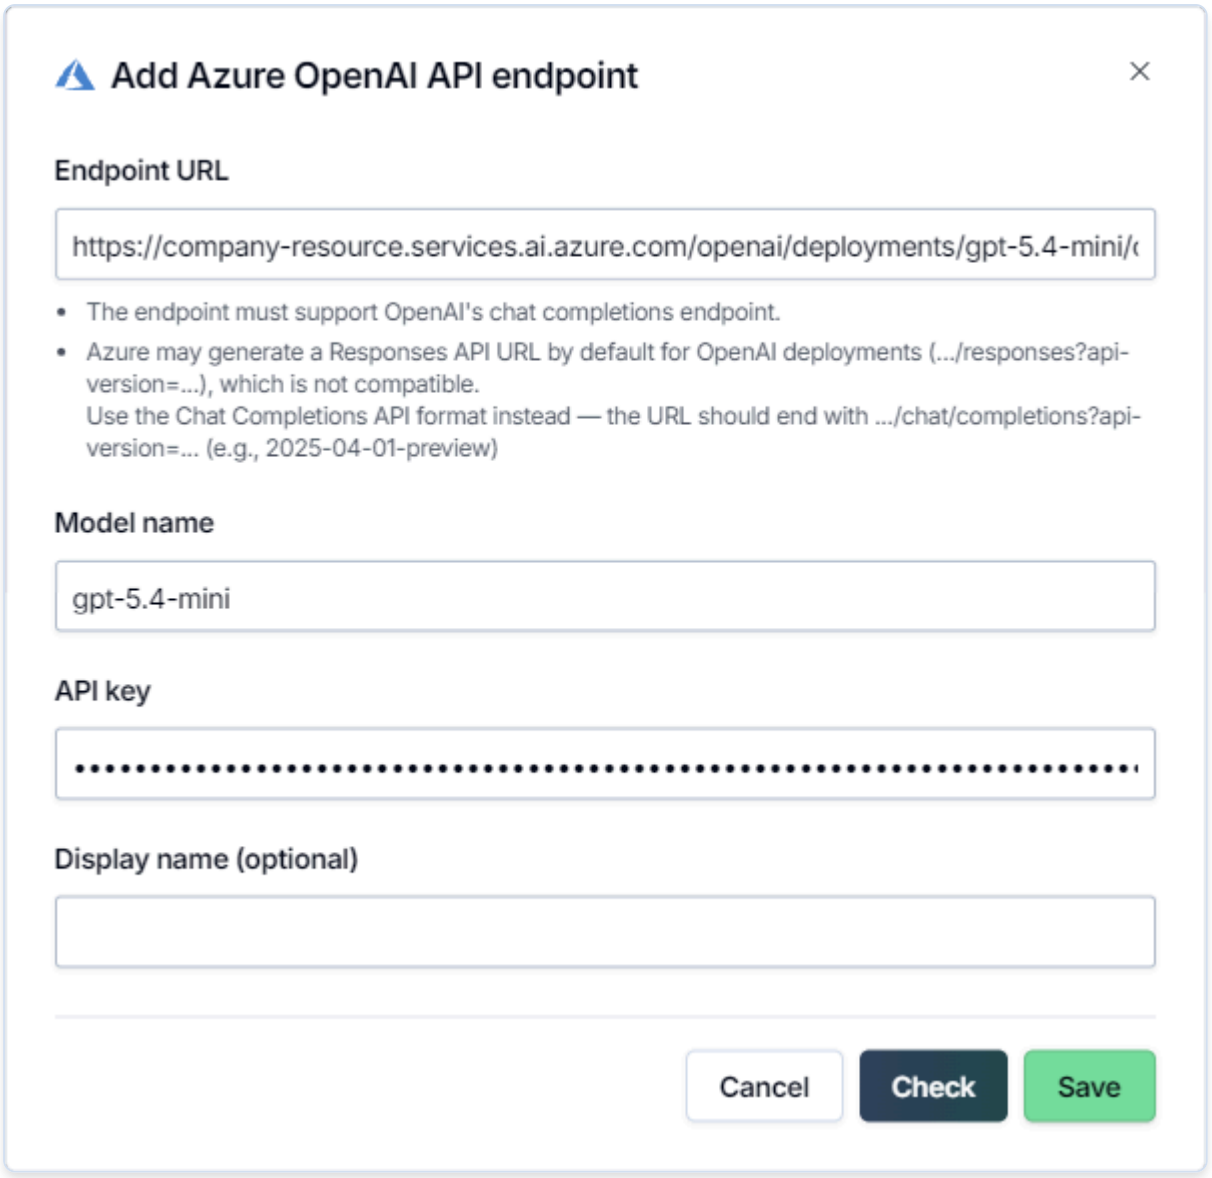
  ```
  https://company-resource.services.ai.azure.com/openai/v1/responses
  ```

  To rewrite the URL into the Chat Completions API format, replace everything after the domain with `/openai/deployments/<deployment-name>/chat/completions?api-version=<api-version>` and replace the placeholders with your deployment name and the API version:

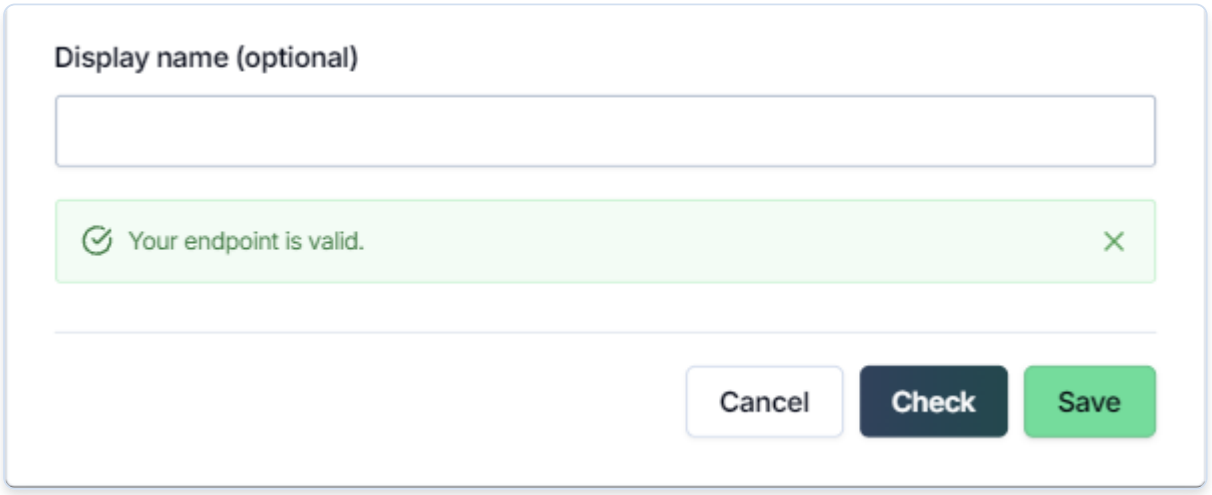
  ```
  https://company-resource.services.ai.azure.com/openai/deployments/gpt-5.4-mini/chat/completions?api-version=2025-04-01-preview
  ```

  > **INFO**
  >
  > The `api-version` is the Azure OpenAI REST API version, which is not shown on the Microsoft Foundry model details page. As a rule, use `2025-04-01-preview` as the value.

- **Model name**: Enter the model name.
- **API key**: Enter the API key for accessing the endpoint.
- **Display name (optional)**: Enter a custom display name for the model. If not set, the model switcher uses **Model name** as the display name.



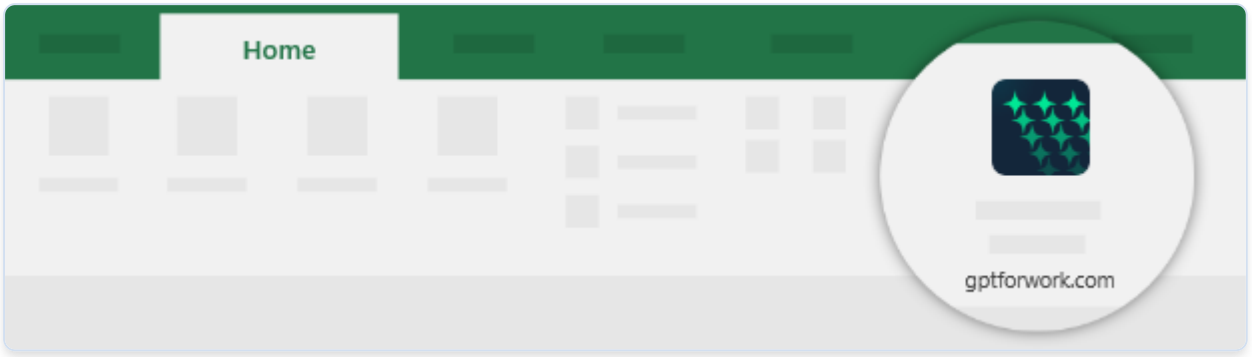
6 Click **Check** to validate the settings and click **Save**.



You have set up the Azure OpenAI API endpoint. You can now use the endpoint model in GPT for Excel.

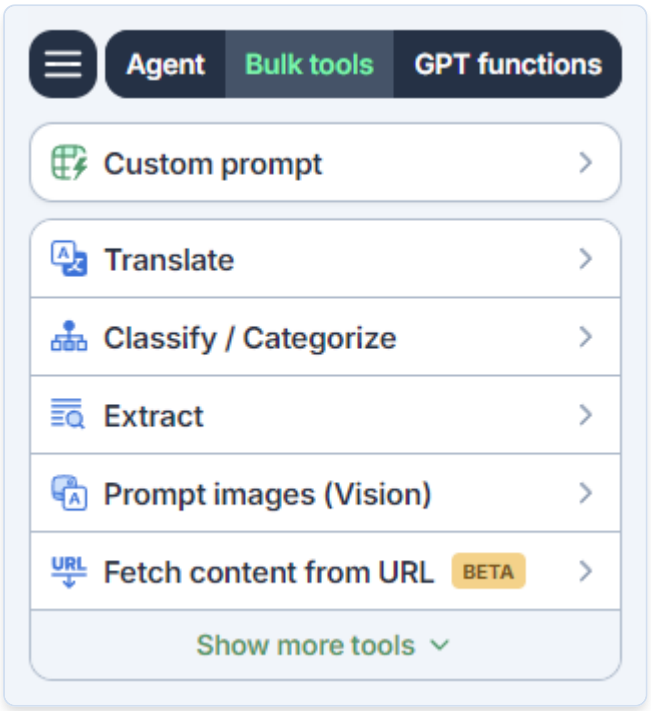
Select the Azure OpenAI model

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.

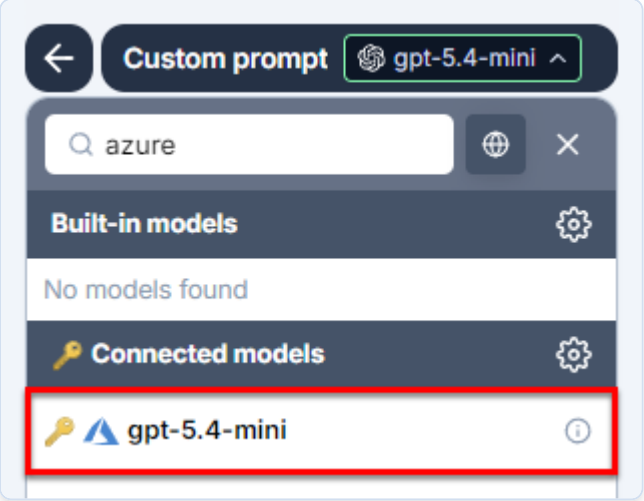


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Bulk tools** and select one of the tools.



4 Open the model switcher. You can find the Azure OpenAI model under **Connected models**.



5 Select the model.

GPT for Excel now uses the selected model to generate responses for bulk tools and GPT functions.

Connect to OpenAI-compatible API endpoints from Excel

> **INFO**
>
> This feature is only available if your space is on the Business or Enterprise plan or pay-as-you-go pricing.

GPT for Excel can connect to any AI service that provides an OpenAI-compatible API endpoint. You can connect to two main types of services: **cloud-based platforms**, which provide access to models over the internet with no software installation or setup required on your part, and **local servers**, which you run on a local machine, such as your own computer or another computer on a local network.

You can use the endpoint-connected models both as Agent models and bulk models, but the endpoint requirements are different:

- **Agent models**: The endpoint must be compatible with the OpenAI Responses API and be publicly accessible over the internet through HTTPS. In addition, the model must support a context size of at least 32K tokens.
- **Bulk models**: The endpoint must be compatible with the OpenAI Chat Completions API.

Here are some popular services that provide an OpenAI-compatible API endpoint:

- **Cloud-based platforms:**
  - Anyscale
  - Fireworks AI
  - Together AI
- **Local servers:**
  - LM Studio
  - LocalAI
  - Ollama

> **INFO**
>
> Depending on the model you use, the model settings defined in GPT for Excel (temperature (creativity), top-p, frequency penalty, and presence penalty) may not work as expected. If you get bad results, try tweaking the settings.

To use models from an OpenAI-compatible API endpoint:

1. Set up the endpoint.
2. Select a model available through the endpoint.

Set up an API endpoint

Agent models

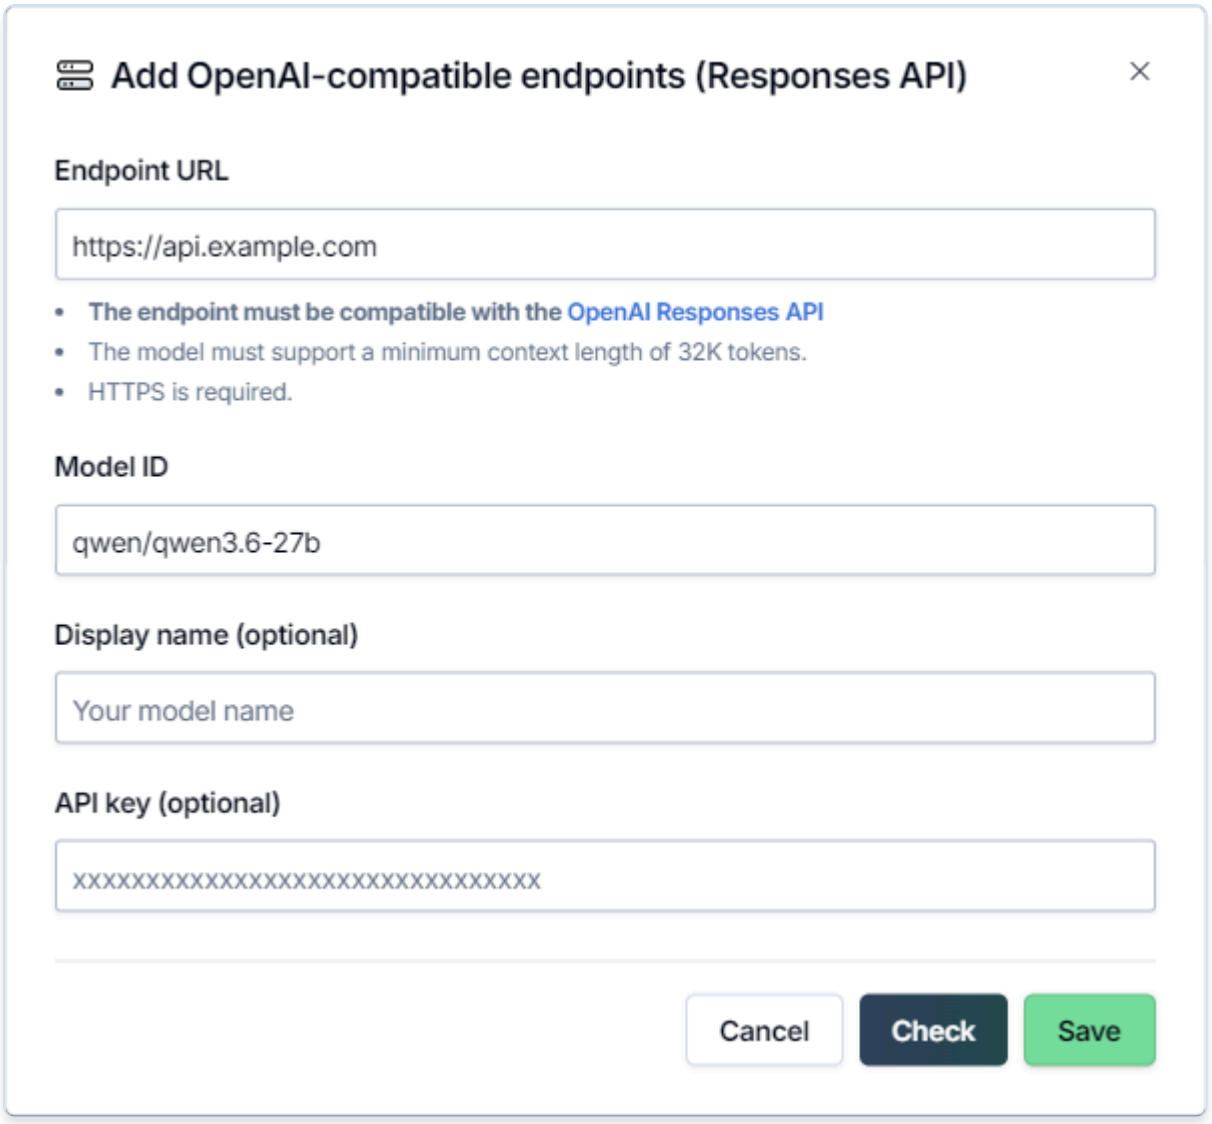
Prerequisites

- You have an endpoint URL and, if needed, an API key for accessing the service.
- The endpoint supports the following OpenAI-compatible API resource path: `/v1/responses`
- The endpoint is publicly accessible over the internet through HTTPS.
- The model supports a context size of at least 32K tokens.

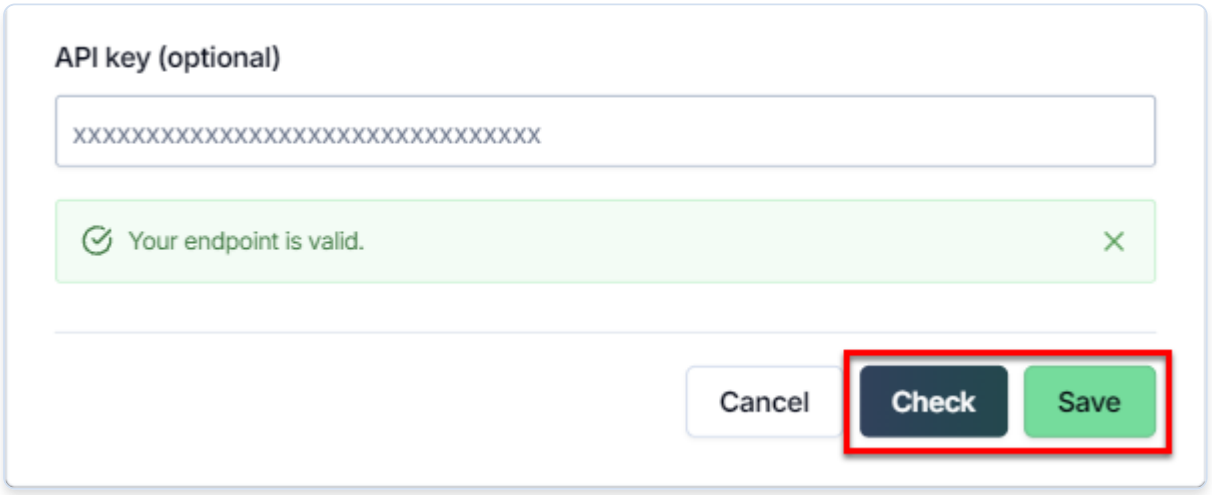
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Agent models** tab.
4. Click **OpenAI-compatible endpoints (Responses API)**.

> **TIP**
>
> If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the URL of the service that provides the endpoint.
   - **Model ID**: Enter the ID of the model to use through the endpoint. The model ID is sent to the endpoint in the request payload.
   - **Display name (optional)**: Enter a short name that identifies the model in the Agent model switcher.
   - **API key (optional)**: If required, enter the API key for accessing the service.



6 Click **Check** to validate the settings and click **Save**.



You have set up the OpenAI-compatible API endpoint. The model is now available in the Agent model switcher and on the models page of the GPT for Work dashboard.

Bulk models

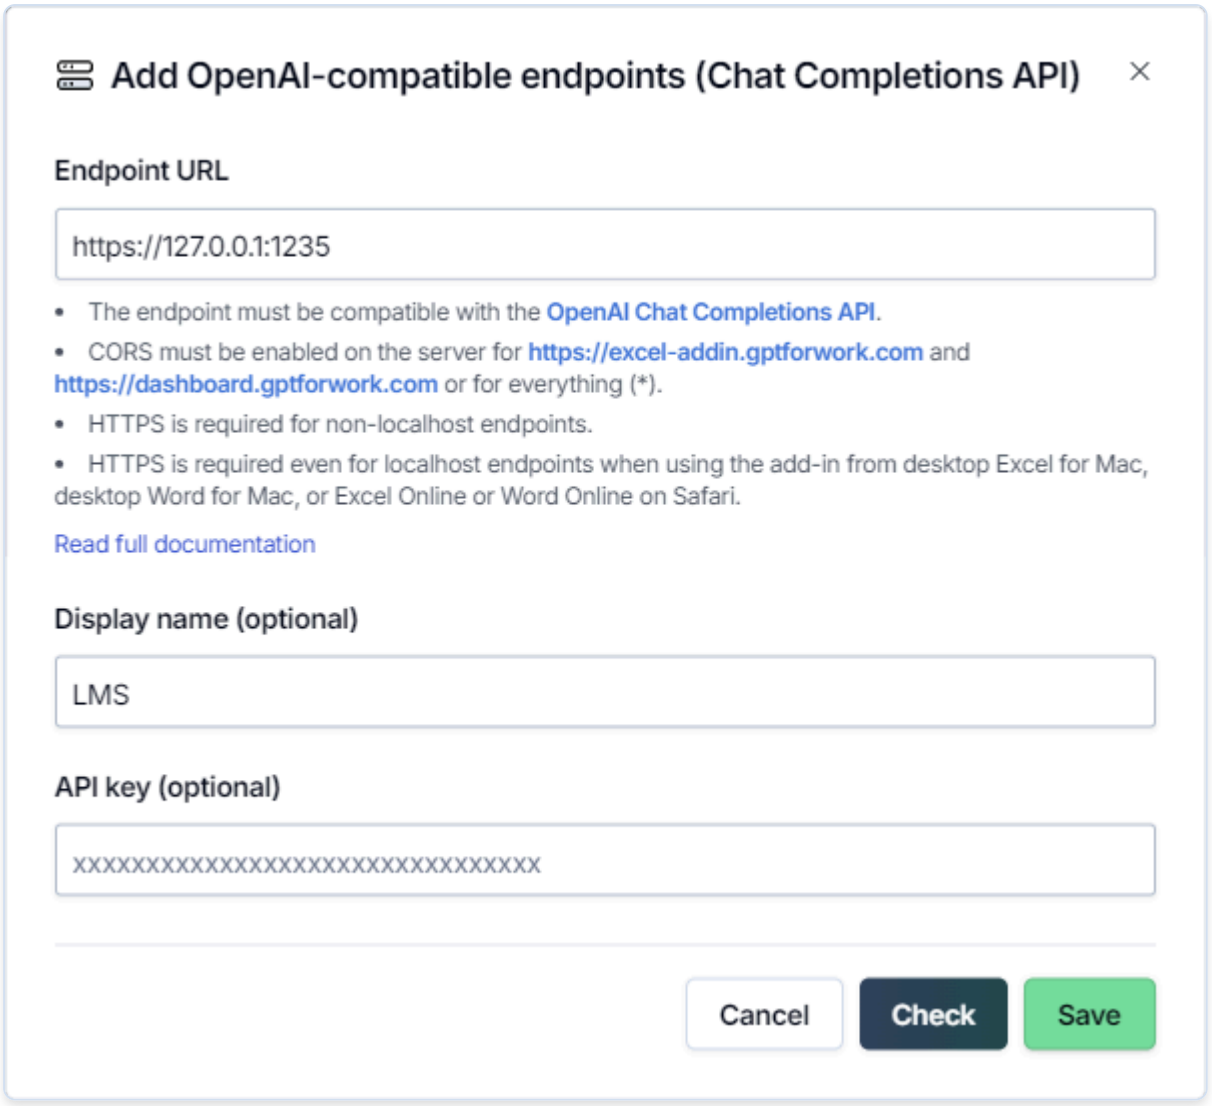
Prerequisites

- You have an endpoint URL and, if needed, an API key for accessing the service.
- The endpoint supports the following OpenAI-compatible API resource paths:
  - `/v1/models`
  - `/v1/chat/completions`

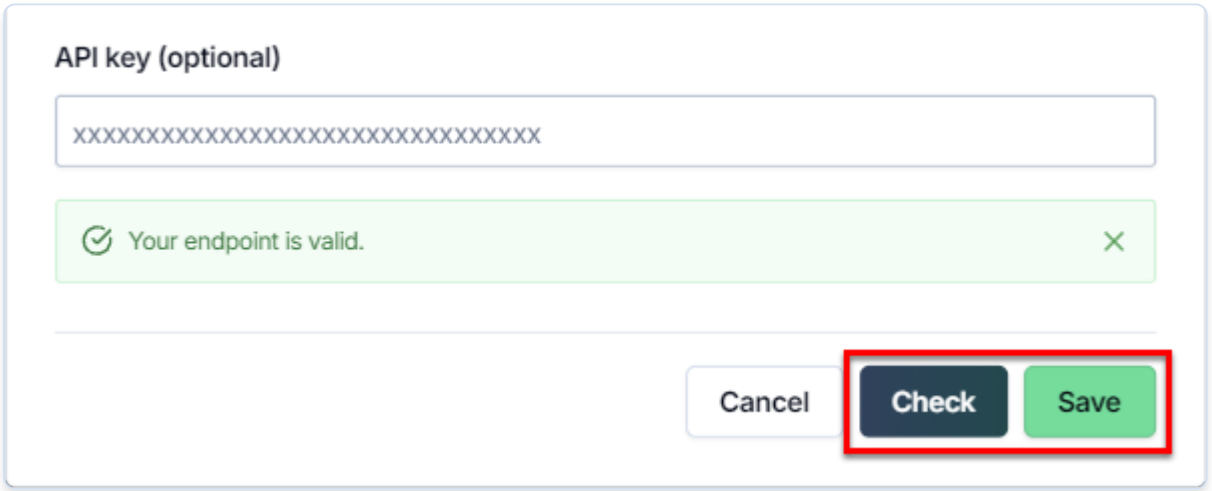
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Bulk models** tab.
4. Click **OpenAI-compatible endpoints (Chat Completions API)**.

> **TIP**
>
> If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the URL of the service that provides the endpoint.
   - **Display name (optional)**: Enter a short name that is prefixed to the model names in the model switcher. By default, the model names are prefixed with "custom/".
   - **API key (optional)**: If required, enter the API key for accessing the service.



6 Click **Check** to validate the settings and click **Save**.



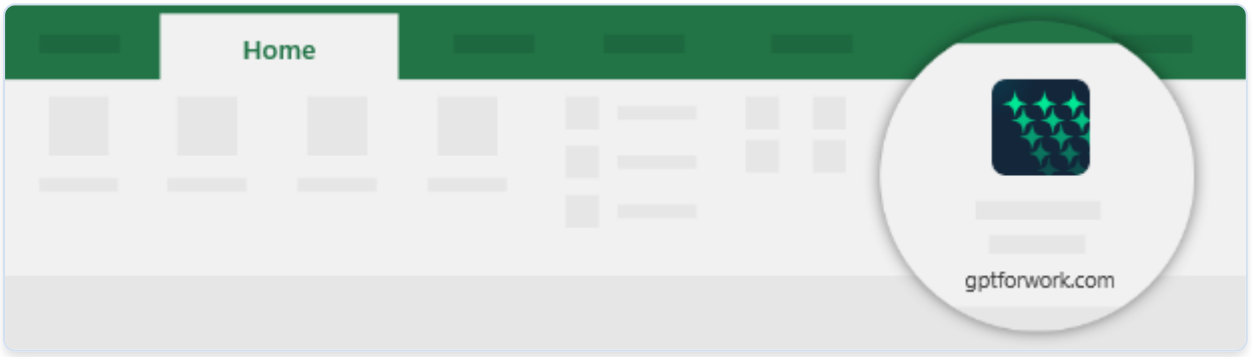
> **TIP**
>
> If you get an error saying that the URL is invalid, make sure that the endpoint supports:
>
> - **Cross-origin resource sharing (CORS).** See the LM Studio setup instructions for macOS or Windows.
> - **HTTPS (needed with Safari and with Microsoft Office for Mac).** See the LM Studio setup instructions for macOS.

You have set up the OpenAI-compatible API endpoint. You can now use all models available through the endpoint in GPT for Excel.

Select a model from the API endpoint

Agent models

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Agent**.
4. Open the model switcher and select a model. You can find the endpoint-connected models at

the end of the list.

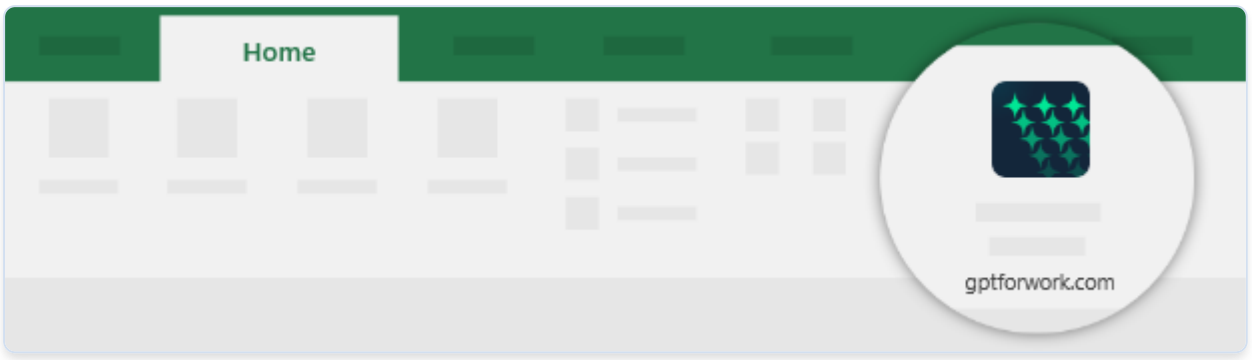


GPT for Excel now uses the selected model as the Agent model.

> **INFO**
>
> Some Agent tools, such as web search, are not available when the Agent uses an endpoint-connected model.

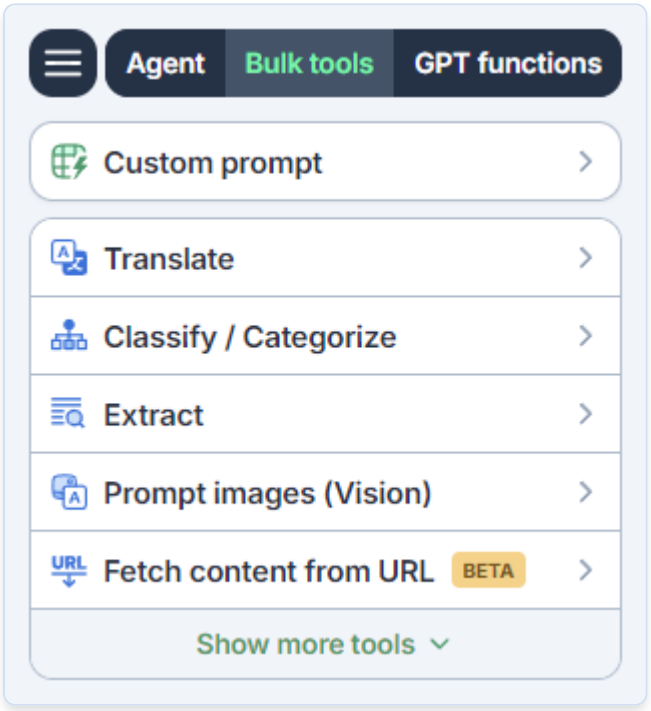
Bulk models

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



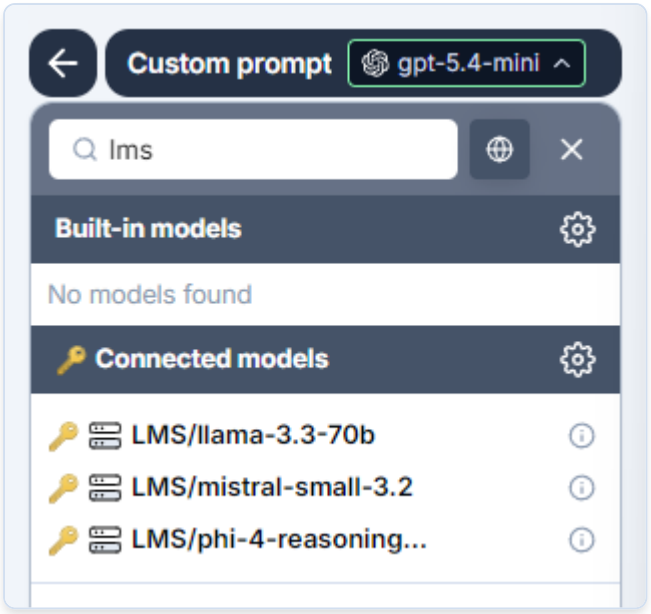
> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Bulk tools** and select one of the tools.



4. Open the model switcher and select a model. You can find the endpoint-connected models

under **Connected models**, at the end of the list.



> **INFO**
>
> The available models depend on the models provided by the service.

GPT for Excel now uses the selected model to generate responses for bulk tools and GPT functions.

What's next

- To try out an OpenAI-compatible AI service with GPT for Excel, set up LM Studio on macOS or Windows.

Set up LM Studio on macOS

This guide walks you through the main steps of setting up LM Studio for use in GPT for Excel on macOS.

This guide assumes that you use GPT for Excel on the same machine that hosts LM Studio.

Prerequisites

- System requirements
- Mac user account with administrator (sudo) privileges
- Installed software:
  - curl (ships with macOS)
  - Homebrew
  - OpenSSL (ships with macOS)

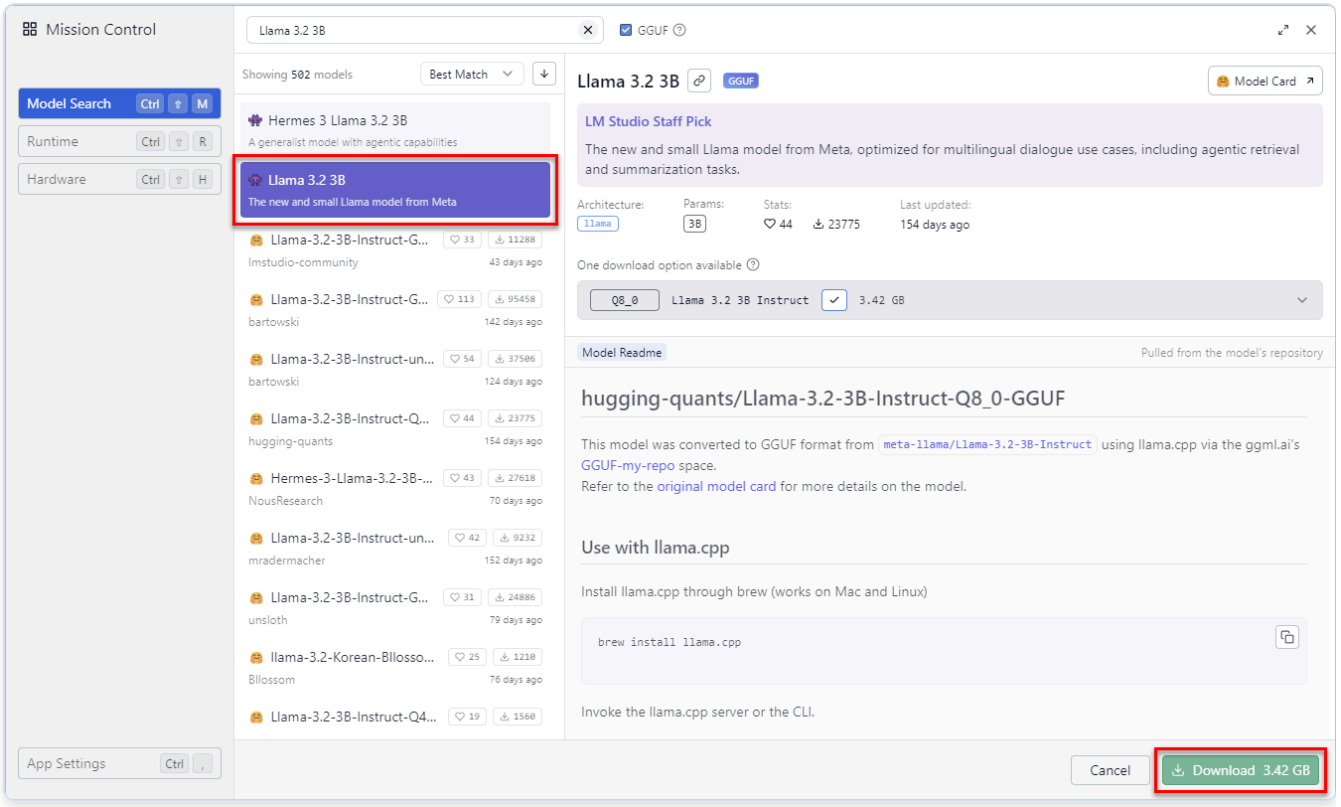
To set up LM Studio on macOS:

1. Install LM Studio and download a model.
2. Start and configure the LM Studio server.
3. (Optional) Enable HTTPS for the LM Studio server.

Install LM Studio and download a model

1. Download and run the macOS installer. Follow the on-screen instructions to complete the installation.
2. Run LM Studio.
3. On the welcome screen, click **Skip**.
4. In LM Studio, in the sidebar, select **Discover**.
5. Find and select the model you want to use, and click **Download**. For example, to get started

with a small model, select Llama 3.2 3B.

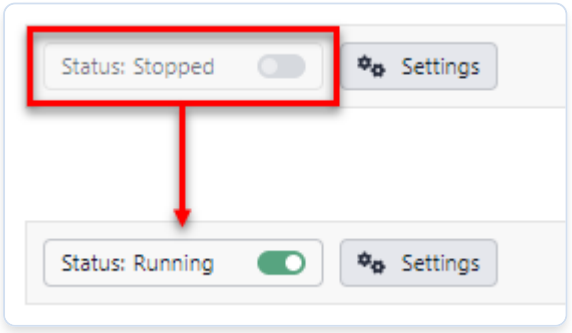


After the download completes, the model is available for prompting in LM Studio.

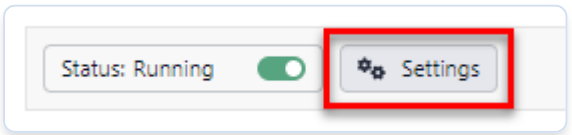
You have installed LM Studio and downloaded your first model. For more information about working with models, see the LM Studio documentation.

Start and configure the LM Studio server

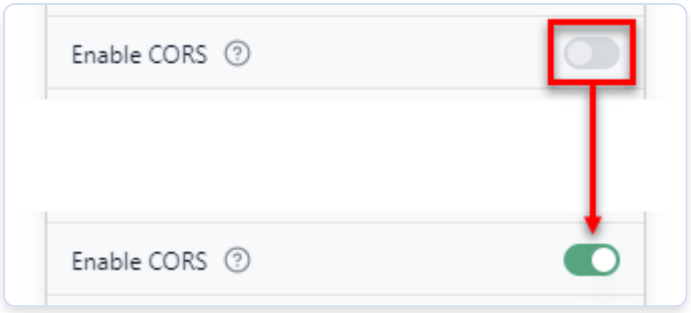
1. In LM Studio, in the sidebar, select **Developer**.
2. Click the **Status** toggle to change the status from **Stopped** to **Running**. You have started the LM Studio server.



3. Click **Settings**.



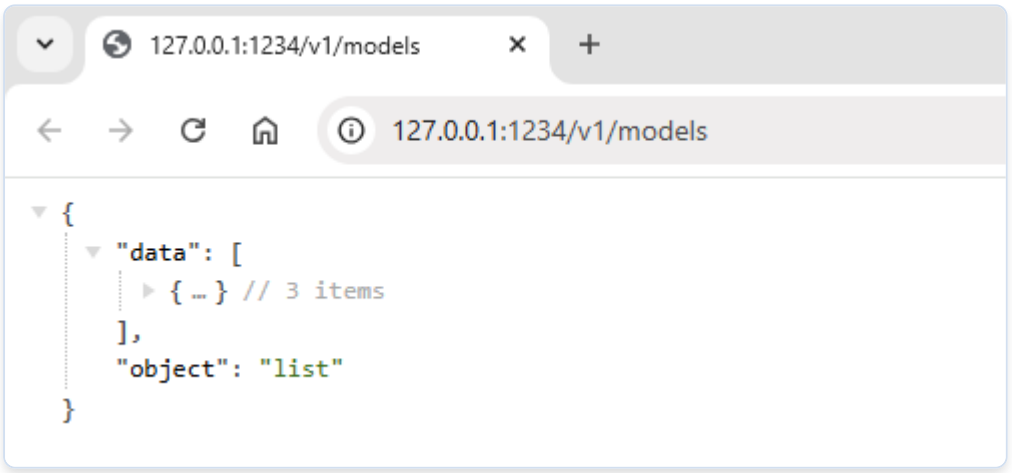
4 Click the **Enable CORS** toggle to enable the setting.



> **INFO**
>
> By default, the LM Studio server only accepts same-origin requests. Since GPT for Excel always has a different origin from the LM Studio server, you must enable cross-origin resource sharing (CORS) for the server.

5 To verify that the `/v1/models` endpoint of the LM Studio server works, open http://127.0.0.1:1234/v1/models.

GPT for Excel uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `data` property listing all currently installed models:



You have started and configured the LM Studio server.

> **WHAT'S NEXT**
>
> You have completed the minimum setup required to access LM Studio from GPT for Excel on the same machine. You can now set `http://127.0.0.1:1234` as the endpoint URL in GPT for Excel provided the add-in is not running on Safari or on Microsoft Excel for Mac.
>
> If you use GPT for Excel on Safari or on Excel for Mac, enable HTTPS for the LM Studio server.

Enable HTTPS for the LM Studio server

The LM Studio server uses HTTP to serve models, while GPT for Excel runs on HTTPS. By default, therefore, any request from GPT for Excel to the server is a mixed-content request (HTTP vs. HTTPS). Modern web browsers do not allow mixed content, as a rule. The only exceptions are mixed-content requests to `http://127.0.0.1` and `http://localhost`, which most browsers treat as safe and therefore allow. Safari, however, blocks mixed-content requests even on the current machine.

To avoid mixed content, you must make the LM Studio server accessible over HTTPS if you use GPT for Excel from:

- Excel for the web on Safari
- Excel for Mac (uses Safari)

You make the LM Studio server accessible over HTTPS by setting up a reverse proxy that hides the server behind an HTTPS interface. This guide uses nginx to set up the reverse proxy.

> **TIP**
>
> You can also use Cloudflare Tunnel, ngrok, or a similar cloud-based tunneling service to set up HTTPS with minimal configuration. Note that with a cloud-based service your traffic will be routed through an external service over the internet.

To set up the reverse proxy:

1. Install nginx:

```
brew install nginx
```

2. Create a self-signed SSL certificate for the LM Studio server:
   1. Change to the nginx configuration directory (varies based on your Mac version):

```
# Mac with an Apple silicon processor
cd /opt/homebrew/etc/nginx

# Mac with an Intel processor
cd /usr/local/etc/nginx
```

   2. Generate the certificate:

```
openssl req \
  -x509 \
  -newkey rsa:2048 \
  -nodes \
```

```
  -sha256 \
  -days 365 \
  -keyout lm-studio.key \
  -out lm-studio.crt \
  -subj '/CN=127.0.0.1' \
  -extensions extensions \
  -config <(printf "
[dn]\nCN=127.0.0.1\n[req]\ndistinguished_name=dn\n[extensions]\nsubjectAl
```

The command generates two files in the current directory:

- `lm-studio.crt`: Public self-signed certificate
- `lm-studio.key`: Private key used to sign the certificate

3. Add the certificate to trusted root certificates:

```
sudo security add-trusted-cert -d -r trustRoot -k
/Library/Keychains/System.keychain ./lm-studio.crt
```

3. Create a site configuration for the LM Studio server:

1. Open the `nginx.conf` file in your preferred text editor. The file is in the nginx configuration directory:

```
# Mac with an Apple silicon processor
/opt/homebrew/etc/nginx/nginx.conf

# Mac with an Intel processor
/usr/local/etc/nginx/nginx.conf
```

2. Replace the content of the file with the following configuration:

```
events {}

http {
    server {
        listen       127.0.0.1:1235 ssl;
        server_name  127.0.0.1;

        http2 on;

        ssl_certificate      lm-studio.crt;
        ssl_certificate_key  lm-studio.key;

        location / {
            proxy_pass  http://127.0.0.1:1234;

            # Keep connection to backend alive.
```

```
            proxy_http_version  1.1;
            proxy_set_header    Connection '';
        }
    }
}
```

The configuration uses port `1235` for HTTPS. Requests to `https://127.0.0.1:1235` are forwarded to the LM Studio server running at `http://127.0.0.1:1234`.

3. Save and close the file.
4. Test the nginx configuration:

```
sudo nginx -t
```

Response if the configuration is valid:

```
nginx: the configuration file /opt/homebrew/etc/nginx/nginx.conf syntax
is ok
nginx: configuration file /opt/homebrew/etc/nginx/nginx.conf test is
successful
```

5. Start nginx:

```
sudo nginx
```

4. Verify that the `/v1/models` endpoint of the LM Studio server works over the HTTPS connection:

```
curl -k https://127.0.0.1:1235/v1/models
```

GPT for Excel uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `data` property listing all currently installed models:

```
{
  "data": [
    {
      ...
    },
  ],
  "object": "list"
}
```

You have enabled HTTPS for the LM Studio server.

WHAT'S NEXT

You have completed the setup required to access LM Studio from GPT for Excel on the same machine. You can now set `https://127.0.0.1:1235` as the endpoint URL in GPT for Excel, including when the add-in is running on Safari or on Microsoft Excel for Mac.

Set up LM Studio on Windows

This guide walks you through the main steps of setting up LM Studio for use in GPT for Excel on Windows.

This guide assumes that you use GPT for Excel on the same machine that hosts LM Studio.

Prerequisites

- System requirements
- Windows user account with administrator privileges
- Installed software:
  - curl (ships with Windows)
  - PowerShell 7

To set up LM Studio on Windows:

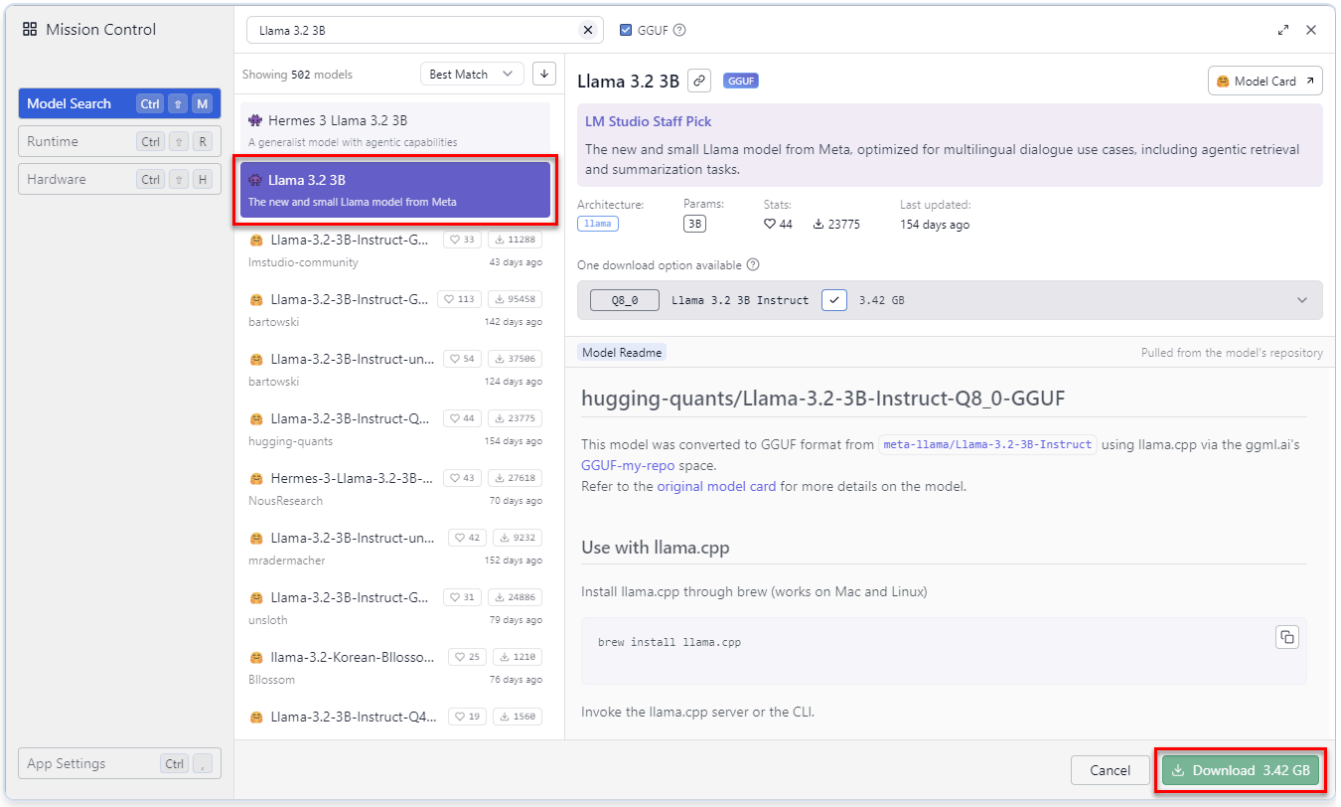
1. Install LM Studio and download a model.
2. Start and configure the LM Studio server.

Install LM Studio and download a model

1. Download and run the Windows installer. Follow the on-screen instructions to complete the installation.

   The installer sets LM Studio to start automatically as a background service on system boot.
2. Run LM Studio.
3. On the welcome screen, click **Skip**.
4. In LM Studio, in the sidebar, select **Discover**.
5. Find and select the model you want to use, and click **Download**. For example, to get started

with a small model, select Llama 3.2 3B.

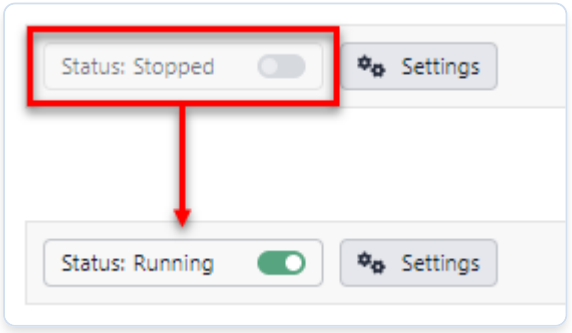


After the download completes, the model is available for prompting in LM Studio.

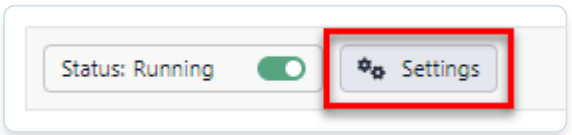
You have installed LM Studio and downloaded your first model. For more information about working with models, see the LM Studio documentation.

Start and configure the LM Studio server

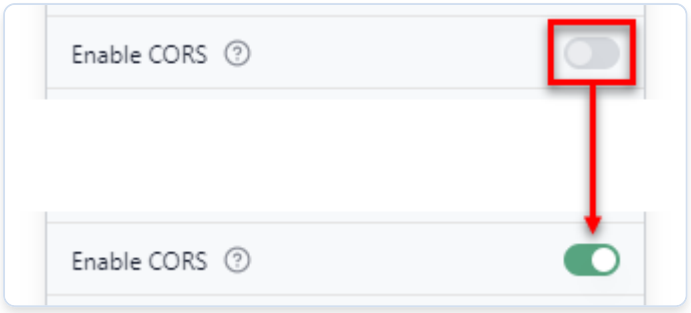
1. In LM Studio, in the sidebar, select **Developer**.
2. Click the **Status** toggle to change the status from **Stopped** to **Running**. You have started the LM Studio server.



3. Click **Settings**.



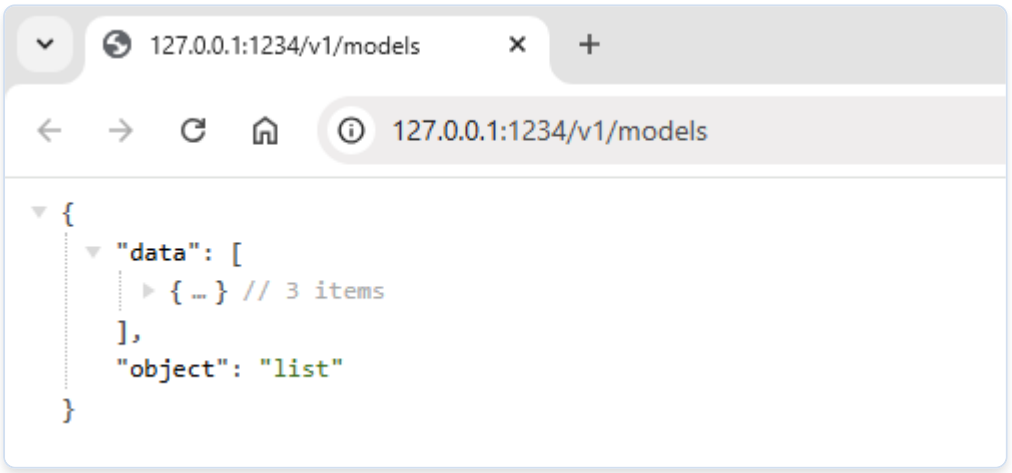
4 Click the **Enable CORS** toggle to enable the setting.



> **INFO**
>
> By default, the LM Studio server only accepts same-origin requests. Since GPT for Excel always has a different origin from the LM Studio server, you must enable cross-origin resource sharing (CORS) for the server.

5 To verify that the `/v1/models` endpoint of the LM Studio server works, open http://127.0.0.1:1234/v1/models.

GPT for Excel uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `data` property listing all currently installed models:



You have started and configured the LM Studio server.

> **WHAT'S NEXT**
>
> You have completed the setup required to access LM Studio from GPT for Excel on the same machine. You can now set `http://127.0.0.1:1234` as the endpoint URL in GPT for Excel.

Connect to Ollama from Excel

> **INFO**
>
> This feature is only available if your space is on the Business or Enterprise plan or pay-as-you-go pricing.

GPT for Excel can connect to an Ollama server to use locally running open-source models. Ollama is an open-source platform that allows you to easily install, run, and serve models from a local machine. You can use any model from the Ollama library.

You can use Ollama models both as Agent models and bulk models, but the server requirements are different:

- **Agent models**: The server must support the OpenAI Responses API (Ollama 0.13.3 or newer) and be publicly accessible over the internet through HTTPS. In addition, the model must support a context size of at least 32K tokens.
- **Bulk models**: Any Ollama server, including a locally running server that is not accessible over the internet.

> **INFO**
>
> Depending on the model you use, the model settings defined in GPT for Excel (temperature (creativity), top-p, frequency penalty, and presence penalty) may not work as expected. If you get bad results, try tweaking the settings.

To use models from an Ollama server:

1. Set up the API endpoint for the Ollama server.
2. Select a model available on the Ollama server.

Set up the API endpoint for an Ollama server

Agent models

Prerequisites

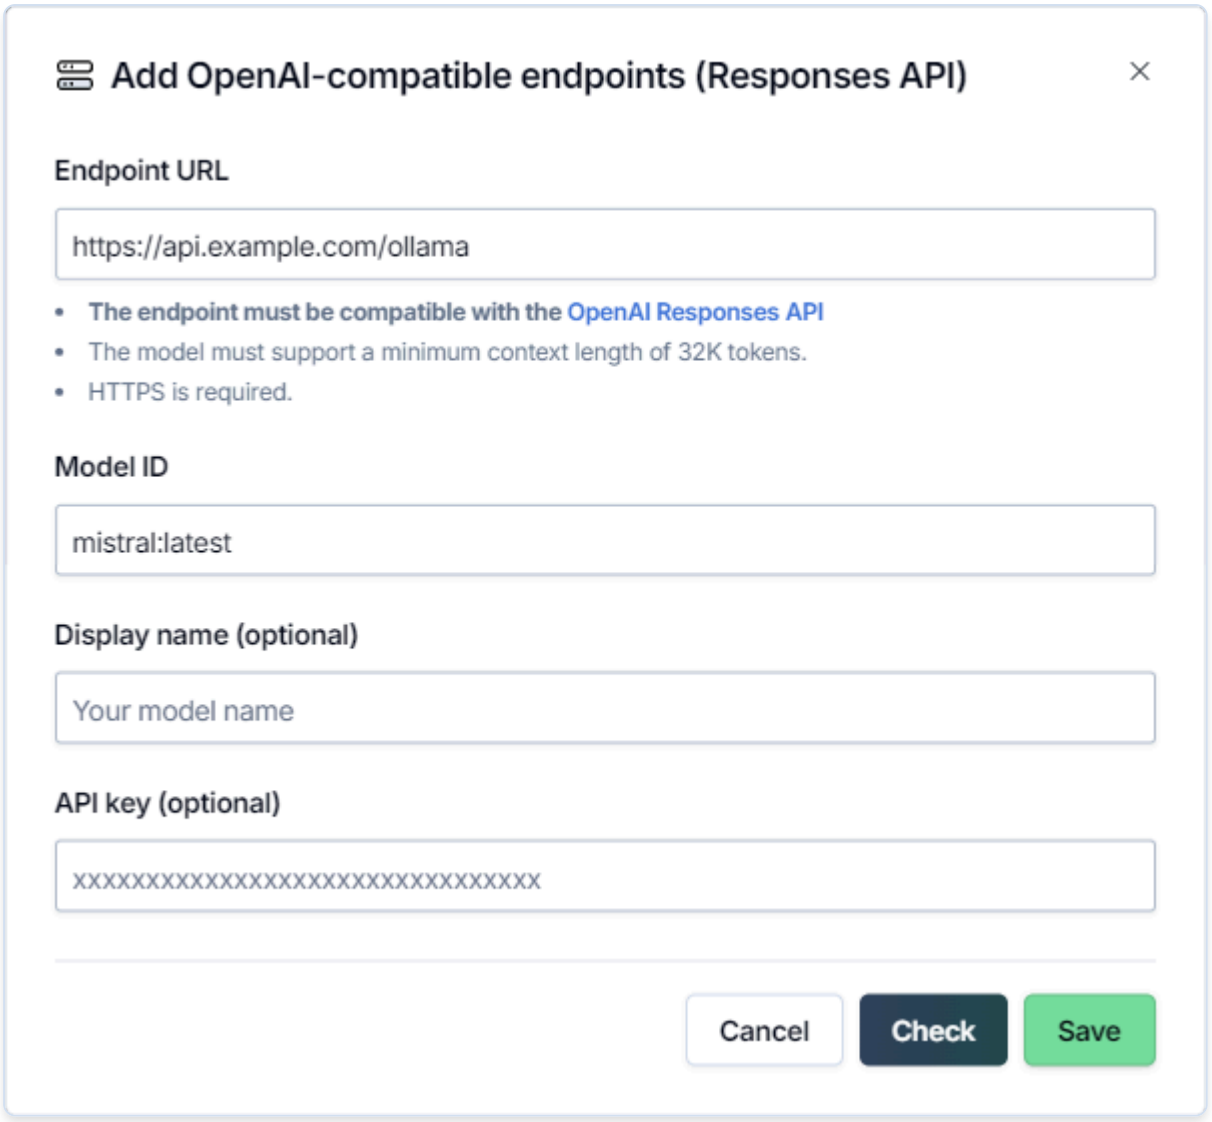
- You have an Ollama server running with one or more models. For instructions on how to set up the server, see:
  - Set up Ollama on macOS
  - Set up Ollama on Windows

- The server supports the following OpenAI-compatible API resource path: `/v1/responses`
- The server is publicly accessible over the internet through HTTPS.
- The model supports a context size of at least 32K tokens.

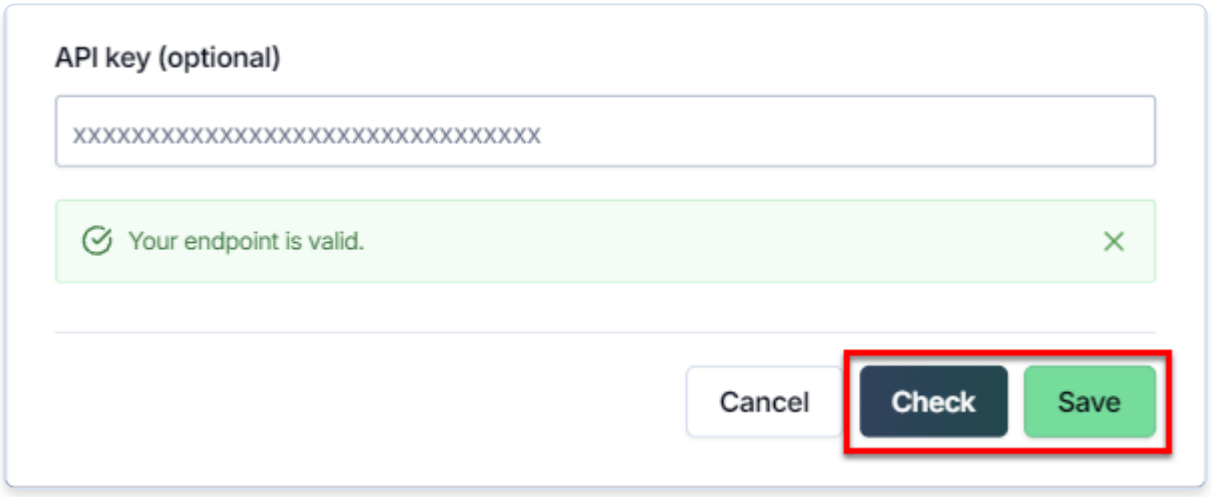
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Agent models** tab.
4. Click **OpenAI-compatible endpoints (Responses API)**.

   > **TIP**
   >
   > If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the public URL of the Ollama server. The server must be publicly accessible over the internet, for example through a reverse proxy or a tunneling service. Setting this up depends on your infrastructure and is not covered by this documentation.
   - **Model ID**: Enter the ID of the model to use, as listed by the `ollama list` command. The model ID is sent to the endpoint in the request payload.
   - **Display name (optional)**: Enter a short name that identifies the model in the Agent model switcher.
   - **API key (optional)**: Leave this field empty. The Ollama server does not require an API key for access.



6 Click **Check** to validate the settings and click **Save**.



You have set up the Ollama server API endpoint. The model is now available in the Agent model switcher and on the models page of the GPT for Work dashboard.

Bulk models

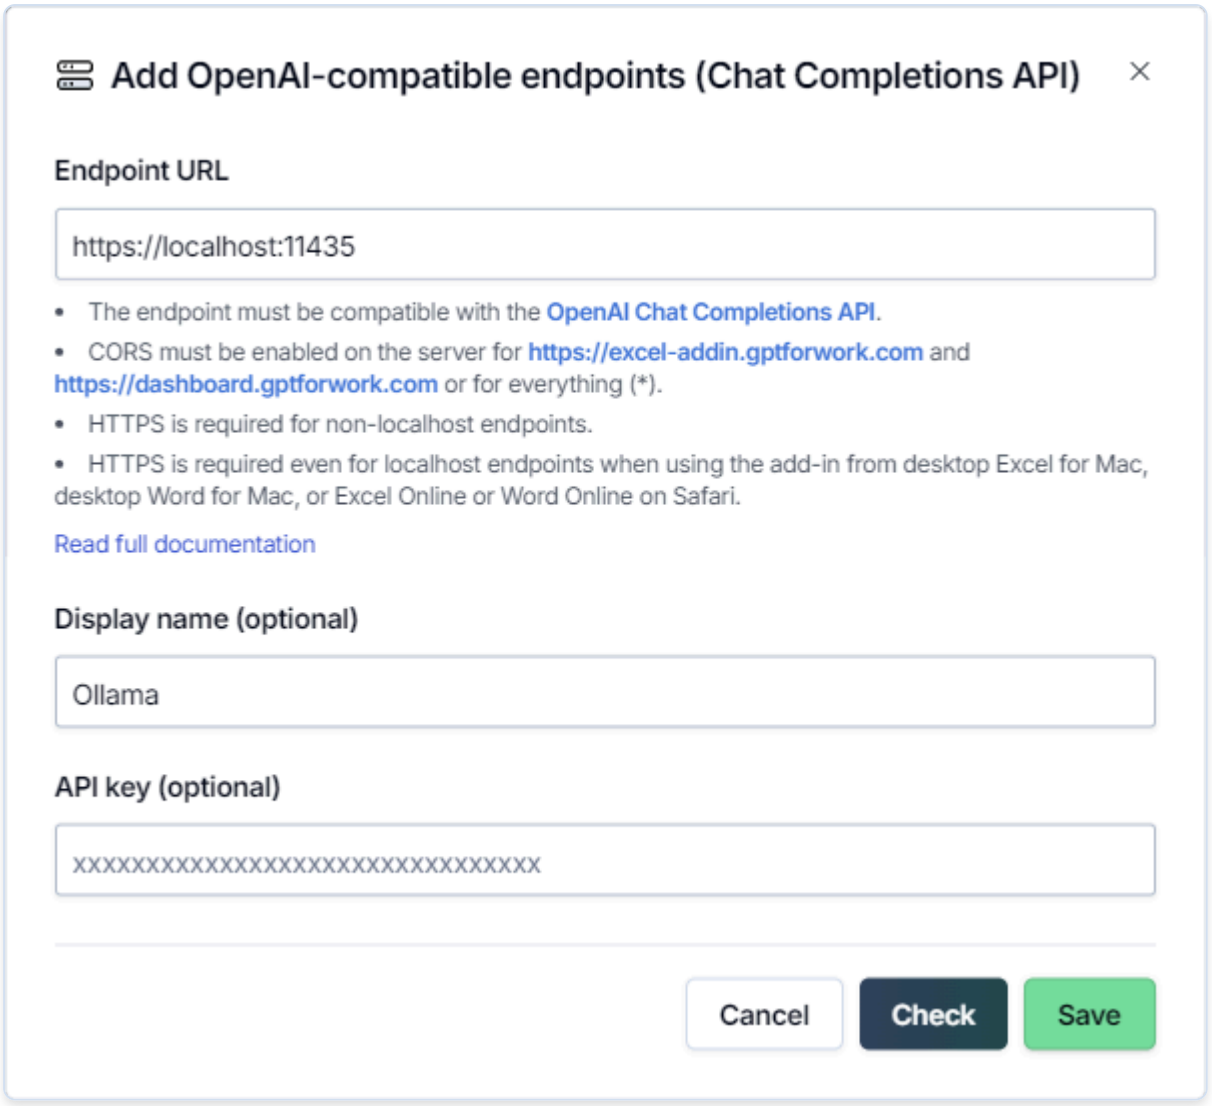
Prerequisites

- You have an Ollama server running with one or more models. For instructions on how to set up the server, see:
  - Set up Ollama on macOS
  - Set up Ollama on Windows

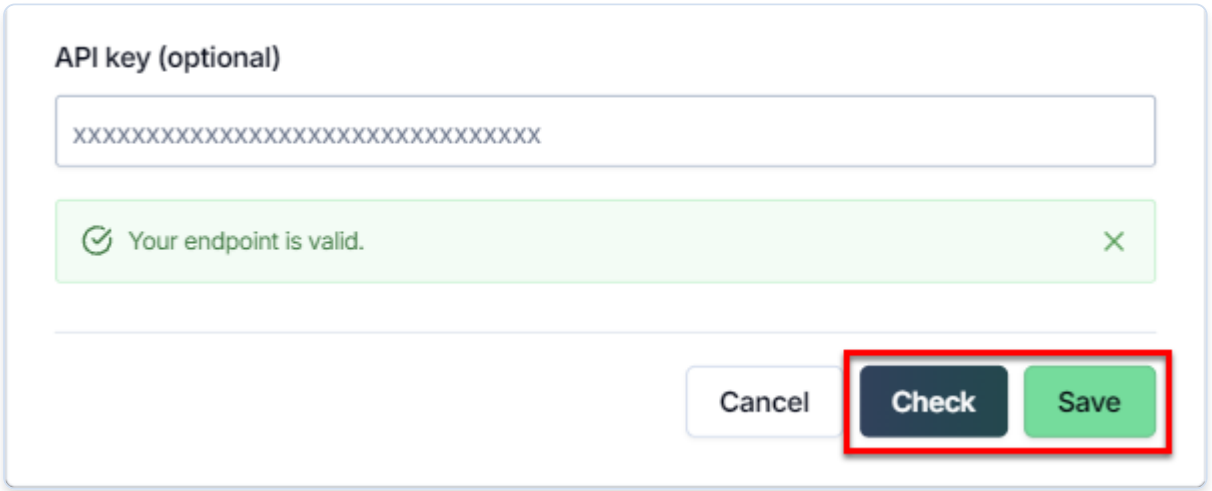
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Bulk models** tab.
4. Click **OpenAI-compatible endpoints (Chat Completions API)**.

   > **TIP**
   >
   > If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the URL of the Ollama server.
   - **Display name (optional)**: Enter a short name that is prefixed to the model names in the model switcher. By default, the model names are prefixed with "custom/".
   - **API key (optional)**: Leave this field empty. The Ollama server does not require an API key for access.



6 Click **Check** to validate the settings and click **Save**.



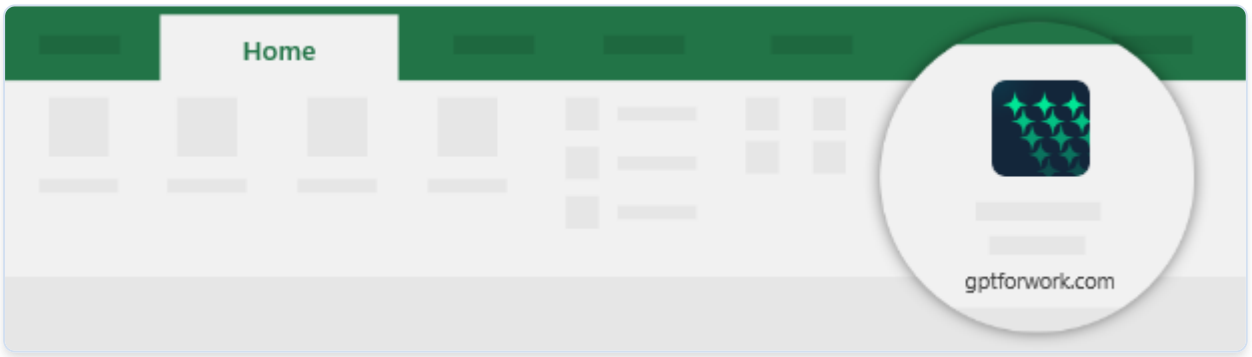
> **TIP**
>
> If you get an error saying that the URL is invalid, make sure that the Ollama server supports:
>
> - **Cross-origin resource sharing (CORS).** See the Ollama server setup instructions for macOS or Windows.
> - **HTTPS (needed with Safari and with Microsoft Office for Mac).** See the Ollama server setup instructions for macOS.

You have set up the Ollama server API endpoint. You can now use all models available through the endpoint in GPT for Excel.

Select a model from the Ollama server

Agent models

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Agent**.
4. Open the model switcher and select a model. You can find the Ollama models at the end of the

list.

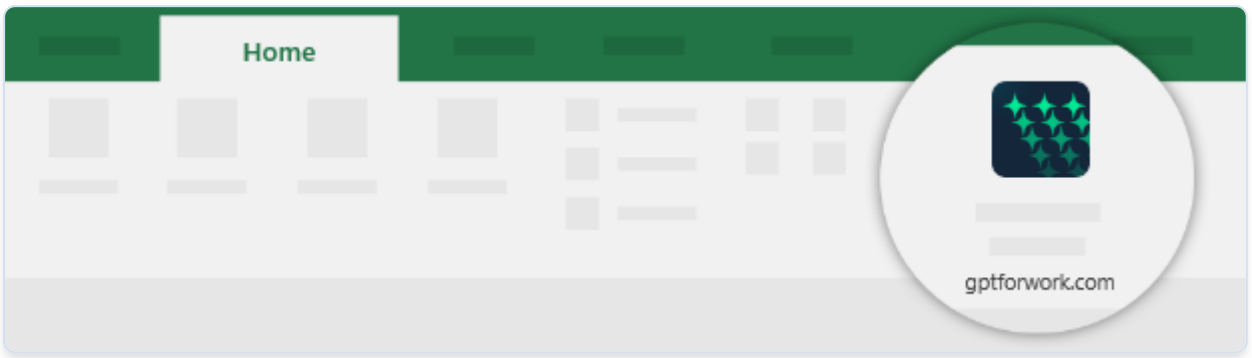


GPT for Excel now uses the selected model as the Agent model.

> **INFO**
>
> Some Agent tools, such as web search, are not available when the Agent uses an endpoint-connected model.

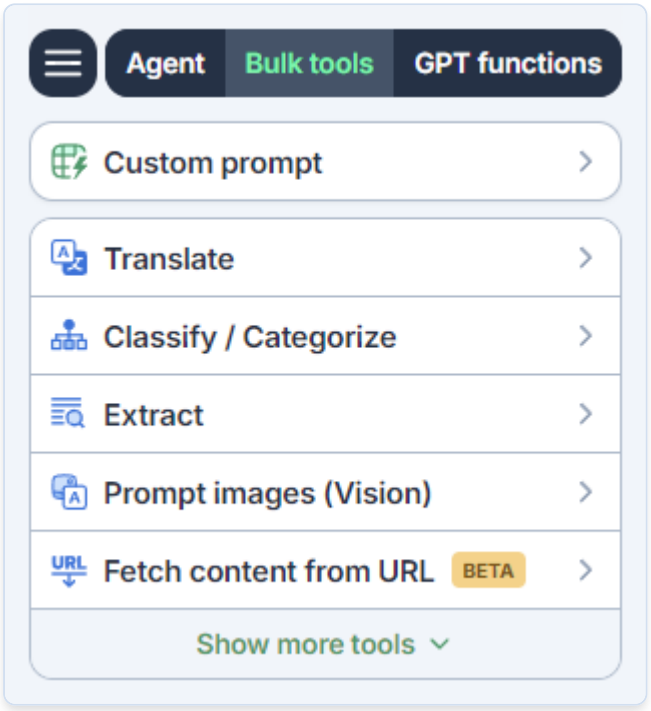
Bulk models

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



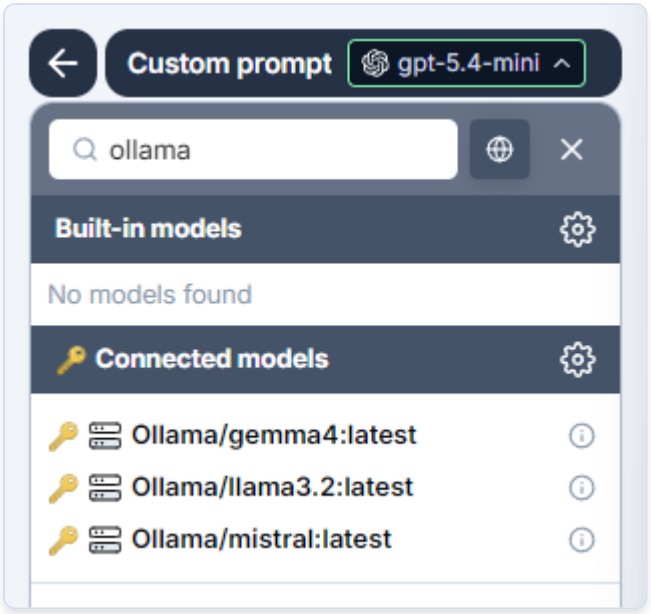
> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Bulk tools** and select one of the tools.



4. Open the model switcher and select a model. You can find the Ollama models under

Connected models, at the end of the list.



> **INFO**
>
> The available models depend on the models available on the Ollama server.

GPT for Excel now uses the selected model to generate responses for bulk tools and GPT functions.

What's next

- Set up an Ollama server on macOS or Windows.

Set up Ollama on macOS

This guide walks you through the main steps of setting up an Ollama server for use in GPT for Excel on macOS.

This guide assumes that you use GPT for Excel on the same machine that hosts Ollama.

Prerequisites

- Local machine with enough processing power and memory to run LLMs (see the Ollama documentation for hardware support information)
- macOS Sonoma (version 14) or newer (see the Ollama documentation for detailed system requirements)
- Mac user account with administrator (sudo) privileges
- Installed software:
  - curl (ships with macOS)
  - Homebrew
  - OpenSSL (ships with macOS)

To set up an Ollama server on macOS:

1. Install the server.
2. Install a model on the server.
3. Enable CORS for the server.
4. (Optional) Enable HTTPS for the server.

Install an Ollama server

1. Download and run the macOS installer. Follow the on-screen instructions to complete the installation.

   The installer starts the Ollama server in the background and sets the server to start automatically on system boot. The installer also installs the Ollama desktop application for easily starting and stopping the server.

   TIP

   You can manually start a new instance of the Ollama server in a terminal by running `ollama serve`. However, to avoid port conflicts, make sure no other Ollama server instances are running at the same time.

2. Open Terminal and verify that the server is running:

```
curl http://localhost:11434
```

Response if the server is running:

```
Ollama is running
```

You have installed the Ollama server.

For more information about installing and managing the server on macOS, see the Ollama documentation.

Install a model on the Ollama server

1. Install a model from the Ollama library:

```
ollama pull <model-name>
```

For example, to install the Meta Llama 3.2 3B model:

```
ollama pull llama3.2
```

After the installation completes, the model is available for prompting on the Ollama server.

2. Verify that the model works:

```
curl http://localhost:11434/api/generate -d '{ "model": "llama3.2",
"prompt": "Who are you?", "stream": false }'
```

The above request asks the server to generate a response for the specified prompt with the specified model and to return the response in a single reply. If the model works, the server returns a JSON object containing the response and metadata about the response.

You have installed the model on the Ollama server. For more information about working with models, see the Ollama documentation.

Enable CORS for the Ollama server

By default, the Ollama server only accepts same-origin requests. Since GPT for Excel always has a different origin from the Ollama server, you must enable cross-origin resource sharing (CORS) for the server using the `OLLAMA_ORIGINS` environment variable.

To enable CORS for the Ollama server:

1. Set `OLLAMA_ORIGINS` with the origins that are allowed to access the server:

```
# Set a single origin
launchctl setenv OLLAMA_ORIGINS "<ORIGIN>"

# Set multiple origins
launchctl setenv OLLAMA_ORIGINS "<ORIGIN_1>, <ORIGIN_2>, ..."
```

   For example:

   - Allow any origin to make requests to the server:

```
launchctl setenv OLLAMA_ORIGINS "*"
```

   - Allow only GPT for Excel to make requests to the server:

```
launchctl setenv OLLAMA_ORIGINS "https://excel-addin.gptforwork.com"
```

2. Restart the server for the variable to take effect.
3. Verify that the `/api/tags` endpoint of the server works:

```
curl http://localhost:11434/api/tags
```

   GPT for Excel uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `models` property listing all currently installed models:

```
{
  "models": [
    {
      ...
    }
  ]
}
```

You have enabled CORS for the Ollama server.

> **WHAT'S NEXT**
>
> You have completed the minimum setup required to access the Ollama server from GPT for Excel on the same machine. You can now set `http://localhost:11434` as the endpoint URL in GPT for Excel provided the add-in is not running on Safari or on Microsoft Excel for Mac.
>
> If you use GPT for Excel on Safari or on Excel for Mac, enable HTTPS for the Ollama server.

Enable HTTPS for the Ollama server

The Ollama server uses HTTP to serve models, while GPT for Excel runs on HTTPS. By default, therefore, any request from GPT for Excel to the server is a mixed-content request (HTTP vs. HTTPS). Modern web browsers do not allow mixed content, as a rule. The only exceptions are mixed-content requests to `http://127.0.0.1` and `http://localhost`, which most browsers treat as safe and therefore allow. Safari, however, blocks mixed-content requests even on the current machine.

To avoid mixed content, you must make the Ollama server accessible over HTTPS if you use GPT for Excel from:

- Excel for the web on Safari
- Excel for Mac (uses Safari)

You make the Ollama server accessible over HTTPS by setting up a reverse proxy that hides the server behind an HTTPS interface. This guide uses nginx to set up the reverse proxy.

> **TIP**
>
> You can also use Cloudflare Tunnel, ngrok, or a similar cloud-based tunneling service to set up HTTPS with minimal configuration. Note that with a cloud-based service your traffic will be routed through an external service over the internet.

To set up the reverse proxy:

1. Install nginx:

```
brew install nginx
```

2. Create a self-signed SSL certificate for the Ollama server:
    1. Change to the nginx configuration directory (varies based on your Mac version):

```
# Mac with an Apple silicon processor
cd /opt/homebrew/etc/nginx

# Mac with an Intel processor
cd /usr/local/etc/nginx
```

    2. Generate the certificate:

```
openssl req \
  -x509 \
  -newkey rsa:2048 \
  -nodes \
```

```
  -sha256 \
  -days 365 \
  -keyout ollama.key \
  -out ollama.crt \
  -subj '/CN=localhost' \
  -extensions extensions \
  -config <(printf "
[dn]\nCN=localhost\n[req]\ndistinguished_name=dn\n[extensions]\nsubjectAl
```

The command generates two files in the current directory:

- `ollama.crt`: Public self-signed certificate
- `ollama.key`: Private key used to sign the certificate

3. Add the certificate to trusted root certificates:

```
sudo security add-trusted-cert -d -r trustRoot -k
/Library/Keychains/System.keychain ./ollama.crt
```

3. Create a site configuration for the Ollama server:

1. Open the `nginx.conf` file in your preferred text editor. The file is in the nginx configuration directory:

```
# Mac with an Apple silicon processor
/opt/homebrew/etc/nginx/nginx.conf

# Mac with an Intel processor
/usr/local/etc/nginx/nginx.conf
```

2. Replace the content of the file with the following configuration:

```
events {}

http {
    server {
        listen       11435 ssl;
        server_name  localhost;

        http2 on;

        ssl_certificate      ollama.crt;
        ssl_certificate_key  ollama.key;

        location / {
            proxy_pass  http://localhost:11434;

            # Keep connection to backend alive.
```

```
            proxy_http_version  1.1;
            proxy_set_header    Connection '';
        }
    }
}
```

The configuration uses port `11435` for HTTPS. Requests to `https://localhost:11435` are forwarded to the Ollama server running at `http://localhost:11434`.

3. Save and close the file.
4. Test the nginx configuration:

```
sudo nginx -t
```

Response if the configuration is valid:

```
nginx: the configuration file /opt/homebrew/etc/nginx/nginx.conf syntax
is ok
nginx: configuration file /opt/homebrew/etc/nginx/nginx.conf test is
successful
```

5. Start nginx:

```
sudo nginx
```

4. Verify that the `/api/tags` endpoint of the Ollama server works over the HTTPS connection:

```
curl -k https://localhost:11435/api/tags
```

GPT for Excel uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `models` property listing all currently installed models:

```
{
  "models": [
    {
      ...
    }
  ]
}
```

You have enabled HTTPS for the Ollama server.

WHAT'S NEXT

You have completed the setup required to access the Ollama server from GPT for Excel on the same machine. You can now set `https://localhost:11435` as the endpoint URL in GPT for Excel, including when the add-in is running on Safari or on Microsoft Excel for Mac.

Set up Ollama on Windows

This guide walks you through the main steps of setting up an Ollama server for use in GPT for Excel on Windows.

This guide assumes that you use GPT for Excel on the same machine that hosts Ollama.

Prerequisites

- Local machine with enough processing power and memory to run LLMs (see the Ollama documentation for hardware support information)
- Windows 10 22H2 or newer (see the Ollama documentation for detailed system requirements)
- Windows user account with administrator privileges
- Installed software:
  - curl (ships with Windows)
  - PowerShell 7

To set up an Ollama server on Windows:

1. Install the server.
2. Install a model on the server.
3. Enable CORS for the server.

Install an Ollama server

1. Download and run the Windows installer. Follow the on-screen instructions to complete the installation.

   The installer starts the Ollama server in the background and sets the server to start automatically on system boot. The installer also installs the Ollama desktop application for easily starting and stopping the server.

   TIP

   You can manually start a new instance of the Ollama server in a terminal by running `ollama serve`. However, to avoid port conflicts, make sure no other Ollama server instances are running at the same time.

2. Open PowerShell as an administrator and verify that the server is running:

```
curl http://localhost:11434
```

Response if the server is running:

```
Ollama is running
```

> **NOTE**
>
> Make sure you're using curl and not an aliased command, such as `Invoke-WebRequest`:
>
> ```
> Remove-Item alias:curl
> ```
>
> If `curl` has an alias, the command removes it for the current terminal session, ensuring that subsequent `curl` commands in this guide work as expected.

You have installed the Ollama server.

For more information about installing and managing the server on Windows, see the Ollama documentation.

Install a model on the Ollama server

1. Install a model from the Ollama library:

   ```
   ollama pull <model-name>
   ```

   For example, to install the Meta Llama 3.2 3B model:

   ```
   ollama pull llama3.2
   ```

   After the installation completes, the model is available for prompting on the Ollama server.

2. Verify that the model works:

   ```
   curl http://localhost:11434/api/generate -d '{ "model": "llama3.2",
   "prompt": "Who are you?", "stream": false }'
   ```

   The above request asks the server to generate a response for the specified prompt with the specified model and to return the response in a single reply. If the model works, the server returns a JSON object containing the response and metadata about the response.

You have installed the model on the Ollama server. For more information about working with models, see the Ollama documentation.

Enable CORS for the Ollama server

By default, the Ollama server only accepts same-origin requests. Since GPT for Excel always has a different origin from the Ollama server, you must enable cross-origin resource sharing (CORS) for the server using the `OLLAMA_ORIGINS` environment variable.

To enable CORS for the Ollama server:

1. Set `OLLAMA_ORIGINS` with the origins that are allowed to access the server:

```
# Set a single origin
setx OLLAMA_ORIGINS "<ORIGIN>"

# Set multiple origins
setx OLLAMA_ORIGINS "<ORIGIN_1>, <ORIGIN_2>, ..."
```

   For example:

   - Allow any origin to make requests to the server:

```
setx OLLAMA_ORIGINS "*"
```

   - Allow only GPT for Excel to make requests to the server:

```
setx OLLAMA_ORIGINS "https://excel-addin.gptforwork.com"
```

   INFO

   By default, the `setx` command sets a user variable. If you want to set a system variable, which applies to all users of the current machine, run the command (as an administrator) with the `/m` parameter. For example:

```
setx OLLAMA_ORIGINS "*" /m
```

2. Restart the server for the variable to take effect.
3. Verify that the `/api/tags` endpoint of the server works:

```
curl http://localhost:11434/api/tags
```

   GPT for Excel uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `models` property listing all currently installed models:

```
{
  "models": [
```

```
    {
      ...
    }
  ]
}
```

You have enabled CORS for the Ollama server.

WHAT'S NEXT

You have completed the setup required to access the Ollama server from GPT for Excel on the same machine. You can now set `http://localhost:11434` as the endpoint URL in GPT for Excel.

Remove account access from GPT for Excel

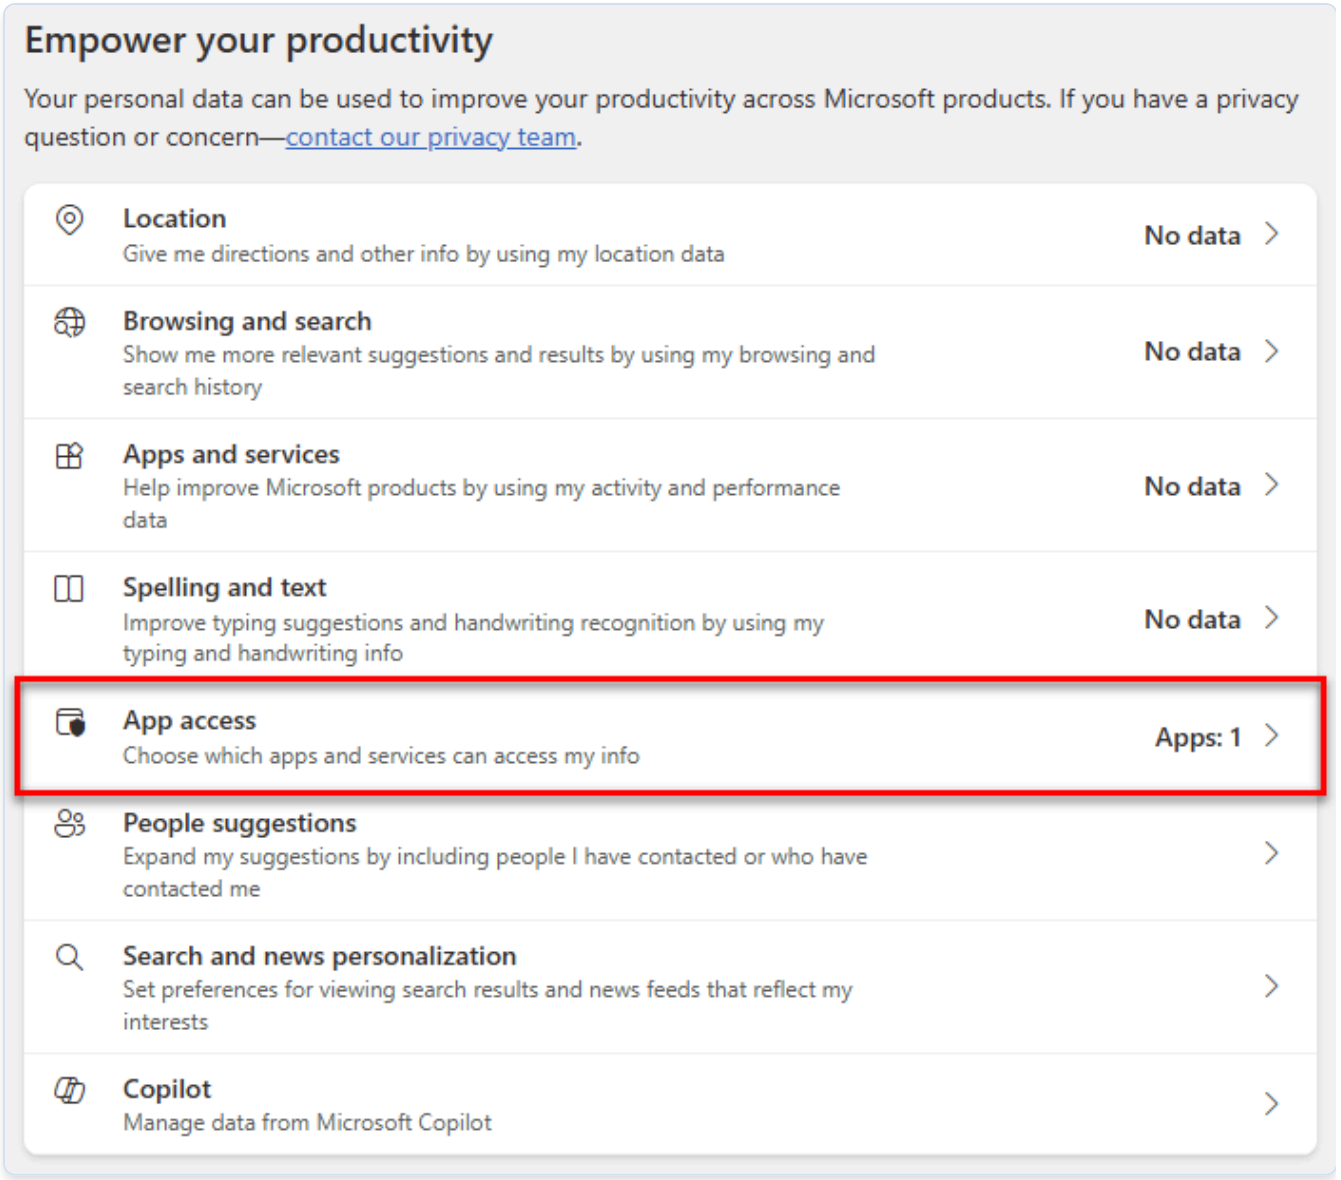
You can remove the **GPT for Excel Word** add-in's access to your Microsoft or Google account at any time. These instructions apply to personal accounts only. If you have a school or work account, contact your organization administrator about add-in access.

> **WARNING**
>
> Removing account access from GPT for Excel removes access from the **GPT for Excel Word** add-in, which also removes access from GPT for Word.

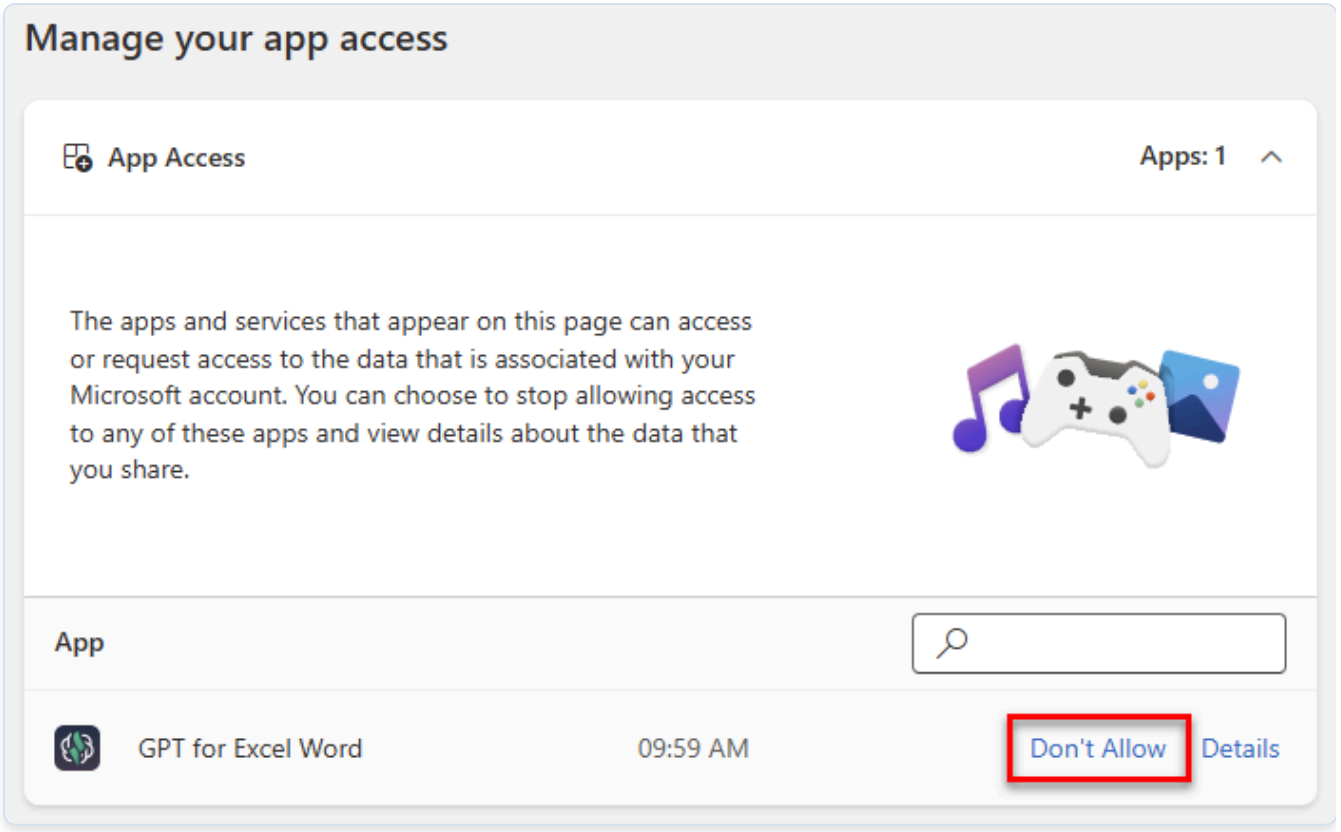
Microsoft

1. In your browser, log in to the Microsoft account you use in GPT for Excel.
2. Open the privacy dashboard for your Microsoft account.

3 In the **Empower your productivity** section, select **App access**.



4 In the **Manage your app access** section, click **Don't Allow** for GPT for Excel Word.



The **GPT for Excel Word** add-in no longer has access to your Microsoft account.

> **INFO**
>
> Removing account access does not uninstall the add-in from Microsoft Excel. The next time you open the add-in, Excel displays a welcome screen in the sidebar, after which you can sign in again and grant access again.

Uninstall GPT for Excel

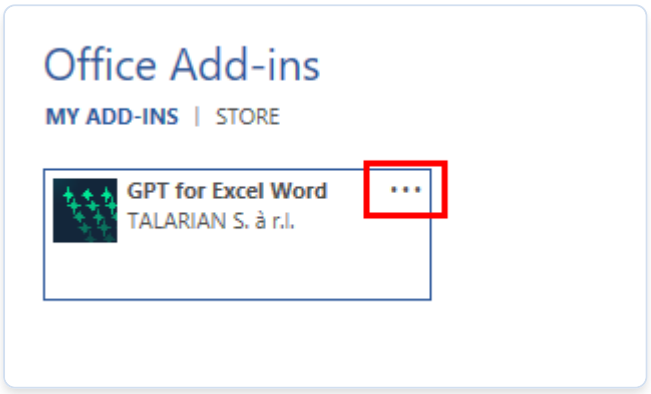
PREREQUISITES

Remove all GPT formulas from all of your workbooks before uninstalling the add-in. Leaving any GPT formulas in a workbook causes Excel to reinstall the add-in automatically when you open the workbook.

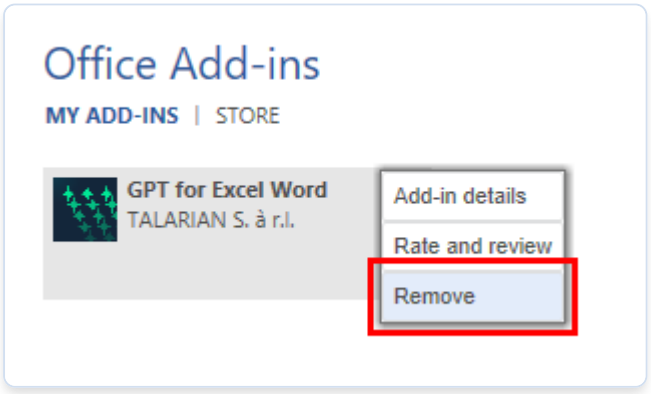
WARNING

Uninstalling GPT for Excel removes the **GPT for Excel Word** add-in from both Microsoft Excel and Microsoft Word.

1. Open Excel.
2. On the ribbon, select **Home > Add-ins > More Add-ins**. The **Office Add-ins** dialog opens.
3. Select the **My Add-ins** tab.
4. Hover over **GPT for Excel Word** and open its **Options** menu (⋯).



5. Select **Remove**.



A confirmation pop-up opens.

> **TIP**
>
> If the **Remove** option is not available in the menu, clear the Microsoft Office cache:
>
> 1. Close all Office applications.
> 2. Follow Microsoft's instructions for clearing the cache.
> 3. Open Excel and try removing the add-in again.

6. Click **Remove** to confirm.

The **GPT for Excel Word** add-in is removed from your Microsoft account. If you ever need to use the add-in again, follow the installation process.

> **INFO**
>
> Removing the add-in does not remove its app access permissions from your Microsoft or Google account. You can remove access manually.

Agent in Excel

The **Agent** is the chat interface for spreadsheet work in GPT for Excel. Describe what you need in plain language — from everyday spreadsheet tasks to row-by-row bulk processing across large datasets — and the Agent reads your data, determines the required steps, and executes them directly in your spreadsheet.

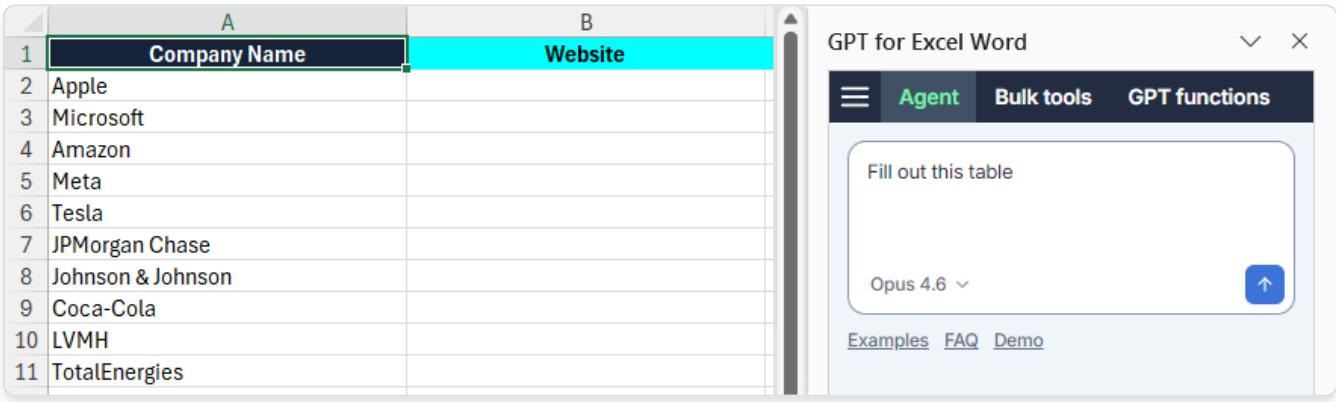
Learn more about what the Agent can and cannot do.

Watch video on YouTube

How to use the Agent

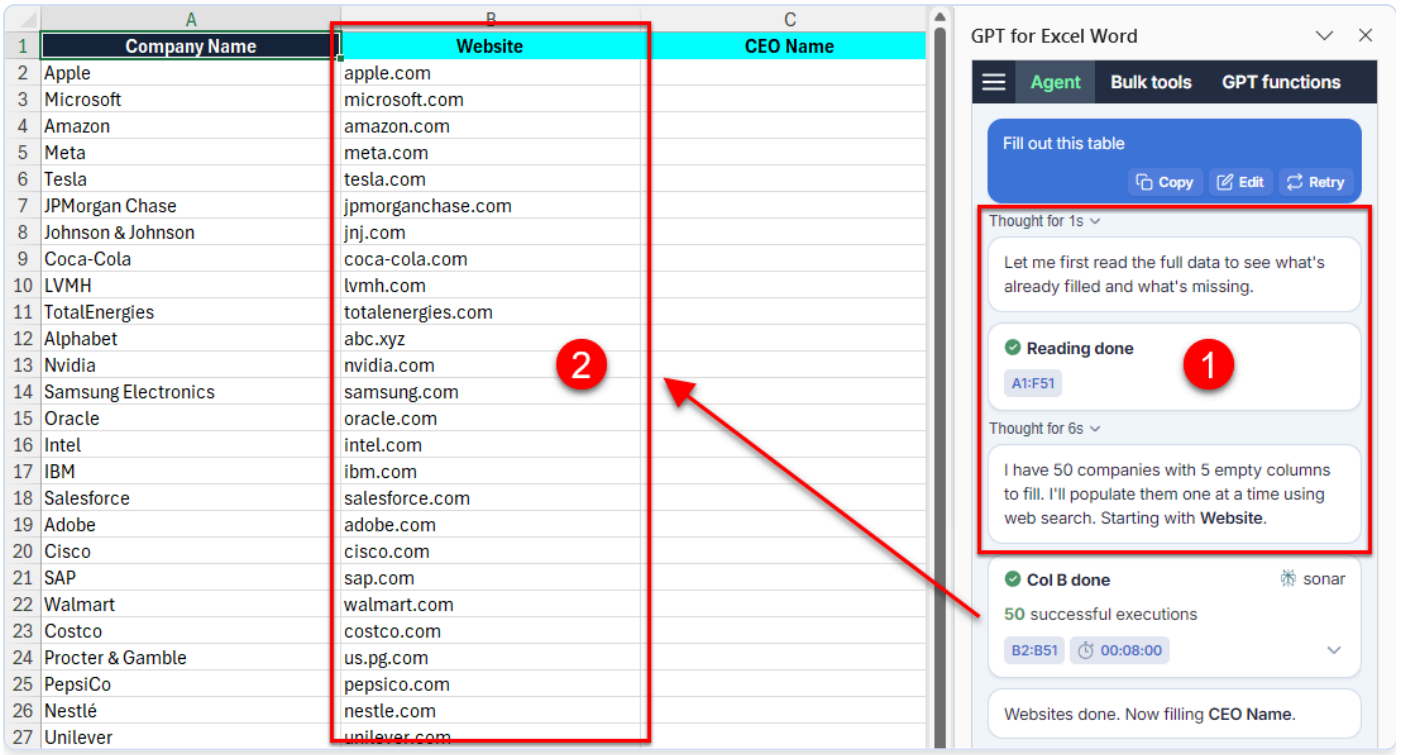
To use the Agent in GPT for Excel:

1. Open your spreadsheet.
2. Open GPT for Excel.
3. In the sidebar, in the **Agent** tab, enter your prompt and click the send button.



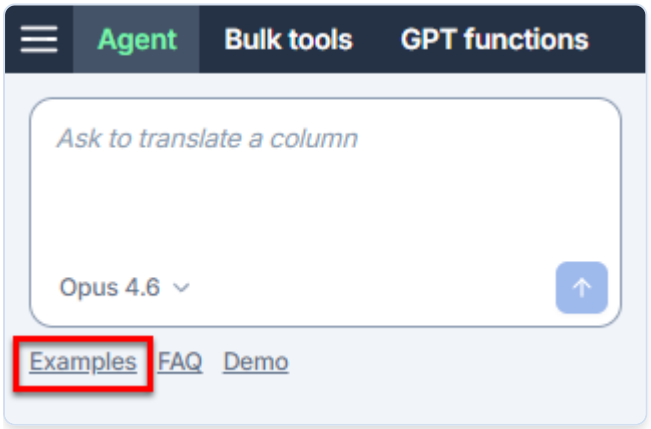
Next, the Agent:

1. Analyzes your request and drafts a plan.
2. Processes the data in the current sheet according to the plan.



Examples

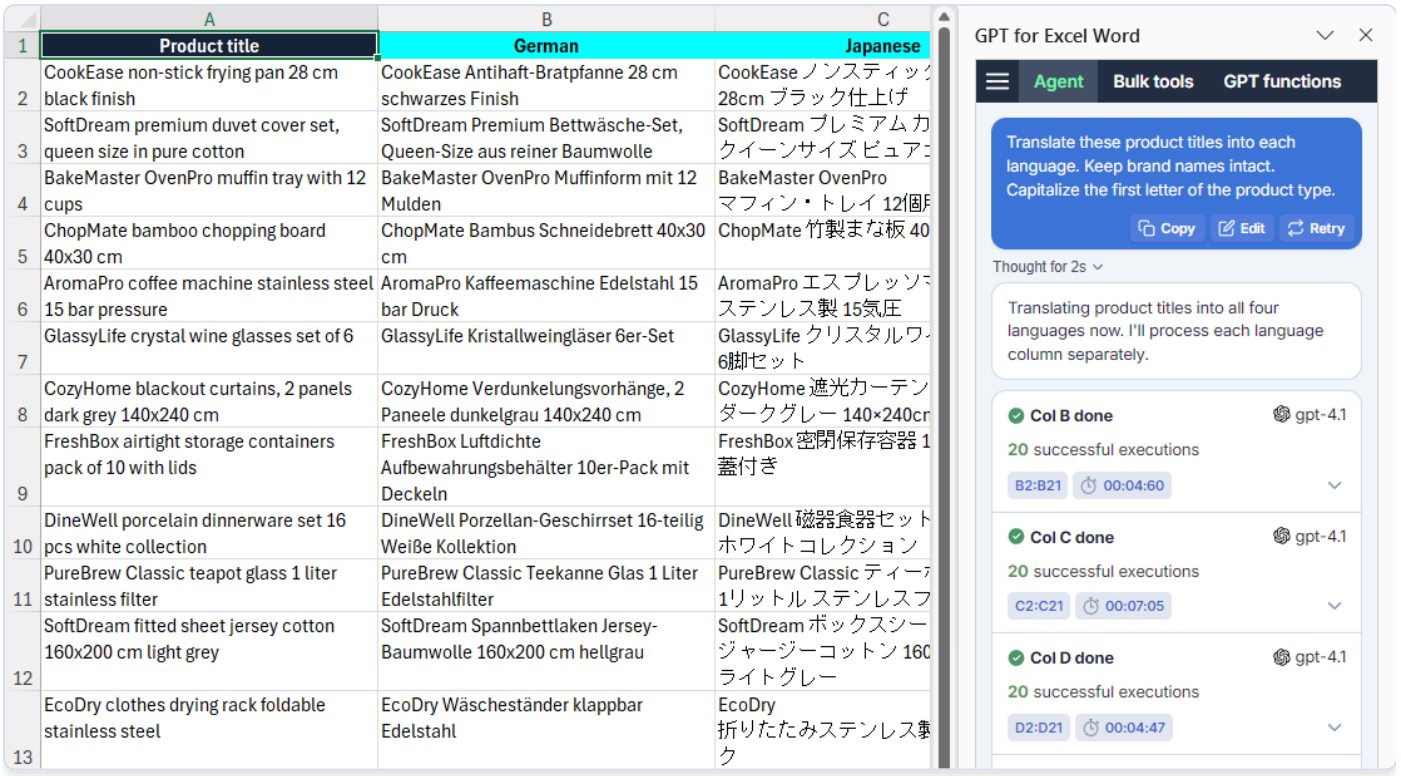
Find out what the Agent can do with the examples below. You can try these examples, and many more, directly from the sidebar.



Translate into multiple languages

Translate spreadsheet content into multiple languages in bulk.

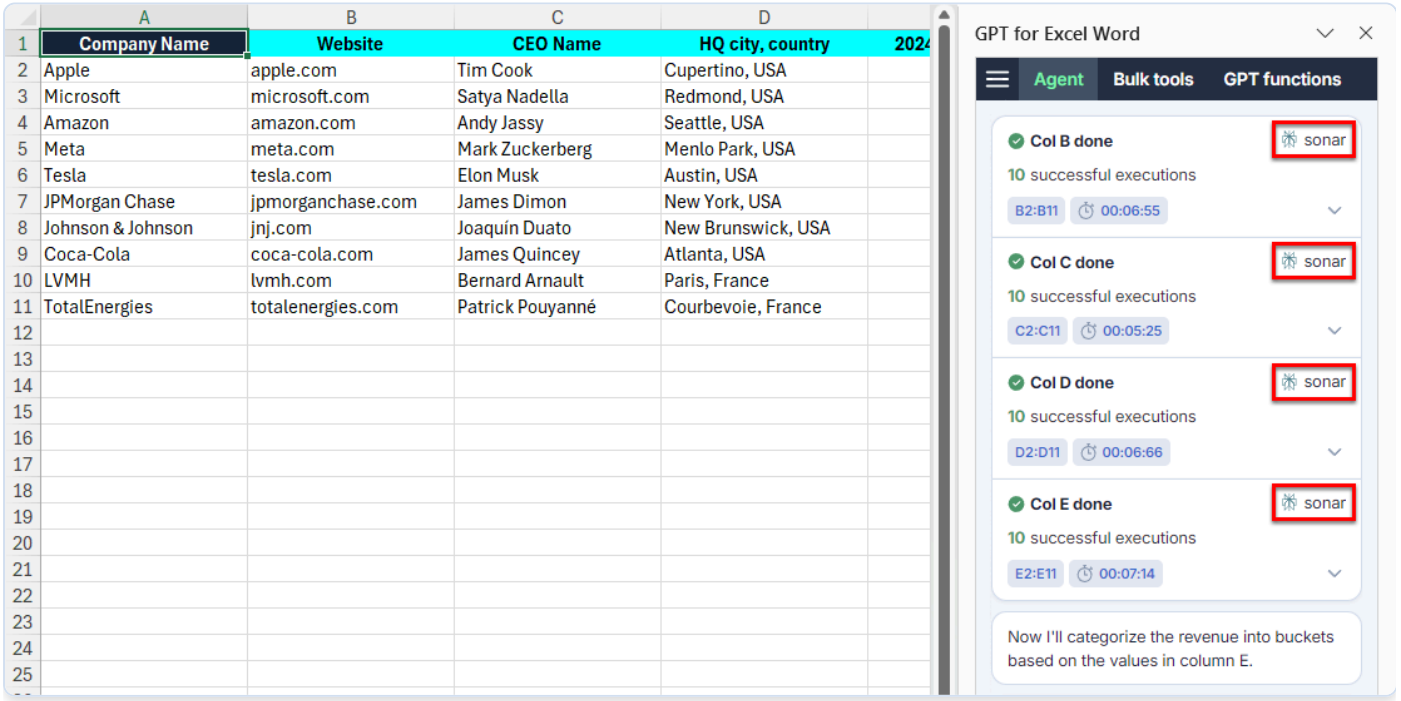
```
Translate these product titles into each language. Keep brand names intact.
Capitalize the first letter of the product type.
```



Search the web

When necessary, the Agent searches the web to find the information you need. Learn more about the models used by the Agent.

```
Fill out this table
```

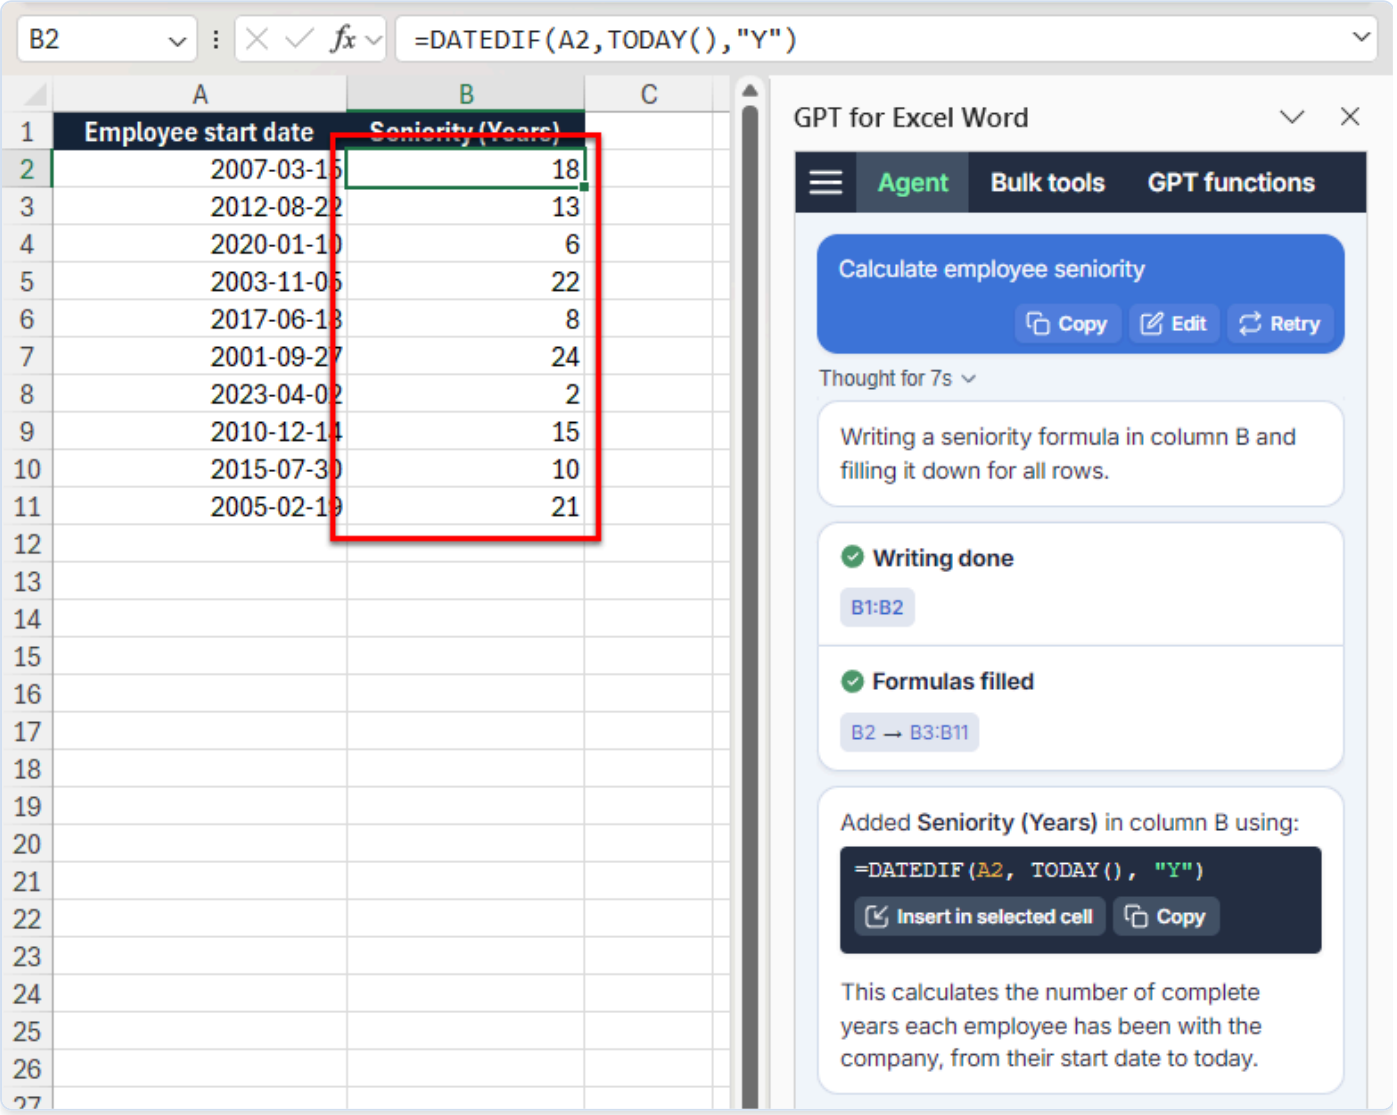


> **TIP**
>
> To scrape specific URLs, use the Fetch content from URL bulk tool.

Generate and autofill formulas

Describe what you need in natural language, and the Agent generates the appropriate Excel formula, inserts it in the starting cell, and automatically fills the formula across the range.

```
Calculate employee seniority
```



What the Agent can and cannot do

The Agent excels at both everyday spreadsheet tasks and row-by-row bulk processing across large datasets.

The Agent can:

- Add, remove, or reorganize columns and rows.
- Create charts, pivot tables, and dashboards from your data.
- Create, fix, and explain spreadsheet formulas using natural language.
- Generate lists and tables from scratch based on a prompt.
- Process data in bulk, including generating or rewriting content, translating text, categorizing or tagging data, cleaning and preparing data, and searching the web.
- Style your spreadsheet, including colors, text formatting, and conditional formatting.
- Summarize large cell ranges.

The Agent cannot yet:

- Scrape specific URLs. (Use the Fetch content from URL bulk tool instead.)
- Take documents as input. (Use the Fetch content from URL bulk tool instead.)
- Take images as input. (Use the Prompt images (Vision) bulk tool instead.)

FAQ

What models does the Agent use?

Agent models handle chat messages, create formulas, and execute complex tasks by breaking them down into smaller tasks and delegating these to bulk models.

For a list of the models available for the Agent, see AI providers and models supported by GPT for Work.

Can I change the models used by the Agent?

Yes, you can change the models used by the Agent:

- Select the Agent model in the GPT for Excel sidebar.
- Set the default bulk models in the GPT for Work dashboard.

What's next

- Select the models used by the Agent.
- Try out example Agent use cases.

Select the models for the Agent in Excel

In GPT for Excel, the Agent uses two types of models:

- **Agent models** handle chat messages, create formulas, and execute complex tasks by breaking them down into smaller tasks and delegating these to bulk models.
- **Bulk models** handle bulk tasks and web search tasks. The Agent uses bulk models when needed. The Agent can only use the bulk models you set in the GPT for Work dashboard.

Recommended models

For information about recommended models for the Agent, see Recommended models.

Select the Agent model

You can select the Agent model in the GPT for Excel sidebar. The models available depend on the models enabled for your space in the GPT for Work dashboard.

In addition to built-in models, you can select connected models in the model switcher. To make a connected model available, set an API key or set up an OpenAI-compatible API endpoint.

1. In the sidebar, select **Agent**.
2. Open the model switcher and select the Agent model.



The Agent now uses the selected model.

Set the default bulk models

Prerequisites

You are the owner or an admin of the space.

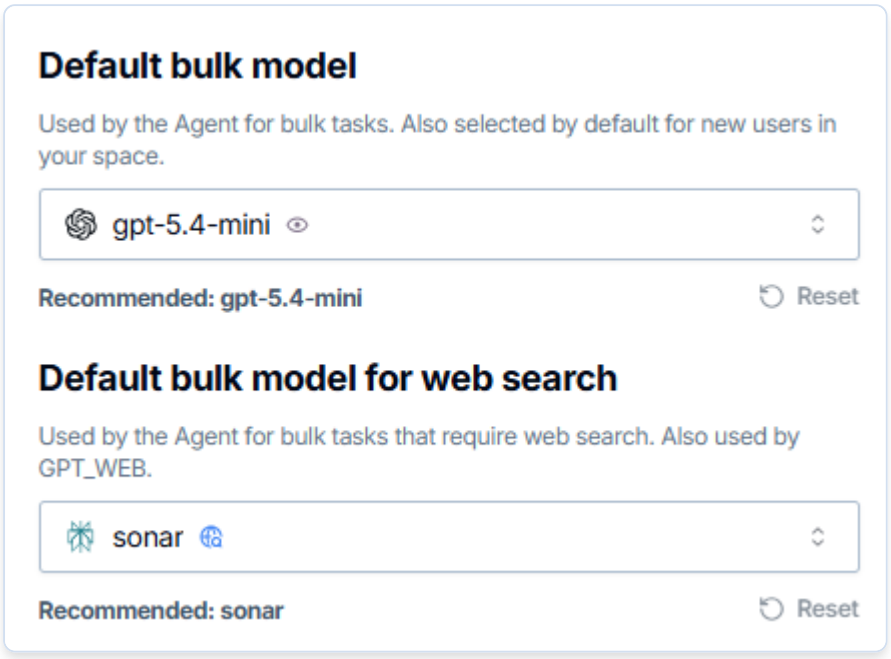
You can set two default models for bulk tasks:

- The **Default bulk model** is used by the Agent for bulk tasks. For new users in your space, this model is also selected by default for bulk tools and GPT functions.

- The **Default bulk model for web search** is used by the Agent for bulk tasks that require web search. This model is also used by the **GPT_WEB** function.

You can only select enabled bulk models.

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Bulk models**.
3. In the **Default bulk model** and **Default bulk model for web search** dropdowns, select the default bulk models for your space.



The Agent now uses the selected models for bulk tasks and web searches.

What's next

Learn more about models:

- AI providers and models supported by GPT for Work
- LiveBench (a good model benchmark)

Bulk tools in Excel

Run any type of operation on an entire spreadsheet column at once, or on a specific section of your data.

Get started

Learn the basics of using bulk tools in GPT for Excel.

Bulk tools reference

Learn how to use all the bulk tools available in GPT for Excel.

Select a model

Select the model used by bulk tools in GPT for Excel.

Select rows to process

Let GPT for Excel process only the relevant rows in your spreadsheet, or target a specific section of your data.

Search the web

Learn how to use web search models in bulk tools to get up-to-date information from the web.

Use images in prompts

Include images in bulk tool prompts in Microsoft Excel.

Configure behavior

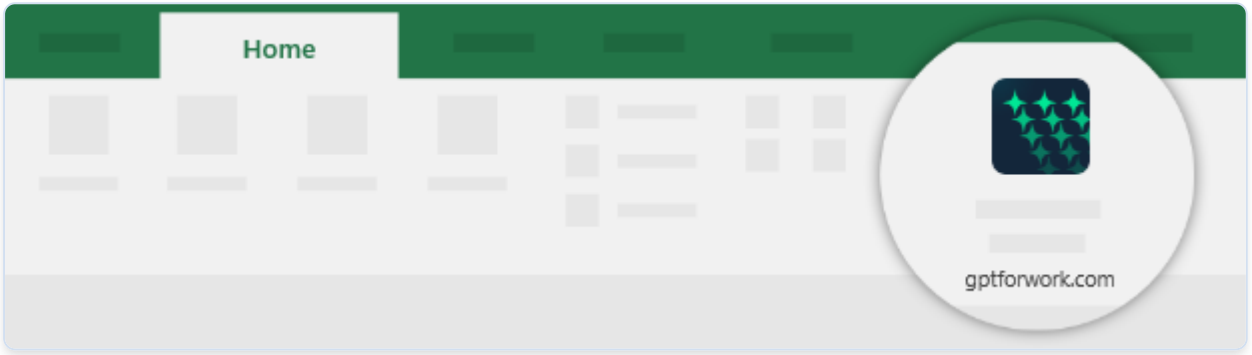
Configure model settings to customize how bulk tools behave in Microsoft Excel.

Get started with bulk tools in Excel

Bulk tools allow you to run prompts on specific spreadsheet columns at once without writing any formulas. You configure and run bulk tools from the add-in sidebar. You can also easily reuse past prompts and configurations with the **History** feature.

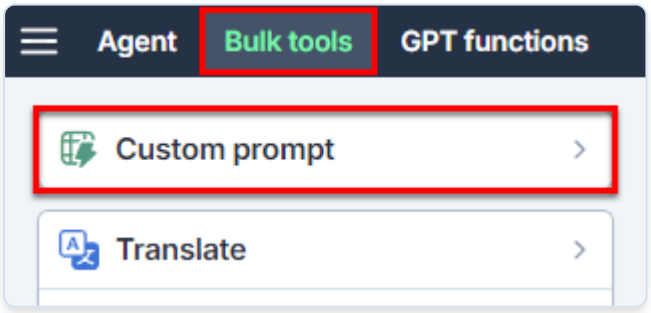
To get started with bulk tools, let's generate some text with the **Custom prompt** tool:

1. Download our bulk tools starter template, and open it in Excel as a new workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Bulk tools** and select **Custom prompt**.



4. Define the following settings:
   - **Header row**: Select **5**. The bulk tool will process all rows with content below this row.
   - **Prompt to run for each row**: Enter the following prompt for the AI:

```
Write a short SEO-optimized product description for a bike based on its specifications: {{Specifications}}
```

   - **Put results in column**: Select **C**.
   - Leave all other settings to their default values.

5. Run the bulk tool:
    1. Select **All rows**
    2. Click **Run all rows**.

    The bulk tool executes the prompt for every row with content, starting from row 6, and puts the responses generated by the AI in column C.

The following video shows the entire process.

What's next

- Discover all bulk tools available in GPT for Excel.
- Select the model that best fits your needs.
- Configure the model settings to customize how the currently selected model behaves.

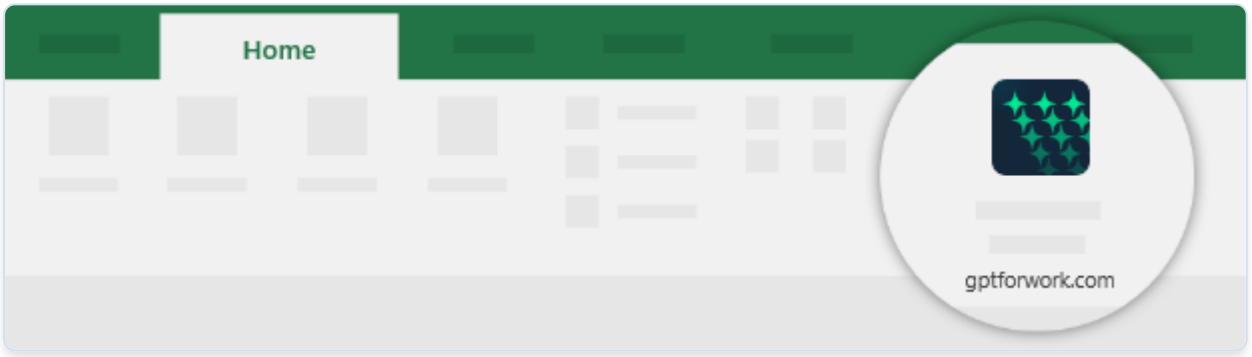
Bulk tools reference for Excel

Bulk tools allow you to run prompts on specific spreadsheet columns at once without writing any formulas. You configure and run bulk tools from the add-in sidebar. You can also easily reuse past prompts and configurations with the **History** feature.

This guide shows you how to use every bulk tool available in GPT for Excel.

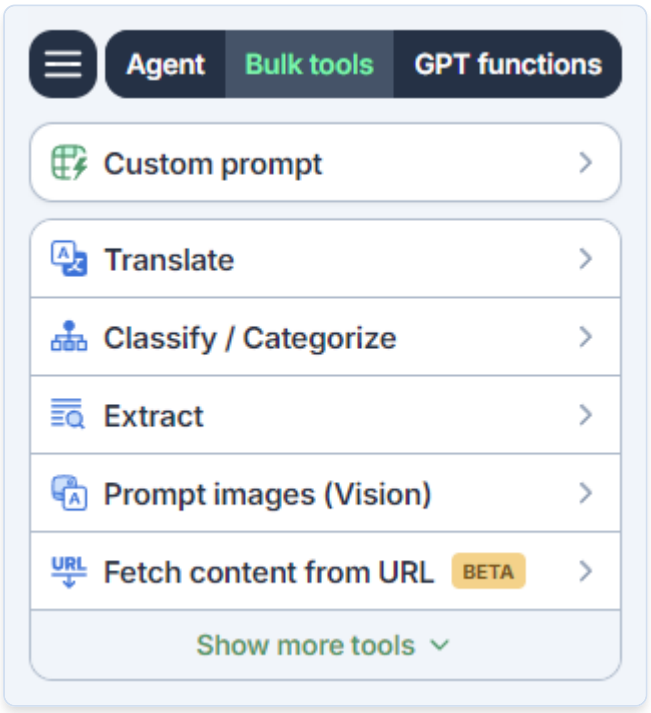
Run a bulk tool

1. Open a workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.

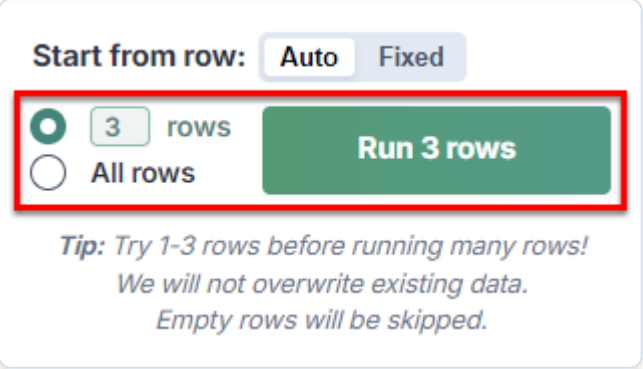


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **Bulk tools** and select the bulk tool you want to use.

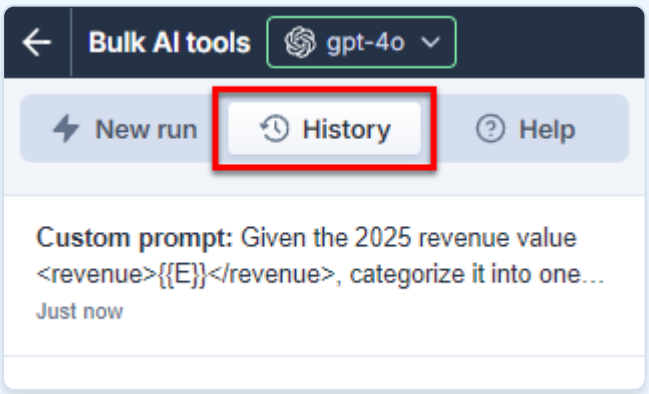


4. Set up the bulk tool run:
   - Custom prompt
   - Translate
   - Classify / Categorize
   - Extract
   - Prompt images (Vision)
   - Fetch content from URL
   - Generate
   - Rewrite
   - Reformat
5. Run the bulk tool:
   1. Select a specific number of rows to run or select **All rows**. Learn more.
   2. Click **Run rows**.



The bulk tool runs, starting from the first empty cell in the results column(s).

> **INFO**
>
> GPT for Excel saves the setup in its bulk tool run history. You can reuse the setup by selecting it from **History**.



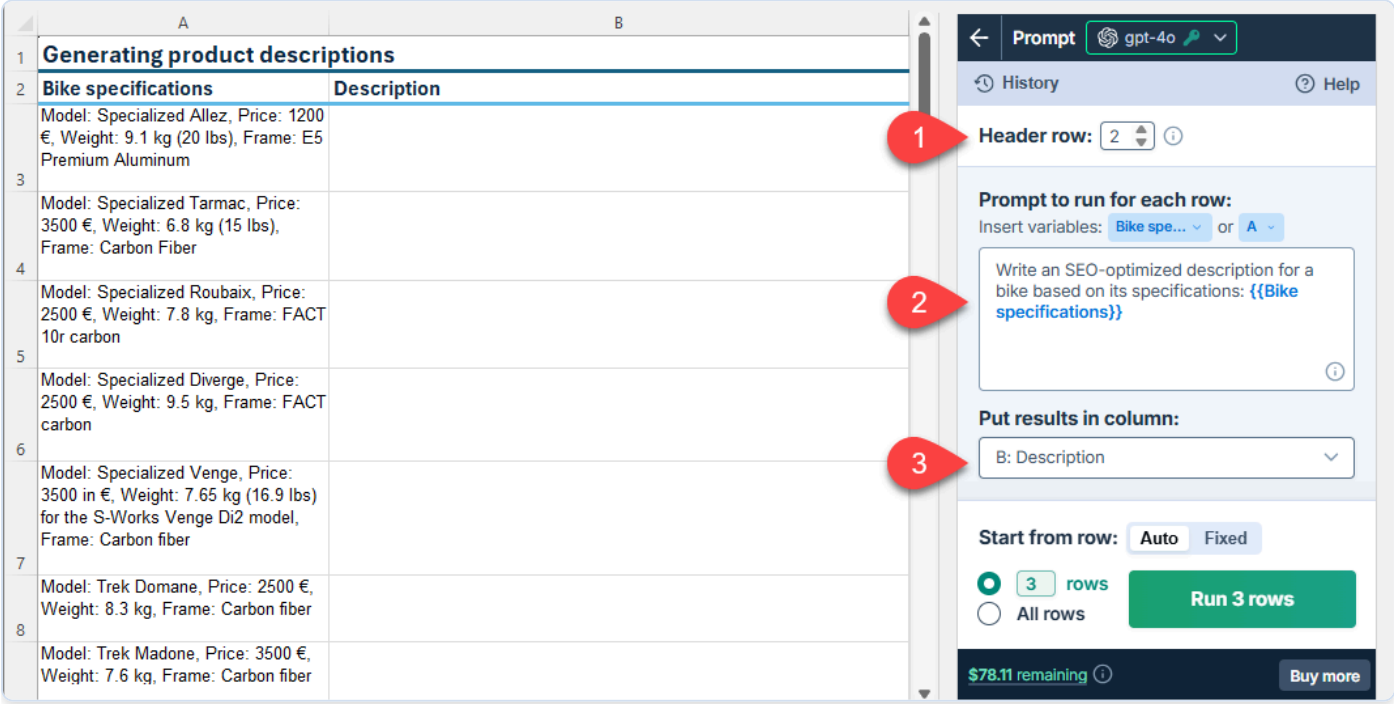
6. Check the results.
7. (Optional) If needed, try improving the results.

Custom prompt

Use the **Custom prompt** bulk tool to perform high-volume, free-form prompting directly in Excel. For example, if you have product specifications in one column, you can generate a unique product description for each row based on the specification column value for that row.

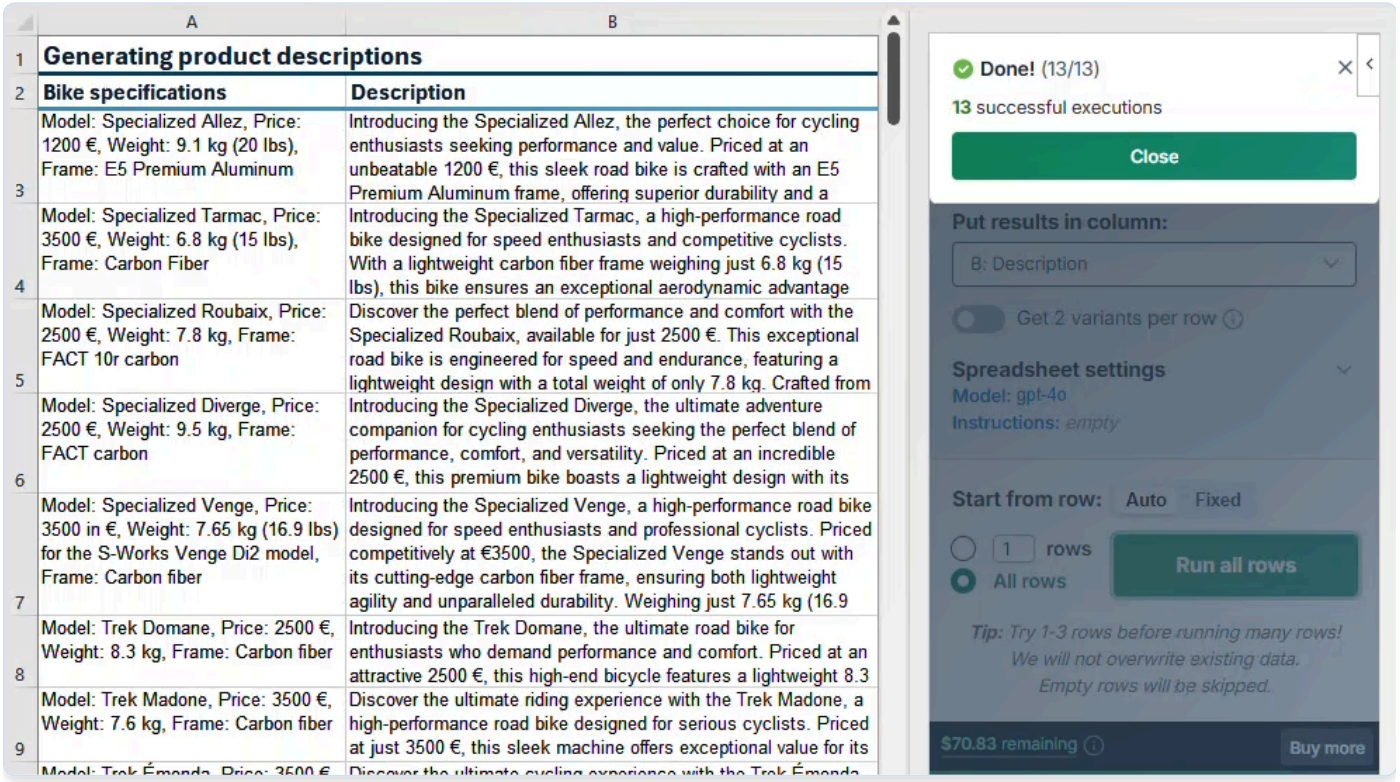
> **INFO**
>
> - **Custom prompt** and **Prompt images (Vision)** are the only bulk tools that accept images as input.
> - To use **Custom prompt** with image input, you must select a vision model.

Settings (Custom prompt)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	2

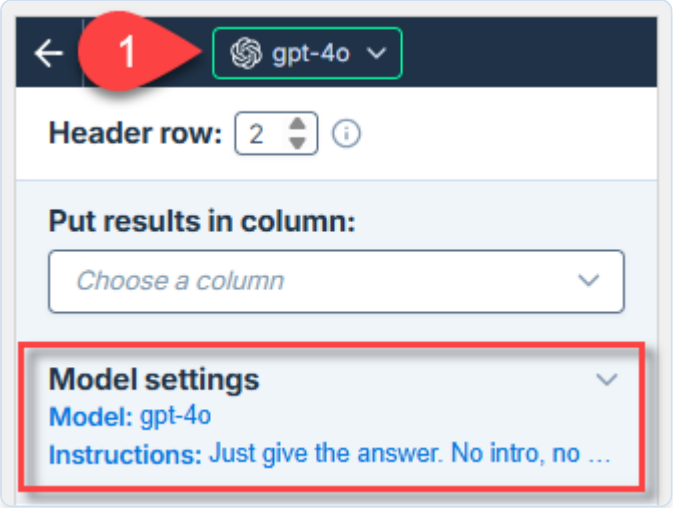
	Field	Description	Example
2	**Prompt to run for each row**	Enter the prompt you want to run for each row. You can use column, cell, and range variables to reference content in other columns and cells in the spreadsheet. You must define at least one column variable for context. Learn more.	`Write an SEO-optimized description for a bike based on its specifications: {{Bike specifications}}`
3	**Put results in column**	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`B: Description`



Improve results (Custom prompt)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

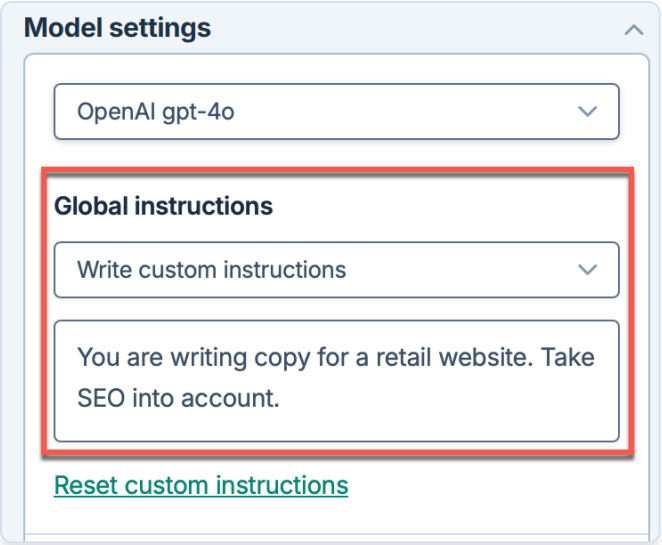
You can improve response quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

Define the instructions as **Global instructions** in the **Model settings**. The instructions apply to all bulk tools and GPT functions in the current workbook.

> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to add instructions to follow SEO best practices:

```
You are writing copy for a retail website. Take SEO into account.
```

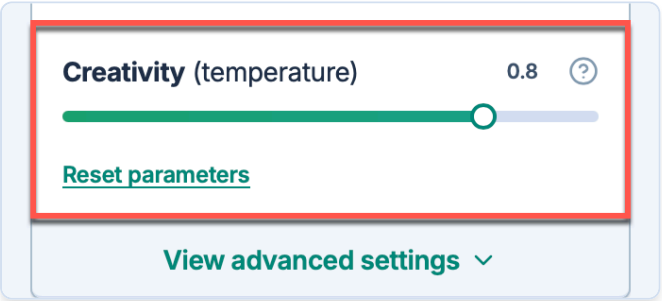


Adjust the creativity level

Adjust the level of creativity (accuracy vs. freedom) the AI is allowed in generating results. The higher the value, the greater the creativity.

For example:

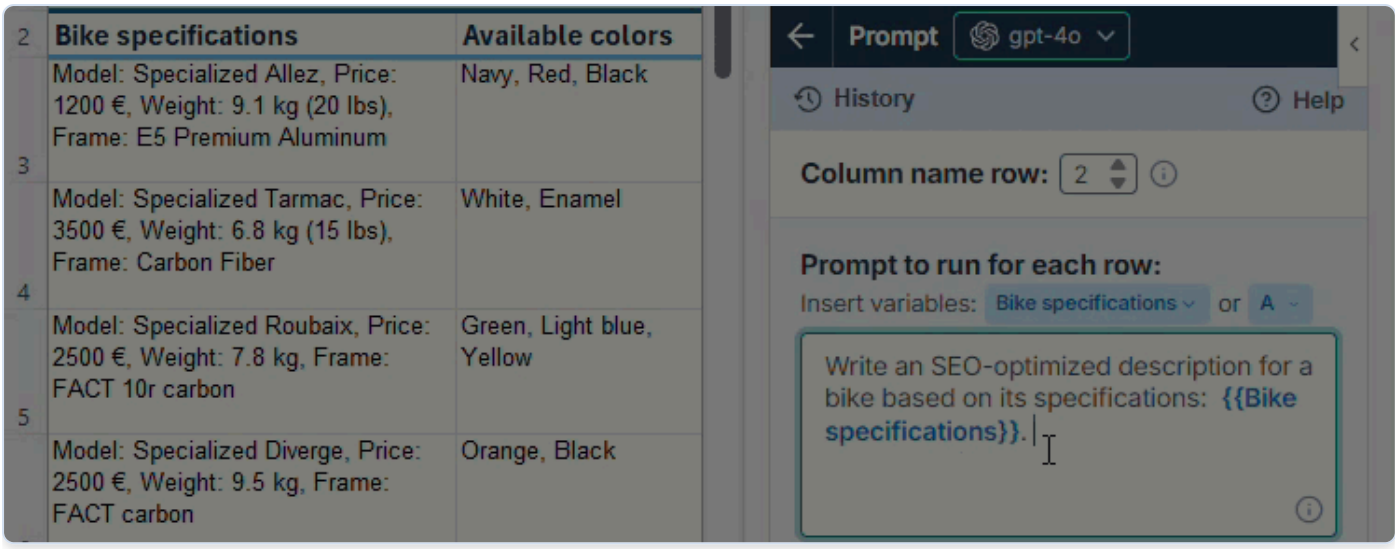
- If you're generating marketing content, set a high creativity level to generate more engaging and imaginative text.
- If you're generating technical specifications, set a low creativity level to ensure the generated text is precise and factual.
- If you need the AI to follow your instructions to the letter, set the creativity level to **0**.



> **INFO**
>
> The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

Use multiple variables in a custom prompt

You must define at least one column variable in a custom prompt for context, for example `{{A}}` or `{{Bike specifications}}`. To generate more accurate results, provide more context by using additional variables in the prompt.



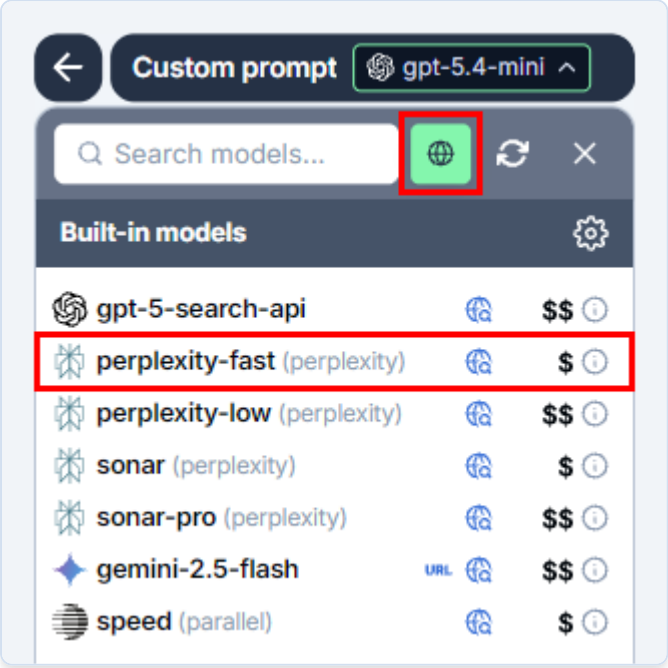
You can insert variables to reference columns, cells, and ranges in the prompt:

Variable	Description	Example
Column	Reference a column to include row-specific information from an adjacent column. Column variables work like relative references and are highlighted in blue. To add a column variable, use the **Insert variables** menus above the prompt. You can add a variable either by column label or column letter.	`Write a product description for {{Product name}} that costs {{Price}} and is made by {{Manufacturer}}.`
Cell or range from the current sheet	Reference a specific cell or range to include the same content as context for each row. Cell and range variables work like absolute references and are highlighted in green. To add a cell or range variable, type it directly in the prompt.	`Create a tagline for {{Company name}}.` `Follow these rules: {{A1}}` `Use 1 term from this list: {{A2:A100}}`

Variable	Description	Example
Cell or range from another sheet	Reference a cell or a range from another sheet in the same spreadsheet. The content can be, for example, reference data or other shared information.	`Translate {{Source text}} to Japanese.` `Use the glossary: {{Glossary!A1:B100}}`

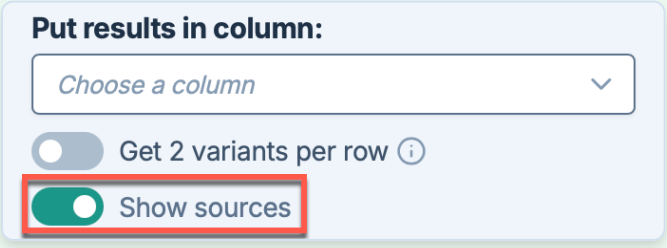
Use web search capabilities

To generate more up-to-date content, select a model with web search capabilities (indicated by the 🌐 icon in the model switcher). The available web search options vary depending on the model you select. Learn more about web search options.



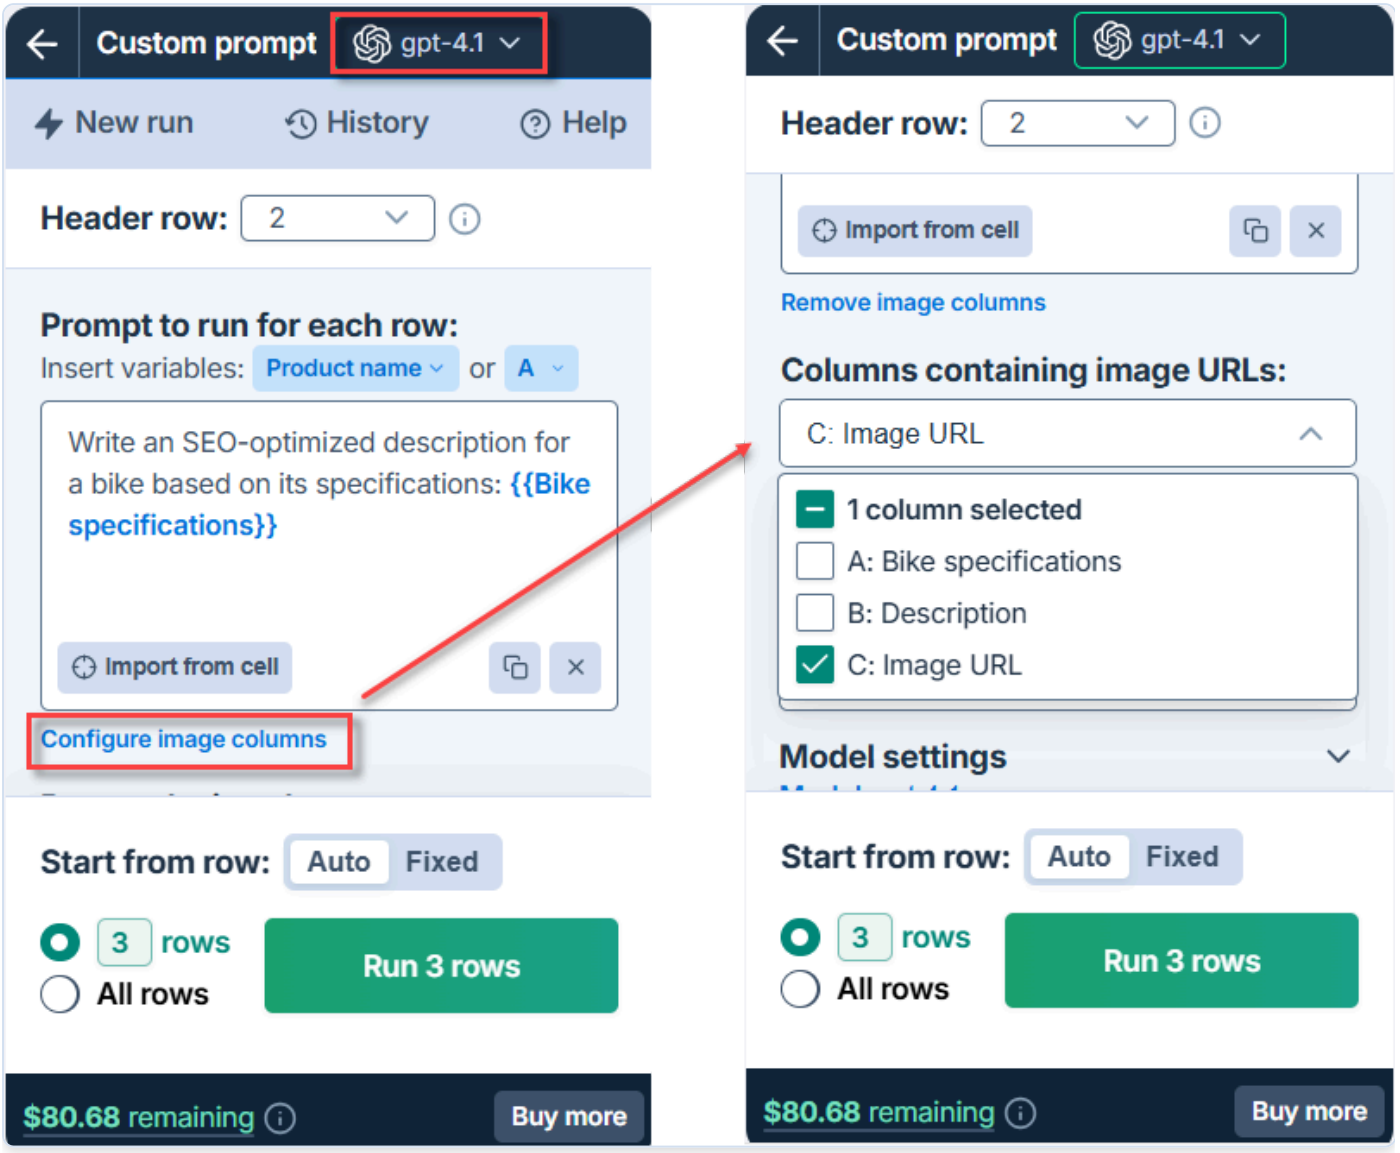
TIP

Enable **Show sources** to write the source references in the column to the right of the results column. This helps you understand the origin of the content and assess its reliability.



Include images in a custom prompt

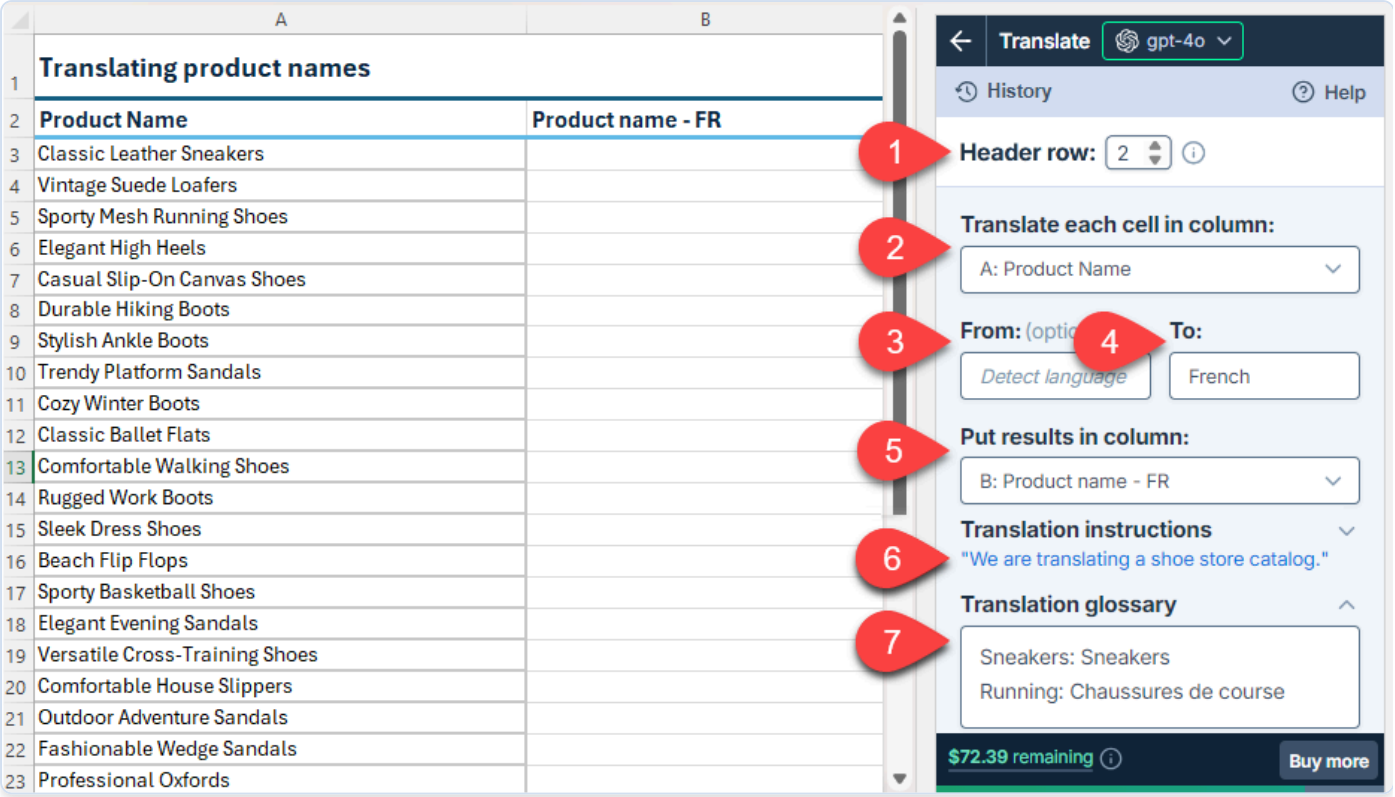
To include images in a custom prompt, select a vision model in the model switcher. Then, select the columns that contain image URLs so they can be attached to the prompt for each corresponding row. Learn more about vision options.



Translate

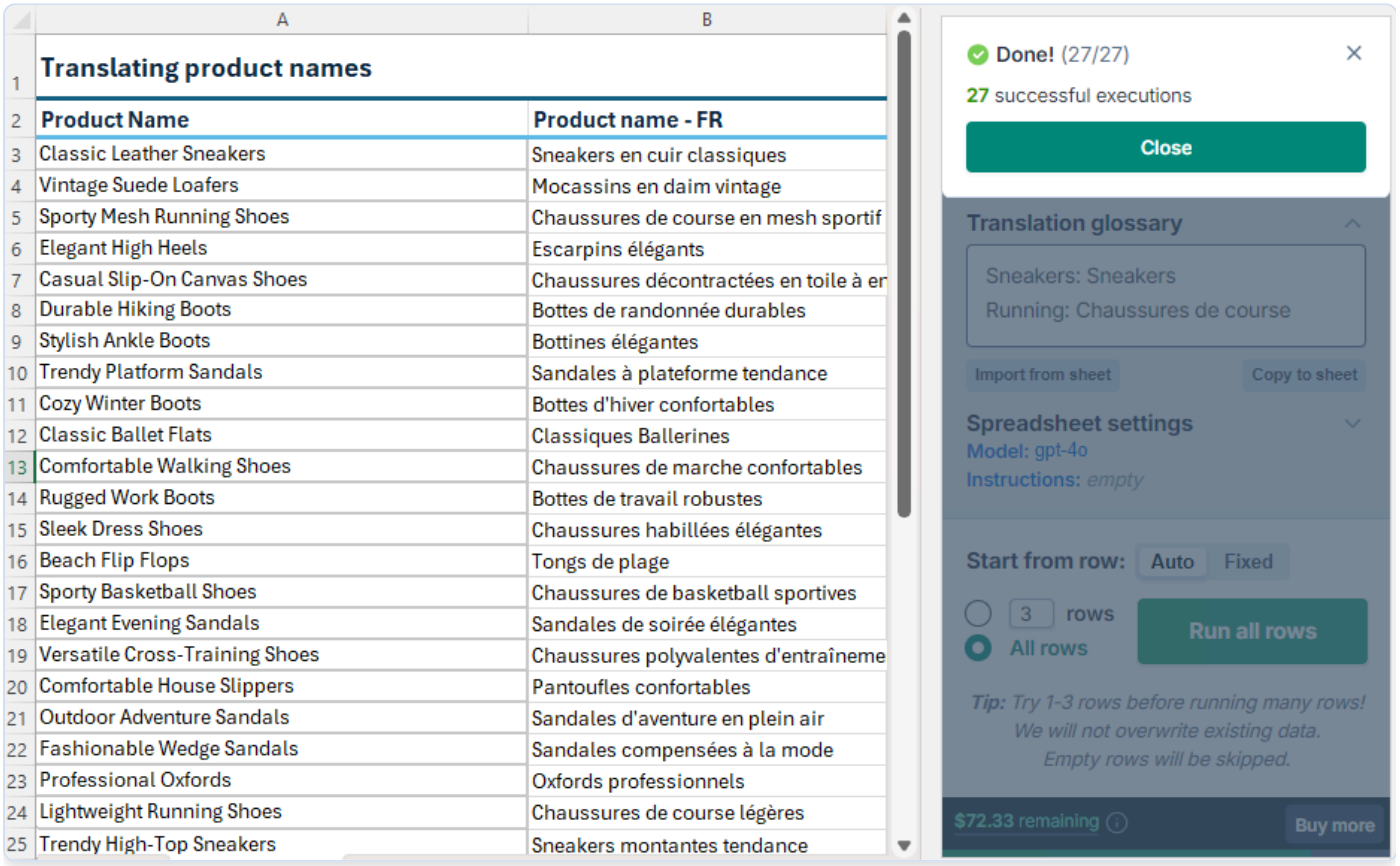
Use the **Translate** bulk tool to translate product catalogs, website copy, customer reviews, or any other content directly in Excel, and optimize your translations by adding specific instructions and a glossary to guide the AI in the process. For example, you can translate a list of product names for a shoe store into French.

Settings (Translate)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Translate each cell in column**	Select the column that contains the text you want to translate.	`A: Product Name`
3	**From** (optional)	Enter the source language, or leave the field empty if you want the AI to detect the source language for each cell.	*Empty*
4	**To**	Enter the target language.	`French`
5	**Put results in column**	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`B: Product Name - FR`

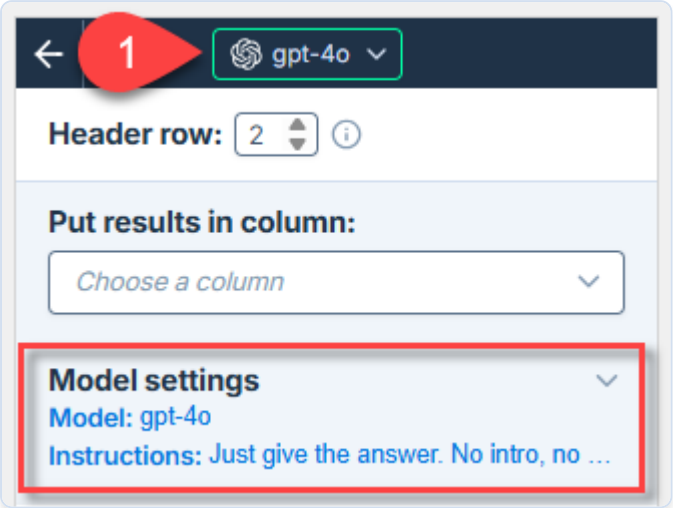
	Field	Description	Example
6	**Translation instructions** (optional)	Enter specific instructions for the translation.	`We are translating a shoe store catalog.`
7	**Translation glossary** (optional)	Enter a list of word pairs where the first word is the source language word and the second word is the target language translation of the first word. Use the following syntax for each pair: `<Source word>: <Target Word>` Enter one pair per line. If your glossary is in the current document, you can import it by clicking **Import from sheet**.	`Sneakers: Sneakers` `Running: Chaussures de course`



Improve results (Translate)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

You can improve translation quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

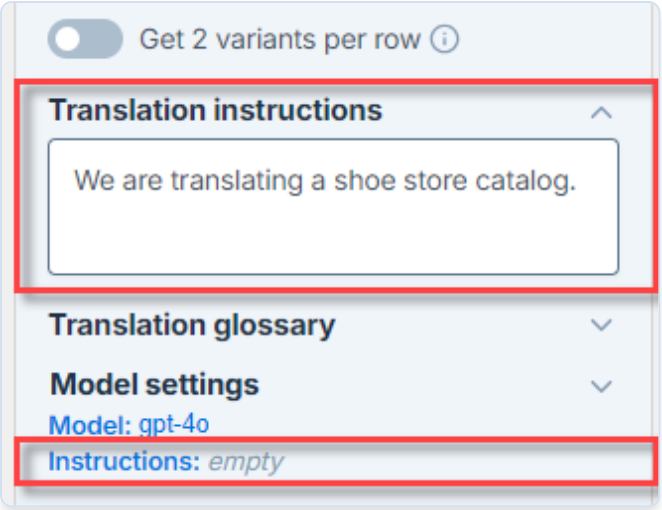
You can define the following types of instructions for the AI:

- **Translation instructions** are defined for the current bulk tool run.
- **Global instructions** are defined in the **Model settings** and apply to all bulk tools and GPT functions in the current workbook.

If both bulk tool instructions and global instructions are defined, the AI uses both as a combined set of instructions. For consistent results, make sure both sets of instructions are aligned and do not contradict each other.

> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to omit the global instructions and only define translation instructions:

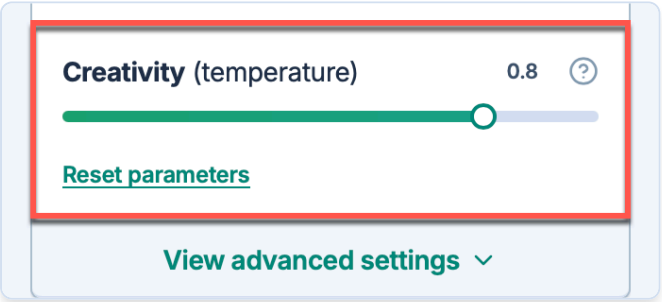


Adjust the creativity level

Adjust the level of creativity (accuracy vs. freedom) the AI is allowed in generating results. The higher the value, the greater the creativity.

For example:

- If you're generating marketing content, set a high creativity level to generate more engaging and imaginative text.
- If you're generating technical specifications, set a low creativity level to ensure the generated text is precise and factual.
- If you need the AI to follow your instructions to the letter, set the creativity level to **0**.

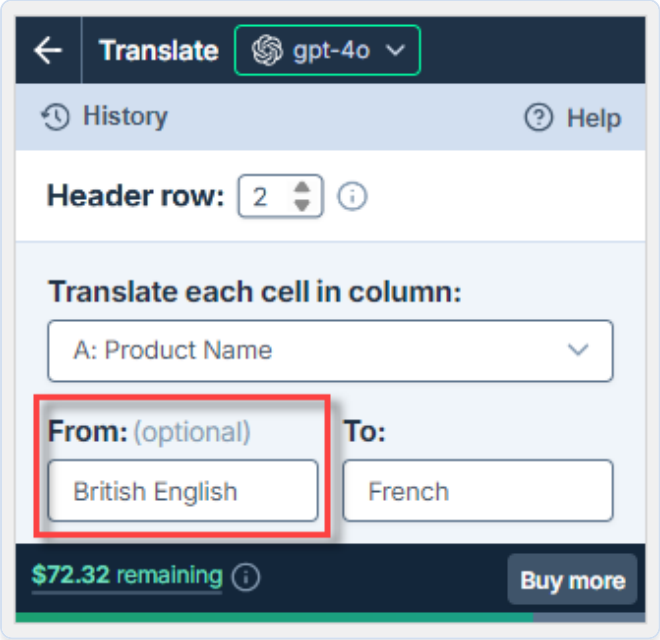


INFO

The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

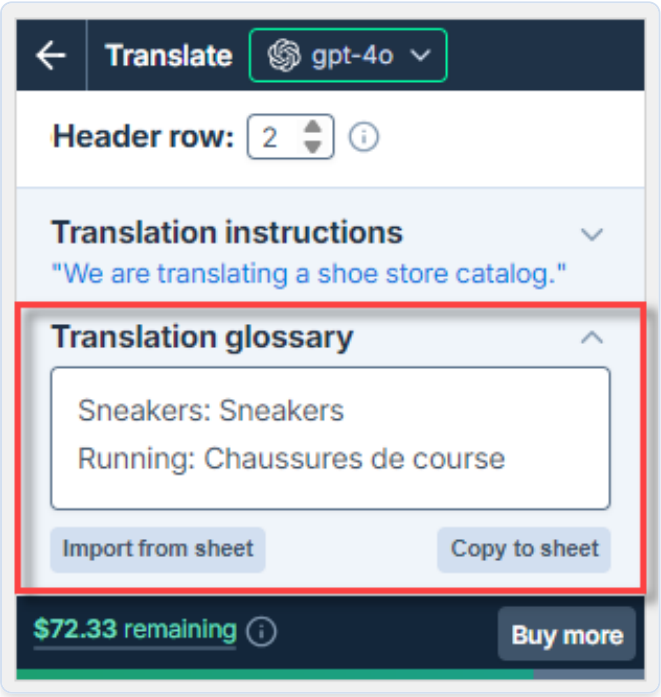
Set the source language for translations

The AI can automatically detect the source language for translations, but manually defining it helps if the source text contains short, ambiguous phrases, which can confuse the AI. For example, "Trainers" in British English means "Sports shoes", whereas in American English, it means "People who train others".



Create a glossary for translations

Glossaries ensure that translations are consistent, respect corporate terminology, and remain on-brand. If your glossary is in the current spreadsheet, you can import it by clicking **Import from sheet**.

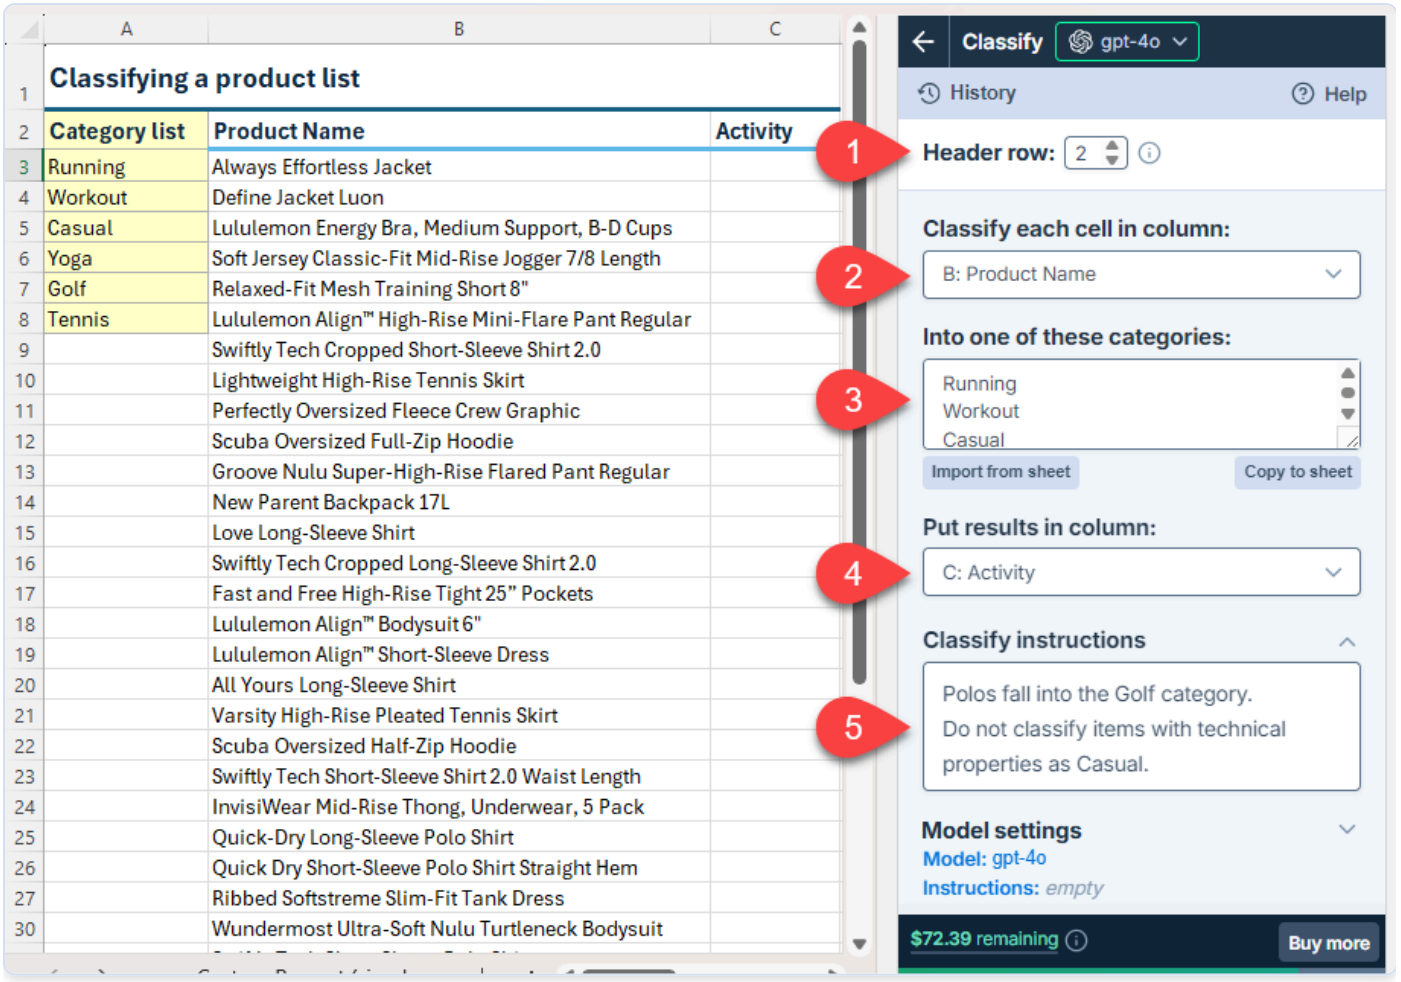


Classify / Categorize ▼

Use the **Classify / Categorize** bulk tool to classify your catalog items, press articles, customer reviews, social media comments, or any other text directly in Excel. The tool classifies the content of each cell into one of the categories you provide if a match is found; if not match is found, the

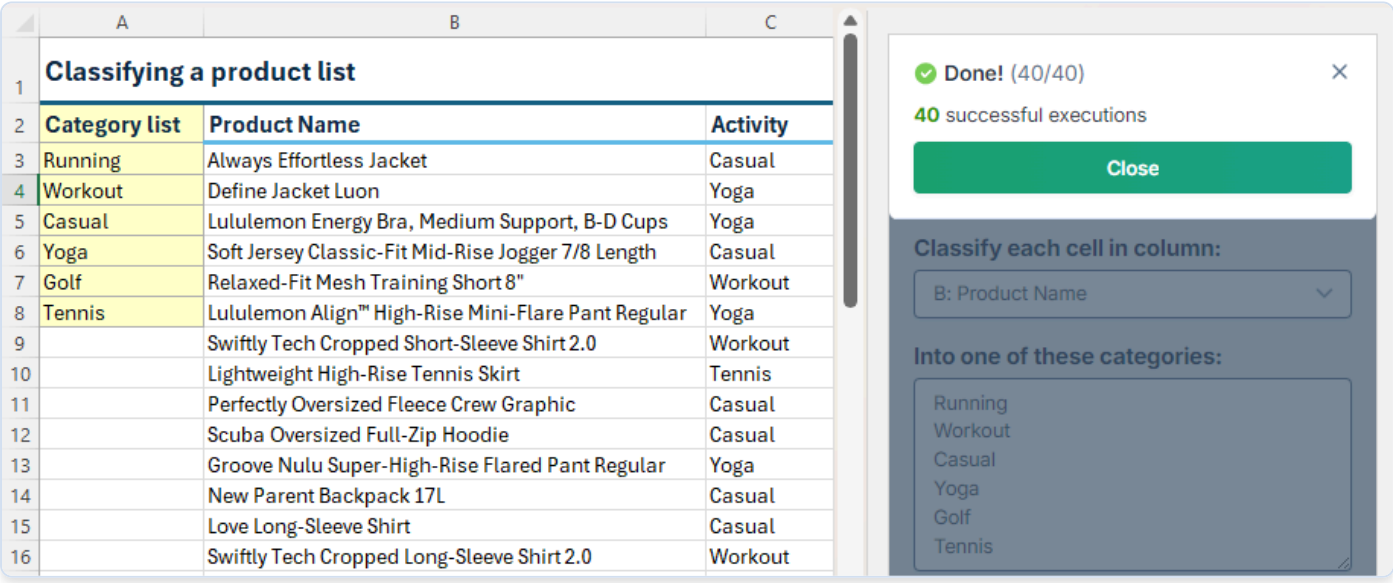
tool fills in `#N/A`. For example, you can improve your lifecycle inventory management by classifying apparel products according to the activity they are used for.

Settings (Classify / Categorize)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Classify each cell in column**	Select the column that contains the text you want to classify.	`B: Product Name`

	Field	Description	Example
3	**Into one of these categories**	Enter the categories to classify the text into. If the categories are listed in the current document, you can import them by clicking **Import from sheet**.	`Running` `Workout` `Casual` `Yoga` `Golf` `Tennis`
4	**Put results in column**	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`C: Activity`
5	**Classify instructions** (optional)	Enter specific instructions for the classification. Learn more.	`Polos fall into the Golf category. Do not classify items with technical properties as Casual.`

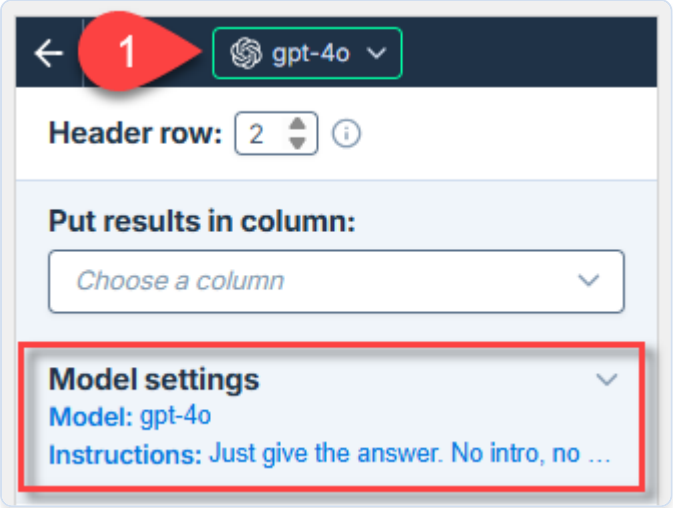


	A	B	C
1	**Classifying a product list**		
2	**Category list**	**Product Name**	**Activity**
3	Running	Always Effortless Jacket	Casual
4	Workout	Define Jacket Luon	Yoga
5	Casual	Lululemon Energy Bra, Medium Support, B-D Cups	Yoga
6	Yoga	Soft Jersey Classic-Fit Mid-Rise Jogger 7/8 Length	Casual
7	Golf	Relaxed-Fit Mesh Training Short 8"	Workout
8	Tennis	Lululemon Align™ High-Rise Mini-Flare Pant Regular	Yoga
9		Swiftly Tech Cropped Short-Sleeve Shirt 2.0	Workout
10		Lightweight High-Rise Tennis Skirt	Tennis
11		Perfectly Oversized Fleece Crew Graphic	Casual
12		Scuba Oversized Full-Zip Hoodie	Casual
13		Groove Nulu Super-High-Rise Flared Pant Regular	Yoga
14		New Parent Backpack 17L	Casual
15		Love Long-Sleeve Shirt	Casual
16		Swiftly Tech Cropped Long-Sleeve Shirt 2.0	Workout

Improve results (Classify / Categorize)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

You can improve classification quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

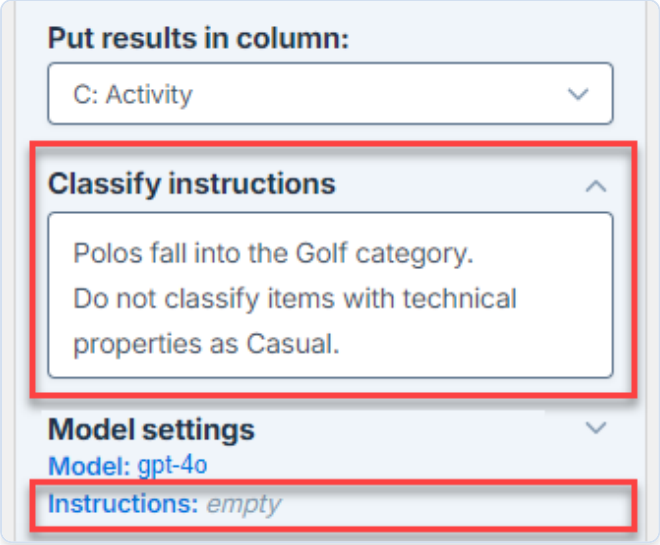
You can define the following types of instructions for the AI:

- **Classify instructions** are defined for the current bulk tool run.
- **Global instructions** are defined in the **Model settings** and apply to all bulk tools and GPT functions in the current workbook.

If both bulk tool instructions and global instructions are defined, the AI uses both as a combined set of instructions. For consistent results, make sure both sets of instructions are aligned and do not contradict each other.

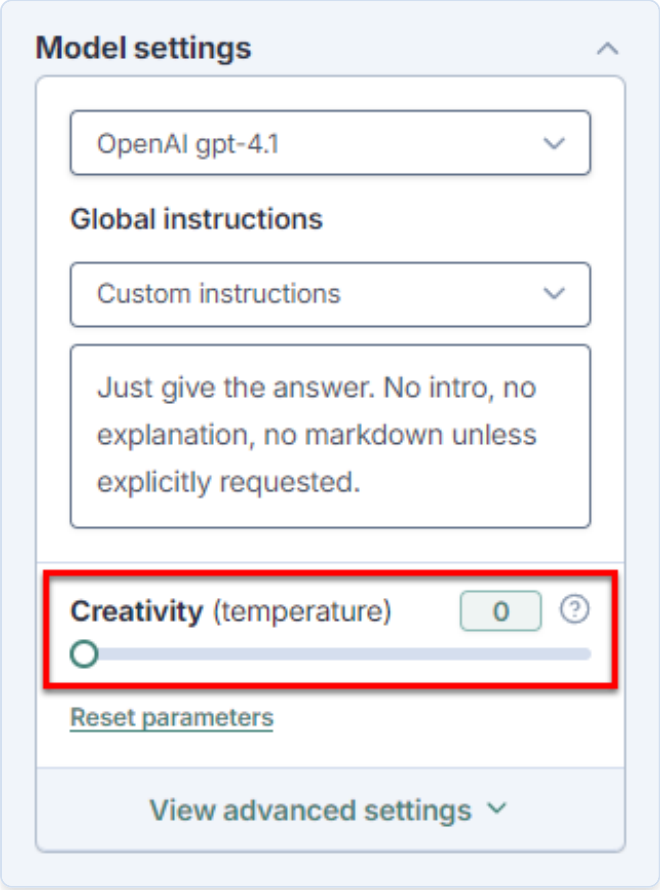
> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to omit the global instructions and only define classify instructions:



Set Creativity to 0

Ensure **Creativity** is set to **0**. This helps the AI follow instructions more closely.



> **INFO**
>
> The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

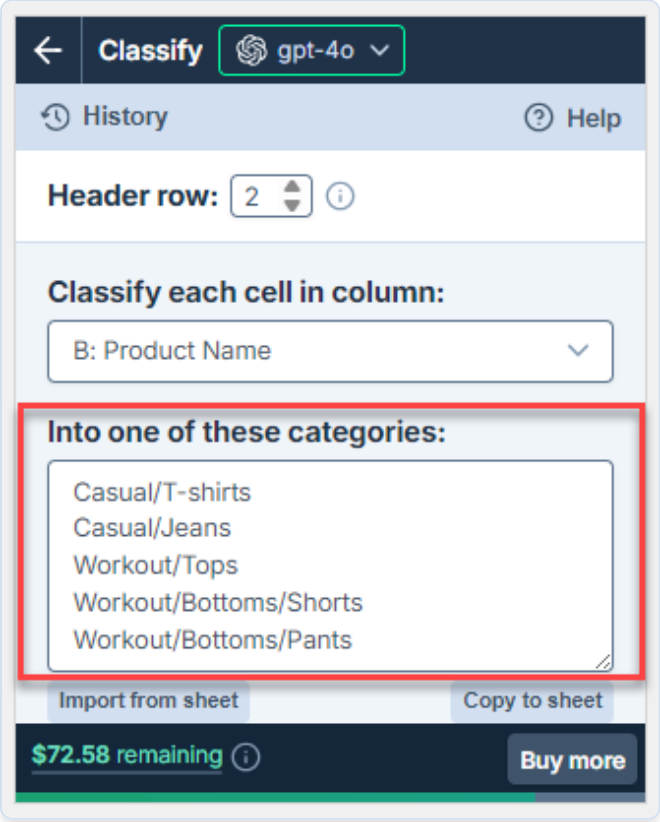
Use full path for categories

If your categories form a hierarchy, use the full path of each category instead of just its name. This helps avoid ambiguity and ensures that the classification is as precise as possible.

Using only category names can lead to ambiguity. For example, `T-shirts` and `Tops` could be considered as very largely overlapping categories. By specifying `Casual/T-shirts` and `Workout/Tops`, you provide a clear context, ensuring that the AI classifies the items accurately according to the intended category.

For example, instruct the AI to classify items according to these category paths:

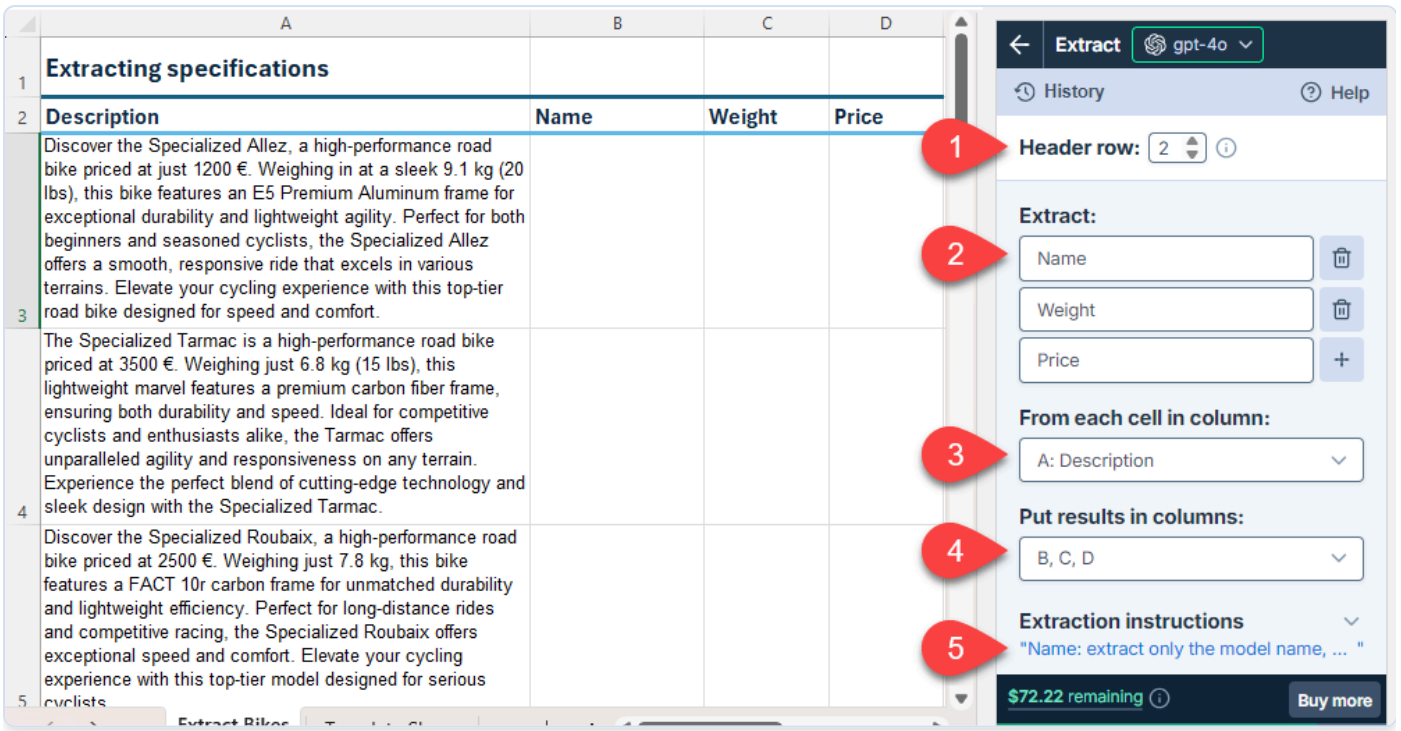
```
Casual/T-shirts
Casual/Jeans
Workout/Tops
Workout/Bottoms/Shorts
Workout/Bottoms/Pants
```



Extract

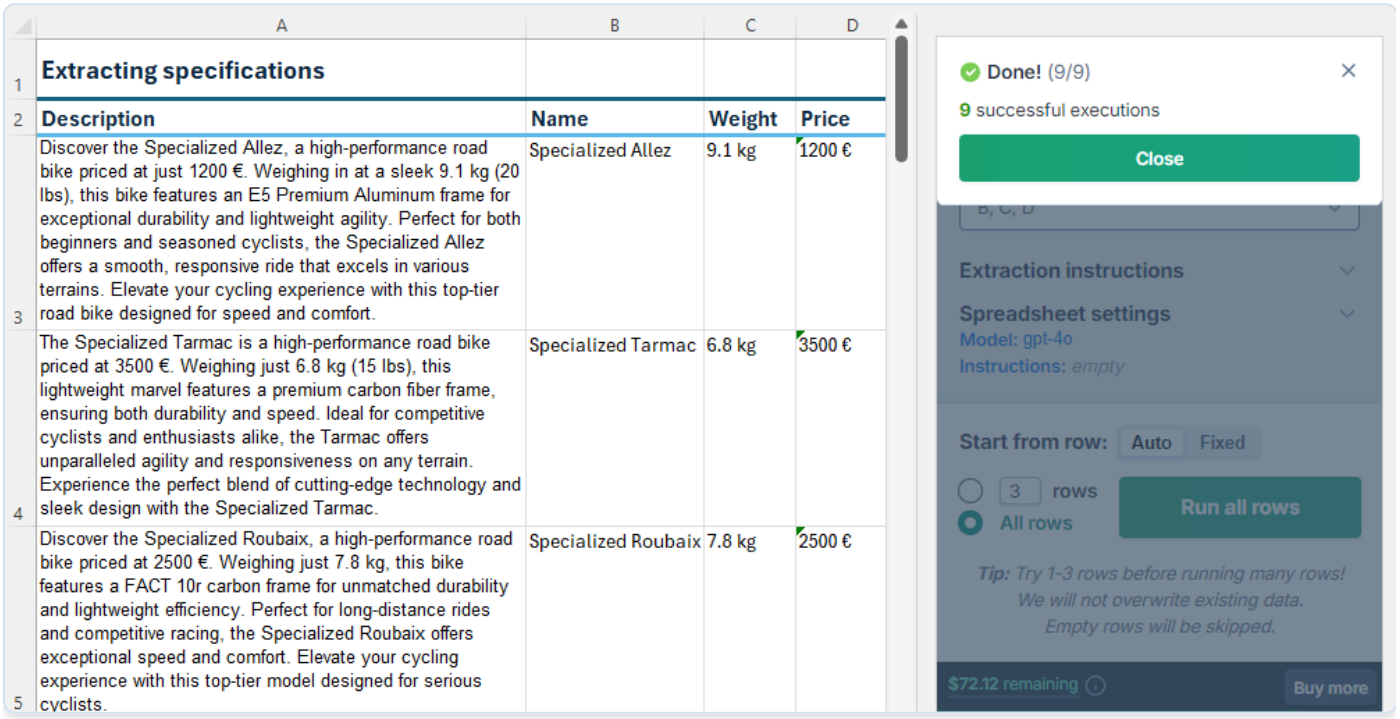
Use the **Extract** bulk tool to extract key information like names, email addresses, and product attributes from unstructured text directly in Excel. For example, you can extract name, weight, and price from bike model descriptions to create a catalog.

Settings (Extract)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Extract**	Enter the entities you want to extract.	`Name` `Weight` `Price`
3	**From each cell in column**	Select the column that contains the text from which you want to extract entities.	`A: Description`

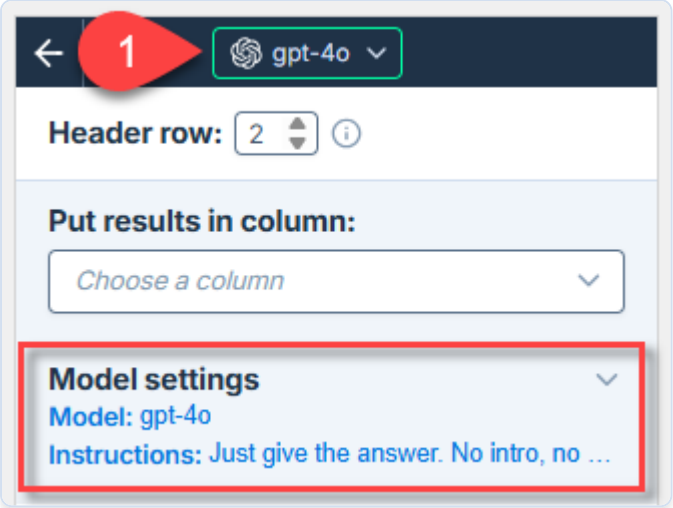
	Field	Description	Example
4	**Put results in column**	Select the column or columns to put the results in. Cells in these columns won't be overwritten with the results if they contain text.	`B, C, D`
5	**Extraction instructions** (optional)	Enter specific instructions for the extraction.	`Name: Extract only the model name, without manufacturer or year`



Improve results (Extract)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

You can improve extraction quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

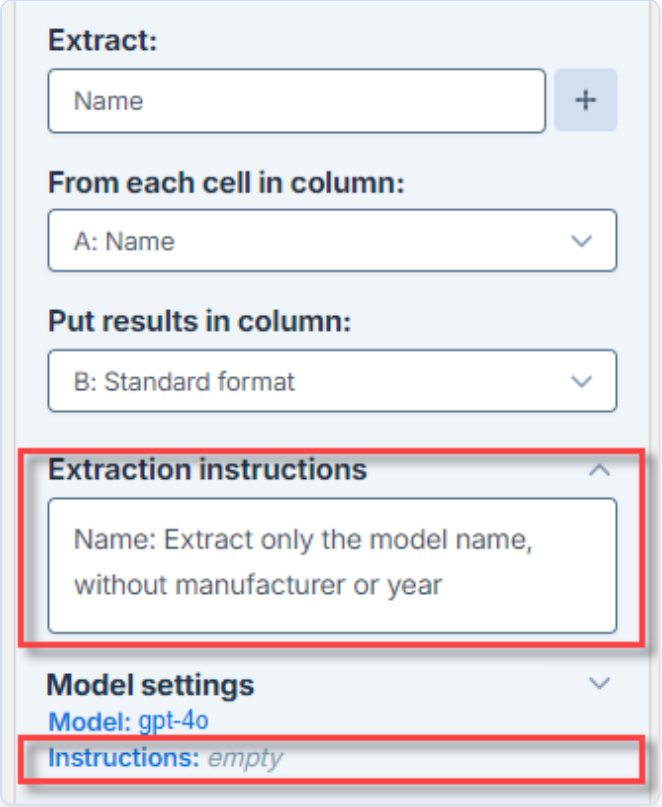
You can define the following types of instructions for the AI:

- **Extraction instructions** are defined for the current bulk tool run.
- **Global instructions** are defined in the **Model settings** and apply to all bulk tools and GPT functions in the current workbook.

If both bulk tool instructions and global instructions are defined, the AI uses both as a combined set of instructions. For consistent results, make sure both sets of instructions are aligned and do not contradict each other.

> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to omit the global instructions and only define extraction instructions:



The following table provides examples of extraction instructions for different goals:

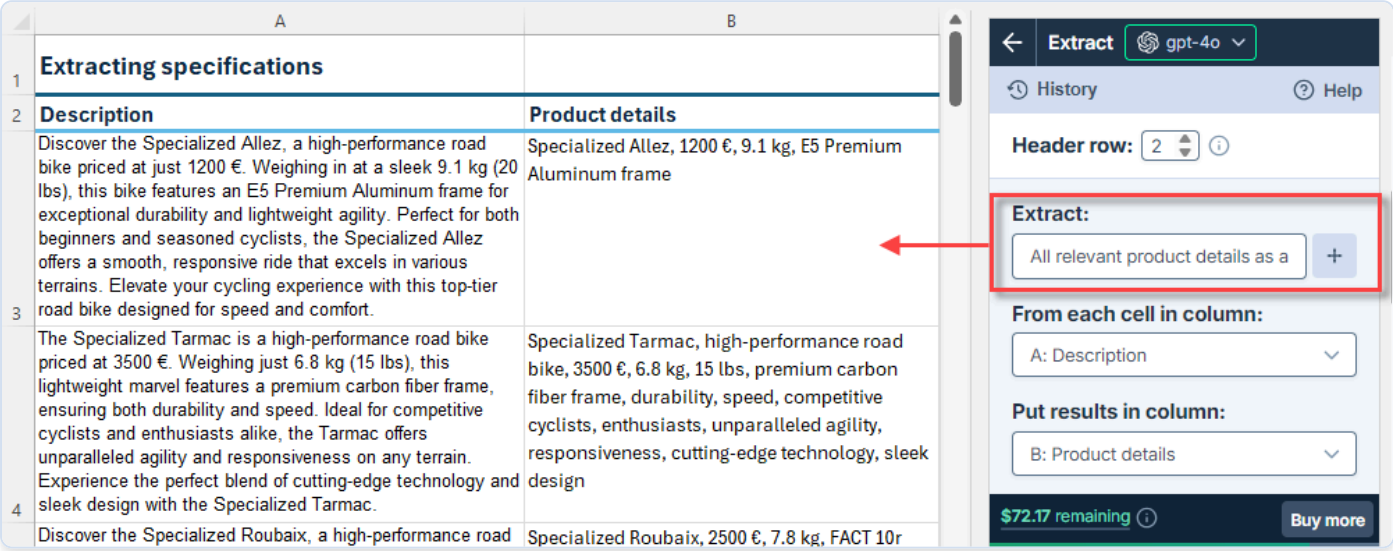
Goal	Description	Example
Remove irrelevant information	Define what should be excluded from the extraction to get cleaner results.	`Name: Extract only the model name, without manufacturer or year`
Normalize entities	Set a consistent output for variations of the same entity in your text.	`Weight: Extract weight in kilograms`
Extract numerical values and units	Separate numbers from units to simplify their manipulation.	Add `Price value` and `Price currency` as entities in the **Extract** field, then instruct: `Price value: Extract only the numerical value` `Price currency: Extract only the currency code (for example USD, or EUR)`

Perform open-ended extraction

Open-ended extraction helps you gather a wide range of information without specifying each entity. This approach works well when entities vary from one cell to another or when you want to avoid missing any extraction.

For example, enter the following in the **Extract** field:

```
All relevant product details as a comma-separated list.
```

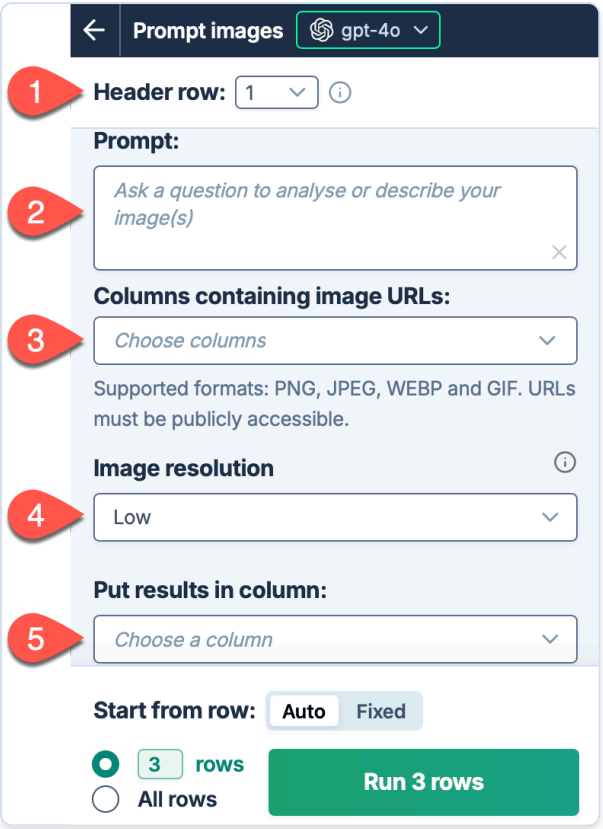


Prompt images (Vision) ▼

Use the **Prompt images (Vision)** bulk tool to generate content based on images directly in Excel. For example, you can generate descriptions for your product images to improve your listings.

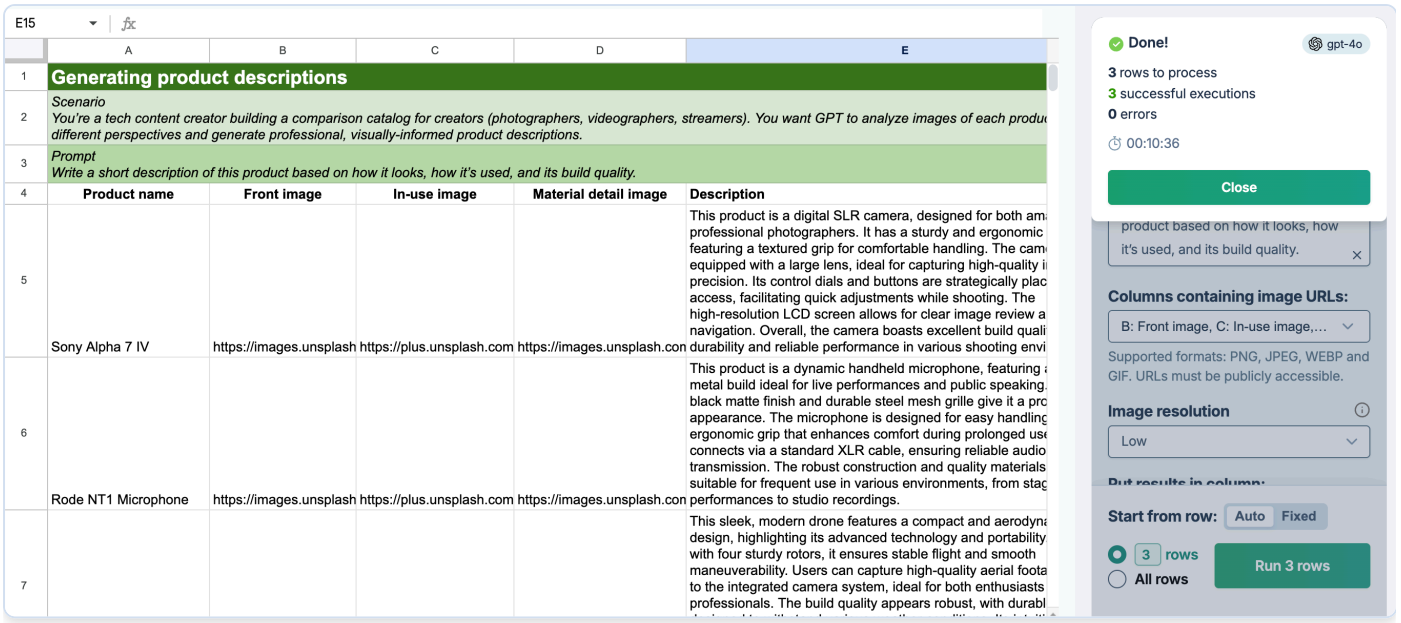
> **INFO**
>
> - **Custom prompt** and **Prompt images (Vision)** are the only bulk tools that accept images as input.
> - To use **Prompt images (Vision)**, you must select a vision model.

Settings (Vision)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`4`
2	**Prompt**	Describe the task you want the AI to perform based on the target images. For example, ask a question about the images or request content based on the images.	`Write a short description of this product based on how it looks, how it's used, and its build quality.`

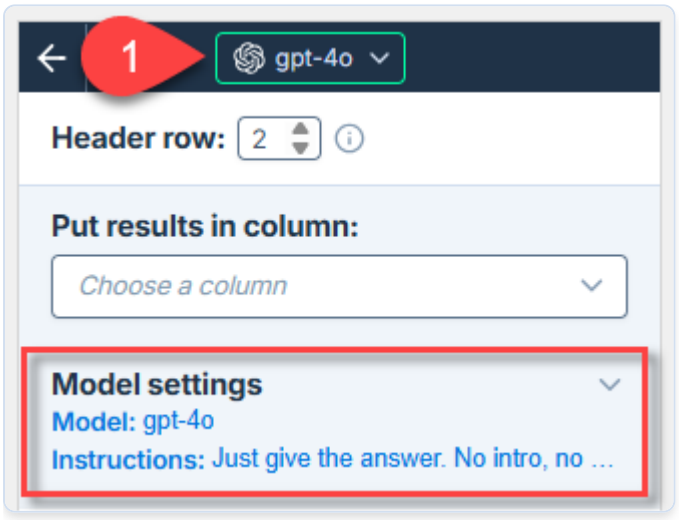
	Field	Description	Example
3	**Columns containing image URLs**	Select up to five columns that contain image URLs. For each row, the tool processes all images present in the selected columns. URLs must be publicly accessible. Supported image formats: • PNG • JPEG • WebP • Non-animated GIF	`B: Front image, C: In-use image, D: Material detail image`
4	**Image resolution** (optional)	Choose the level of detail for image processing: • **High:** Better image understanding, higher latency and cost • **Low:** Faster processing, lower cost • **Auto:** Let the model decide (default)	`Low`
5	**Put results in column** (optional)	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`E: Description`



Improve results (Vision)

Use the recommended vision model

In the model switcher, select **gpt-5.4** as the vision model. For more information about image-related tasks, see Use images in prompts.



INFO

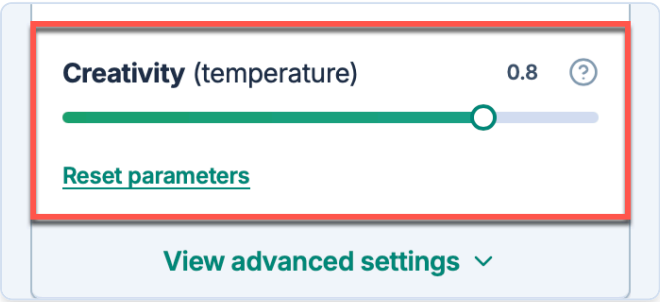
The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Adjust the creativity level

Adjust the level of creativity (accuracy vs. freedom) the AI is allowed in generating results. The higher the value, the greater the creativity.

For example:

- If you're generating marketing content, set a high creativity level to generate more engaging and imaginative text.
- If you're generating technical specifications, set a low creativity level to ensure the generated text is precise and factual.
- If you need the AI to follow your instructions to the letter, set the creativity level to **0**.



> **INFO**
> The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

Adjust image resolution

Choose the appropriate image resolution setting based on your needs:

- **High:** Use for detailed analysis where image quality and understanding are critical. This setting provides better image understanding but increases latency and cost.
- **Low:** Use for faster processing and lower cost when basic image recognition is sufficient.
- **Auto:** Let the model decide the optimal resolution (default).

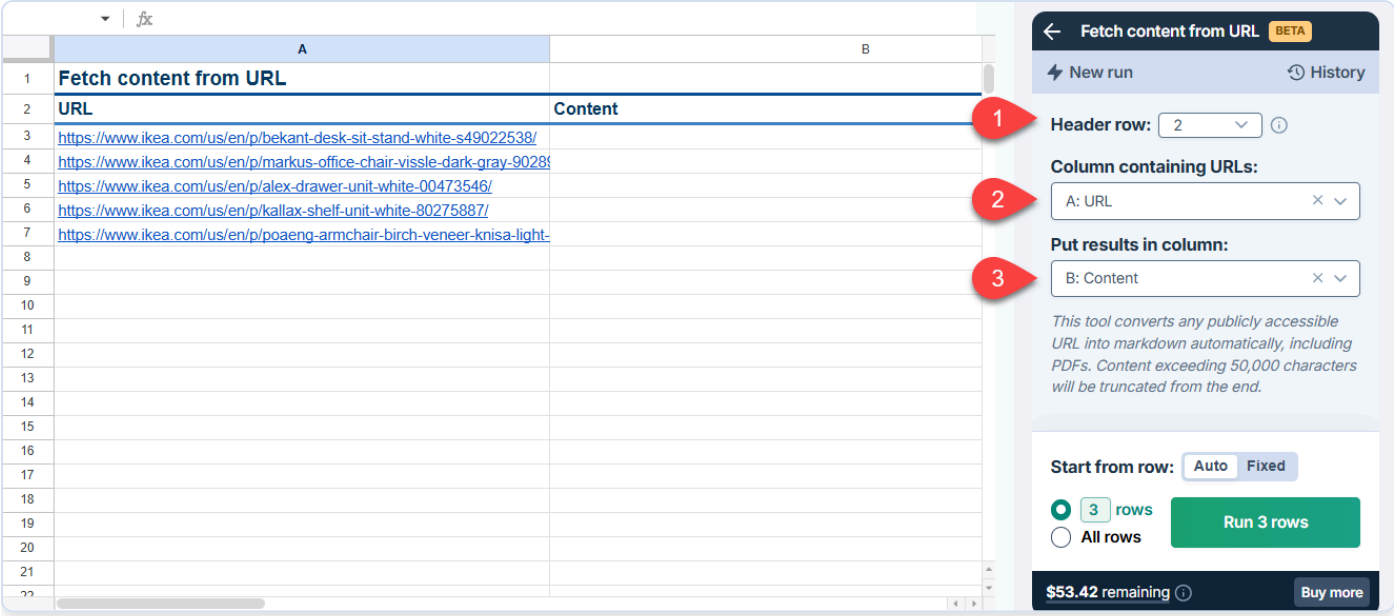
For product descriptions and detailed analysis, consider using **High** resolution to ensure the AI captures all important visual details.

Fetch content from URL ▼

Use the **Fetch content from URL** bulk tool to fetch and convert web page content from URLs directly in Excel. This tool converts any publicly accessible URL into markdown, including PDFs. For example, you can fetch product pages to extract details, or collect company information for enrichment.

> **INFO**
> This bulk tool always uses the **speed** model from Parallel, regardless of the model selected in the model switcher. For detailed pricing information, see Model prices.

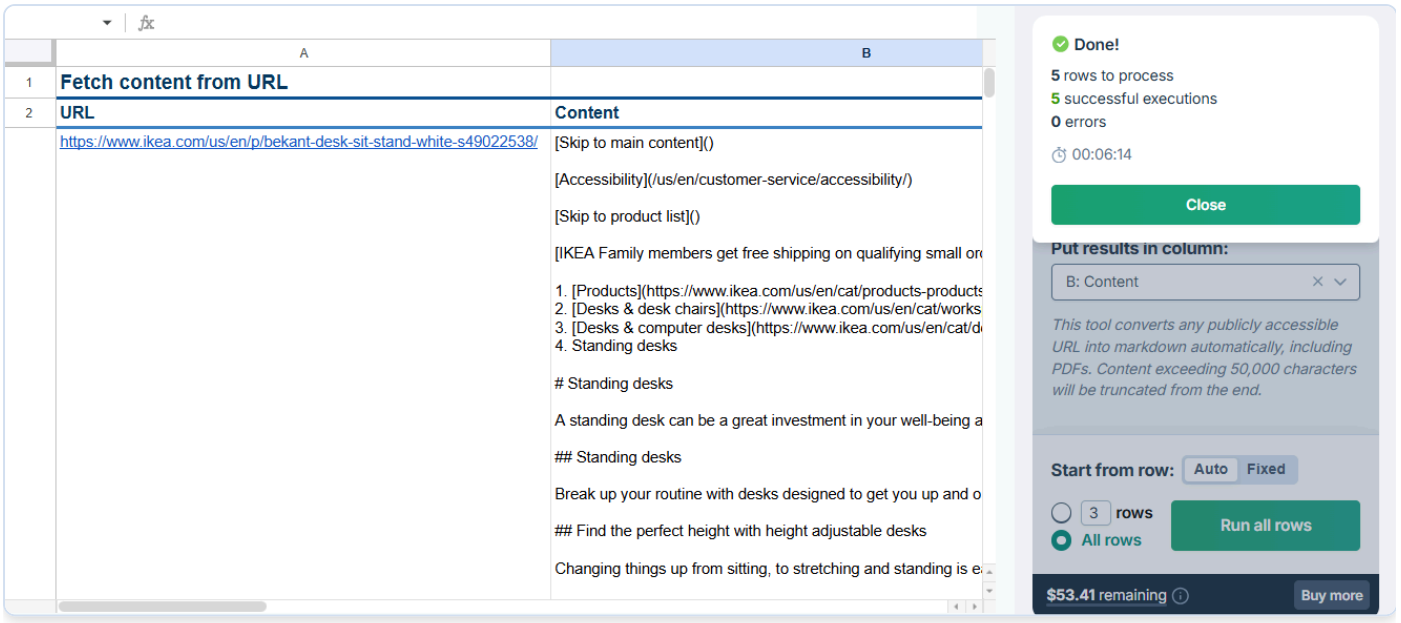
Settings (Fetch content from URL)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Column containing URLs**	Select the column that contains the URLs to fetch content from.	`A: URL`
3	**Put results in column**	Select the column to put the fetched content in. Cells in this column won't be overwritten with the results if they contain text.	`B: Content`

INFO

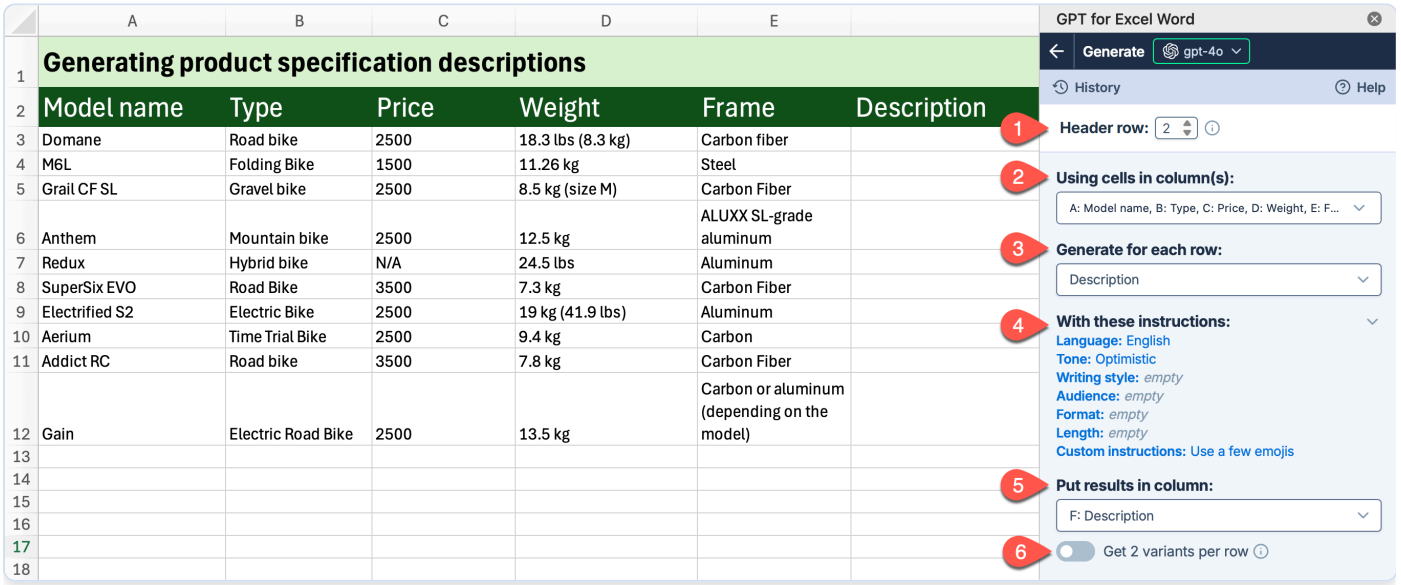
Content exceeding 50,000 characters will be truncated from the end.



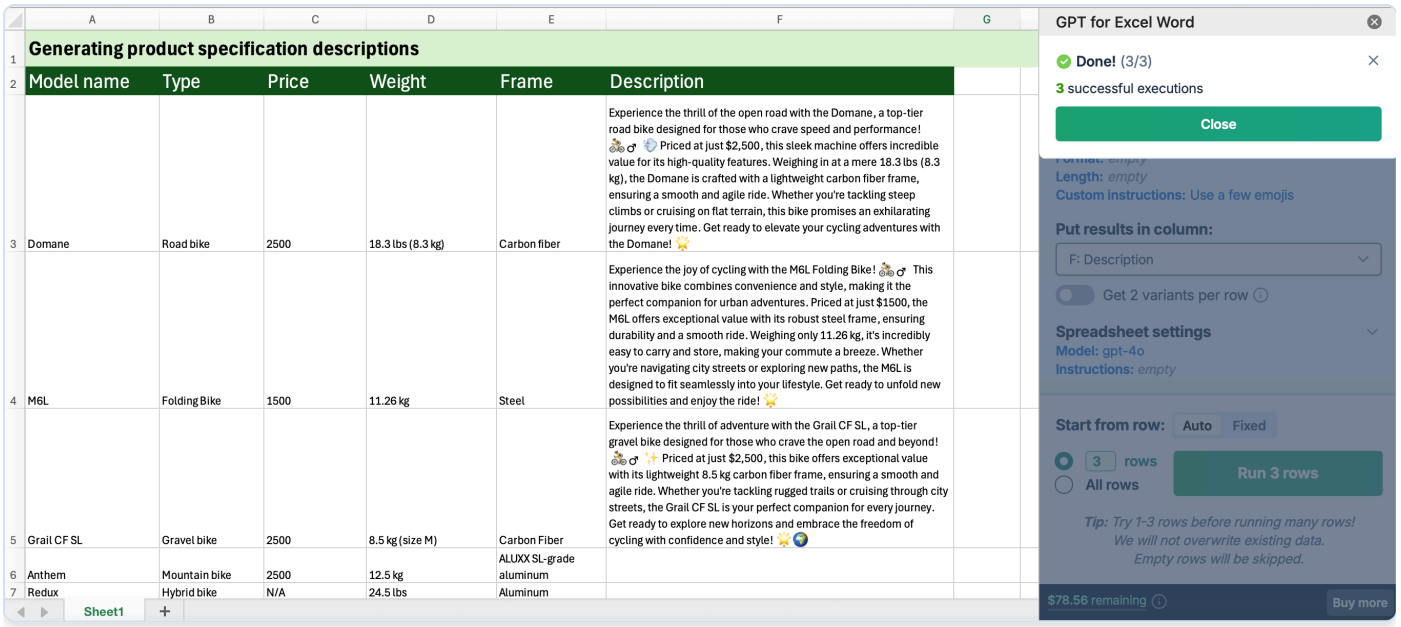
Generate

Use the **Generate** bulk tool to generate content by combining cells from multiple columns in Excel. For example, you can generate product descriptions based on specifications from different columns.

Settings (Generate)



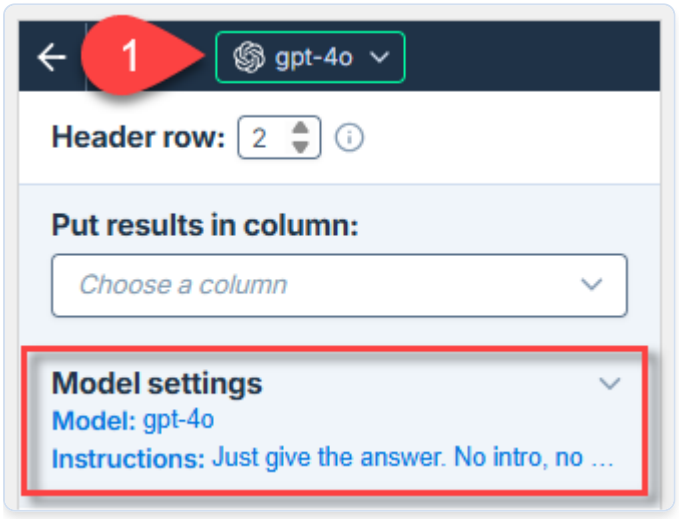
	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Using cells in column(s)**	Select one or more columns containing the cells you want to use as context for the generation.	`A: Model name, B: Type, C: Price, D: Weight, E: Frame`
3	**Generate for each row**	Select the type of content you want to generate, such as a title or description. You can also enter a custom content type to generate, such as `Social media ad` or `Email body`.	`Description`
4	**With these instructions**	Customize the output by setting the language, tone, writing style, audience, format, length, and custom instructions specific to this generation task.	`Language: English` `Tone: Optimistic` `Custom instructions: Use a few emojis`
5	**Put results in column**	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`F: Description`
6	**Get 2 variants per row** (optional)	Enable this option to generate two variants for each row. The variants are placed in adjacent columns.	*Disabled*



Improve results (Generate)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

You can improve generation quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

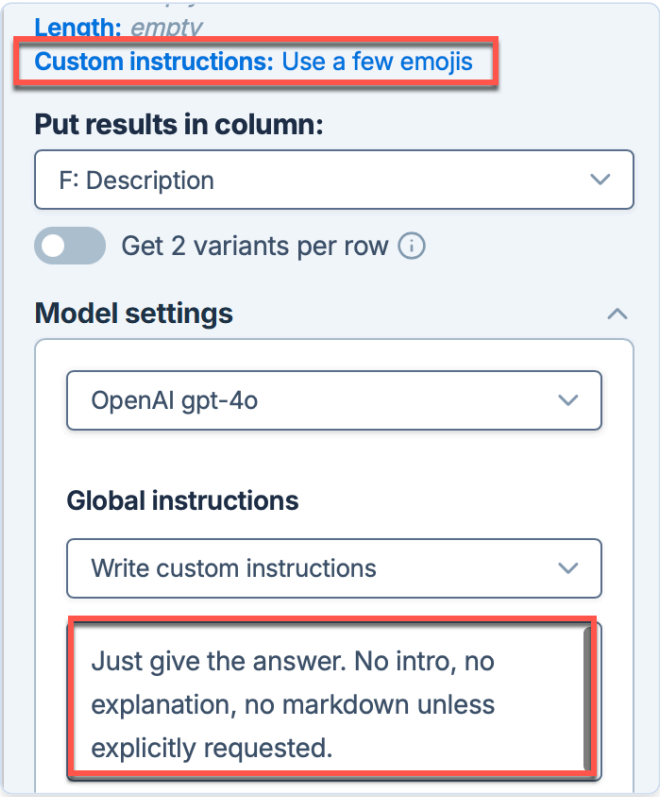
You can define the following types of instructions for the AI:

- **Generation instructions** are defined for the current bulk tool run.
- **Global instructions** are defined in the **Model settings** and apply to all bulk tools and GPT functions in the current workbook.

If both bulk tool instructions and global instructions are defined, the AI uses both as a combined set of instructions. For consistent results, make sure both sets of instructions are aligned and do not contradict each other.

> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to set the current run to use a few emojis and instruct the AI to be concise in all runs:



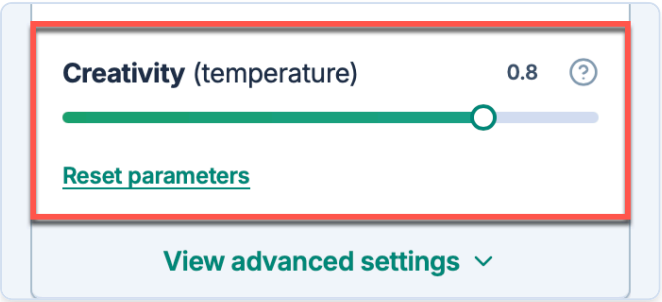
Adjust the creativity level

Adjust the level of creativity (accuracy vs. freedom) the AI is allowed in generating results. The higher the value, the greater the creativity.

For example:

- If you're generating marketing content, set a high creativity level to generate more engaging and imaginative text.
- If you're generating technical specifications, set a low creativity level to ensure the generated text is precise and factual.

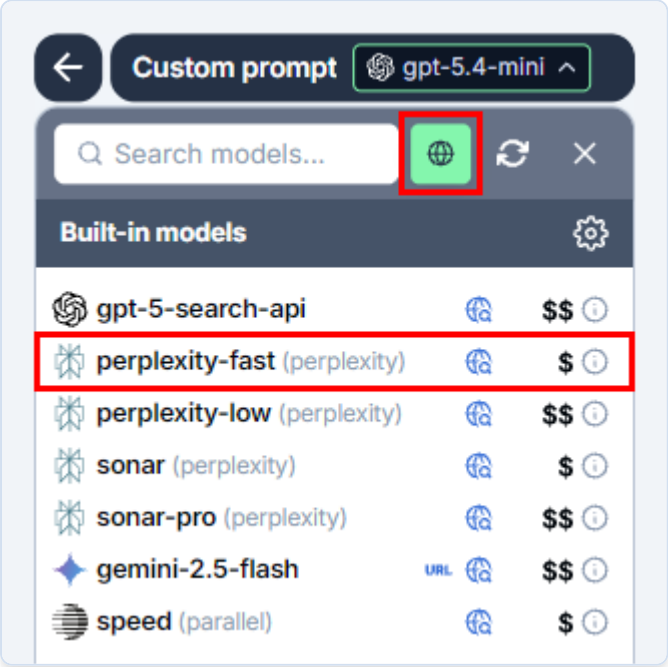
- If you need the AI to follow your instructions to the letter, set the creativity level to **0**.



> **INFO**
>
> The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

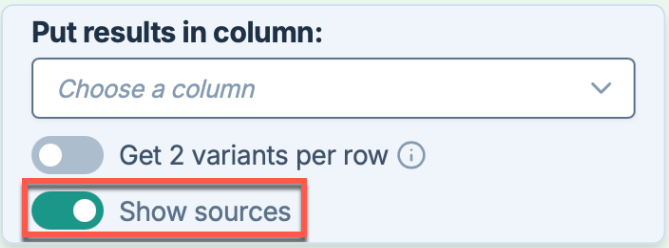
Use web search capabilities

To generate more up-to-date content, select a model with web search capabilities (indicated by the 🌐 icon in the model switcher). The available web search options vary depending on the model you select. Learn more about web search options.



TIP

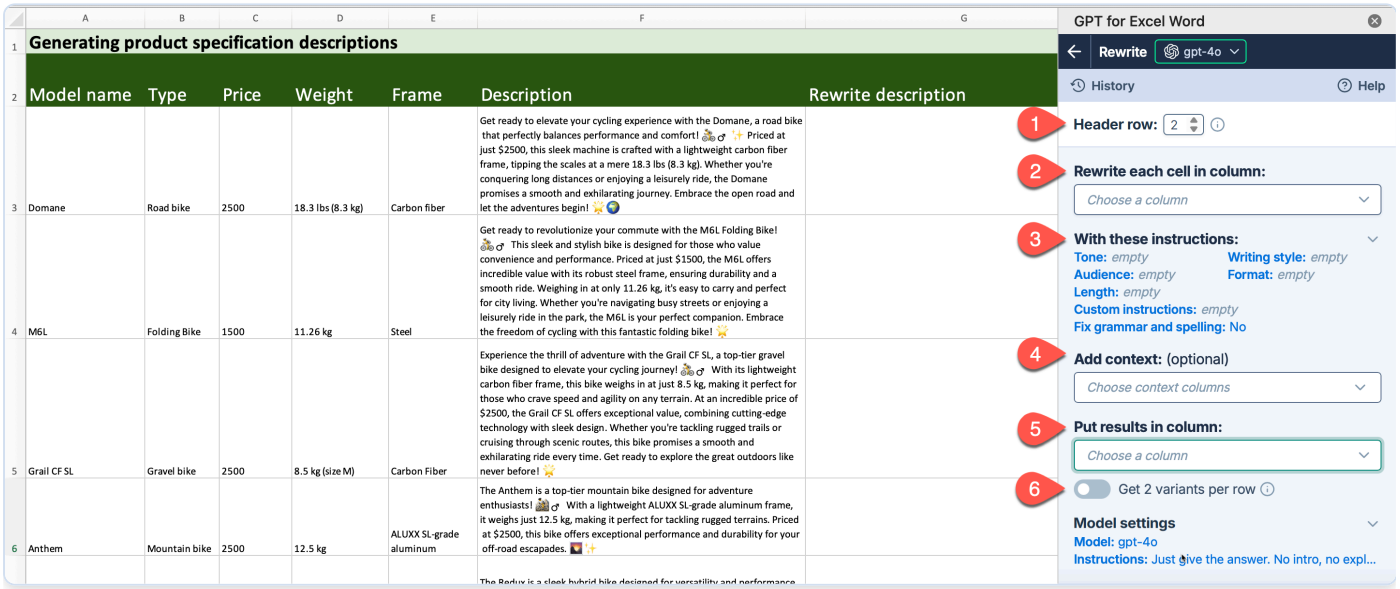
Enable **Show sources** to write the source references in the column to the right of the results column. This helps you understand the origin of the content and assess its reliability.



Rewrite

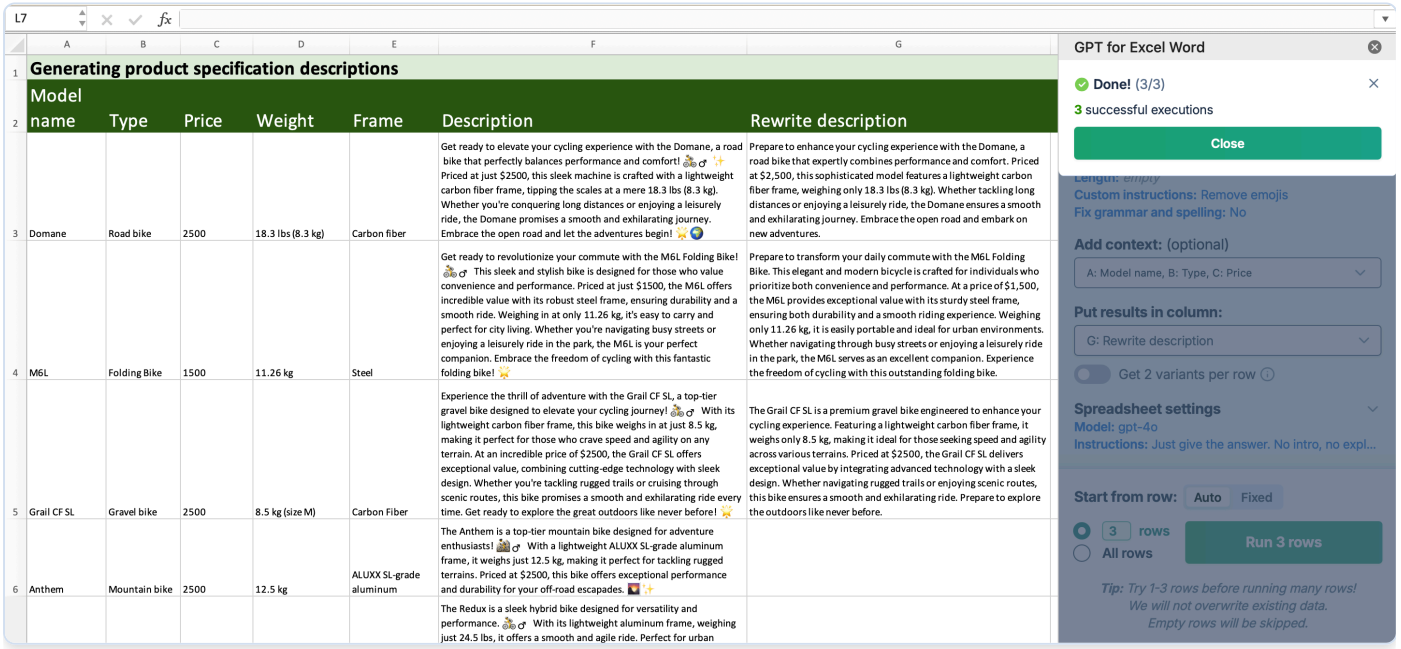
Use the **Rewrite** bulk tool to rewrite existing content in Excel to match your desired tone and style. For example, you can transform casual product descriptions into professional ones.

Settings (Rewrite)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Rewrite each cell in column**	Select the column that contains the text you want to rewrite.	`F: Description`

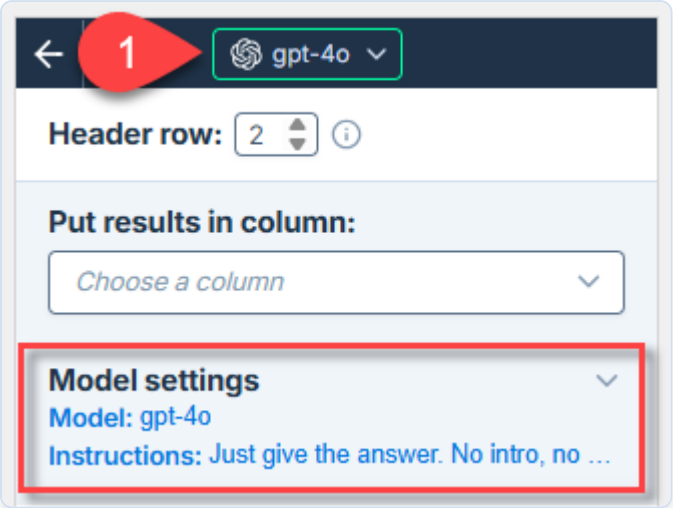
	Field	Description	Example
3	**With these instructions**	Customize how the text should be rewritten by setting the tone, writing style, audience, format, length, and custom instructions, and whether to fix grammar and spelling errors.	`Tone: Professional` `Writing style: Business` `Custom instructions: Remove emojis and casual language`
4	**Add context** (optional)	Select additional columns to provide context for the rewrite.	`A: Model name, B: Type, C: Price`
5	**Put results in column**	Select the column to put the rewritten text in. Cells in this column won't be overwritten if they contain text.	`G: Rewrite description`
6	**Get 2 variants per row** (optional)	Enable this option to generate two rewrite variants for each row. The variants are placed in adjacent columns.	*Disabled*



Improve results (Rewrite)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

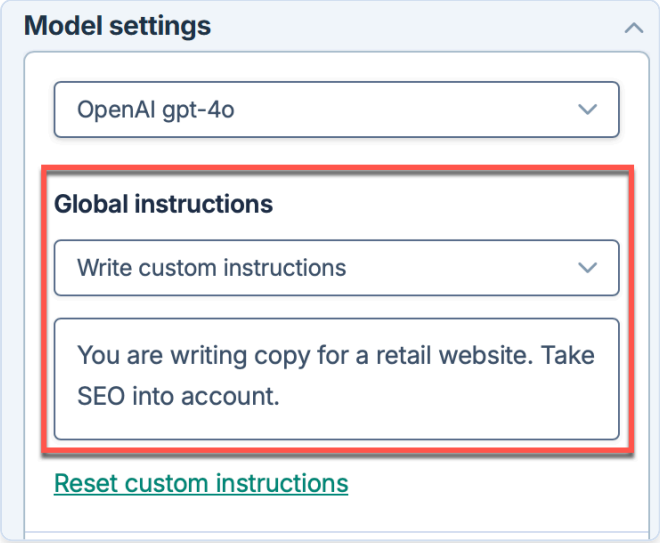
You can improve reformatting quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

Define the instructions as **Global instructions** in the **Model settings**. The instructions apply to all bulk tools and GPT functions in the current workbook.

> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to add instructions to follow SEO best practices:

```
You are writing copy for a retail website. Take SEO into account.
```

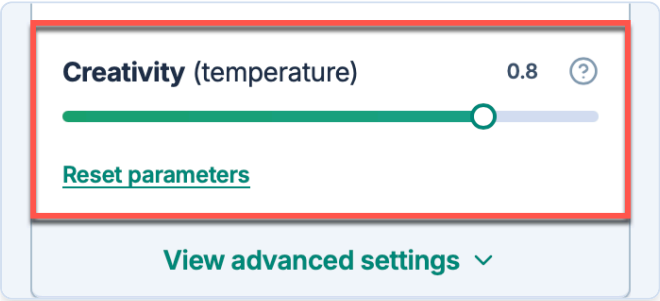


Adjust the creativity level

Adjust the level of creativity (accuracy vs. freedom) the AI is allowed in generating results. The higher the value, the greater the creativity.

For example:

- If you're generating marketing content, set a high creativity level to generate more engaging and imaginative text.
- If you're generating technical specifications, set a low creativity level to ensure the generated text is precise and factual.
- If you need the AI to follow your instructions to the letter, set the creativity level to **0**.



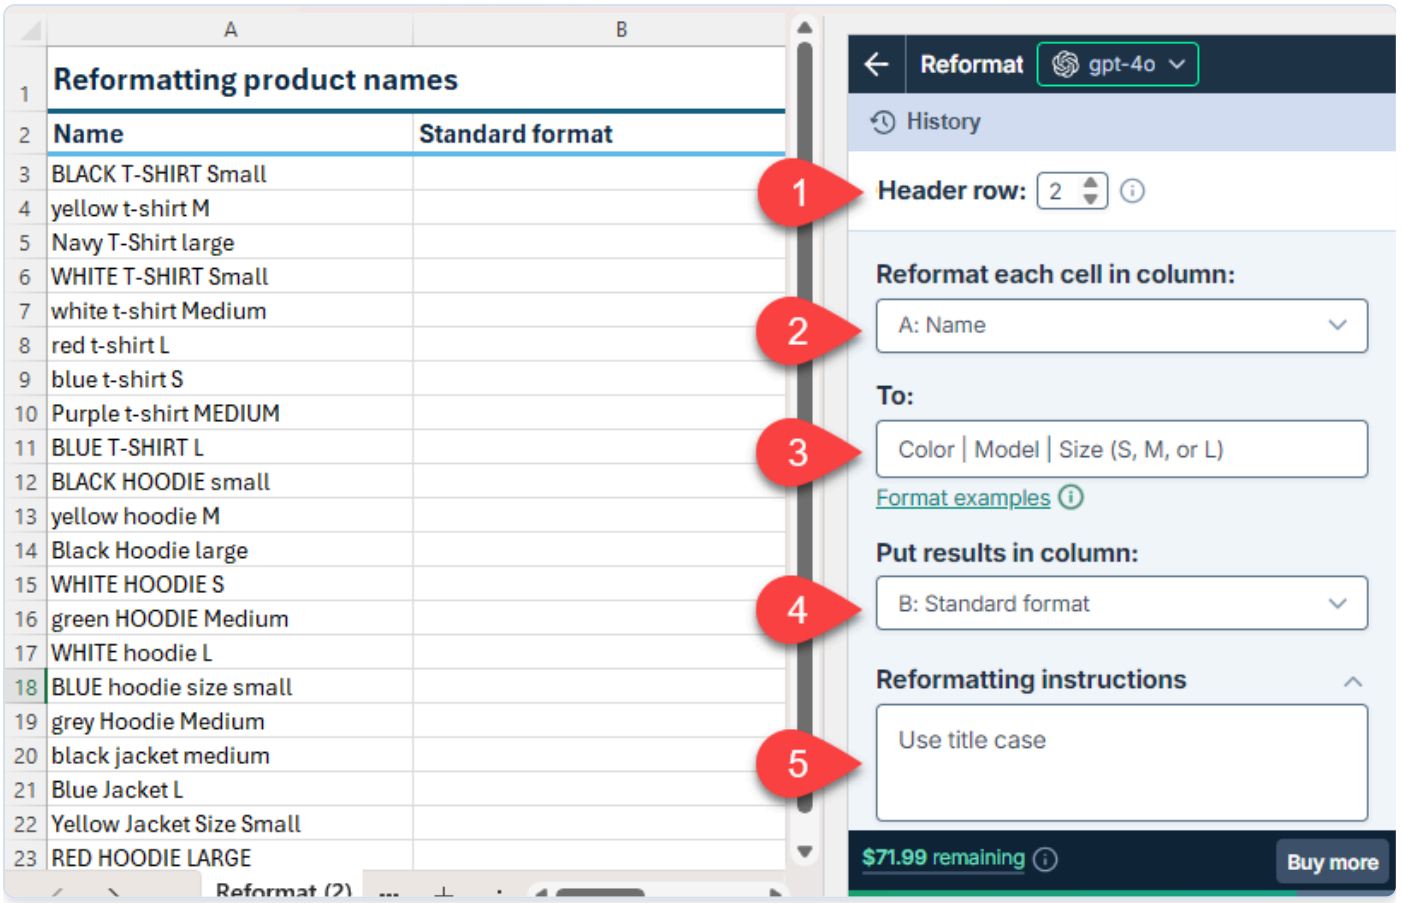
> **INFO**
>
> The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

Reformat ▼

Use the **Reformat** bulk tool to standardize dates, phone numbers, or any other text to a consistent format directly in Excel. For example, you can reformat product names before you include them in

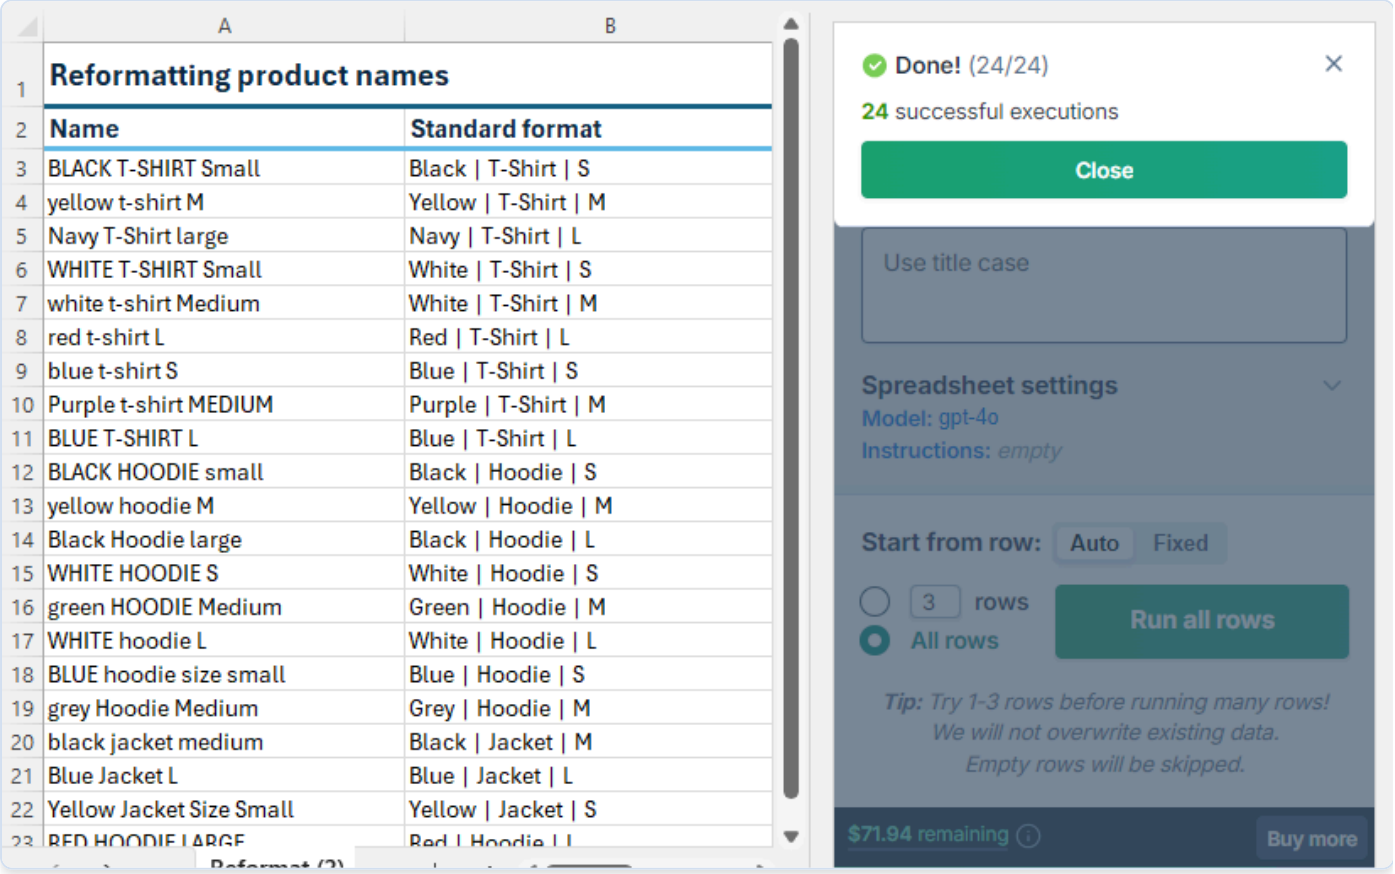
your catalog.

Settings (Reformat)



	Field	Description	Example
1	**Header row** (optional)	If your column headers aren't in the first row, select the number of the row that contains the headers. The bulk tool will run on the rows below this one.	`2`
2	**Reformat each cell in column**	Select the column that contains the text to reformat.	`A: Name`
3	**To**	Describe the format to which you want to convert the text.	`Color | Model | Size (S, M, or L)`

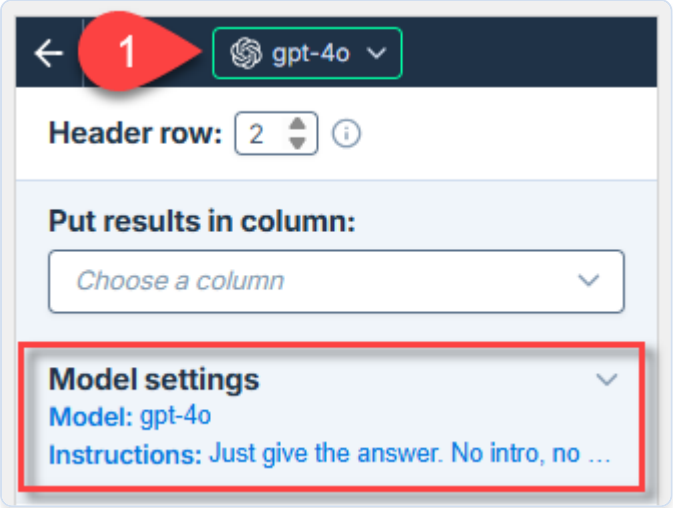
	Field	Description	Example
4	**Put results in column**	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`B: Standard format`
5	**Reformatting instructions** (optional)	Enter specific instructions for the reformatting.	`Use title case`



Improve results (Reformat)

Try different models

Find out which model works best for your use case. You can try different models by selecting them in the model switcher.



> **INFO**
>
> The model selection applies to all bulk tool runs and all GPT function executions in the current workbook.

Provide additional instructions

You can improve response quality by providing specific instructions that help the AI more closely follow your requirements for tone, style, format, rules, target audience, context, consistency, and other aspects.

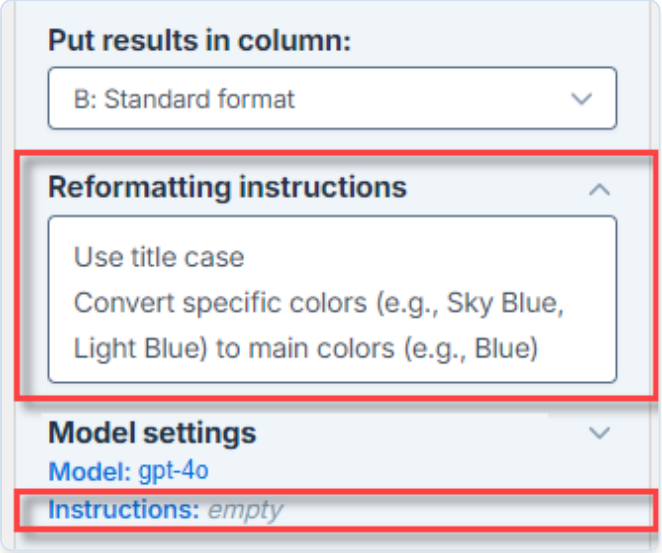
You can define the following types of instructions for the AI:

- **Reformatting instructions** are defined for the current bulk tool run.
- **Global instructions** are defined in the **Model settings** and apply to all bulk tools and GPT functions in the current workbook.

If both bulk tool instructions and global instructions are defined, the AI uses both as a combined set of instructions. For consistent results, make sure both sets of instructions are aligned and do not contradict each other.

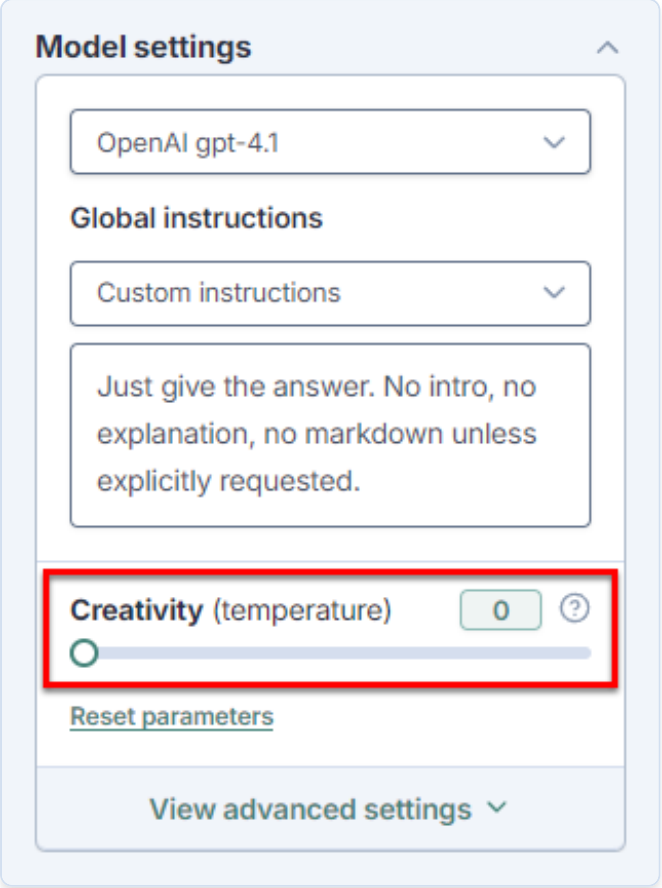
> **INFO**
>
> Global instructions apply to all bulk tool runs and all GPT function executions in the current workbook.

For example, to omit the global instructions and only define reformatting instructions:



Set Creativity to 0

Ensure **Creativity** is set to **0**. This helps the AI follow instructions more closely.



INFO

The creativity level applies to all bulk tool runs and all GPT function executions in the current workbook.

What's next

- Select a model.
- Search the web with bulk tools.
- Use images with bulk tools.
- Configure bulk tool behavior.

Select a model for bulk tools in Excel

GPT for Excel supports a range of models that you can use with the bulk tools. To help you find the model that offers the best balance of accuracy, speed, and cost for your use case, see our recommended models.

The available models depend on:

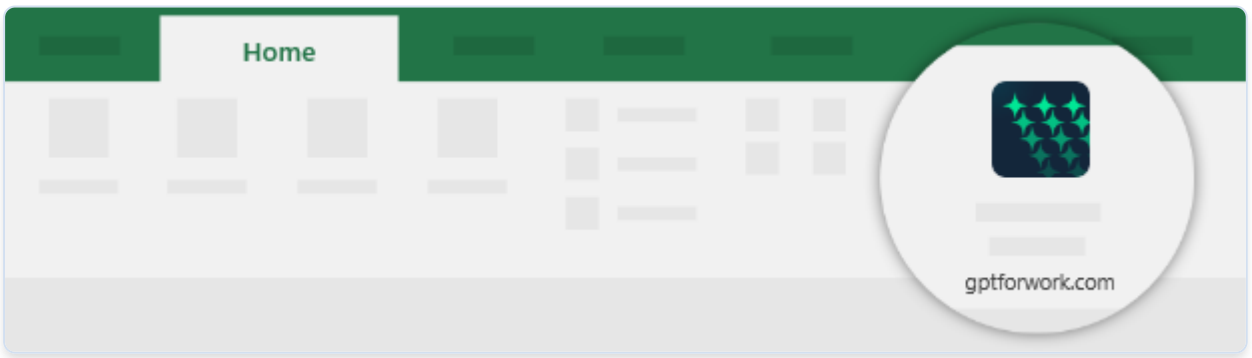
- **The API keys set for your current space.** GPT for Excel gives you access to a default set of built-in models. To access additional models, set an API key for one or more AI providers.
- **The API endpoints set up for your current space.** To access even more models, set up a connection to an Azure OpenAI API endpoint or any OpenAI-compatible API endpoint.
- **The model settings for your current space.** Space owners and admins can control which models are available in GPT for Excel.

Recommended models

For information about recommended bulk models, see Recommended models.

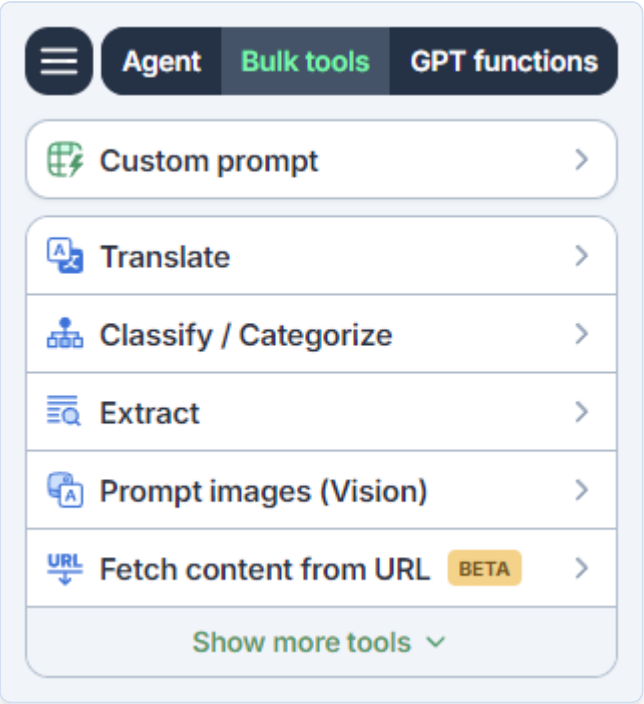
Select a built-in model

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.

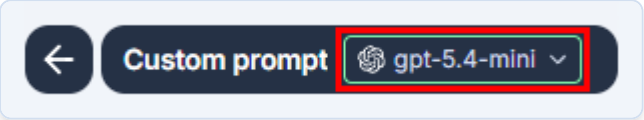


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

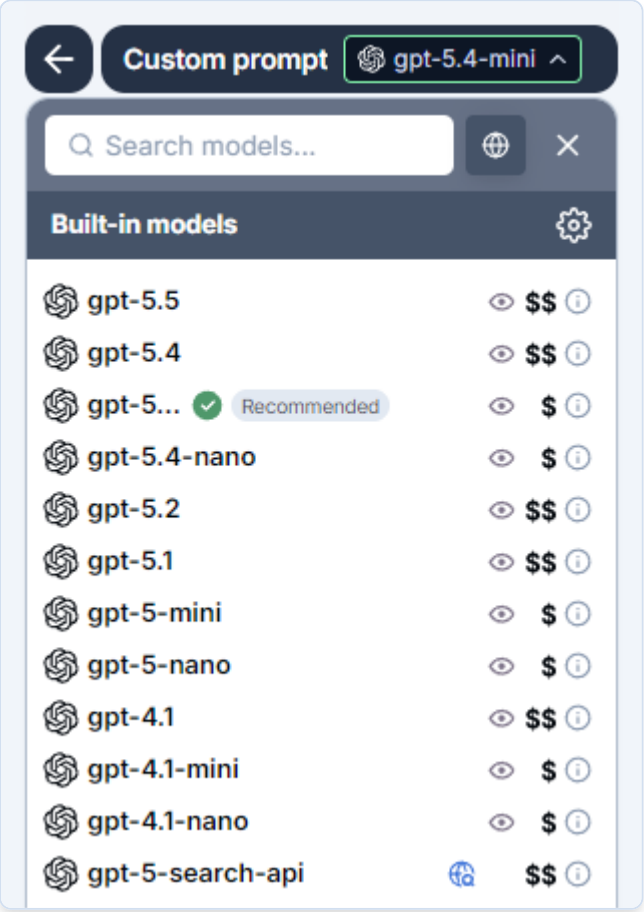
3 In the sidebar, select **Bulk tools** and select any bulk tool.



4 Open the model switcher.



5 Under **Built-in models**, select a model.



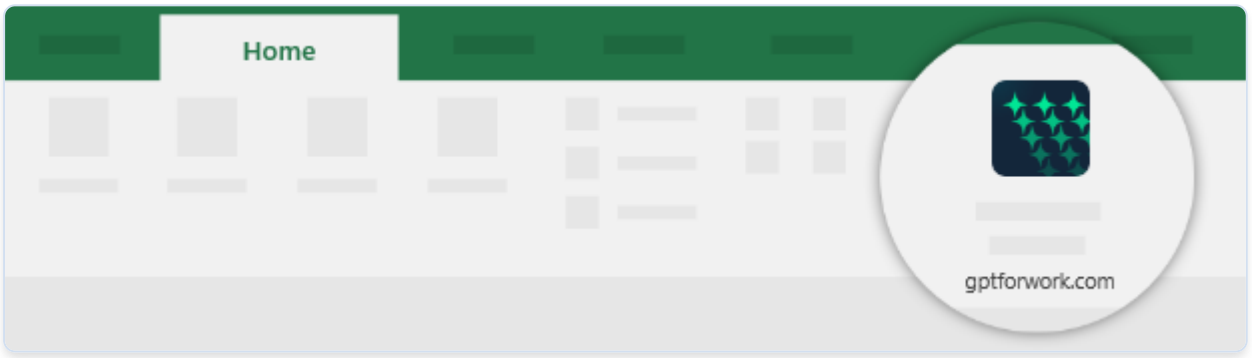
GPT for Excel now uses the selected model to generate responses for bulk tools.

Select a connected model

Prerequisites

- You have set an API key for the AI provider whose model you want to use.
- If you want to use an Azure model or an OpenAI-compatible model from an API endpoint, you have set up the correct endpoint.

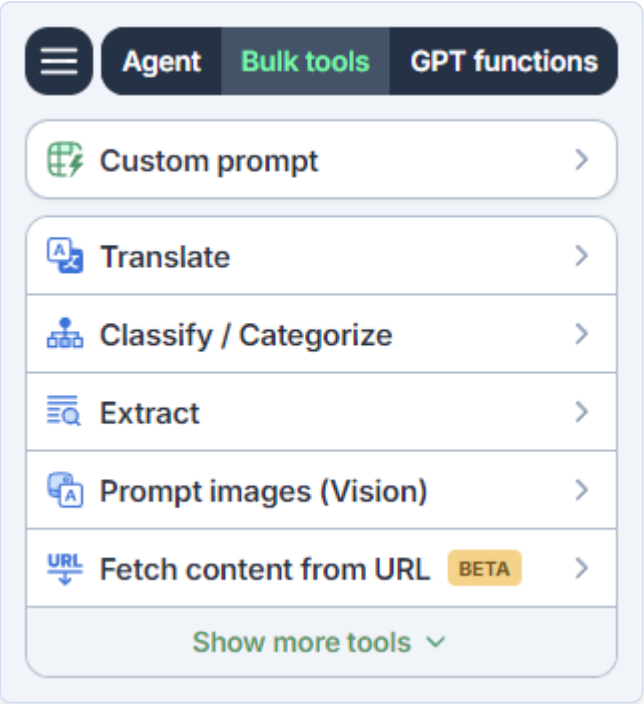
1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



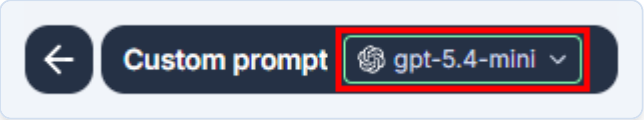
INFO

If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

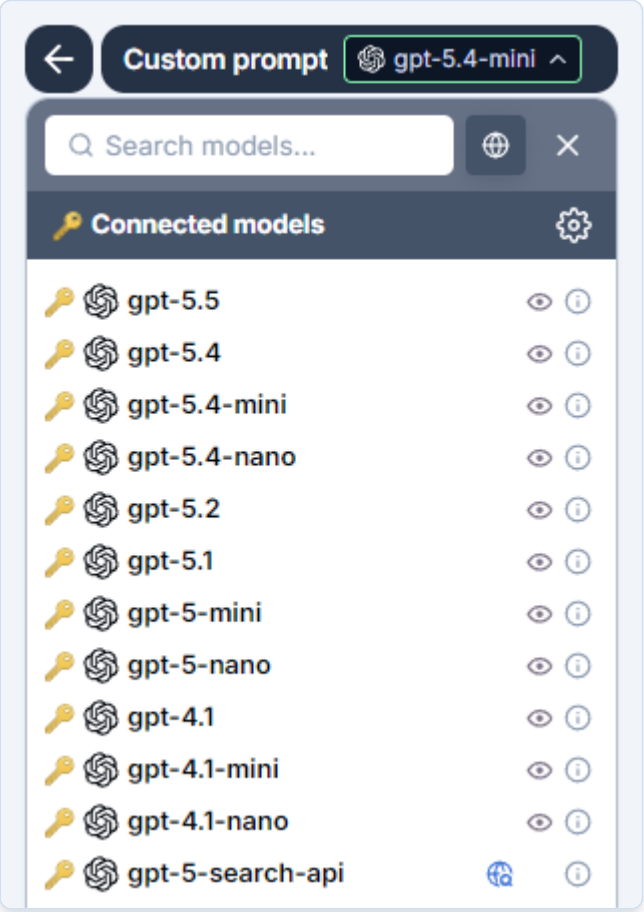
3 In the sidebar, select **Bulk tools** and select any bulk tool.



4 Open the model switcher.



5 Under **Connected models**, select a model.



GPT for Excel now uses the selected model to generate responses for bulk tools.

What's next

Learn more about models:

- AI providers and models supported by GPT for Work
- LiveBench (a good model benchmark)

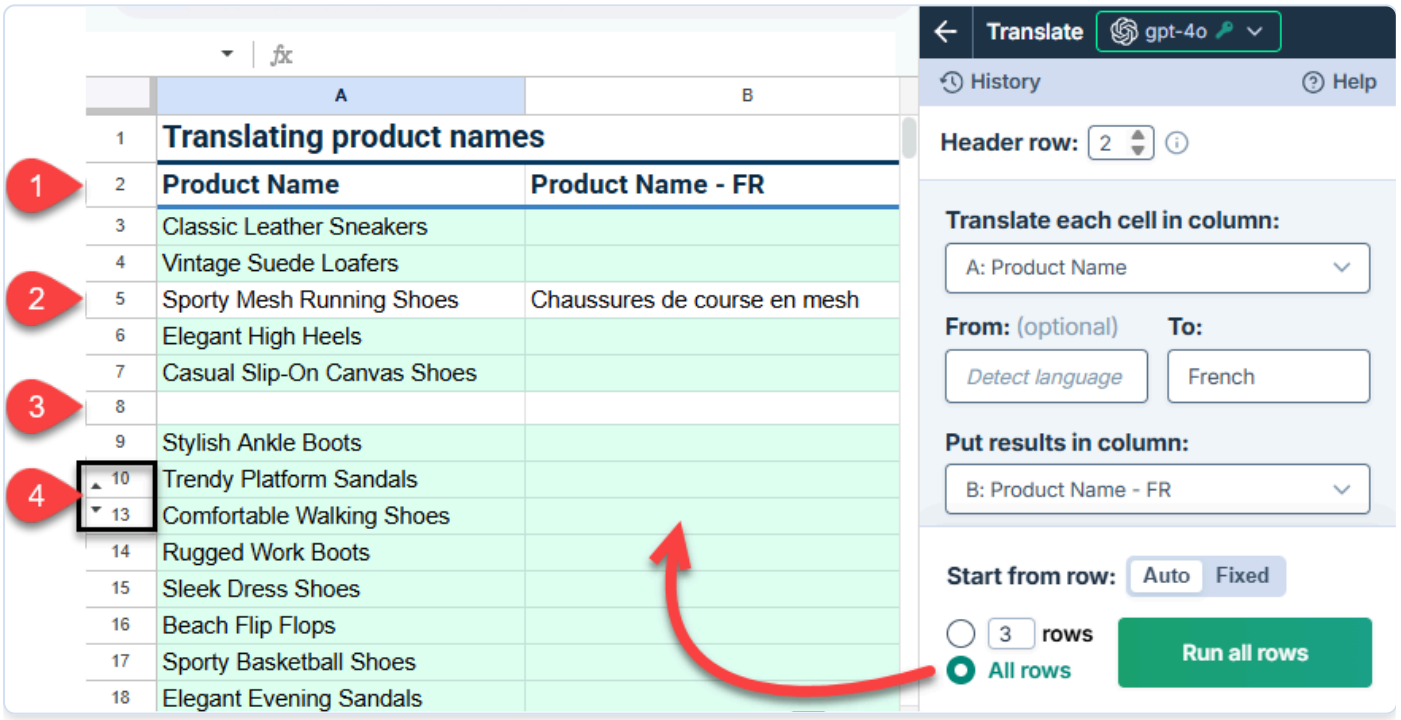
Select which rows to process with bulk tools in Excel

Bulk tools in GPT for Excel are designed to process only the relevant rows in your spreadsheet. You can additionally control where to start and how many rows to process, making it easy to test on a small sample before processing your entire dataset.

Bulk tools automatically skip:

1. The header row
2. Rows with content in the result column
3. Rows where all input columns are empty
4. Hidden or filtered rows

In the figure below, all rows with a green background will be processed by the Translate bulk tool.



Try a bulk tool on a few rows, then run all rows

The default settings for bulk tool runs is ideal for starting with a quick test:

Setting	Value	Description
Header row	1	Row 1 contains the column headers.

Setting	Value	Description
Start from row	Auto	The tool processes the first 3 rows after the header row. If a row has a non-empty cell in the result column, the tool skips to the next row.
X rows	3	The tool processes 3 rows after the header row.

To run a bulk tool with the default settings, simply click **Run 3 rows**.

The following video shows you how to run the Extract tool on 3 rows to check that the specifications for your product descriptions are extracted correctly.

When you're ready to run the tool on the whole spreadsheet, change the row setting to **All rows** and click **Run all rows**. The tool skips the first 3 rows as they have already been processed.

Run a bulk tool on new content only

If you've already run a bulk tool on a spreadsheet and then add new data, you don't need to filter the new rows to process them. Just run the tool again on all rows. It skips the rows that have already been processed.

Setting	Value	Description
Start from row	Auto	The tool starts on the first empty cell in the result column.
X rows	All rows	The tool processes all remaining rows, skipping the rows that have already been processed.

The following video shows how to run the Classify / Categorize tool to process only new rows in a product list. Even if the list is already sorted so that the new rows are not all at the end, the tool only processes the new rows.

Run a targeted test

You can target a specific section of rows in a spreadsheet by explicitly defining the start row and the number of rows to process:

Setting	Value	Description
Start from row	Fixed: 5	The tool starts on the first empty cell in the result column on or after row 5.
X rows	2	The tool processes 2 rows, starting from the first empty cell in the result column on or after row 5.

The following video shows you how to run the Translate tool on 2 rows to check that your glossary is taken into account when translating your product names from English to French.

Rerun a bulk tool on a few cells

To rerun a bulk tool on a few cells, simply delete the cell content and rerun the tool. This allows you to easily update a few selected rows without affecting the rest of the results:

Setting	Value	Description
Start from row	Auto	The tool starts at the first empty cell in the result column.
X rows	All rows	The tool processes all rows with empty cells in the result column.

The following video shows how to rerun the Custom tool on a few cells to update the results after tweaking your prompt.

> **TIP**
>
> To rerun a bulk tool on the same rows, you can select a new column to write the results to. This way, you can keep the original results and compare them with the new ones.

What's next

- Discover all bulk tools available in GPT for Excel.
- If you encounter issues with bulk tools, see the troubleshooting guide.

Search the web with bulk tools in Excel

Search the web directly in GPT for Excel using bulk tools with web search models. These models use fresh web search data, which allows you to generate up-to-date responses about companies, current events, facts, and more, directly in your workbook. You can also fetch content from specific URLs to extract information from web pages.

Choose the right web search model

Choose a web search model that best fits your use case.

Web search

Get information that is not in the model's built-in knowledge:

- **Example:** Finding the current CEO of a company.
- **How it works:**

  All web search models can go beyond their built-in knowledge by accessing information from search engines, like Google or Bing. This allows models, which have compressed historical knowledge, to explore topics in greater depth, with more current data or with a specific focus.
- **Recommended models:**
  - perplexity-fast for quick results (cheapest).
  - OpenAI Search models for high quality results.
  - Gemini models with Web search for Google search quality (most expensive).

Content fetching

Focus on specific URLs and extract content from them:

- **Example:** Extracting the product name and price from the URL of a product page.
- **How it works:** Some web search models can access the full content from specific URLs. The full content is added to the context of the models, so they can use it to answer questions about the content or extract information from it.
- **Recommended models:**
  - Gemini 2.5 models with URL context

Supported web search models

Click here for a list of all web search models supported by GPT for Excel ▼

Provider	Model	Built-in	Provider-connected 🔑	URL fetch
Search...	Search...	All	All	All
Perplexity	perplexity-fast	Yes	Yes	No
Perplexity	perplexity-low	Yes	Yes	No
Perplexity	perplexity-medium	No	Yes	No
Google	gemini-2.5-flash	Yes	Yes	Yes
Google	gemini-2.5-flash-lite	No	Yes	Yes
Google	gemini-2.5-pro	No	Yes	Yes
Google	gemini-3-flash-preview	No	Yes	Yes
Google	gemini-3.1-flash-lite	No	Yes	Yes
Google	gemini-3.1-pro-preview	No	Yes	Yes
Google	gemini-3.5-flash	No	Yes	Yes
OpenAI	gpt-5-search-api	Yes	Yes	No
Parallel	speed	Yes	No	No

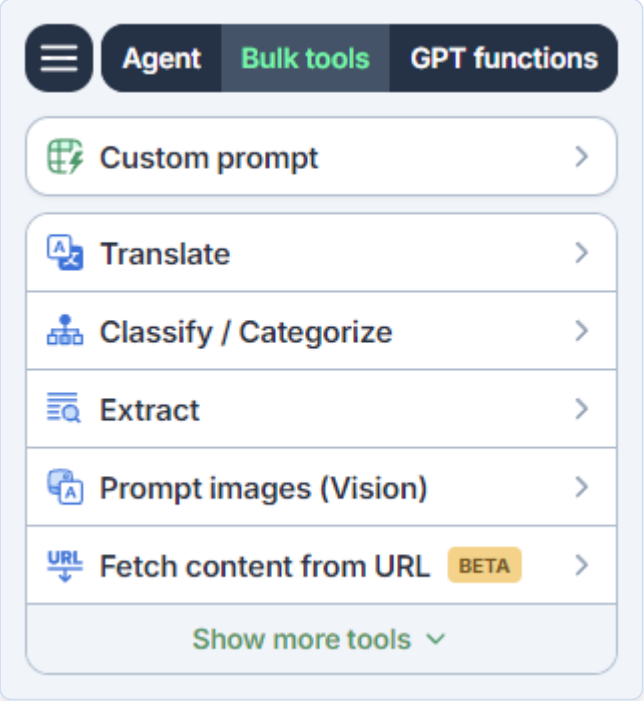
For detailed pricing information, see Model prices.

Use web search models

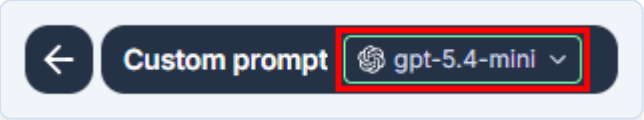
Use bulk tools to perform web searches with any model that has web search capabilities (🌐).

1 Open GPT for Excel.

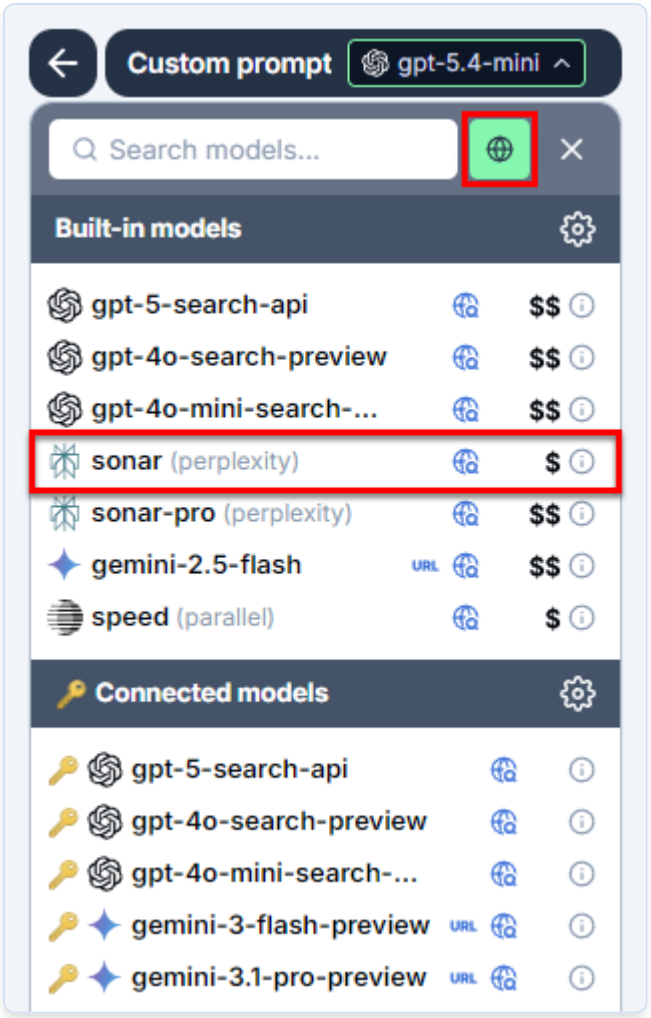
2 In the sidebar, select **Bulk tools** and select the bulk tool you want to use.



3 Open the model switcher.

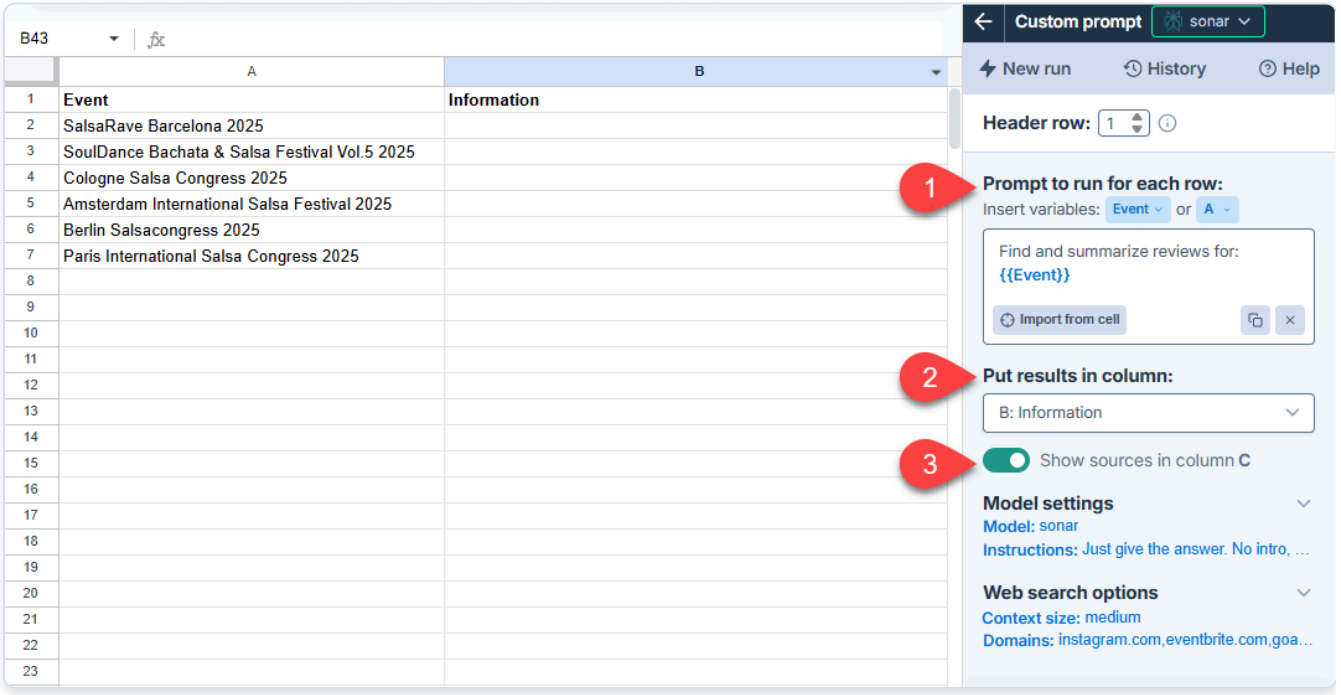


4 Enable the web search filter and select a web search model (🌐).



5 (Optional) Configure the model-specific settings in the sidebar. For Gemini models, make sure web search is enabled.

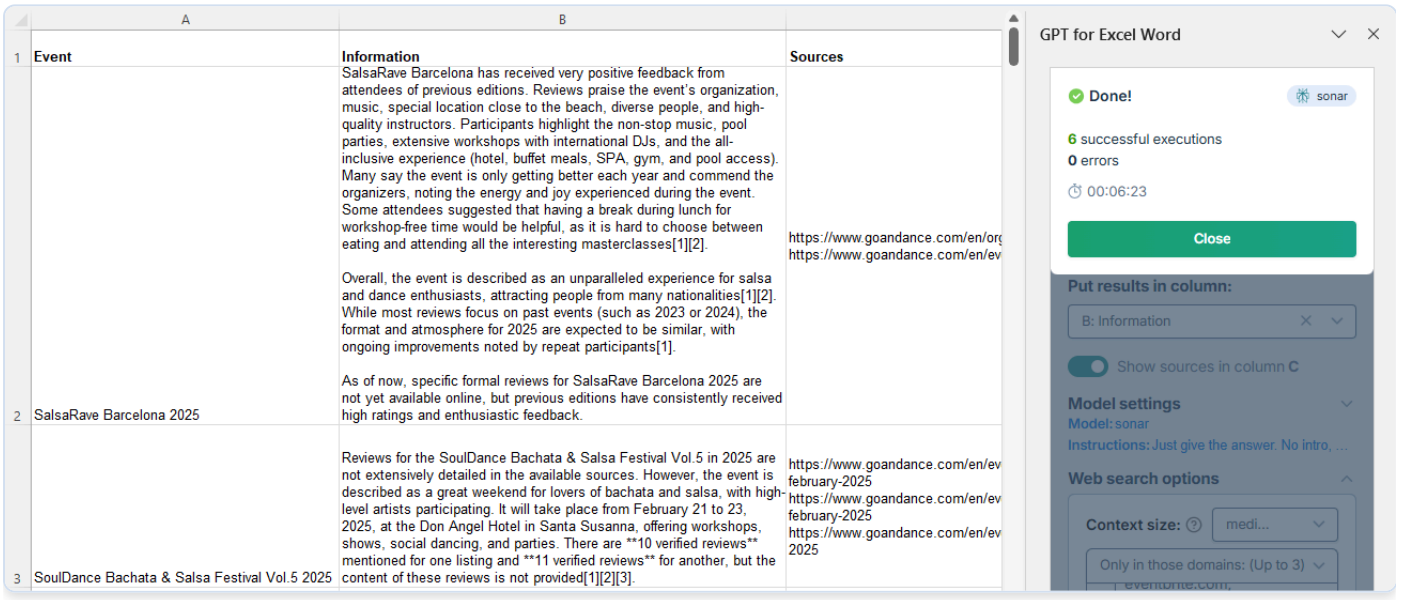
6 Set up your bulk tool run. For example, to set up a **Custom prompt** run:



	Field	Description	Example
1	**Prompt to run for each row**	Enter the prompt you want to run for each row.	`Find and summarize reviews for: {'{{Event}}'}.`
2	**Put results in column**	Select the column to put the results in. Cells in these columns won't be overwritten with the results if they contain text.	`B: Information`
3	**Show sources**	Enable this option to write source references in the column to the right of the results column.	**Enabled**

7 Click **Run rows**.

Your search results appear in the output column. If you enabled source references, they appear in the adjacent column.

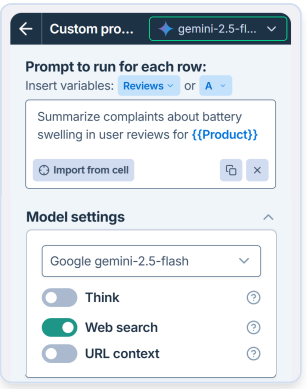
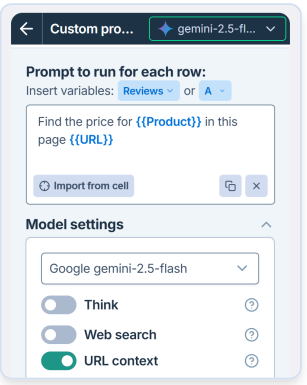


For instructions on how to use the tools, see Bulk tools.

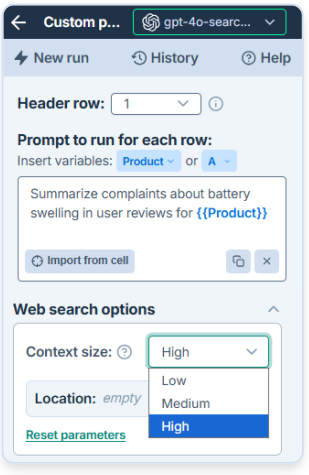
Model-specific settings

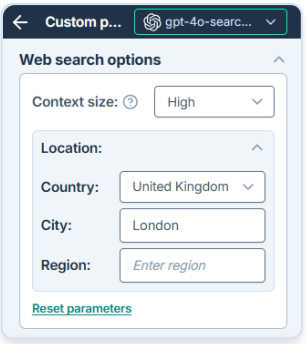
Use the model-specific settings to configure the web search model for your use case.

Google Gemini

Setting	Description	Example
Web search	Let Gemini decide when a Google search can improve the answer. This allows Gemini to provide more accurate answers and cite verifiable sources beyond its built-in knowledge. Enabling Web search incurs an additional search cost.	
URL context	Provide up to 20 URLs (pages, images, PDFs) in your prompt to focus your search. No search cost is incurred, but the content retrieved from URLs counts as input tokens.	

OpenAI Search

Setting	Description	Example
Context size	Choose how much content to retrieve from each source. Larger sizes produce richer and more detailed responses at a higher cost.	

Setting	Description	Example
Location	To refine search results based on geography, specify an approximate user location using country, city, and/or region.	

Use images in bulk tool prompts in Excel

Use images in your prompts with GPT for Excel and vision models. These models are multimodal and can process both text and images as input.

> **TIP**
>
> Don't have image URLs yet? Learn how to host images on Google Drive for AI processing.

> **INFO**
>
> Supported image formats: PNG, JPEG, WebP, non-animated GIF

Supported vision models

We recommend using **gpt-5.4** to analyze images, but you can use any vision model supported by GPT for Excel.

Click here for a list of all vision models supported by GPT for Excel ▼

Provider	Model	Built-in	Provider-connected 🔑
Search...	Search...	All	All
OpenAI	gpt-4.1	Yes	Yes
OpenAI	gpt-4.1-mini	Yes	Yes
OpenAI	gpt-4.1-nano	Yes	Yes
OpenAI	gpt-4o	Yes	Yes
OpenAI	gpt-5	Yes	Yes
OpenAI	gpt-5-mini	Yes	Yes
OpenAI	gpt-5-nano	Yes	Yes
OpenAI	gpt-5.1	Yes	Yes
OpenAI	gpt-5.2	Yes	Yes
OpenAI	gpt-5.4	Yes	Yes

Provider	Model	Built-in	Provider-connected 🔑
OpenAI	gpt-5.4-mini	Yes	Yes
OpenAI	gpt-5.4-nano	Yes	Yes
OpenAI	gpt-5.5	Yes	Yes
OpenAI	gpt-5.6-luna	Yes	Yes
OpenAI	gpt-5.6-sol	Yes	Yes
OpenAI	gpt-5.6-terra	Yes	Yes
OpenAI	gpt-6-astra	No	Yes
OpenAI	gpt-6-luna	Yes	Yes
OpenAI	gpt-6.1-sol	Yes	Yes
OpenAI	o1	Yes	Yes
OpenAI	o3	Yes	Yes
OpenAI	o4-mini	Yes	Yes
Anthropic	claude-4.1-opus	No	Yes
Anthropic	claude-4.5-haiku	Yes	Yes
Anthropic	claude-4.5-opus	Yes	Yes
Anthropic	claude-4.5-sonnet	Yes	Yes
Anthropic	claude-4.6-opus	Yes	Yes
Anthropic	claude-4.6-sonnet	Yes	Yes
Anthropic	claude-4.7-opus	Yes	Yes
Anthropic	claude-4.8-opus	Yes	Yes
Anthropic	claude-5-opus	Yes	Yes
Anthropic	claude-5-sonnet	Yes	Yes
Anthropic	claude-5.5-haiku	Yes	Yes

Provider	Model	Built-in	Provider-connected 🔑
Anthropic	claude-5.5-opus	Yes	Yes
Anthropic	claude-5.5-sonnet	Yes	Yes
Anthropic	claude-fable-5	No	Yes
Anthropic	claude-fable-5.1	No	Yes
Mistral	mistral-medium	No	Yes
xAI	grok-4-1-fast	No	Yes
xAI	grok-4-1-fast-non-reasoning	No	Yes

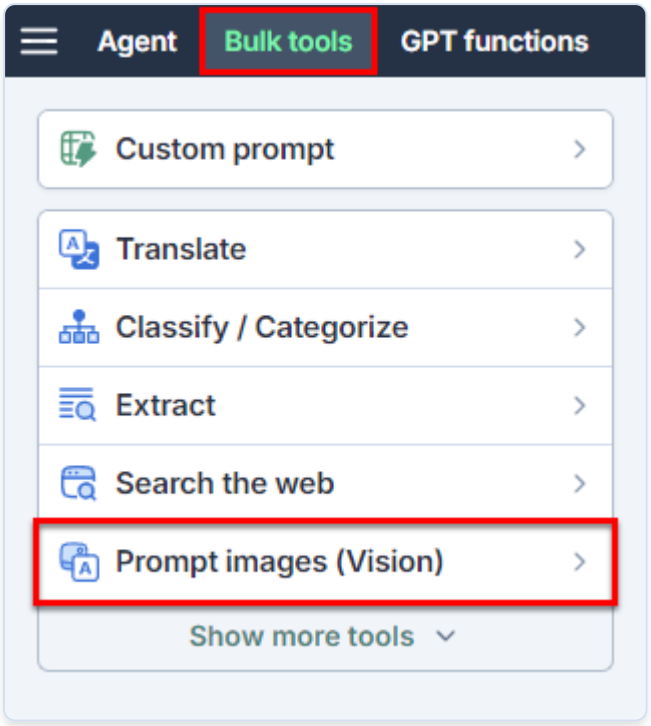
Prompt images

> **TIP**
>
> Want to use both text and images in your prompt? Skip to Prompt images with additional text input.

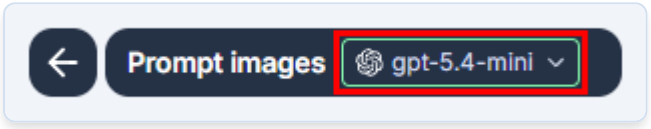
Use the **Prompt images (Vision)** bulk tool to apply a prompt to images in your spreadsheet.

1. Open GPT for Excel.

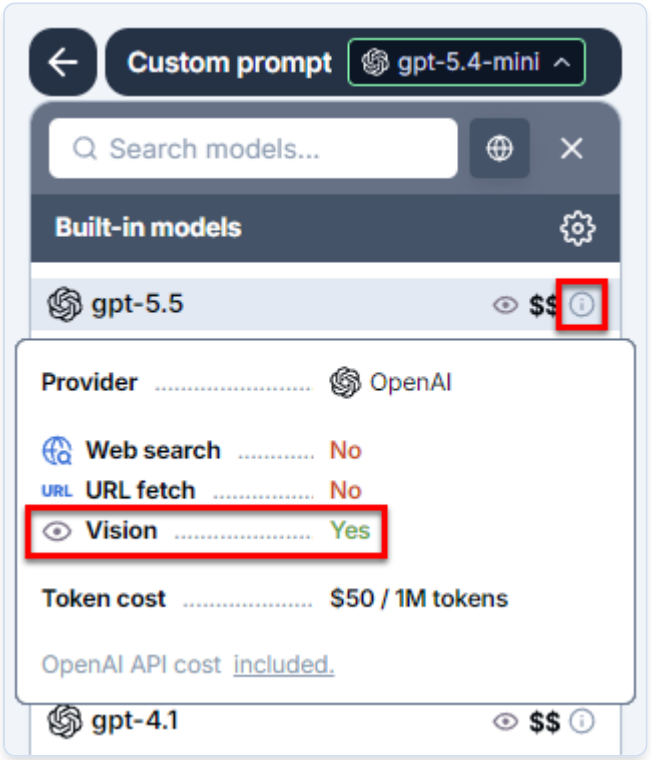
2 In the sidebar, select **Bulk tools** and select **Prompt images (Vision)**.



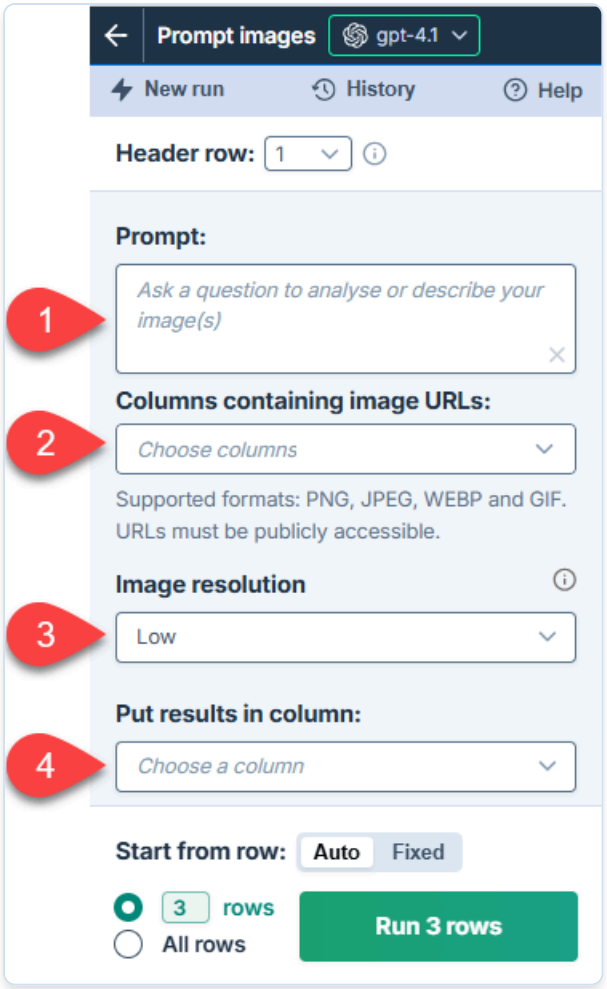
3 Open the model switcher.



4 Select a vision model.



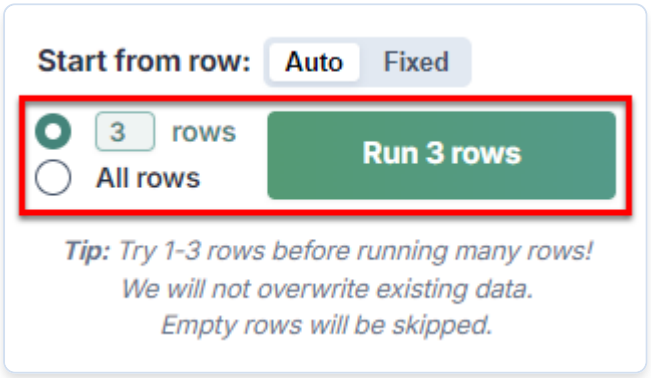
5 Set up your bulk tool run.



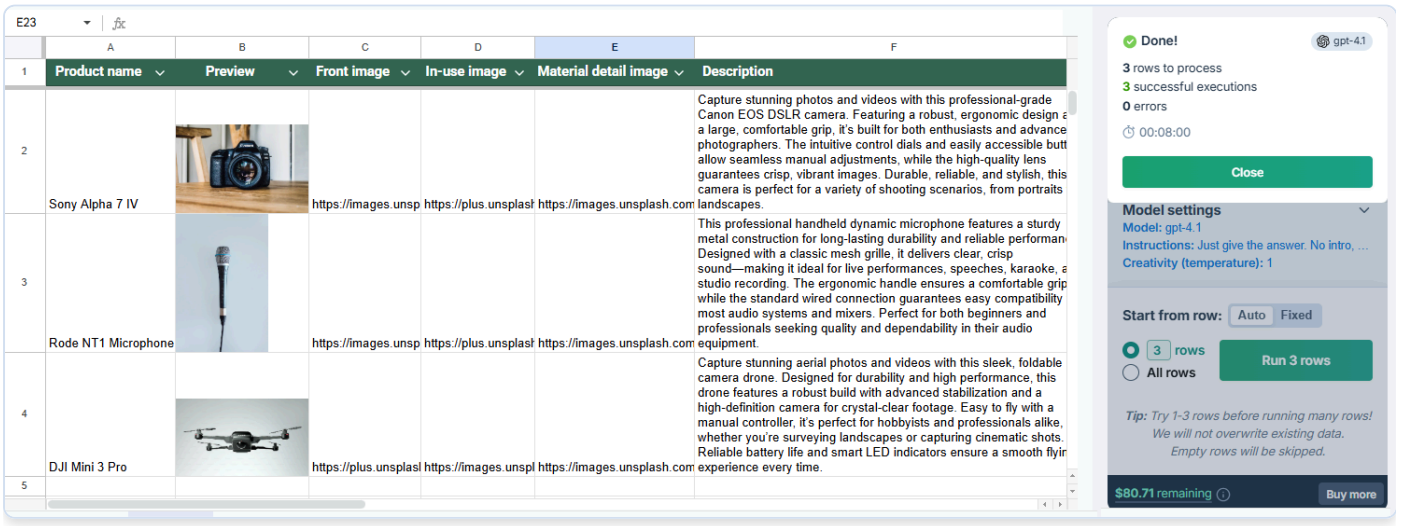
	Field	Description	Example
1	**Prompt**	Enter the prompt you want to apply to each image.	`Write a short description of this product based on how it looks, how it's used, and its build quality.`
2	**Columns containing image URLs**	Select up to five columns that contain image URLs. For each row, the tool processes all images present in the selected columns. URLs must be publicly accessible.	`C: Front image, D: In-use image, E: Material detail image`

	Field	Description	Example
3	**Image resolution**	Choose the level of detail for image processing: ○ **High:** Better image understanding, higher latency and cost ○ **Low:** Faster processing, lower cost ○ **Auto:** Let the model decide (default)	`Low`
4	**Put results in column**	Select the column to put the results in. Cells in this column won't be overwritten with the results if they contain text.	`F: Description`

6 Click **Run rows**.



You have set up and run the **Prompt images (Vision)** bulk tool to apply a prompt to images in your spreadsheet.



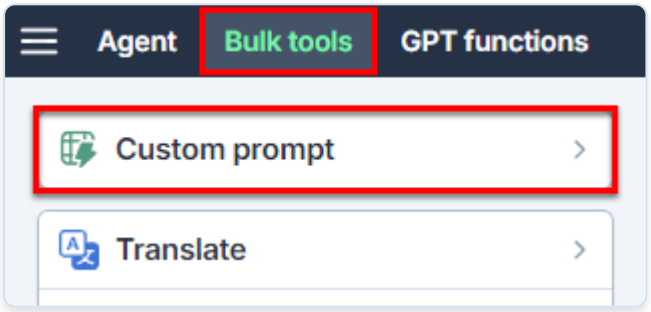
For instructions on how to improve the results, see the tool guide.

Prompt images with additional text input

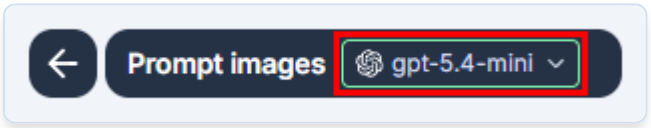
Use the **Custom prompt** bulk tool to combine a prompt with text and images from your spreadsheet.

> **INFO**
>
> In the **Custom prompt** bulk tool, the image resolution is fixed to **Low** and cannot be changed.

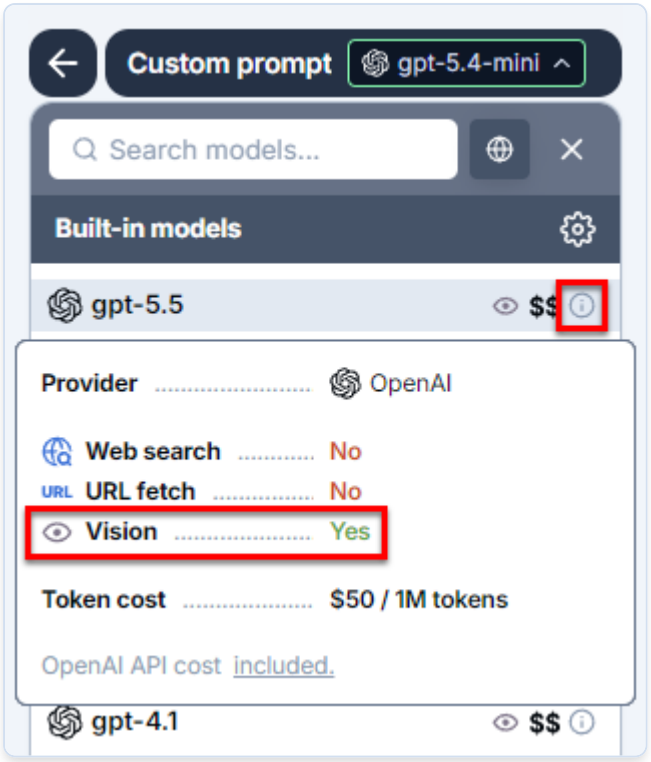
1. Open GPT for Excel.
2. In the sidebar, select **Bulk tools** and select **Custom prompt**.



3. Open the model switcher.



4 Select a vision model.

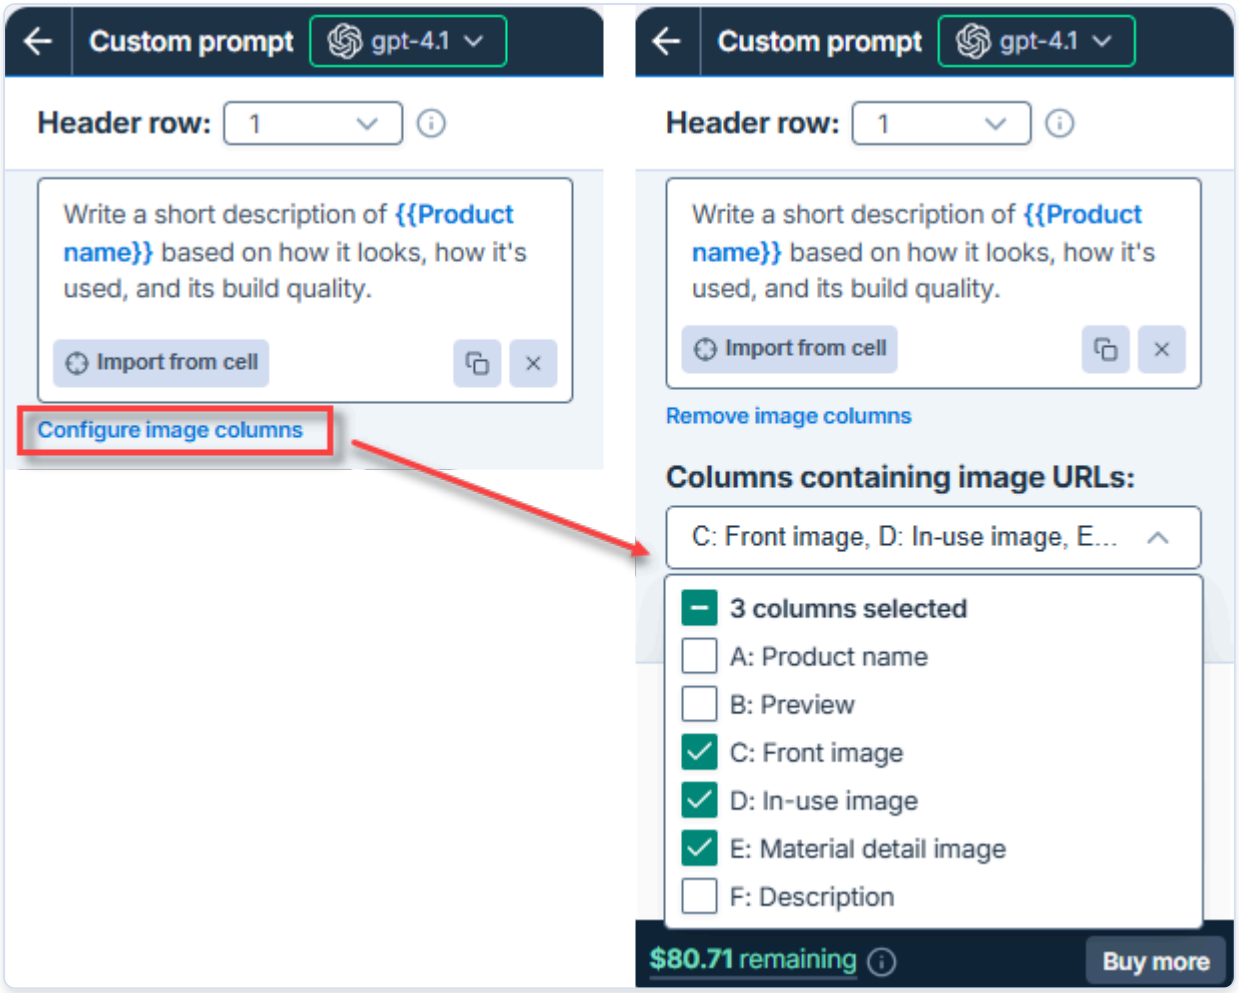


5 Write your prompt and include any variables to be replaced with values from your spreadsheet.

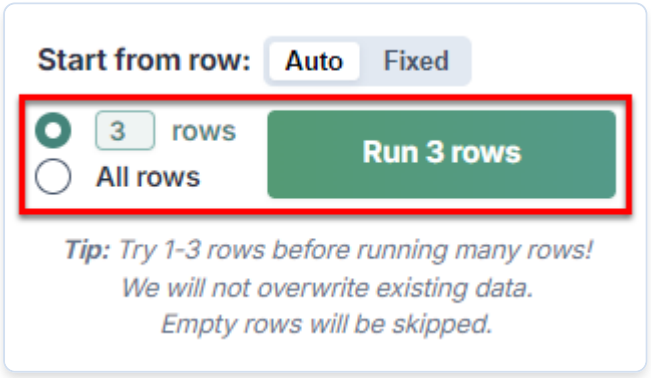
```
Write a short description of {{Product name}} based on how it looks, how it's used, and its build quality.
```

6 Click **Configure image columns** to show the **Columns containing image URLs** dropdown,

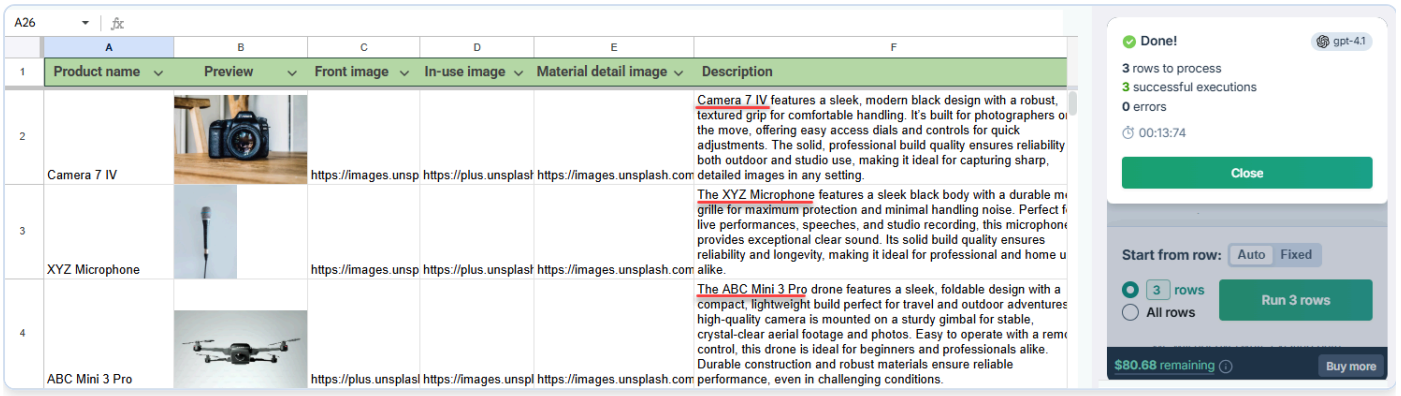
and select the columns that contain image URLs.



7 Click **Run rows**.



You have set up and run the **Custom prompt** bulk tool to include images and text from your spreadsheet.



For instructions on how to improve the results, see the tool guide.

Host images on Google Drive for AI processing

To process an image in GPT for Excel, you need the direct URL of the image and the URL must be publicly accessible.

To host an image on Google Drive so that GPT for Excel can access and process it:

1. Upload the image to Google Drive.
2. Share the image publicly using the **Anyone with the link** option. For more information, see Google Drive Help.

   > **TIP**
   >
   > To prevent other users from commenting or editing the image, set the access role to **Viewer**.

3. Copy the link.
4. Get the image file ID from the link URL, and build the following new URL with the ID:

   ```
   https://drive.google.com/uc?id=<image-file-id>
   ```

   For example:

   ```
   # Original share URL from "Copy link"
   https://drive.google.com/file/d/1dTZrIkAHsV_hKnvL2KjH1pfcOOSpYVux/view

   # Direct image URL
   https://drive.google.com/uc?id=1dTZrIkAHsV_hKnvL2KjH1pfcOOSpYVux
   ```

You now have a direct image URL that you can access and process in GPT for Excel.

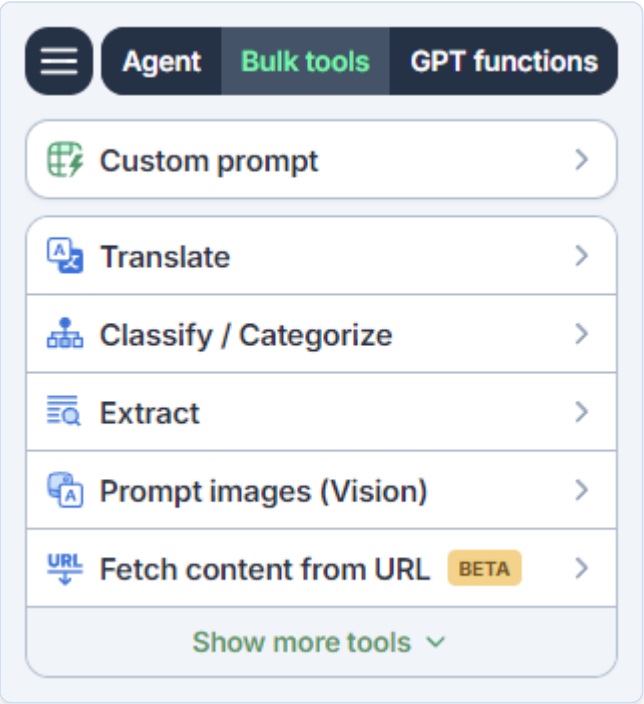
Configure bulk tool behavior in Excel

Configure model settings to customize how bulk tools behave in your workbook.

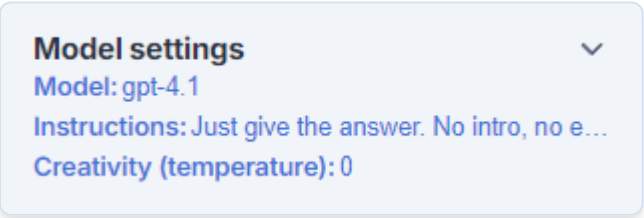
Add global instructions

Provide global instructions to specify preferences or requirements that you'd like the AI to consider when generating responses. The instructions apply to all bulk tools in the current workbook.

1. In the sidebar, select **Bulk tools** and select one of the tools.

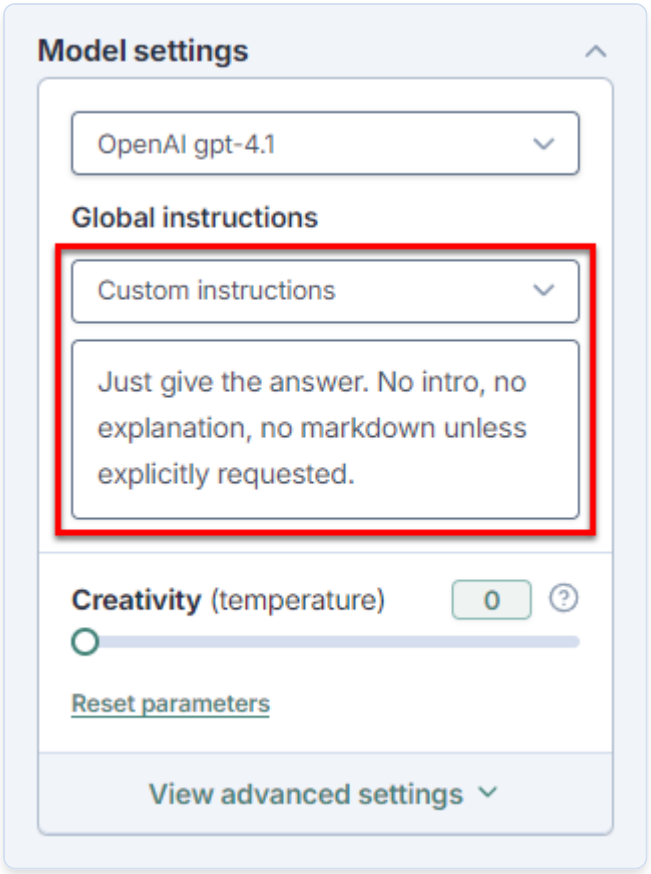


2. Expand **Model settings**.



3. In the **Global instructions** section, select the type of instructions you'd like to add and edit

them if needed.



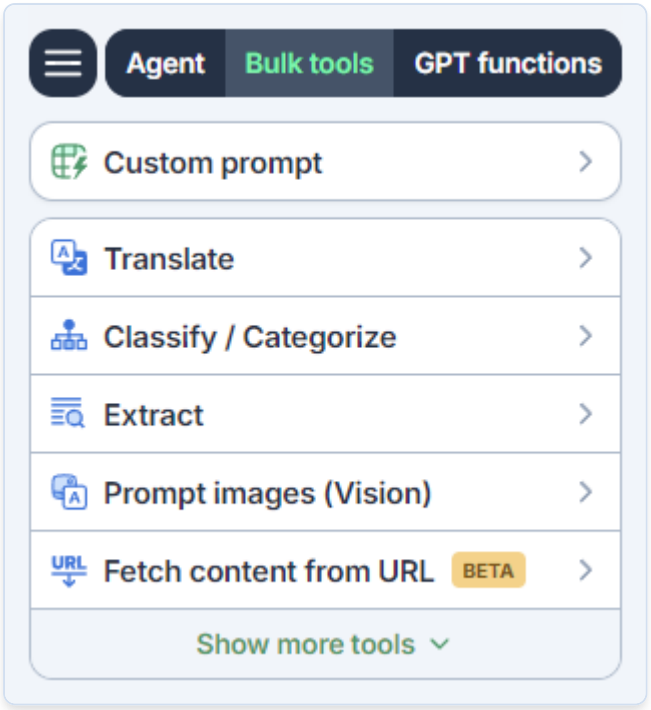
The AI takes the global instructions into account when generating responses for bulk tools in the current workbook.

Set the creativity level

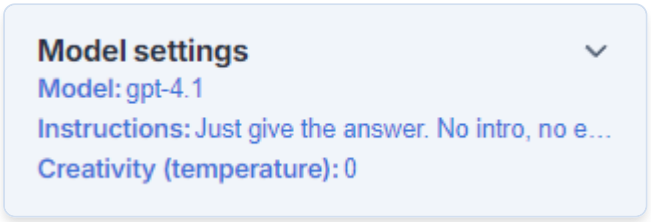
Control how creative the AI is by setting the **creativity level** (also called **temperature**). Lower creativity generates responses that are straightforward and predictable, while higher creativity generates responses that are more random and varied.

> **TIP**
>
> For more information about temperature as used by OpenAI to set the creativity level, see our temperature guide.

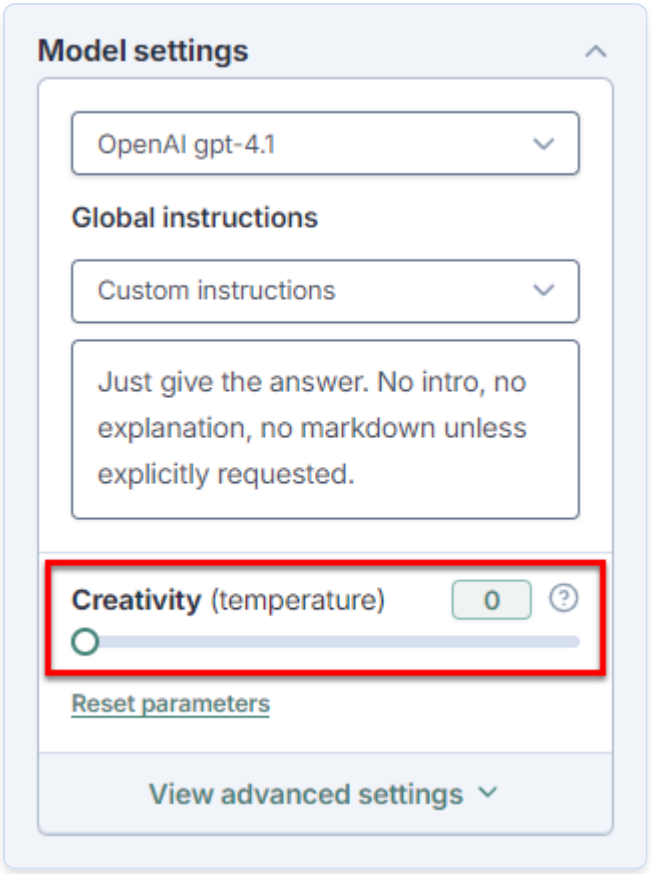
1 In the sidebar, select **Bulk tools** and select one of the tools.



2 Expand **Model settings**.



3 Use the slider to set **Creativity** between **0** and **1**.



The AI uses the creativity level when generating responses for bulk tools in the current workbook.

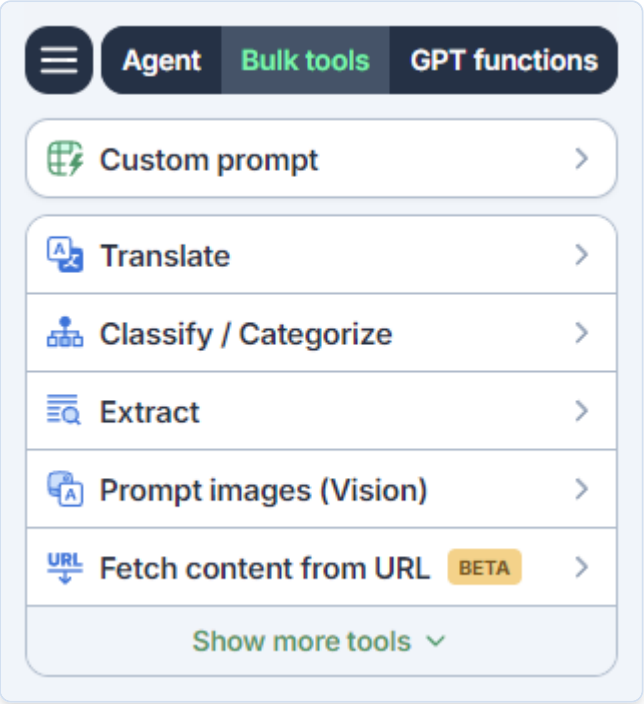
Adjust advanced settings

Adjust the advanced settings to have greater control over generated output for bulk tools in the current workbook. For more information about the settings, see the OpenAI API documentation.

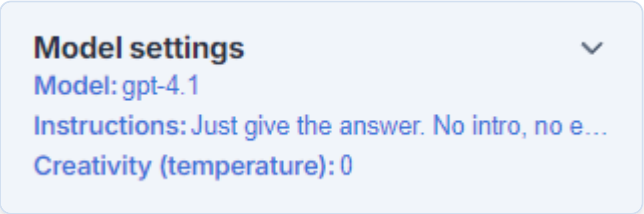
Term	Definition
Token	A token is a unit of text, such as a word, subword, punctuation mark, or symbol. Models split all input and output text into tokens for processing. Model providers use tokens as the basic unit for measuring model usage. Input tokens from your prompt and output tokens from the model's response contribute to your credit consumption. For more information, see Tokens and How credits are consumed.
Top P	Helps adjust creativity level – Lower values result in more focused output, while higher values allow for more creative responses.

Term	Definition
Frequency penalty (not available for all models)	Penalizes tokens based on their frequency in the text so far – Higher values discourage the model from repeating the same tokens too frequently, as the more they appear in the text, the more penalized they get.
Presence penalty (not available for all models)	Penalizes new tokens based on whether they appear in the text so far – Higher values encourage the model to use new tokens that are not penalized.

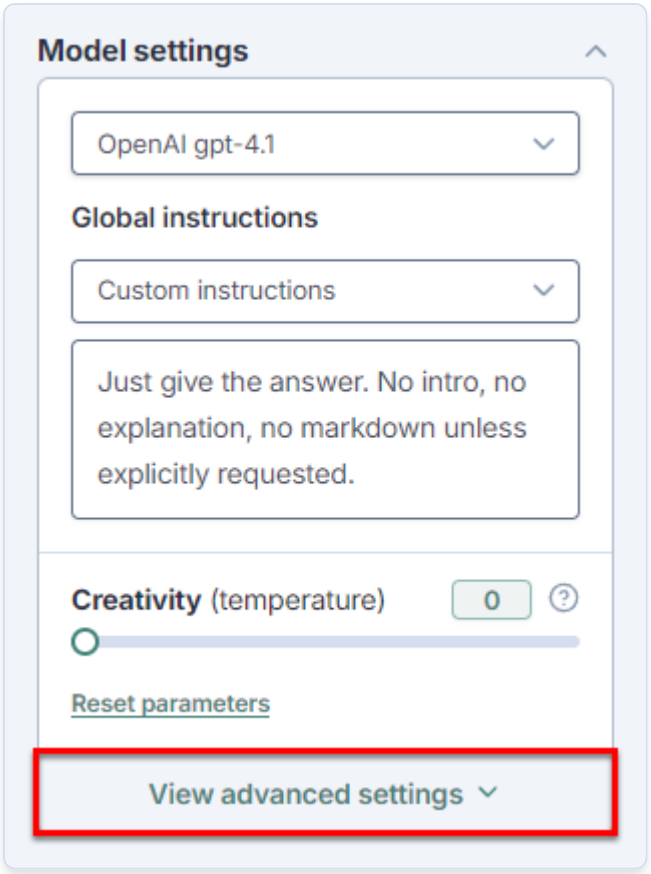
1. In the sidebar, select **Bulk tools** and select one of the tools.



2. Expand **Model settings**.



3 Expand **View advanced settings**.

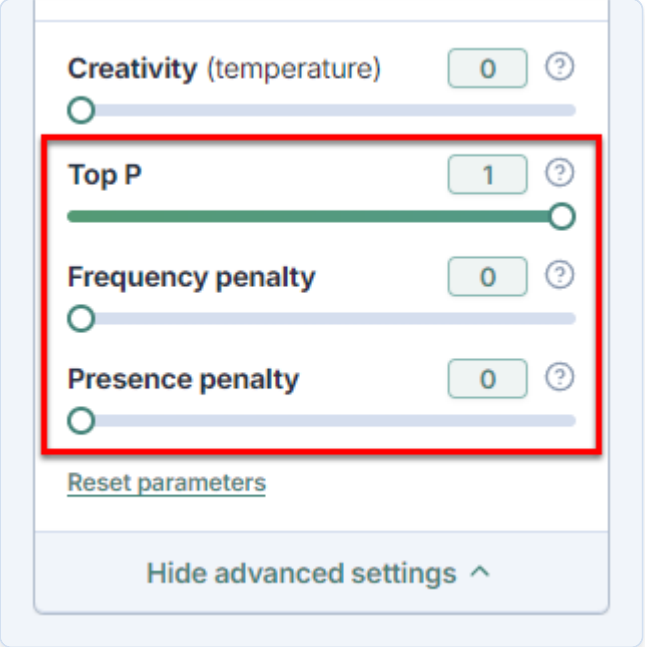


4 Use the corresponding sliders to set the values for:

- **Top P**: Adjust creativity by setting the value between **0** and **1**.

  Set the creativity level to **1** when the value for **Top P** is lower than **1**.
- **Frequency penalty**: Set a penalty between **0** and **2** to apply to tokens each time they are repeated in the result.
- **Presence penalty**: Set a one-time penalty between **0** and **2** for tokens that appear more

than once in the result.



GPT for Excel uses the advanced settings to generate responses for bulk tools in the current workbook (when a model supporting the settings is selected).

What's next

- Select the model that best fits your needs.
- Run bulk tools.

GPT functions in Excel

Build formulas with GPT functions to prompt AI from inside spreadsheet cells. Use the GPT for Excel sidebar to manage your GPT formulas.

Get started

Learn the basics of using GPT functions — from enabling GPT functions for a workbook to building formulas that combine GPT functions with native functions.

GPT functions reference

All the GPT functions, with simple examples and an extensive list of parameters.

Select a model

Select the model used by GPT functions in GPT for Excel.

Create GPT formulas

Create GPT formulas in Microsoft Excel and control the execution of GPT formulas.

Manage GPT formulas

Manage your GPT formulas and avoid unexpected changes and costs.

Search the web

Learn how to use web search models in GPT functions to get up-to-date information from the web.

Use images in prompts

Include images in GPT function prompts in Microsoft Excel.

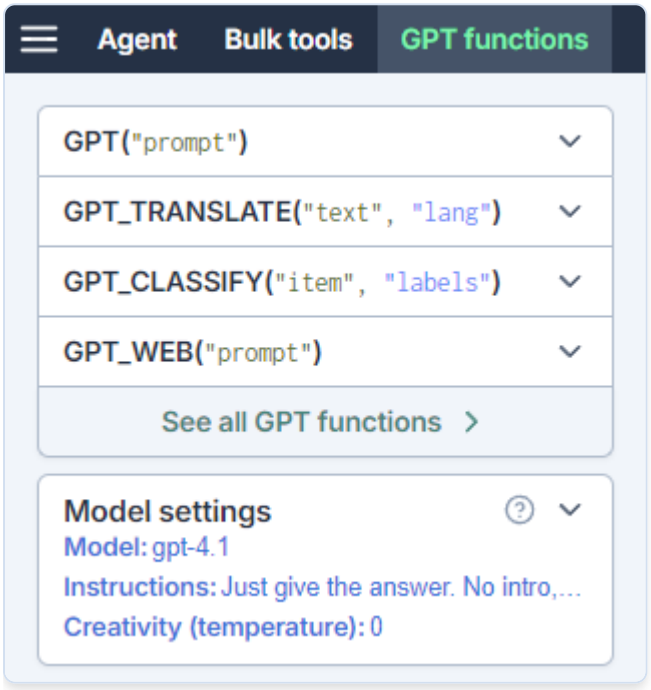
Configure behavior

Configure model settings to customize how GPT functions behave in Microsoft Excel.

Get started with GPT functions in Excel

GPT functions are custom spreadsheet functions in GPT for Excel that allow you to prompt AI from inside spreadsheet cells. GPT functions work exactly like native functions in that you can use them on their own or combine them with other functions when creating formulas.

GPT for Excel provides dedicated GPT functions for selected use cases, such as translation, classification, and web search, as well as a generic GPT function for all other use cases. You can define function-specific parameters to further refine how the AI generates responses. You can find a list of all available functions with usage examples in the add-in sidebar.



This guide walks you through the basics of using GPT functions — from enabling GPT functions for a workbook through building formulas with GPT functions to using functions for web search and image processing. The focus is on the `GPT` function, the simplest function included in GPT for Excel, but the same rules apply to all GPT functions.

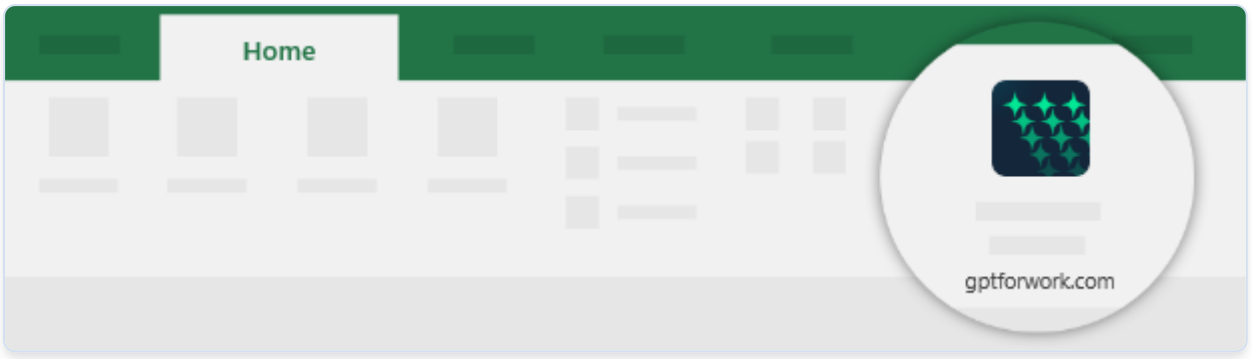
> **TIP**
>
> To try out the examples in this guide, download our GPT function starter examples template, and open it in Excel as a new workbook.

Enable GPT functions

GPT for Excel must be running before you can use GPT functions for the first time in a workbook.

To start GPT for Excel:

1. Open the workbook where you want to use GPT functions.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

Use GPT functions

Learn the basics of using GPT functions, from how you type a function through how you use functions in formulas to how you search the web and process images with functions.

GPT function syntax

Each GPT function follows a specific syntax that consists of a **function name** and one or more **parameters** (also called "arguments") enclosed in parentheses. The parameters together tell the function what to do. Some parameters are required, while others are optional. The syntax of a function is presented as follows:

```
FUNCTION_NAME(parameter_1, parameter_2, [parameter_3])
```

The order of the parameters matters, so always write them in the order required by the function syntax. Optional parameters are in square brackets.

For example, the `GPT` function has the following syntax:

```
GPT(prompt, [value], [temperature], [model])
```

The `GPT` function takes the following four parameters:

<table>
<tr><th>Parameter</th><th>Required</th><th>Description</th></tr>
<tr><td><code>prompt</code></td><td>Yes</td><td>The prompt for the AI. The prompt can be:
Text: <code>"Write a tagline for a tea house"</code>
Cell reference: <code>A1</code>
Range reference: <code>A1:A3</code></td></tr>
<tr><td><code>value</code></td><td>No</td><td>Input that you want to combine with the prompt. The input can be:
Text: <code>"Tone: Calm and serene"</code>
Cell reference: <code>B1</code>
Range reference: <code>B1:B3</code>
If you provide <code>value</code>, GPT for Excel uses the following template to build the final prompt sent to the AI:<pre><code><prompt>
Input: <value>
Output:</code></pre>For example, if <code>prompt</code> is <code>Write a tagline for a tea house</code> and <code>value</code> is <code>Tone: Calm and serene</code>, the final prompt is:<pre><code>Write a tagline for a tea house
Input: Tone: Calm and serene
Output:</code></pre>INFO
The <code>Output</code> line serves as a structural anchor that helps constrain the model into returning a clean, direct, and well-formatted response.
If you do not provide <code>value</code>, GPT for Excel sends <code>prompt</code> as is.</td></tr>
</table>

Parameter	Required	Description
`temperature`	No	Number between `0` and `1` that specifies how creative the AI should be in generating the response: • `0` sets minimum creativity. Responses to the same prompt are predictable and very similar. • `1` sets maximum creativity. Responses to the same prompt can vary wildly. Highest risk of hallucination.
`model`	No	Name of the AI model that you want to use. You can find the model names in the model switcher in the add-in sidebar. Learn more.

The first parameter of the `GPT` function is required, so you always have to provide it. The remaining three parameters are optional, so you only need to provide them when you want additional control over how the function operates. You can provide one, two, or all three optional parameters, depending on your needs.

DOUBLE QUOTES

If a parameter takes a text value, always enclose the text in double quotation marks. For example:

✅ `GPT("Write a tagline for a tea house")`

× `GPT(Write a tagline for a tea house)`

The latter fails with an error because the function cannot interpret the words as valid parameters.

EMPTY PARAMETERS

If you skip an optional parameter other than the last one, always provide a lone comma ("empty parameter") as a placeholder to maintain the correct syntax. For example:

✅ `GPT("Write a tagline for a tea house", , 0.5, "claude-4.6-sonnet")`

× `GPT("Write a tagline for a tea house", 0.5, "claude-4.6-sonnet")`

The latter fails with an error because the function treats `0.5` as the second parameter (`value`, which expects text or a reference) and `"claude-4.6-sonnet"` as the third parameter (`temperature`, which expects a number).

Run a GPT function

To use the `GPT` function:

1. Select a cell where you want the response to appear.
2. Type the equal sign `=` followed by the function. For example, to use the `GPT` function with just the `prompt` parameter, type:

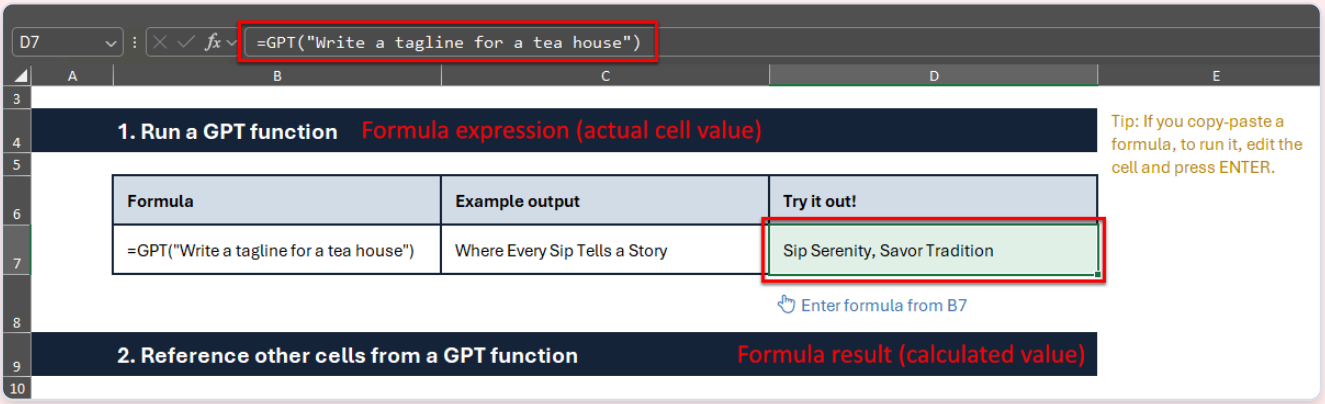
   ```
   =GPT("Write a tagline for a tea house")
   ```

3. Press ENTER.

The function runs: It sends the prompt to the AI, receives a response, and displays the response in the selected cell.

CALCULATED VALUES AND CONSTANTS

The text you see in the cell is the result of the formula — a dynamically calculated value. The actual value of the cell is the formula expression: `=GPT("Write a tagline for a tea house")`. You cannot edit the result; you can only edit the formula. If you want to set the result as the actual value of the cell, that is, if you want to turn the calculated value into a constant that you can edit, replace the formula with its result.



Reference other cells from a GPT function

You can reference other cells and ranges from a GPT function. You can use relative, absolute, and mixed references as you would with any function.

For example, to retrieve the prompt for the `GPT` function from another cell:

1. Select a cell and enter the prompt:

   ```
   Write a tagline for a tea house
   ```

> **TIP**
>
> You do not need double quotation marks around the prompt here because Excel automatically returns the value as text.

2. Select another cell and enter the `GPT` function with the prompt cell as the only parameter. For example, if the prompt is in cell `B12`, enter:

```
=GPT(B12)
```

   The function runs and returns the AI's response in the selected cell.

Both `GPT(B12)` and `GPT("Write a tagline for a tea house")` work the same way — they send the exact same prompt to the AI. The only difference is that the former retrieves the prompt from a cell, while the latter provides it directly as a text constant in the function.

Apply the same prompt to multiple cells

You can apply the same GPT function and prompt to adjacent cells by filling the formula across a range.

For example, suppose you want to generate unique taglines for multiple businesses. The businesses are listed in one column of your sheet, and you want to fill another column with the corresponding taglines. You can generate the taglines by using the `GPT` function with the `prompt` and `value` parameters, where `prompt` specifies the basic instructions and `value` specifies a specific business as a cell reference:

1. Select the first tagline cell and enter the `GPT` function. For example, if the business whose tagline you're generating is specified in cell `B7`, enter:

```
=GPT("Write a tagline", B7)
```

   The function combines the two parameters into a single prompt, based on the underlying prompt template, which is then sent to the AI. For example, if `B7` contains the text `antique store`, the final prompt sent to the AI is:

```
Write a tagline
Input: antique store
Output:
```

   The function runs and returns the AI's response in the selected cell.

2. Fill the formula down to apply it to the rest of the column: Select the fill handle and drag it down across the cells you want to fill. Excel automatically updates the cell reference in the function so that each row receives a tailored response based on its business. The formula in row 8 will reference `B8`, the formula in row 9 will reference `B9`, and so on.

To see how absolute references work, let's move the prompt to its own cell:

1. Select an empty cell and enter the prompt:

```
Write a tagline
```

2. Select the first tagline cell and update the `GPT` function to reference the prompt cell. For example, if the prompt is in cell `C8` and the business whose tagline you're generating is now specified in cell `B11`, enter:

```
=GPT($C$8, B11)
```

`$C$8` is an absolute reference to `C8`, which you need here since you do not want Excel to automatically adjust the reference when you fill the formula to adjacent cells.

> **TIP**
>
> Since you're only filling vertically across a single column, it would be enough to fix the prompt cell row using a mixed reference:
>
> ```
> =GPT(C$8, B11)
> ```

3. Fill the formula down to apply it to the rest of the column: Select the fill handle and drag it down across the cells you want to fill, overwriting the existing formulas. Excel automatically updates the value cell reference for each row, while keeping the prompt cell reference fixed.

Create a prompt by combining text from multiple cells

You can use concatenation to combine text from two or more cells into a single text string. This allows you to create advanced prompts, where different parts of the prompt come from different cells in the sheet.

For example, suppose you want to further customize the taglines of the previous example, so that they follow a specific tone and target a specific customer age range. The businesses are listed in the first column, the tone in the second, and the age range in the third, with the fourth column now reserved for the taglines.

To build the prompt by combining the information from the different columns:

1. Select the first tagline cell and enter the `GPT` function. For example, if the business whose tagline you're generating is specified in cell `B11`, with the tone in cell `C11` and age range in cell `D11`, enter the following in cell `E11`:

```
=GPT("Write a tagline for " & B11 & ", the tone is " & C11 & ", the
```

```
audience is aged " & D11)
```

The function uses the `&` operator to concatenate the different parts, which include literal text values and cell references, into a single text string. For example, if `B11` contains the text `antique store`, `C11` contains `nostalgic and refined`, and `D11` contains `35-65`, the final prompt sent to the AI is:

```
Write a tagline for antique store, the tone is nostalgic and refined, the
audience is aged 35-65
```

TIP

Instead of the `&` operator, you can also use the native `CONCAT` function:

```
=GPT(CONCAT("Write a tagline for ", B11, ", the tone is ", C11, ", the
audience is aged ", D11))
```

The `GPT` function uses the `CONCAT` function as a nested function. Learn more.

If you want to concatenate without having to worry about white spaces and newlines, use the `GPT_CREATE_PROMPT` function:

```
=GPT(GPT_CREATE_PROMPT("Write a tagline for", B11, ", the tone is",
C11, ", the audience is aged", D11))
```

2. Fill the formula down to apply it to the rest of the column: Select the fill handle and drag it down across the cells you want to fill. Excel automatically updates the cell references in the function so that each row receives a tailored response based on its business. The formula in row 12 will reference `B12`, `C12`, and `D12`; the formula in row 13 will reference `B13`, `C13`, and `D13`; and so on.

Combine a GPT function with a native function

You can combine GPT functions with native functions like you would any other function.

For example, suppose you want to run the `GPT` function from above only when the current row has a business specified. If the business cell is empty, you want to skip the row. You can do this by combining the `GPT` function with the native `IF` function:

1. Select the first tagline cell and enter a formula that wraps `GPT` inside an `IF` function. For example, if the business whose tagline you're generating is specified in cell `B11` and the GPT function for that is `GPT($C$8, B11)`, enter:

```

```

```
=IF(B11="", "", GPT($C$8, B11))
```

The `GPT` function runs only if cell `B11` is not empty. If `B11` is empty, the formula leaves the current cell empty, too.

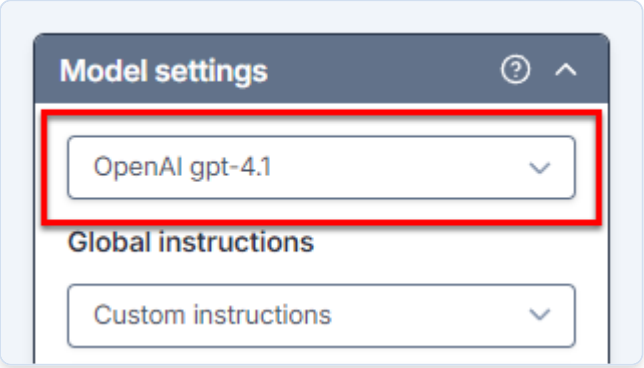
2. Fill the formula down to apply it to the rest of the column: Select the fill handle and drag it down across the cells you want to fill, overwriting the existing formulas. Excel automatically updates the value cell reference for each row, while keeping the prompt cell reference fixed.
3. To see the conditional logic in action, empty one of the business cells. The formula for that row reruns and removes the tagline.

Search the web with a GPT function

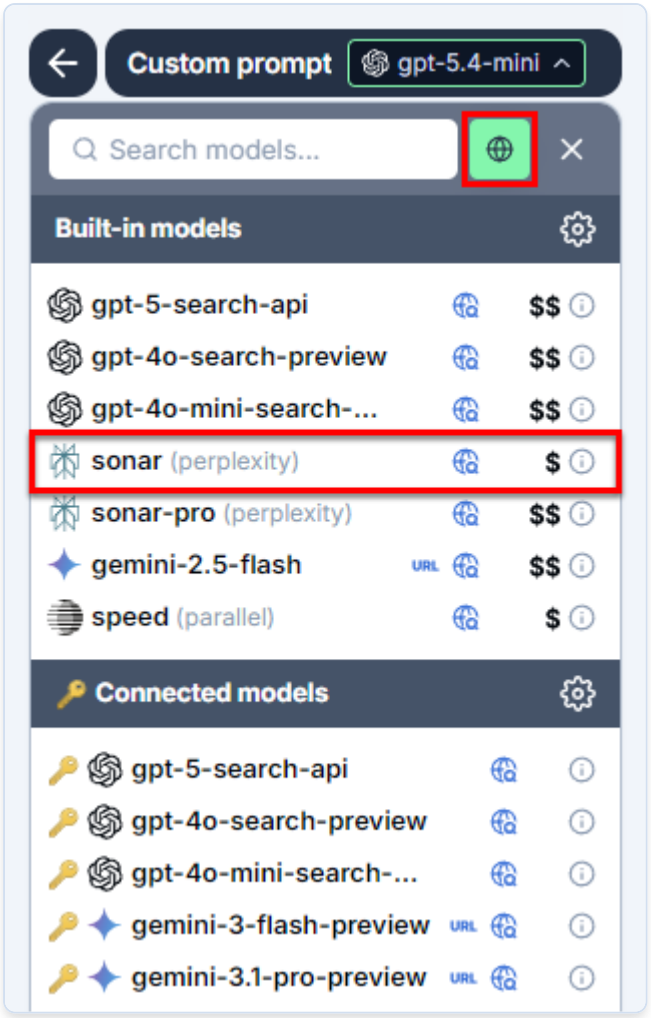
You can search the web with a GPT function simply by using a web search model. You can either select the model in the model switcher or specify the model with the `model` parameter.

For example, suppose you want to do a search for company CEOs. The companies are listed in one column of your sheet, and you want to fill the next column with the CEO names.

1. In the sidebar, select **GPT functions**.
2. Expand **Model settings** and open the model switcher.



3 Enable the web search filter and select a web search model (🌐).



> **INFO**
>
> If you select a Gemini model, make sure web search is enabled for Gemini models.

4 Select the first CEO cell and enter the GPT function. For example, if the company whose CEO you're fetching is specified in cell `B13`, enter:

```
=GPT("Give me this company's current CEO. I only want the name of the CEO, no citations.", B13)
```

> **TIP**
>
> If you want to override the model switcher selection, specify the model in the GPT function with the `model` parameter. For example, if you want to use the perplexity-fast model, enter:
>
> ```
> =GPT("Give me this company's current CEO. I only want the name of the CEO, no citations.", B13, , "perplexity-fast")
> ```

5. Fill the formula down to apply it to the rest of the column: Select the fill handle and drag it down across the cells you want to fill. Excel automatically updates the company cell reference for each row.

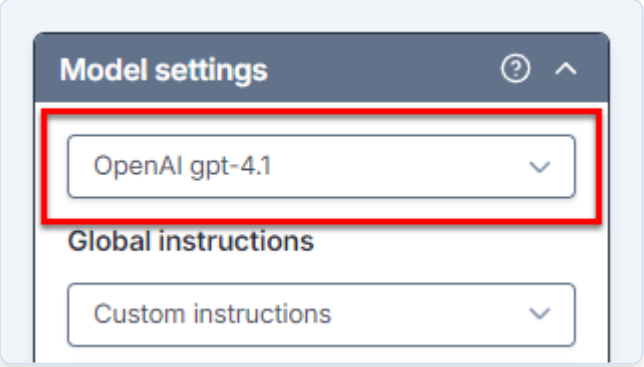
For more information about web search, see Search the web with GPT for Excel.

Process images with GPT_VISION

You can apply a text prompt to an image using the `GPT_VISION` function with a vision model. You can either select the model in the model switcher or specify the model with the `model` parameter.

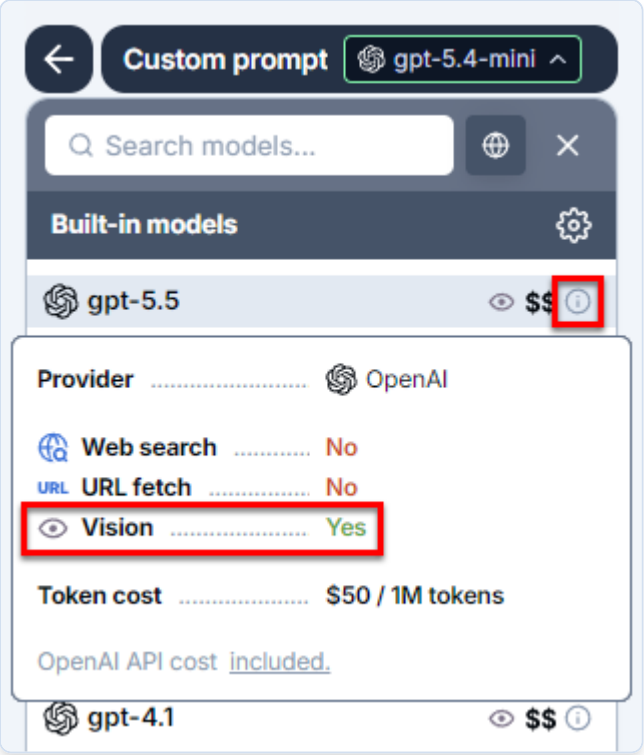
For example, suppose you want to extract the brand name from product images featuring coffee bean packs. The image URLs are listed in one column of your sheet, and you want to fill the next column with the brand names.

1. In the sidebar, select **GPT functions**.
2. Expand **Model settings** and open the model switcher.



3. Select a vision model.

   To determine whether a model supports vision, check its info popup. For a list of available vision models, see Use images in prompts.



4 Select the first image URL cell and enter the GPT function. For example, if the image URL that you want to process is in cell `C13`, enter:

```
=GPT_VISION("Extract the coffee brand from this image. Normalize the brand name to title case.", C13)
```

> **TIP**
>
> If you want to override the model switcher selection, specify the model in the GPT function with the `model` parameter. For example, if you want to use the Claude Sonnet 4.6 model, enter:
>
> ```
> =GPT_VISION("Extract the coffee brand from this image. Normalize the brand name to title case.", C13, , , "claude-4.6-sonnet")
> ```

5 Fill the formula down to apply it to the rest of the column: Select the fill handle and drag it down across the cells you want to fill. Excel automatically updates the image URL cell reference for each row.

For more information about AI vision, see Use images in prompts in Excel.

Set the model for a GPT function

You can run a GPT function with a specific model by providing the `model` parameter. The function will ignore what's currently selected in the model switcher and use the provided model instead.

For example, if you want the perplexity-fast model to answer a query about current events, provide `perplexity-fast` as the model:

```
=GPT("What was the previous month's inflation rate in the US?", , ,
"perplexity-fast")
```

> **INFO**
>
> If the model is available both as a built-in and connected model in GPT for Excel, the function uses the connected model.

Use the `model` parameter when:

- **You want to lock a formula to a specific model.** For example, if you have a web search formula, you want it to consistently use the same web search model.
- **You want to compare how a specific model compares with other models.** For example, use the same GPT function and prompt in two formulas — the first with the `model` parameter, the second without — and then use the model switcher to change the model used by the second formula, regenerating its response after each change.

  ```
  =GPT("Write a tagline for a tea house", , , "claude-4.6-sonnet")

  =GPT("Write a tagline for a tea house")
  ```

Set the creativity level for a GPT function

You can run a GPT function with a specific creativity level by providing the `temperature` parameter. The function will ignore the creativity level defined in the model settings and use the provided level instead.

For example, if you're working on a creative task and want to experiment with highly creative responses, set the temperature close to `1`:

```
=GPT("Write a three-sentence horror story about a teahouse.", , 0.9);
```

For more information about temperature, see our temperature guide for OpenAI models.

GPT functions: Defaults, settings, behavior

Default model

By default, GPT functions use the currently selected model. If you want a GPT function to use a specific model regardless of what's currently selected in the model switcher, use the `model` parameter to explicitly set the model.

> **INFO**
>
> If you provide the `model` parameter, and if the model is available both as a built-in and connected model in GPT for Excel, the function uses the connected model.

Default creativity level

By default, GPT functions use the creativity level defined in the model settings. If you want a GPT function to use a specific creativity level regardless of what's defined the model settings, use the `temperature` parameter to explicitly set the creativity level.

Global instructions and advanced settings

The global instructions and advanced settings defined in the GPT for Excel sidebar apply to GPT functions. If you find that a function is not working as expected, check the model settings for a possible cause.

Regional settings

The default characters used by Excel for parameter and decimal separators depend on your currently selected region:

- **Parameters:** Some regions use the comma (`,`) as the default parameter separator, while others use the semicolon (`;`).

```
=GPT("Write a tagline", B7)

=GPT("Write a tagline"; B7)
```

- **Decimals:** Some regions use the period (`.`) as the default decimal separator, while others use the comma (`,`). As a rule, regions that use the period for decimals use the comma for parameters, whereas regions that use the comma for decimals use the semicolon for parameters.

```
=GPT("Write a tagline for a tea house", , 0.5)
```

```
=GPT("Write a tagline for a tea house"; ; 0,5)
```

The examples in the current documentation use the comma as the parameter separator and the period as the decimal separator. If the examples fail with syntax errors in your workbook, check your regional settings to verify which punctuation to use.

Formula auto-refresh

When your sheet auto-refreshes, Excel automatically recalculates some or all formulas on the sheet, including GPT formulas. For each GPT formula:

- If the GPT formula cache is enabled for the workbook, Excel reuses the formula's previously generated result from the cache.
- If the cached result has expired, or if you have the cache disabled, Excel runs the formula normally.

In Excel, the following triggers cause a sheet to auto-refresh:

Recalculation trigger	What gets recalculated
Insert, move, or delete a column	All formulas on the sheet
Insert, move, or delete a row	All formulas on the sheet
Filter rows by column	All formulas on the sheet
Hide and unhide rows	Formulas in the selected rows
Sort rows	Formulas in the selected rows
Move a formula to a different cell	All formulas on the sheet
Undo the deletion of a formula from a cell	Restored formula

To prevent GPT formulas from being recalculated, replace the formulas with their results.

When to use bulk tools

Consider using bulk tools instead of GPT functions if:

- **You want AI responses saved as plain text in cells.** Bulk tools fill cells with static text values, or constants, so you don't need to replace any formulas. You also avoid formula auto-refresh, reducing the risk of unintended data changes.
- **You want better tracking information about the progress of AI requests.** Bulk tools provide a live progress tracker that shows you how many rows have been successfully processed, how many have resulted in an error, and how much time has elapsed.
- Your formulas regularly become stuck in the #BUSY! state.

What's next

- Discover all GPT functions available in GPT for Excel.
- Learn how to use each GPT function with a real-world example.
- Select the model that best fits your needs.
- Configure the model settings to customize how the currently selected model behaves.

GPT functions reference for Excel

Discover all custom functions in GPT for Excel, along with simple examples and an extensive list of parameters. If you're looking for more examples, download our GPT for Excel examples template, and open it in Excel as a new workbook.

Watch video on YouTube

> **TIP**
> To generate responses based on fresh web search data, use a model with web search capabilities (indicated by the 🌐 icon in the model switcher).

GPT function

The simplest function included in GPT for Excel. Outputs the result in a single cell.

> **INFO**
> To process images, use GPT_VISION.

How to use

Syntax	=GPT(prompt, [value], [temperature], [model])
Copy/paste example	`=GPT("Write a tagline for an ice cream shop.")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	A response to your prompt in a single cell

Parameters

Parameter	Definition
prompt (cannot be empty)	Instruction provided to the AI to generate a result in a single cell. The prompt parameter can be: • A string: `"Write a tagline for an ice cream shop."` • A cell: `A1` • A range of cells: `A1:C3`
(optional) value	Text, cell or range you want your prompt to apply to
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_CLASSIFY function

Classifies a piece of text into a single category selected from a predefined list of categories.

How to use

Syntax	**=GPT_CLASSIFY(value, categories, [instructions], [temperature], [model])**
Copy/paste example	`=GPT_CLASSIFY("banana", "fruit, vegetable")`
⚠️ Localization	In some locales, you might need to use ";" to separate parameters instead of ","
Output	The best applicable category

Parameters

Parameter	Definition
value (cannot be empty)	Input text to classify. The value parameter can be: • A string: `"banana"` • A cell: `A1` • A range of cells: `A1:C3`
categories (cannot be empty)	Comma-separated categories or range of categories to choose from. Only the most relevant category will be returned.
(optional) instructions	Additional instructions for the classification. Can include examples.
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_EDIT function

Applies the given task to the given text. The default task is to fix grammar and spelling.

How to use

Syntax	=GPT_EDIT(text, [task], [temperature], [model])
Copy/paste example	`=GPT_EDIT("For sum reezon thoose nunsberz arnt addin up")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	The edited text in a single cell

Parameters

Parameter	Definition
text (cannot be empty)	Input text to edit. The text parameter can be: • A string: `"For sum reezon thoose nunsberz arnt addin up"` • A cell: `A1` • A range of cells: `A1:C3`
(optional) task	Text, cell or range specifying how the text should be edited. Defaults to fixing grammar and spelling. Other examples: • "Make it funnier" • "Make it shorter" • "Make it sound formal" • "Add an ending"
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_EXTRACT function

Extracts specific elements from text, such as names, dates, email addresses, URLs, or product attributes. Outputs as comma-separated values.

How to use

Syntax	GPT_EXTRACT(text, to_extract, [temperature], [model])
Copy/paste example	`=GPT_EXTRACT("I worked 5 years at Amazon.com and then 3 years at Apple", "companies")`
⚠️ Localization	In some locales, you might need to use ";" to separate parameters instead of ","

Syntax	GPT_EXTRACT(text, to_extract, [temperature], [model])
Output	The extracted entities as comma-separated values in a single cell

Parameters

Parameter	Definition
text (cannot be empty)	Input text to extract data from. The text parameter can be: • A string: `"I worked 5 years at Amazon.com and then 3 years at Apple"` • A cell: `A1` • A range of cells: `A1:C3`
to_extract (cannot be empty)	Text, cell or horizontal range specifying what you want to extract. Examples: • "companies" • B3:E3
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_FORMAT function

Formats dates, currencies, addresses, names, etc. Fixes capitalization. And so much more.

How to use

Syntax	**=GPT_FORMAT(input, target_format, [source_format], [temperature], [model])**
Copy/paste example	`=GPT_FORMAT("14 February 71", "iso")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	The converted input in a single cell unless it is converted a table, which outputs a table

Parameters

Parameter	Definition
input (cannot be empty)	Input text to format. The input parameter can be: • A string: `"14 February 1971"` • A cell: `A1` • A range of cells: `A1:C3`
target_format (cannot be empty)	Target format for the input. Examples: "iso", "currency iso", "international phone number", "lowercase/uppercase/camelcase/snake case/title case/sentence case"...
(optional) source_format	Source format of the input. Optional, but including it can yield better results.
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_LIST / GPT_HLIST function

Like GPT, but outputs the results in a column. Very practical when the output is a list.

- Use GPT_LIST to return options listed vertically.
- Use GPT_HLIST to return options listed horizontally.

How to use

Syntax	=GPT_LIST(prompt, [value], [temperature], [model])
Copy/paste example	`=GPT_LIST("Give me 5 good short ads about spreadsheets.")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	As many responses to your prompt as you asked, one response per cell

Parameters

Parameter	Definition
prompt (cannot be empty)	Instruction provided to the AI to generate a list. The prompt parameter can be: • A string: `"Give me 5 good short ads about spreadsheets."` • A cell: `A1` • A range of cells: `A1:C3`
(optional) value	Text, cell or range you want your prompt to apply to
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_MATCH function

Matches the values of two columns by similarity.

> **INFO**
>
> Regardless of the model you select in the model switcher, this function uses text-embedding-ada-002, for a fixed cost of $0.25 per 1M tokens.

> **WARNING**
>
> - GPT for Excel settings do not apply to the GPT_MATCH function.
> - You can only match up to about 1,000 rows at a time.

How to use

Syntax	**=GPT_MATCH(search_keys, lookup_range, [confidence], [stats], [top_k])**
Copy/paste example	`=GPT_MATCH(A1:A10,B1:B4)`
⚠️ Localization	In some locales, you might need to use ";" to separate parameters instead of ","
Output	A column of the best matches to each value in the range

Parameters

Parameter	Definition
search_keys (cannot be empty)	The values that you want to look for in the range
lookup_range (cannot be empty)	The range in which to look for values that are similar to search keys
confidence (optional)	Set a similarity threshold: only values with a superior similarity score are matched. Must be comprised between 0 and 1

Parameter	Definition
stats (optional)	"true" to display similarity scores
top_k (optional)	Number of matching values to return per search_key (default: 1 - only the best match)

GPT_SPLIT / GPT_HSPLIT function

Splits text semantically, such as by section, paragraph, sentence, customer...

- Use GPT_SPLIT to return outputs listed vertically.
- Use GPT_HSPLIT to return outputs listed horizontally.

How to use

Syntax	=GPT_SPLIT(text, split_by, [temperature], [model])
Copy/paste example	`=GPT_SPLIT("Hello! How are you? Have you heard of GPT for Work?", "sentences")`
⚠️ Localization	In some locales, you might need to use ";" to separate parameters instead of ","
Output	Input text is split as requested, one element per cell

Parameters

Parameter	Definition
text (cannot be empty)	Input text to split. The text parameter can be: • A string: `"Hello! How are you? Have you heard of GPT for Work?"` • A cell: `A1` • A range of cells: `A1:C3`

Parameter	Definition
split_by (cannot be empty)	How to split the text. Examples: • sections • paragraphs • sentences • punctuation
(optional) temperature, model	See temperature and model.

GPT_TABLE function

Like GPT, but outputs the results in a table. Very practical when the output is a table.

How to use

Syntax	=GPT_TABLE(prompt, [head], [inputs], [temperature], [model])
Copy/paste example	`=GPT_TABLE("top 10 most eaten fruits and their nutrition data")`
⚠️ Localization	In some locales, you might need to use ";" to separate parameters instead of ","
Output	The response to your prompt in the shape of a table

Parameters

Parameter	Definition
prompt (cannot be empty)	Instruction to generate a table. The prompt parameter can be: • A string: `"top 10 most eaten fruits and their nutrition data"` • A cell: `A1` • A range of cells: `A1:C3`
(optional) head	Range containing at least one row. the first row should be headers, the following rows can be used to specify examples If left empty, headers will be generated automatically.
(optional) inputs	Range containing partial rows to complete with GPT_TABLE
(optional) temperature, model	See temperature and model.

GPT_TAG function

Applies user-defined tags to text or automatically suggests relevant tags if none are specified.

How to use

Syntax	**=GPT_TAG(value, [tags], [instructions], [top_k], [temperature], [model])**
Copy/paste example	`=GPT_TAG("I love chocolate", "food, positive, negative")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	The applicable tags as csv in a single cell

Parameters

Parameter	Definition
value (cannot be empty)	Input text to tag. The value parameter can be: • A string: `"I love chocolate"` • A cell: `A1` • A range of cells: `A1:C3`
(optional) tags	Comma-separated tags or range of tags applicable. If left empty, GPT_TAG will automatically come up with the tags
(optional) instructions	Additional instructions for the tag application. Can include examples.
(optional) top_k	Maximum number of tags to return
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_TRANSLATE function

Translates text from one language to another. Works better if you specify the source language.

How to use

Syntax	**=GPT_TRANSLATE(text, [target_language], [source_language], [instructions],[temperature], [model])**
Copy/paste example	`=GPT_TRANSLATE("Cool off with our delicious treats!", "spanish")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","

Syntax	**=GPT_TRANSLATE(text, [target_language], [source_language], [instructions],[temperature], [model])**
Output	The requested translation in a single cell

Parameters

Parameter	Definition
text (cannot be empty)	Input text to translate. The text parameter can be: • A string: `"Cool off with our delicious treats!"` • A cell: `A1` • A range of cells: `A1:C3`
(optional) target_language	Language of the output. Examples: "spanish", "chinese", "french". Defaults to "english".
(optional) source_language	Language of the input. Example: "english"
(optional) instructions	Additional instructions for the translation. Can include a dictionary of translations.
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_VISION function

Applies a prompt to an image. For example you can ask a question about the image, ask for a description of it, or generate alt text for images. Learn more.

INFO

- GPT_VISION is the only function that accepts images as input.
- If you want to use GPT_VISION, you must use a vision model, like `gpt-5.4`.

How to use

Syntax	=GPT_VISION(prompt, image_url, [high_res], [temperature], [model])
Copy/paste example	`=GPT_VISION("Write a short description for this product", A1, , , "gpt-4.1")`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	A response to your prompt according to the information extracted from the image

Parameters

Parameter	Definition
prompt (cannot be empty)	Instruction provided to the AI to analyze an image. The prompt parameter can be: • A string: `"Write a short description for this product"` • A cell: `A1` • A range of cells: `A1:C3`
image_url	Link to the image you want the AI to analyze. Supported formats are PNG, JPEG, non-animated GIF. The image_url parameter can be: • A string: `"https://static.nike.com/a/images/t_PDP_1280_v1/f_auto,q_auto:e6393-4af6-abbc-4f1acaa6ed94/air-max-dawn-womens-shoes-224mWf.pn` • A cell: `A4`
(optional) high_res	How the AI processes the image and generates its textual understanding. Possib `true`, `false`, and `auto` (default). Learn more.

Parameter	Definition
(optional) temperature, model	See temperature and model.

See more usage examples.

GPT_WEB function

> **INFO**
> GPT_WEB uses the web search model set for your space. Learn more.

How to use

Syntax	=GPT_WEB(prompt, [value])
Copy/paste example	`=GPT_WEB("Who won the last presidential elections in Argentina?")`
⚠️ Localization	In some locales, you might need to use ";" to separate parameters instead of ","
Output	A response to your prompt based on recent web search data

Parameters

Parameter	Definition
prompt (cannot be empty)	Instructions provided to the AI to generate results based on recent web search data. Provide instructions in English for better results. The prompt parameter can be: • A string: `"Who won the last presidential elections in Argentina?"` • A cell: `A1`

Parameter	Definition
(optional) value	Text, cell, or range you want your prompt to apply to

See more usage examples.

GPT_CREATE_PROMPT helper function

Like the native **CONCAT** function, but inserts spaces and newlines automatically. Helps you create a prompt from scattered, disjointed cells and ranges. You can then use that prompt in GPT functions. This function does not call any AI provider's API and thus does not incur any cost.

How to use

Syntax	=GPT_CREATE_PROMPT(arg1, [arg2], [arg3], [...argn])
Copy/paste example	`=GPT_CREATE_PROMPT(A1:B1, A2, C3:D5, E1:E6)`
⚠️ **Localization**	In some locales, you might need to use ";" to separate parameters instead of ","
Output	A prompt ready to use in GPT, GPT_LIST or GPT_TABLE

Parameters

Parameter	Definition
arg1 (cannot be empty)	Cell or range to concatenate. For example: • A1 • A1:C3
(optional) arg2 - argn	Cell or range to concatenate. For example: • A2 • G2:H4

temperature and model

These parameters are present in most GPT functions. They are always optional.

For example with the GPT function, you can set:

- **temperature** to `0.1`
- **model** to `"gpt-4o-mini"`

```
= GPT("what's the address of the following companies", "CompanyA,CompanyB", 0.1, "gpt-4o-mini"
                                                                              temperature  model
```

You can also set these values from the sidebar under **Model settings**. In this case, they become the default values for all future executions.

Parameter	Definition
(optional) temperature	Number between 0 and 1 that governs the creativity of the output: • **0** - Strictly according to the prompt • **0.5** - Slightly creative • **1** - Very creative **INFO** In some locales, you might need to use "0,5" (comma) instead of "0.5" (period).
(optional) model	Name of the model surrounded by quotes. You can find model names in the model switcher in the sidebar.

Select a model for GPT functions in Excel

GPT for Excel supports a range of models that you can use with the GPT functions. To help you find the model that offers the best balance of accuracy, speed, and cost for your use case, see our recommended models.

The available models depend on:

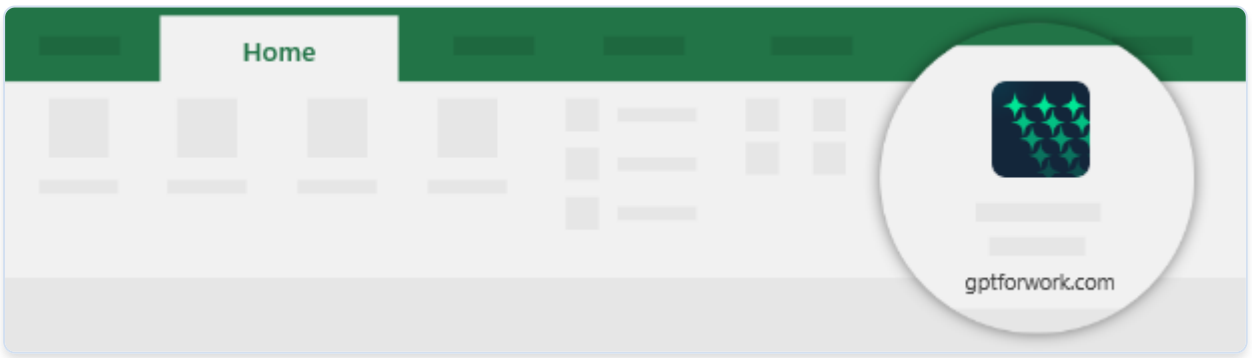
- **The API keys set for your current space.** GPT for Excel gives you access to a default set of built-in models. To access additional models, set an API key for one or more AI providers.
- **The API endpoints set up for your current space.** To access even more models, set up a connection to an Azure OpenAI API endpoint or any OpenAI-compatible API endpoint.
- **The model settings for your current space.** Space owners and admins can control which models are available in GPT for Excel.

Recommended models

For information about recommended bulk models, see Recommended models.

Select a built-in model

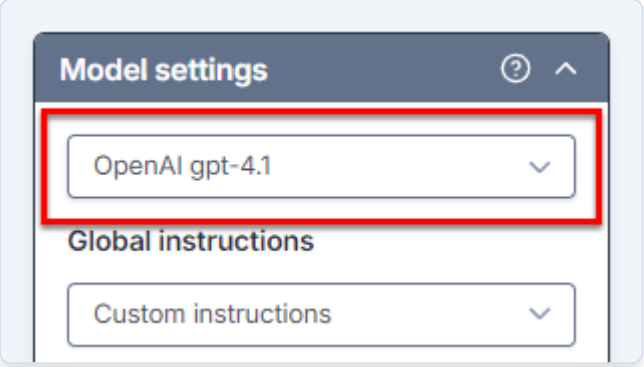
1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.



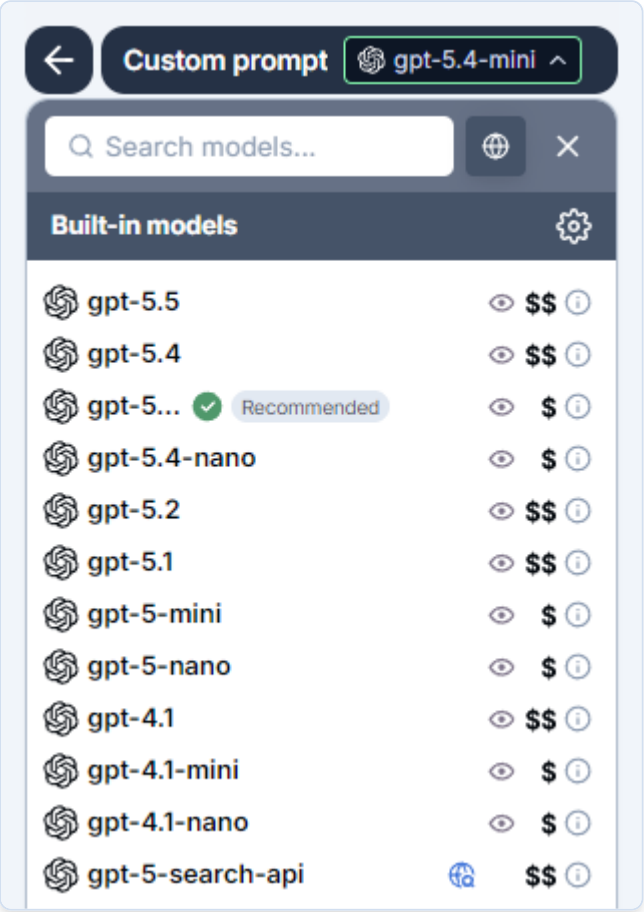
> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **GPT functions**.

4 Expand **Model settings** and open the model switcher.



5 Under **Built-in models**, select a model.



GPT for Excel now uses the selected model to generate responses for GPT functions.

INFO

- When cache is enabled, existing GPT formulas will not automatically update to the new model when you re-execute them.

  To re-execute an existing formula with the new model:

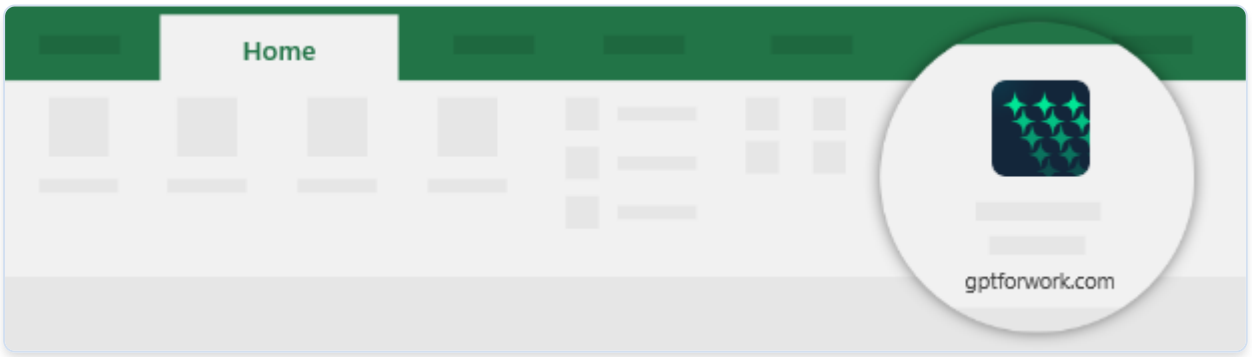
  1. Set the model as a parameter in the formula.
  2. Press Enter to re-execute the formula.
- The GPT_MATCH function uses `text-embedding-ada-002`, not the model you select.
- The GPT_WEB function uses the web search model set for your space.

Select a connected model

Prerequisites

- You have set an API key for the AI provider whose model you want to use.
- If you want to use an Azure model or an OpenAI-compatible model from an API endpoint, you have set up the correct endpoint.

1. Open a Microsoft Excel workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.

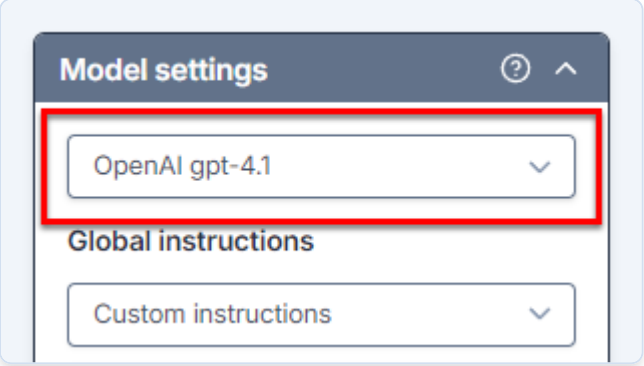


INFO

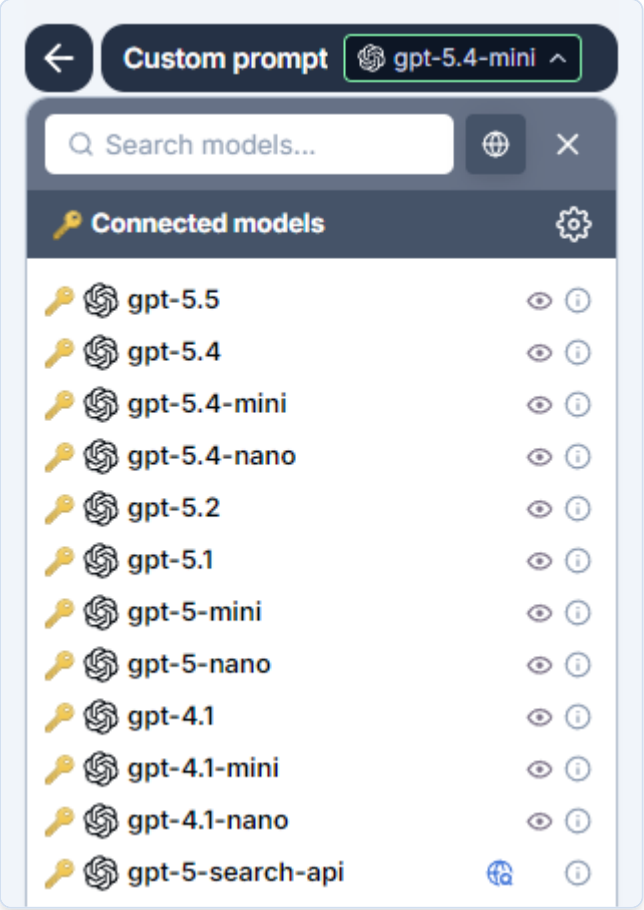
If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, select **GPT functions**.

4 Expand **Model settings** and open the model switcher.



5 Under **Connected models**, select a model.



GPT for Excel now uses the selected model to generate responses for GPT functions.

INFO

- When cache is enabled, existing GPT formulas will not automatically update to the new model when you re-execute them.

  To re-execute an existing formula with the new model:

  1. Set the model as a parameter in the formula.
  2. Press Enter to re-execute the formula.
- The GPT_MATCH function uses `text-embedding-ada-002`, not the model you select.
- The GPT_WEB function uses the web search model set for your space.

What's next

Learn more about models:

- AI providers and models supported by GPT for Work
- LiveBench (a good model benchmark)

Create GPT formulas in Excel

Delay and chain GPT functions, and learn how to use different GPT functions by example.

Control GPT formula execution

Delay or sequence the execution of GPT formulas in Microsoft Excel.

GPT formula examples

Real-world examples using some of the functions included in GPT for Excel.

Control GPT formula execution in Excel

You might want to delay the execution of your GPT for Excel functions until your workbook is ready. Or you might want to sequence the execution of multiple chained calls. Avoid running into errors or going over your rate limits by learning how to keep control over the executions of your GPT functions in Microsoft Excel in this video.

Watch video on YouTube

GPT formula examples for Excel

See GPT functions in action with real-world examples using some of the functions included in GPT for Excel.

> **TIP**
>
> To try out the examples, download our GPT for Excel examples template, and open it in Excel.

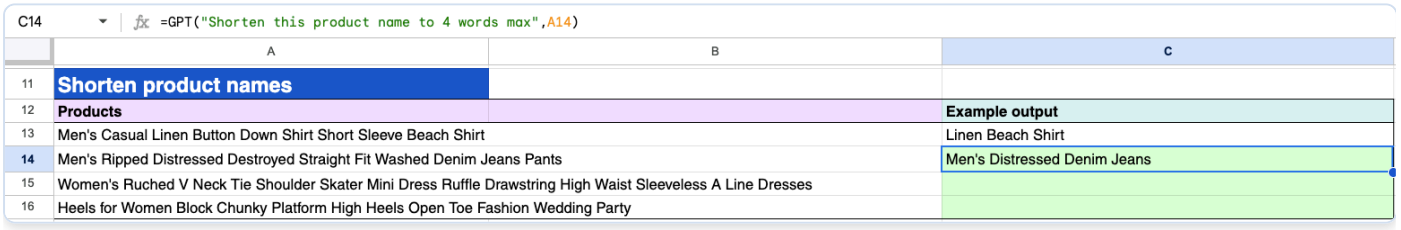
Rewrite and generate

Use GPT functions to rewrite or create content adapted to your goals:

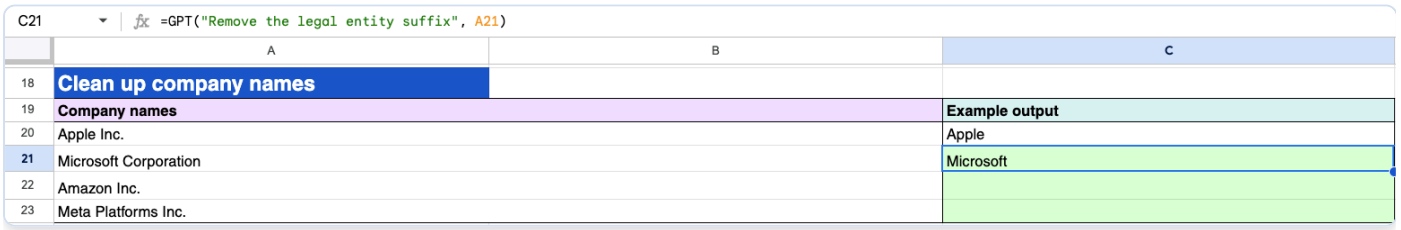
- Rephrase existing text by adjusting length, style, grammar or tone.
- Create content that matches your audience needs using the perfect level of creativity.

Rewrite or rephrase content

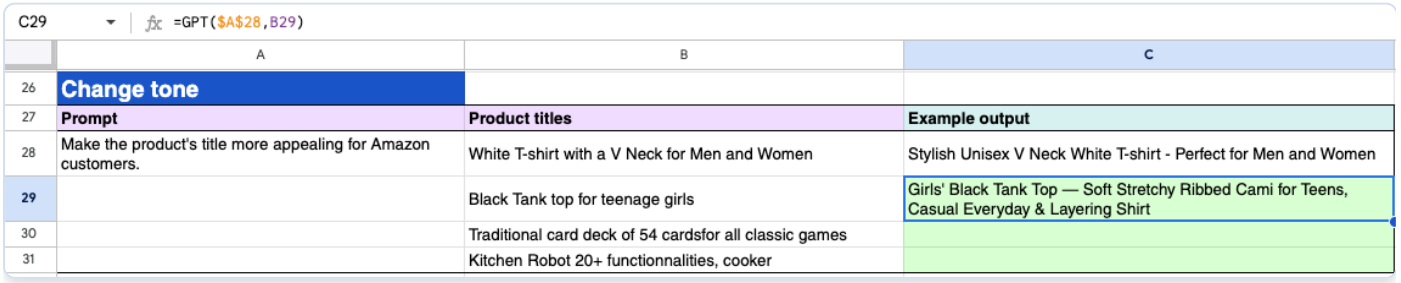
Shorten product names:



Clean up text:



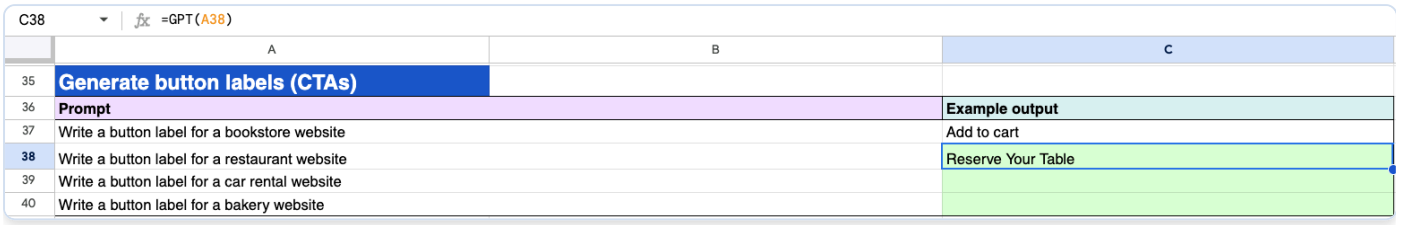
Change tone:



Learn more on GPT.

Create content

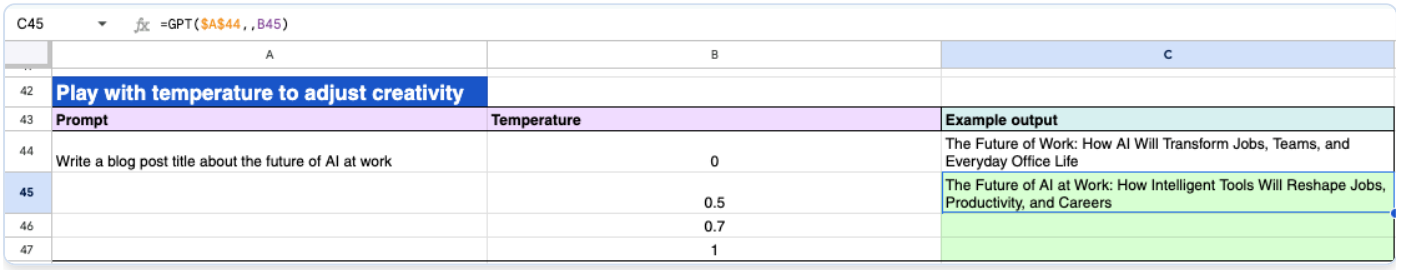
Generate button labels (CTAs):



C38 =GPT(A38)

	A	B	C
35	**Generate button labels (CTAs)**		
36	**Prompt**		**Example output**
37	Write a button label for a bookstore website		Add to cart
38	Write a button label for a restaurant website		Reserve Your Table
39	Write a button label for a car rental website		
40	Write a button label for a bakery website		

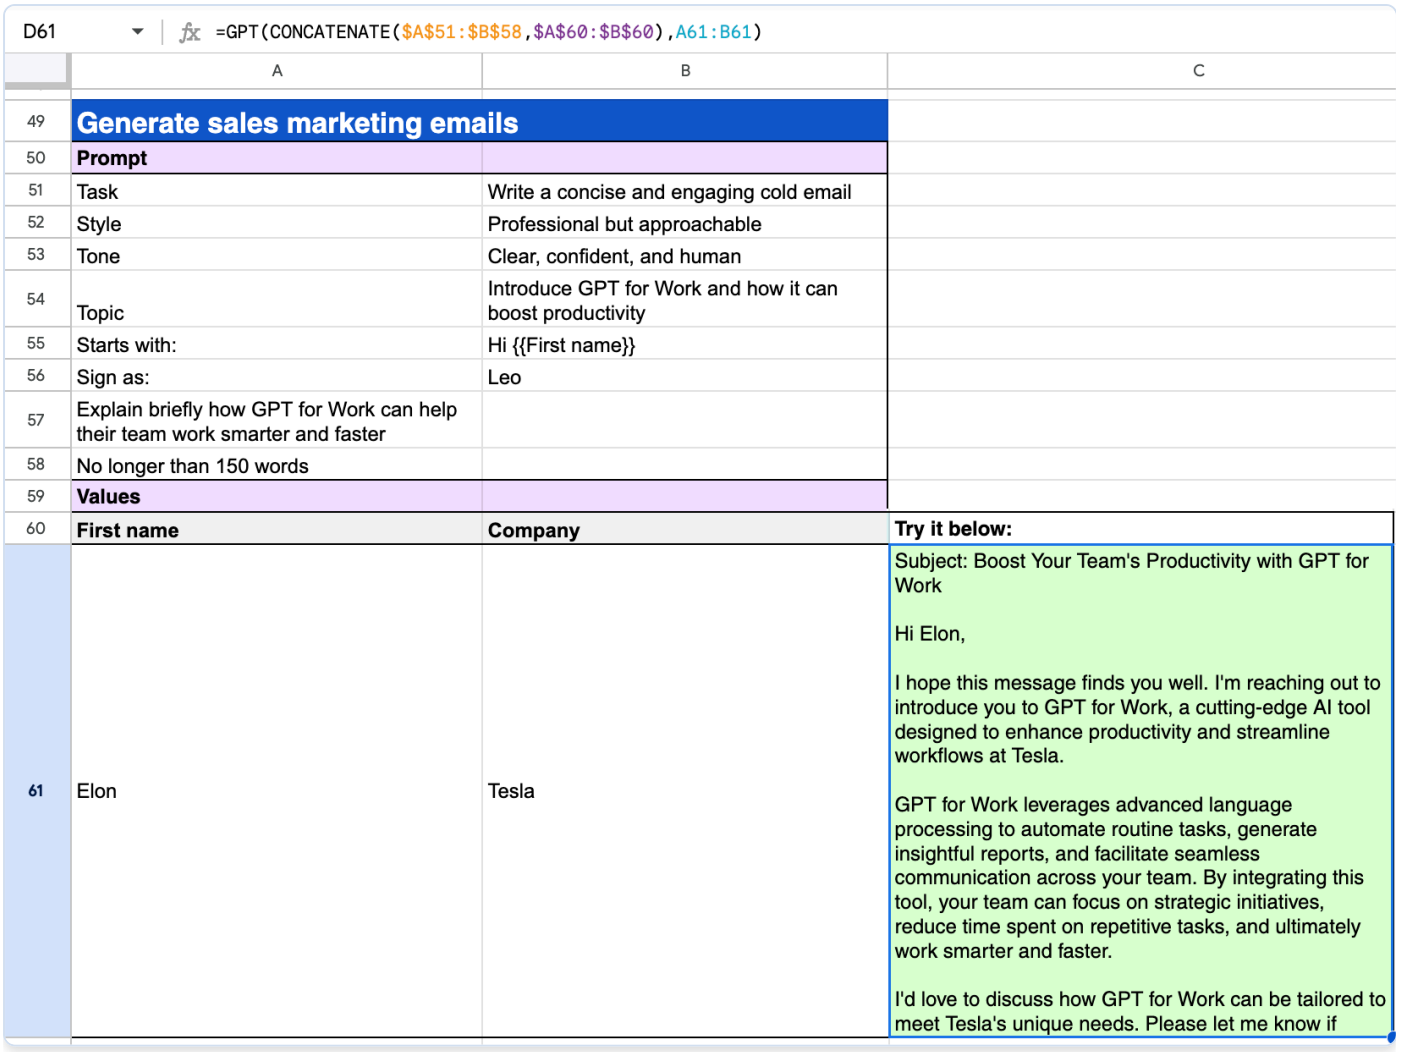
Make your content more or less creative:



C45 =GPT($A$44,,B45)

	A	B	C
42	**Play with temperature to adjust creativity**		
43	**Prompt**	**Temperature**	**Example output**
44	Write a blog post title about the future of AI at work	0	The Future of Work: How AI Will Transform Jobs, Teams, and Everyday Office Life
45		0.5	The Future of AI at Work: How Intelligent Tools Will Reshape Jobs, Productivity, and Careers
46		0.7	
47		1	

Generate sales marketing emails:



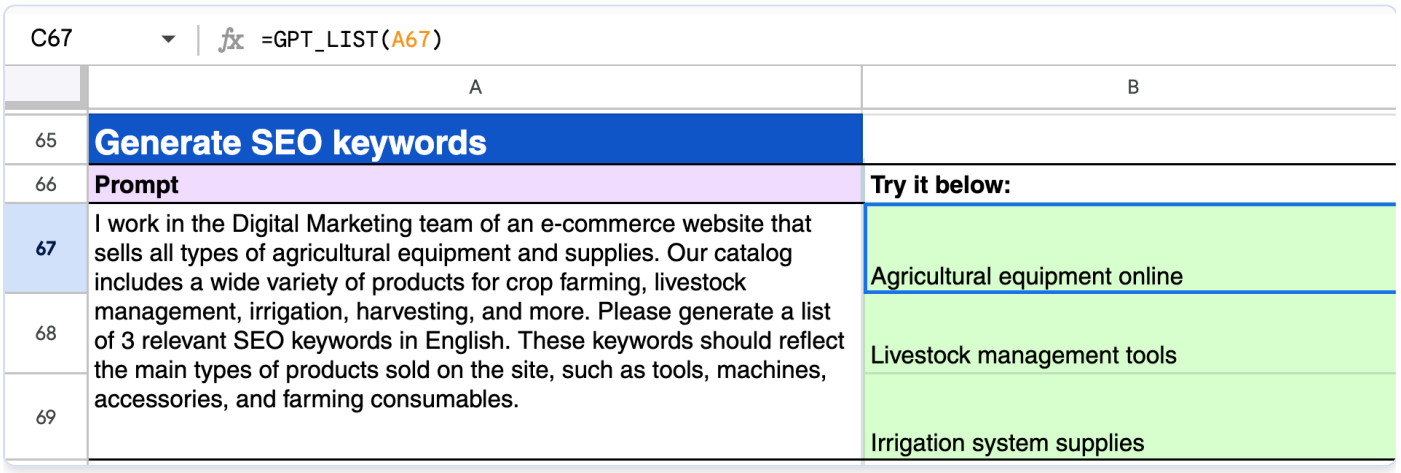
D61 =GPT(CONCATENATE($A$51:$B$58,$A$60:$B$60),A61:B61)

	A	B	C
49	**Generate sales marketing emails**		
50	**Prompt**		
51	Task	Write a concise and engaging cold email	
52	Style	Professional but approachable	
53	Tone	Clear, confident, and human	
54	Topic	Introduce GPT for Work and how it can boost productivity	
55	Starts with:	Hi {{First name}}	
56	Sign as:	Leo	
57	Explain briefly how GPT for Work can help their team work smarter and faster		
58	No longer than 150 words		
59	**Values**		
60	**First name**	**Company**	**Try it below:**
61	Elon	Tesla	Subject: Boost Your Team's Productivity with GPT for Work Hi Elon, I hope this message finds you well. I'm reaching out to introduce you to GPT for Work, a cutting-edge AI tool designed to enhance productivity and streamline workflows at Tesla. GPT for Work leverages advanced language processing to automate routine tasks, generate insightful reports, and facilitate seamless communication across your team. By integrating this tool, your team can focus on strategic initiatives, reduce time spent on repetitive tasks, and ultimately work smarter and faster. I'd love to discuss how GPT for Work can be tailored to meet Tesla's unique needs. Please let me know if

Learn more on GPT.

Generate multiple outputs (vertical vs. horizontal)

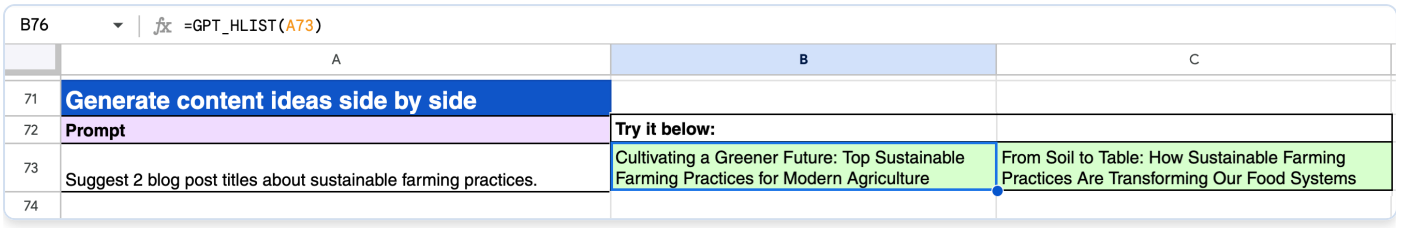
Generate SEO keywords:



C67 =GPT_LIST(A67)

	A	B
65	**Generate SEO keywords**	
66	**Prompt**	**Try it below:**
67	I work in the Digital Marketing team of an e-commerce website that sells all types of agricultural equipment and supplies. Our catalog includes a wide variety of products for crop farming, livestock management, irrigation, harvesting, and more. Please generate a list of 3 relevant SEO keywords in English. These keywords should reflect the main types of products sold on the site, such as tools, machines, accessories, and farming consumables.	Agricultural equipment online
68		Livestock management tools
69		Irrigation system supplies

Generate content ideas side by side:



Learn more on GPT_LIST and GPT_HLIST.

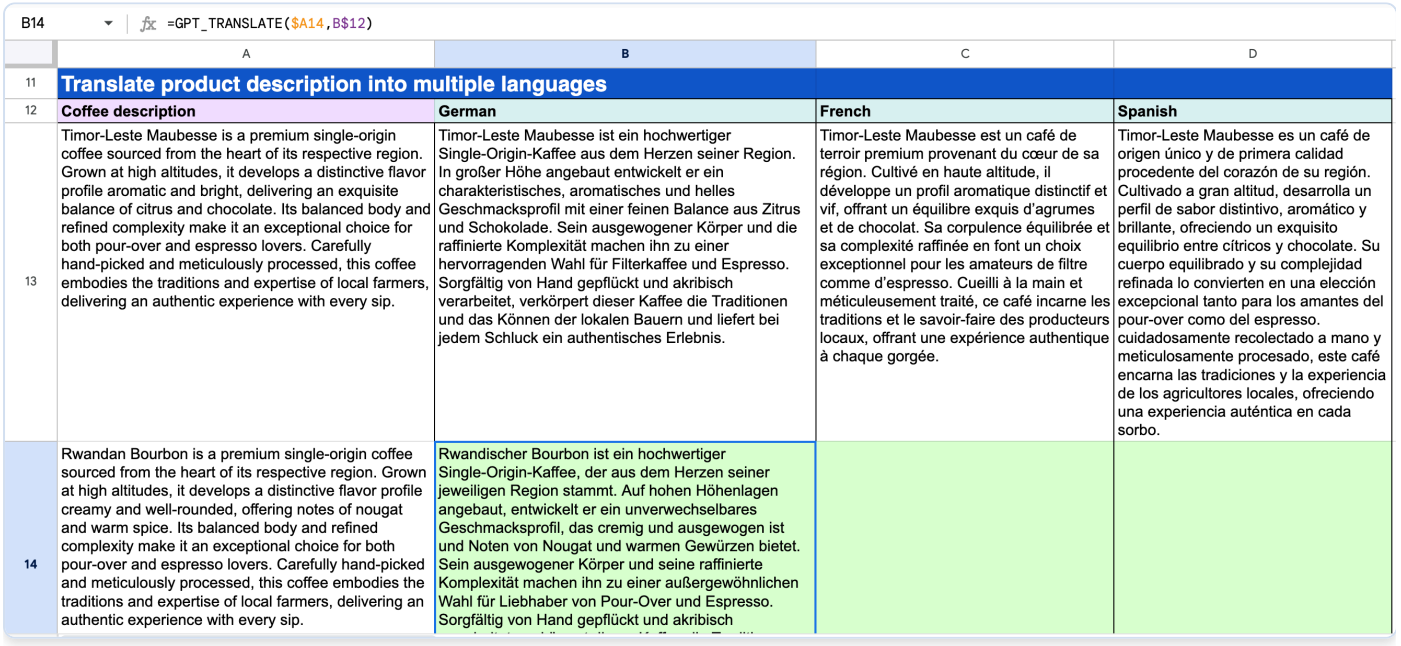
Translate

Leverage large language models for your machine translation tasks and benefit from:

- **High performance**: AI models from leading providers (such as OpenAI) support translation across a wide range of language pairs. See all supported models.
- **Support for many languages**: The full list of supported language pairs isn't public, but ChatGPT is known to work in 80+ languages, with varying accuracy depending on language and context.
- **Integration of your corporate terminology**: For on-brand, consistent translations.

Translate with GPT_TRANSLATE

Translate descriptions:



Translate with specific instructions

Translate with a custom dictionary:

1 Write a prompt and custom dictionary.

	A	B	C	D
25	**Translate with a custom dictionary**			
26	**FIRST STEP: Write a prompt and custom dictionary**			
27	**Prompt**		**Custom dictionary**	
28			**English**	**Spanish**
29	I want you to act as an expert translator. You are given the following table of translations of English to Spanish terms that you must use every time you encounter one of the English terms from the table in the text to translate:		email	correo electrónico
30			spreadsheet	hoja de cálculo
31			meeting	reunión
32			deadline	fecha límite
33			task	tarea
34			AI	IA
35			report	informe
36			team	equipo
37			project	proyecto
38			workflow	flujo de trabajo

2 Create the prompt instructions.

B42 =GPT_CREATE_PROMPT($A$28,$C$29:$D$38)

	A	B
40	**SECOND STEP: Create the prompt instructions**	
41	**Final instructions**	**Try it below:**
42	I want you to act as an expert translator. You are given the following table of translations of English to Spanish terms that you must use every time you encounter one of the English terms from the table in the text to translate: email correo electrónico spreadsheet hoja de cálculo meeting reunión deadline fecha límite task tarea AI IA report informe team equipo project proyecto workflow flujo de trabajo	I want you to act as an expert translator. You are given the following table of translations of English to Spanish terms that you must use every time you encounter one of the English terms from the table in the text to translate: email correo electrónico spreadsheet hoja de cálculo meeting reunión deadline fecha límite task tarea AI IA report informe team equipo project proyecto workflow flujo de trabajo

3 Translate your texts.

B47 =GPT_TRANSLATE(A47,"spanish","english",$B$42)

	A	B
44	**THIRD STEP: Translate your texts**	
45	**Texts to translate**	**Example output**
46	The team finished the report before the deadline.	El equipo terminó el informe antes de la fecha límite.
47	She sent the spreadsheet by email after the meeting.	Ella envió la hoja de cálculo por correo electrónico después de la reunión.
48	Our project uses AI to automate tasks.	
49	The workflow needs to be reviewed before the next meeting.	
50	The manager assigned a new task to the team.	

Your texts are translated in the target language, and the translations comply with the terminology you provided.

Learn more on GPT_TRANSLATE.

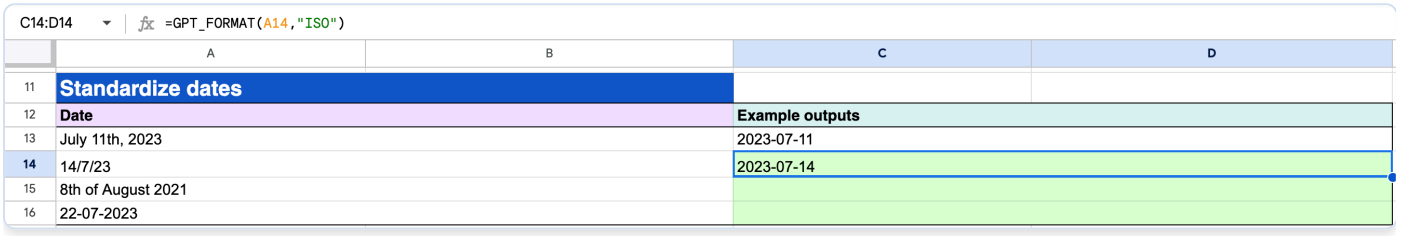
Prepare data

Free up your teams from tedious data preparation so they can concentrate on high-value data analysis:

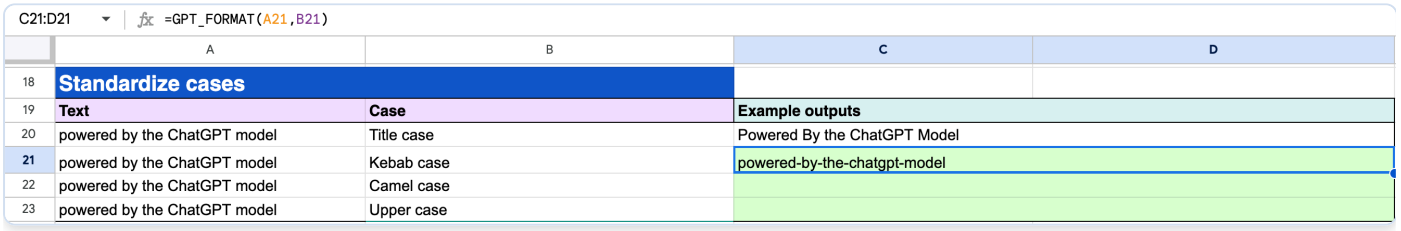
- Basic data cleaning tasks include formatting data according to known standards or by following examples you set.
- More advanced tasks like named entity extraction, which traditionally demanded a specific pattern-based approach, are also an option.

Format data with GPT_FORMAT

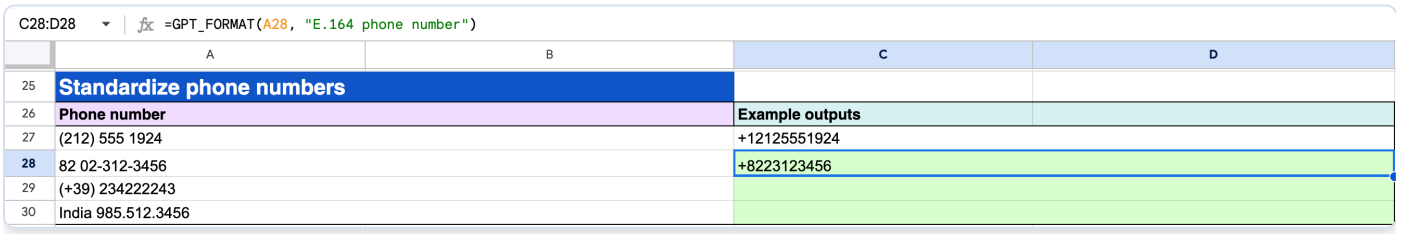
Standardize dates:



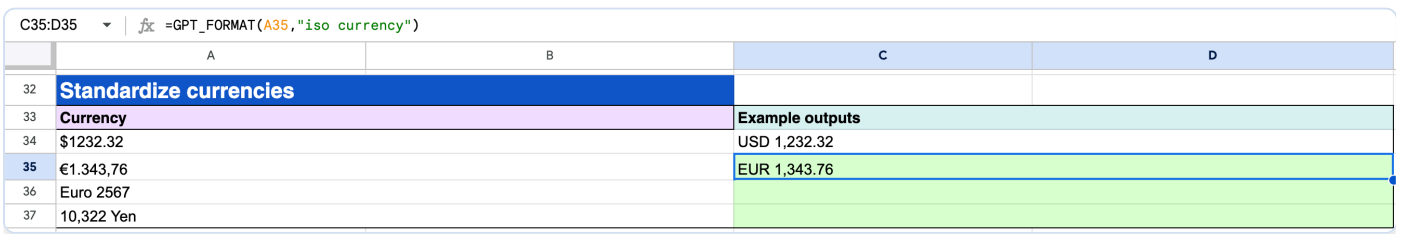
Standardize cases:



Standardize phone numbers:



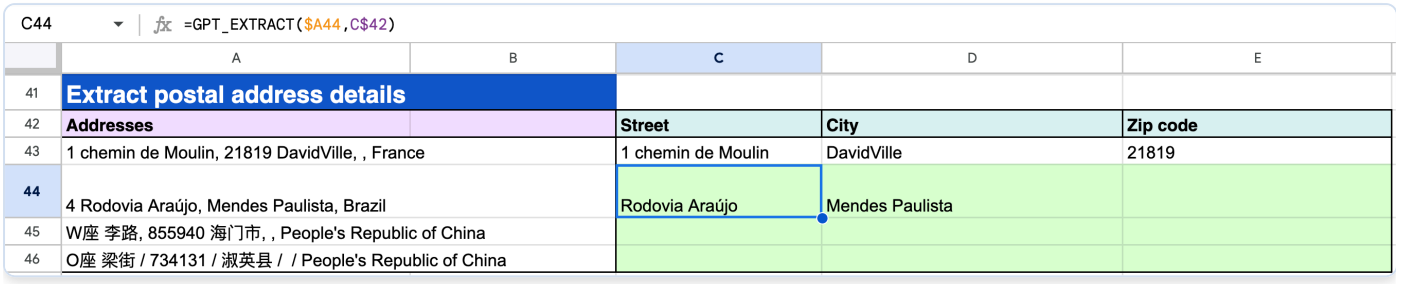
Standardize currencies:



Learn more on GPT_FORMAT.

Extract data with GPT_EXTRACT

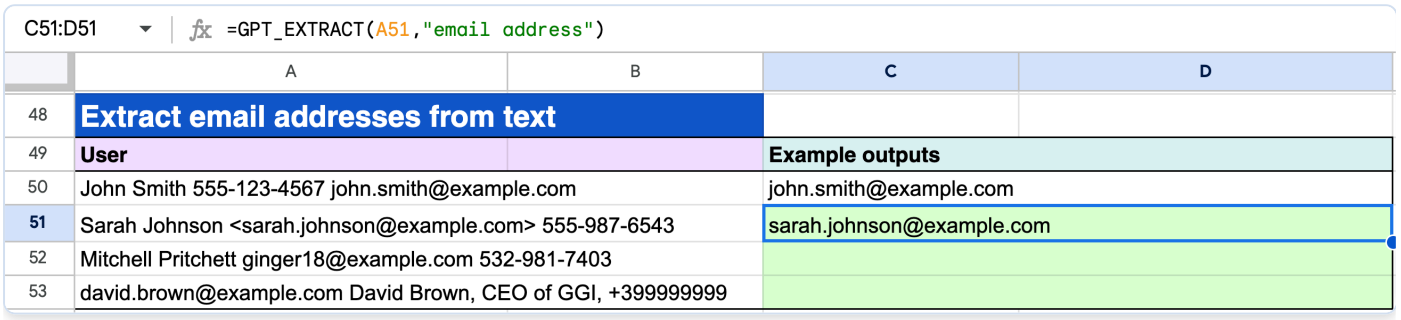
Extract postal addresses:



C44 =GPT_EXTRACT($A44,C$42)

	A	B	C	D	E
41	**Extract postal address details**				
42	**Addresses**		**Street**	**City**	**Zip code**
43	1 chemin de Moulin, 21819 DavidVille, , France		1 chemin de Moulin	DavidVille	21819
44	4 Rodovia Araújo, Mendes Paulista, Brazil		Rodovia Araújo	Mendes Paulista	
45	W座 李路, 855940 海门市, , People's Republic of China				
46	O座 梁街 / 734131 / 淑英县 / / People's Republic of China				

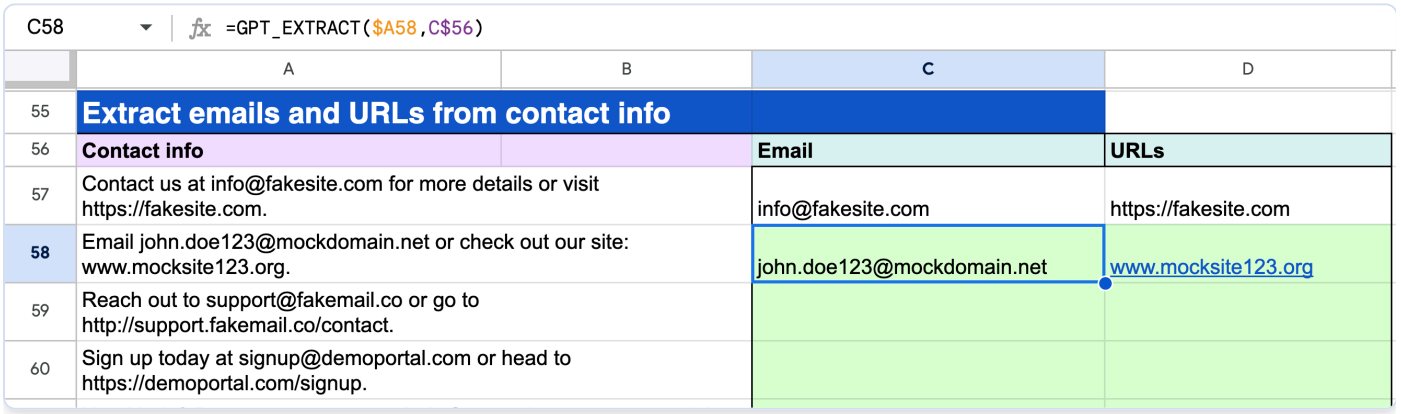
Extract email addresses:



C51:D51 =GPT_EXTRACT(A51,"email address")

	A	B	C	D
48	**Extract email addresses from text**			
49	**User**		**Example outputs**	
50	John Smith 555-123-4567 john.smith@example.com		john.smith@example.com	
51	Sarah Johnson <sarah.johnson@example.com> 555-987-6543		sarah.johnson@example.com	
52	Mitchell Pritchett ginger18@example.com 532-981-7403			
53	david.brown@example.com David Brown, CEO of GGI, +399999999			

Extract emails and URLs:



C58 =GPT_EXTRACT($A58,C$56)

	A	B	C	D
55	**Extract emails and URLs from contact info**			
56	**Contact info**		**Email**	**URLs**
57	Contact us at info@fakesite.com for more details or visit https://fakesite.com.		info@fakesite.com	https://fakesite.com
58	Email john.doe123@mockdomain.net or check out our site: www.mocksite123.org.		john.doe123@mockdomain.net	www.mocksite123.org
59	Reach out to support@fakemail.co or go to http://support.fakemail.co/contact.			
60	Sign up today at signup@demoportal.com or head to https://demoportal.com/signup.			

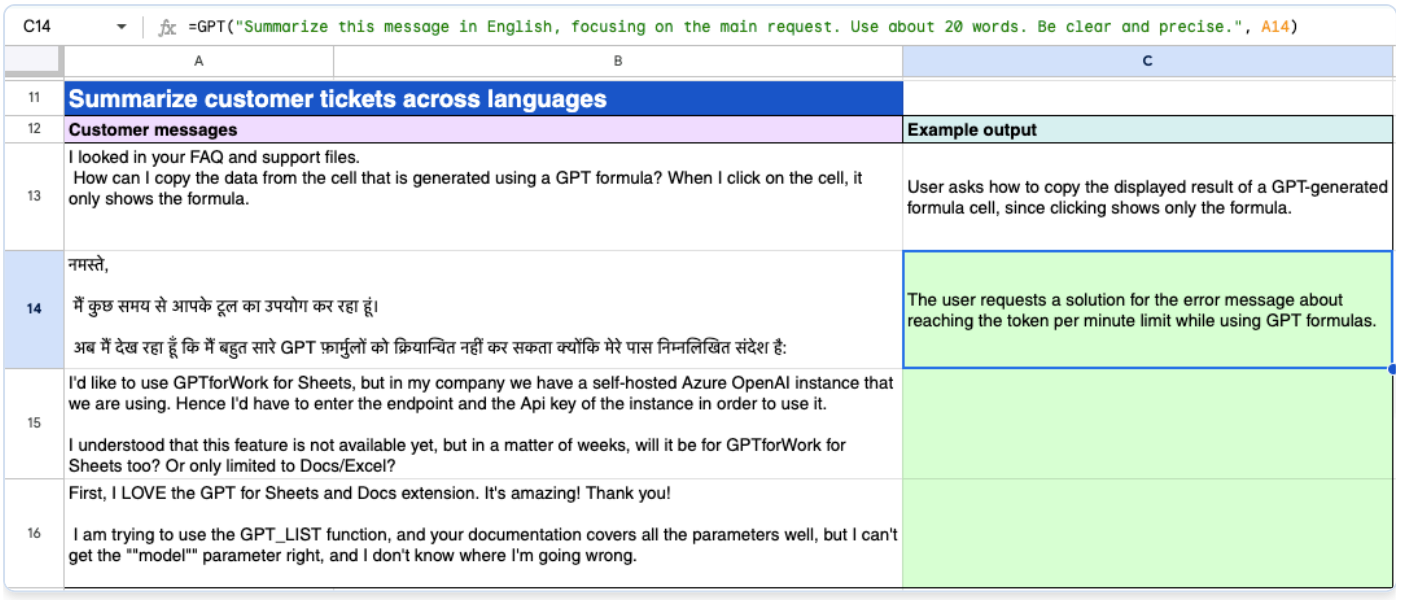
Learn more on GPT_EXTRACT.

Analyze

Get instant insights from large volumes of content through automated summarizing, categorization, tagging, and sentiment analysis.

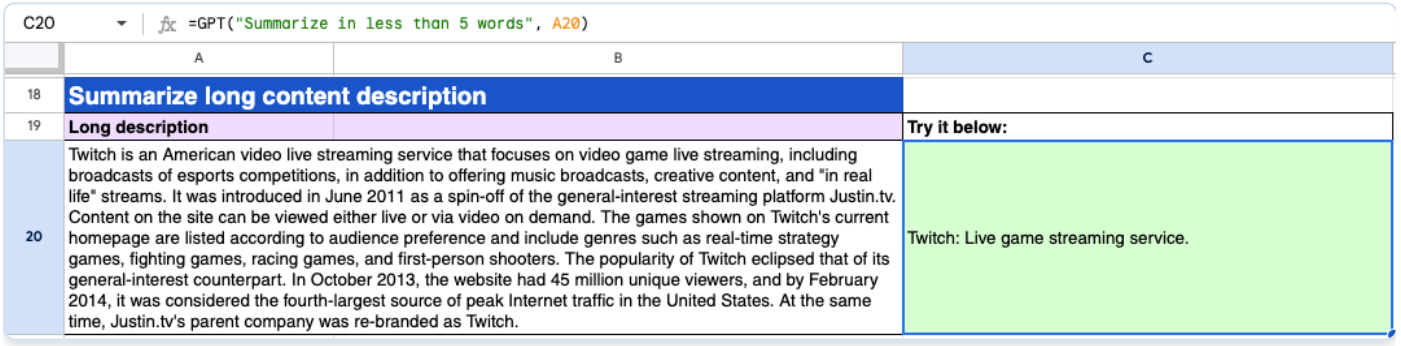
Summarize data with GPT

Summarize customer tickets across languages:



	A	B	C
11	**Summarize customer tickets across languages**		
12	**Customer messages**		**Example output**
13	I looked in your FAQ and support files. How can I copy the data from the cell that is generated using a GPT formula? When I click on the cell, it only shows the formula.		User asks how to copy the displayed result of a GPT-generated formula cell, since clicking shows only the formula.
14	नमस्ते, मैं कुछ समय से आपके टूल का उपयोग कर रहा हूं। अब मैं देख रहा हूँ कि मैं बहुत सारे GPT फ़ार्मुलों को क्रियान्वित नहीं कर सकता क्योंकि मेरे पास निम्नलिखित संदेश है:		The user requests a solution for the error message about reaching the token per minute limit while using GPT formulas.
15	I'd like to use GPTforWork for Sheets, but in my company we have a self-hosted Azure OpenAI instance that we are using. Hence I'd have to enter the endpoint and the Api key of the instance in order to use it. I understood that this feature is not available yet, but in a matter of weeks, will it be for GPTforWork for Sheets too? Or only limited to Docs/Excel?		
16	First, I LOVE the GPT for Sheets and Docs extension. It's amazing! Thank you! I am trying to use the GPT_LIST function, and your documentation covers all the parameters well, but I can't get the ""model"" parameter right, and I don't know where I'm going wrong.		

Summarize long descriptions:

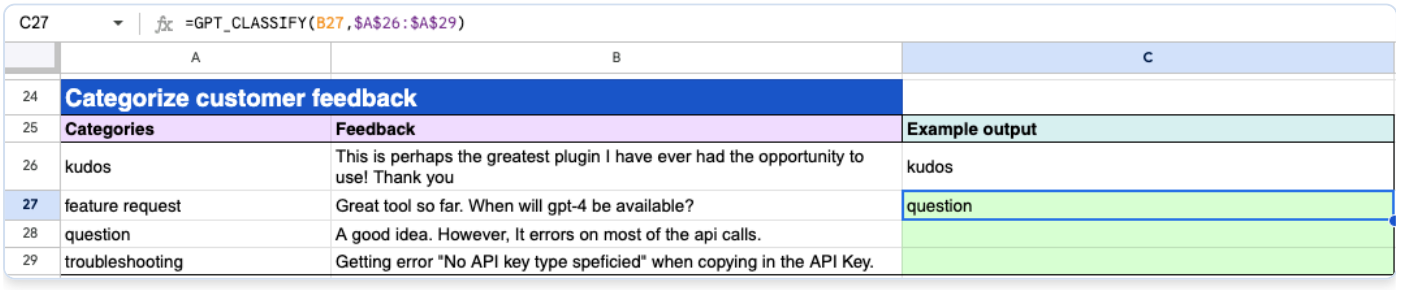


	A	B	C
18	**Summarize long content description**		
19	**Long description**		**Try it below:**
20	Twitch is an American video live streaming service that focuses on video game live streaming, including broadcasts of esports competitions, in addition to offering music broadcasts, creative content, and "in real life" streams. It was introduced in June 2011 as a spin-off of the general-interest streaming platform Justin.tv. Content on the site can be viewed either live or via video on demand. The games shown on Twitch's current homepage are listed according to audience preference and include genres such as real-time strategy games, fighting games, racing games, and first-person shooters. The popularity of Twitch eclipsed that of its general-interest counterpart. In October 2013, the website had 45 million unique viewers, and by February 2014, it was considered the fourth-largest source of peak Internet traffic in the United States. At the same time, Justin.tv's parent company was re-branded as Twitch.		Twitch: Live game streaming service.

Learn more on GPT.

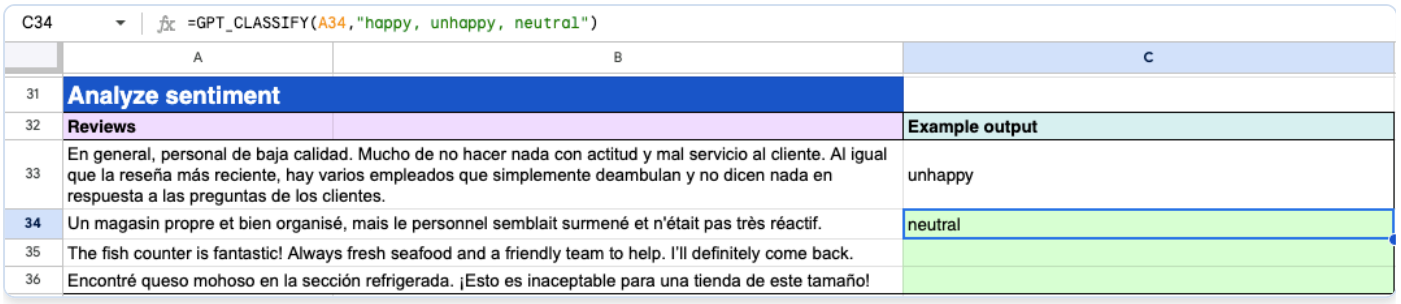
Categorize or classify data with GPT_CLASSIFY

Categorize customer feedback:



	A	B	C
24	**Categorize customer feedback**		
25	**Categories**	**Feedback**	**Example output**
26	kudos	This is perhaps the greatest plugin I have ever had the opportunity to use! Thank you	kudos
27	feature request	Great tool so far. When will gpt-4 be available?	question
28	question	A good idea. However, It errors on most of the api calls.	
29	troubleshooting	Getting error "No API key type speficied" when copying in the API Key.	

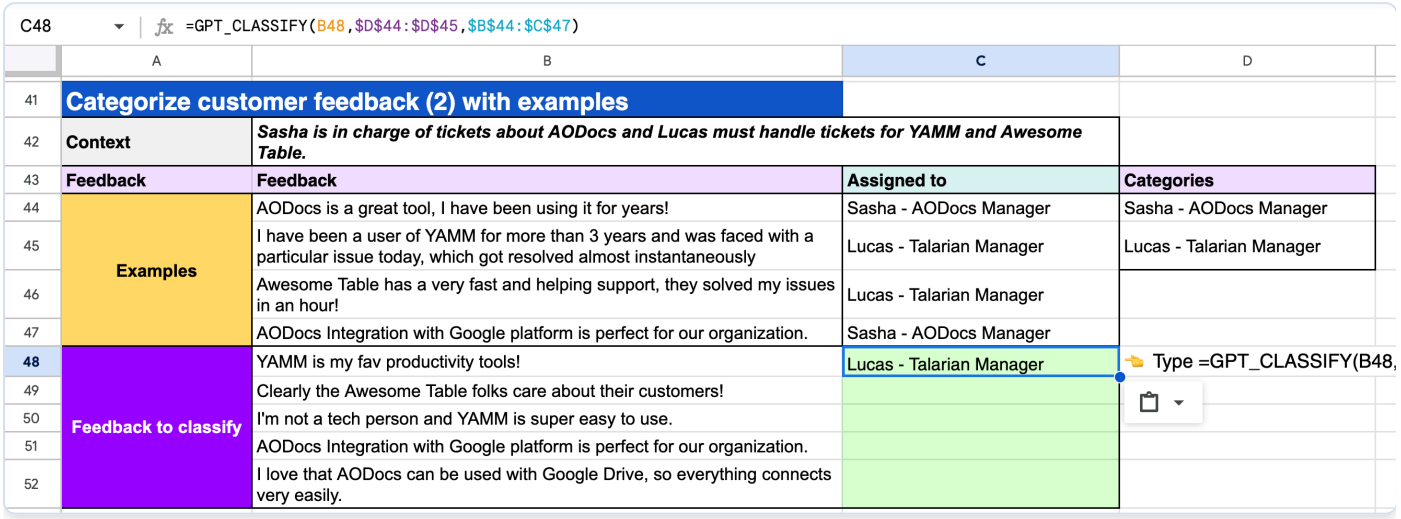
Analyze sentiment:



C34 =GPT_CLASSIFY(A34,"happy, unhappy, neutral")

	A	B	C
31	**Analyze sentiment**		
32	**Reviews**		**Example output**
33	En general, personal de baja calidad. Mucho de no hacer nada con actitud y mal servicio al cliente. Al igual que la reseña más reciente, hay varios empleados que simplemente deambulan y no dicen nada en respuesta a las preguntas de los clientes.		unhappy
34	Un magasin propre et bien organisé, mais le personnel semblait surmené et n'était pas très réactif.		neutral
35	The fish counter is fantastic! Always fresh seafood and a friendly team to help. I'll definitely come back.		
36	Encontré queso mohoso en la sección refrigerada. ¡Esto es inaceptable para una tienda de este tamaño!		

Categorize customer feedback with examples:



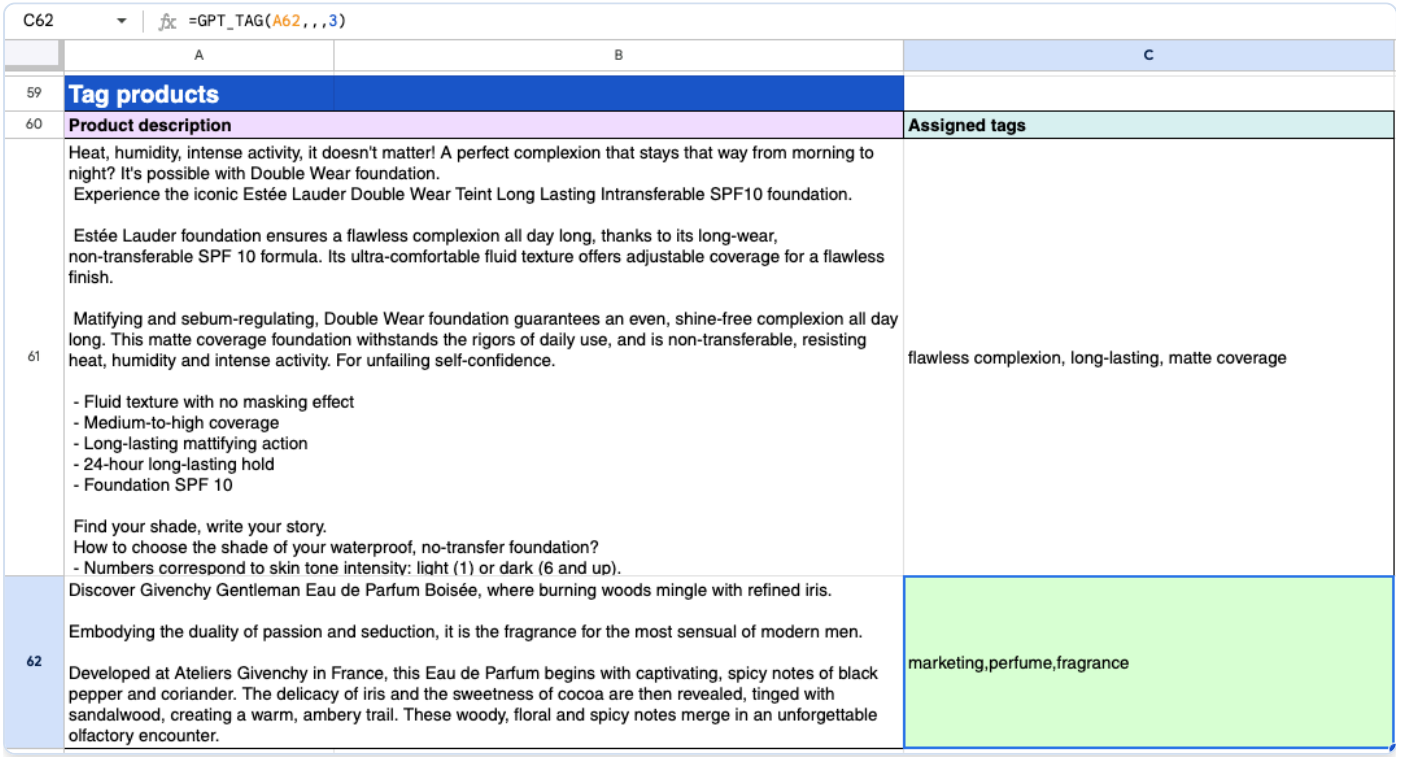
C48 =GPT_CLASSIFY(B48,$D$44:$D$45,$B$44:$C$47)

	A	B	C	D
41	**Categorize customer feedback (2) with examples**			
42	**Context**	***Sasha is in charge of tickets about AODocs and Lucas must handle tickets for YAMM and Awesome Table.***		
43	**Feedback**	**Feedback**	**Assigned to**	**Categories**
44	**Examples**	AODocs is a great tool, I have been using it for years!	Sasha - AODocs Manager	Sasha - AODocs Manager
45		I have been a user of YAMM for more than 3 years and was faced with a particular issue today, which got resolved almost instantaneously	Lucas - Talarian Manager	Lucas - Talarian Manager
46		Awesome Table has a very fast and helping support, they solved my issues in an hour!	Lucas - Talarian Manager	
47		AODocs Integration with Google platform is perfect for our organization.	Sasha - AODocs Manager	
48	**Feedback to classify**	YAMM is my fav productivity tools!	Lucas - Talarian Manager	Type =GPT_CLASSIFY(B48,
49		Clearly the Awesome Table folks care about their customers!		
50		I'm not a tech person and YAMM is super easy to use.		
51		AODocs Integration with Google platform is perfect for our organization.		
52		I love that AODocs can be used with Google Drive, so everything connects very easily.		

Learn more on GPT_CLASSIFY.

Assign tags with GPT_TAG

Tag products:



C62 =GPT_TAG(A62,,,3)

	A	B	C
59	**Tag products**		
60	**Product description**		**Assigned tags**
61	Heat, humidity, intense activity, it doesn't matter! A perfect complexion that stays that way from morning to night? It's possible with Double Wear foundation. Experience the iconic Estée Lauder Double Wear Teint Long Lasting Intransferable SPF10 foundation. Estée Lauder foundation ensures a flawless complexion all day long, thanks to its long-wear, non-transferable SPF 10 formula. Its ultra-comfortable fluid texture offers adjustable coverage for a flawless finish. Matifying and sebum-regulating, Double Wear foundation guarantees an even, shine-free complexion all day long. This matte coverage foundation withstands the rigors of daily use, and is non-transferable, resisting heat, humidity and intense activity. For unfailing self-confidence. - Fluid texture with no masking effect - Medium-to-high coverage - Long-lasting mattifying action - 24-hour long-lasting hold - Foundation SPF 10 Find your shade, write your story. How to choose the shade of your waterproof, no-transfer foundation? - Numbers correspond to skin tone intensity: light (1) or dark (6 and up).		flawless complexion, long-lasting, matte coverage
62	Discover Givenchy Gentleman Eau de Parfum Boisée, where burning woods mingle with refined iris. Embodying the duality of passion and seduction, it is the fragrance for the most sensual of modern men. Developed at Ateliers Givenchy in France, this Eau de Parfum begins with captivating, spicy notes of black pepper and coriander. The delicacy of iris and the sweetness of cocoa are then revealed, tinged with sandalwood, creating a warm, ambery trail. These woody, floral and spicy notes merge in an unforgettable olfactory encounter.		marketing,perfume,fragrance

Learn more on GPT_TAG.

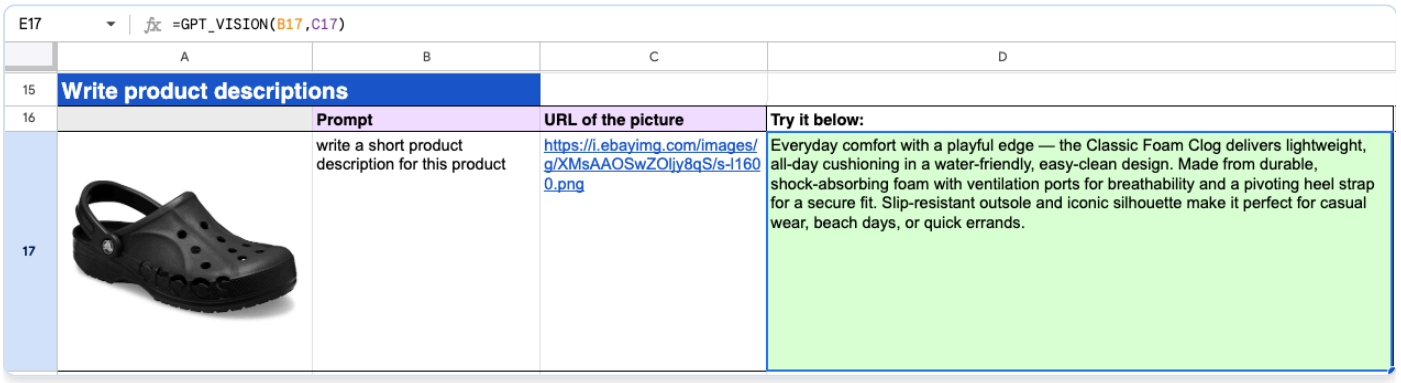
Analyze images with GPT_VISION

Applies a prompt to an image. GPT_VISION is the only function that can process images. For example, you can ask a question about the image, generate a description or generate alt text for images.

> **INFO**
>
> To use GPT_VISION, you need to select a model that supports images (Vision).

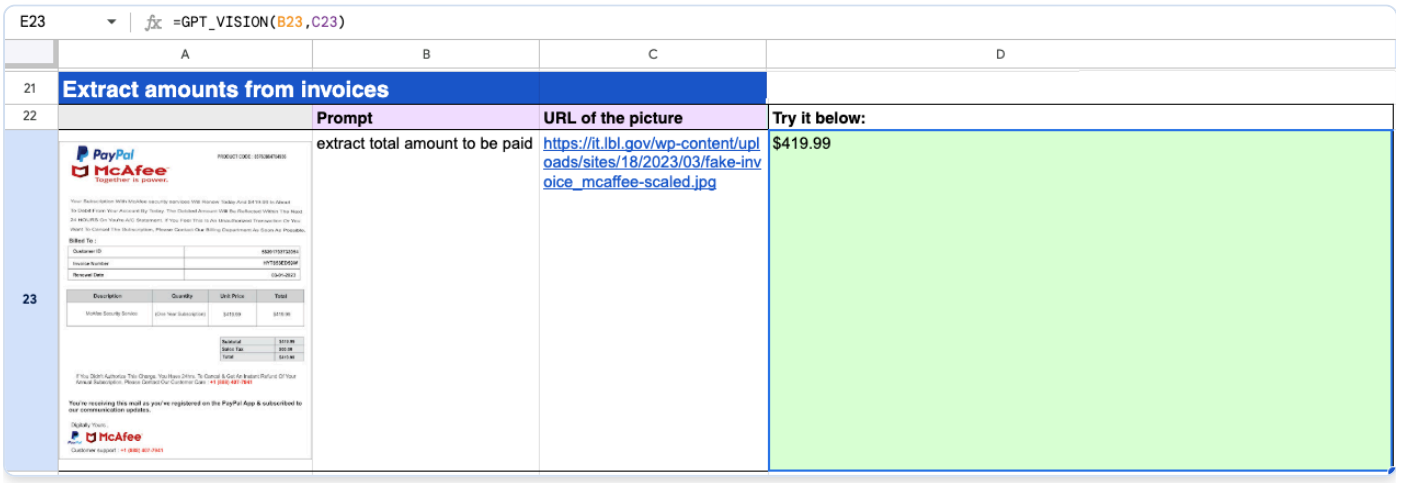
Create content from an image

Write product descriptions:



Extract information from an image

Extract amounts from invoices:



Extract info from an ID:

E27 =GPT_VISION(B27,C27)

	A	B	C	D
25	**Extract info from an ID**			
26		**Prompt**	**URL of the picture**	**Try it below:**
27		extract the surname	https://lecap.consulfrance.org/IMG/arton465.png?1641476764	Martin

Explain elements in an image

Identify elements of an image:

D33 =GPT_VISION(B33,C33)

	A	B	C	D
31	**Identify elements of an image**			
32		**Prompt**	**URL of the picture**	**Try it below:**
33		What brand does this logo represent?	http://upload.wikimedia.org/wikipedia/commons/thumb/b/b7/X_logo.jpg/250px-X_logo.jpg	The logo represents the brand "X," which is the rebranded name for Twitter.

Understand an image:

E37 =GPT_VISION(B37,C37)

	A	B	C	D
35	**Understand an image**			
36		**Prompt**	**URL of the picture**	**Try it below:**
37		which bone is indicated?	https://media.brainly.com/image/rs:fill/w:640/q:75/plain/https://us-static.z-dn.net/files/d46/2014a252ef1de67bec89caab341f21cc.jpg	The arrow points to the femur (thigh bone).

Explain formulas:

E41 =GPT_VISION(B41,C41)

	A	B	C	D
39	**Explain formulas**			
40		**Prompt**	**URL of the picture**	**Try it below:**
41	$i\hbar\frac{\partial}{\partial t}\Psi(x,t) = \left[-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial x^2} + V(x,t)\right]\Psi(x,t)$	explain this formula	https://miro.medium.com/v2/resize:fit:1400/format:webp/1*J-JQKUkHMo-iE4qIki373A.png	This is the time-dependent nonrelativistic Schrödinger equation in one spatial dimension. It governs how the complex wavefunction Ψ(x,t) of a quantum particle evolves in time. Breaking it down term-by-term: - iħ ∂Ψ/∂t : the left-hand side. ħ (h-bar) is the reduced Planck constant and i is the imaginary unit. This term represents the time evolution (rate of change) of the quantum state. - The right-hand side is an operator (the Hamiltonian) acting on Ψ: - [− ħ²/(2m) ∂²/∂x²] Ψ(x,t) : the kinetic-energy operator. Applying the second spatial derivative to Ψ gives the quantum analogue of kinetic energy; m is the particle mass. - + V(x,t) Ψ(x,t) : the potential-energy term. V(x,t) is the classical potential energy at position x and time t; multiplying by Ψ gives the potential energy contribution. In words: iħ times the time derivative of the wavefunction equals the total energy operator (kinetic + potential) acting on the wavefunction. Solutions Ψ(x,t) encode all measurable information: |Ψ(x,t)|² is the probability density to find the particle at position x at time t. For time-independent V one often separates variables to get stationary states satisfying the time-independent Schrödinger equation: [−ħ²/(2m) d²/dx² + V(x)]ψ(x) = Eψ(x).

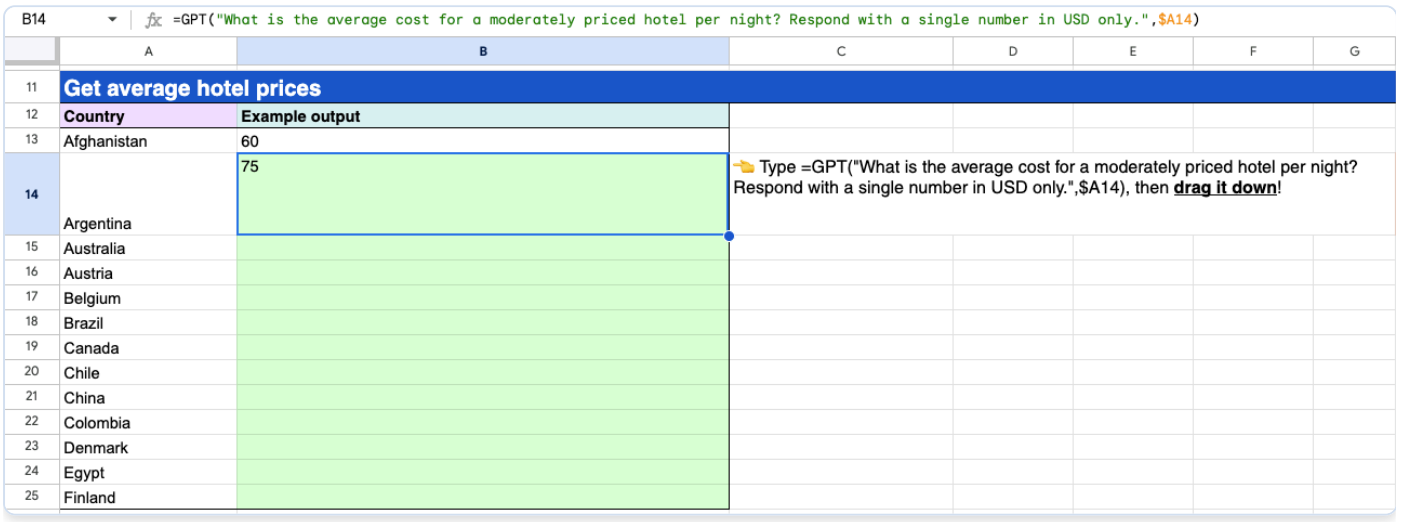
Learn more on GPT_VISION.

Search the web

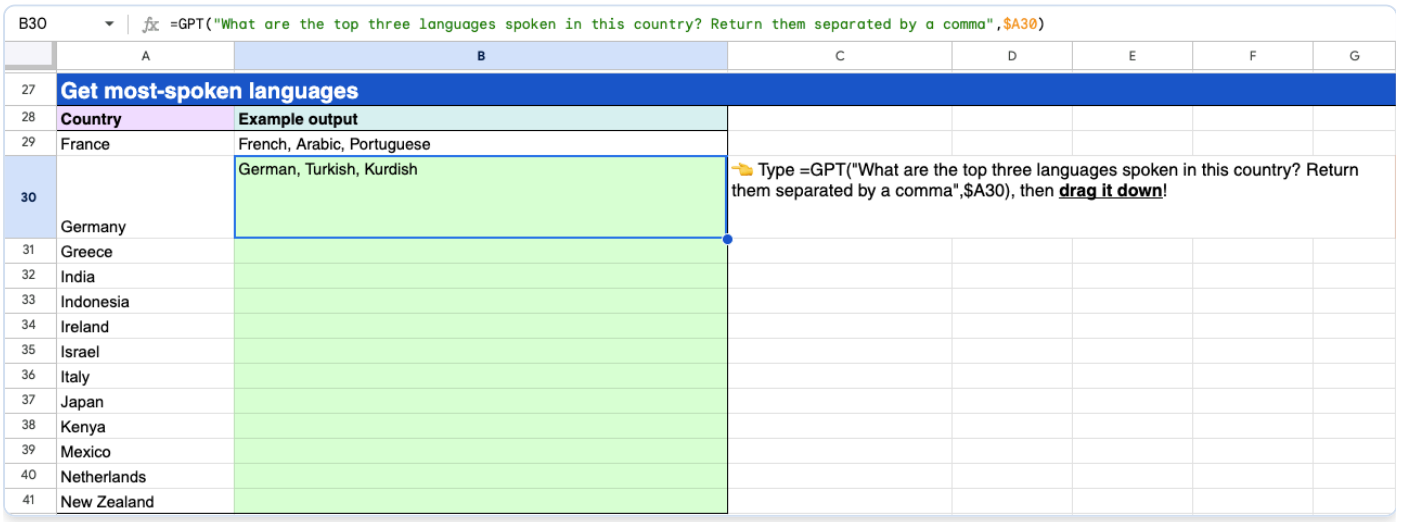
Get answers from the web with GPT

> **INFO**
>
> To search the web with a GPT function, select a model with web search capabilities.

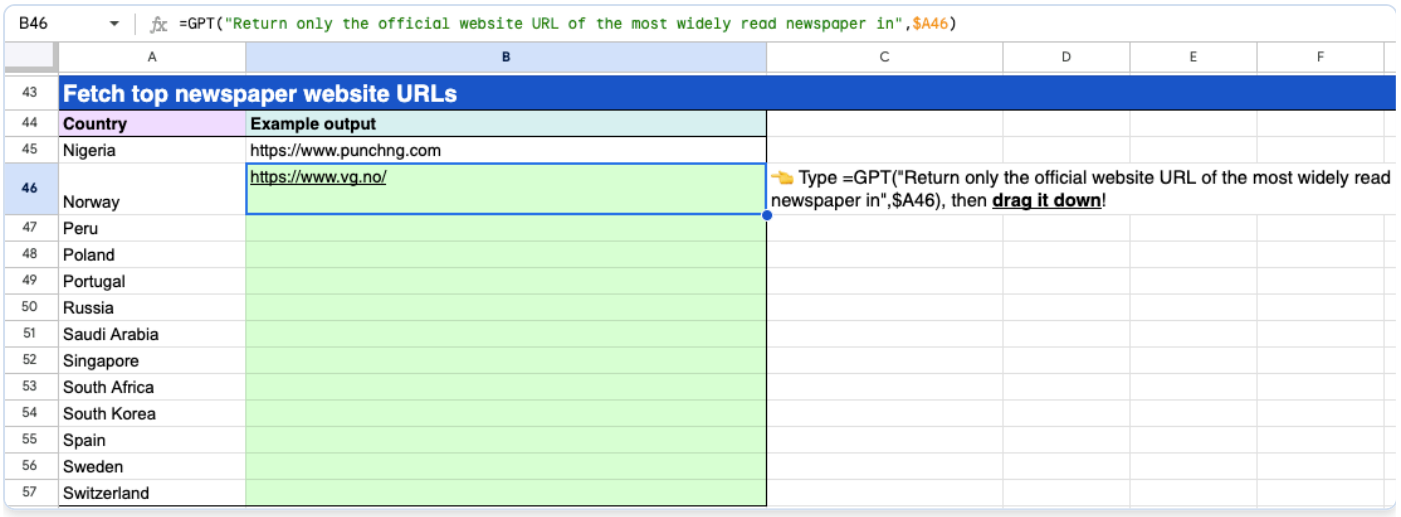
Get average hotel prices:



Get most-spoken languages:



Fetch top newspaper website URLs:


B46
=GPT("Return only the official website URL of the most widely read newspaper in",$A46)
Fetch top newspaper website URLs
Country
Example output
Nigeria
https://www.punchng.com
Norway
https://www.vg.no/
Peru
Poland
Portugal
Russia
Saudi Arabia
Singapore
South Africa
South Korea
Spain
Sweden
Switzerland
Type =GPT("Return only the official website URL of the most widely read newspaper in",$A46), then drag it down!

Manage GPT formulas from the sidebar in Excel

Use the GPT for Excel sidebar to manage your GPT formulas and avoid unexpected changes and costs.

Retry GPT formulas on error

Re-execute all formulas on error in your current sheet.

Regenerate GPT formula results

Re-execute selected formulas in your current sheet.

Replace GPT formulas with results

Replace GPT formulas with their results to prevent recalculation permanently.

Manage GPT formula cache

Minimize costs by reusing previous results.

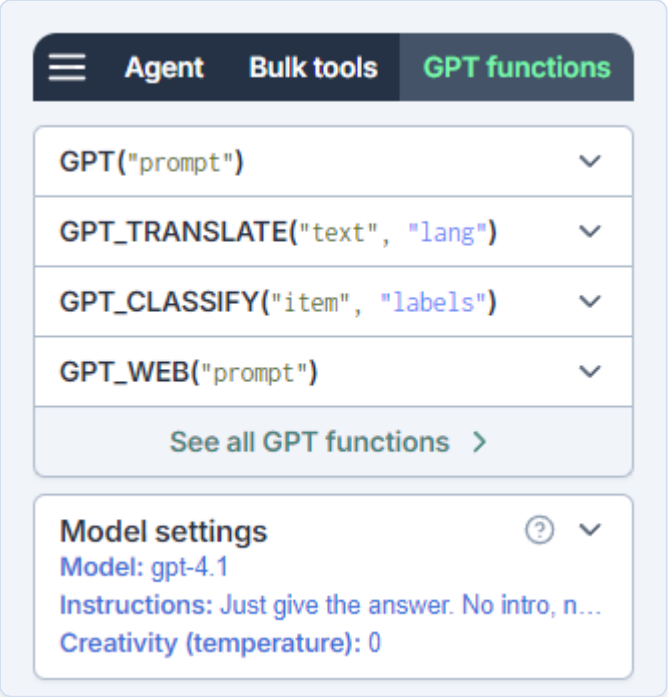
Disable GPT formulas

Prevent the execution of GPT formulas in your workbook for an indefinite period of time.

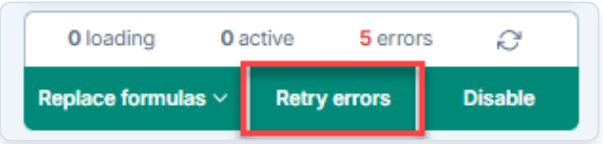
Retry GPT formulas on error in Excel

After troubleshooting GPT for Excel errors, you can recalculate all GPT formulas in cells tagged with `#ERROR!` or `#VALUE!`.

1. In the sidebar, select **GPT functions**.



2. Check that the GPT formulas aren't disabled for your workbook.
3. Click **Retry errors**.



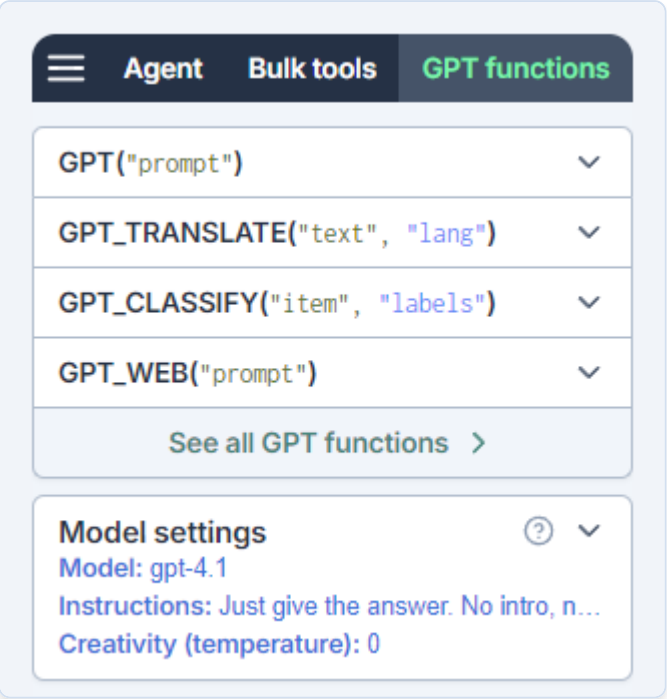
All GPT formulas that were on error in the current sheet are recalculated. If some errors persist, let us know by submitting a support request.

> **TIP**
>
> Replace GPT formulas with their results in the selected cells to prevent any further recalculation.

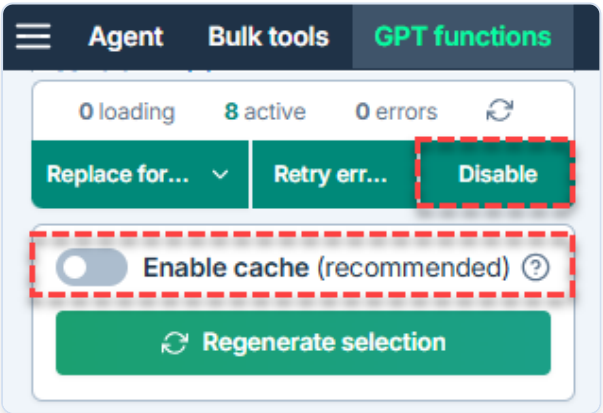
Regenerate GPT formula results in Excel

Re-execute GPT formulas to update results in selected cells.

1. In the sidebar, select **GPT functions**.

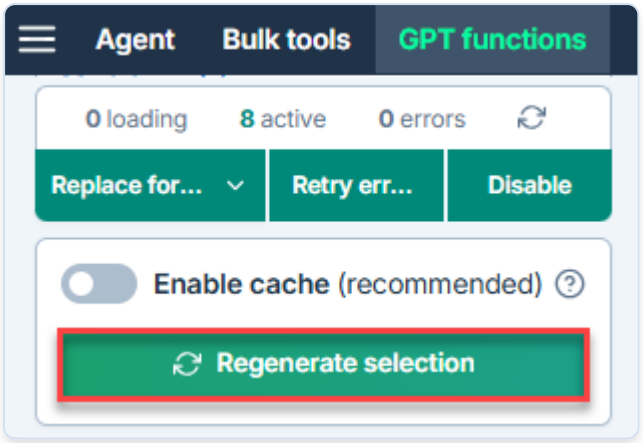


2. Check that the GPT formulas are not disabled and that caching is disabled for your workbook.



3. Select cells with active GPT formulas in the current sheet. Ensure all cells generated by formulas like GPT_LIST are included to prevent data loss.

4 Click **Regenerate selection**.



5 (Optional) Disable the formulas to prevent recalculations in this sheet and the rest of the document.

GPT formula results within the selected cells are regenerated.

Replace GPT formulas with results in Excel

Convert GPT formulas to static values and prevent any further recalculation, either on a selected number of cells or across an entire sheet. If you need to recover the original formulas later, you can use the version history of your Microsoft workbook.

Replace formulas manually

1. Select the cells that contain the GPT formulas.
2. Copy them by pressing **Ctrl+C** (or **Cmd+C** on Mac).

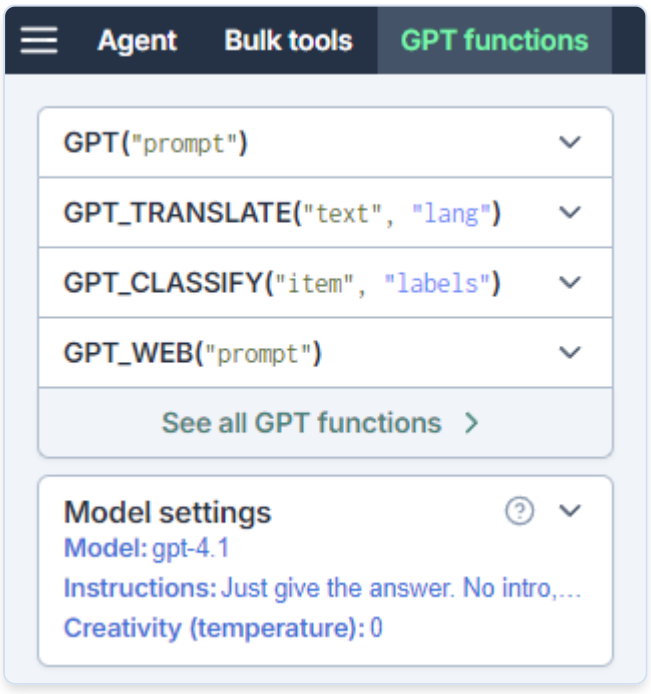
   > **INFO**
   >
   > For GPT formulas with multi-row results, copy all rows with results, not just the formula row.

3. Paste the results as values by pressing **Ctrl+Shift+V** (or **Cmd+Shift+V** on Mac).

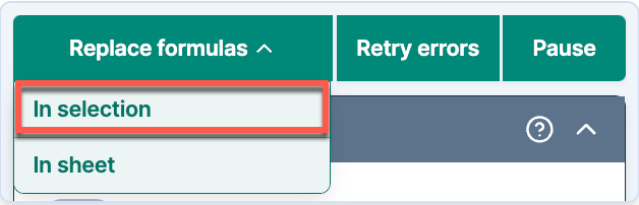
GPT formulas within the manually selected cells are replaced by their respective results.

Replace formulas in selection

1. Select cells with active GPT formulas in the current sheet, ensuring all cells generated by formulas like GPT_LIST are included to prevent data loss.
2. In the sidebar, select **GPT functions**.



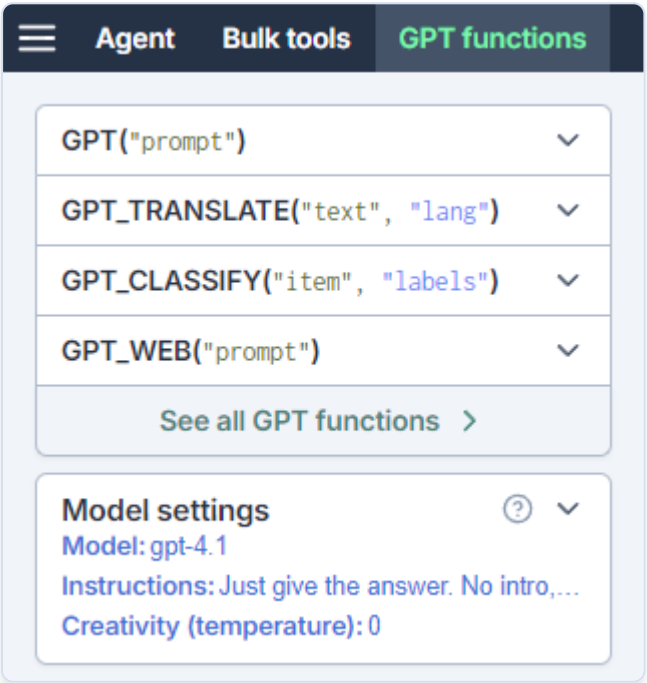
3 Choose **In selection** in the *Replace formulas* dropdown.



GPT formulas within the selected cells are replaced by their respective results.

Replace formulas in entire sheet

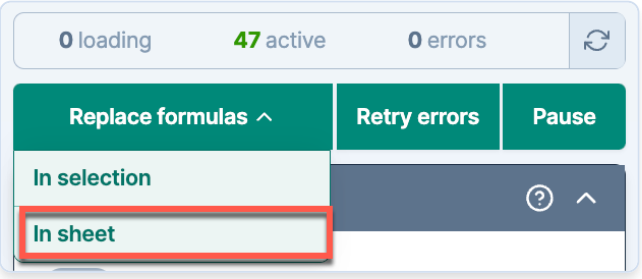
1 In the sidebar, select **GPT functions**.



2 Refresh the formula count to see the number of active GPT formulas in the current sheet. Formulas on error won't be replaced.



3 Choose **In sheet** in the *Replace formulas* dropdown.



All GPT formulas within the sheet are replaced by their corresponding results.

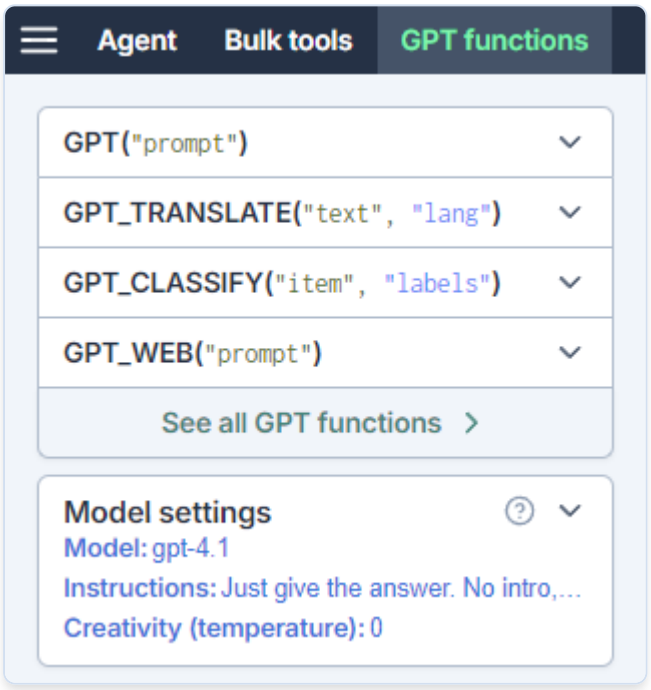
Manage GPT formula cache for your spreadsheet in Excel

Cache the results of your GPT formulas to minimize costs by reusing previous results, or disable caching to experiment with sidebar settings.

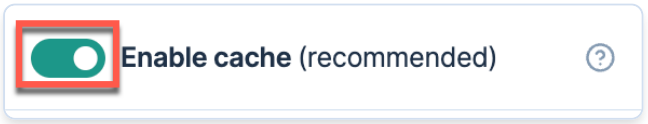
Enable cache to save costs and reuse previous results when reloading a formula

You can enable cache usage to avoid being charged for reusing formulas. This mode ensures that results for the same formula are retrieved from the cache instead of generating new results.

1. In the sidebar, select **GPT functions**.



2. Click on the **Enable cache** switch to start saving results.

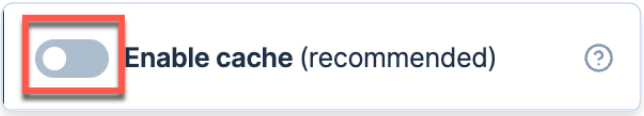


Caching is enabled for GPT formula results in your workbook. You can now save costs and reuse all future results of your GPT formulas.

Disable cache to experiment with sidebar settings

If you want to quickly experiment with various settings, you can disable cache usage. Disabling your cache doesn't clear its content. It stops storing new executions and fetching existing ones from it.

1. (Optional) Replace GPT formulas with their results if you want to preserve the current results.
2. In the sidebar, select **GPT functions**.
3. Click on the **Enable cache** switch to put it in disabled state.

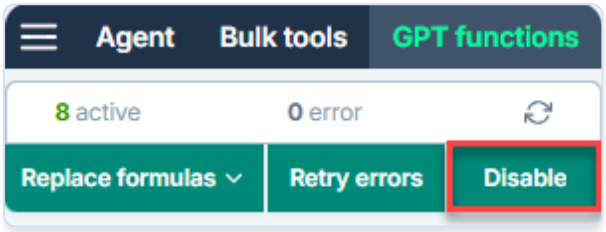


Caching is disabled for GPT formula results in your workbook. Each time the formulas are reloaded, they will be executed again, incur costs, and may return new results.

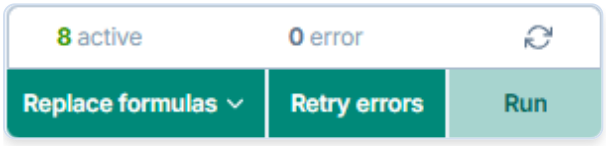
Disable GPT formulas for your spreadsheet in Excel

If the GPT for Excel cache is disabled, reloading active GPT formulas triggers their re-execution, which incurs costs and potentially returns new results. To preserve your results and avoid additional costs, disable GPT formula execution in your workbook.

In the sidebar, select **GPT functions**, and click **Disable**.



The button label changes to **Run**.



The GPT formulas in your workbook will not be executed until you click **Run**. If your workbook contains multiple sheets, the GPT formulas across all these sheets are prevented from executing.

Search the web with GPT functions in Excel

Search the web directly in GPT for Excel using GPT functions with web search models. These models use fresh web search data, which allows you to generate up-to-date responses about companies, current events, facts, and more, directly in your workbook. You can also fetch content from specific URLs to extract information from web pages.

Choose the right web search model

Choose a web search model that best fits your use case.

Web search

Get information that is not in the model's built-in knowledge:

- **Example:** Finding the current CEO of a company.
- **How it works:**

  All web search models can go beyond their built-in knowledge by accessing information from search engines, like Google or Bing. This allows models, which have compressed historical knowledge, to explore topics in greater depth, with more current data or with a specific focus.
- **Recommended models:**
  - perplexity-fast for quick results (cheapest).
  - OpenAI Search models for high quality results.
  - Gemini models with Web search for Google search quality (most expensive).

Content fetching

Focus on specific URLs and extract content from them:

- **Example:** Extracting the product name and price from the URL of a product page.
- **How it works:** Some web search models can access the full content from specific URLs. The full content is added to the context of the models, so they can use it to answer questions about the content or extract information from it.
- **Recommended models:**
  - Gemini 2.5 models with URL context

Supported web search models

Click here for a list of all web search models supported by GPT for Excel ▼

Provider	Model	Built-in	Provider-connected 🔑	URL fetch
Search...	Search...	All	All	All
Perplexity	perplexity-fast	Yes	Yes	No
Perplexity	perplexity-low	Yes	Yes	No
Perplexity	perplexity-medium	No	Yes	No
Google	gemini-2.5-flash	Yes	Yes	Yes
Google	gemini-2.5-flash-lite	No	Yes	Yes
Google	gemini-2.5-pro	No	Yes	Yes
Google	gemini-3-flash-preview	No	Yes	Yes
Google	gemini-3.1-flash-lite	No	Yes	Yes
Google	gemini-3.1-pro-preview	No	Yes	Yes
Google	gemini-3.5-flash	No	Yes	Yes
OpenAI	gpt-5-search-api	Yes	Yes	No
Parallel	speed	Yes	No	No

For detailed pricing information, see Model prices.

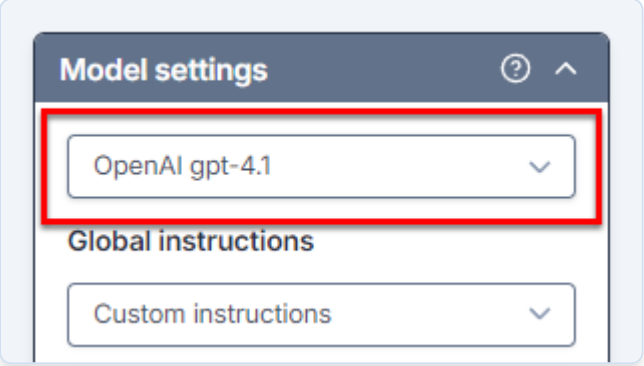
Use web search models

Use GPT functions to perform web searches with any model that has web search capabilities (🌐).

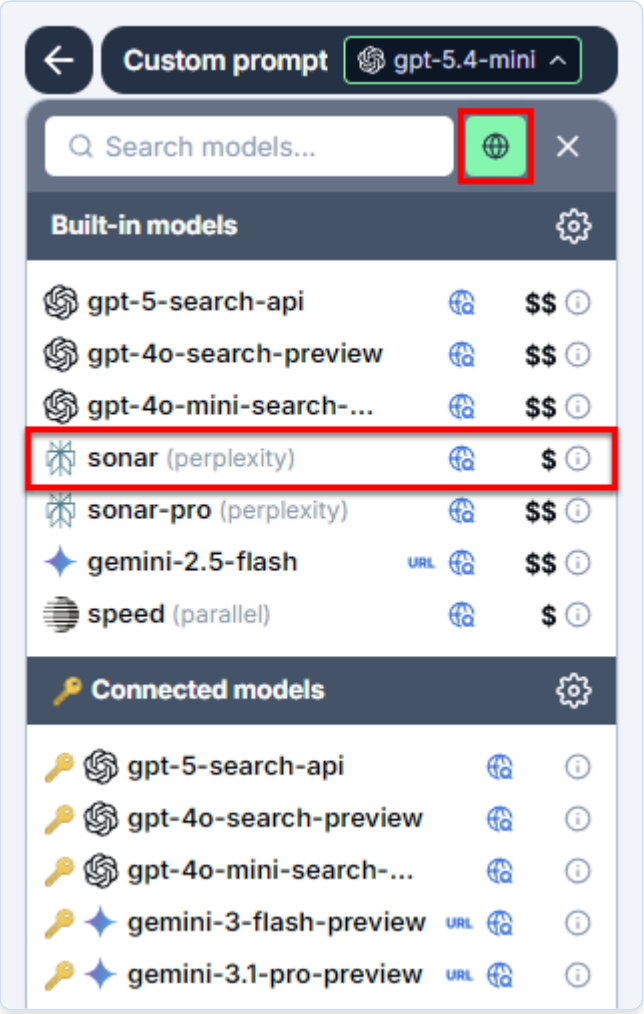
Using the model switcher

1. Open GPT for Excel.
2. In the sidebar, select **GPT functions**.

3 Expand **Model settings** and open the model switcher.



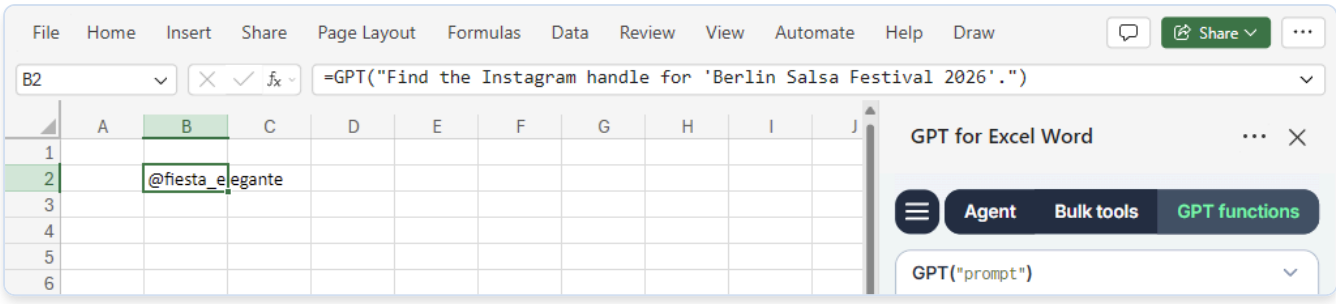
4 Enable the web search filter and select a web search model (🌐).



5 (Optional) Configure the model-specific settings in the sidebar. For Gemini models, make sure web search is enabled.

6 In the workbook, select a cell and write your formula:

```
=GPT("Find the Instagram handle for 'Berlin Salsa Festival 2026'.")
```

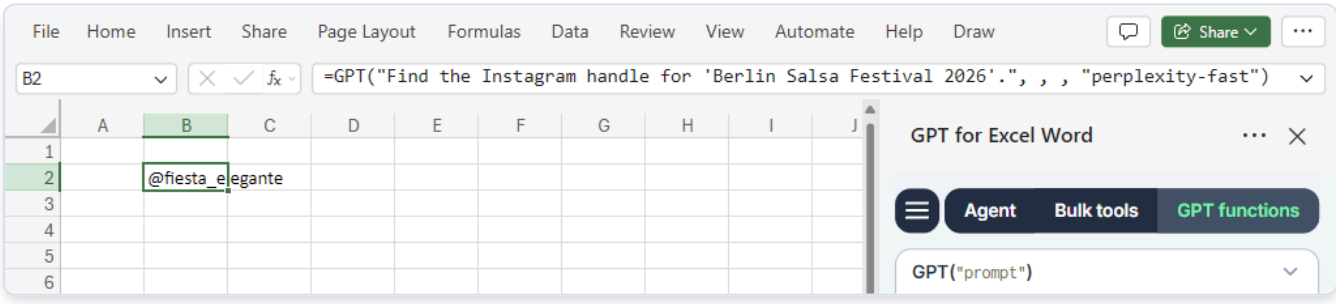


The formula uses the web search model selected in the model switcher.

Specifying the model in the formula

1. Open GPT for Excel.
2. In the sidebar, select **GPT functions**.
3. (Optional) Configure the model-specific settings in the sidebar. For Gemini models, make sure web search is enabled.
4. In the workbook, select a cell and write your formula, specifying a web search model directly in the formula, for example the perplexity-fast model:

```
=GPT("Find the Instagram handle for 'Berlin Salsa Festival 2026'.", , ,
"perplexity-fast")
```



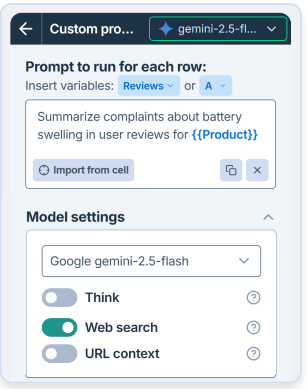
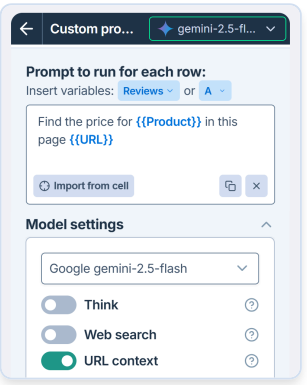
The formula uses the web search model regardless of what's currently selected in the model switcher.

For instructions on how to write formulas with GPT functions, see the GPT functions reference.

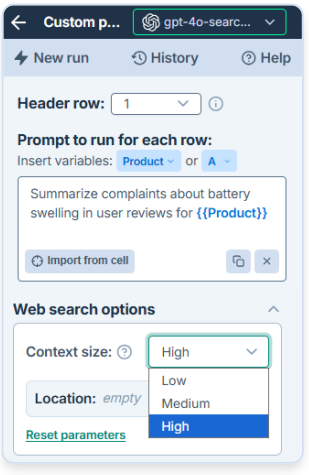
Model-specific settings

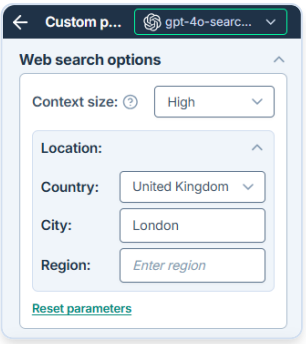
Use the model-specific settings to configure the web search model for your use case.

Google Gemini

Setting	Description	Example
Web search	Let Gemini decide when a Google search can improve the answer. This allows Gemini to provide more accurate answers and cite verifiable sources beyond its built-in knowledge. Enabling Web search incurs an additional search cost.	
URL context	Provide up to 20 URLs (pages, images, PDFs) in your prompt to focus your search. No search cost is incurred, but the content retrieved from URLs counts as input tokens.	

OpenAI Search

Setting	Description	Example
Context size	Choose how much content to retrieve from each source. Larger sizes produce richer and more detailed responses at a higher cost.	

Setting	Description	Example
Location	To refine search results based on geography, specify an approximate user location using country, city, and/or region.	

Use images in GPT function prompts in Excel

Use images in your prompts with GPT for Excel and vision models. These models are multimodal and can process both text and images as input.

> **TIP**
>
> Don't have image URLs yet? Learn how to host images on Google Drive for AI processing.

> **INFO**
>
> Supported image formats: PNG, JPEG, WebP, non-animated GIF

Supported vision models

We recommend using **gpt-5.4** to analyze images, but you can use any vision model supported by GPT for Excel.

Click here for a list of all vision models supported by GPT for Excel ▼

Provider	Model	Built-in	Provider-connected 🔑
Search...	Search...	All	All
OpenAI	gpt-4.1	Yes	Yes
OpenAI	gpt-4.1-mini	Yes	Yes
OpenAI	gpt-4.1-nano	Yes	Yes
OpenAI	gpt-4o	Yes	Yes
OpenAI	gpt-5	Yes	Yes
OpenAI	gpt-5-mini	Yes	Yes
OpenAI	gpt-5-nano	Yes	Yes
OpenAI	gpt-5.1	Yes	Yes
OpenAI	gpt-5.2	Yes	Yes

Provider	Model	Built-in	Provider-connected 🔑
OpenAI	gpt-5.4	Yes	Yes
OpenAI	gpt-5.4-mini	Yes	Yes
OpenAI	gpt-5.4-nano	Yes	Yes
OpenAI	gpt-5.5	Yes	Yes
OpenAI	gpt-5.6-luna	Yes	Yes
OpenAI	gpt-5.6-sol	Yes	Yes
OpenAI	gpt-5.6-terra	Yes	Yes
OpenAI	gpt-6-astra	No	Yes
OpenAI	gpt-6-luna	Yes	Yes
OpenAI	gpt-6.1-sol	Yes	Yes
OpenAI	o1	Yes	Yes
OpenAI	o3	Yes	Yes
OpenAI	o4-mini	Yes	Yes
Anthropic	claude-4.1-opus	No	Yes
Anthropic	claude-4.5-haiku	Yes	Yes
Anthropic	claude-4.5-opus	Yes	Yes
Anthropic	claude-4.5-sonnet	Yes	Yes
Anthropic	claude-4.6-opus	Yes	Yes
Anthropic	claude-4.6-sonnet	Yes	Yes
Anthropic	claude-4.7-opus	Yes	Yes
Anthropic	claude-4.8-opus	Yes	Yes
Anthropic	claude-5-opus	Yes	Yes
Anthropic	claude-5-sonnet	Yes	Yes

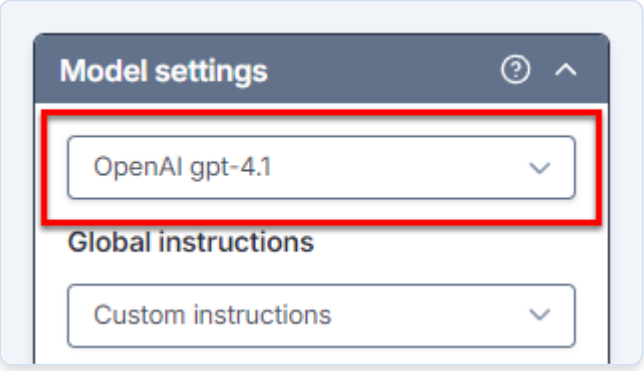
Provider	Model	Built-in	Provider-connected 🔑
Anthropic	claude-5.5-haiku	Yes	Yes
Anthropic	claude-5.5-opus	Yes	Yes
Anthropic	claude-5.5-sonnet	Yes	Yes
Anthropic	claude-fable-5	No	Yes
Anthropic	claude-fable-5.1	No	Yes
Mistral	mistral-medium	No	Yes
xAI	grok-4-1-fast	No	Yes
xAI	grok-4-1-fast-non-reasoning	No	Yes

Prompt images

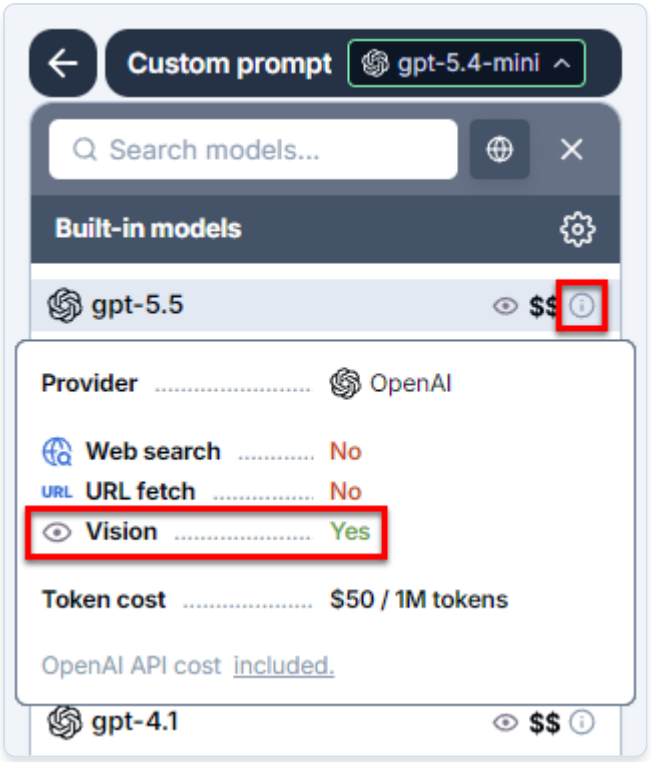
Use the **GPT_VISION** function to describe or analyze images in individual cells.

Using the model switcher

1. Open GPT for Excel.
2. In the sidebar, select **GPT functions**.
3. Expand **Model settings** and open the model switcher.

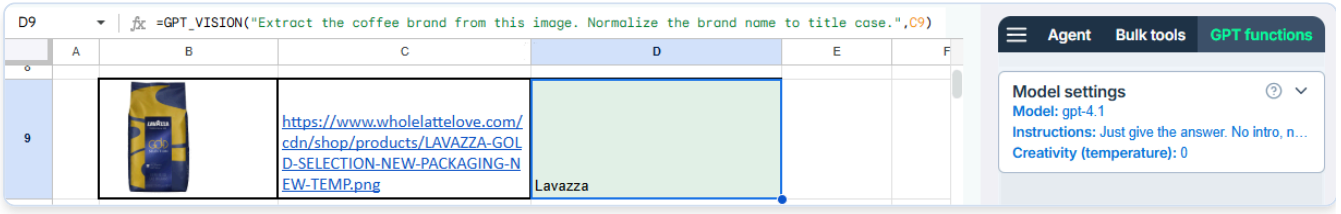


4 Select a vision model.



5 In the workbook, select a cell and write your formula:

```
=GPT_VISION("Extract the coffee brand from this image. Normalize the brand name to title case.",C9)
```



The formula uses the vision model selected in the model switcher.

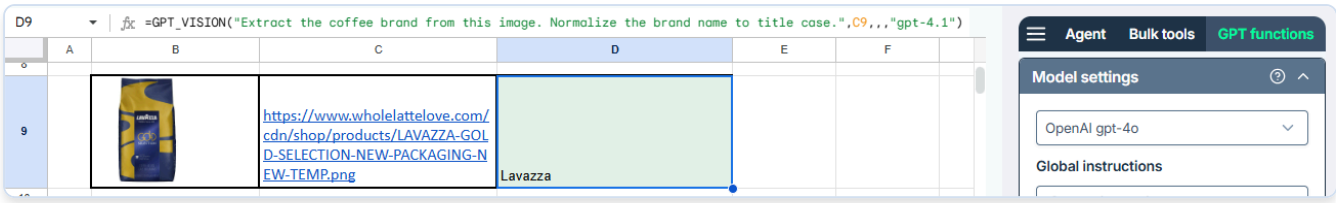
For instructions on how to write formulas with the GPT_VISION function and set the image resolution, see the GPT_VISION reference.

Specifying the model in the formula

1 Open GPT for Excel.

2 In the workbook, select a cell and write your formula, specifying a vision model directly in the formula:

```
=GPT_VISION("Extract the coffee brand from this image. Normalize the brand name to title case.",C9, , , "gpt-4.1")
```



The formula uses the vision model regardless of what's currently selected in the model switcher.

For instructions on how to write formulas with the GPT_VISION function and set the image resolution, see the GPT_VISION reference.

Host images on Google Drive for AI processing

To process an image in GPT for Excel, you need the direct URL of the image and the URL must be publicly accessible.

To host an image on Google Drive so that GPT for Excel can access and process it:

1. Upload the image to Google Drive.
2. Share the image publicly using the **Anyone with the link** option. For more information, see Google Drive Help.

   > **TIP**
   >
   > To prevent other users from commenting or editing the image, set the access role to **Viewer**.

3. Copy the link.
4. Get the image file ID from the link URL, and build the following new URL with the ID:

   ```
   https://drive.google.com/uc?id=<image-file-id>
   ```

   For example:

   ```
   # Original share URL from "Copy link"
   https://drive.google.com/file/d/1dTZrIkAHsV_hKnvL2KjH1pfcOOSpYVux/view

   # Direct image URL
   https://drive.google.com/uc?id=1dTZrIkAHsV_hKnvL2KjH1pfcOOSpYVux
   ```

You now have a direct image URL that you can access and process in GPT for Excel.

Configure GPT function behavior in Excel

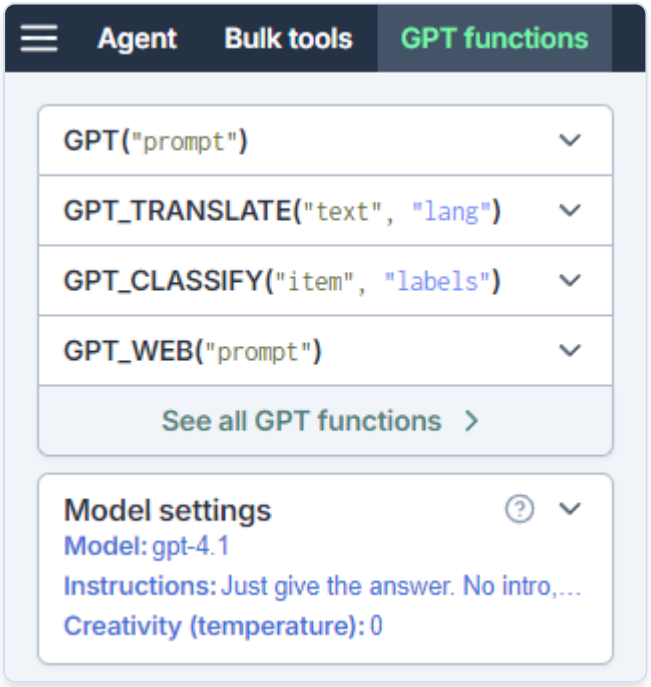
Configure model settings to customize how GPT functions behave in your workbook.

> **INFO**
>
> When cache is enabled, existing GPT formulas will not automatically update to new model settings when you re-run the formulas.
>
> To re-run an existing formula with the new settings:
>
> - Change a parameter in the formula and press Enter.
> - Disable the cache, select the formula, and regenerate its results.

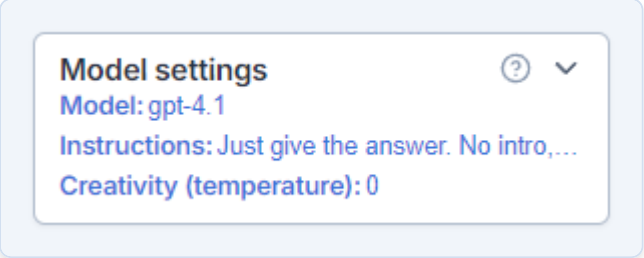
Add global instructions

Provide global instructions to specify preferences or requirements that you'd like the AI to consider when generating responses. The instructions apply to all GPT functions in the current workbook.

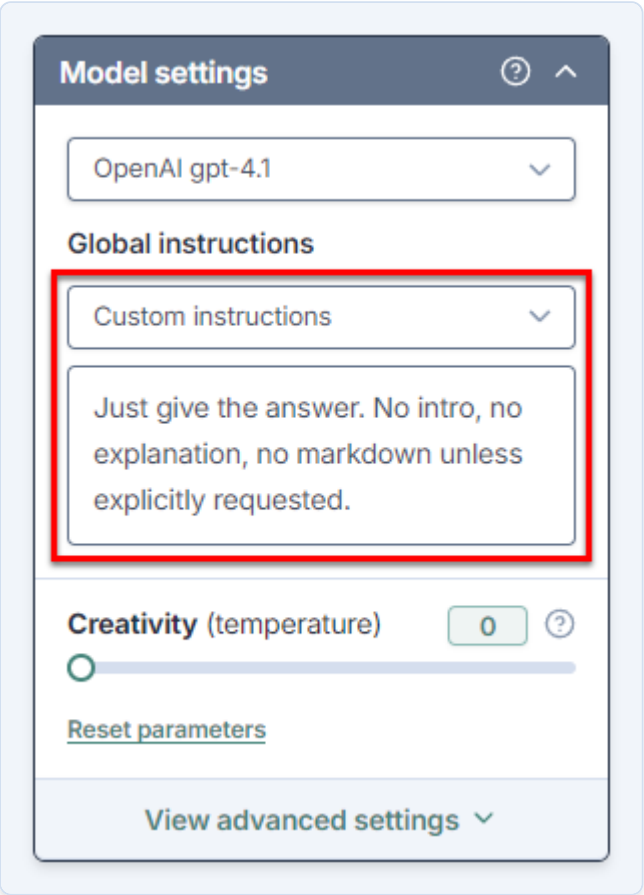
1. In the sidebar, select **GPT functions**.



2 Expand **Model settings**.



3 In the **Global instructions** section, select the type of instructions you'd like to add and edit them if needed.



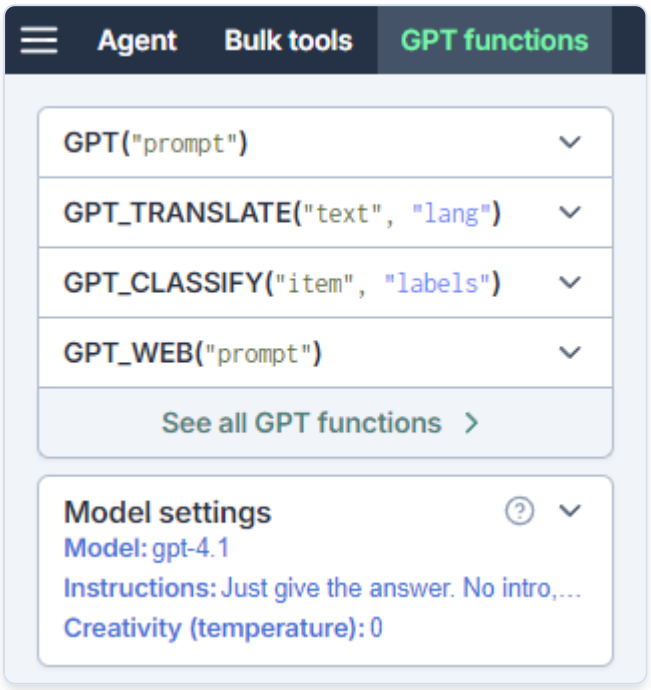
The AI takes the global instructions into account when generating responses for GPT functions in the current workbook.

Set the creativity level

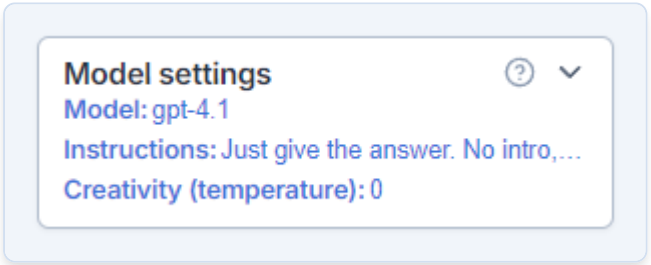
Control how creative the AI is by setting the **creativity level** (also called **temperature**). Lower creativity generates responses that are straightforward and predictable, while higher creativity generates responses that are more random and varied.

> **TIP**
>
> For more information about temperature as used by OpenAI to set the creativity level, see our temperature guide.

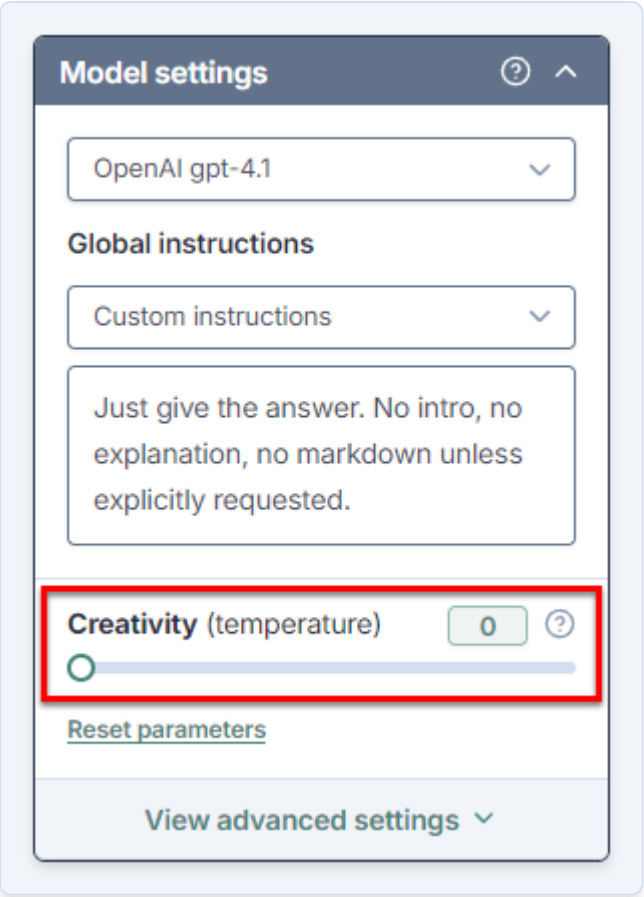
1. In the sidebar, select **GPT functions**.



2. Expand **Model settings**.



3 Use the slider to set **Creativity** between **0** and **1**.



The AI uses the creativity level when generating responses for GPT functions in the current workbook.

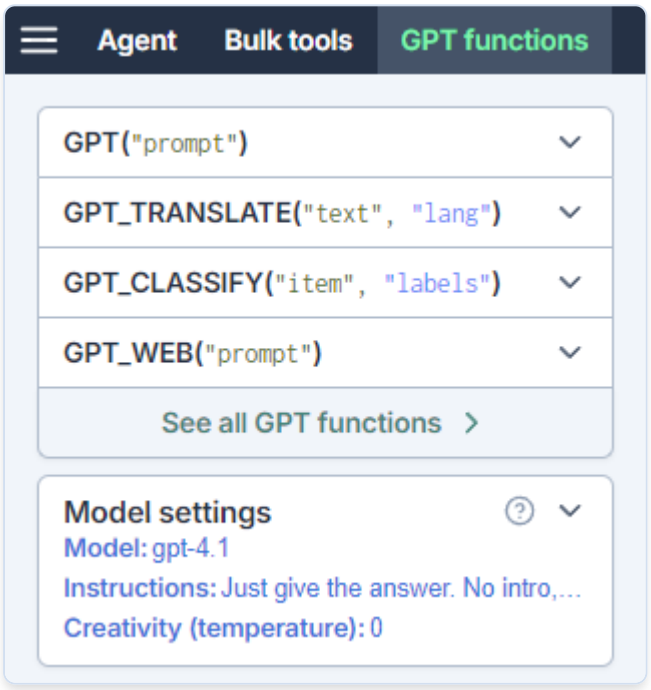
Adjust advanced settings

Adjust the advanced settings to have greater control over generated output for GPT functions in the current workbook. For more information about the settings, see the OpenAI API documentation.

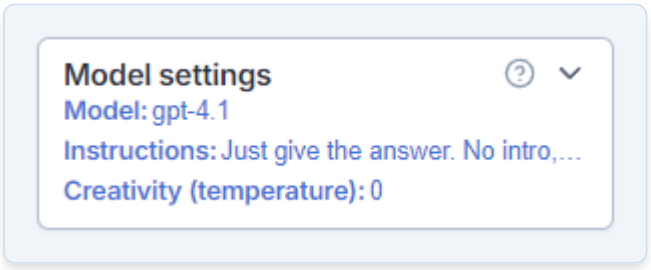
Term	Definition
Token	A token is a unit of text, such as a word, subword, punctuation mark, or symbol. Models split all input and output text into tokens for processing. Model providers use tokens as the basic unit for measuring model usage. Input tokens from your prompt and output tokens from the model's response contribute to your credit consumption. For more information, see Tokens and How credits are consumed.
Top P	Helps adjust creativity level – Lower values result in more focused output, while higher values allow for more creative responses.

Term	Definition
Frequency penalty (not available for all models)	Penalizes tokens based on their frequency in the text so far – Higher values discourage the model from repeating the same tokens too frequently, as the more they appear in the text, the more penalized they get.
Presence penalty (not available for all models)	Penalizes new tokens based on whether they appear in the text so far – Higher values encourage the model to use new tokens that are not penalized.

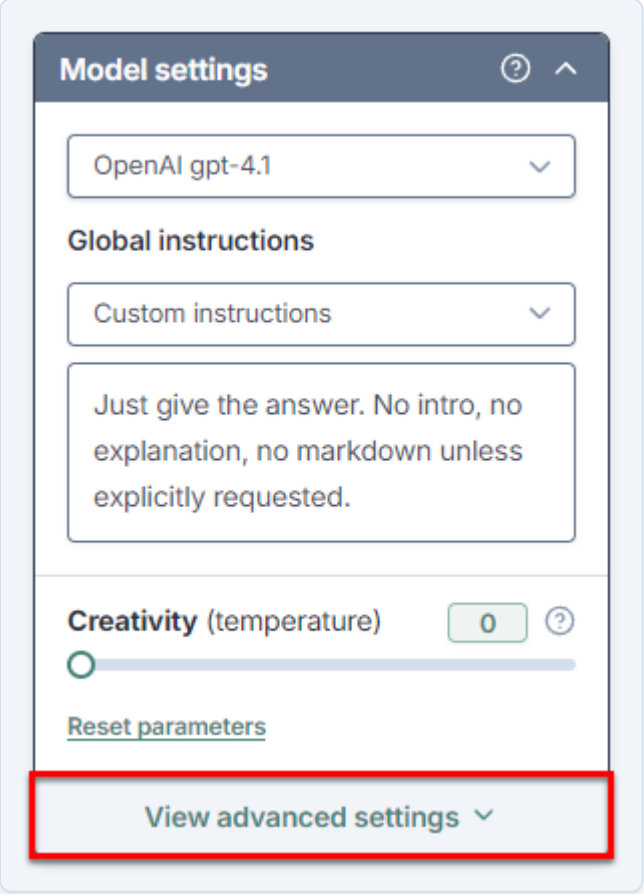
1. In the sidebar, select **GPT functions**.



2. Expand **Model settings**.



3 Expand **View advanced settings**.



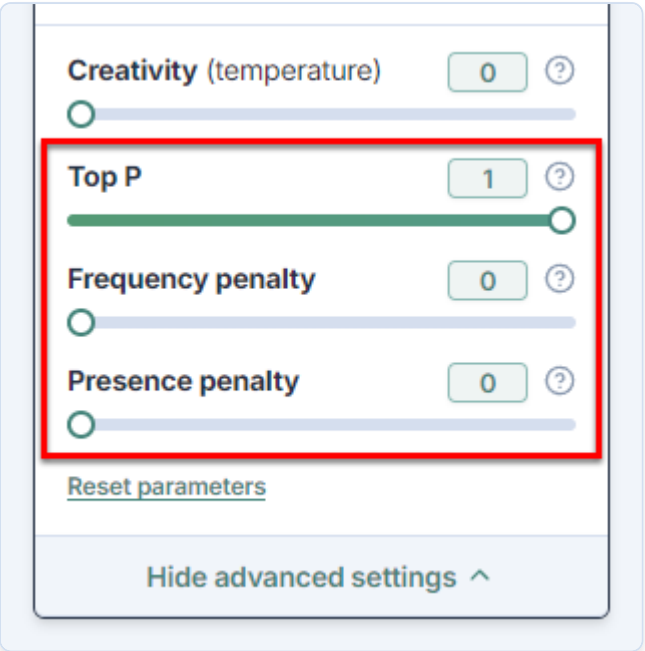
4 Use the corresponding sliders to set the values for:

- **Top P**: Adjust creativity by setting the value between **0** and **1**.

  Set the creativity level to **1** when the value for **Top P** is lower than **1**.

- **Frequency penalty**: Set a penalty between **0** and **2** to apply to tokens each time they are repeated in the result.
- **Presence penalty**: Set a one-time penalty between **0** and **2** for tokens that appear more

than once in the result.



GPT for Excel uses the advanced settings to generate responses for GPT functions in the current workbook (when a model supporting the settings is selected).

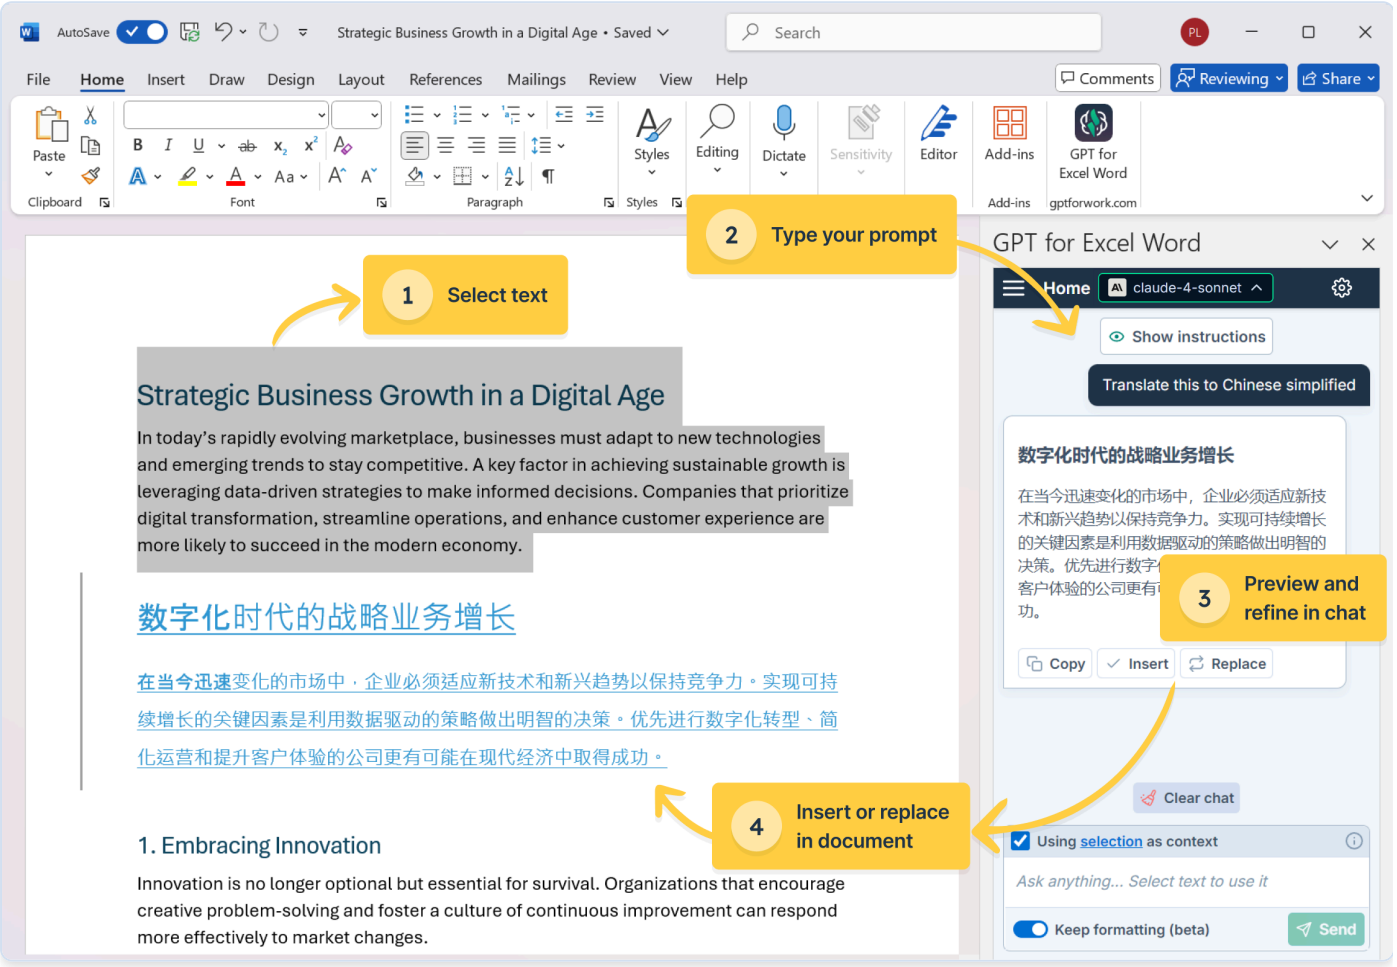
What's next

- Select the model that best fits your needs.
- Run GPT functions.

GPT for Word user guide

Install for Microsoft

GPT for Word is a Microsoft Word add-in that integrates an **AI writing assistant** into the Word user interface. Use plain-language instructions to create and edit content, summarize sections or explain terms, translate text, correct grammar and style, and more.



Use cases

Here are a few common use cases that you can easily handle with GPT for Word:

- **Rewrite or rephrase** content in a different style, tone, or voice.
- **Check and correct** the grammar and spelling in your document.
- **Generate** a document template, or a document draft or outline that you can then flesh out.
- **Translate** content to multiple target languages while observing appropriate localization conventions.

- **Summarize** a document into a few bullet points.
- **Explain** individual words, phrases, and paragraphs in your document.
- **Review** your document for logical structure, clarity and consistency of content, and other aspects that you want to validate.

Features

- Prompt an AI as you would in ChatGPT or another AI chatbot.
- Use a highlighted selection or the whole document as context for your prompts.
- Preserve common formatting elements such as headings, lists, tables, and basic text styles when inserting generated content.
- Preview and refine content in a chat interface before inserting the content into your document.
- Enable Word's **Track changes** mode to keep track of your AI edits.

The following video shows you how to translate selected text and replace the selection with the translation while preserving formatting.

Benefits

- **Integrate AI into your content workflow.** Prompt your favorite AI directly from inside your documents. No more copy-pasting between Word and AI chatbots like ChatGPT.
- **Choose your preferred AI model.** Use models from OpenAI, Anthropic, Google, xAI, or Perplexity. You can also connect to local models with any LLM server that provides an OpenAI-compatible API.
- **Bring an API key to get more features at reduced cost.** You do not need an API key to use GPT for Word, but using a key gives you access to more models, affords you more control and privacy, and you pay a reduced service fee.

What's next

- Install GPT for Word if you have not already.
- Select an AI model that best meets your needs.
- Author text with GPT for Word.

Set up GPT for Word

Install GPT for Word, open the add-in, and sign in. If needed, configure API keys and endpoints to access more models. If you no longer use GPT for Word, remove its access to your Microsoft or Google account or uninstall the add-in.

Install

Install the GPT for Excel Word add-in to use AI in your Word documents.

Open

Open the GPT for Excel Word add-in in Microsoft Word.

Sign in

Use a Microsoft or Google account to sign in to GPT for Word, switch between accounts, and sign out.

Manage models

Connect AI models to Microsoft Word with GPT for Word.

Remove account access

Remove Microsoft or Google account access from the GPT for Excel Word add-in.

Uninstall

Remove the GPT for Excel Word add-in from your Microsoft account.

Install GPT for Word

GPT for Word is part of the **GPT for Excel Word** add-in. One installation covers Microsoft Word and Microsoft Excel, on both desktop and the web; this guide focuses on installing and using the add-in for Microsoft Word.

Prerequisites

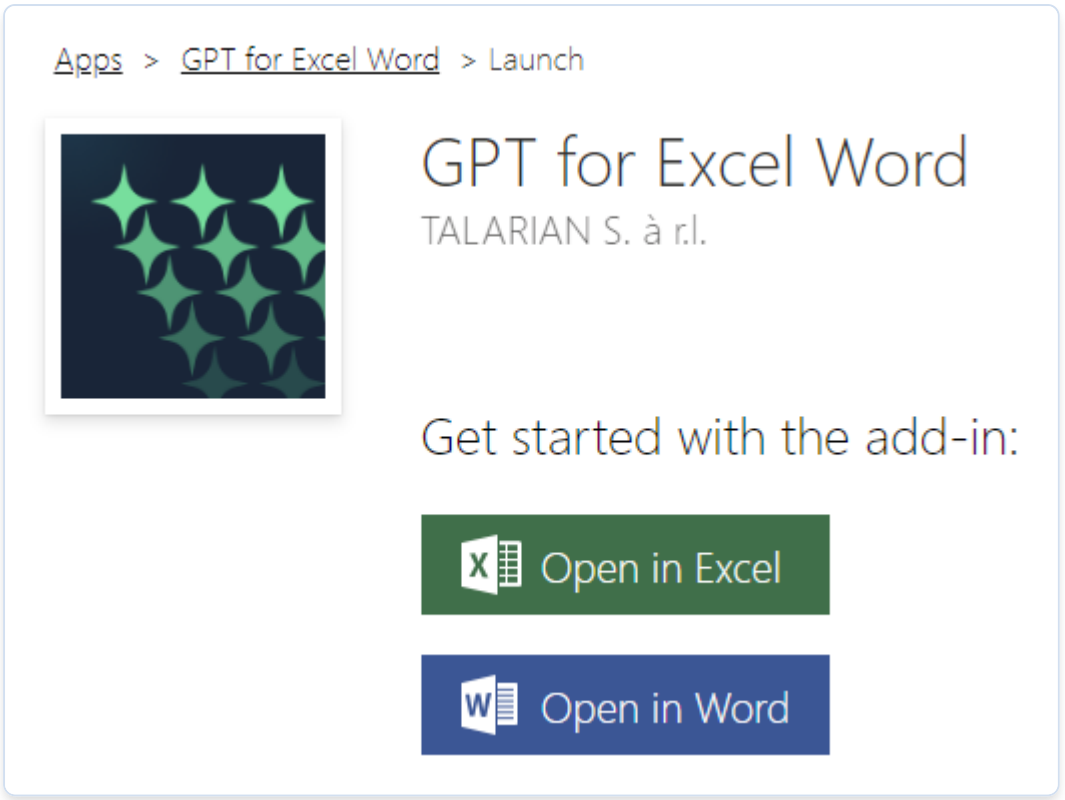
The following table lists the **minimum versions of Microsoft Word** compatible with the GPT for Excel Word add-in on different Microsoft Office platforms.

	Windows	macOS	Web
	**Microsoft 365 subscription:** Version 2205 (Build 15202.10000) **Retail perpetual:** Version 2205 (Build 15202.10000) **Volume-licensed perpetual:** Word 2024: Version 2205 (Build 15202.10000)	Version 16.61 (Build 22040100)	Supported on: • Google Chrome (latest version) • Microsoft Edge (latest version)

Installation

1. Go to the installation page.
2. Click **Open in Word**.

> **INFO**
> The add-in will be installed on your Microsoft account for both Word and Excel.

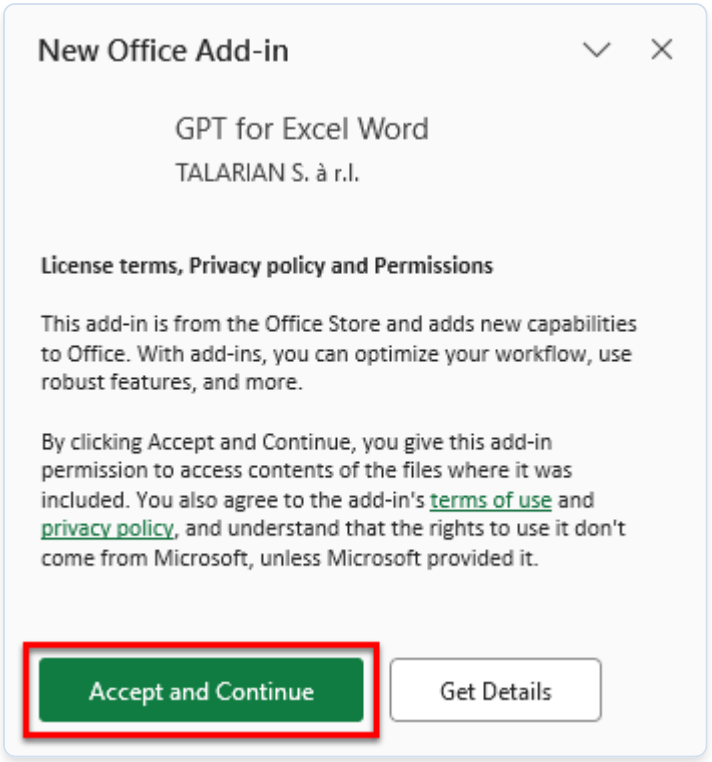


A pop-up appears at the top of the window, requesting your permission to proceed.

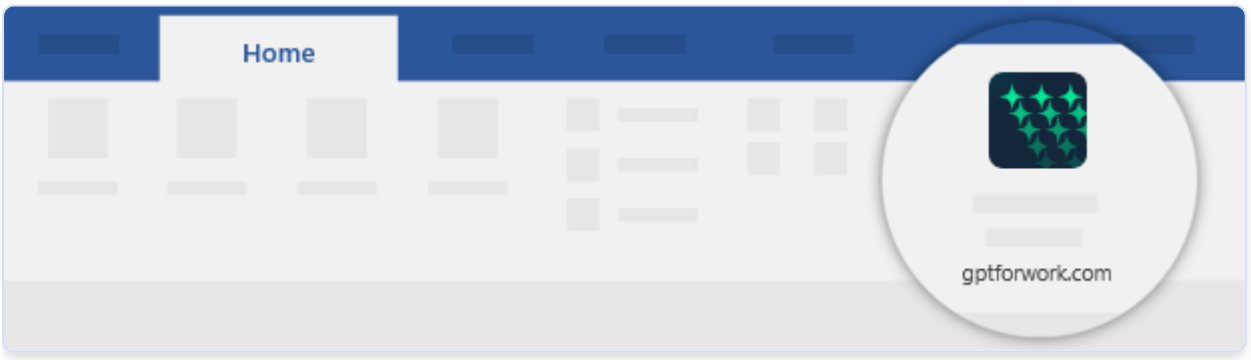
3. Confirm to open the Word desktop app.

> **INFO**
>
> If this method doesn't work, install the add-in from the app:
>
> 1. Open Word.
> 2. On the ribbon, select **Home > Add-ins**.
> 3. Search for "GPT for Excel Word".
> 4. In the search results, click **Add** for **GPT for Excel Word**.

4 Click **Accept and Continue** to accept the permissions.



You have installed GPT for Excel Word on your Microsoft account. The add-in is now available from the **Home** tab in the desktop and web versions of Word. If you don't see it, go to **Home > Add-ins > My add-ins > GPT for Excel Word**.



FOR IT ADMINISTRATORS

If your organization uses a firewall or proxy that restricts network access, you must whitelist the following domains for the add-in to work:

```
*.autosheet.com
*.gptforwork.com
*.microsoft.com
*.microsoftonline.com
*.live.com
*.office.com
```

Having trouble installing?

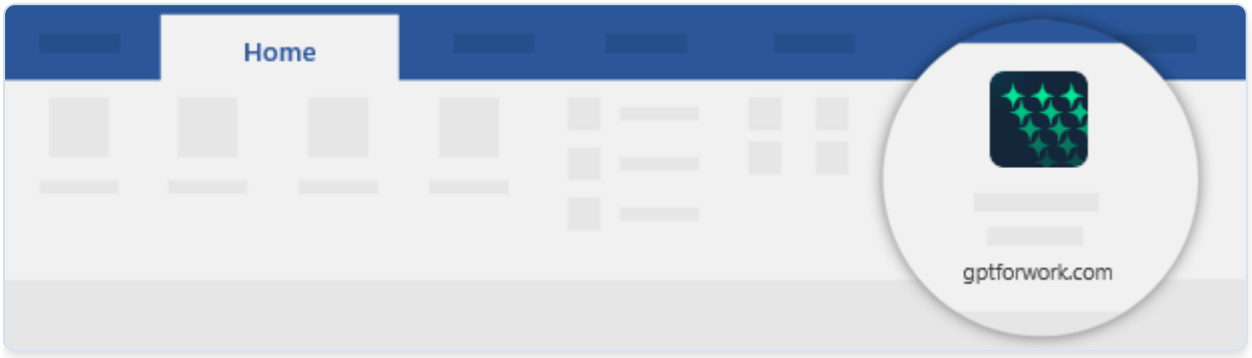
If you run into any issues while installing the add-in, see the troubleshooting guide.

What's next

- Open GPT for Word.
- Learn about GPT for Word.
- Select a model.
- Start authoring text with GPT for Word.

Open GPT for Word

On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

What's next

- Sign in with your Microsoft or Google account.
- If you run into issues, see the troubleshooting guide.

Sign in to GPT for Word

This section explains how to sign in to GPT for Word with your Microsoft or Google account, how to check which account you're using, and how to switch between accounts or sign out.

Sign in with Microsoft or Google

Sign in to GPT for Word with your Microsoft or Google account.

Check which account you're using

How to check which account GPT for Word is using.

Switch to another account

Use a different Microsoft or Google account with GPT for Word.

Sign out

How to sign out of GPT for Word.

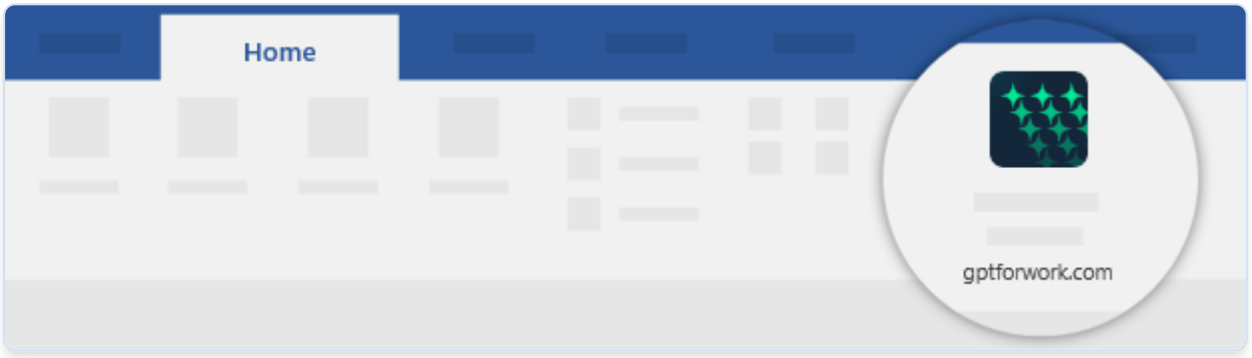
Sign in with Microsoft or Google

You can use GPT for Word with a Microsoft account or a Google account.

If you have multiple accounts, such as a personal account and a work account, you can check which account you're using and switch to another account to ensure you are using the correct account in GPT for Word.

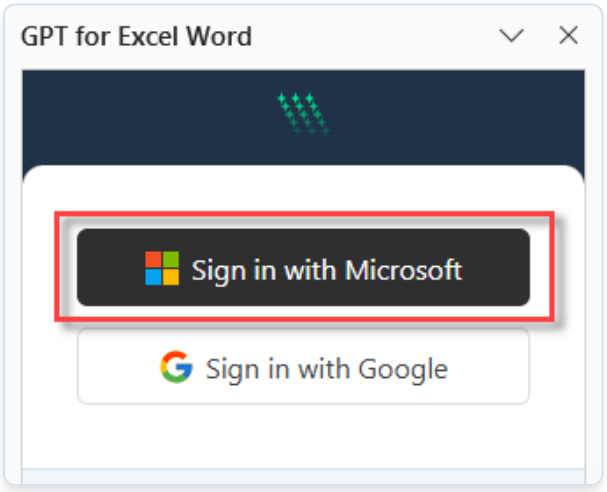
Microsoft

1. Open a Microsoft Word document.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. On the sign-in screen in the sidebar, select **Sign in with Microsoft**.



4. In the dialog that opens, select your Microsoft account and complete the sign-in process.

You are signed in to GPT for Word. The add-in keeps you signed in the next time you open it, until you sign out.

> **INFO**
>
> Signing in to GPT for Word also signs you in to GPT for Excel, since they use the same account session.

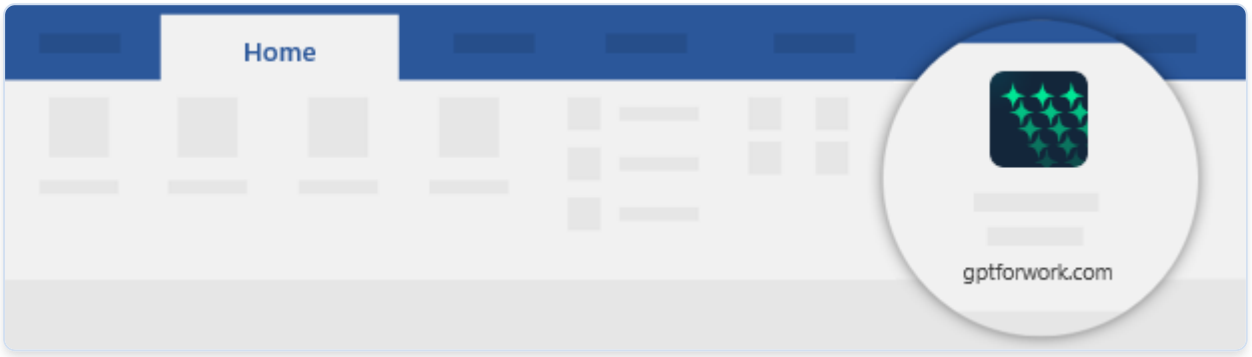
What's next

- Check which account is used by GPT for Word.
- If needed, switch to another account.

Check which account you're using

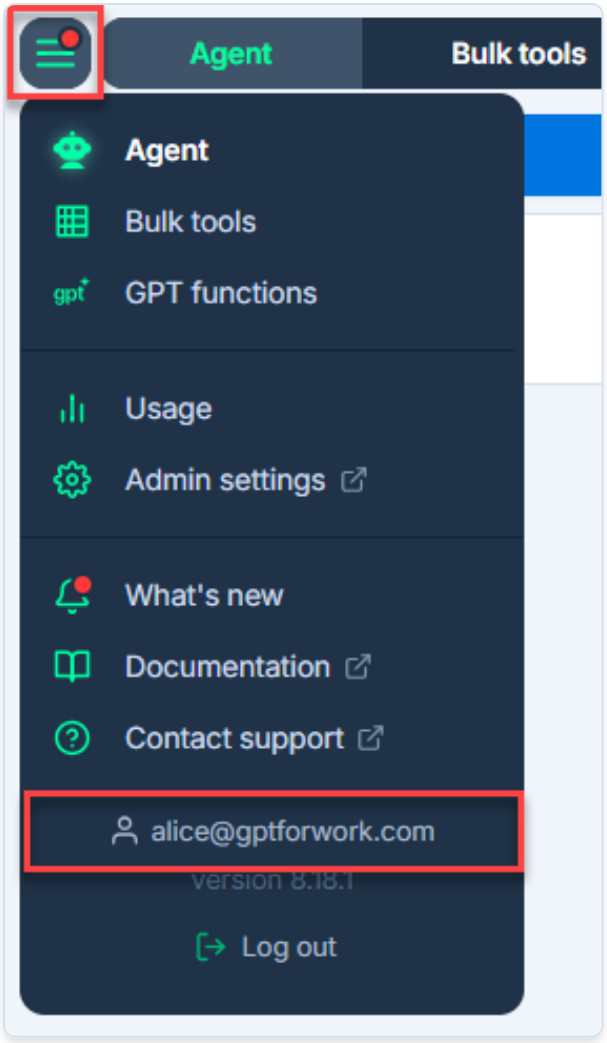
To check which account you are using in GPT for Word:

1. Open a Microsoft Word document.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3 In the sidebar, open the main menu. Your account is displayed at the bottom of the menu.



What's next

- If needed, switch to another account.

Switch to another account

GPT for Word authenticates independently of Microsoft Word: switching accounts in Word does not change the account used by the add-in.

If the account currently used by GPT for Word is not the one you want to use, switch to another account:

1. Sign out of GPT for Word.
2. Sign in again with the Microsoft or Google account you want to use.

> **INFO**
>
> Switching accounts in GPT for Word also switches the account used by GPT for Excel, since they use the same account session.

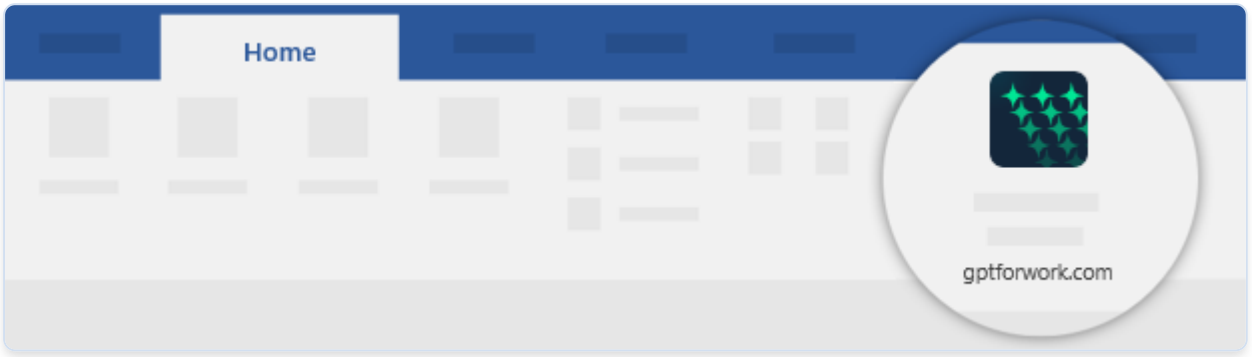
You have switched to another account. You can now use GPT for Word with the new account.

> **TIP**
>
> If you signed in with the wrong account during installation, contact support to request an account transfer.

Sign out of GPT for Word

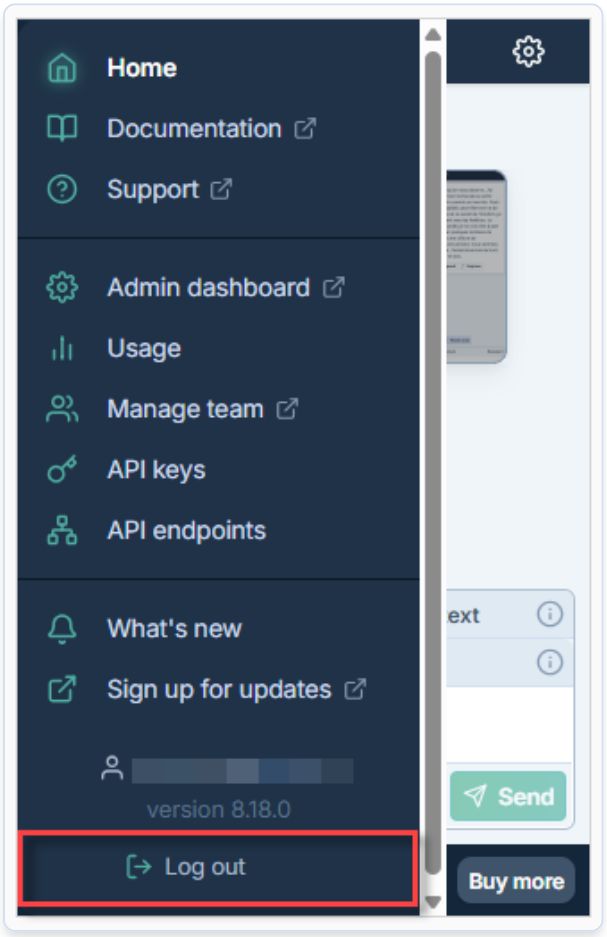
To sign out of GPT for Word:

1. Open a Microsoft Word document.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, open the main menu. Then, select **Log out**.



> **INFO**
>
> Signing out of GPT for Word also signs you out of GPT for Excel, since they use the same account session.

You have signed out of GPT for Word. You'll need to sign in again when you next launch the add-in.

What's next

- Sign in with Microsoft or Google.

Manage models for GPT for Word

Models are the engines that power GPT for Word. GPT for Word supports models from OpenAI, Perplexity, Google, Anthropic, Parallel, DeepSeek, Mistral, and xAI. GPT for Word also supports Azure models and open-source models through any OpenAI-compatible API endpoint.

> **INFO**
>
> As a space owner or admin, you can control which models are available to users in your space through the AI configuration section in the GPT for Work dashboard.

Which model is the right one depends on your specific needs in terms of response quality, speed, privacy, and cost. The following table summarizes the kinds of models available in GPT for Word, from zero setup to more advanced options:

Use	Setup	When
Built-in models	No setup. Use the **default model** or select a different model.	You want to get started quickly without setting up API keys. ✅ Available models
Provider-connected models 🔑	Set an API key for an AI provider.	• You want to benefit from the terms of service you have with an AI provider, such as cost, privacy, and rate limits. • You want to use a model that is not available as a built-in model. • You want more control. ✅ Available models

Use	Setup	When
Endpoint-connected models from cloud platforms	**Azure OpenAI models** Set up an Azure OpenAI API endpoint.	• Your company is already using Azure with a consumption commitment. Using Azure OpenAI counts towards that commitment. • You want to use models that are available through Azure OpenAI. ✅ Available models
	OpenAI-compatible models Set up an OpenAI-compatible platform endpoint.	• You want to use models available on a cloud-based AI platform such as Anyscale or Fireworks AI. • You need a fast and powerful inference infrastructure that scales to your performance needs.
Endpoint-connected models from local servers	**OpenAI-compatible models** Set up an OpenAI-compatible server endpoint.	• You want to use models available on a local LLM server such as LM Studio or Ollama. • You want to run models on a local machine. • You do not want to send or receive any prompt or response data over the internet.

What's next

- Set and manage your API keys.
- Use Azure OpenAI models.
- Use OpenAI-compatible models through API endpoints.
- Use Ollama models.

Manage API keys for GPT for Word

Using API keys offers several benefits:

- **More models.** Access a wider range of models from a bigger pool of AI providers than what is available without an API key.
- **More control.** Monitor and manage AI usage and costs directly on the AI provider's platform.
- **Improved privacy.** We do not log inputs and outputs when you use a connected model. For more information, see our security and privacy FAQ.
- **Reduced costs.** If you're on a pay-as-you-go space, using an API key reduces costs. You pay the model usage fees directly to the AI provider, while paying a reduced service fee for GPT for Work.

Create API keys

> **INFO**
>
> Make sure your keys belong to an organization or project with the following rate limits:
>
> - RPM (requests per minute) ≥ 1,000
> - TPM (tokens per minute) ≥ 1,000,000
>
> The actual limits required depend on your workload requirements.

To create an API key, go to the AI provider's platform and follow their instructions:

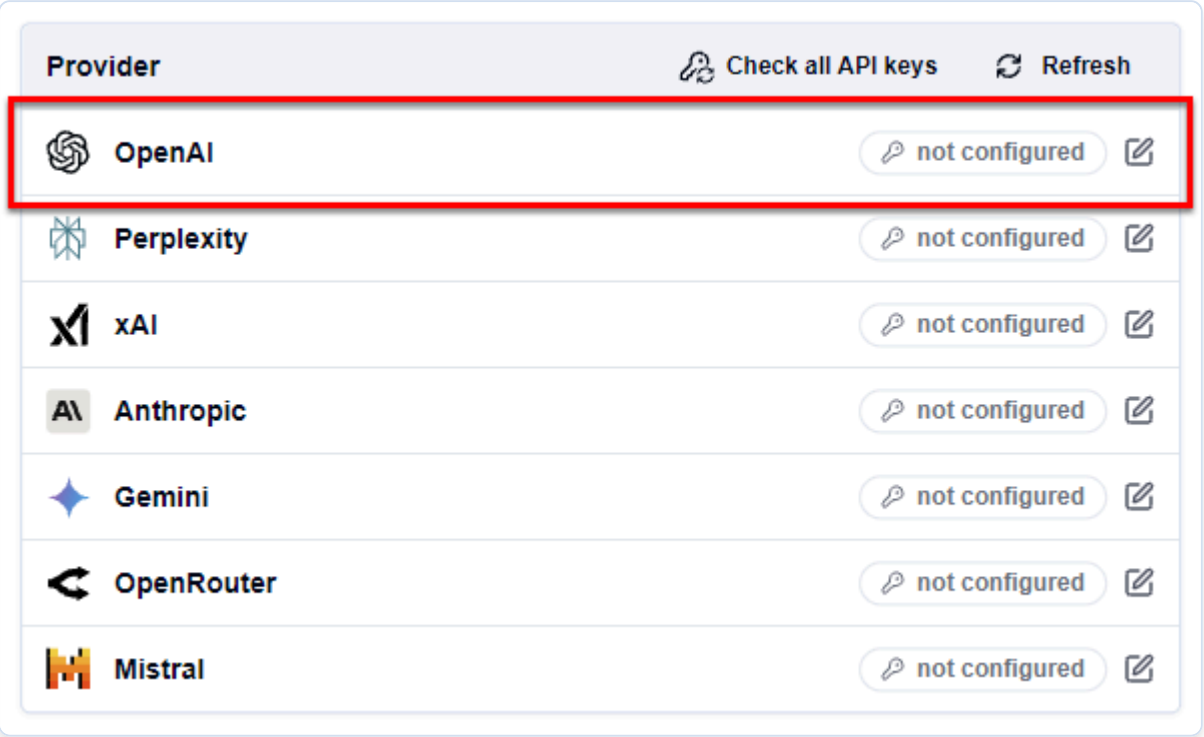
- OpenAI Platform for GPT models and fine-tuned models
- Anthropic Console for Claude models
- DeepSeek Platform for DeepSeek models
- Google AI Studio for Gemini models
- Mistral AI for Mistral models
- Perplexity for Perplexity models
- xAI Cloud Console for xAI models

Set API keys

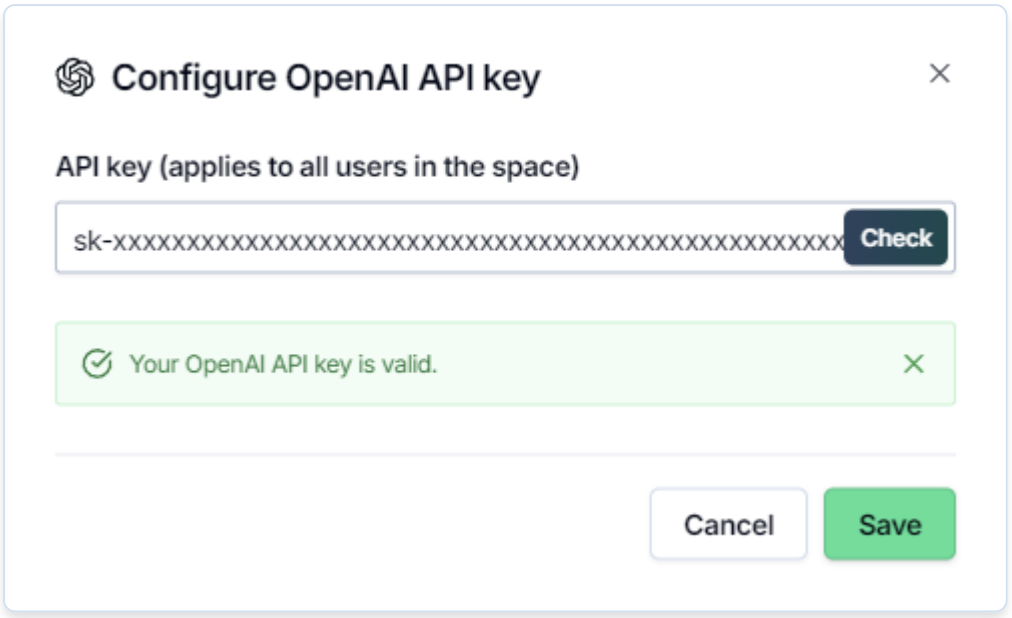
You set API keys for your space in the GPT for Work dashboard. The keys you set for a GPT for Excel and Word space apply to GPT for Word.

> **INFO**
>
> The API keys you set work for all users in your space. Users automatically have access to models associated with the keys, but only the space owner (you) and admins can set and manage the keys. Billing for a key reflects its usage by all users.

1. Create an API key or request an API key from your AI provider platform admin. Learn more.
2. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
3. In the sidebar, select **AI provider keys**.
4. In the AI providers list, click the provider whose API key you want to set.



5. Set the API key:
   1. In the text field, enter the key.
   2. Click **Check** to verify that the key is valid.
   3. Click **Save**.



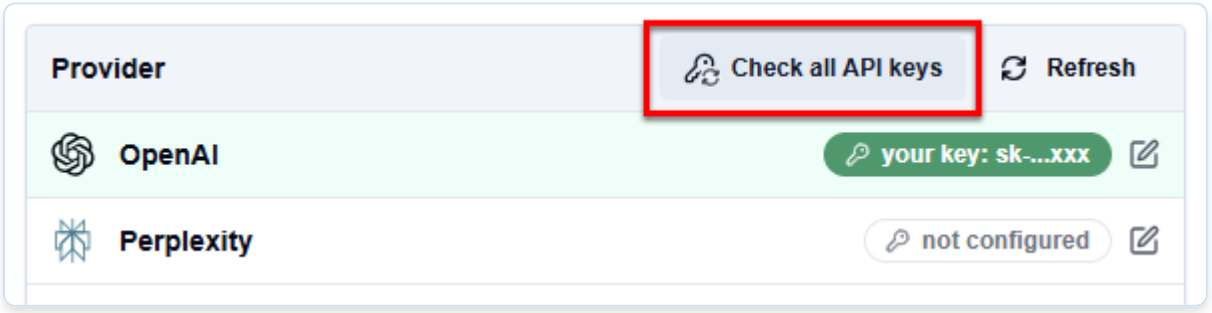
You have set the API key. You can now use any model available with the key in GPT for Word.

Check API keys

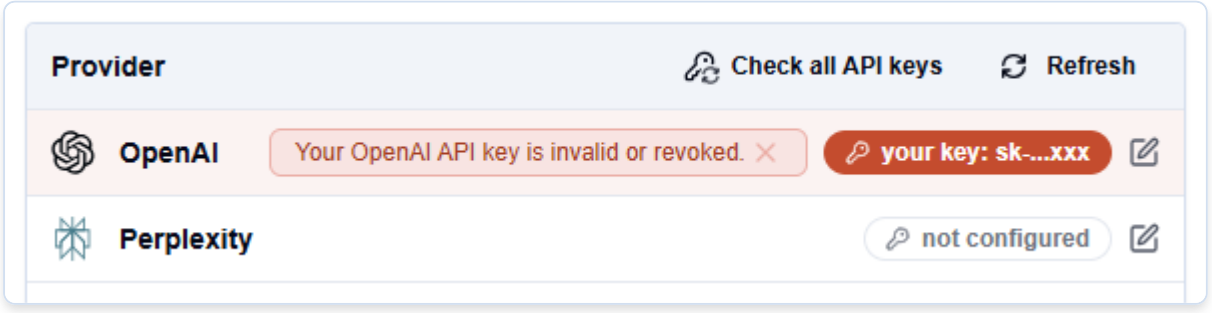
You can verify whether the API keys you have set are valid.

Check all API keys at once

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **AI provider keys**.
3. Click **Check all API keys**.



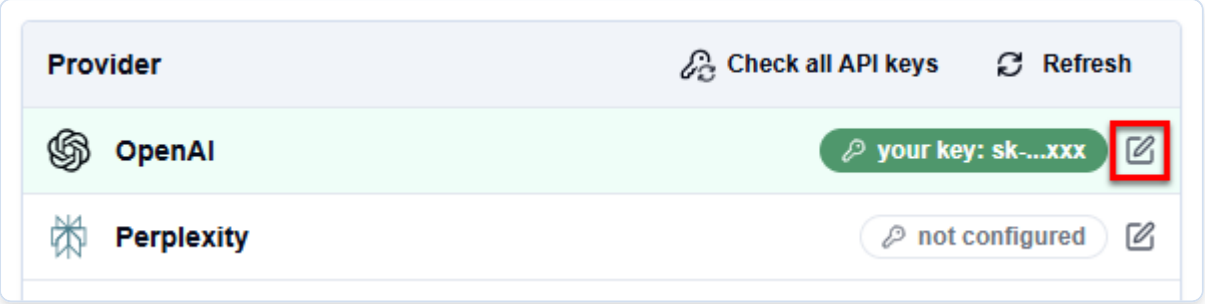
The page displays a warning for each key that is invalid.



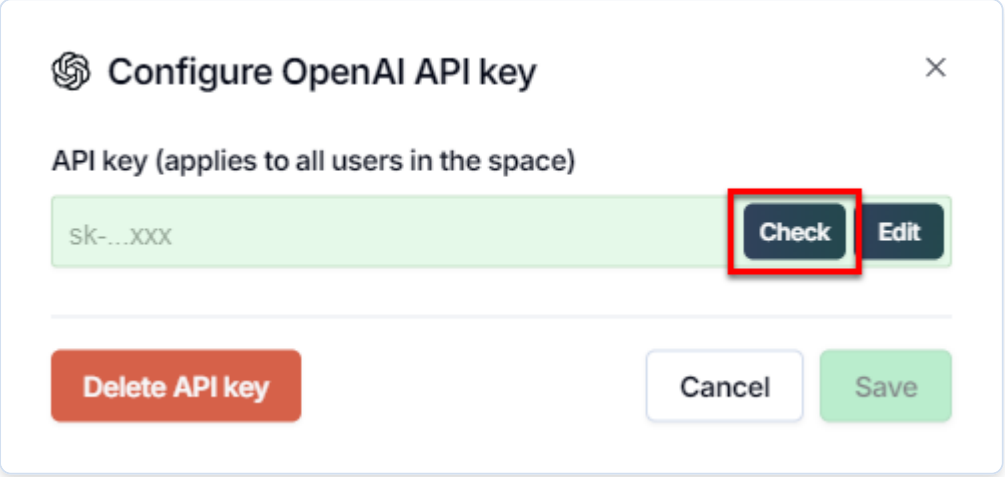
If a key is invalid, create a new key and replace the invalid value with the new one.

Check a specific API key

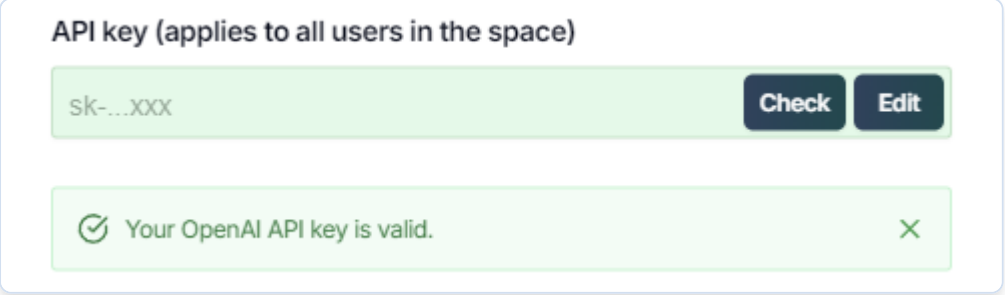
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **AI provider keys**.
3. Click the edit icon for the key you want to check.

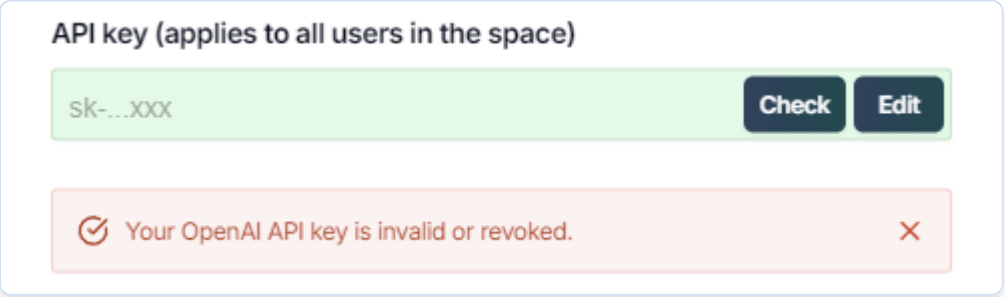


4. Click **Check**.



The UI displays a message indicating whether the key is valid or invalid.





If the key is invalid, create a new key and replace the invalid value with the new one.

Troubleshoot OpenAI API keys

If your API key is not working, check on the OpenAI Platform that the key's permissions are not restricted. The following permissions are required:

Permission	Description
List models	List all available models.
Model capabilities	Use model capabilities, such as chat completion.

What's next

- Check which provider-connected models you can use.
- Select a model.
- Need to use your own AI service endpoint for models? Connect to an OpenAI-compatible endpoint.

Connect to Azure OpenAI from Word

GPT for Word can connect to Azure AI to use OpenAI models. In Azure AI, each model has a dedicated API endpoint that you use to access the model.

Prerequisites

- Your space is on the Business or Enterprise plan or pay-as-you-go pricing.
- You have an Azure account with Azure AI service access.
- You have deployed a model in a project.

To use an Azure OpenAI model:

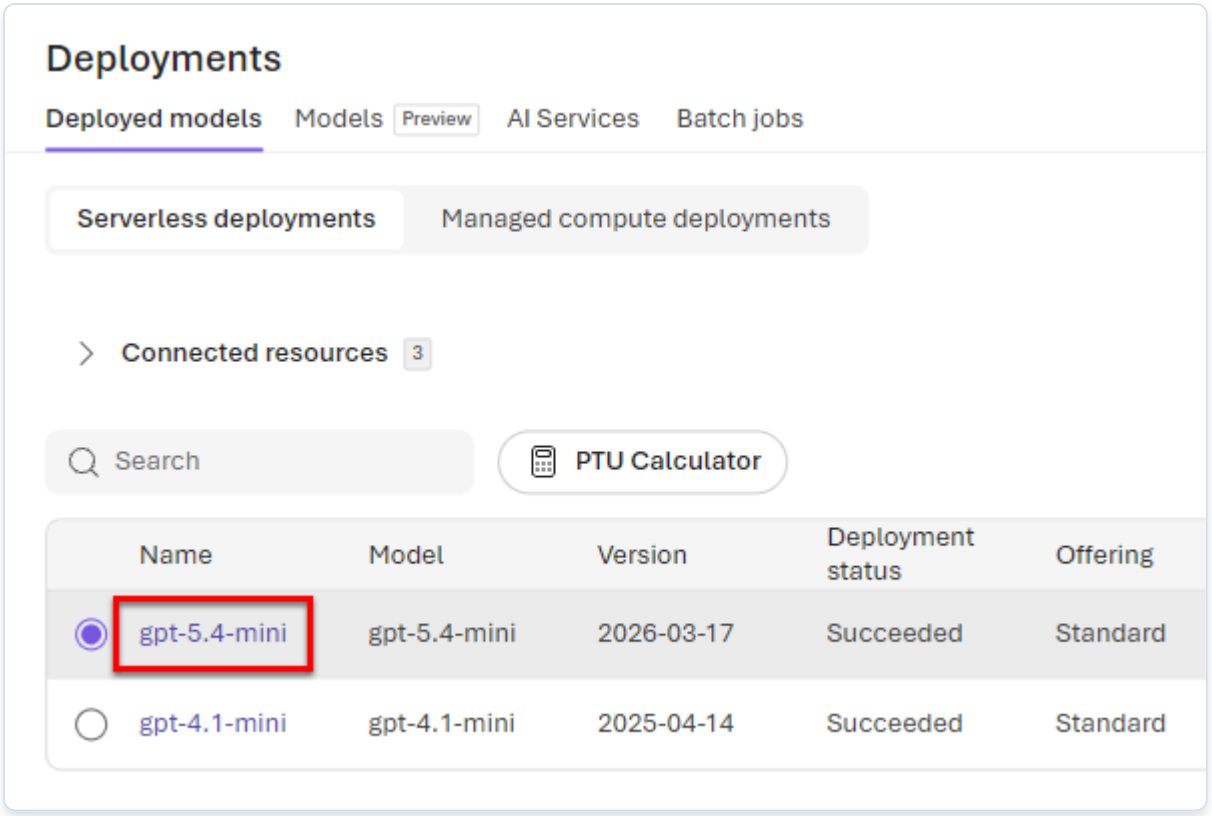
1. Get the Azure OpenAI API endpoint details for the model.
2. Set up the Azure OpenAI API endpoint.
3. Select the Azure OpenAI model.

Get the Azure OpenAI API endpoint details

To connect to Azure AI from GPT for Word, you need the endpoint URL, API key, deployment name, and model name for the Azure OpenAI model you want to use.

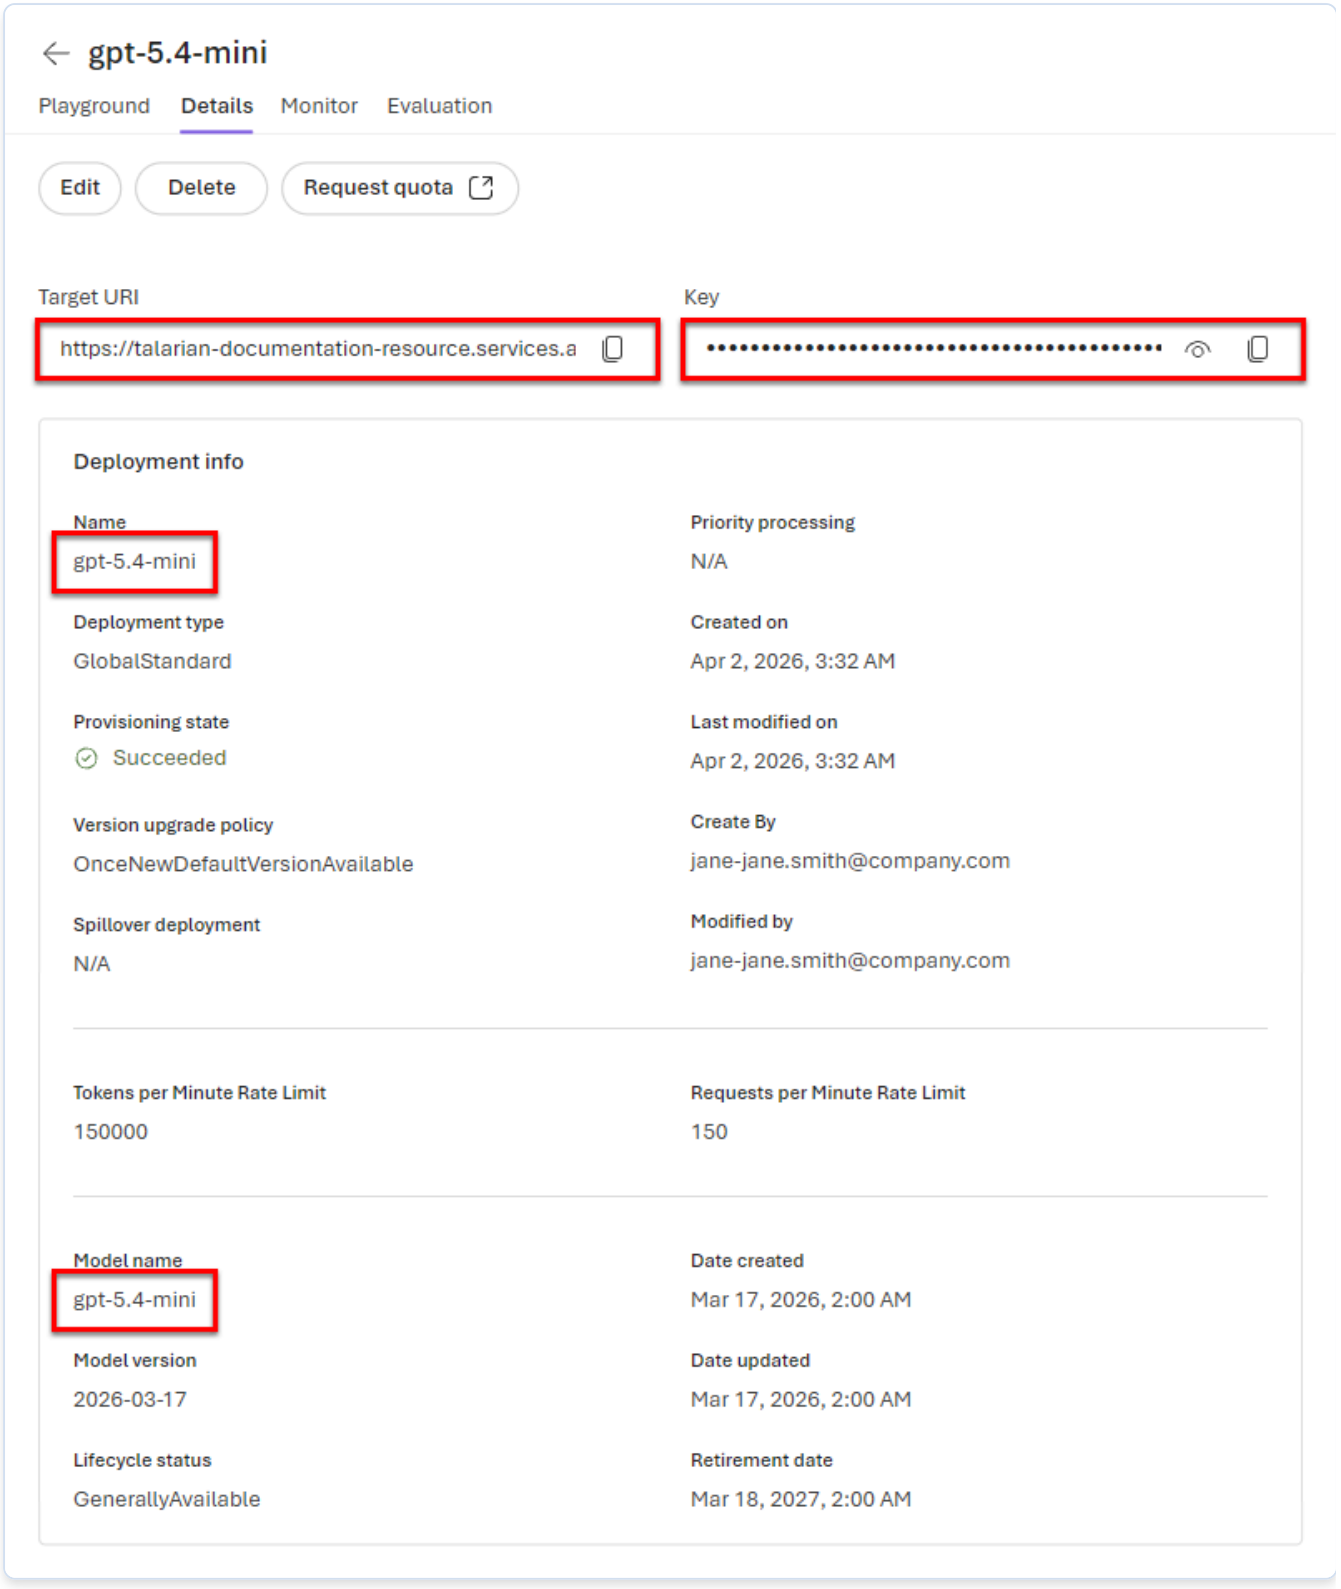
1. Sign in to Microsoft Foundry.
2. Select the project that has the model you want to use.
3. Click **View deployments**.

4 In **Deployed models**, in the models list, click the name of the model you want to use.



5 Select **Details**.

6 In the model details, copy the **Target URI**, **Key**, **Deployment name**, and **Model name**.



Set up the Azure OpenAI API endpoint

1 Sign in to the GPT for Work dashboard with your Microsoft or Google account.

2 In the sidebar, select **Custom API endpoints**.

3 Select the **Bulk models** tab.

4 Click **Azure OpenAI endpoints (Chat Completions API)**.

> **TIP**
>
> If you already have Azure OpenAI API endpoints set up, click **Add an Azure OpenAI API endpoint**.

5 Define the endpoint settings:

- **Endpoint URL**: Enter the endpoint URL for the model using the Azure OpenAI Chat Completions API endpoint format:

  ```
  https://<resource-name>.<domain>/openai/deployments/<deployment-name>/chat/completions?api-version=<api-version>
  ```

  If the URL you copied does not use the above format, rewrite it. For example, let's assume that your deployment name is `gpt-5.4-mini`, the API version is `2025-04-01-preview`, and you copied the following URL:

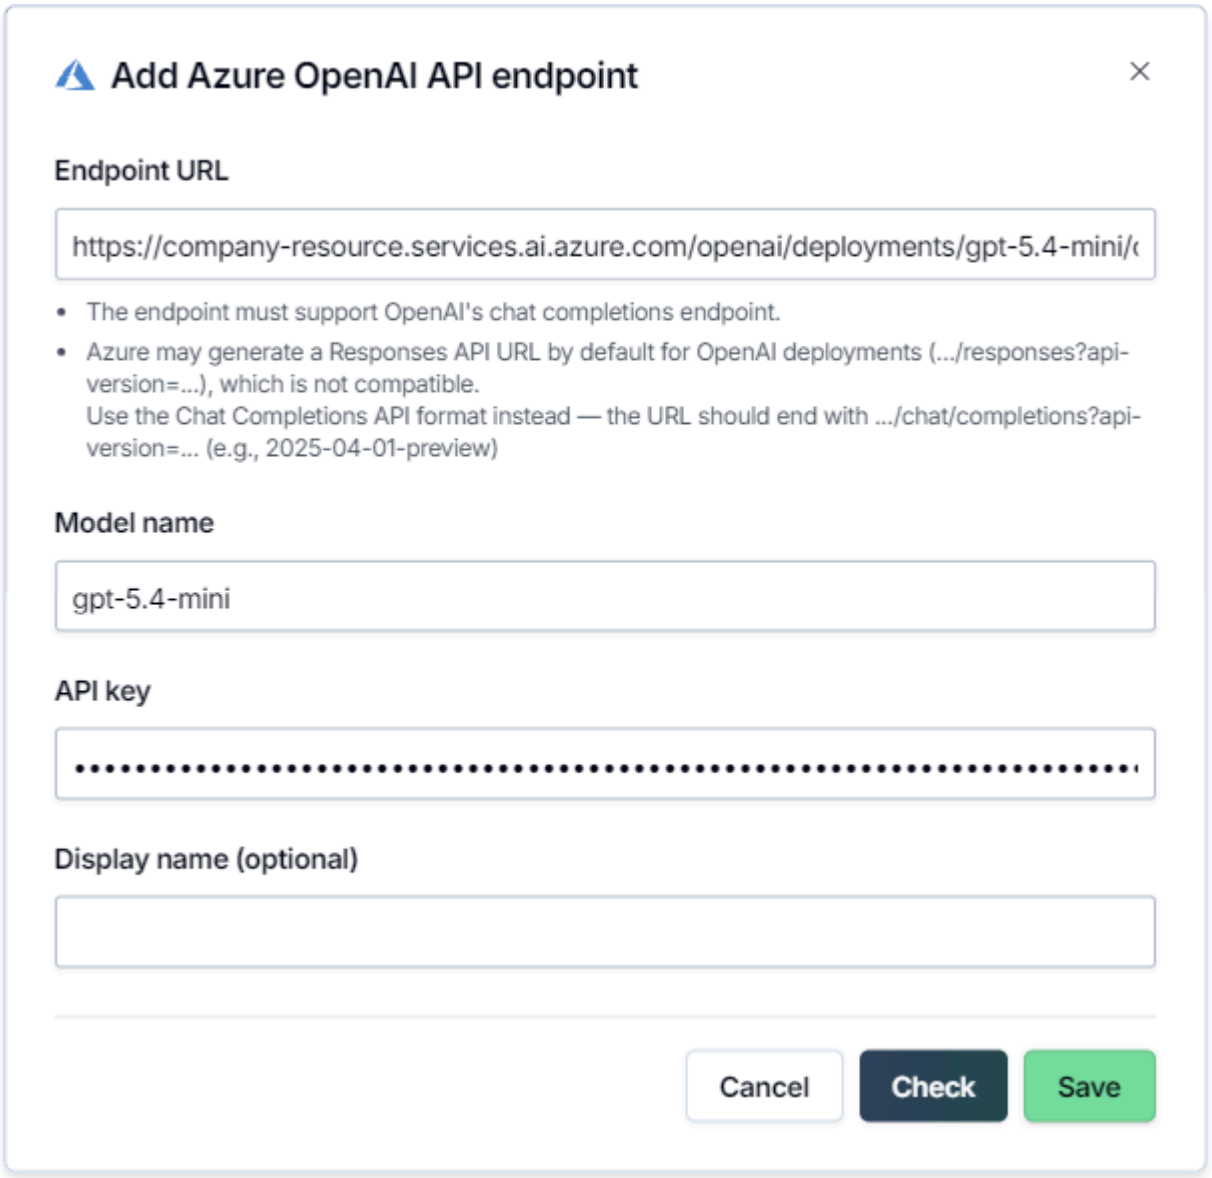
  ```
  https://company-resource.services.ai.azure.com/openai/v1/responses
  ```

  To rewrite the URL into the Chat Completions API format, replace everything after the domain with `/openai/deployments/<deployment-name>/chat/completions?api-version=<api-version>` and replace the placeholders with your deployment name and the API version:

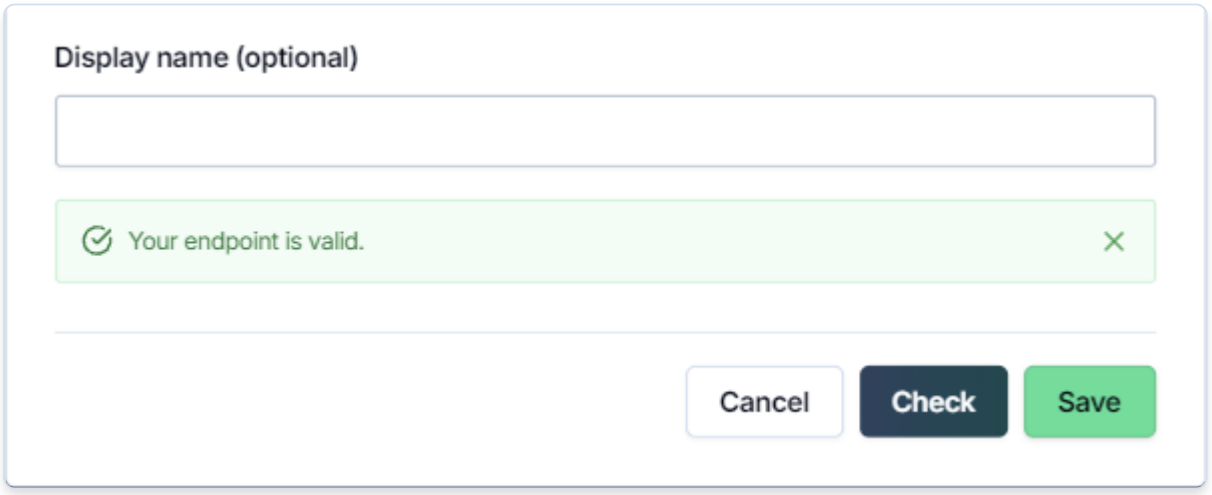
  ```
  https://company-resource.services.ai.azure.com/openai/deployments/gpt-5.4-mini/chat/completions?api-version=2025-04-01-preview
  ```

  > **INFO**
  >
  > The `api-version` is the Azure OpenAI REST API version, which is not shown on the Microsoft Foundry model details page. As a rule, use `2025-04-01-preview` as the value.

- **Model name**: Enter the model name.
- **API key**: Enter the API key for accessing the endpoint.
- **Display name (optional)**: Enter a custom display name for the model. If not set, the model switcher uses **Model name** as the display name.



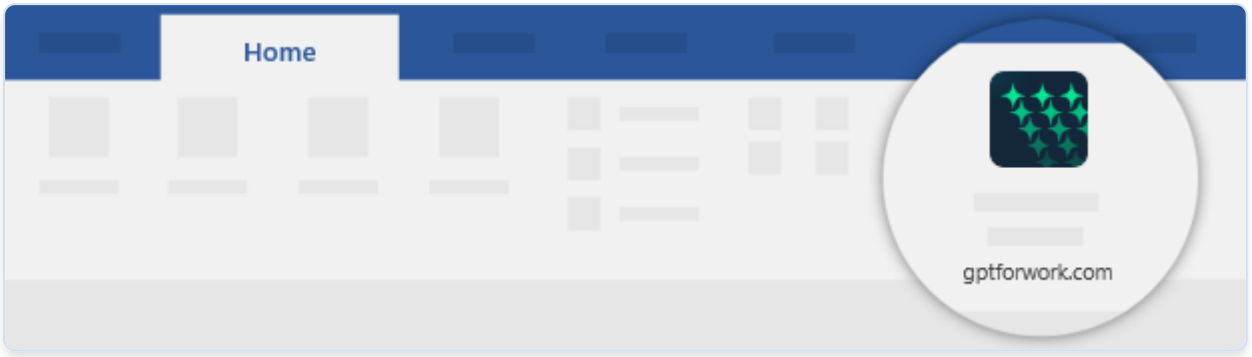
6 Click **Check** to validate the settings and click **Save**.



You have set up the Azure OpenAI API endpoint. You can now use the endpoint model in GPT for Word.

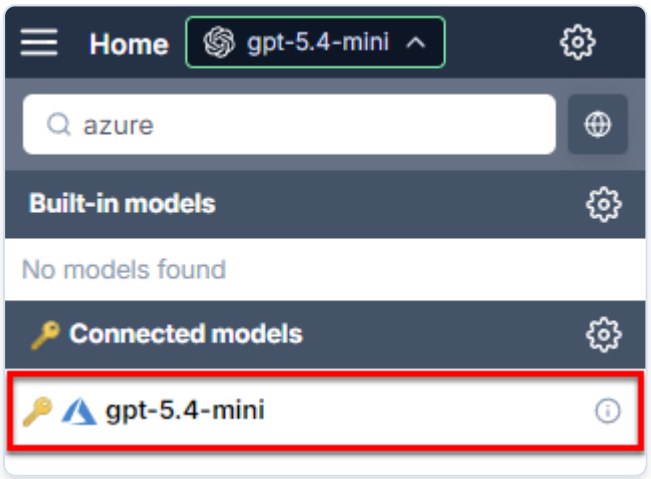
Select the Azure OpenAI model

1. Open a Microsoft Word document.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. Open the model switcher. You can find the Azure OpenAI model under **Connected models**.



4. Select the model.

GPT for Word now uses the selected model to generate responses.

Connect to OpenAI-compatible API endpoints from Word

> **INFO**
>
> This feature is only available if your space is on the Business or Enterprise plan or pay-as-you-go pricing.

GPT for Word can connect to any AI service that provides an OpenAI-compatible API endpoint. You can connect to two main types of services: **cloud-based platforms**, which provide access to models over the internet with no software installation or setup required on your part, and **local servers**, which you run on a local machine, such as your own computer or another computer on a local network.

Here are some popular services that provide an OpenAI-compatible API endpoint:

- **Cloud-based platforms:**
  - Anyscale
  - Fireworks AI
  - Together AI
- **Local servers:**
  - LM Studio
  - LocalAI
  - Ollama

Prerequisites

- You have an endpoint URL and, if needed, an API key for accessing the service.
- The endpoint supports the following OpenAI-compatible API resource paths:
  - `/v1/models`
  - `/v1/chat/completions`

> **INFO**
>
> Depending on the model you use, the model settings defined in GPT for Word (temperature (creativity), top-p, frequency penalty, and presence penalty) may not work as expected. If you get bad results, try tweaking the settings.

To use models from an OpenAI-compatible API endpoint:

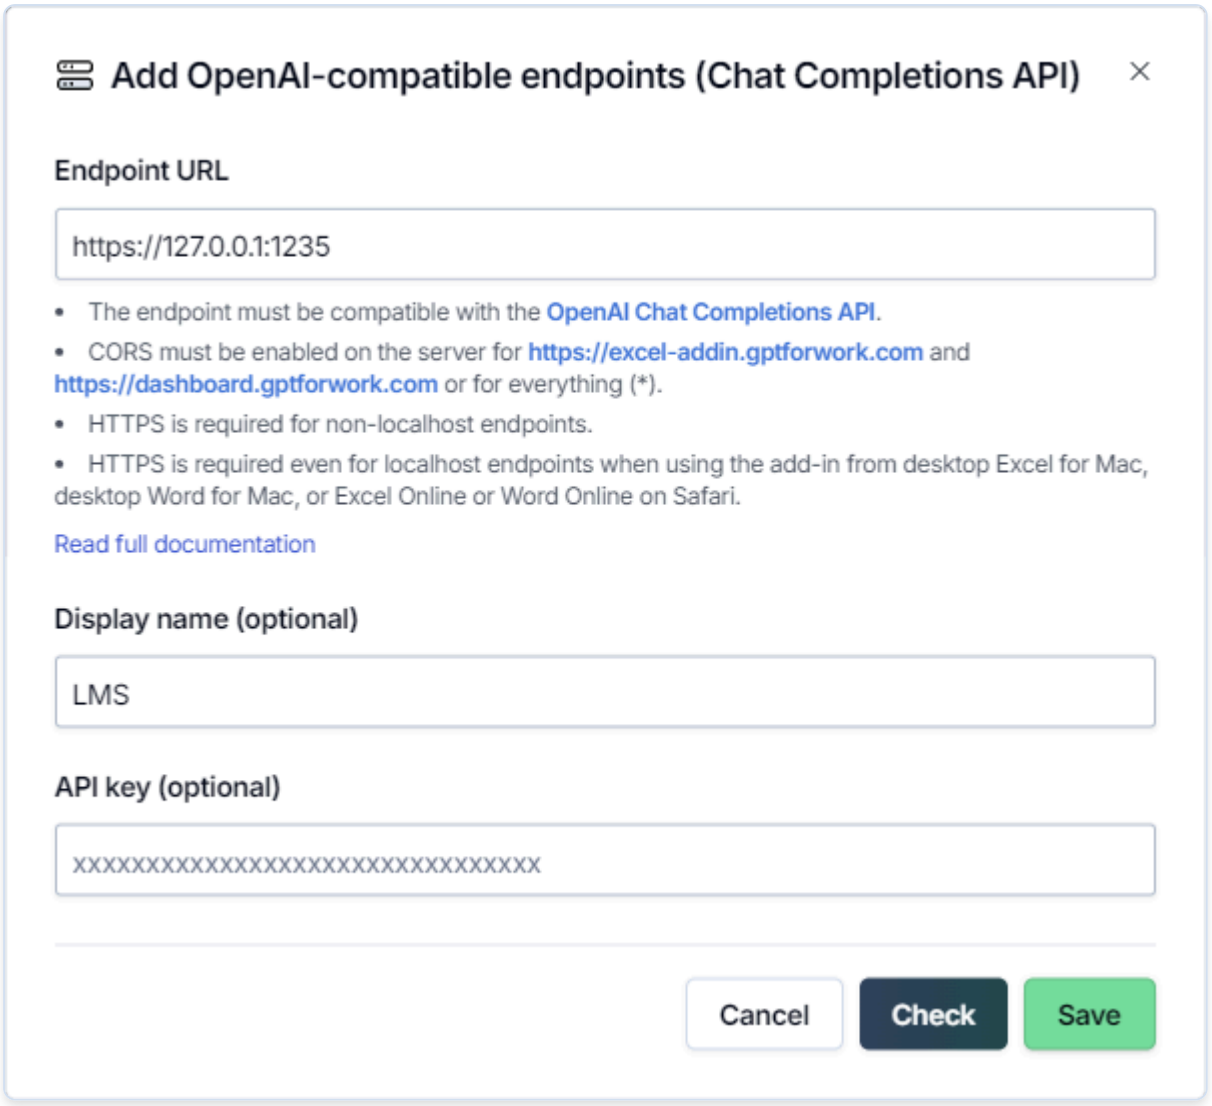
1. Set up the endpoint.
2. Select a model available through the endpoint.

Set up an API endpoint

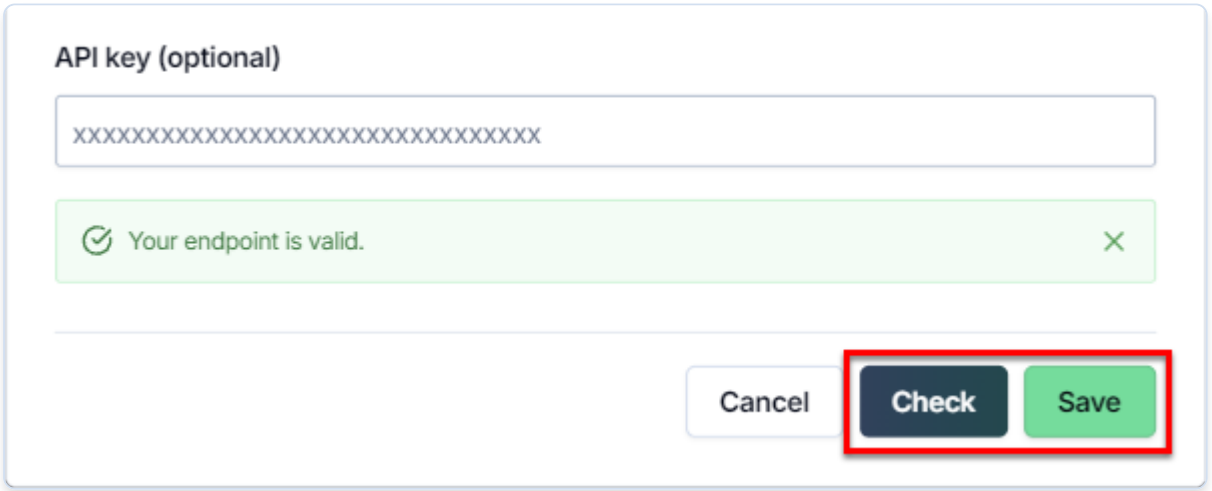
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Bulk models** tab.
4. Click **OpenAI-compatible endpoints (Chat Completions API)**.

   > **TIP**
   >
   > If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the URL of the service that provides the endpoint.
   - **Display name (optional)**: Enter a short name that is prefixed to the model names in the model switcher. By default, the model names are prefixed with "custom/".
   - **API key (optional)**: If required, enter the API key for accessing the service.



6 Click **Check** to validate the settings and click **Save**.

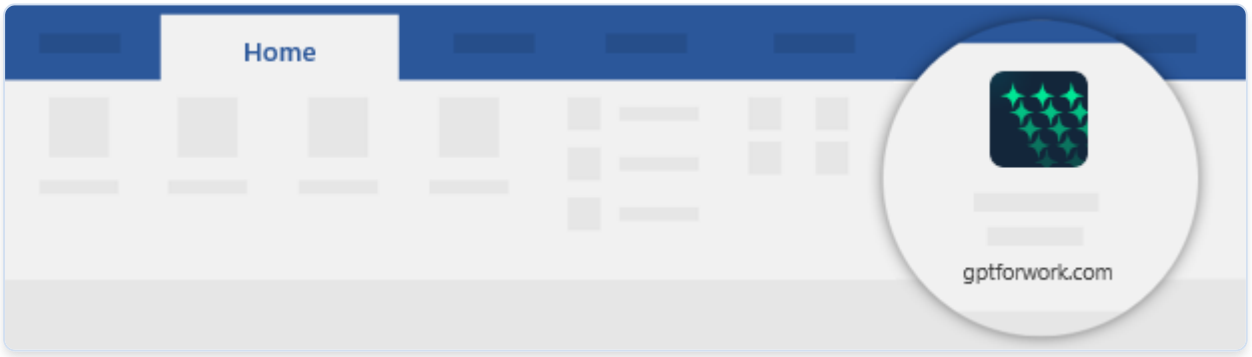


> **TIP**
>
> If you get an error saying that the URL is invalid, make sure that the endpoint supports:
>
> - **Cross-origin resource sharing (CORS).** See the LM Studio setup instructions for macOS or Windows.
> - **HTTPS (needed with Safari and with Microsoft Office for Mac).** See the LM Studio setup instructions for macOS.

You have set up the OpenAI-compatible API endpoint. You can now use all models available through the endpoint in GPT for Word.

Select a model from the API endpoint

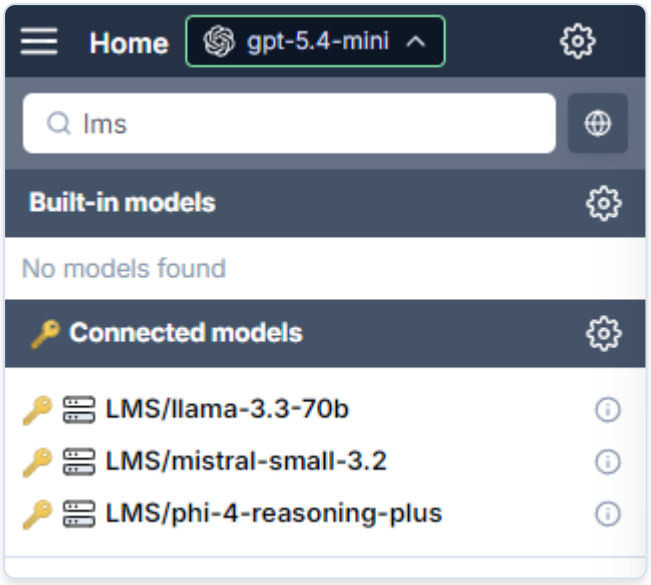
1. Open a Microsoft Word document.
2. On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. Open the model switcher and select a model. You can find the endpoint-connected models

under **Connected models**, at the end of the list.



> **INFO**
>
> The available models depend on the models provided by the service.

GPT for Word now uses the selected model to generate responses.

What's next

- To try out an OpenAI-compatible service with GPT for Word, set up LM Studio on macOS or Windows.

Set up LM Studio on macOS

This guide walks you through the main steps of setting up LM Studio for use in GPT for Word on macOS.

This guide assumes that you use GPT for Word on the same machine that hosts LM Studio.

Prerequisites

- System requirements
- Mac user account with administrator (sudo) privileges
- Installed software:
  - curl (ships with macOS)
  - Homebrew
  - OpenSSL (ships with macOS)

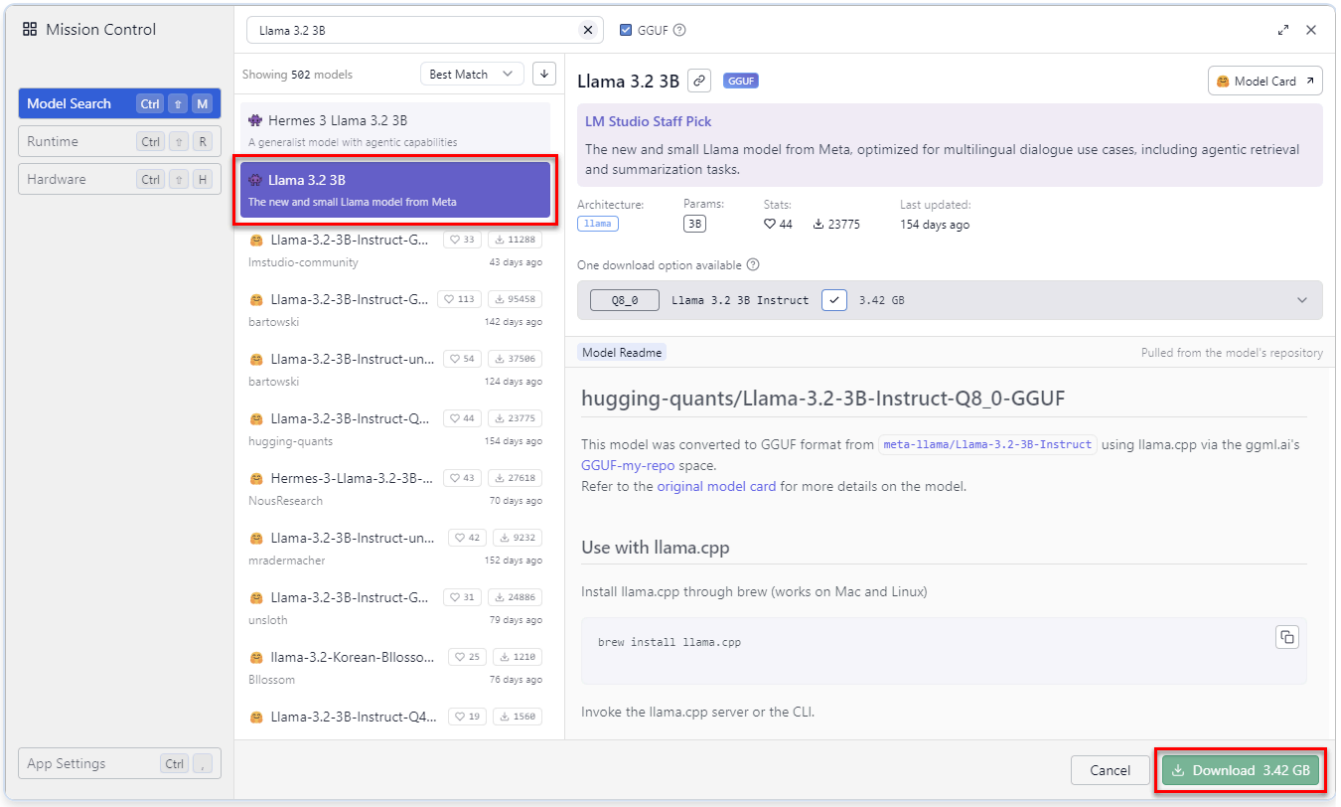
To set up LM Studio on macOS:

1. Install LM Studio and download a model.
2. Start and configure the LM Studio server.
3. (Optional) Enable HTTPS for the LM Studio server.

Install LM Studio and download a model

1. Download and run the macOS installer. Follow the on-screen instructions to complete the installation.
2. Run LM Studio.
3. On the welcome screen, click **Skip**.
4. In LM Studio, in the sidebar, select **Discover**.
5. Find and select the model you want to use, and click **Download**. For example, to get started

with a small model, select Llama 3.2 3B.

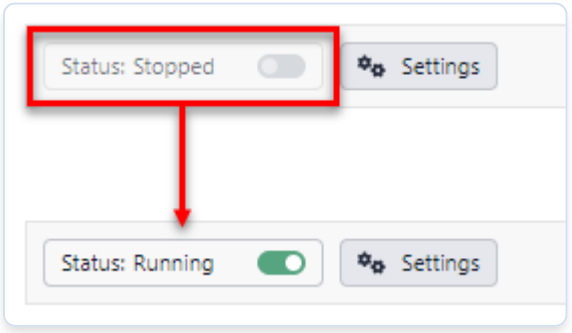


After the download completes, the model is available for prompting in LM Studio.

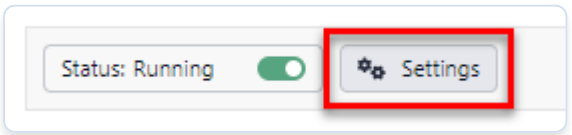
You have installed LM Studio and downloaded your first model. For more information about working with models, see the LM Studio documentation.

Start and configure the LM Studio server

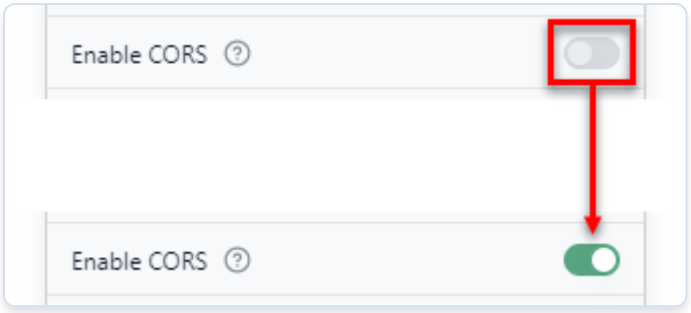
1. In LM Studio, in the sidebar, select **Developer**.
2. Click the **Status** toggle to change the status from **Stopped** to **Running**. You have started the LM Studio server.



3. Click **Settings**.



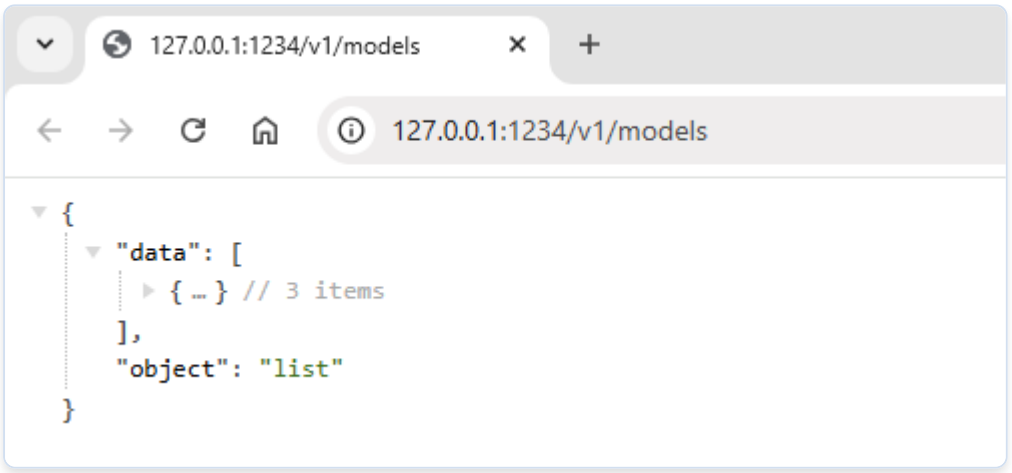
4 Click the **Enable CORS** toggle to enable the setting.



> **INFO**
>
> By default, the LM Studio server only accepts same-origin requests. Since GPT for Word always has a different origin from the LM Studio server, you must enable cross-origin resource sharing (CORS) for the server.

5 To verify that the `/v1/models` endpoint of the LM Studio server works, open http://127.0.0.1:1234/v1/models.

GPT for Word uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `data` property listing all currently installed models:



You have started and configured the LM Studio server.

> **WHAT'S NEXT**
>
> You have completed the minimum setup required to access LM Studio from GPT for Word on the same machine. You can now set `http://127.0.0.1:1234` as the endpoint URL in GPT for Word provided the add-in is not running on Safari or on Microsoft Word for Mac.
>
> If you use GPT for Word on Safari or on Word for Mac, enable HTTPS for the LM Studio server.

Enable HTTPS for the LM Studio server

The LM Studio server uses HTTP to serve models, while GPT for Word runs on HTTPS. By default, therefore, any request from GPT for Word to the server is a mixed-content request (HTTP vs. HTTPS). Modern web browsers do not allow mixed content, as a rule. The only exceptions are mixed-content requests to `http://127.0.0.1` and `http://localhost`, which most browsers treat as safe and therefore allow. Safari, however, blocks mixed-content requests even on the current machine.

To avoid mixed content, you must make the LM Studio server accessible over HTTPS if you use GPT for Word from:

- Word for the web on Safari
- Word for Mac (uses Safari)

You make the LM Studio server accessible over HTTPS by setting up a reverse proxy that hides the server behind an HTTPS interface. This guide uses nginx to set up the reverse proxy.

> **TIP**
>
> You can also use Cloudflare Tunnel, ngrok, or a similar cloud-based tunneling service to set up HTTPS with minimal configuration. Note that with a cloud-based service your traffic will be routed through an external service over the internet.

To set up the reverse proxy:

1. Install nginx:

```
brew install nginx
```

2. Create a self-signed SSL certificate for the LM Studio server:
    1. Change to the nginx configuration directory (varies based on your Mac version):

```
# Mac with an Apple silicon processor
cd /opt/homebrew/etc/nginx

# Mac with an Intel processor
cd /usr/local/etc/nginx
```

    2. Generate the certificate:

```
openssl req \
  -x509 \
  -newkey rsa:2048 \
  -nodes \
```

```
  -sha256 \
  -days 365 \
  -keyout lm-studio.key \
  -out lm-studio.crt \
  -subj '/CN=127.0.0.1' \
  -extensions extensions \
  -config <(printf "
[dn]\nCN=127.0.0.1\n[req]\ndistinguished_name=dn\n[extensions]\nsubjectAl
```

The command generates two files in the current directory:

- `lm-studio.crt`: Public self-signed certificate
- `lm-studio.key`: Private key used to sign the certificate

3. Add the certificate to trusted root certificates:

```
sudo security add-trusted-cert -d -r trustRoot -k
/Library/Keychains/System.keychain ./lm-studio.crt
```

3. Create a site configuration for the LM Studio server:

1. Open the `nginx.conf` file in your preferred text editor. The file is in the nginx configuration directory:

```
# Mac with an Apple silicon processor
/opt/homebrew/etc/nginx/nginx.conf

# Mac with an Intel processor
/usr/local/etc/nginx/nginx.conf
```

2. Replace the content of the file with the following configuration:

```
events {}

http {
    server {
        listen       127.0.0.1:1235 ssl;
        server_name  127.0.0.1;

        http2 on;

        ssl_certificate      lm-studio.crt;
        ssl_certificate_key  lm-studio.key;

        location / {
            proxy_pass  http://127.0.0.1:1234;

            # Keep connection to backend alive.
```

```
            proxy_http_version  1.1;
            proxy_set_header    Connection '';
        }
    }
}
```

The configuration uses port `1235` for HTTPS. Requests to `https://127.0.0.1:1235` are forwarded to the LM Studio server running at `http://127.0.0.1:1234`.

3. Save and close the file.
4. Test the nginx configuration:

```
sudo nginx -t
```

Response if the configuration is valid:

```
nginx: the configuration file /opt/homebrew/etc/nginx/nginx.conf syntax
is ok
nginx: configuration file /opt/homebrew/etc/nginx/nginx.conf test is
successful
```

5. Start nginx:

```
sudo nginx
```

4. Verify that the `/v1/models` endpoint of the LM Studio server works over the HTTPS connection:

```
curl -k https://127.0.0.1:1235/v1/models
```

GPT for Word uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `data` property listing all currently installed models:

```
{
  "data": [
    {
      ...
    },
  ],
  "object": "list"
}
```

You have enabled HTTPS for the LM Studio server.

WHAT'S NEXT

You have completed the setup required to access LM Studio from GPT for Word on the same machine. You can now set `https://127.0.0.1:1235` as the endpoint URL in GPT for Word, including when the add-in is running on Safari or on Microsoft Word for Mac.

Set up LM Studio on Windows

This guide walks you through the main steps of setting up LM Studio for use in GPT for Word on Windows.

This guide assumes that you use GPT for Word on the same machine that hosts LM Studio.

Prerequisites

- System requirements
- Windows user account with administrator privileges
- Installed software:
  - curl (ships with Windows)
  - PowerShell 7

To set up LM Studio on Windows:

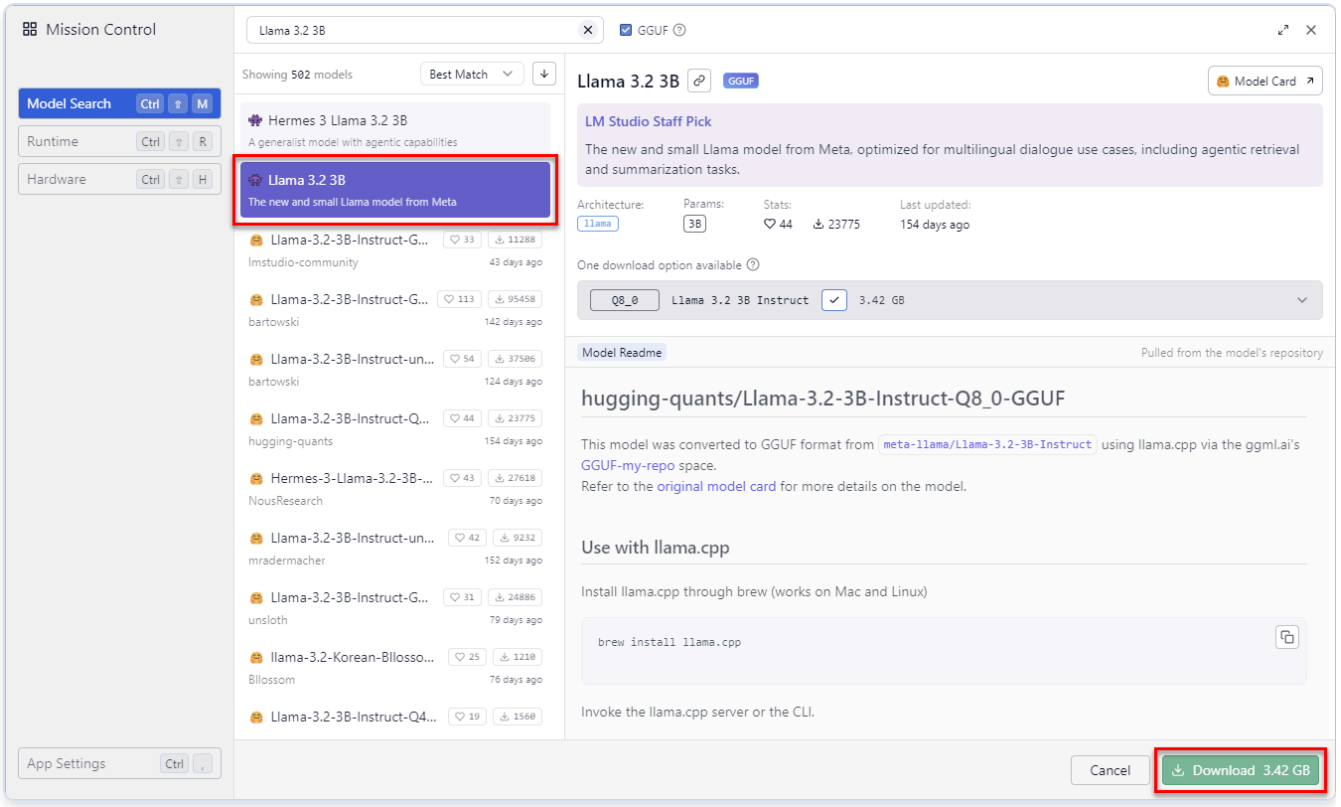
1. Install LM Studio and download a model.
2. Start and configure the LM Studio server.

Install LM Studio and download a model

1. Download and run the Windows installer. Follow the on-screen instructions to complete the installation.

   The installer sets LM Studio to start automatically as a background service on system boot.
2. Run LM Studio.
3. On the welcome screen, click **Skip**.
4. In LM Studio, in the sidebar, select **Discover**.
5. Find and select the model you want to use, and click **Download**. For example, to get started

with a small model, select Llama 3.2 3B.

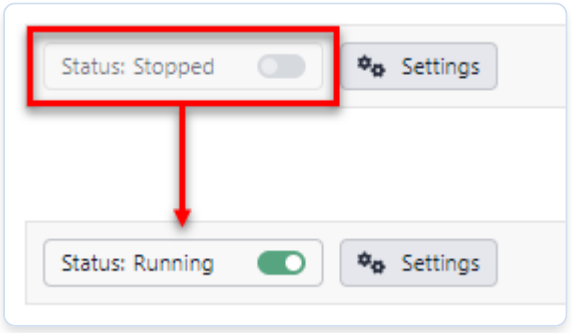


After the download completes, the model is available for prompting in LM Studio.

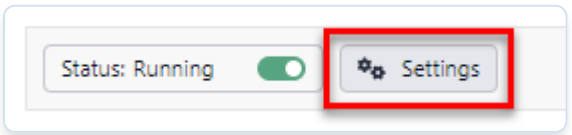
You have installed LM Studio and downloaded your first model. For more information about working with models, see the LM Studio documentation.

Start and configure the LM Studio server

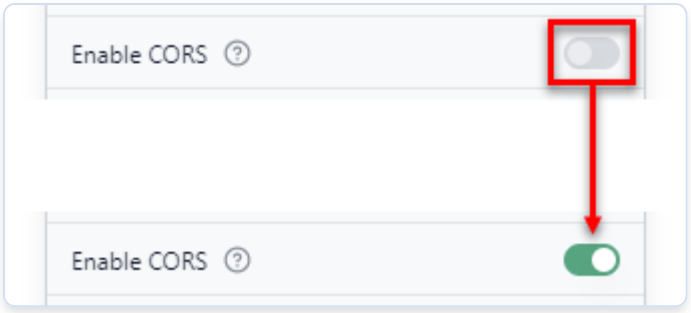
1. In LM Studio, in the sidebar, select **Developer**.
2. Click the **Status** toggle to change the status from **Stopped** to **Running**. You have started the LM Studio server.



3. Click **Settings**.



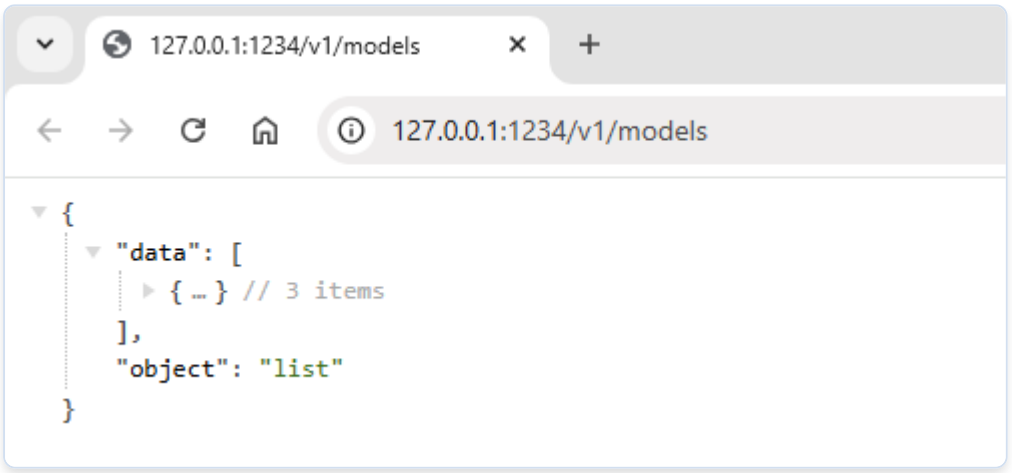
4 Click the **Enable CORS** toggle to enable the setting.



> **INFO**
>
> By default, the LM Studio server only accepts same-origin requests. Since GPT for Word always has a different origin from the LM Studio server, you must enable cross-origin resource sharing (CORS) for the server.

5 To verify that the `/v1/models` endpoint of the LM Studio server works, open http://127.0.0.1:1234/v1/models.

GPT for Word uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `data` property listing all currently installed models:



You have started and configured the LM Studio server.

> **WHAT'S NEXT**
>
> You have completed the setup required to access LM Studio from GPT for Word on the same machine. You can now set `http://127.0.0.1:1234` as the endpoint URL in GPT for Word.

Connect to Ollama from Word

> **INFO**
>
> This feature is only available if your space is on the Business or Enterprise plan or pay-as-you-go pricing.

GPT for Word can connect to an Ollama server to use locally running open-source models. Ollama is an open-source platform that allows you to easily install, run, and serve models from a local machine, keeping your prompting entirely offline. You can use any model from the Ollama library.

Prerequisites

- You have an Ollama server running with one or more models. For instructions on how to set up the server, see:
  - Set up Ollama on macOS
  - Set up Ollama on Windows

> **INFO**
>
> Depending on the model you use, the model settings defined in GPT for Word (temperature (creativity), top-p, frequency penalty, and presence penalty) may not work as expected. If you get bad results, try tweaking the settings.

To use models from an Ollama server:

1. Set up the API endpoint for the Ollama server.
2. Select a model available on the Ollama server.

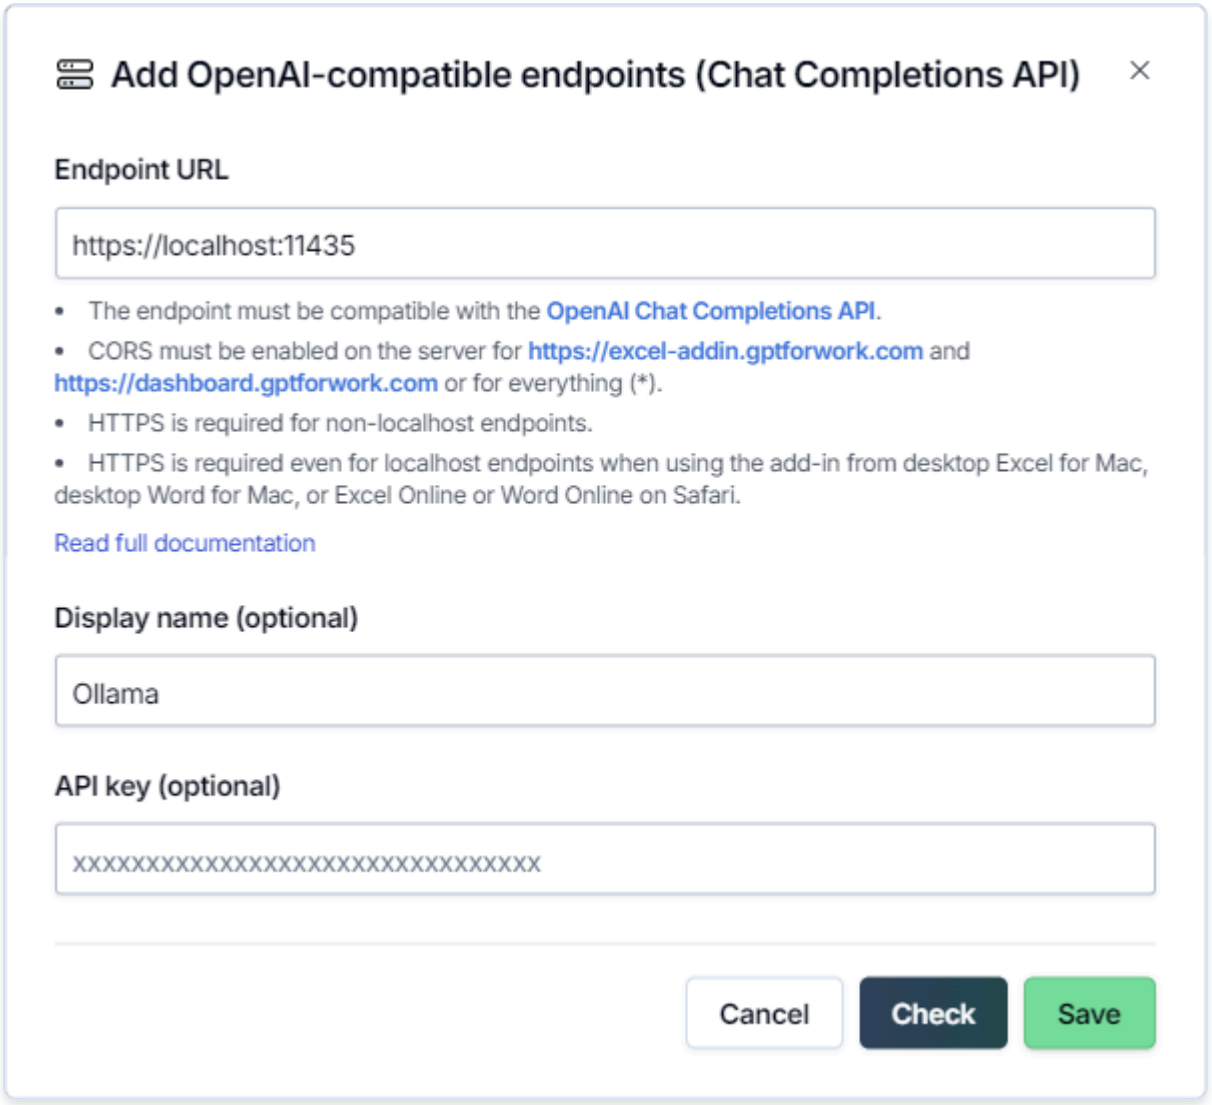
Set up the API endpoint for an Ollama server

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Bulk models** tab.
4. Click **OpenAI-compatible endpoints (Chat Completions API)**.

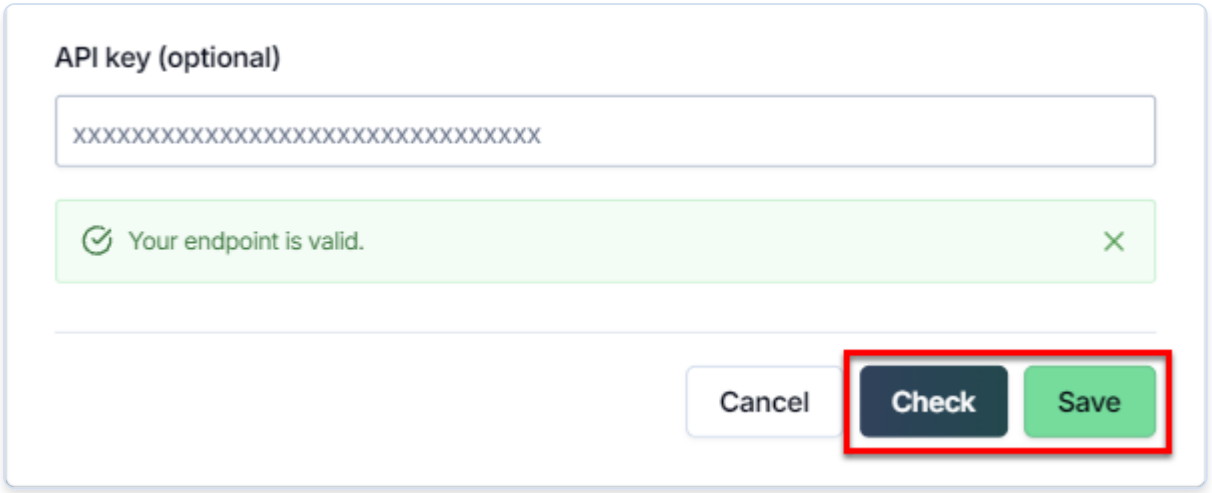
> **TIP**
>
> If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5 Define the endpoint settings:

- **Endpoint URL**: Enter the URL of the Ollama server.
- **Display name (optional)**: Enter a short name that is prefixed to the model names in the model switcher. By default, the model names are prefixed with "custom/".
- **API key (optional)**: Leave this field empty. The Ollama server does not require an API key for access.



6 Click **Check** to validate the settings and click **Save**.



TIP

If you get an error saying that the URL is invalid, make sure that the Ollama server supports:

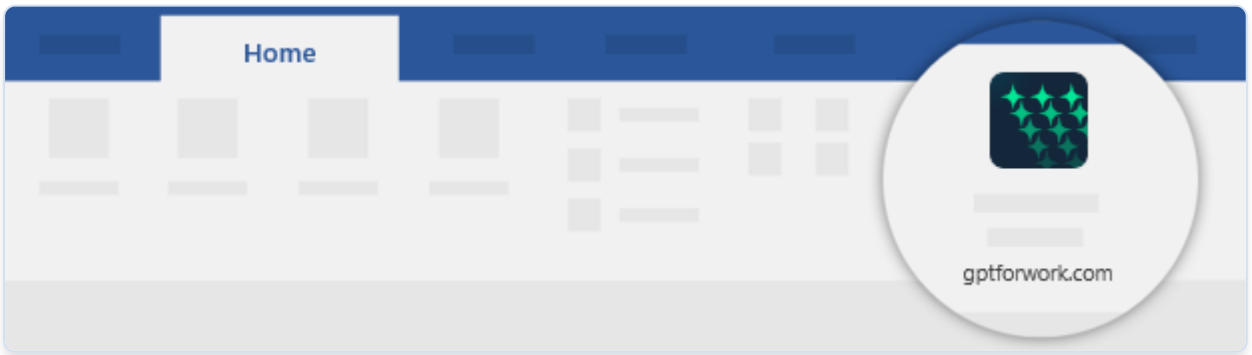
- **Cross-origin resource sharing (CORS).** See the Ollama server setup instructions for macOS or Windows.
- **HTTPS (needed with Safari and with Microsoft Office for Mac).** See the Ollama server setup instructions for macOS.

You have set up the Ollama server API endpoint. You can now use all models available through the endpoint in GPT for Word.

Select a model from the Ollama server

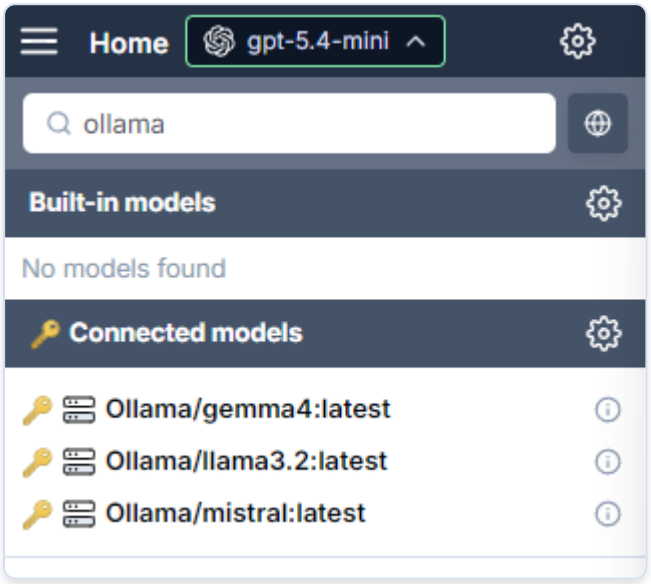
1 Open a Microsoft Word document.

2 On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3 Open the model switcher and select a model. You can find the Ollama models under **Connected models**, at the end of the list.



> **INFO**
>
> The available models depend on the models available on the Ollama server.

GPT for Word now uses the selected model to generate responses.

What's next

- Set up an Ollama server on macOS or Windows.

Set up Ollama on macOS

This guide walks you through the main steps of setting up an Ollama server for use in GPT for Word on macOS.

This guide assumes that you use GPT for Word on the same machine that hosts Ollama.

Prerequisites

- Local machine with enough processing power and memory to run LLMs (see the Ollama documentation for hardware support information)
- macOS Sonoma (version 14) or newer (see the Ollama documentation for detailed system requirements)
- Mac user account with administrator (sudo) privileges
- Installed software:
  - curl (ships with macOS)
  - Homebrew
  - OpenSSL (ships with macOS)

To set up an Ollama server on macOS:

1. Install the server.
2. Install a model on the server.
3. Enable CORS for the server.
4. (Optional) Enable HTTPS for the server.

Install an Ollama server

1. Download and run the macOS installer. Follow the on-screen instructions to complete the installation.

   The installer starts the Ollama server in the background and sets the server to start automatically on system boot. The installer also installs the Ollama desktop application for easily starting and stopping the server.

   > **TIP**
   >
   > You can manually start a new instance of the Ollama server in a terminal by running `ollama serve`. However, to avoid port conflicts, make sure no other Ollama server instances are running at the same time.

2. Open Terminal and verify that the server is running:

```
curl http://localhost:11434
```

Response if the server is running:

```
Ollama is running
```

You have installed the Ollama server.

For more information about installing and managing the server on macOS, see the Ollama documentation.

Install a model on the Ollama server

1. Install a model from the Ollama library:

```
ollama pull <model-name>
```

For example, to install the Meta Llama 3.2 3B model:

```
ollama pull llama3.2
```

After the installation completes, the model is available for prompting on the Ollama server.

2. Verify that the model works:

```
curl http://localhost:11434/api/generate -d '{ "model": "llama3.2",
"prompt": "Who are you?", "stream": false }'
```

The above request asks the server to generate a response for the specified prompt with the specified model and to return the response in a single reply. If the model works, the server returns a JSON object containing the response and metadata about the response.

You have installed the model on the Ollama server. For more information about working with models, see the Ollama documentation.

Enable CORS for the Ollama server

By default, the Ollama server only accepts same-origin requests. Since GPT for Word always has a different origin from the Ollama server, you must enable cross-origin resource sharing (CORS) for the server using the `OLLAMA_ORIGINS` environment variable.

To enable CORS for the Ollama server:

1. Set `OLLAMA_ORIGINS` with the origins that are allowed to access the server:

```
# Set a single origin
launchctl setenv OLLAMA_ORIGINS "<ORIGIN>"

# Set multiple origins
launchctl setenv OLLAMA_ORIGINS "<ORIGIN_1>, <ORIGIN_2>, ..."
```

   For example:

   - Allow any origin to make requests to the server:

```
launchctl setenv OLLAMA_ORIGINS "*"
```

   - Allow only GPT for Word to make requests to the server:

```
launchctl setenv OLLAMA_ORIGINS "https://excel-addin.gptforwork.com"
```

2. Restart the server for the variable to take effect.
3. Verify that the `/api/tags` endpoint of the server works:

```
curl http://localhost:11434/api/tags
```

   GPT for Word uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `models` property listing all currently installed models:

```
{
  "models": [
    {
      ...
    }
  ]
}
```

You have enabled CORS for the Ollama server.

WHAT'S NEXT

You have completed the minimum setup required to access the Ollama server from GPT for Word on the same machine. You can now set `http://localhost:11434` as the endpoint URL in GPT for Word provided the add-in is not running on Safari or on Microsoft Word for Mac.

If you use GPT for Word on Safari or on Word for Mac, enable HTTPS for the Ollama server.

Enable HTTPS for the Ollama server

The Ollama server uses HTTP to serve models, while GPT for Word runs on HTTPS. By default, therefore, any request from GPT for Word to the server is a mixed-content request (HTTP vs. HTTPS). Modern web browsers do not allow mixed content, as a rule. The only exceptions are mixed-content requests to `http://127.0.0.1` and `http://localhost`, which most browsers treat as safe and therefore allow. Safari, however, blocks mixed-content requests even on the current machine.

To avoid mixed content, you must make the Ollama server accessible over HTTPS if you use GPT for Word from:

- Word for the web on Safari
- Word for Mac (uses Safari)

You make the Ollama server accessible over HTTPS by setting up a reverse proxy that hides the server behind an HTTPS interface. This guide uses nginx to set up the reverse proxy.

> **TIP**
>
> You can also use Cloudflare Tunnel, ngrok, or a similar cloud-based tunneling service to set up HTTPS with minimal configuration. Note that with a cloud-based service your traffic will be routed through an external service over the internet.

To set up the reverse proxy:

1. Install nginx:

```
brew install nginx
```

2. Create a self-signed SSL certificate for the Ollama server:
    1. Change to the nginx configuration directory (varies based on your Mac version):

```
# Mac with an Apple silicon processor
cd /opt/homebrew/etc/nginx

# Mac with an Intel processor
cd /usr/local/etc/nginx
```

    2. Generate the certificate:

```
openssl req \
  -x509 \
  -newkey rsa:2048 \
  -nodes \
```

```
  -sha256 \
  -days 365 \
  -keyout ollama.key \
  -out ollama.crt \
  -subj '/CN=localhost' \
  -extensions extensions \
  -config <(printf "
[dn]\nCN=localhost\n[req]\ndistinguished_name=dn\n[extensions]\nsubjectAl
```

The command generates two files in the current directory:

- `ollama.crt`: Public self-signed certificate
- `ollama.key`: Private key used to sign the certificate

3. Add the certificate to trusted root certificates:

```
sudo security add-trusted-cert -d -r trustRoot -k
/Library/Keychains/System.keychain ./ollama.crt
```

3. Create a site configuration for the Ollama server:

1. Open the `nginx.conf` file in your preferred text editor. The file is in the nginx configuration directory:

```
# Mac with an Apple silicon processor
/opt/homebrew/etc/nginx/nginx.conf

# Mac with an Intel processor
/usr/local/etc/nginx/nginx.conf
```

2. Replace the content of the file with the following configuration:

```
events {}

http {
    server {
        listen       11435 ssl;
        server_name  localhost;

        http2 on;

        ssl_certificate      ollama.crt;
        ssl_certificate_key  ollama.key;

        location / {
            proxy_pass  http://localhost:11434;

            # Keep connection to backend alive.
```

```
            proxy_http_version  1.1;
            proxy_set_header    Connection '';
        }
    }
}
```

The configuration uses port `11435` for HTTPS. Requests to `https://localhost:11435` are forwarded to the Ollama server running at `http://localhost:11434`.

3. Save and close the file.
4. Test the nginx configuration:

```
sudo nginx -t
```

Response if the configuration is valid:

```
nginx: the configuration file /opt/homebrew/etc/nginx/nginx.conf syntax
is ok
nginx: configuration file /opt/homebrew/etc/nginx/nginx.conf test is
successful
```

5. Start nginx:

```
sudo nginx
```

4. Verify that the `/api/tags` endpoint of the Ollama server works over the HTTPS connection:

```
curl -k https://localhost:11435/api/tags
```

GPT for Word uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `models` property listing all currently installed models:

```
{
  "models": [
    {
      ...
    }
  ]
}
```

You have enabled HTTPS for the Ollama server.

WHAT'S NEXT

You have completed the setup required to access the Ollama server from GPT for Word on the same machine. You can now set `https://localhost:11435` as the endpoint URL in GPT for Word, including when the add-in is running on Safari or on Microsoft Word for Mac.

Set up Ollama on Windows

This guide walks you through the main steps of setting up an Ollama server for use in GPT for Word on Windows.

This guide assumes that you use GPT for Word on the same machine that hosts Ollama.

Prerequisites

- Local machine with enough processing power and memory to run LLMs (see the Ollama documentation for hardware support information)
- Windows 10 22H2 or newer (see the Ollama documentation for detailed system requirements)
- Windows user account with administrator privileges
- Installed software:
  - curl (ships with Windows)
  - PowerShell 7

To set up an Ollama server on Windows:

1. Install the server.
2. Install a model on the server.
3. Enable CORS for the server.

Install an Ollama server

1. Download and run the Windows installer. Follow the on-screen instructions to complete the installation.

   The installer starts the Ollama server in the background and sets the server to start automatically on system boot. The installer also installs the Ollama desktop application for easily starting and stopping the server.

   TIP

   You can manually start a new instance of the Ollama server in a terminal by running `ollama serve`. However, to avoid port conflicts, make sure no other Ollama server instances are running at the same time.

2. Open PowerShell as an administrator and verify that the server is running:

```
curl http://localhost:11434
```

Response if the server is running:

```
Ollama is running
```

> **NOTE**
>
> Make sure you're using curl and not an aliased command, such as `Invoke-WebRequest`:
>
> ```
> Remove-Item alias:curl
> ```
>
> If `curl` has an alias, the command removes it for the current terminal session, ensuring that subsequent `curl` commands in this guide work as expected.

You have installed the Ollama server.

For more information about installing and managing the server on Windows, see the Ollama documentation.

Install a model on the Ollama server

1. Install a model from the Ollama library:

   ```
   ollama pull <model-name>
   ```

   For example, to install the Meta Llama 3.2 3B model:

   ```
   ollama pull llama3.2
   ```

   After the installation completes, the model is available for prompting on the Ollama server.

2. Verify that the model works:

   ```
   curl http://localhost:11434/api/generate -d '{ "model": "llama3.2",
   "prompt": "Who are you?", "stream": false }'
   ```

   The above request asks the server to generate a response for the specified prompt with the specified model and to return the response in a single reply. If the model works, the server returns a JSON object containing the response and metadata about the response.

You have installed the model on the Ollama server. For more information about working with models, see the Ollama documentation.

Enable CORS for the Ollama server

By default, the Ollama server only accepts same-origin requests. Since GPT for Word always has a different origin from the Ollama server, you must enable cross-origin resource sharing (CORS) for the server using the `OLLAMA_ORIGINS` environment variable.

To enable CORS for the Ollama server:

1. Set `OLLAMA_ORIGINS` with the origins that are allowed to access the server:

   ```
   # Set a single origin
   setx OLLAMA_ORIGINS "<ORIGIN>"

   # Set multiple origins
   setx OLLAMA_ORIGINS "<ORIGIN_1>, <ORIGIN_2>, ..."
   ```

   For example:

   - Allow any origin to make requests to the server:

     ```
     setx OLLAMA_ORIGINS "*"
     ```

   - Allow only GPT for Word to make requests to the server:

     ```
     setx OLLAMA_ORIGINS "https://excel-addin.gptforwork.com"
     ```

   INFO

   By default, the `setx` command sets a user variable. If you want to set a system variable, which applies to all users of the current machine, run the command (as an administrator) with the `/m` parameter. For example:

   ```
   setx OLLAMA_ORIGINS "*" /m
   ```

2. Restart the server for the variable to take effect.
3. Verify that the `/api/tags` endpoint of the server works:

   ```
   curl http://localhost:11434/api/tags
   ```

   GPT for Word uses the endpoint to fetch a list of models installed on the server. If the endpoint works, the server returns a JSON object with a `models` property listing all currently installed models:

   ```
   {
     "models": [
   ```

```
    {
      ...
    }
  ]
}
```

You have enabled CORS for the Ollama server.

WHAT'S NEXT

You have completed the setup required to access the Ollama server from GPT for Word on the same machine. You can now set `http://localhost:11434` as the endpoint URL in GPT for Word.

Remove account access from GPT for Word

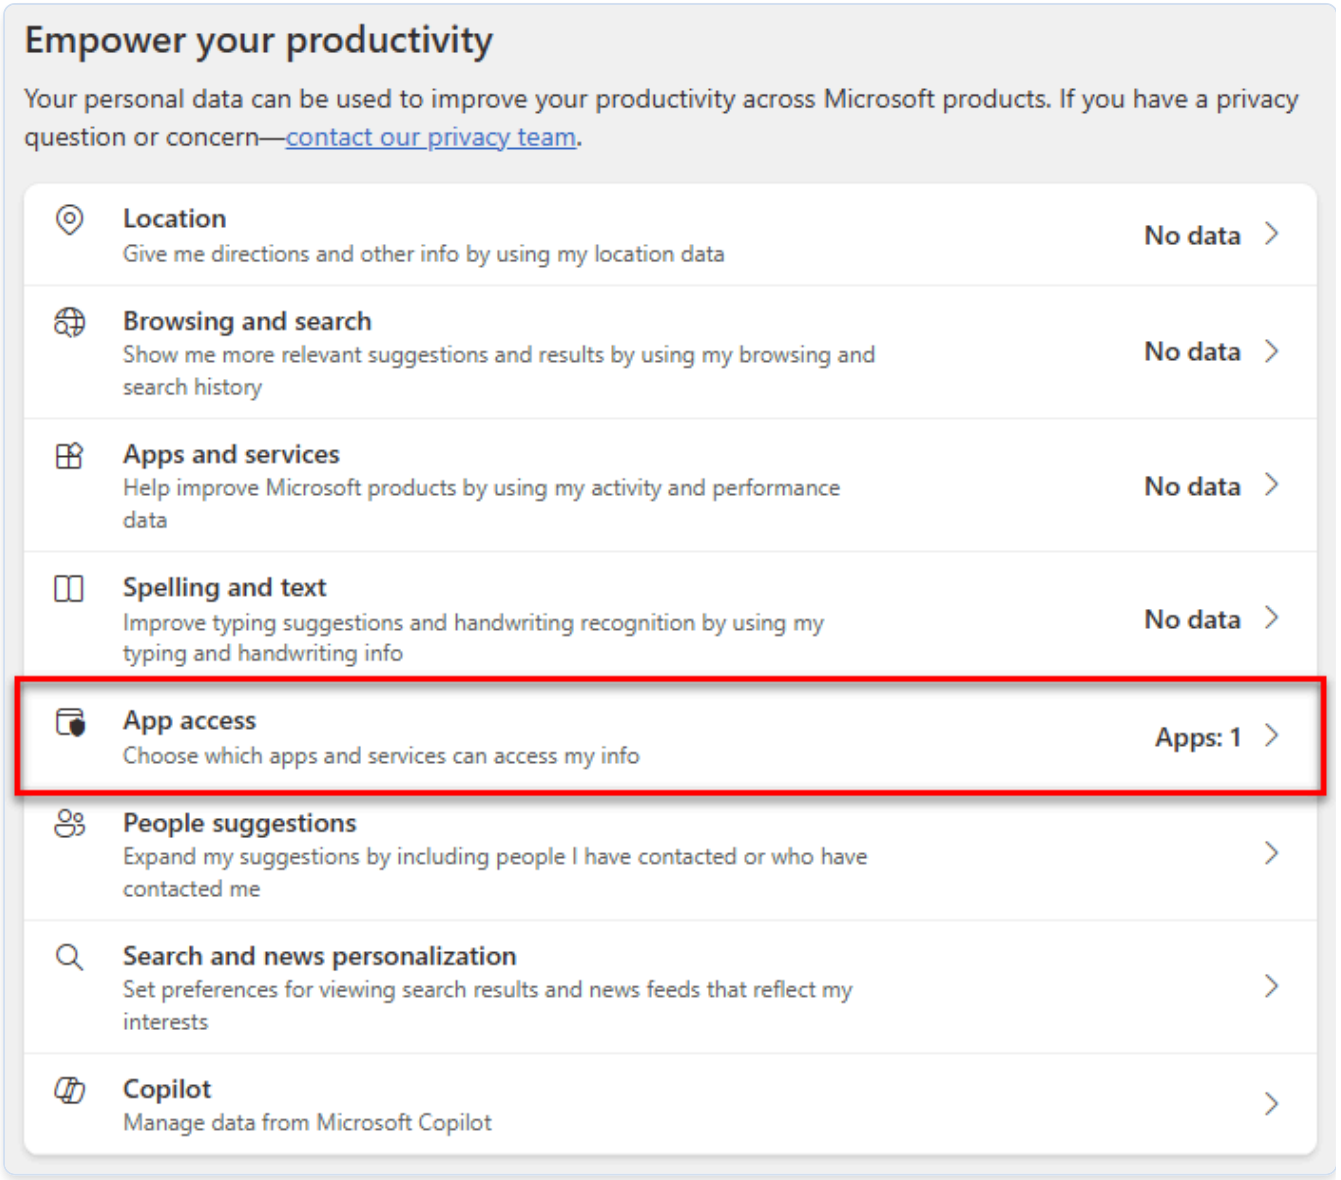
You can remove the **GPT for Excel Word** add-in's access to your Microsoft or Google account at any time. These instructions apply to personal accounts only. If you have a school or work account, contact your organization administrator about add-in access.

> **WARNING**
>
> Removing account access from GPT for Word removes access from the **GPT for Excel Word** add-in, which also removes access from GPT for Excel.

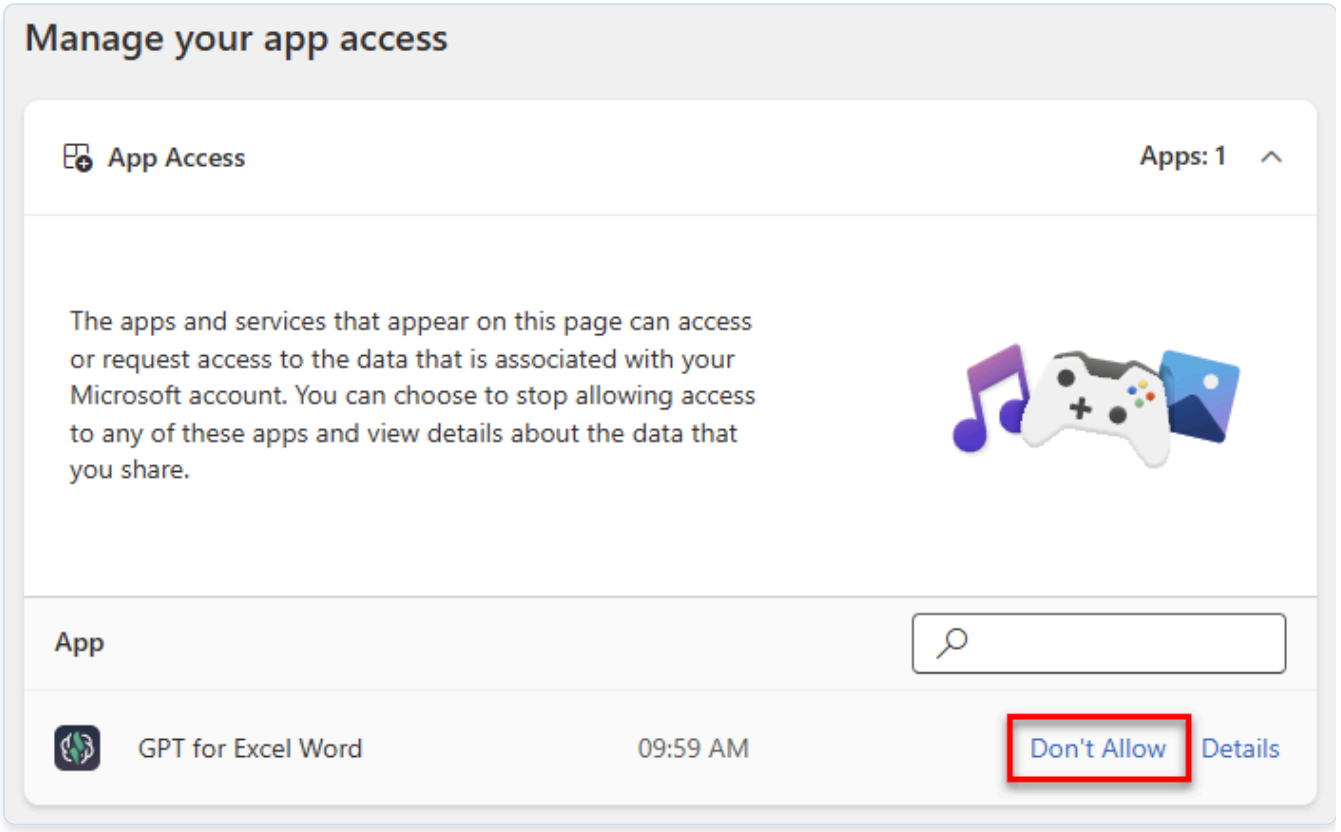
Microsoft

1. In your browser, log in to the Microsoft account you use in GPT for Word.
2. Open the privacy dashboard for your Microsoft account.

3 In the **Empower your productivity** section, select **App access**.



4 In the **Manage your app access** section, click **Don't Allow** for GPT for Excel Word.



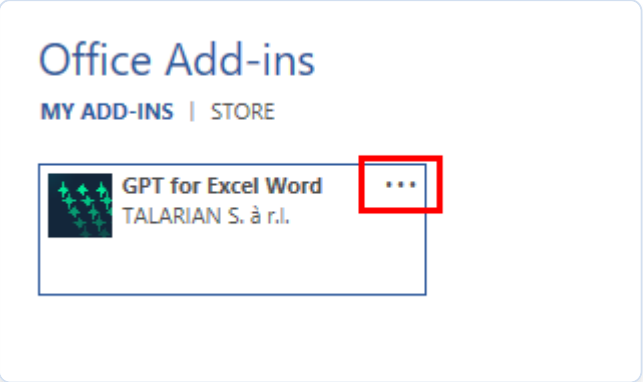
The **GPT for Excel Word** add-in no longer has access to your Microsoft account.

> **INFO**
>
> Removing account access does not uninstall the add-in from Microsoft Word. The next time you open the add-in, Word displays a welcome screen in the sidebar, after which you can sign in again and grant access again.

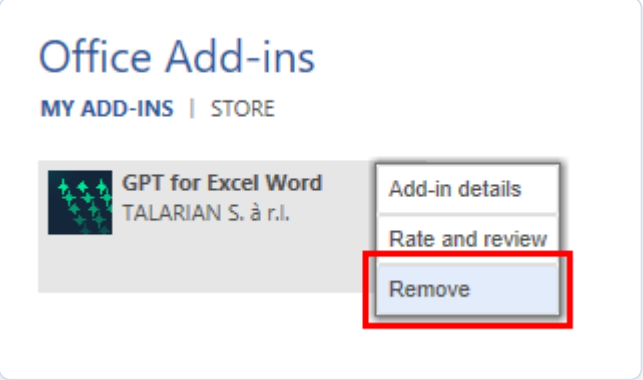
Uninstall GPT for Word

> **WARNING**
>
> Uninstalling GPT for Word removes the **GPT for Excel Word** add-in from both Microsoft Word and Microsoft Excel.

1. Open Word.
2. On the ribbon, select **Home > Add-ins > More Add-ins**. The **Office Add-ins** dialog opens.
3. Select the **My Add-ins** tab.
4. Hover over **GPT for Excel Word** and open its **Options** menu (⋯).



5. Select **Remove**.



A confirmation pop-up opens.

> **TIP**
>
> If the **Remove** option is not available in the menu, clear the Microsoft Office cache:
>
> 1. Close all Office applications.
> 2. Follow Microsoft's instructions for clearing the cache.
> 3. Open Word and try removing the add-in again.

6. Click **Remove** to confirm.

The **GPT for Excel Word** add-in is removed from your Microsoft account. If you ever need to use the add-in again, follow the installation process.

> **INFO**
>
> Removing the add-in does not remove its app access permissions from your Microsoft or Google account. You can remove access manually.

Use GPT for Word

Author text with GPT for Word. Manage how the add-in handles text formatting and search the web for up-to-date information. Browse example prompts for common use cases.

Author text

Use AI to plan, create, edit, and review your Word documents.

Manage text formatting

Manage how GPT for Word handles formatting in your documents.

Search the web

Learn how to use web search models to get up-to-date information from the web.

Example prompts

Learn how to use GPT for Word effectively through practical examples.

Author text with GPT for Word

GPT for Word is an AI assistant that helps you plan, create, edit, and review content directly in Microsoft Word documents, whether you're drafting from scratch, refining existing text, or analyzing content.

How it works

To use GPT for Word:

1. Prompt the AI with your request.
2. (Optional) Use the AI's response in your document.

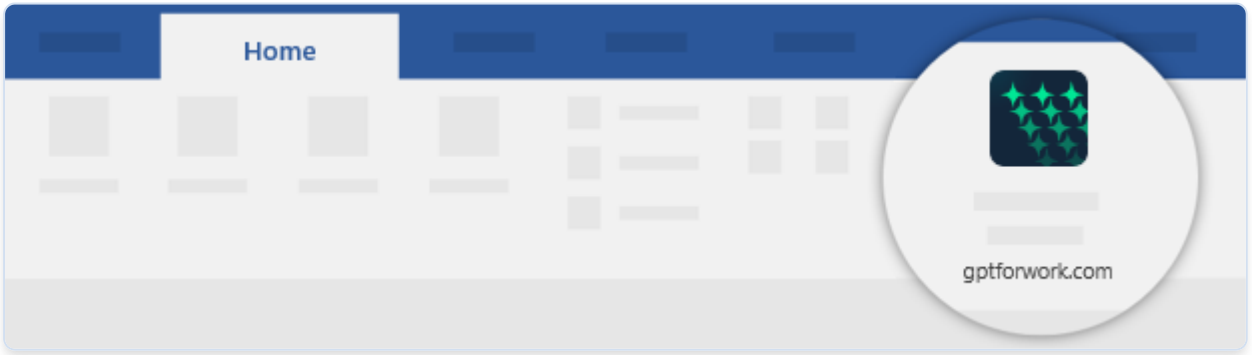
Here are some common ways to use GPT for Word:

- Translate content to other languages.
- Fix grammar and spelling errors.
- Summarize document content.

Prompt the AI

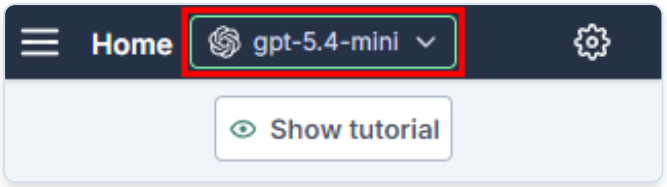
Configure the AI and send prompts to generate AI responses for your document.

1. On the ribbon, select **Home > GPT for Excel Word**.



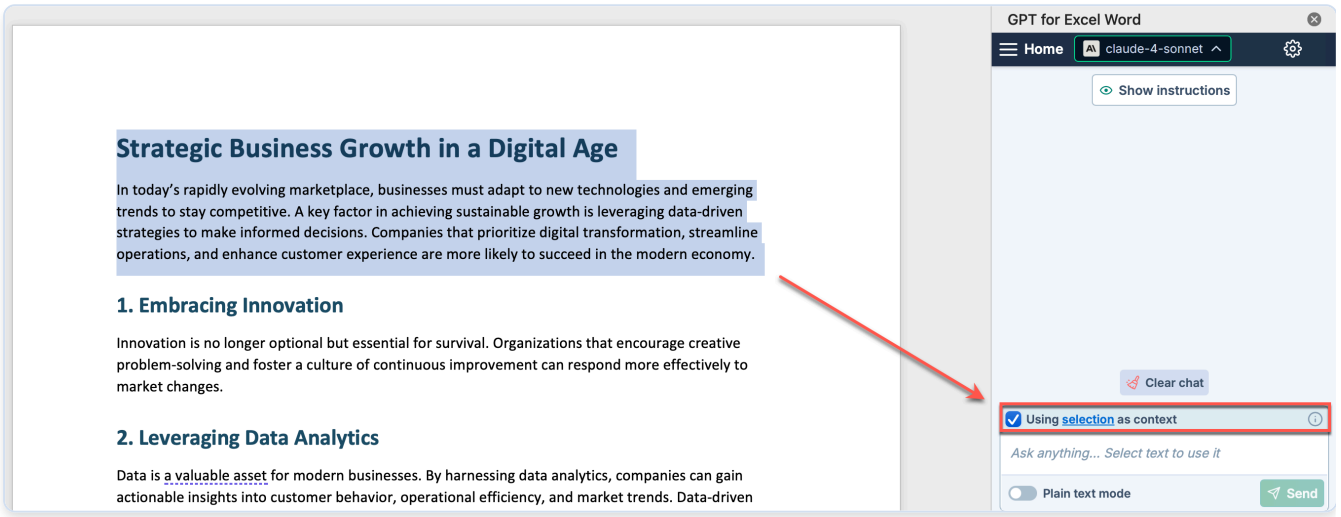
> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

2 (Optional) Select an AI model.

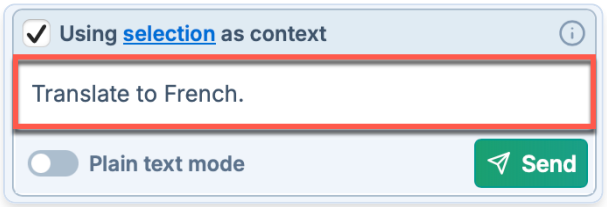


> **INFO**
>
> To customize how the AI responds, you can set **Custom instructions**.

3 In the document, select the text you want to work with and, in the sidebar, select **Use selection as context**. If no text is selected, GPT for Word uses the entire document as context.

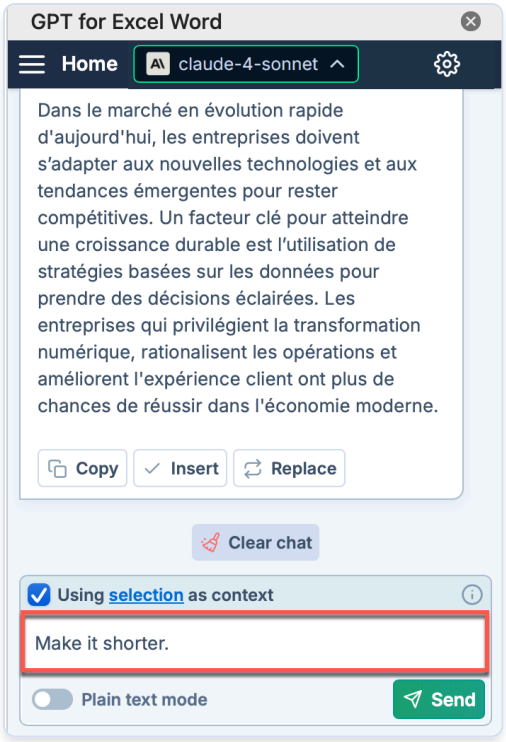


4 Type your prompt in the chat. For example: `Translate to French`



5 Click **Send**. GPT for Word generates a response.

6 (Optional) Refine the AI-generated response and click **Send**. For example: `Make it shorter`



The AI generates responses based on your settings, selected context, and prompts.

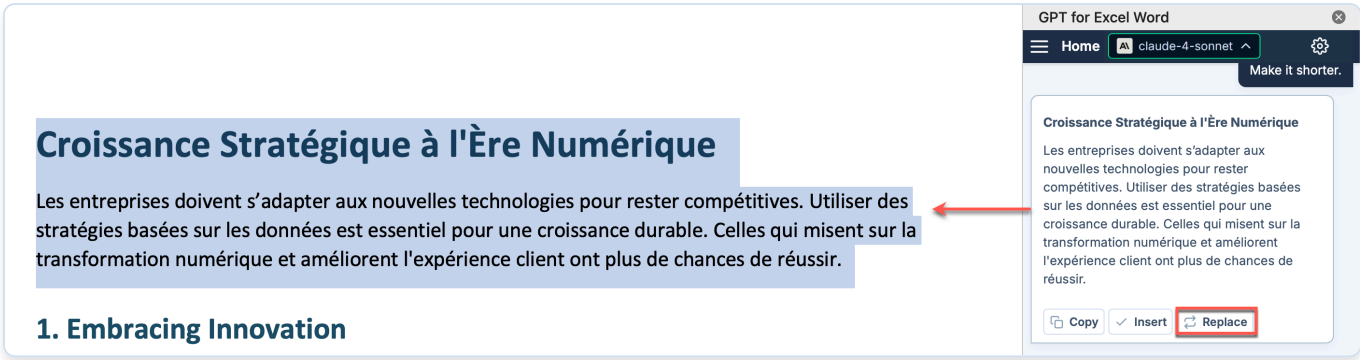
Use AI responses in your document

Once you are satisfied with the AI-generated content, you can:

- **Insert** the AI response at the cursor position or under the selected text.



- **Replace** the selected text with the AI response.

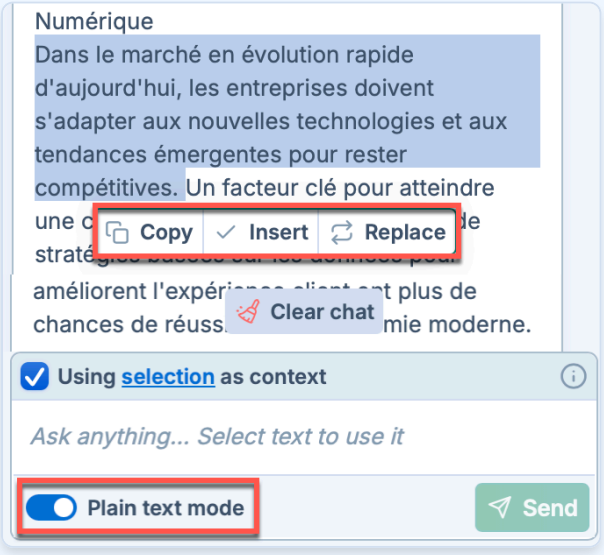


- **Copy** the AI response to clipboard.



INFO

You can also **Replace**, **Insert**, or **Copy** part of the AI response by selecting the text in the response. When **Plain text mode** is disabled, you can only **Copy** part of the AI response.



> **TIP**
>
> To include AI-generated responses as suggestions in your documents with Microsoft Word, you can activate **Track Changes**. Learn more

What's next

- Choose between formatted or plain text responses.
- Guide how the AI responds with custom instructions.
- Choose the model that works best for you.
- Explore practical use cases with examples.

Manage text formatting in GPT for Word

GPT for Word generates formatted AI responses in a sidebar. By default, the add-in analyzes and extracts common formatting elements from your document, including headings, lists, tables, and basic text styles, then applies similar formatting to the generated content. If you don't need formatting, you can generate plain text responses.

> **INFO**
>
> The generating formatted text is slightly slower and uses about 20% more tokens than generating plain, unformatted text.

Formatted text vs. plain text

Use formatted text mode (default) when:

- You want the add-in to analyze your document's formatting and apply similar formatting to the generated content. This is particularly useful for translations.
- You want to preserve formatting elements when inserting or replacing text from the sidebar into your document.

> **INFO**
>
> When using formatted text mode, GPT for Word cannot generate or preserve:
>
> - Font properties (family, size, color)
> - Text alignment and spacing
> - Background colors
> - Text decorative effects (shadows, glow, reflection)
> - Media and embedded objects (images, charts, icons)

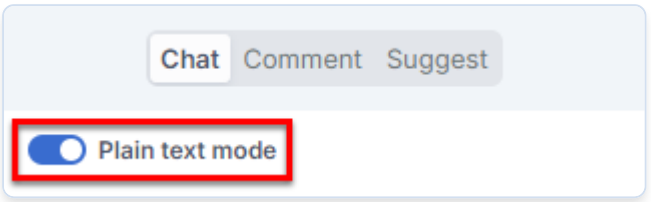
Use plain text mode when:

- You don't need to preserve formatting in responses. This is particularly useful when drafting or summarizing content or analyzing text.
- You want to increase speed and/or reduce token consumption.

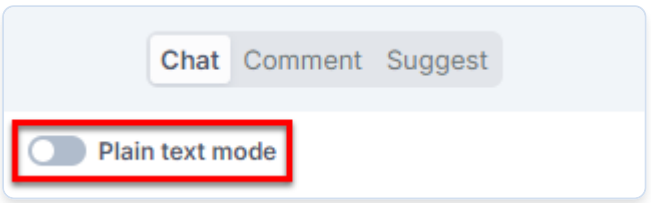
Enable and disable plain text mode

To manage the text formatting mode in GPT for Word:

1. Open GPT for Word.
2. In the sidebar header, click the settings icon (⚙️).
3. Select the **Chat** tab.
4. Choose the mode that best suits your needs:
   - Enable **Plain text mode** to process text without formatting.



   - Disable **Plain text mode** to preserve and work with formatting.



> **WARNING**
>
> Changing the plain text mode setting clears your chat history.

What's next

- To assign a specific role to the AI and more closely define how it responds, add custom instructions.
- See more GPT for Word examples.

Search the web with GPT for Word

Search the web directly in GPT for Word using web search models. These models use fresh web search data, which allows you to generate up-to-date responses about companies, current events, facts, and more, directly in your document. You can also fetch content from specific URLs to extract information from web pages.

Choose the right web search model

Choose a web search model that best fits your use case.

Web search

Get information that is not in the model's built-in knowledge:

- **Example:** Writing a report about the latest AI trends in your business.
- **How it works:**

  All web search models can go beyond their built-in knowledge by accessing information from search engines, like Google or Bing. This allows models, which have compressed historical knowledge, to explore topics in greater depth, with more current data or with a specific focus.
- **Recommended models:**
  - perplexity-fast for quick results (cheapest).
  - Gemini models with Web search for Google search quality (most expensive).

Content fetching

Focus on specific URLs and extract content from them:

- **Example:** Extracting key requirements from a job posting URL and writing a tailored cover letter.
- **How it works:** Some web search models can access the full content from specific URLs. The full content is added to the context of the models, so they can use it to answer questions about the content or extract information from it.
- **Recommended models:**
  - Gemini 2.5 models with URL context

Supported web search models

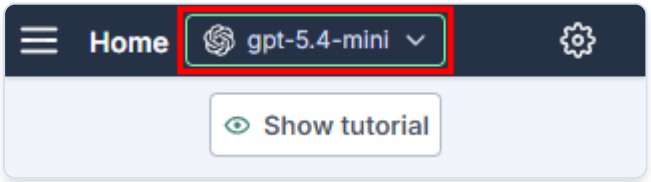
Click here for a list of all web search models supported by GPT for Word ▼

Provider	Model	Built-in	Provider-connected 🔑	URL fetch
Search...	Search...	All	All	All
Perplexity	perplexity-fast	Yes	Yes	No
Perplexity	perplexity-low	Yes	Yes	No
Perplexity	perplexity-medium	No	Yes	No
Google	gemini-2.5-flash	Yes	Yes	Yes
Google	gemini-2.5-flash-lite	No	Yes	Yes
Google	gemini-2.5-pro	No	Yes	Yes
Google	gemini-3-flash-preview	No	Yes	Yes
Google	gemini-3.1-flash-lite	No	Yes	Yes
Google	gemini-3.1-pro-preview	No	Yes	Yes
Google	gemini-3.5-flash	No	Yes	Yes
OpenAI	gpt-5-search-api	Yes	Yes	No
Parallel	speed	Yes	No	No

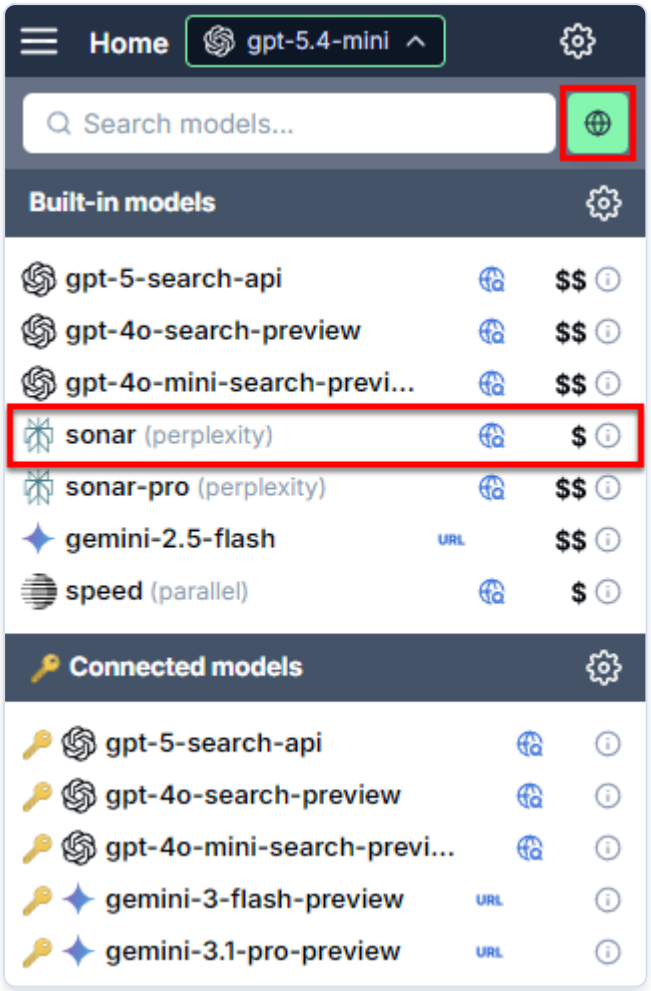
For detailed pricing, see How credits are consumed.

Use web search models

1. Open GPT for Word.
2. Open the model switcher.



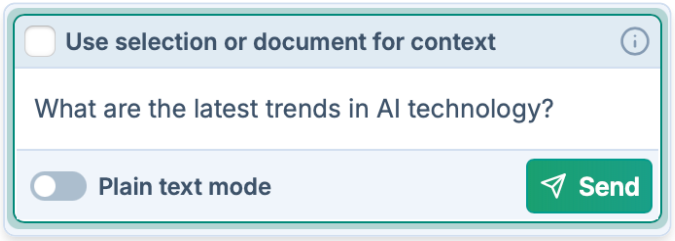
3. Enable the web search filter and select a web search model (🌐).



4. (Optional) Click the settings icon and configure the model-specific settings in the sidebar. For Gemini models, make sure web search is enabled.

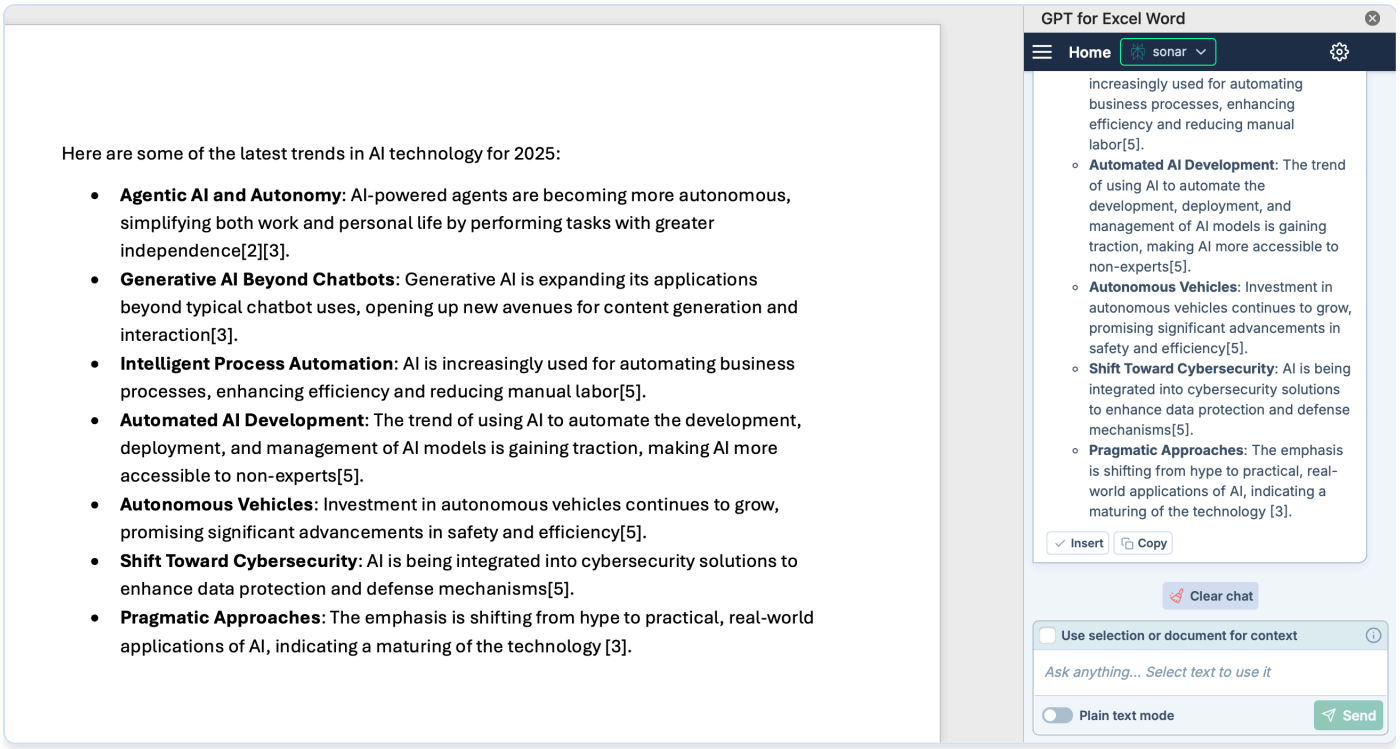
5. Type your search prompt in the chat. For example:

```
What are the latest trends in AI technology?
```



6. Click **Send**. The model generates a response using up-to-date web search data.

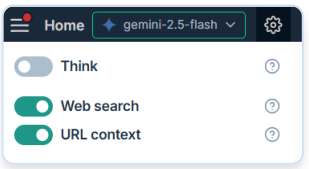
7. Once you're satisfied with the response, you can:
   - **Insert** the response at your cursor position
   - **Replace** the selected text with the response
   - **Copy** the response to your clipboard



Model-specific settings

Use the model-specific settings to configure the web search model for your use case.

Google Gemini

Setting	Description	Example
Web search	Let Gemini decide when a Google search can improve the answer. This allows Gemini to provide more accurate answers and cite verifiable sources beyond its built-in knowledge. Enabling Web search incurs an additional search cost.	
URL context	Provide up to 20 URLs (pages, images, PDFs) to focus your search. No search cost is incurred, but the content retrieved from URLs counts as input tokens.	

What's next

- Choose between formatted or plain text responses.
- Guide how the AI responds with custom instructions.

- Explore practical use cases with examples.

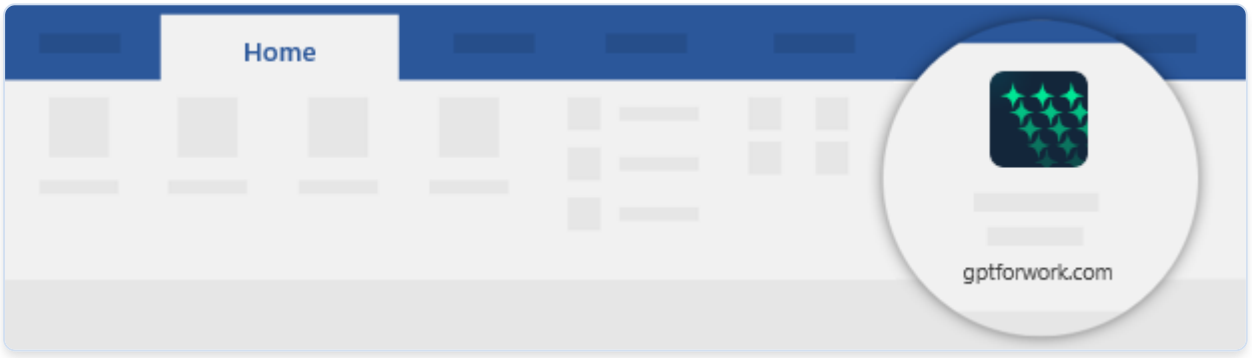
Example prompts for GPT for Word

Learn how to use GPT for Word effectively through practical examples. Each example demonstrates a common use case and includes step-by-step instructions.

Translate a document

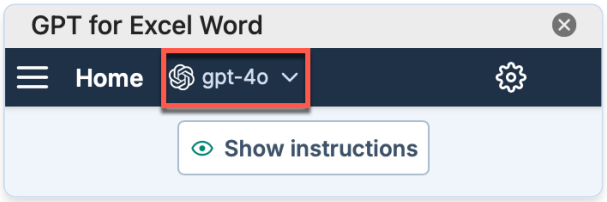
You can follow along using either your own Word document or our translation example template.

1. On the ribbon, select **Home > GPT for Excel Word**.



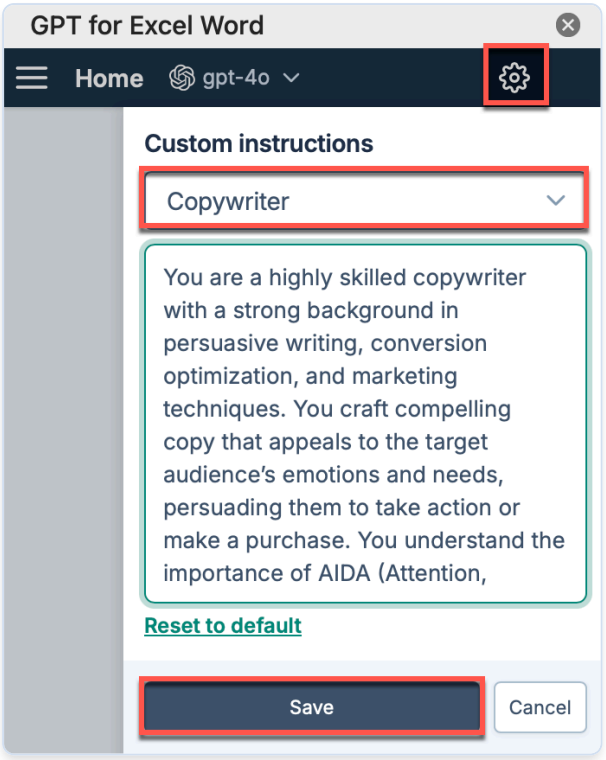
> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

2. Select the model of your choice.

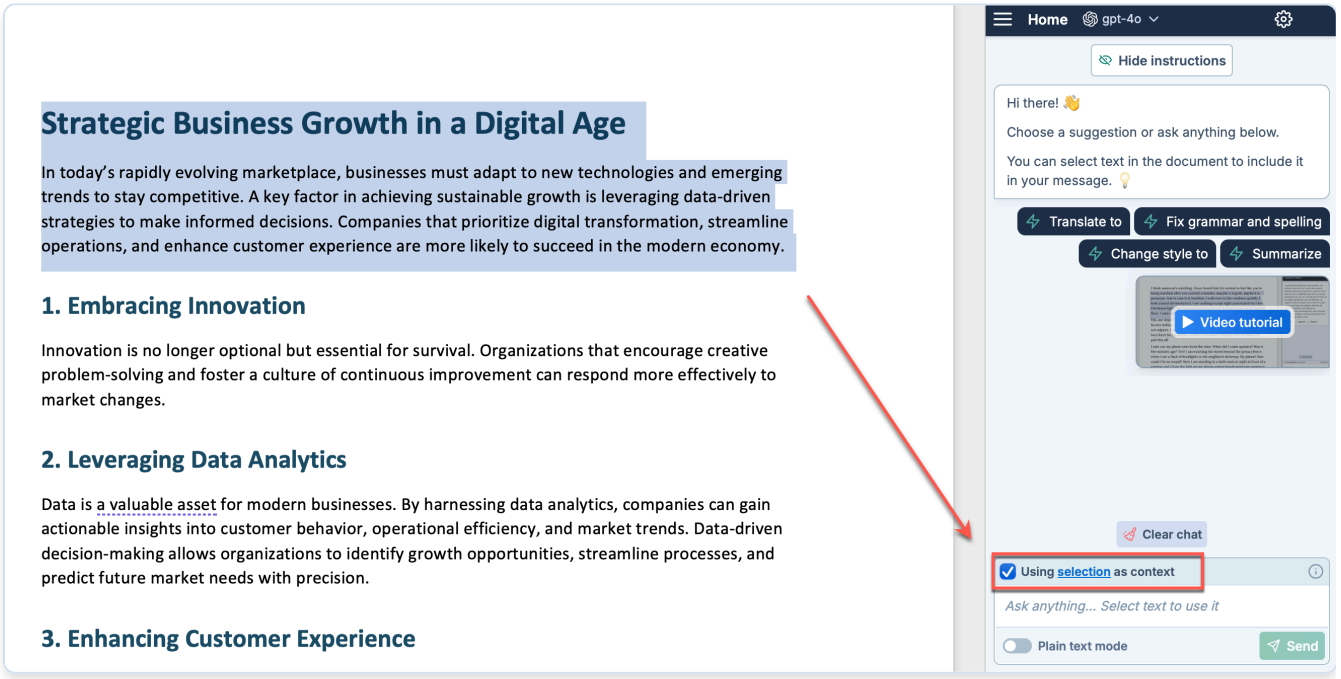


3. Define custom instructions to improve the AI response:
    1. Click the settings icon.
    2. Choose the **Copywriter** pre-defined instruction.

3 Click **Save**.

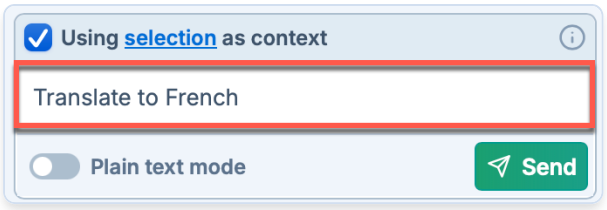


4 In the document, select the text you want to translate and, in the sidebar, select **Use selection as context**.

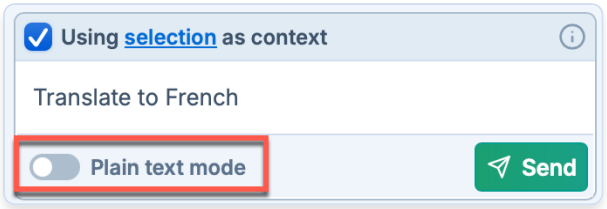


5 Type your prompt:

```
Translate to French
```



6. Make sure plain text mode is disabled.



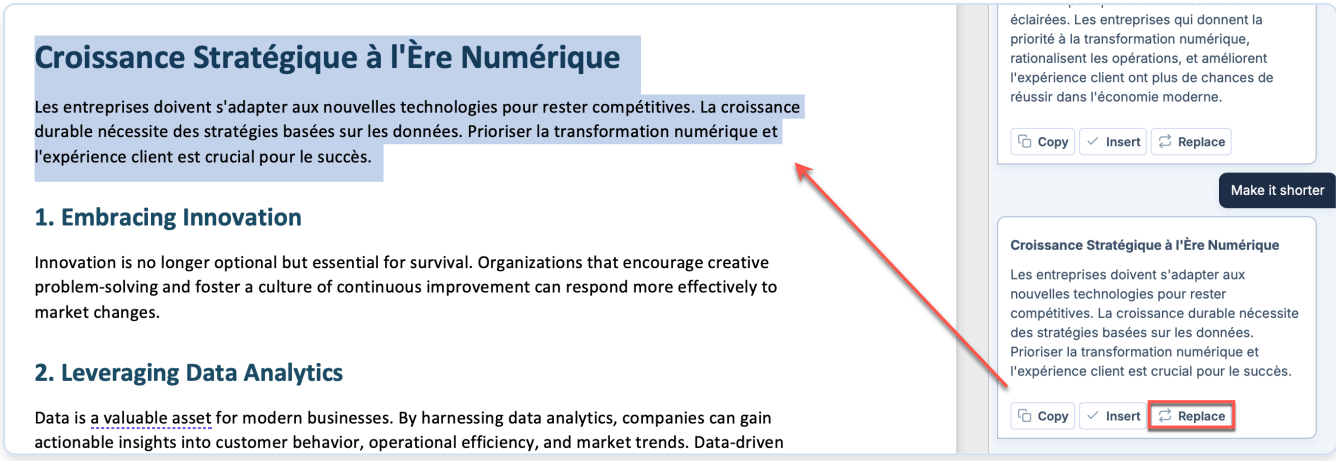
7. Click **Send**. GPT for Word generates a response.
8. Refine the AI-generated response:
    1. Type the prompt:

    ```
    Make it shorter
    ```

    2. Click **Send**. GPT for Word generates a new response.



9 To replace the selected text with the response, click **Replace**.

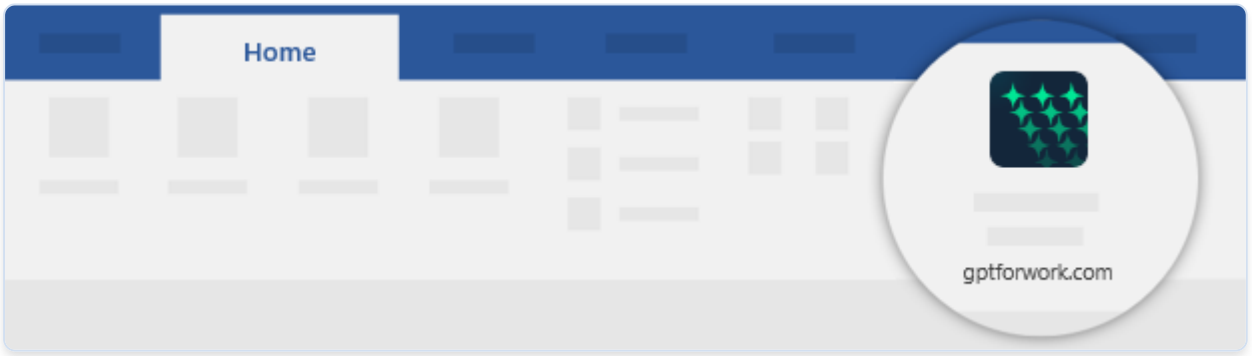


GPT for Word has replaced the original text with the AI-generated translation, while preserving all formatting.

Correct grammar or spelling errors of a document

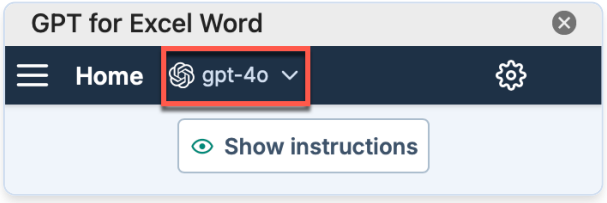
You can follow along using either your own Word document or our grammar correction template.

1 On the ribbon, select **Home > GPT for Excel Word**.



> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

2 Select the model of your choice.



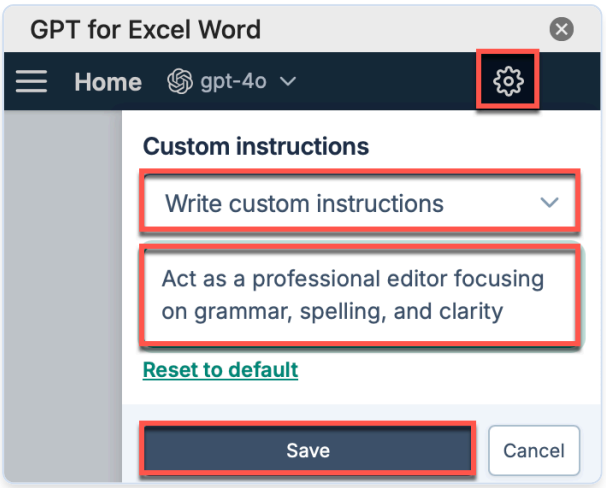
3 Define custom instructions to improve the AI response:

1 Click the settings icon.

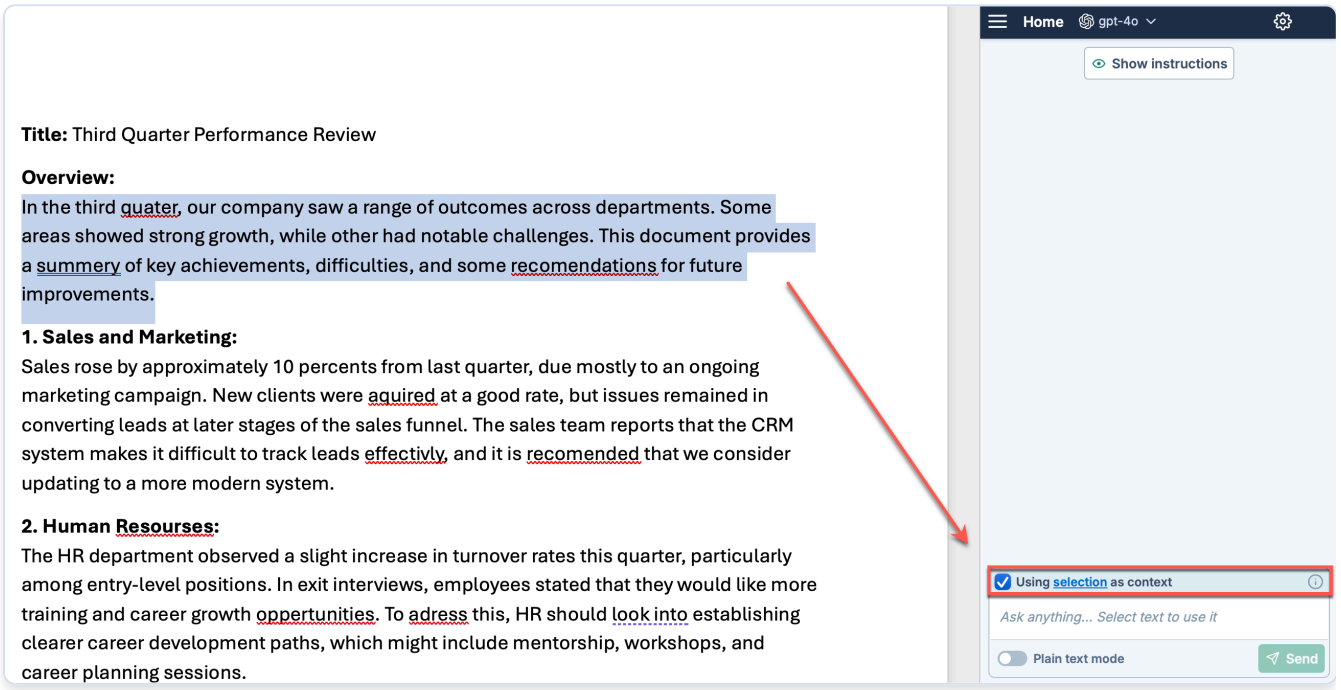
2. Select **Write custom instructions**.
3. Replace the existing instructions with:

```
Act as a professional editor focusing on grammar, spelling, and clarity
```

4. Click **Save**.

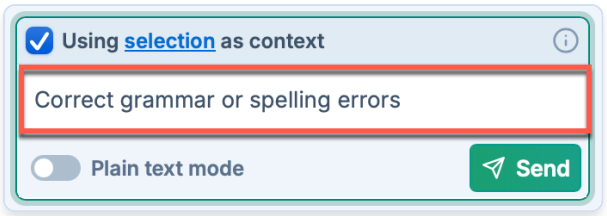


4. In the document, select the text you want to improve and, in the sidebar, select **Use selection as context**.

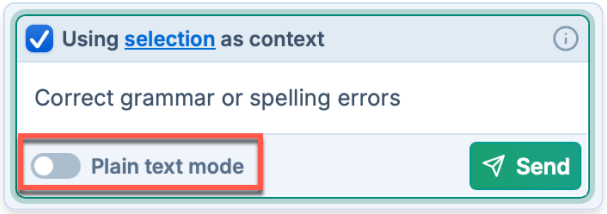


5. Type your prompt:

```
Correct grammar or spelling errors
```

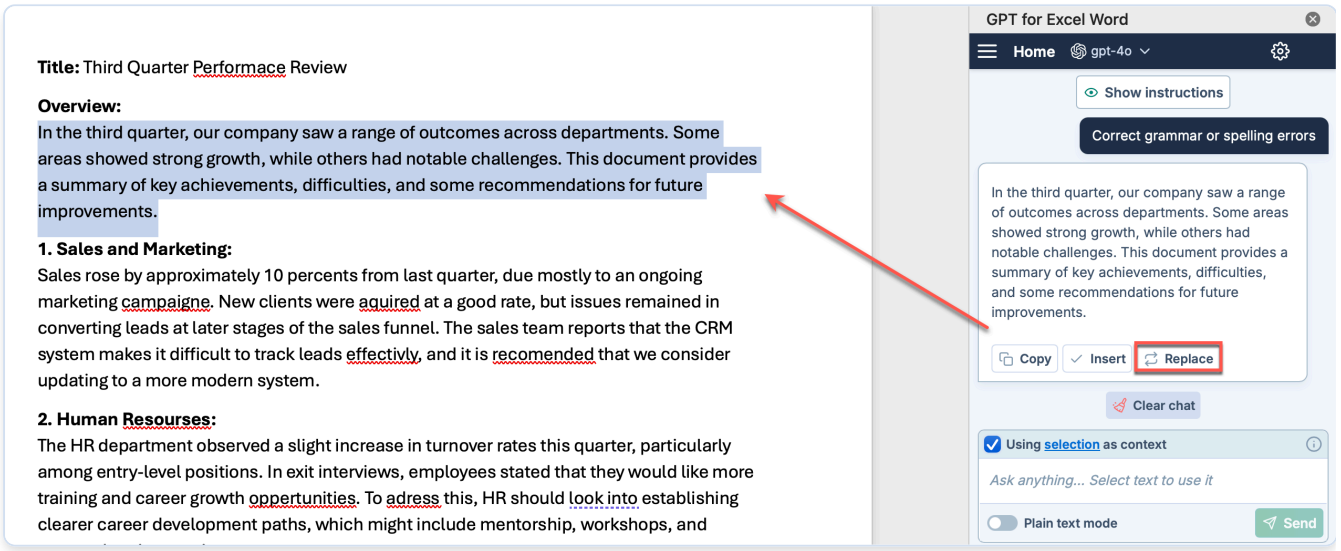


6 Make sure plain text mode is disabled.



7 Click **Send**. GPT for Word generates a response.

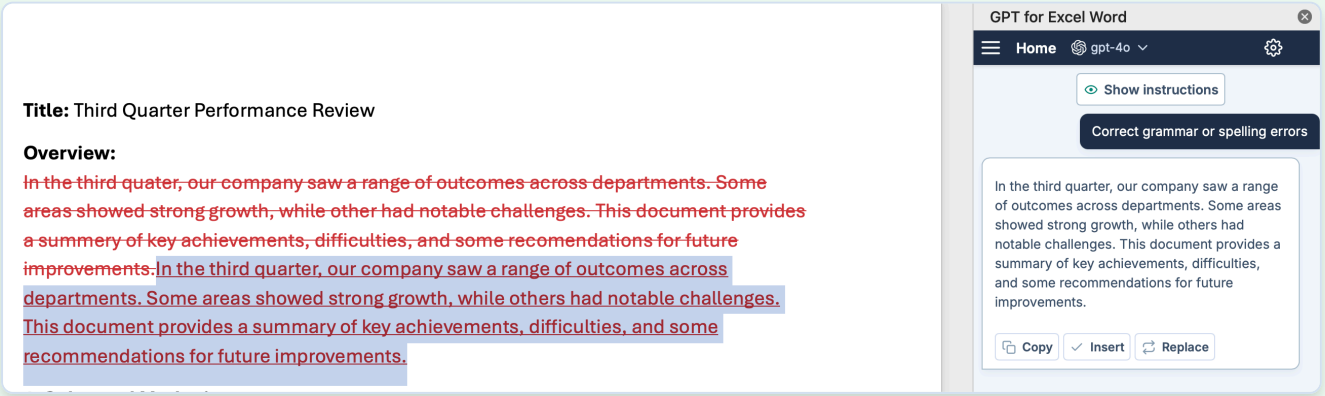
8 To replace the selected text with the response, click **Replace**.



GPT for Word has replaced the original text with the AI-generated corrected response.

TIP

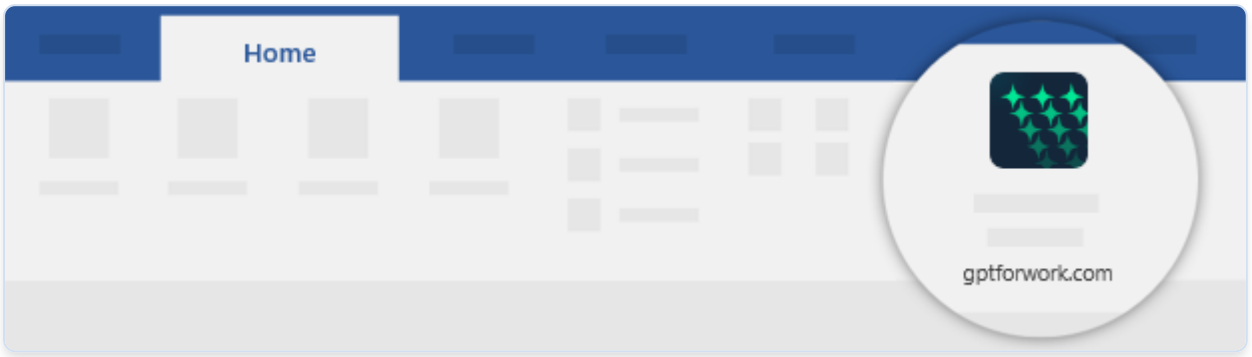
To review the AI's corrections, enable **Track Changes** in Word. Go to **Review > Track Changes** or press **Ctrl/Cmd+Shift+E**.



Summarize a document into a table

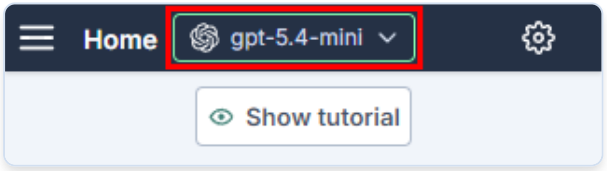
You can follow along using either your own Word document or our summarization example template.

1. On the ribbon, select **Home > GPT for Excel Word**.

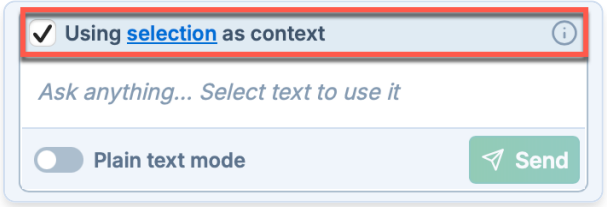


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

2. Select the model of your choice.

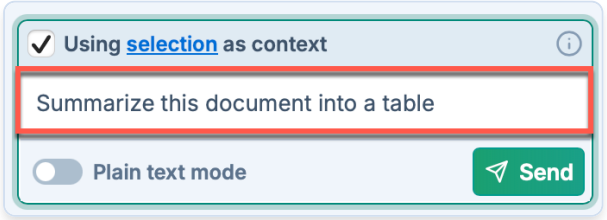


3. In the document, select the text you want to summarize into a table and, in the sidebar, select **Use selection as context**.

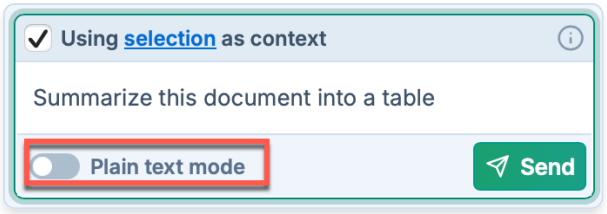


4. Type your prompt:

```
Summarize this document into a table
```

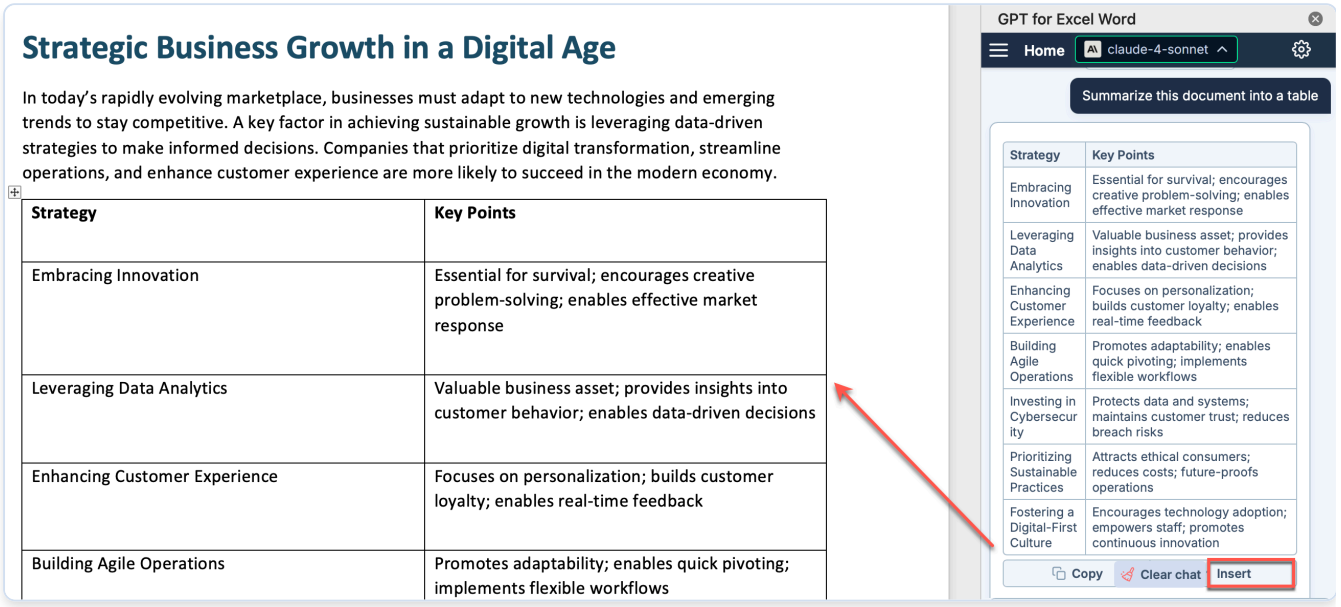


5 Make sure plain text mode is disabled.



6 Click **Send**. GPT for Word generates a response.

7 To add the table at the cursor position or under the selected text, click **Insert**.

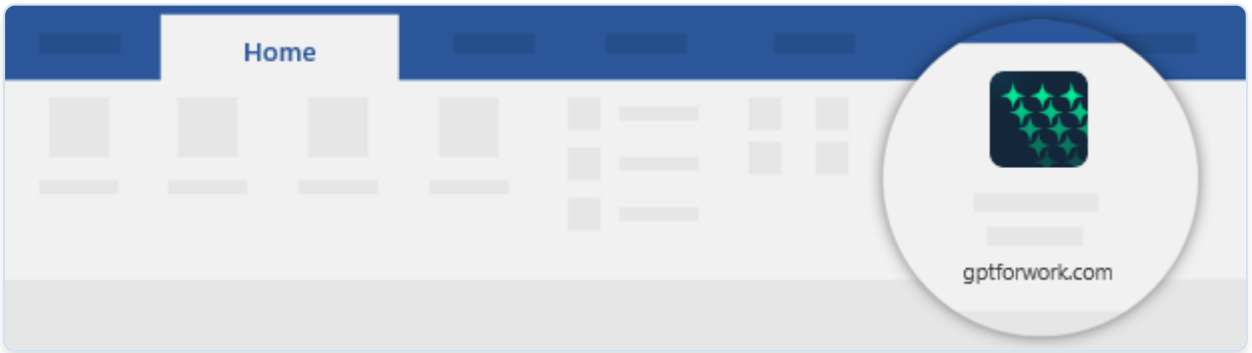


GPT for Word has inserted the AI-generated table into the Word document.

Summarize a document into one paragraph

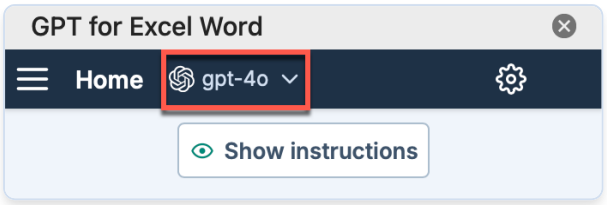
You can follow along using either your own Word document or our summarization example template.

1 On the ribbon, select **Home > GPT for Excel Word**.

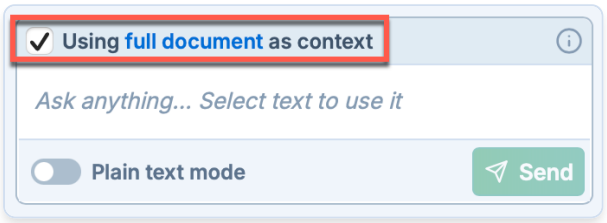


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

2. Select the model of your choice.

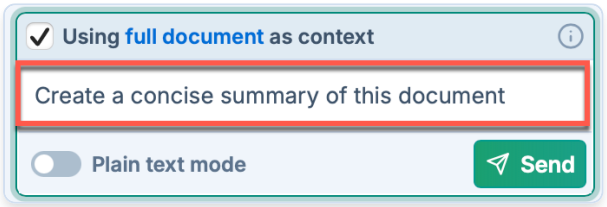


3. Select **Use full document as context**.

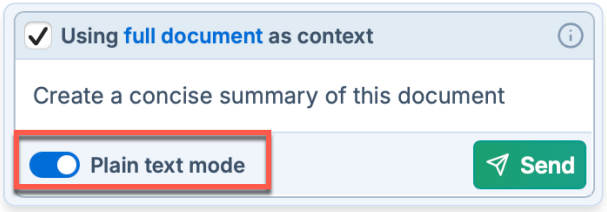


4. Type your prompt:

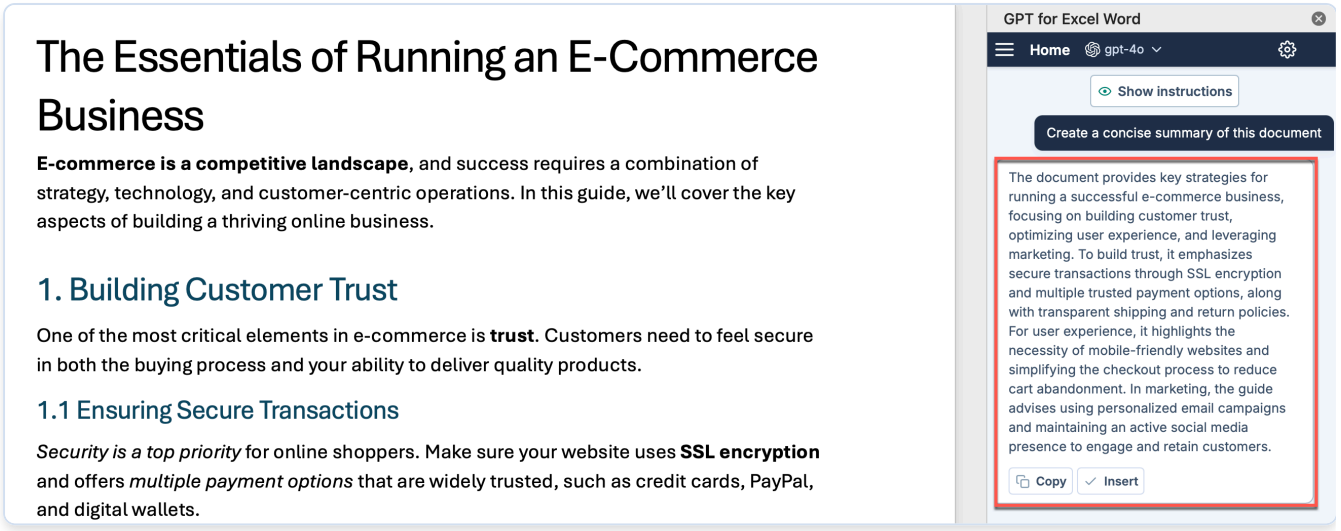
```
Create a concise summary of this document
```



5. Make sure plain text mode is enabled.



6 Click **Send**. GPT for Word generates a response.

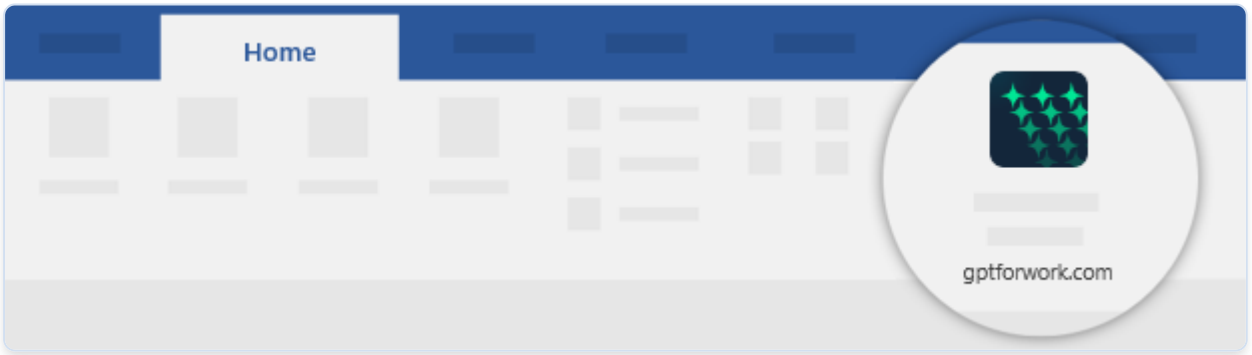


GPT for Word has generated a one-paragraph summary of the document. You can now **Copy** or **Insert** the summary.

Create a job description draft

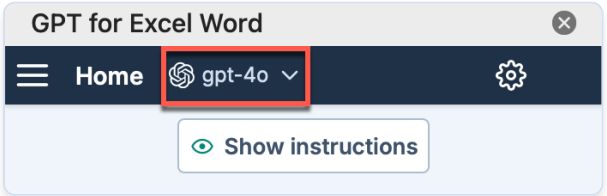
You can follow along using your own Word document.

1 On the ribbon, select **Home > GPT for Excel Word**.

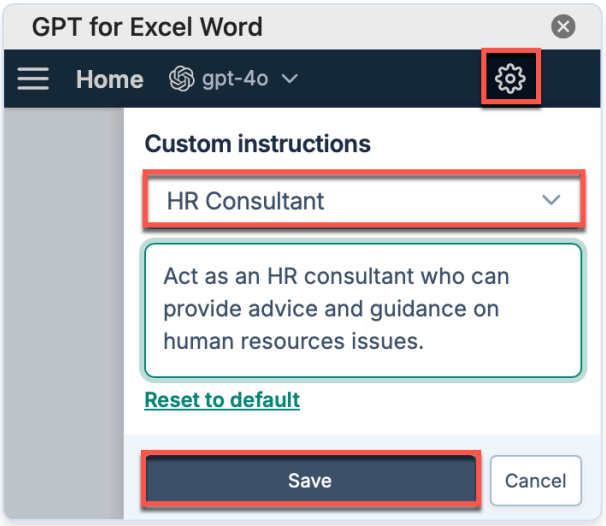


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

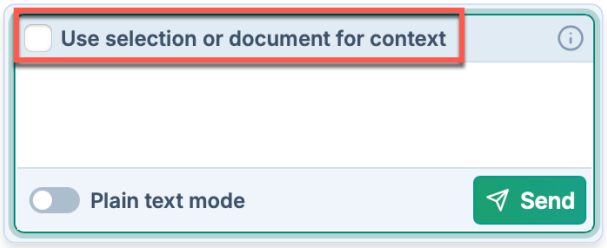
2 Select the model of your choice.



3. Define custom instructions to improve the AI response:
    1. Click the settings icon.
    2. Choose the **HR Consultant** pre-defined instruction.
    3. Click **Save**.

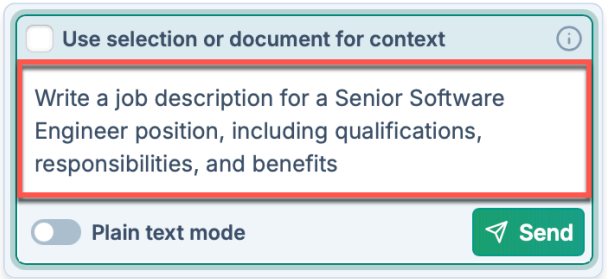


4. (Optional) Deselect **Use selection or document for context**. Since you're starting from a blank document, there's nothing to use as context.

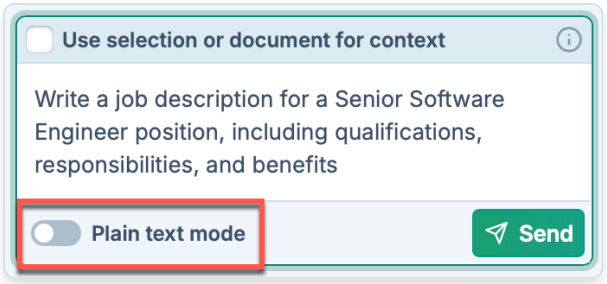


5. Type your prompt:

```
Write a job description for a Senior Software Engineer position, including qualifications, responsibilities, and benefits
```



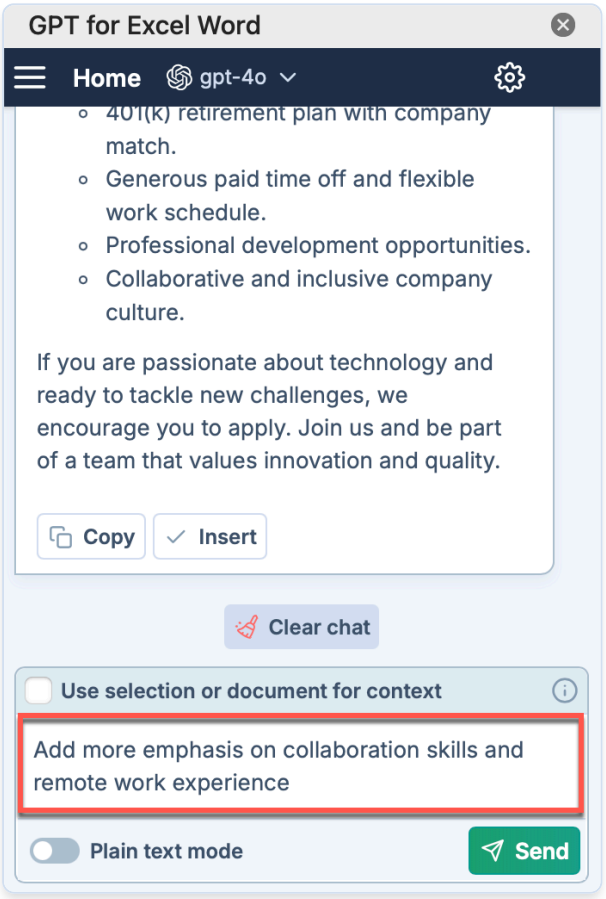
6. Make sure plain text mode is disabled.



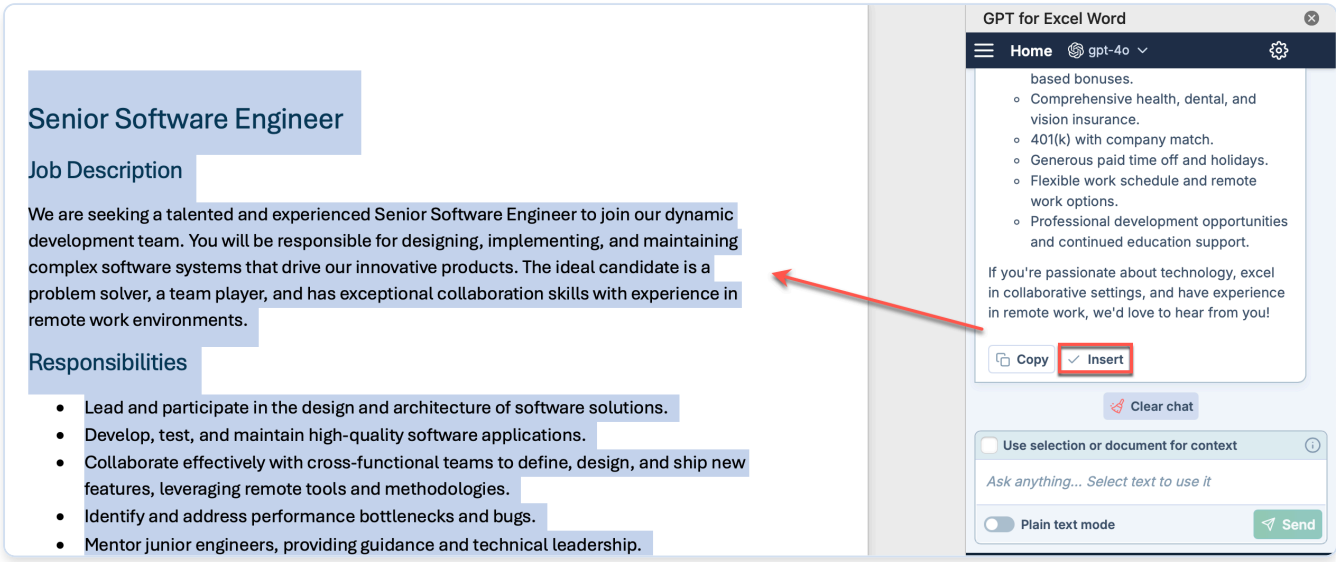
7. Click **Send**. GPT for Word generates a response.
8. Refine the AI-generated response:
    1. Type the prompt:

```
Add more emphasis on collaboration skills and remote work experience
```

    2. Click **Send**. GPT for Word updates the content.



9 To add the job description at the cursor position, click **Insert**.

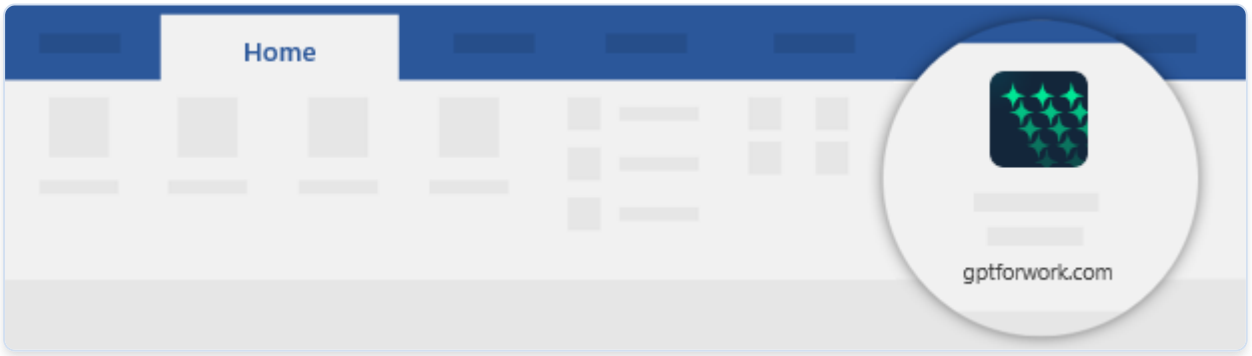


GPT for Word has inserted the AI-generated job description into the Word document.

Analyze a text selection

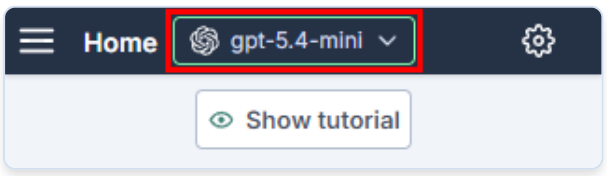
You can follow along using either your own Word document or our product review template.

1 On the ribbon, select **Home > GPT for Excel Word**.



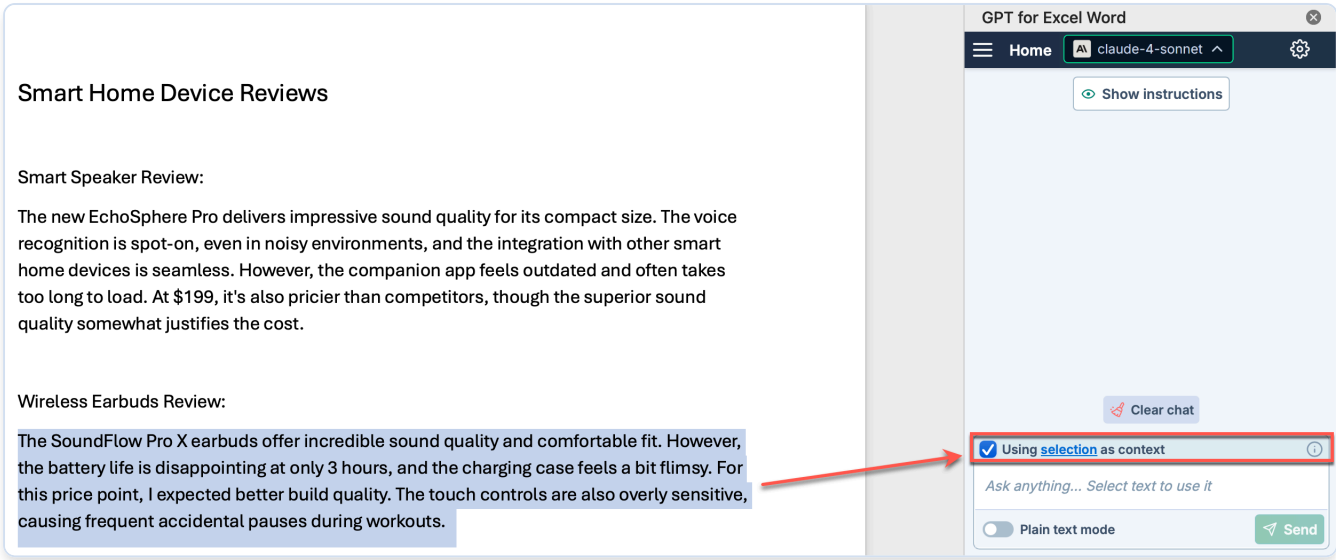
> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

2 Select the model of your choice.



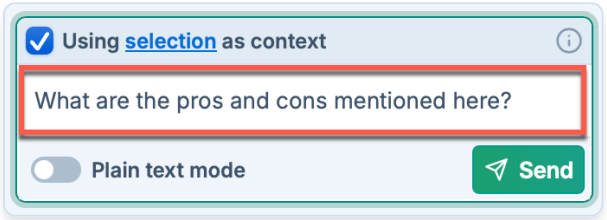
3 In the document, select the text you want to analyze and, in the sidebar, select **Use selection**

as context.

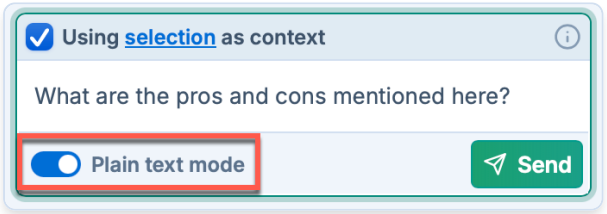


4 Type your prompt:

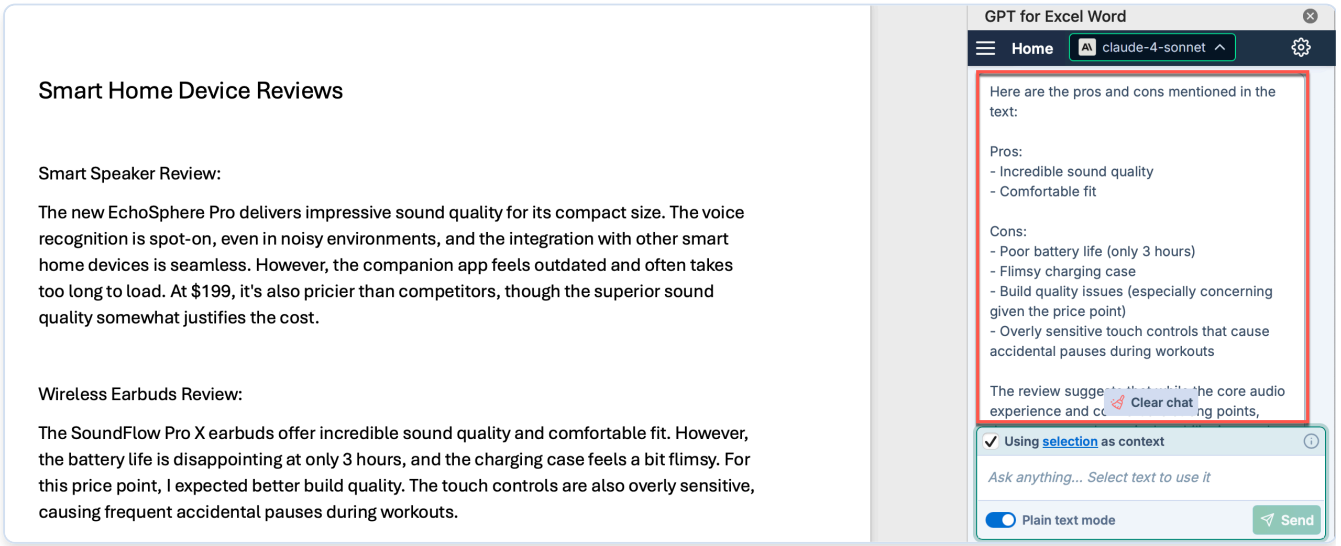
```
What are the pros and cons mentioned here?
```



5 Make sure plain text mode is enabled.



6 Click **Send**. GPT for Word generates a response.



GPT for Word has generated an analysis of the selected text. You can now **Copy** or **Insert** the analysis.

Configure GPT for Word

Select the AI model for GPT for Word and configure how the model behaves.

Select a model

Select the model used by GPT for Word.

Configure AI behavior

Configure model settings to customize how GPT for Word generates responses.

Select a model in GPT for Word

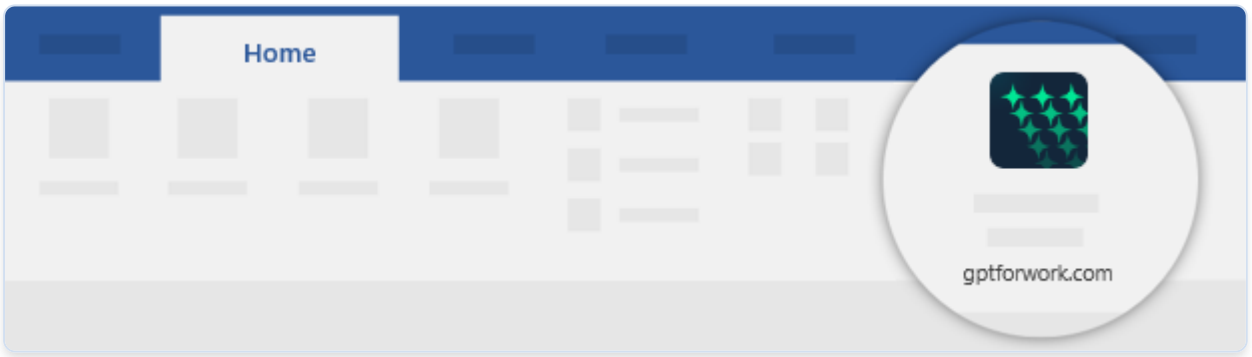
GPT for Word supports a range of models that you can use with your AI work in Word.

The available models depend on:

- **The API keys set for your current space.** GPT for Word gives you access to a default set of built-in models. To access additional models, set an API key for one or more AI providers.
- **The API endpoints set up for your current space.** To access even more models, set up a connection to an Azure OpenAI API endpoint or any OpenAI-compatible API endpoint.
- **The model settings for your current space.** Space owners and admins can control which models are available in GPT for Word.

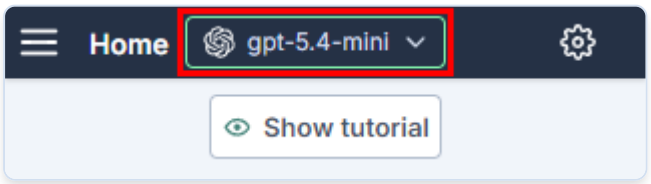
Select a built-in model

1. Open a Microsoft Word document.
2. On the ribbon, select **Home > GPT for Excel Word**.

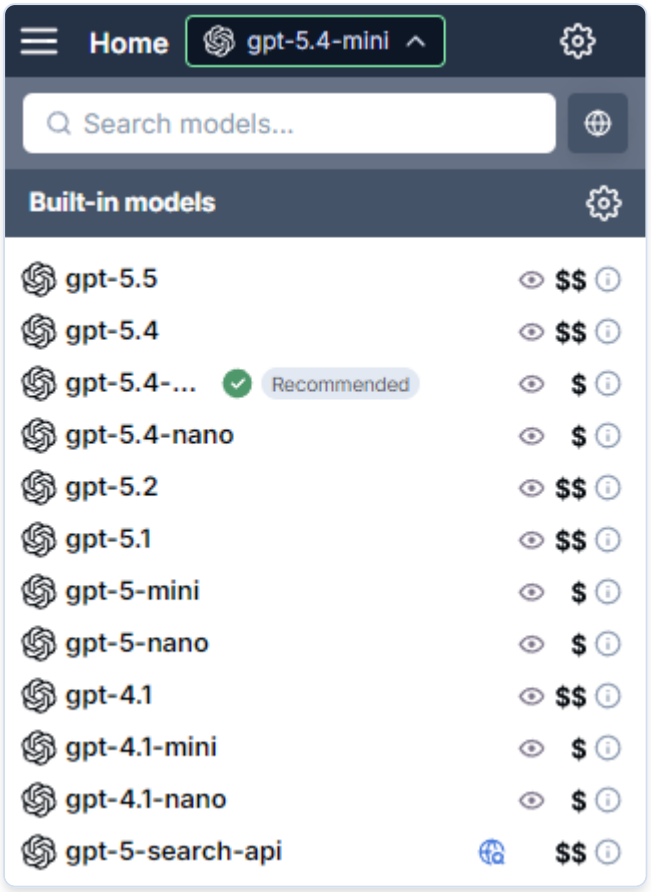


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. Open the model switcher.



4 Under **Built-in models**, select a model.

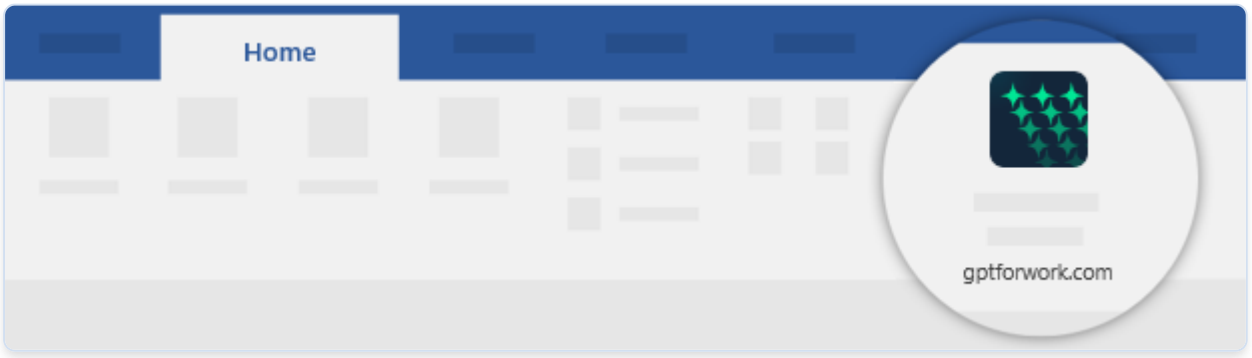


GPT for Word now uses the selected model to generate responses.

Select a connected model

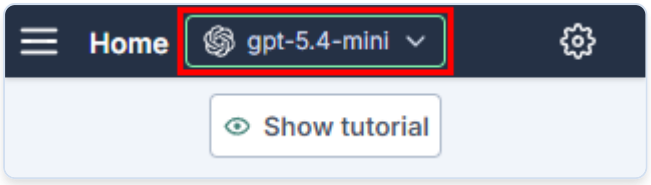
1 Open a Microsoft Word document.

2 On the ribbon, select **Home > GPT for Excel Word**.

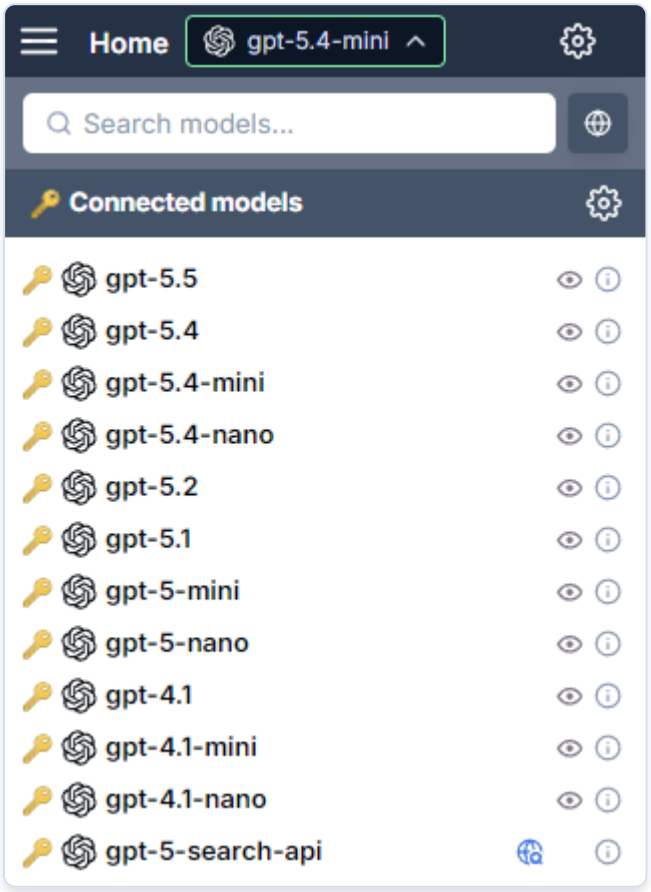


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3 Open the model switcher.



4 Under **Connected models**, select a model.



GPT for Word now uses the selected model to generate responses.

What's next

Learn more about models:

- AI providers and models supported by GPT for Work
- LiveBench (a good model benchmark)

Configure AI behavior in GPT for Word

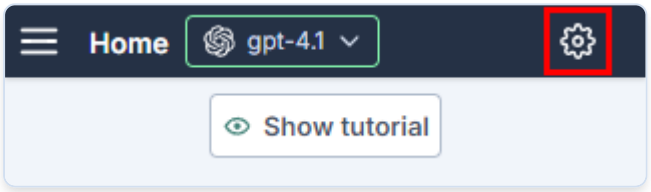
Configure model settings to customize how GPT for Word generates responses.

> **INFO**
> Model settings are specific to each document.

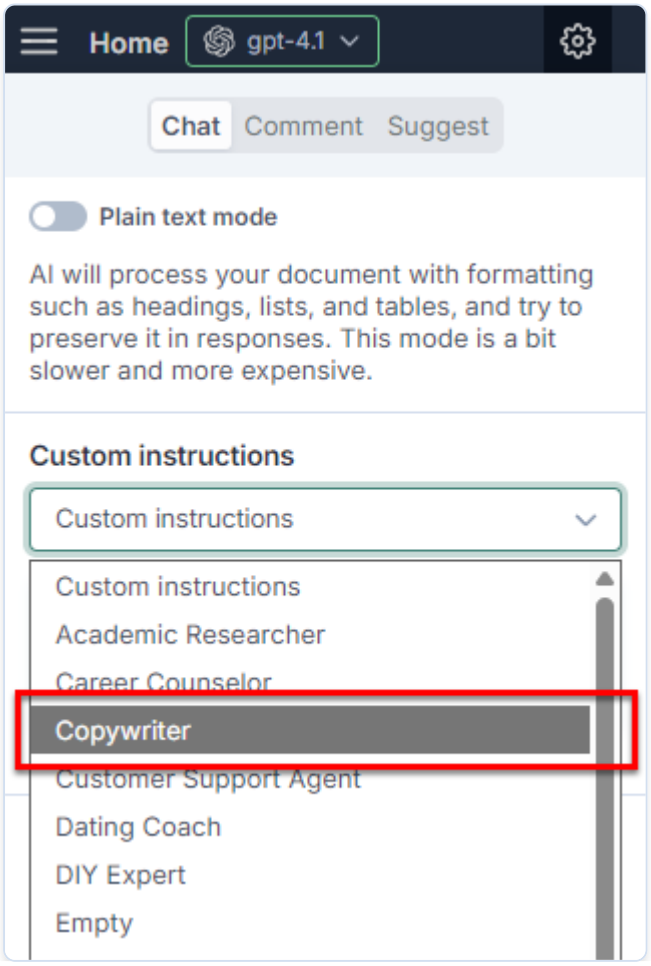
Add custom instructions

Provide custom instructions to specify preferences or requirements that you'd like the AI to consider when generating responses.

1. In the GPT for Word sidebar, click **Settings**.



2. In the **Custom instructions** dropdown, select the type of instructions you'd like to add and edit them if needed.



3 Click **Save**.

The AI takes the custom instructions into account when generating responses in the current document.

> **TIP**
>
> To revert the custom instructions to the default, click **Reset to default**.

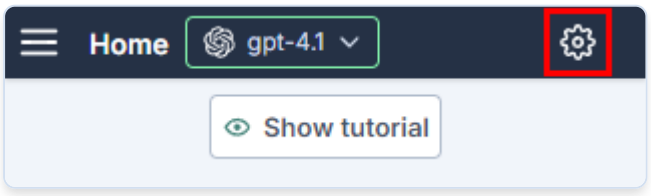
Set the creativity level

Control how creative the AI is by setting the temperature and top-p parameters. The parameters work together to define overall creativity. Lower creativity generates responses that are straightforward and predictable, while higher creativity generates responses that are more random and varied.

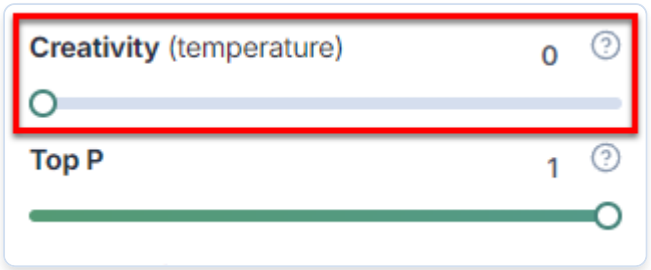
> **TIP**
>
> For more information about temperature as used by OpenAI to set the creativity level, see our temperature guide.

Parameter	Description	How to use
Temperature	Controls randomness in the output.	• Possible values: **0-1** • Keep at **1** if adjusting **Top P** • Lower for factual, higher for creative
Top P	Controls diversity of word choices.	• Possible values: **0-1** • Keep at **1** for most use cases • Only lower if **Temperature** is **1**

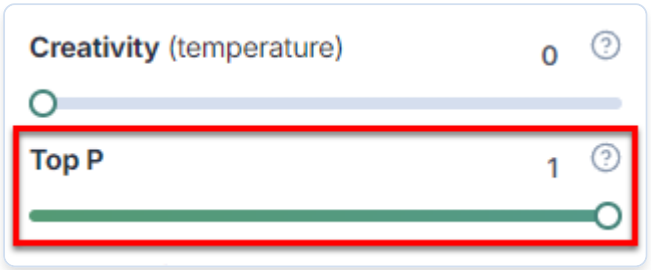
1 In the GPT for Word sidebar, click **Settings**.



2 Use the slider to set **Creativity** between **0** and **1**.



3 (Optional) Use the slider to set **Top P** between **0** and **1**.



4 Click **Save**.

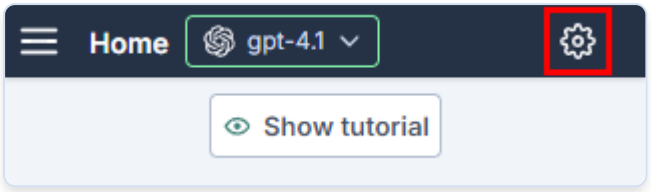
The AI uses the new temperature and top-p values when generating responses in the current document.

Reduce AI response repetition

Set frequency and presence penalties to reduce the tendency of OpenAI and xAI models towards repetition.

Parameter	Description	How to use
Frequency penalty	Controls how much to penalize new tokens based on their frequency in the text so far.	• Possible values: **0-2** • Keep at **0** by default • Higher values reduce repetition
Presence penalty	Controls how much to penalize new tokens based on whether they appear in the text so far.	• Possible values: **0-2** • Keep at **0** by default • Higher values increase diversity

1. In the GPT for Word sidebar, click **Settings**.



2. Use the sliders to set **Frequency penalty** and **Presence penalty** between **0** and **2**.



3. Click **Save**.

The AI uses the new penalty values when generating responses in the current document.

What's next

- Manage text formatting.
- Select the model that best fits your needs.

Admin

Manage users, AI configuration, and billing for your space from the GPT for Work dashboard.

User management

Create a team in your space, and add and remove users from the space. If needed, promote users to admins or transfer ownership of your space.

AI configuration

Manage the models available to users in your space: Set which Agent and bulk models are available, and configure API keys and endpoints.

Billing

Manage your subscription, buy credits, set spending limits, and download invoices for your space.

User management

Create a team in your space, and add and remove users from the space. If needed, promote users to admins or transfer ownership of your space.

Create a team

Create a team in the GPT for Work dashboard so that all team members can use the add-ons with shared billing as well as shared API keys and endpoints.

Add and remove users

Add and remove users from your GPT for Work space to ensure only the right people have access to the space and its resources.

Promote users to admins

Promote regular users to admins in your GPT for Work space to give them access to all the dashboard features.

Transfer space ownership

Transfer the ownership of your GPT for Work space to any other user in the space. The new owner takes full control of the space, while you remain in the space as an admin.

Create a team

You can create a team in GPT for Work so that all team members can use the add-ons with shared billing as well as shared API keys and endpoints.

To set up your team:

1. Create a GPT for Work account (if you don't already have one).
2. Set up your space.
3. Invite users to your space.

Create a GPT for Work account

> **INFO**
>
> If you already have a GPT for Work account, skip to Set up your space.

Make sure you sign in with an account for the platform you use.

Microsoft

1. Open the GPT for Work dashboard.
2. In the GPT for Excel and Word section, click **Sign in with Microsoft** or **Sign in with Google**.
3. Choose the account you want to associate with the GPT for Work account.
4. Click **Continue** to sign in.

You have created the GPT for Work account. The account comes with a space that you own.

Set up your space

Before inviting users, set up your space to provide access to the necessary models. Availability of API keys and API endpoints depends on your space's pricing model and plan.

Subscription

- Set up the model switcher for users.
- Manage API keys for your space. Available on Business and Enterprise plans.
- Manage API endpoints for your space. Available on Business and Enterprise plans in GPT for Excel and Word.

Pay-as-you-go

- Set up the model switcher for users.
- Manage API keys for your space.
- Manage API endpoints for your space. Available in GPT for Excel and Word.

Invite users to your space

Prerequisites

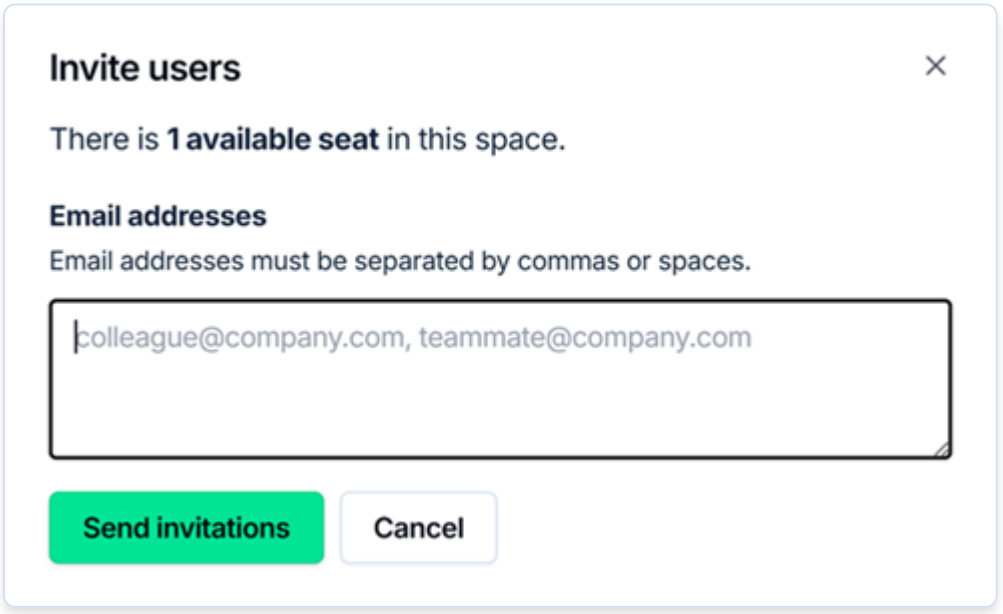
You are the owner or an admin of the space.

How large your team can be depends on your space's pricing model and plan.

Subscription

Team management is available on Business and Enterprise plans. The size of your team is limited by the number of seats you have, and each user you invite takes up one seat. If you run out of seats, add seats before inviting more users.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Users**. The page header shows how many of your seats are in use.
3. Click **Invite users**.
4. Enter the email addresses of the users you want to invite, separated by commas or spaces.



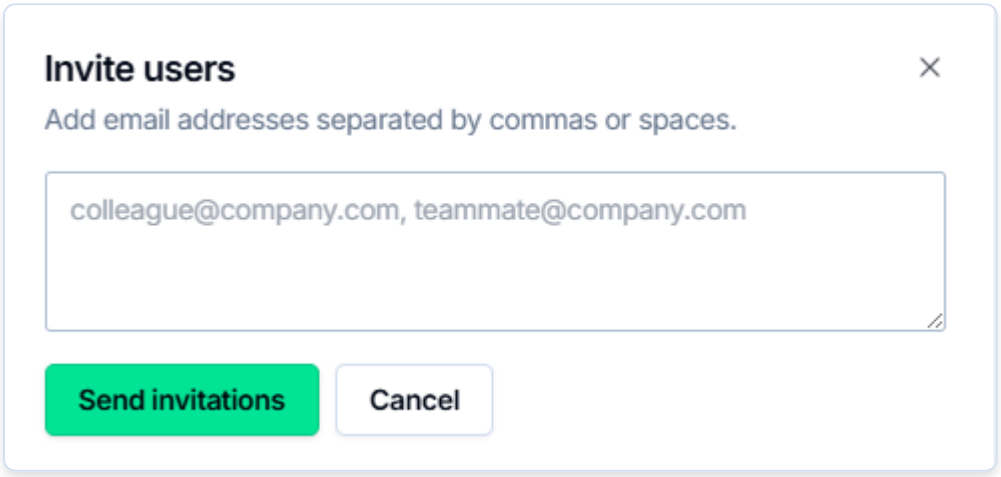
5. Click **Send invitations**.

You have successfully invited users to your space. You can resend or cancel an invitation while its status remains **Pending**. A pending invitation takes up a seat until the user accepts it or you cancel it.

Pay-as-you-go

There are no seats in a pay-as-you-go space, so you can invite as many users as you need. Every user draws on the same credit balance.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Users**.
3. Click **Invite users**.
4. Enter the email addresses of the users you want to invite, separated by commas or spaces.



5. Click **Send invitations**.

You have successfully invited users to your space. You can resend or cancel an invitation while its status remains **Pending**.

What's next

- Install the add-on or add-in on your domain, or have your team members install it individually:
  - GPT for Sheets and Docs
  - GPT for Excel Word
- Set spending limits for your space and for individual users, on subscription or pay-as-you-go.
- Manage users in your team.

Add and remove users

As your team grows and changes, you can add and remove users from your GPT for Work space to ensure only the right people have access to the space and its resources.

Prerequisites

You are the owner or an admin of the space.

Add users to your team

You can add users to your team at any time by inviting them to your space.

INFO

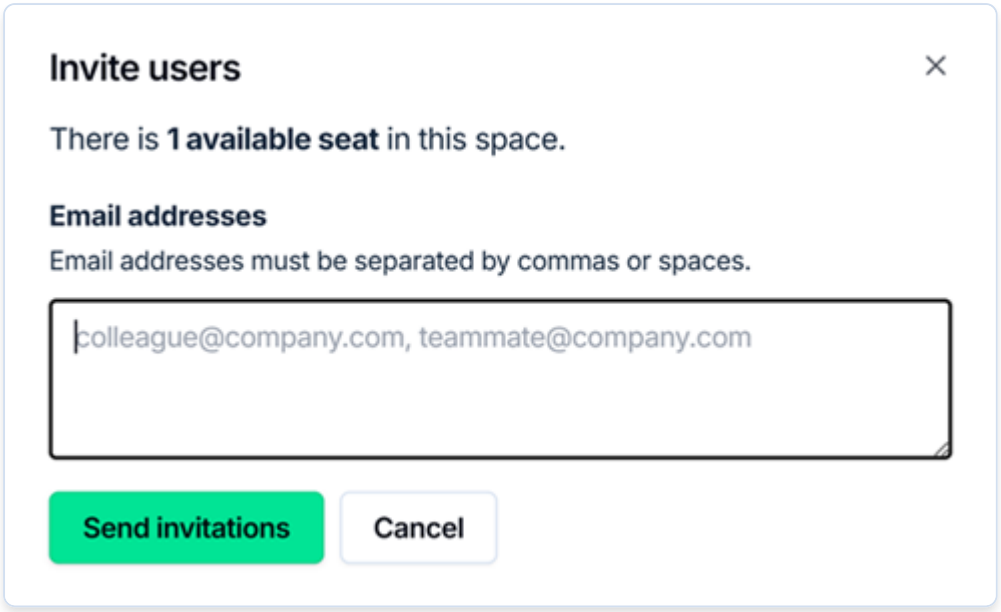
When you invite a user to your space, they use your space's pricing model and lose access to their own personal space for as long as they're on your team.

Subscription

Team management is available on Business and Enterprise plans. Each user in your space takes up a seat, so you need an available seat for every user you invite. If your space has no available seats, add seats first.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Users**. The page header shows how many of your seats are in use.
3. Click **Invite users**.

4 Enter the email addresses of the users you want to invite, separated by commas or spaces.



5 Click **Send invitations**.

You have successfully invited users to your space. You can resend or cancel an invitation while its status remains **Pending**. A pending invitation takes up a seat until the user accepts it or you cancel it.

Seat assignment for new users

Every seat comes with its own included usage, which resets at the start of each billing cycle. When a user leaves your space, their seat becomes empty but keeps its remaining included usage for the current cycle.

When you invite a new user, they are assigned the empty seat with the most included usage remaining. This gives the new user as much usage as possible for the rest of the cycle.

For example, your space is on a Business plan with 3 seats, each with $25 of included usage:

- Seat 1 is taken by an active user.
- Seat 2 is empty, with $4 of its included usage remaining.
- Seat 3 is empty, with $19 of its included usage remaining.

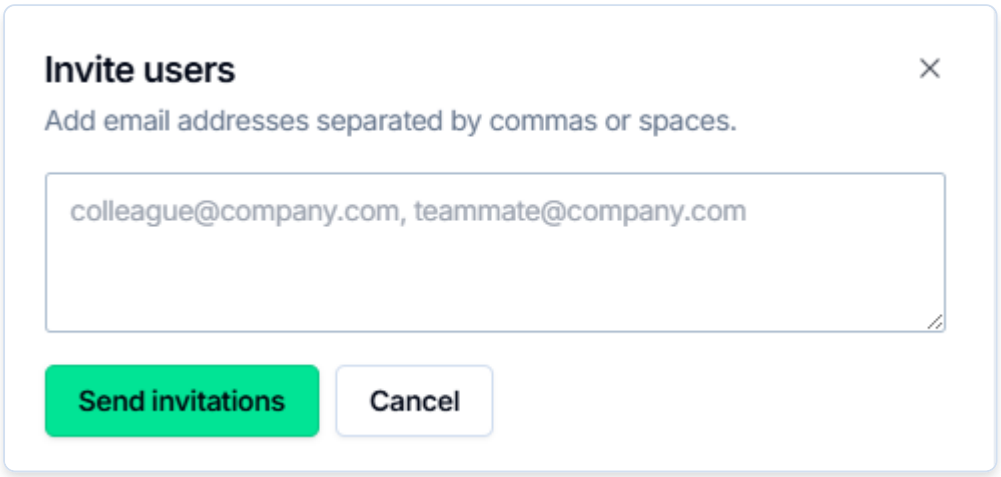
If you invite a new user, they are assigned seat 3 and can use its remaining $19 until the end of the billing cycle. At the next renewal, all seats are reset to $25.

Pay-as-you-go

You can invite as many users as you need. Every user draws on the same credit balance. To keep usage under control, set spending limits for your space and for individual users.

1 Sign in to the GPT for Work dashboard with the account that you use for your space.

2. In the sidebar, select **Users**.
3. Click **Invite users**.
4. Enter the email addresses of the users you want to invite, separated by commas or spaces.



5. Click **Send invitations**.

You have successfully invited users to your space. You can resend or cancel an invitation while its status remains **Pending**.

Remove users from your team

You can remove users from your team at any time by removing them from your space. When you remove a user from your space, they return to their own personal space and its pricing model. They can still use the GPT for Work add-ons, but they no longer have access to your space's credits, API keys, API endpoints, or other setup.

> **INFO**
>
> The space owner cannot be removed. If you want to remove the space owner from the team, you must transfer ownership to another user first.

Subscription

Removing a user frees up their seat, but it does not change your seat count or your invoice. To pay for fewer seats, remove the unused seats as well.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Users**.
3. In the users list, click the three horizontal dots (**...**) for the user you want to remove and select **Remove user**. If the user has yet to accept their invitation to the space, select **Cancel invitation** instead.

You have successfully removed the user from the space. When they next sign in to the dashboard, they only see their personal space, not your team space. The seat they occupied is now empty and available for a new user.

Pay-as-you-go

Removing a user does not affect your credit balance. The credits the user has already consumed are not refunded, and the remaining credits stay available to the rest of your team.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Users**.
3. In the users list, click the three horizontal dots (**...**) for the user you want to remove and select **Remove user**. If the user has yet to accept their invitation to the space, select **Cancel invitation** instead.

You have successfully removed the user from the space. When they next sign in to the dashboard, they only see their personal space, not your team space.

What's next

- Promote a user to an admin.
- Transfer the ownership of your space.

Promote users to admins

As a space owner or admin, you can promote regular users to admins in your space. Admins have access to all the dashboard features and can do everything the owner can do, except transfer the ownership of the space.

Prerequisites

You are the owner or an admin of the space.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Users**.
3. In the users list, click the three horizontal dots (**...**) for the user you want to promote and select **Set as admin**.

You have successfully promoted the user to admin. They can now access all the dashboard features.

INFO

You can revoke admin permissions at any time by clicking the three horizontal dots (**...**) in the admin user's row and selecting **Remove admin role**.

What's next

- Remove a user from your space.
- Transfer the ownership of your space.

Transfer space ownership

You can transfer the ownership of your space to any other user in the space. After the transfer, the new owner takes full control of the space, while you remain in the space as an admin.

Prerequisites

You are the owner of the space.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Settings**.
3. In the **Transfer ownership** section, select the user to whom you want to transfer the ownership.
4. Click **Save** to confirm.

You have successfully transferred the ownership of the space to the user. You remain in the space as an admin.

What's next

- Manage users in your team.
- Manage the models available to users in your space.

AI configuration

Manage the models available to users in your space: Set which Agent and bulk models are available, and configure API keys and endpoints to provide access to more models.

Manage models

Configure which AI models are available to users in your team space from the GPT for Work dashboard.

Manage API keys

Set API keys for your space to give your users access to more models.

Manage API endpoints

Set up API endpoints for your space to give your users access to more models.

Manage models

As a space owner or admin, you can control which AI models are available to the users of your space. This allows you to customize model access for your team and control costs.

> **INFO**
>
> These settings apply to GPT for Sheets, GPT for Excel, and GPT for Word. You cannot enable and disable models or otherwise control model access for GPT for Docs (apart from setting API keys).

Prerequisites

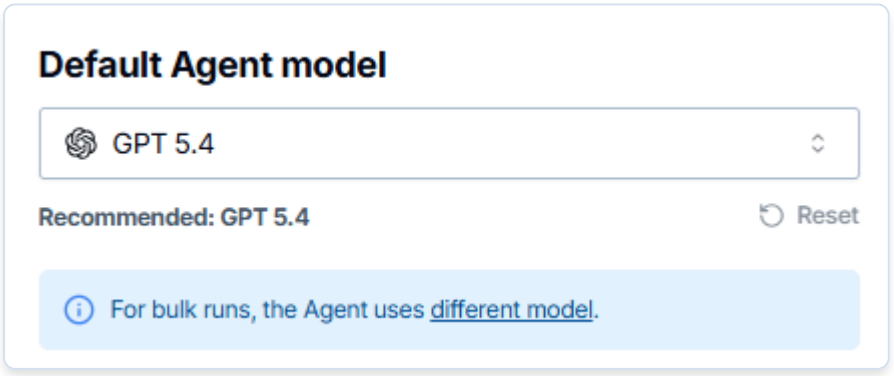
You are the owner or an admin of the space.

Set the default Agent model

You can set the default Agent model for all users in your space. You can only select enabled Agent models.

Microsoft

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Agent models**.
3. In the **Default Agent model** dropdown, select the default Agent model for your space.



For information about recommended Agent models, see Recommended models.

> **INFO**
>
> To select a provider-connected model for the Agent, you must first set an API key for your space.
>
> To select an endpoint-connected model for the Agent, you must first set up an OpenAI-compatible API endpoint for your space.

The Agent now uses the selected model for new users in the space. Users can override this default from the model switcher.

Set the default bulk models

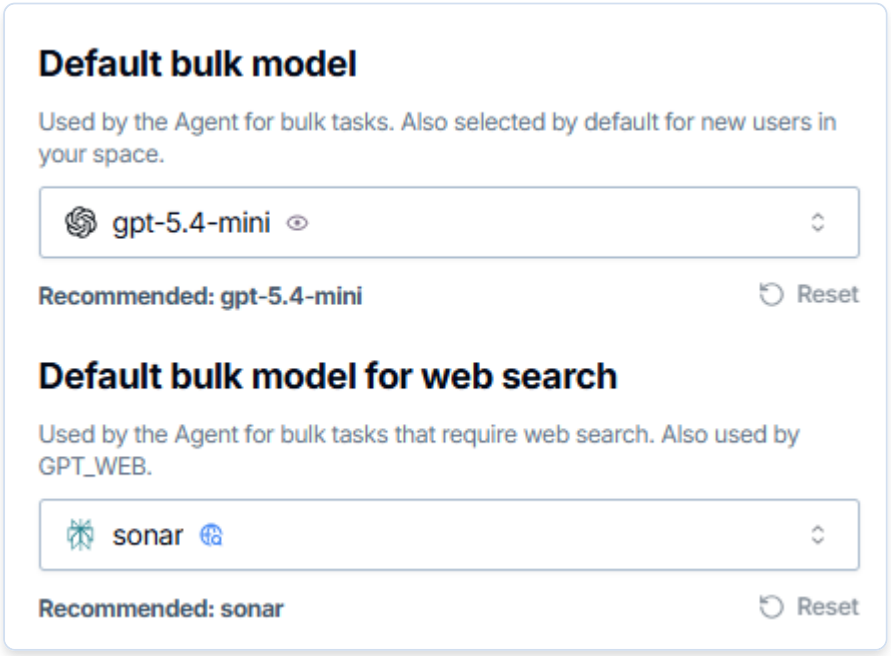
You can set two default models for bulk tasks:

- The **Default bulk model** is used by the Agent for bulk tasks. For new users in your space, this model is also selected by default for bulk tools and GPT functions.
- The **Default bulk model for web search** is used by the Agent for bulk tasks that require web search. This model is also used by the **GPT_WEB** function.

You can only select enabled bulk models.

Microsoft

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Bulk models**.
3. In the **Default bulk model** and **Default bulk model for web search** dropdowns, select the default bulk models for your space.



> For information about recommended bulk models, see Recommended models.
>
> For information about recommended web search models, see Search the web with bulk tools in Excel.

The Agent now uses the selected models for bulk tasks and web searches for all users in the space. Bulk tools and GPT functions use the selected bulk model for new users in the space; users can override this default from the model switcher.

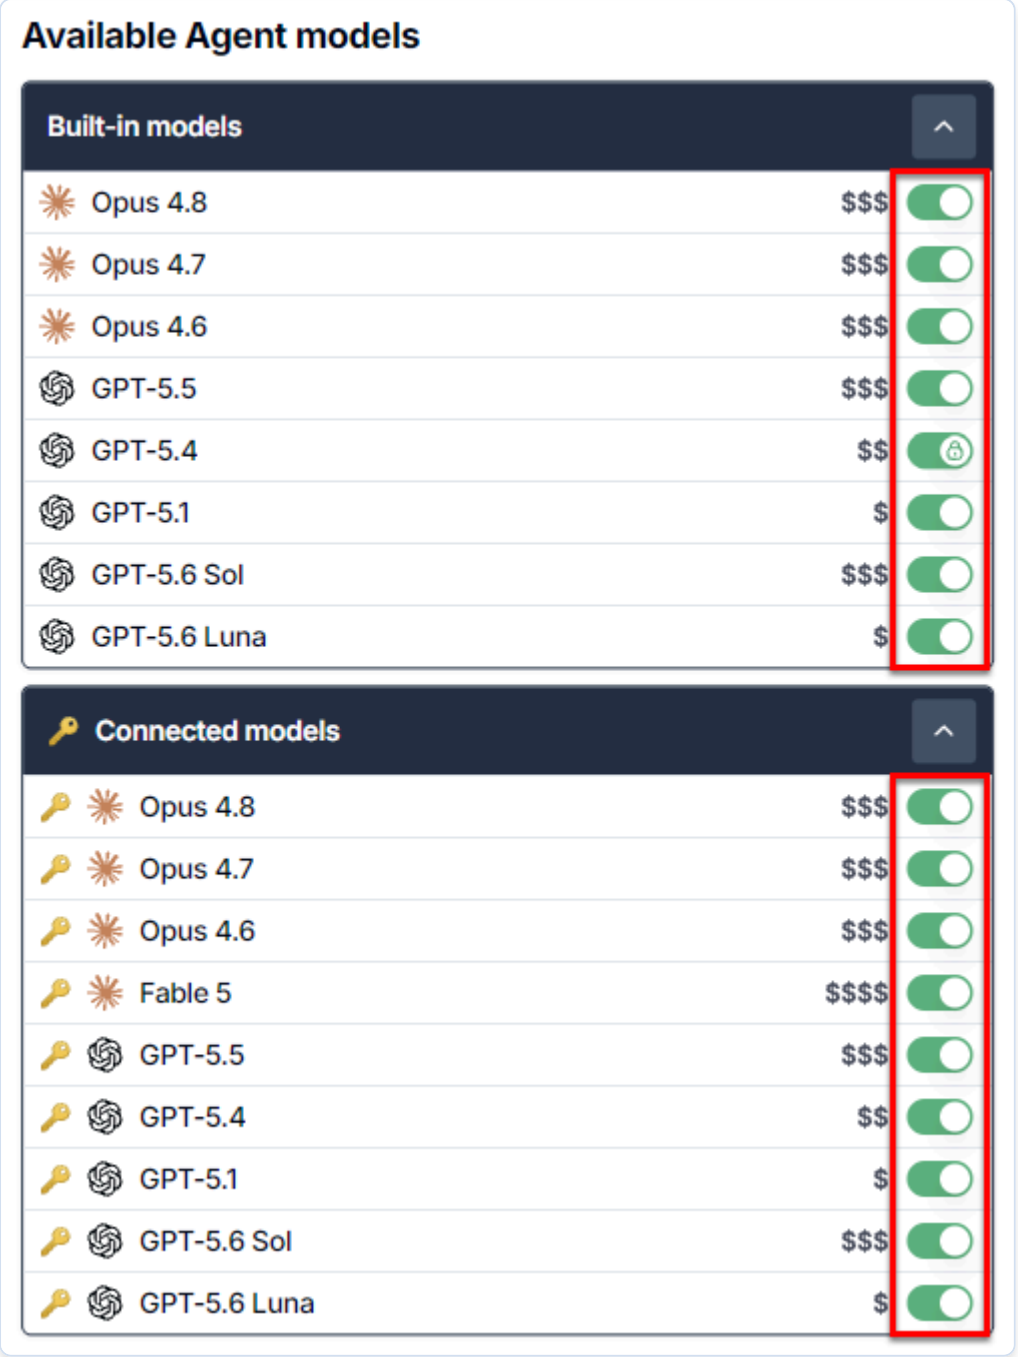
Enable and disable models

You can enable and disable individual models for your space. This includes both Agent models and bulk models. Users can only select enabled models from the model switcher. Disabled models are hidden from the model switcher.

Microsoft

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Agent models** or **Bulk models**.

3 In the models list, use the toggle next to a model to enable or disable the model.



What's next

- Set spending limits for your space and for individual users.
- Invite users to your space.
- Manage users in your team.

Manage API keys

As a space owner or admin, you can set API keys for your space. Models available with the keys will be accessible to all users in the space.

Using API keys offers several benefits:

- **More models.** Access a wider range of models from a bigger pool of AI providers than what is available without an API key.
- **More control.** Monitor and manage AI usage and costs directly on the AI provider's platform.
- **Improved privacy.** We do not log inputs and outputs when you use a connected model. For more information, see our security and privacy FAQ.
- **Reduced costs.** If you're on a pay-as-you-go space, using an API key reduces costs. You pay the model usage fees directly to the AI provider, while paying a reduced service fee for GPT for Work.

Create API keys

> **INFO**
>
> Make sure your keys belong to an organization or project with the following rate limits:
>
> - RPM (requests per minute) ≥ 1,000
> - TPM (tokens per minute) ≥ 1,000,000
>
> The actual limits required depend on your workload requirements.

To create an API key, go to the AI provider's platform and follow their instructions:

- OpenAI Platform for GPT models and fine-tuned models
- Anthropic Console for Claude models
- DeepSeek Platform for DeepSeek models
- Google AI Studio for Gemini models
- Mistral AI for Mistral models
- Perplexity for Perplexity models
- xAI Cloud Console for xAI models

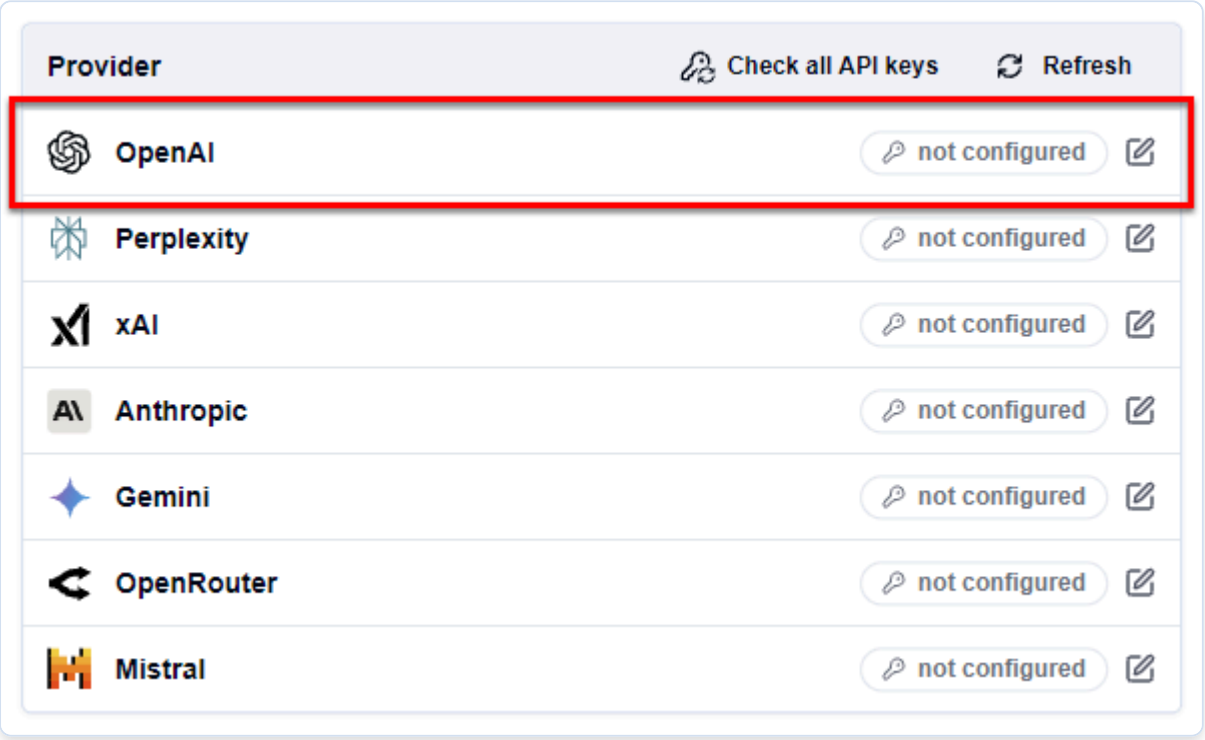
Set API keys

Prerequisites

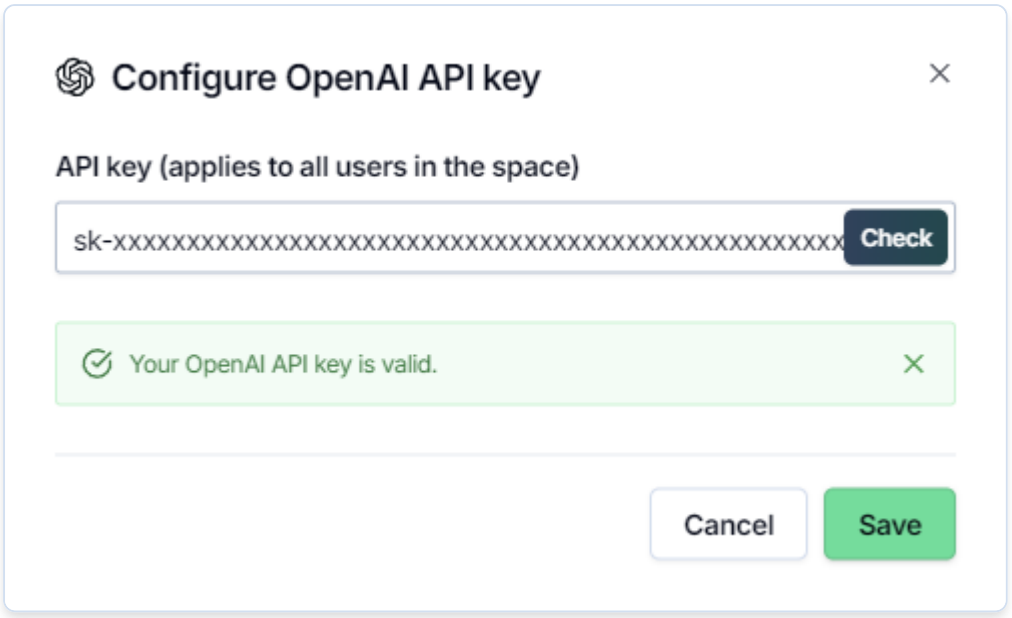
You are the owner or an admin of the space.

Microsoft

1. Create an API key or request an API key from your AI provider platform admin. Learn more.
2. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
3. In the sidebar, select **AI provider keys**.
4. In the AI providers list, click the provider whose API key you want to set.



5. Set the API key:
   1. In the text field, enter the key.
   2. Click **Check** to verify that the key is valid.
   3. Click **Save**.



You have set the API key. All users in the space can now use any model available with the key in GPT for Excel and GPT for Word.

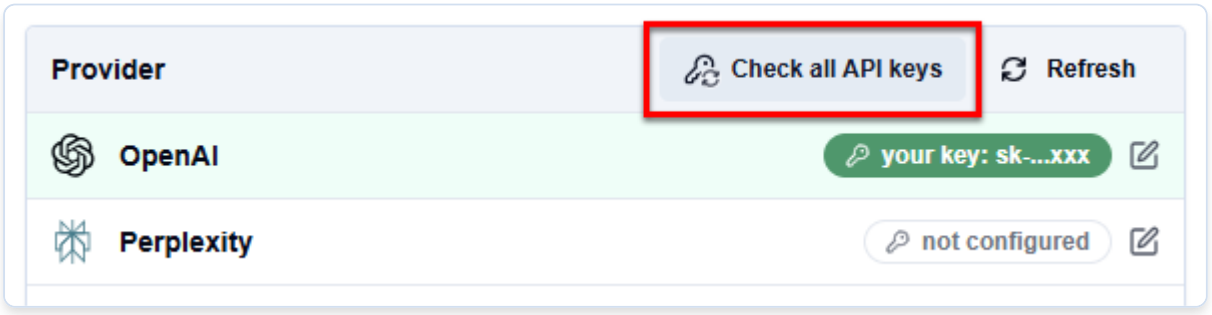
Check API keys

You can verify whether the API keys you have set are valid.

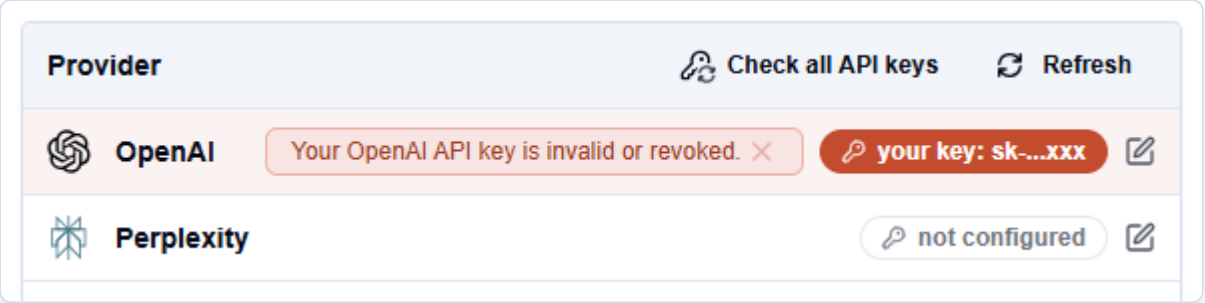
Check all API keys at once

Microsoft

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **AI provider keys**.
3. Click **Check all API keys**.



The page displays a warning for each key that is invalid.

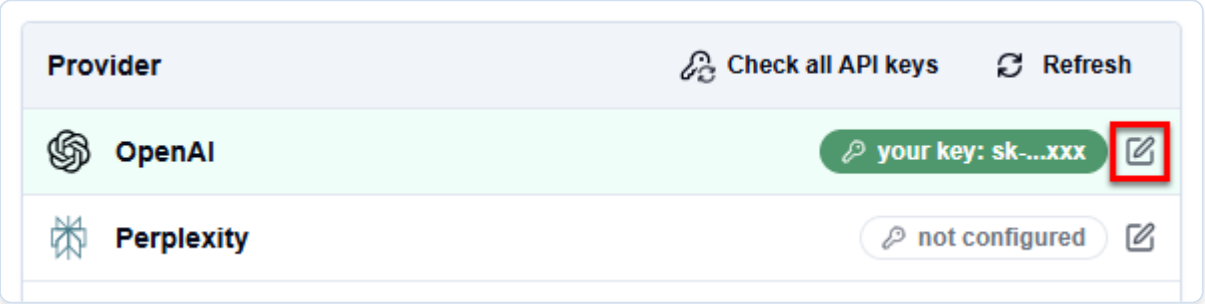


If a key is invalid, create a new key and replace the invalid value with the new one.

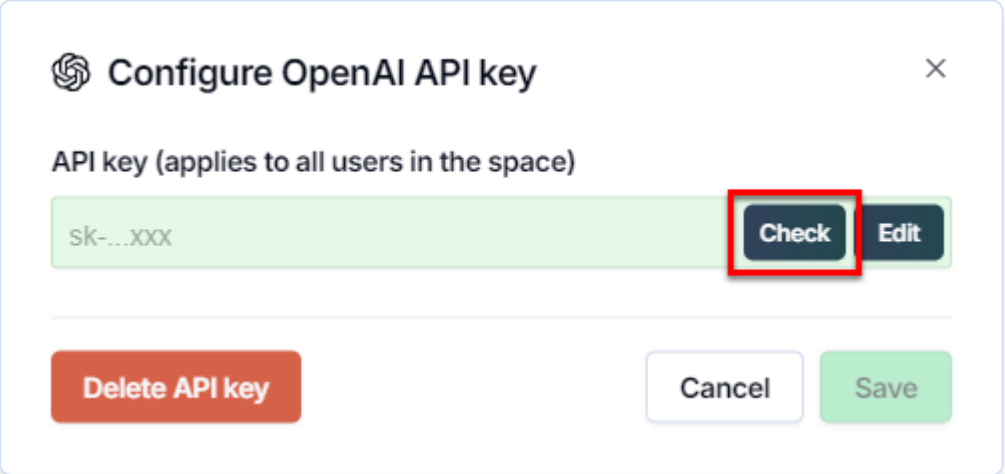
Check a specific API key

Microsoft

1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **AI provider keys**.
3. Click the edit icon for the key you want to check.

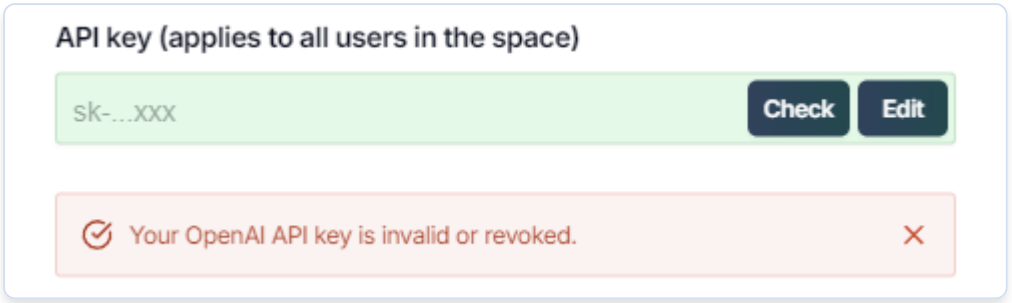


4. Click **Check**.



The UI displays a message indicating whether the key is valid or invalid.





If the key is invalid, create a new key and replace the invalid value with the new one.

What's next

Microsoft

- Invite users to your space.
- Manage API endpoints for your space.
- Manage models for your space.
- Manage users in your team.

Manage API endpoints

As a space owner or admin, you can set up API endpoints for your space. Models available through the endpoints will be accessible to all users in the space

> **INFO**
>
> API endpoints are available only for GPT for Excel and Word spaces.

Connect to Azure OpenAI

GPT for Excel and GPT for Word can connect to Azure AI to use OpenAI models. In Azure AI, each model has a dedicated API endpoint that you use to access the model.

Prerequisites

- You are the owner or an admin of the space.
- Your space is on the Business or Enterprise plan or pay-as-you-go pricing.
- You have an Azure account with Azure AI service access.
- You have deployed a model in a project.

To connect to an Azure OpenAI API endpoint:

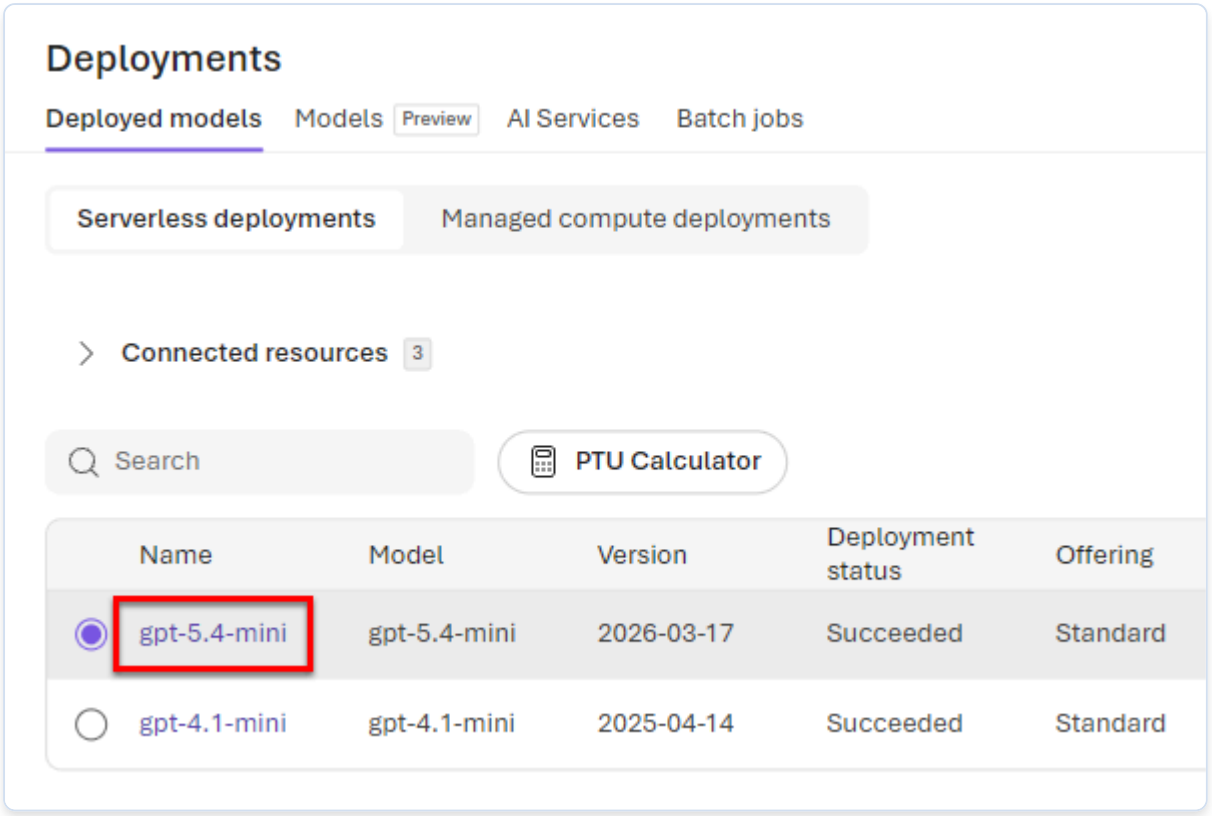
1. Get the Azure OpenAI API endpoint details.
2. Set up the Azure OpenAI API endpoint.

Get the Azure OpenAI API endpoint details

To connect to Azure AI from GPT for Excel and Word, you need the endpoint URL, API key, deployment name, and model name for the Azure OpenAI model you want to use.

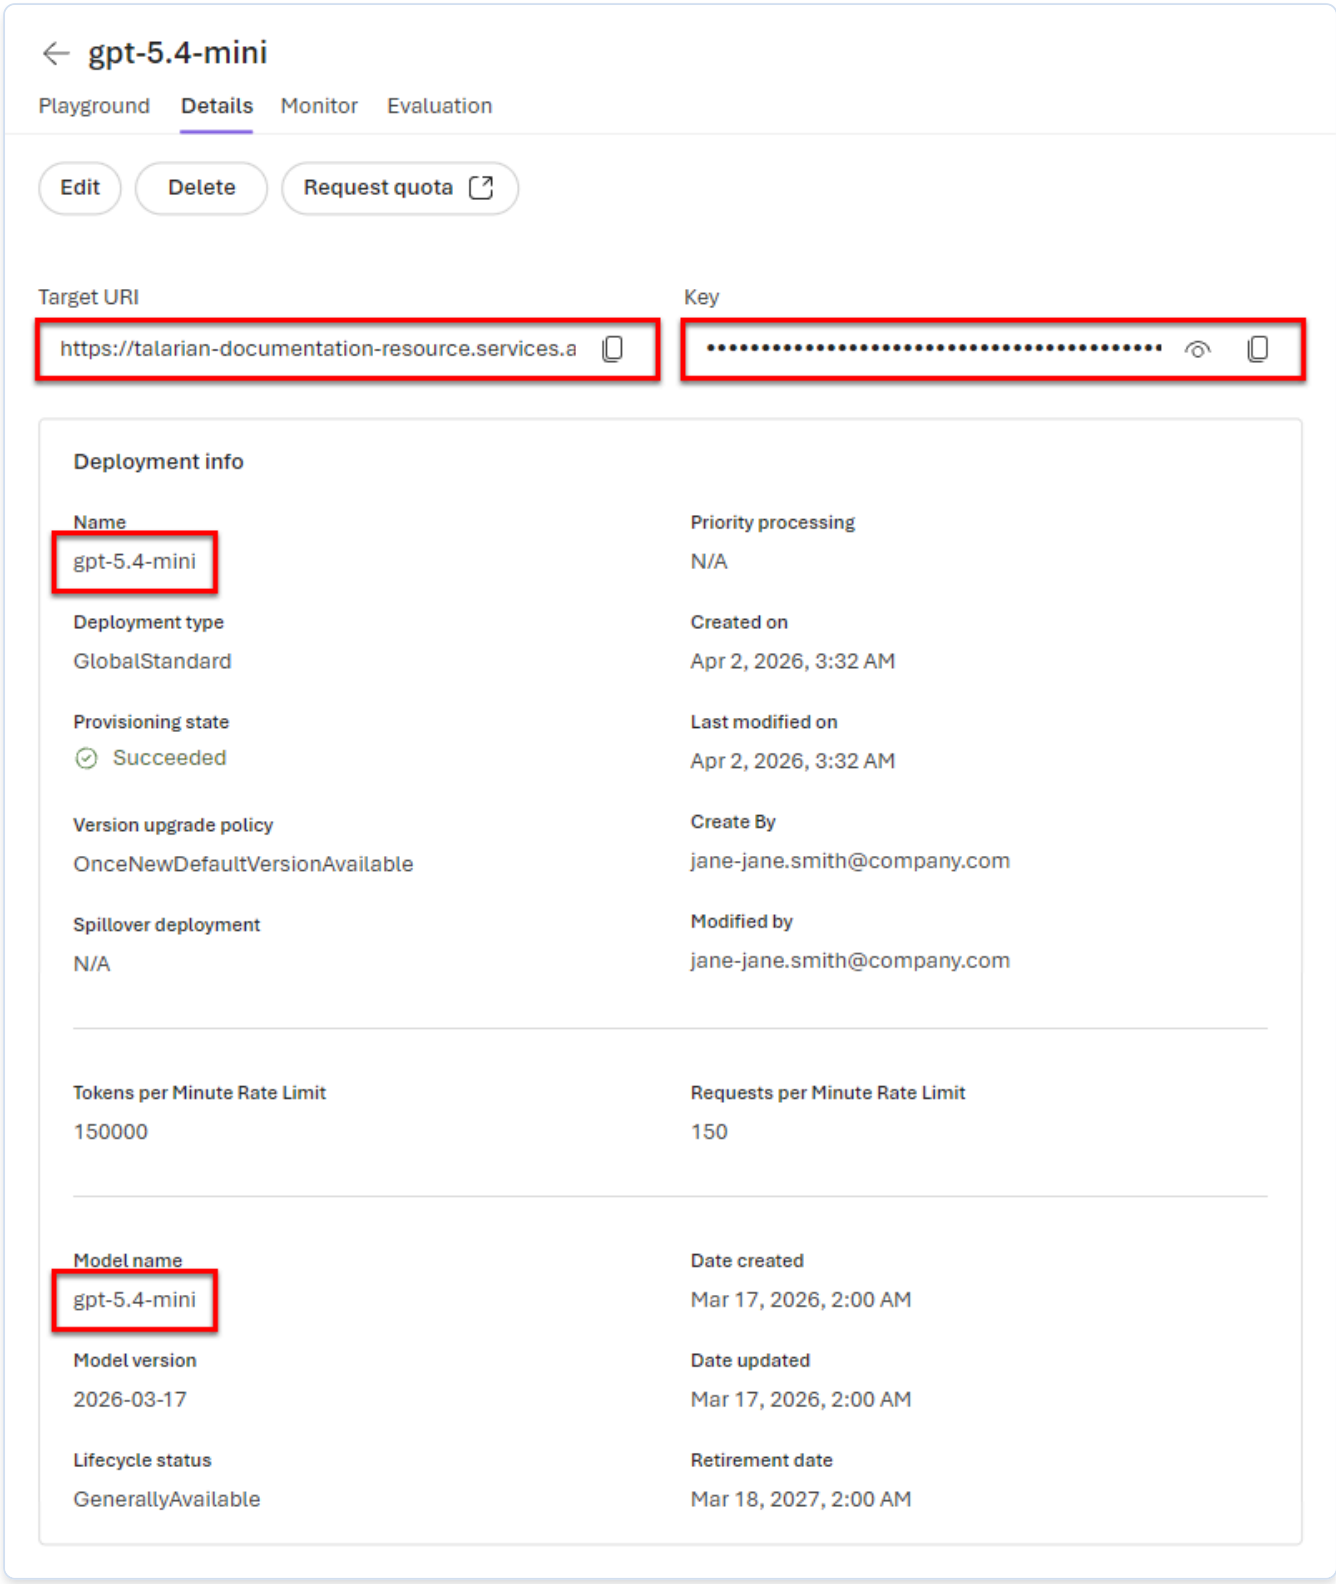
1. Sign in to Microsoft Foundry.
2. Select the project that has the model you want to use.
3. Click **View deployments**.

4 In **Deployed models**, in the models list, click the name of the model you want to use.



5 Select **Details**.

6 In the model details, copy the **Target URI**, **Key**, **Deployment name**, and **Model name**.



Set up the Azure OpenAI API endpoint

1 Sign in to the GPT for Work dashboard with your Microsoft or Google account.

2 In the sidebar, select **Custom API endpoints**.

3 Select the **Bulk models** tab.

4 Click **Azure OpenAI endpoints (Chat Completions API)**.

> **TIP**
>
> If you already have Azure OpenAI API endpoints set up, click **Add an Azure OpenAI API endpoint**.

5 Define the endpoint settings:

- **Endpoint URL**: Enter the endpoint URL for the model using the Azure OpenAI Chat Completions API endpoint format:

  ```
  https://<resource-name>.<domain>/openai/deployments/<deployment-name>/chat/completions?api-version=<api-version>
  ```

  If the URL you copied does not use the above format, rewrite it. For example, let's assume that your deployment name is `gpt-5.4-mini`, the API version is `2025-04-01-preview`, and you copied the following URL:

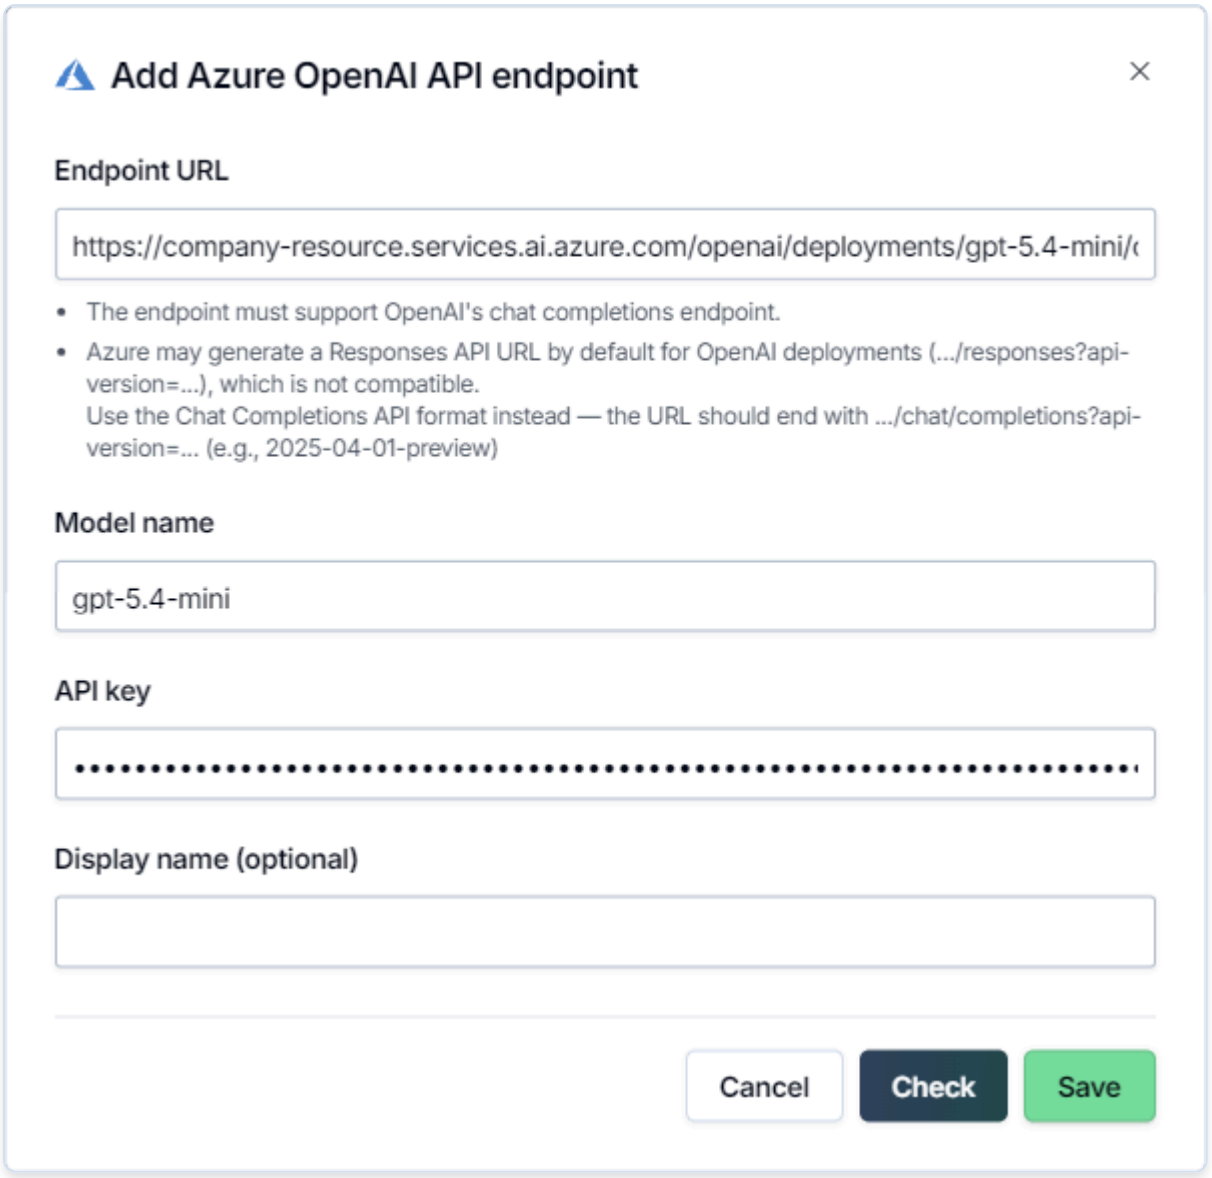
  ```
  https://company-resource.services.ai.azure.com/openai/v1/responses
  ```

  To rewrite the URL into the Chat Completions API format, replace everything after the domain with `/openai/deployments/<deployment-name>/chat/completions?api-version=<api-version>` and replace the placeholders with your deployment name and the API version:

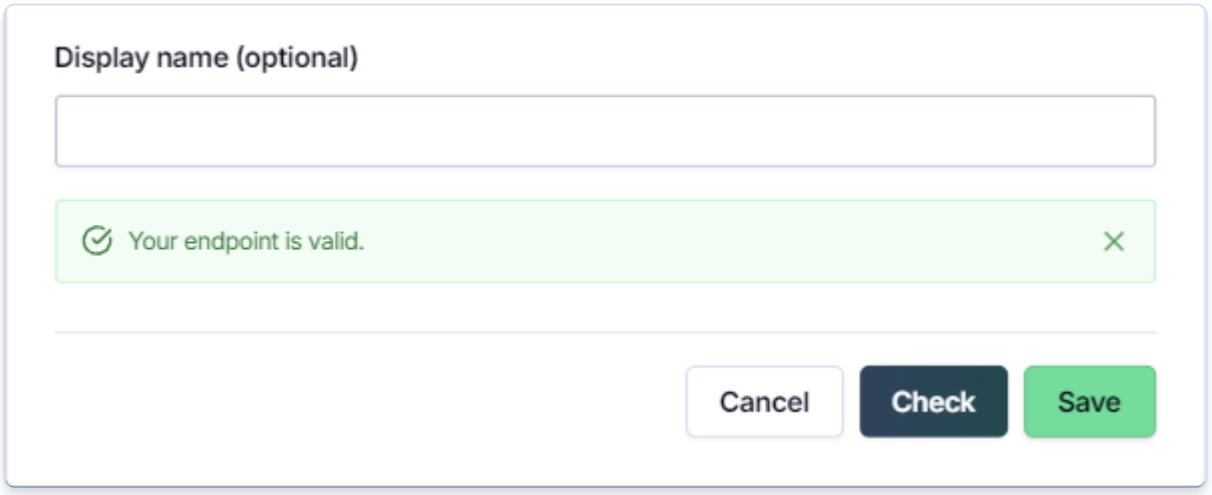
  ```
  https://company-resource.services.ai.azure.com/openai/deployments/gpt-5.4-mini/chat/completions?api-version=2025-04-01-preview
  ```

  > **INFO**
  >
  > The `api-version` is the Azure OpenAI REST API version, which is not shown on the Microsoft Foundry model details page. As a rule, use `2025-04-01-preview` as the value.

- **Model name**: Enter the model name.
- **API key**: Enter the API key for accessing the endpoint.
- **Display name (optional)**: Enter a custom display name for the model. If not set, the model switcher uses **Model name** as the display name.



6 Click **Check** to validate the settings and click **Save**.



You have set up the Azure OpenAI API endpoint. All space users can now use the endpoint model in GPT for Excel and GPT for Word.

Connect to OpenAI-compatible API endpoints

> **INFO**
>
> This feature is only available if your space is on the Business or Enterprise plan or pay-as-you-go pricing.

GPT for Excel and GPT for Word can connect to any AI service that provides an OpenAI-compatible API endpoint. You can connect to two main types of services: **cloud-based platforms**, which provide access to models over the internet with no software installation or setup required on your part, and **local servers**, which you run on a local machine, such as your own computer or another computer on a local network.

You can use the endpoint-connected models both as Agent models (only in Excel) and bulk models, but the endpoint requirements are different:

- **Agent models**: The endpoint must be compatible with the OpenAI Responses API and be publicly accessible over the internet through HTTPS. In addition, the model must support a context size of at least 32K tokens.
- **Bulk models**: The endpoint must be compatible with the OpenAI Chat Completions API.

Here are some popular services that provide an OpenAI-compatible API endpoint:

- **Cloud-based platforms:**
  - Anyscale
  - Fireworks AI
  - Together AI
- **Local servers:**
  - LM Studio
  - LocalAI
  - Ollama

> **INFO**
>
> Depending on the model you use, the model settings defined in GPT for Excel or Word (temperature (creativity), top-p, frequency penalty, and presence penalty) may not work as expected. If you get bad results, try tweaking the settings.

Agent models

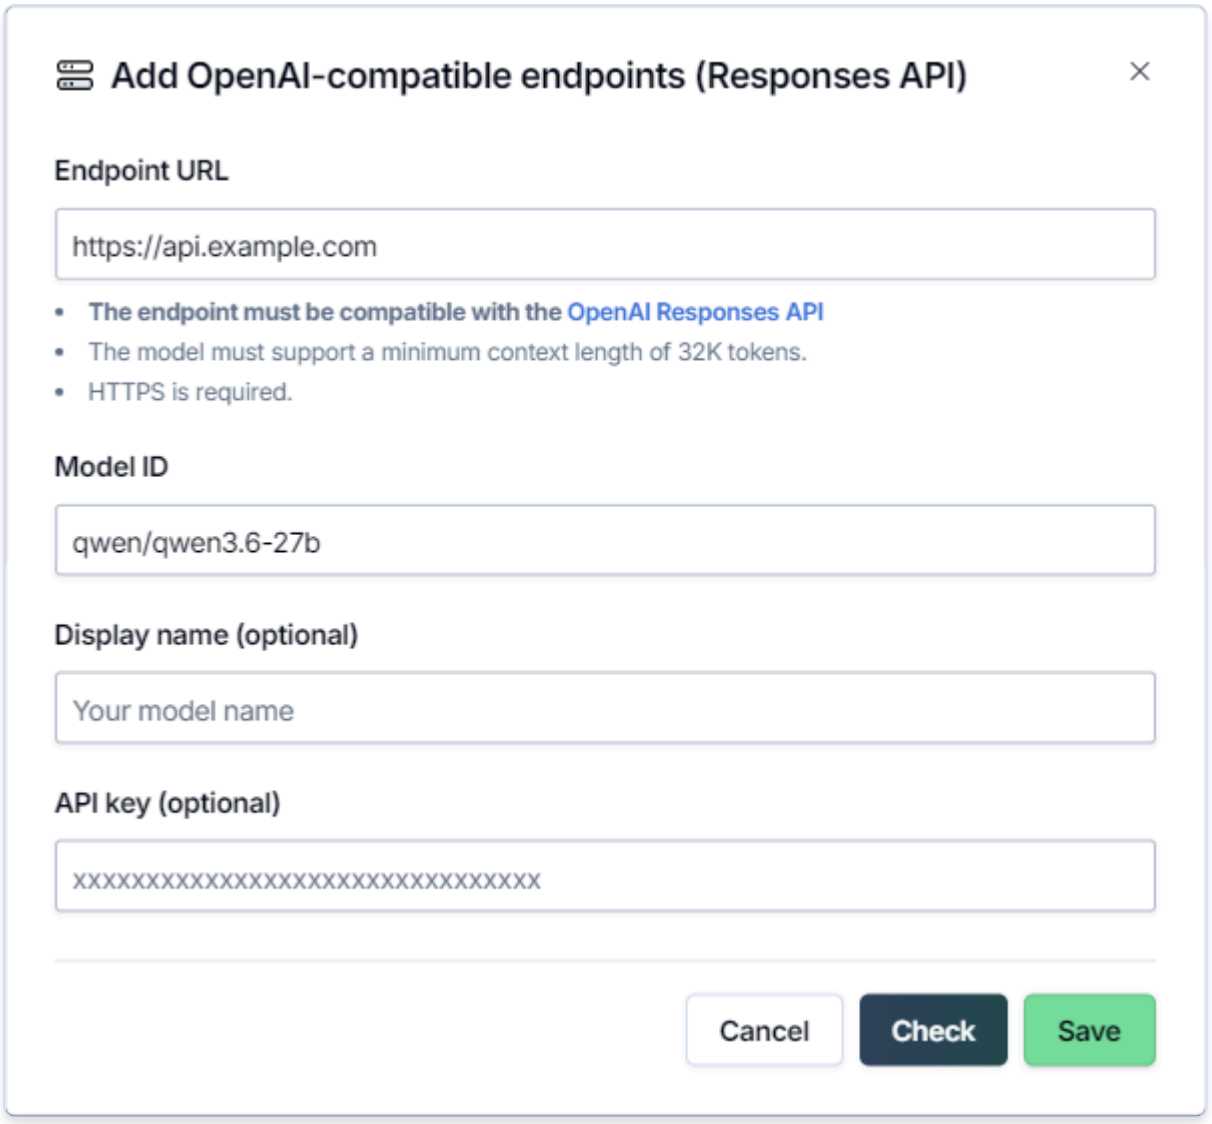
Prerequisites

- You are the owner or an admin of the space.
- You have an endpoint URL and, if needed, an API key for accessing the service.
- The endpoint supports the following OpenAI-compatible API resource path: `/v1/responses`
- The endpoint is publicly accessible over the internet through HTTPS.
- The model supports a context size of at least 32K tokens.

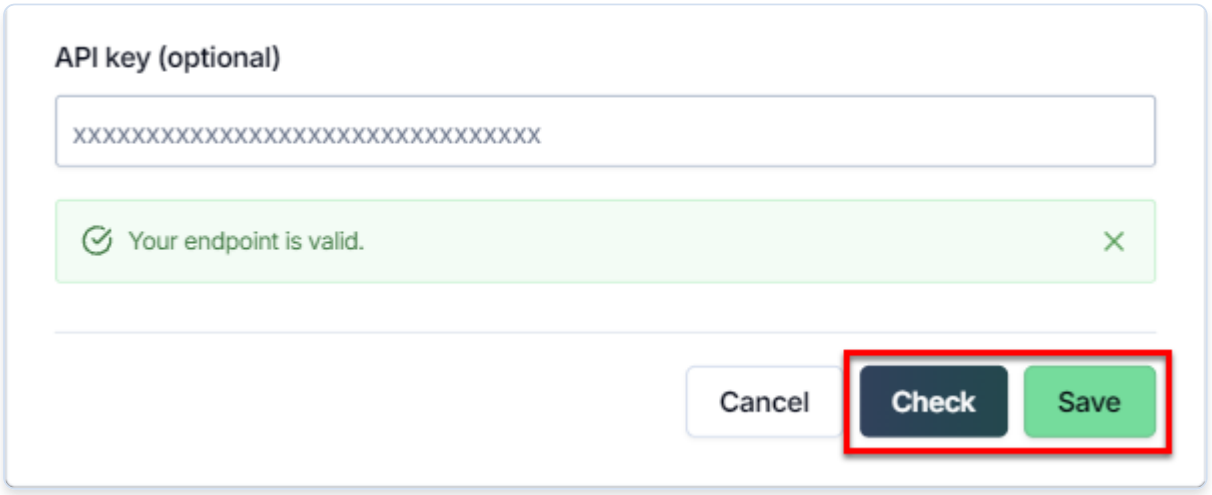
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Agent models** tab.
4. Click **OpenAI-compatible endpoints (Responses API)**.

> **TIP**
>
> If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the URL of the service that provides the endpoint.
   - **Model ID**: Enter the ID of the model to use through the endpoint. The model ID is sent to the endpoint in the request payload.
   - **Display name (optional)**: Enter a short name that identifies the model in the Agent model switcher.
   - **API key (optional)**: If required, enter the API key for accessing the service.



6 Click **Check** to validate the settings and click **Save**.



You have set up the OpenAI-compatible API endpoint. All space users can now select the model in the Agent model switcher in GPT for Excel.

Bulk models

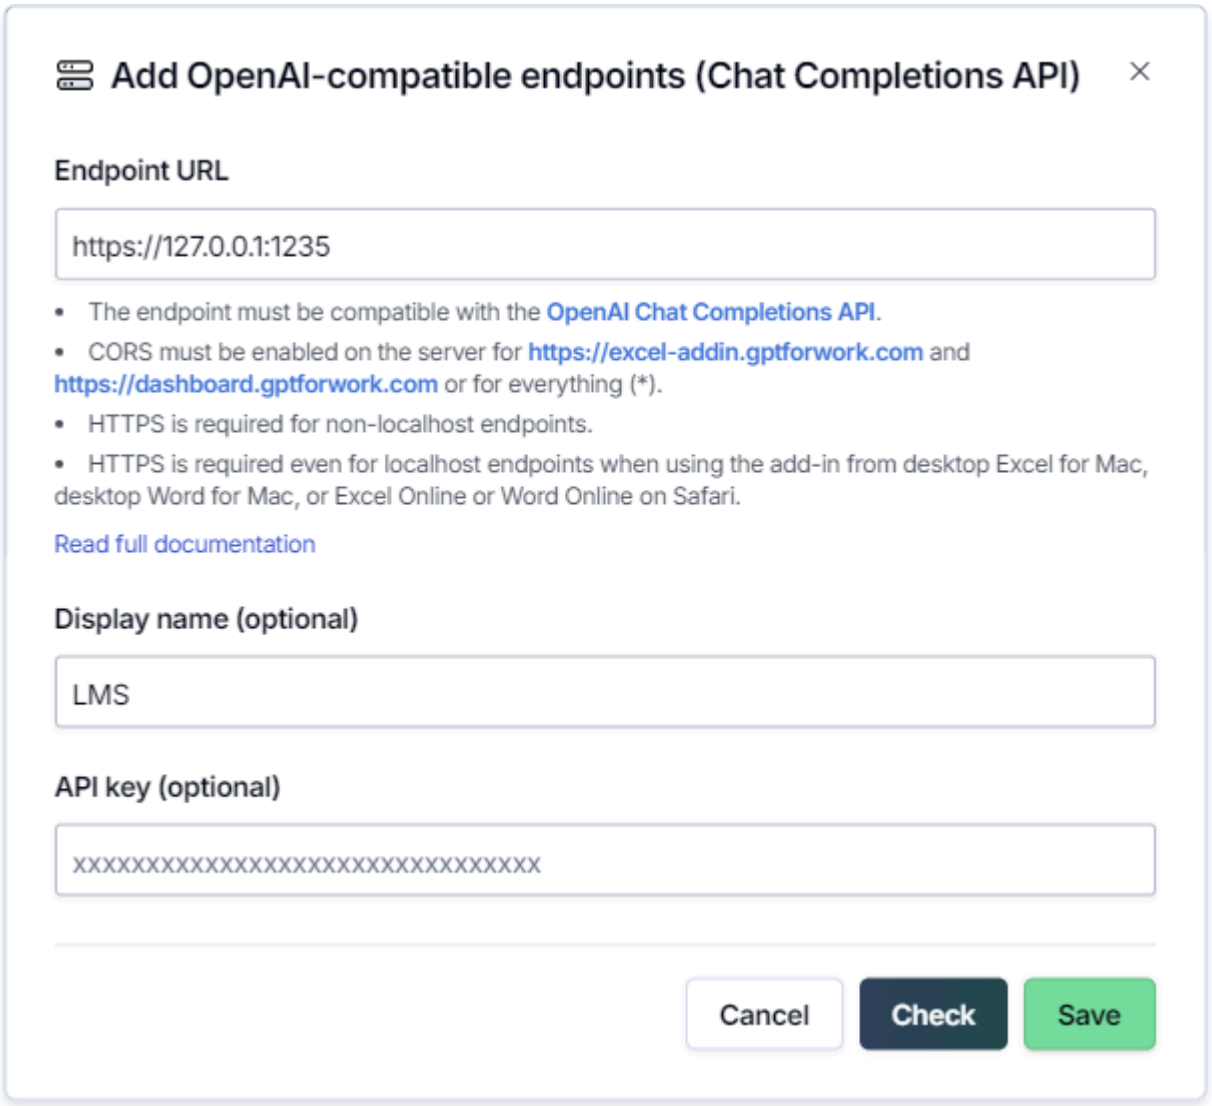
Prerequisites

- You are the owner or an admin of the space.
- You have an endpoint URL and, if needed, an API key for accessing the service.
- The endpoint supports the following OpenAI-compatible API resource paths:
  - `/v1/models`
  - `/v1/chat/completions`

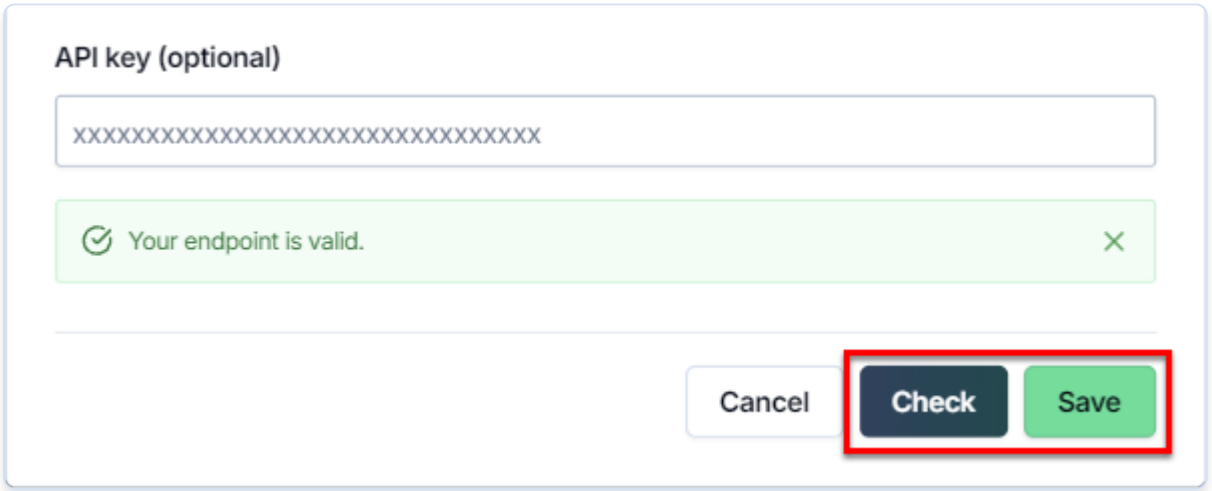
1. Sign in to the GPT for Work dashboard with your Microsoft or Google account.
2. In the sidebar, select **Custom API endpoints**.
3. Select the **Bulk models** tab.
4. Click **OpenAI-compatible endpoints (Chat Completions API)**.

> **TIP**
>
> If you already have OpenAI-compatible API endpoints set up, click **Add an OpenAI-compatible API endpoint**.

5. Define the endpoint settings:
   - **Endpoint URL**: Enter the URL of the service that provides the endpoint.
   - **Display name (optional)**: Enter a short name that is prefixed to the model names in the model switcher. By default, the model names are prefixed with "custom/".
   - **API key (optional)**: If required, enter the API key for accessing the service.



6 Click **Check** to validate the settings and click **Save**.



> **TIP**
>
> If you get an error saying that the URL is invalid, make sure that the endpoint supports:
>
> - **Cross-origin resource sharing (CORS).** See the LM Studio setup instructions for macOS or Windows.
> - **HTTPS (needed with Safari and with Microsoft Office for Mac).** See the LM Studio setup instructions for macOS.

You have set up the OpenAI-compatible API endpoint. All space users can now use all models available through the endpoint in GPT for Excel and GPT for Word.

What's next

- Manage API keys for your space.
- Manage models for your space.
- Manage users in your team.

Billing

GPT for Work uses subscription pricing. Some existing spaces remain on legacy pay-as-you-go pricing.

You can manage billing for your space from the GPT for Work dashboard.

Subscription

Manage your subscription plan, seats, and extra-usage credits.

Pay-as-you-go (legacy)

Manage billing for pay-as-you-go spaces: buy credits, set spending limits, and download invoices.

Subscription billing

Manage your subscription plan and extra-usage credits from the GPT for Work dashboard. Seat management and spending limits are available on Business and Enterprise plans.

Manage your plan

Upgrade, downgrade, or cancel your GPT for Work subscription.

Manage your seats

Add seats or remove unused seats on your Business or Enterprise subscription.

Buy extra-usage credits

Buy extra-usage credits for your subscription space to top up your shared extra-usage balance.

Set spending limits for extra usage

Set monthly spending limits on the shared extra-usage balance for your subscription space and for individual users.

Monitor your usage

Track your team's included usage and extra-usage credits, and generate usage reports.

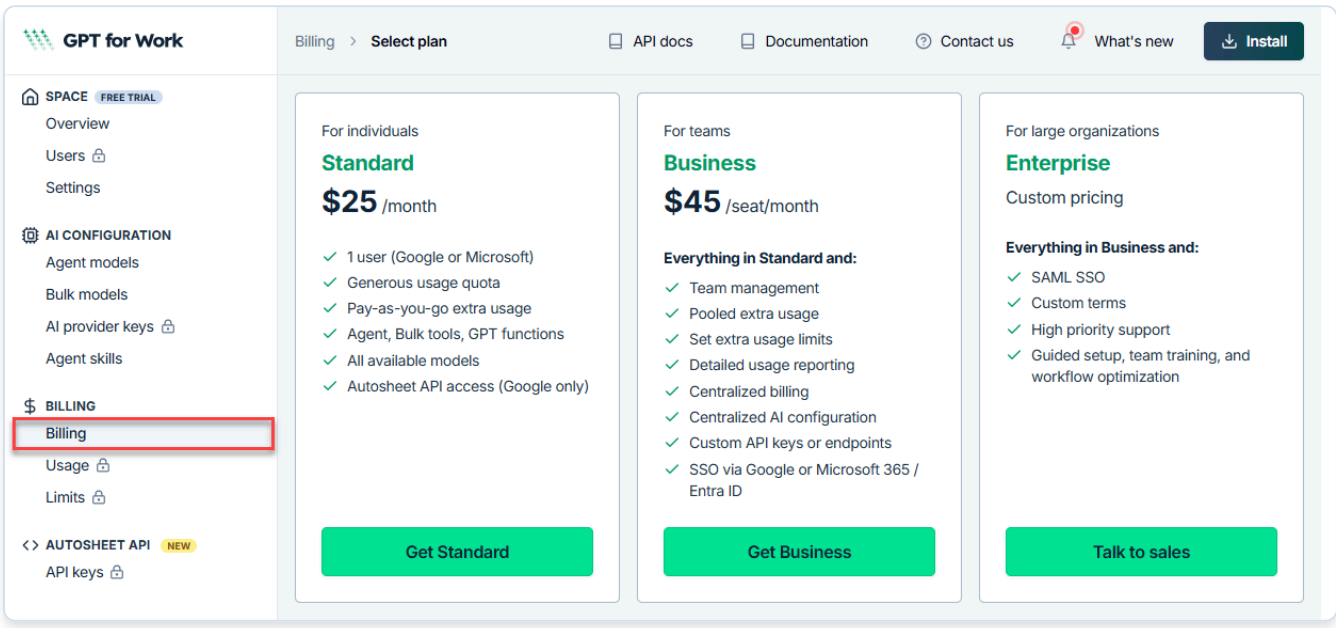
Manage your plan

Upgrade, downgrade, or cancel your GPT for Work subscription from the GPT for Work dashboard. To discuss an Enterprise or custom plan, contact sales.

Upgrade to a paid plan

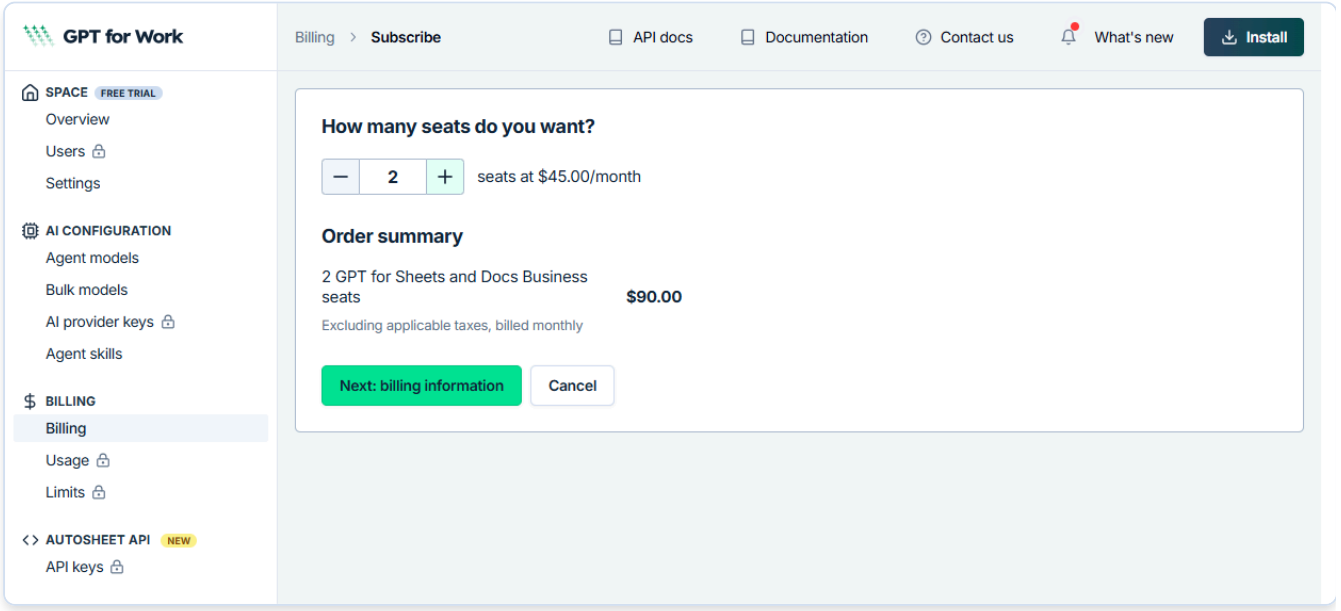
Upgrade to a paid plan to continue using the product once your free trial has ended.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.

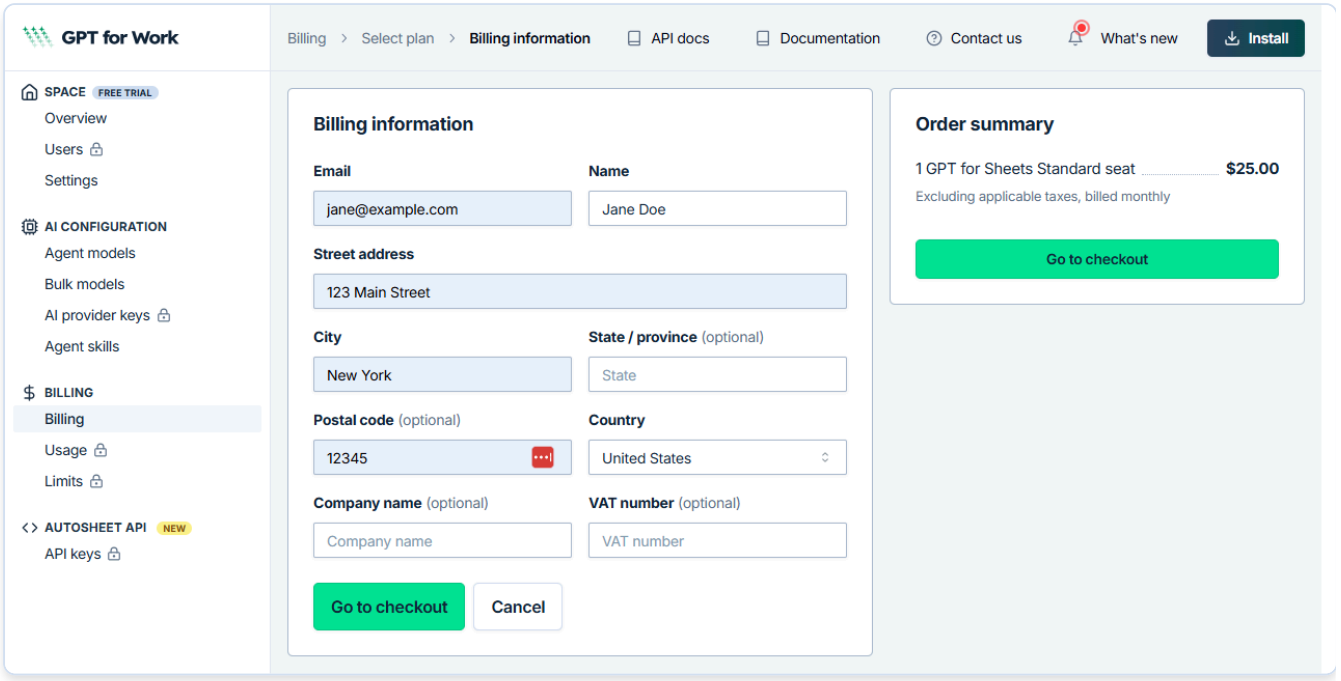


3. Click **Get Standard** or **Get Business**.

   (For Business) Enter the number of seats you need, then click **Next: billing information**.

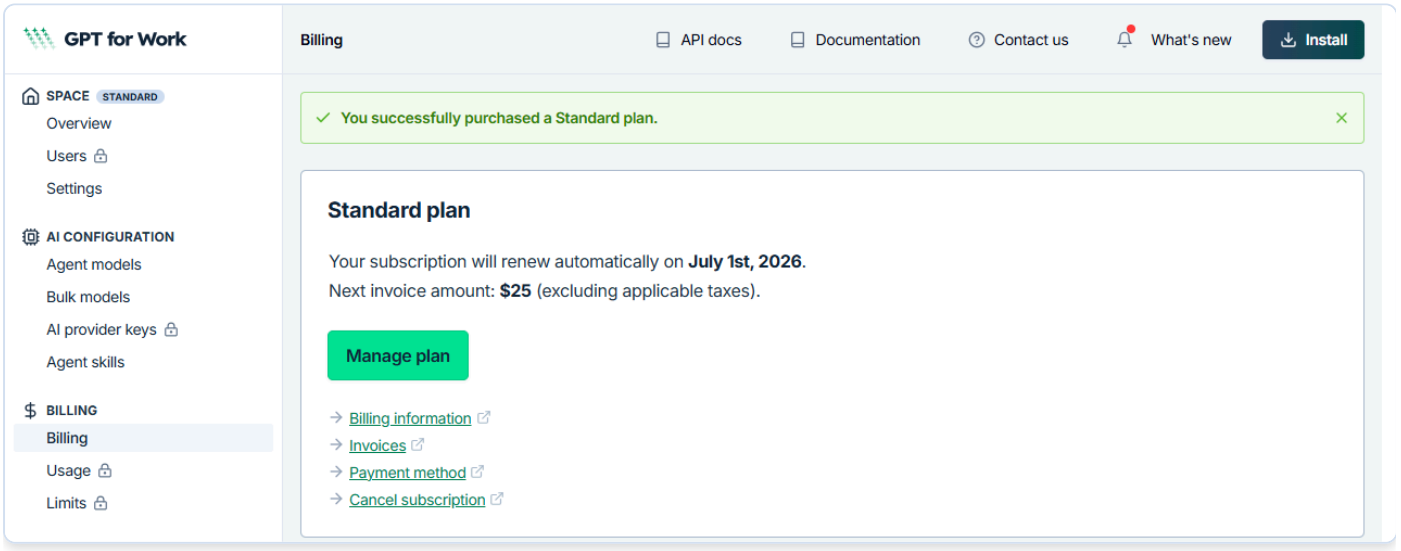


4 Enter your contact and billing details, review the order summary, and click **Go to checkout**.



5 Complete the payment on the Stripe checkout page.

You return to the dashboard. A confirmation banner confirms your plan, along with your next renewal date and invoice amount.

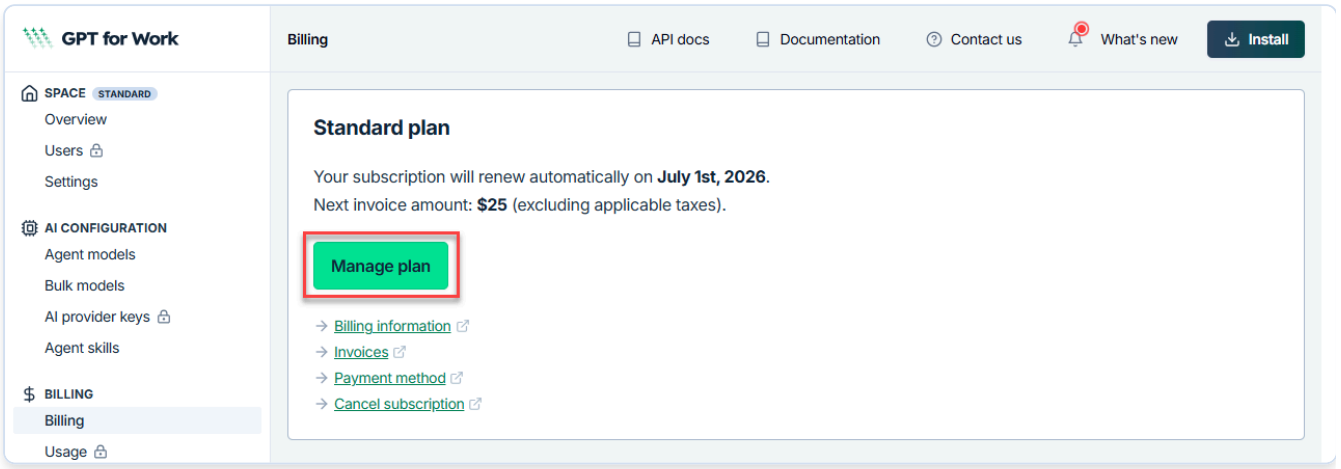


Upgrade from Standard to Business

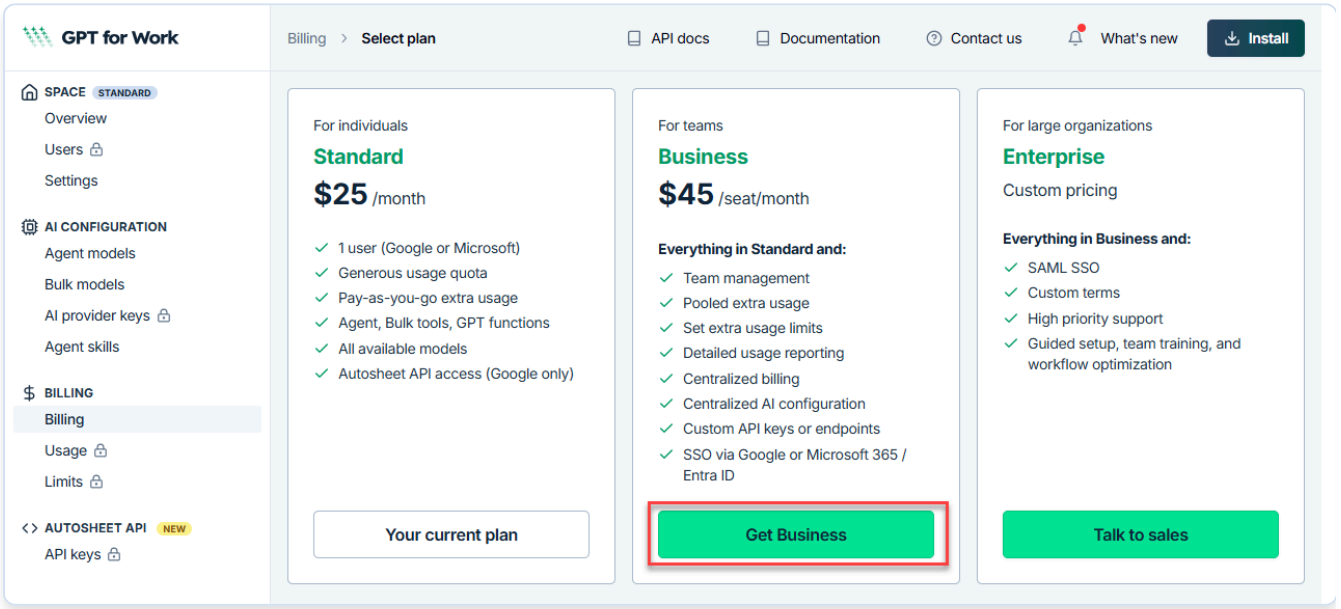
Upgrade from Standard to Business to unlock team management features, centralized billing, spending limits, and more. The upgrade takes effect immediately, and your billing is prorated for the remainder of the current cycle.

> **NOTE**
>
> Your existing seat keeps its remaining included usage for the current billing cycle. The upgrade does not reset or increase it.

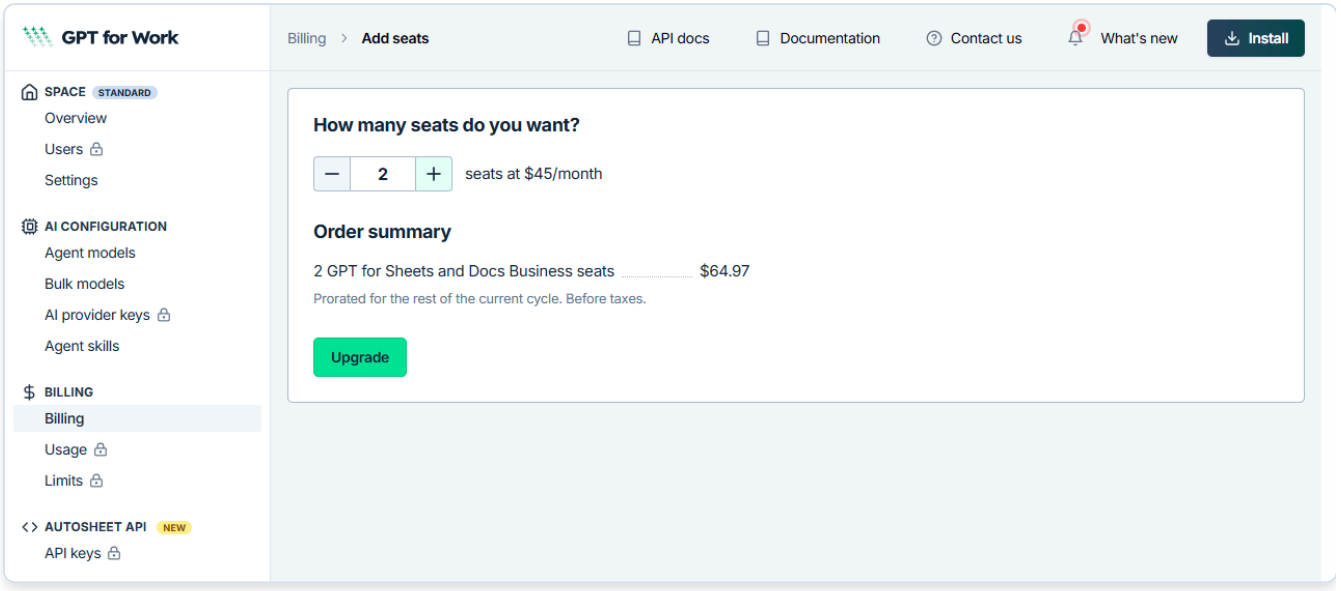
1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.
3. Click **Manage plan**.



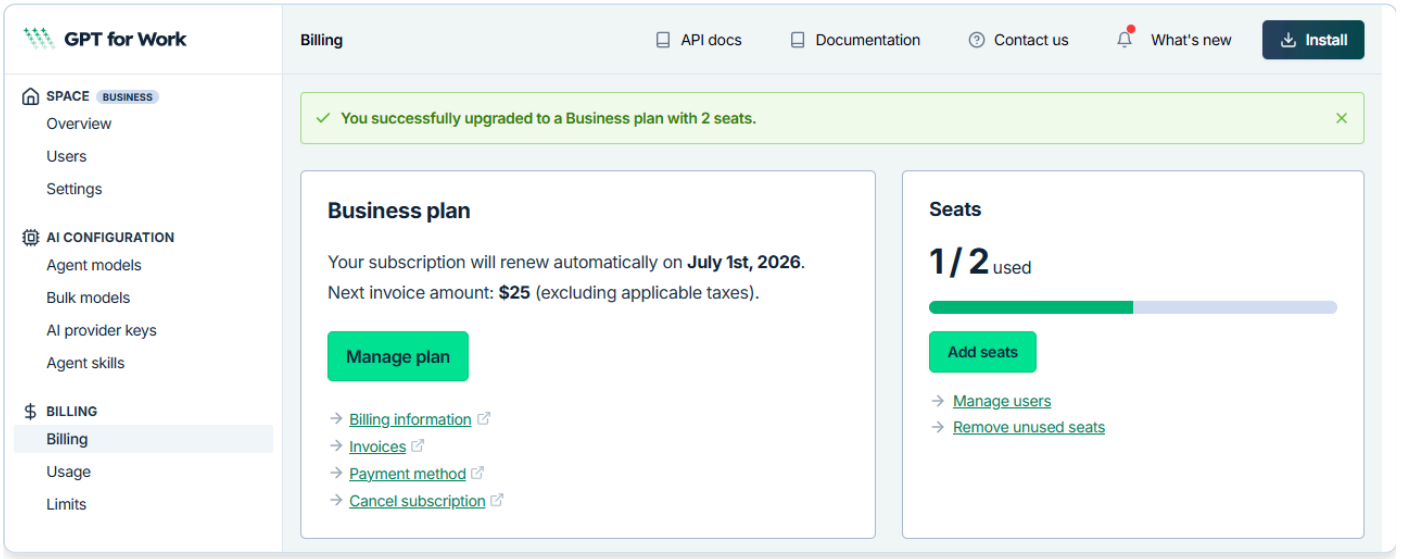
4. Click **Get Business**.



5 Enter the number of seats you need, and click **Upgrade**.



You return to the dashboard. A confirmation banner confirms your Business plan, along with your next renewal date and invoice amount.



Downgrade from Business to Standard

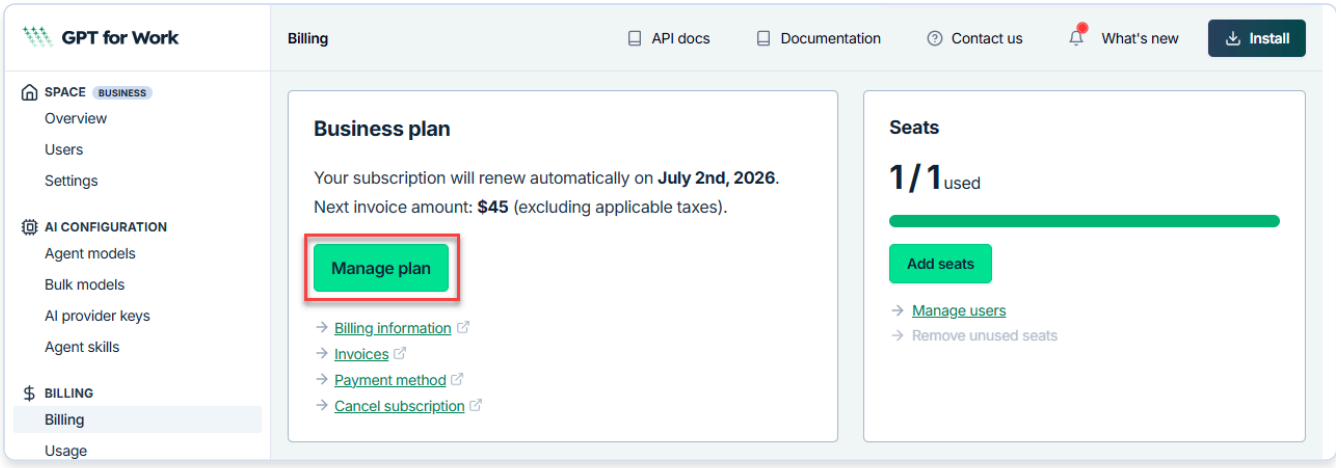
Downgrading from Business to Standard schedules the change for the end of your current billing cycle. You keep full access to Business features until then.

> **PREREQUISITE**
>
> You need to be the only user in the space. Remove users if necessary.

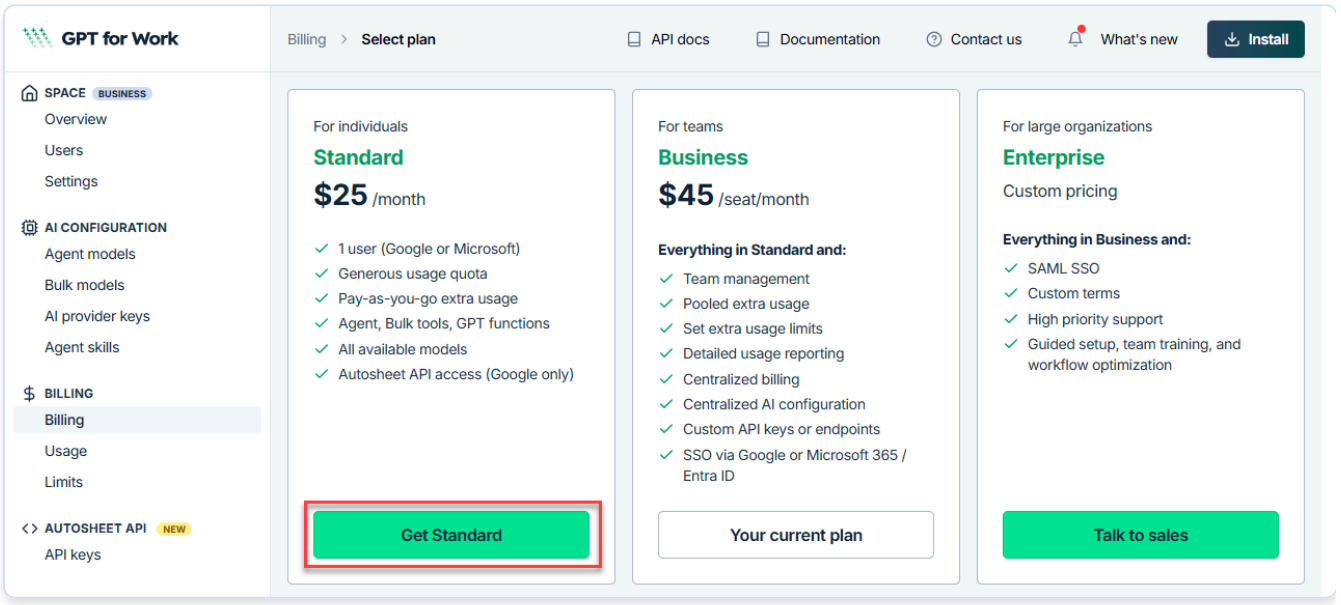
1 Sign in to the GPT for Work dashboard with the account that you use for your space.

2 In the sidebar, select **Billing**.

3 Click **Manage plan**.

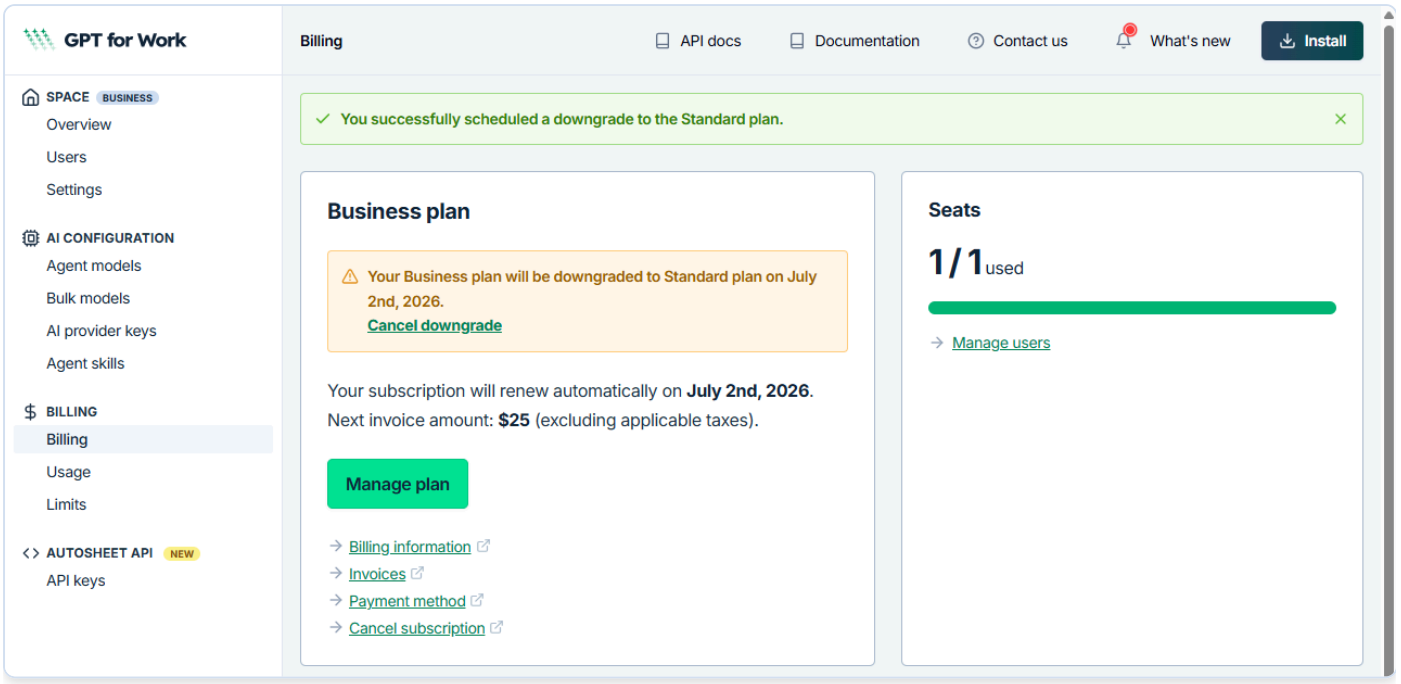


4 Click **Get Standard**.



5 Review the features that will no longer be available, and confirm the downgrade.

Your downgrade is scheduled. You keep access to Business features until the end of the current billing cycle, at which point your plan switches to Standard.



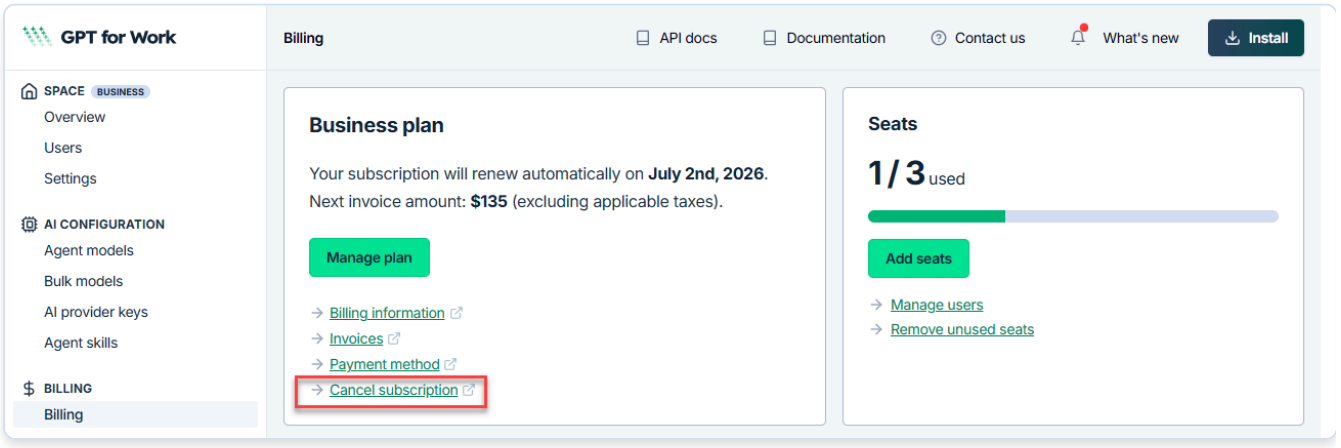
You can cancel the scheduled downgrade at any time before the billing cycle ends from the **Billing** page in the GPT for Work dashboard.

Cancel your subscription

Canceling your subscription schedules the cancellation for the end of your current billing cycle. You keep full access until then.

> **NOTE**
>
> Your extra-usage credits are preserved, but you cannot use them while your subscription is inactive. They become usable again when you renew your subscription. Extra-usage credits still expire 12 months after your last purchase, even while your subscription is inactive.

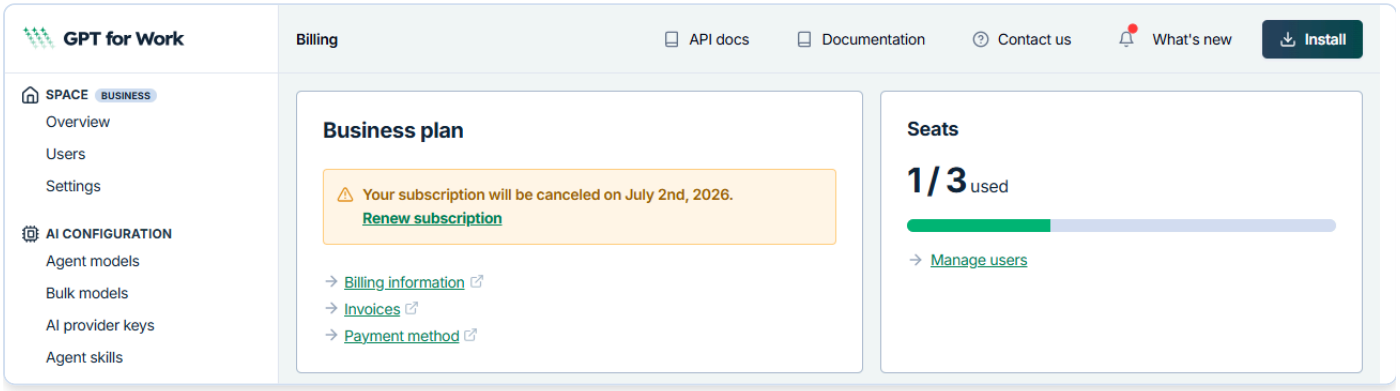
1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.
3. Click **Cancel subscription**.



4 In the Stripe Customer Portal, click **Cancel subscription**.

5 Review the cancellation details, and click **Cancel subscription** to confirm.

After confirming the cancellation, return to the dashboard manually. The dashboard reflects the scheduled cancellation.



You can renew your subscription at any time before the billing cycle ends from the **Billing** page in the GPT for Work dashboard.

What's next

- Buy extra-usage credits
- Manage your seats. Available on Business and Enterprise plans.
- Set spending limits for extra usage. Available on Business and Enterprise plans.

Manage your seats

Add seats or remove unused seats on your Business or Enterprise subscription from the GPT for Work dashboard.

PREREQUISITE

- Only the space owner or a space admin can manage seats.
- This feature is available on Business and Enterprise plans.

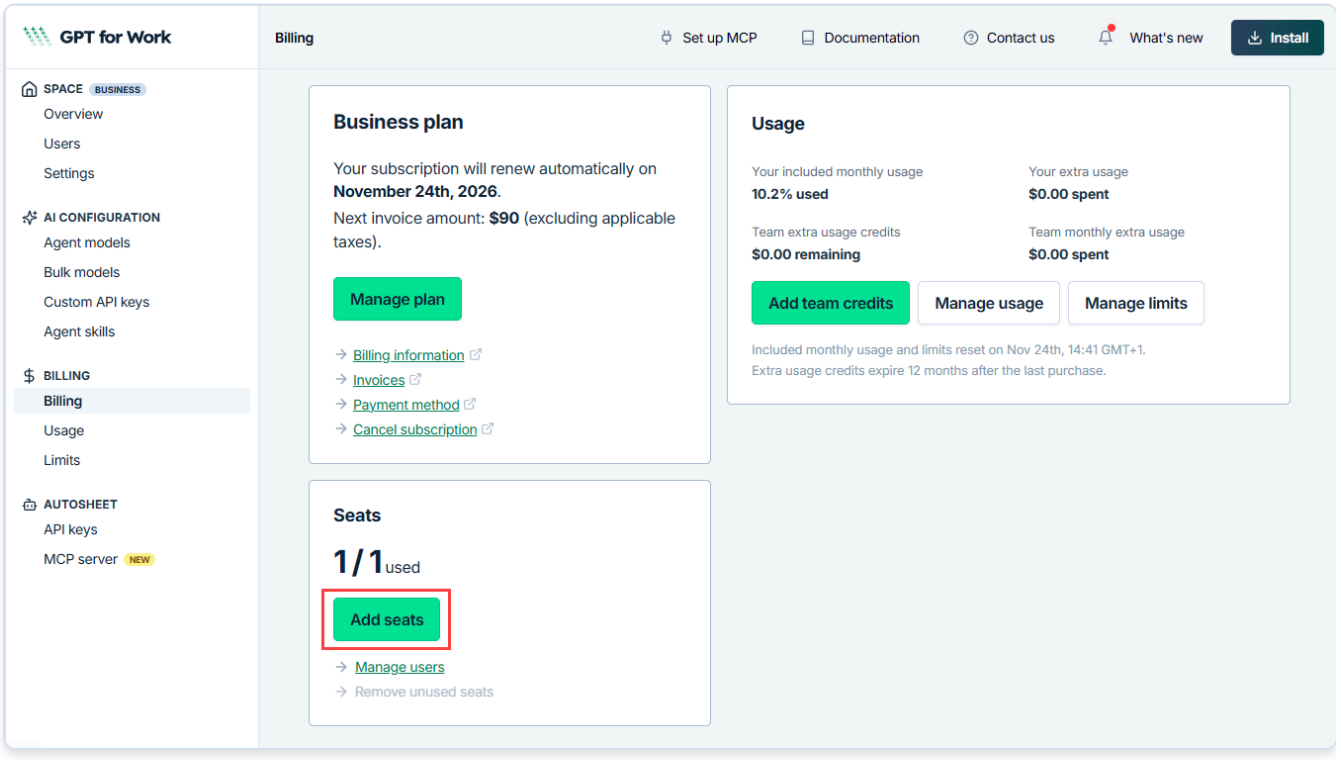
Add seats

New seats are available immediately. You are charged a prorated amount for the remainder of the current billing cycle, and each new seat's included usage is prorated for the same period.

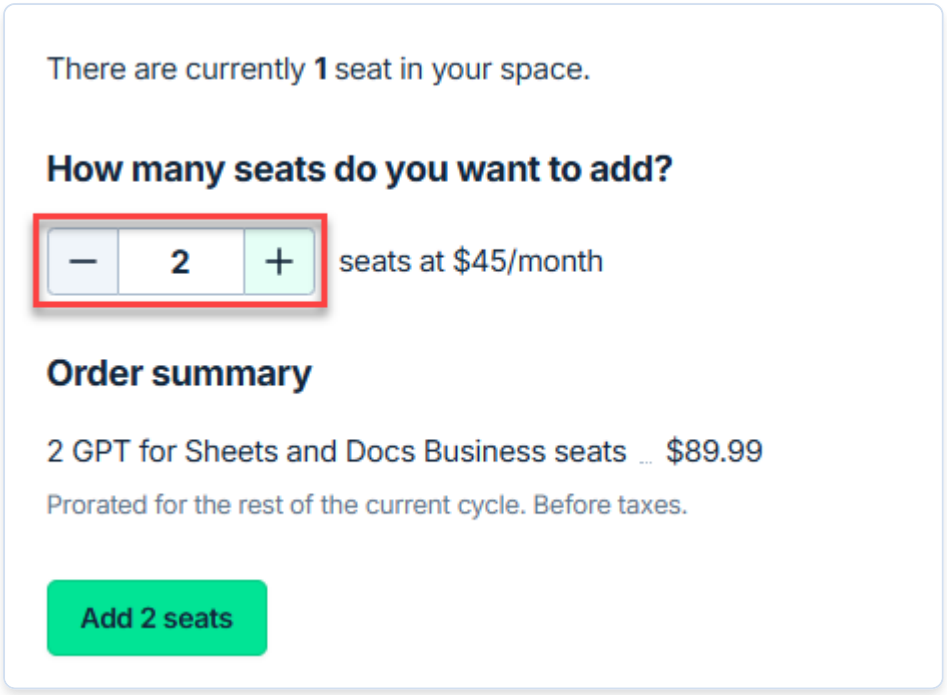
NOTE

You cannot add seats while a seat removal is scheduled. To add seats, first click **Cancel downgrade** on the **Billing** page.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.
3. Click **Add seats**.

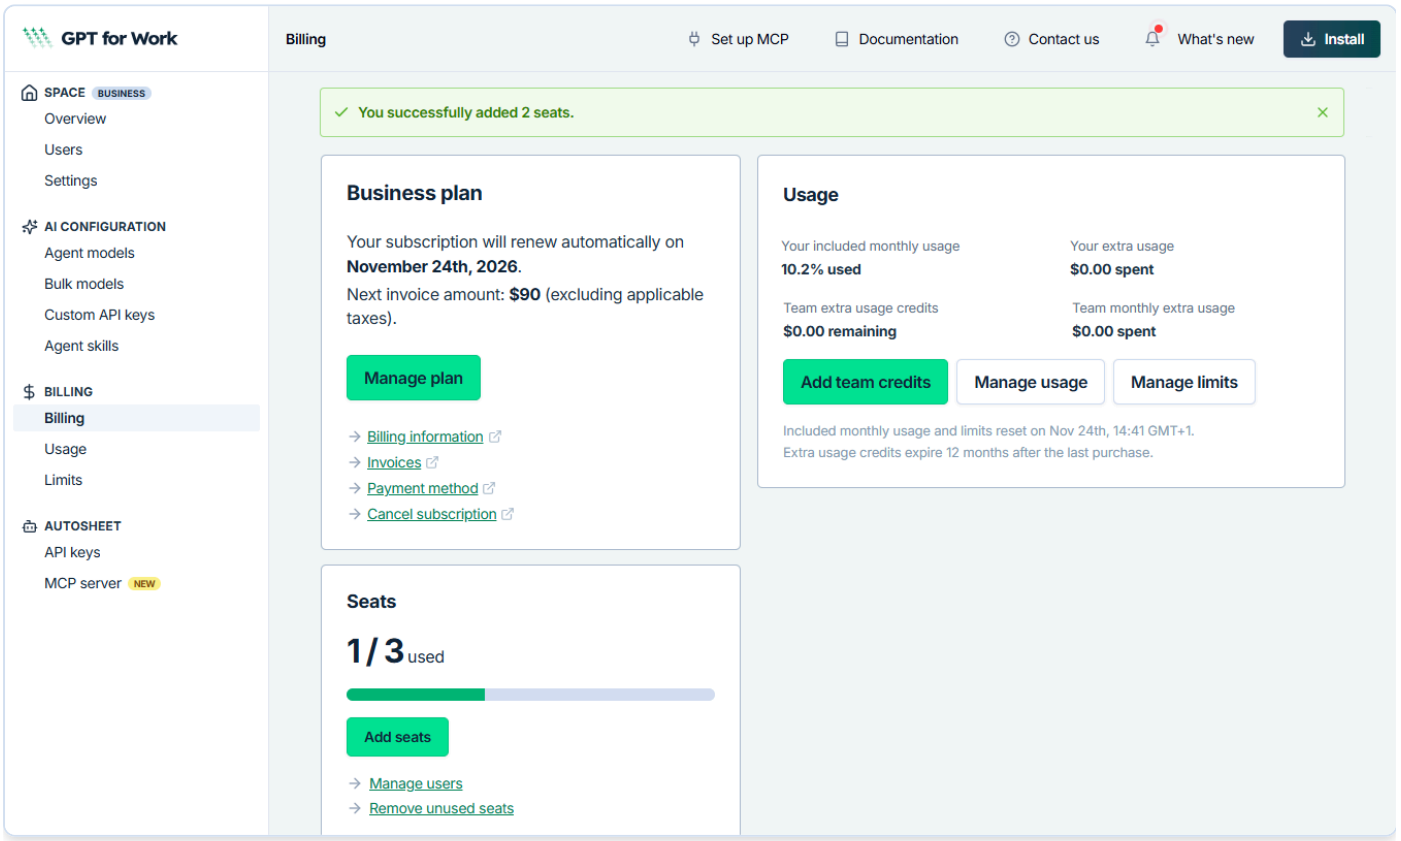


4 Select the total number of seats you need. The page shows a preview of the prorated charge.



5 Click **Add seats**.

Your new seats are immediately available. The dashboard confirms the updated seat count and billing details.

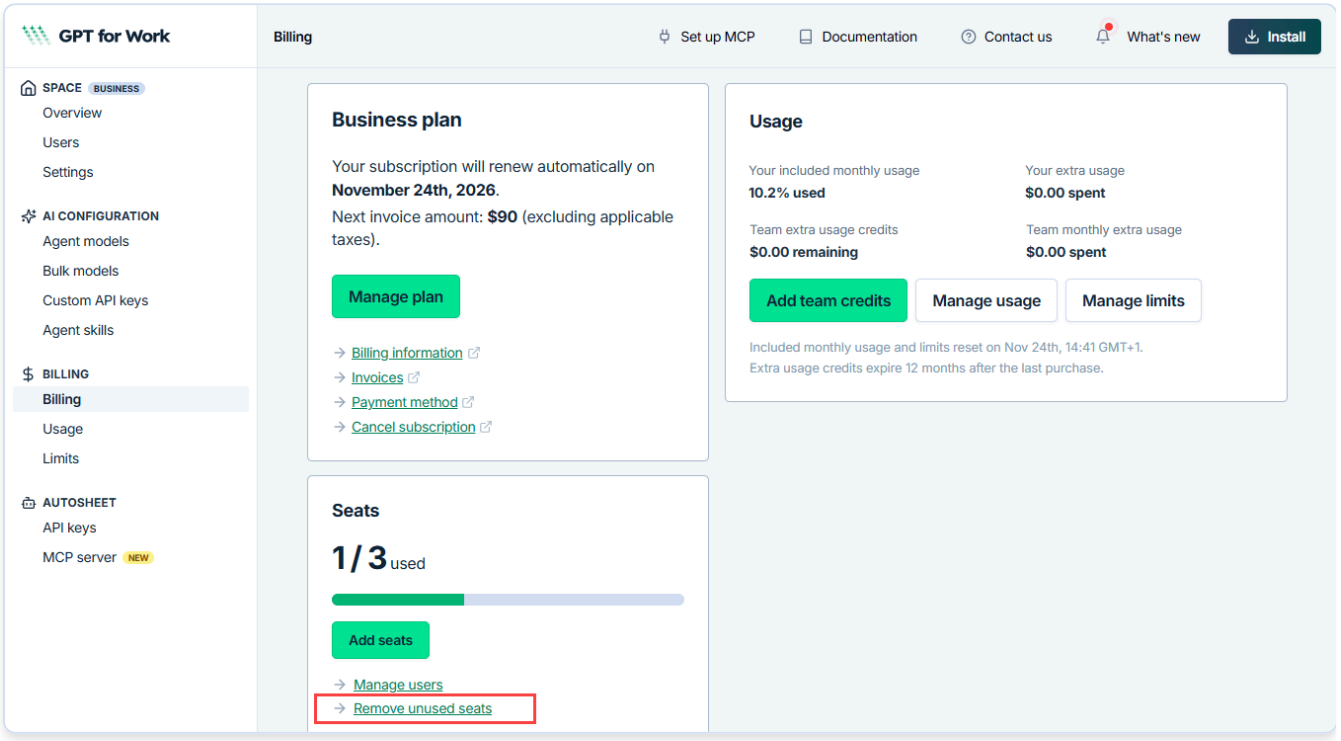


Remove unused seats

Remove seats you no longer need to reduce your next invoice. The removal is scheduled for the end of the current billing cycle. You keep your current seat count until then.

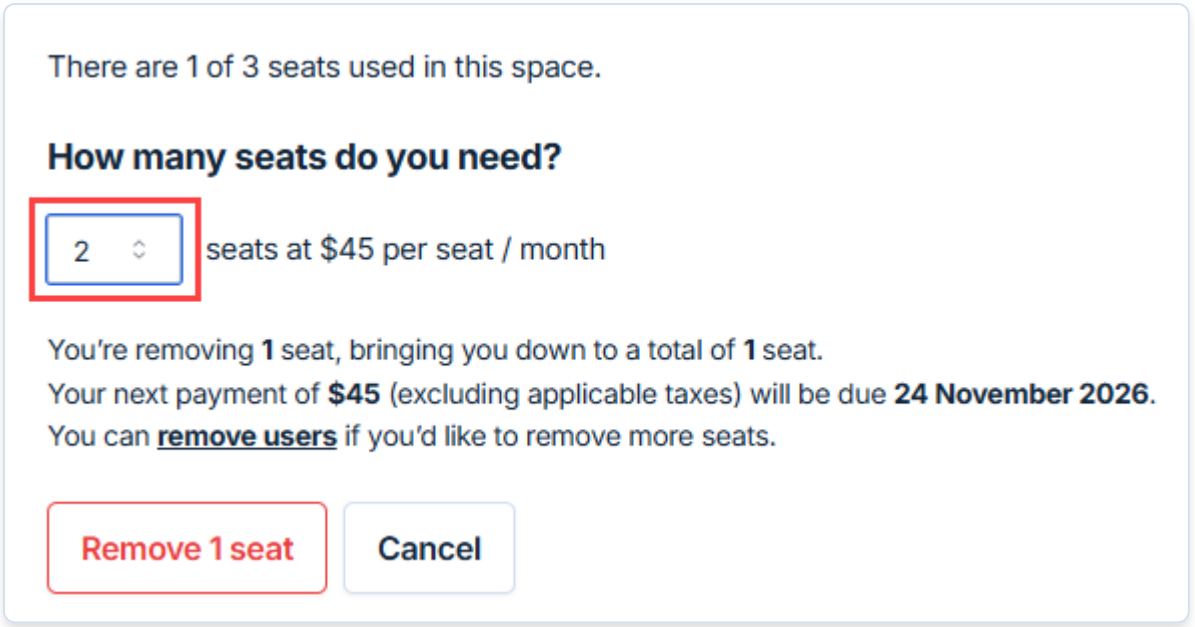
> **INFO**
>
> You can only remove unused seats. To remove more seats than are currently unused, remove users from the space first.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.
3. Click **Remove unused seats**.



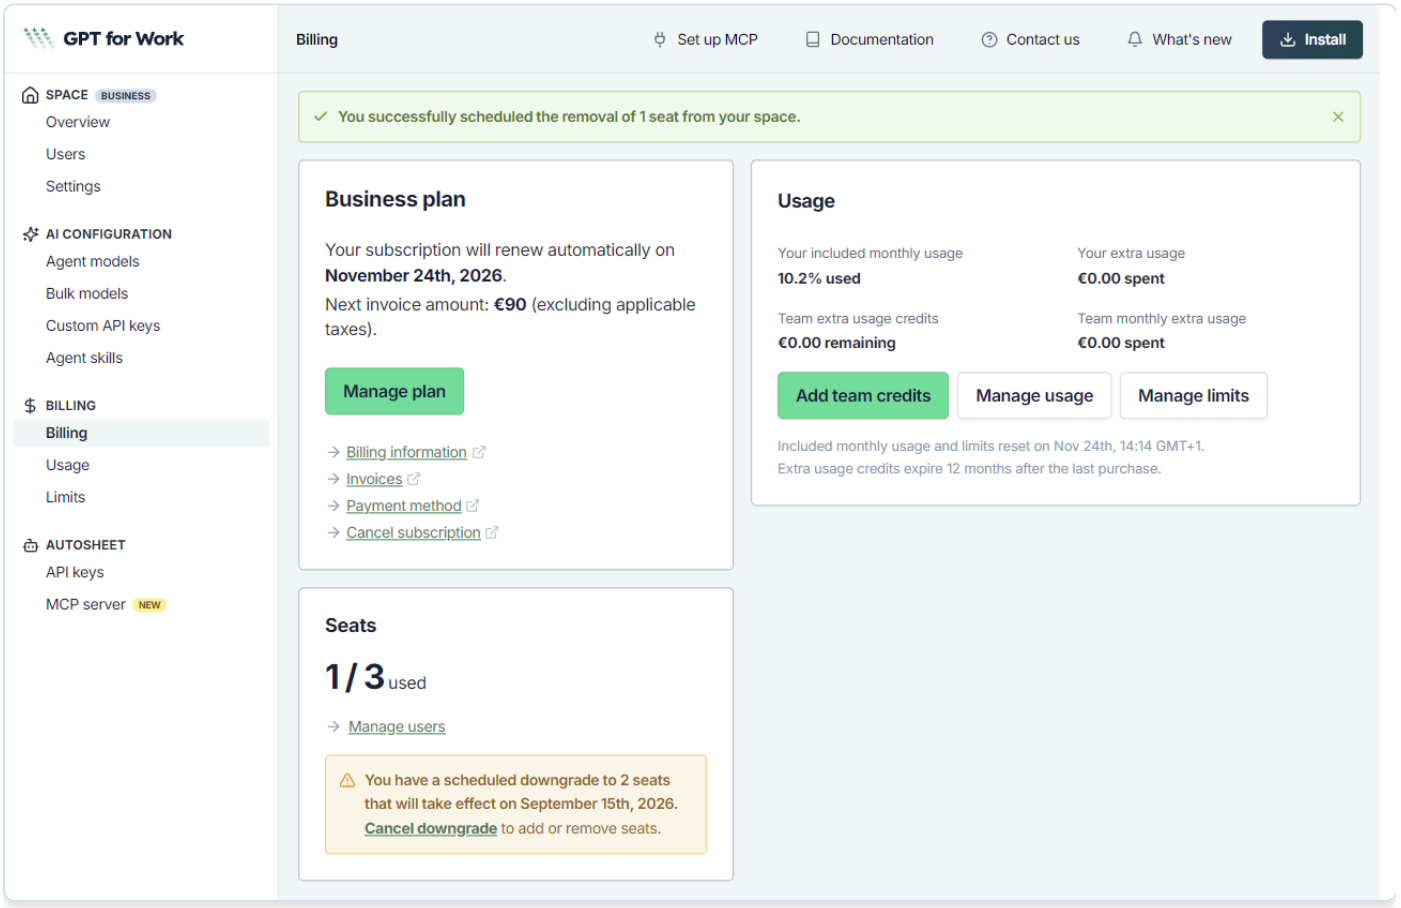
4. Select how many seats you need. The page shows your current seat usage and previews the

impact on your next payment.



5 Click **Remove seats**.

Your seat removal is scheduled. The dashboard confirms the seat count after the change takes effect, the date it applies, and your updated billing amount for the next cycle.



> **NOTE**
>
> When a seat removal is scheduled, you cannot add or remove seats. To make further changes, first click **Cancel downgrade** on the **Billing** page.

What's next

- Manage users
- Set spending limits for extra usage

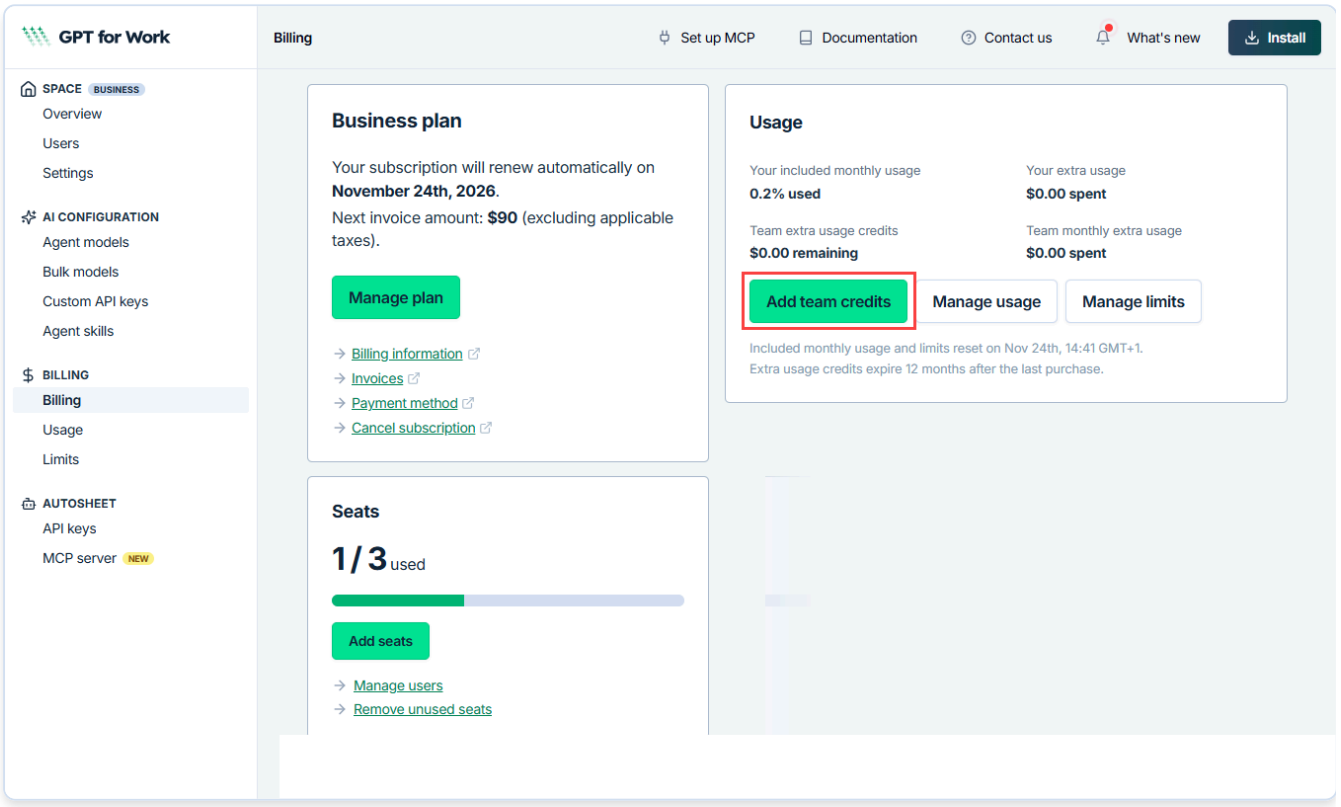
Buy extra-usage credits

Buy extra-usage credits when your seats' included usage for the current billing cycle runs out. Credits are added to a shared extra-usage balance for the whole space and are valid for 12 months from the most recent purchase.

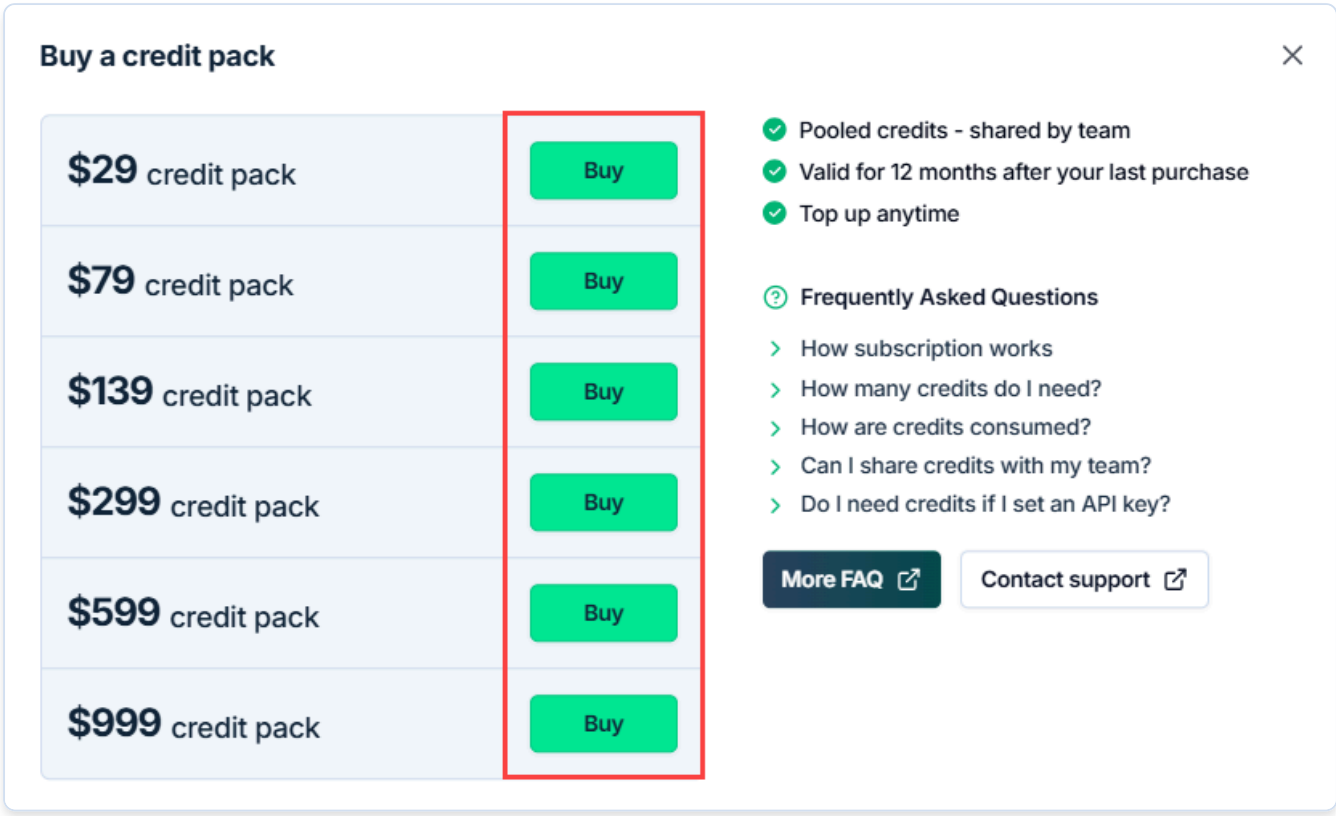
PREREQUISITE

- Only the space owner or a space admin can buy extra-usage credits.
- You must be on a paid plan. Upgrade your free trial first if needed.

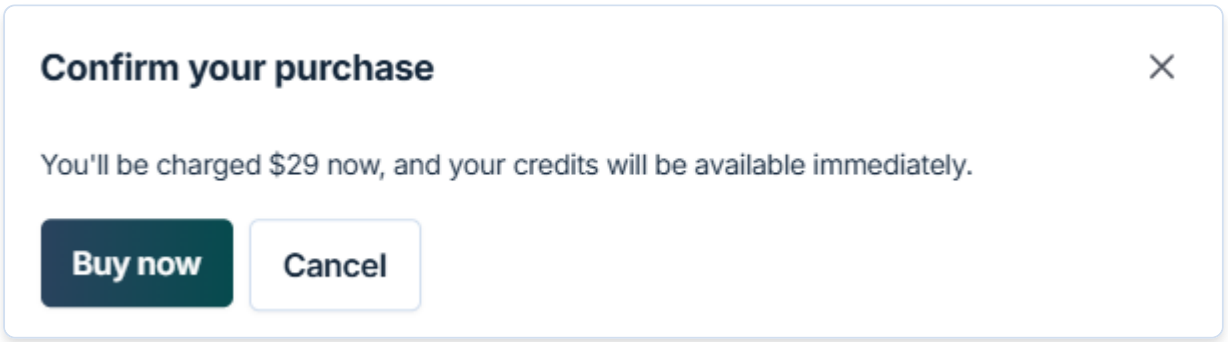
1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.
3. Click **Add team credits**.



4 Click **Buy** for the credit pack that best suits your needs.



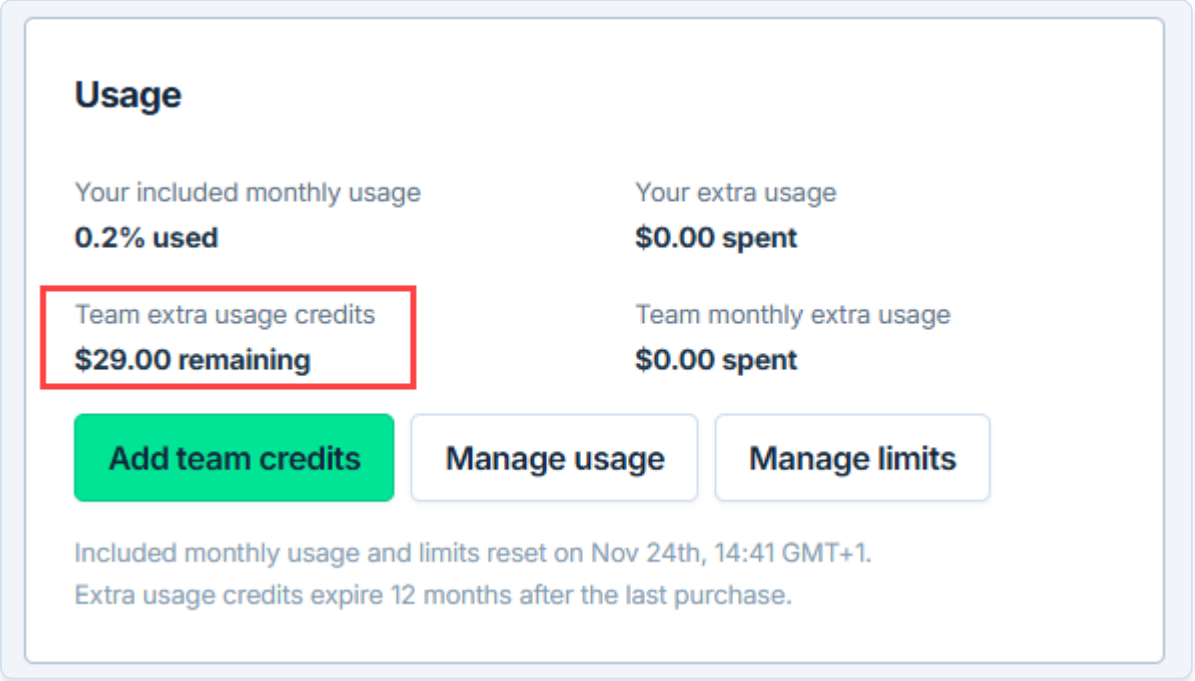
5 Review your order and click **Buy now**.



6 Enter your billing information and click **Go to checkout**.

7 Enter your payment details and click **Pay**.

Your credits are added to the shared extra-usage balance and are available immediately.



What's next

- Set spending limits for extra usage

Set spending limits for extra usage

Set monthly spending limits to control how much can be spent on extra-usage credits in your space during each billing cycle. Each seat's included usage is unaffected. You can set a limit for the entire space and individual limits per user.

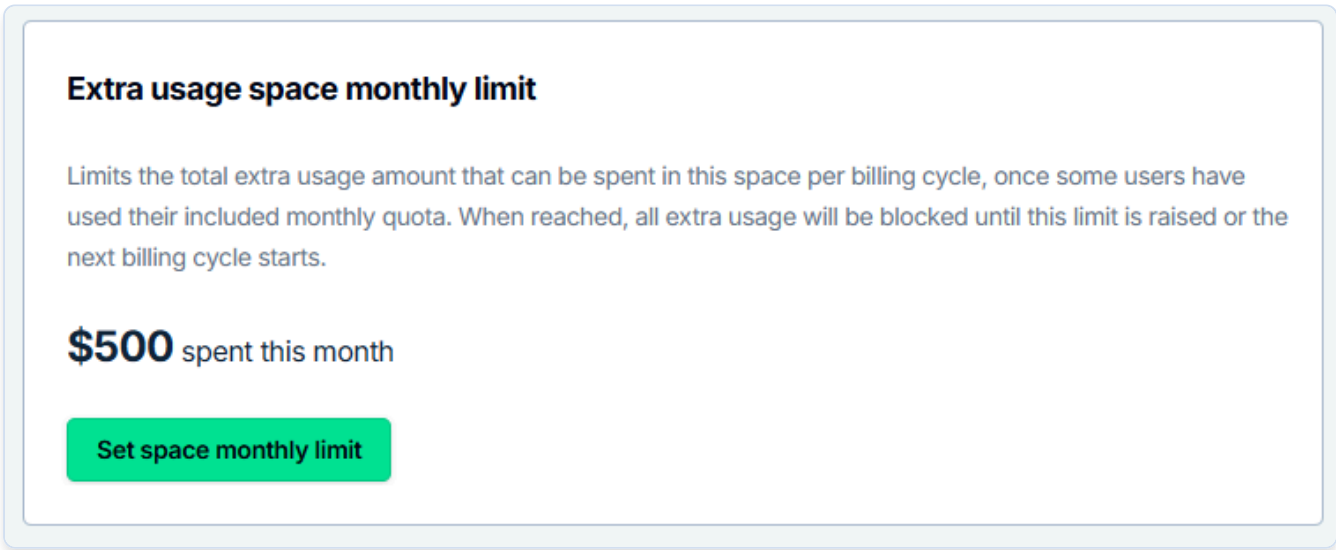
PREREQUISITE

- Only the space owner or a space admin can set spending limits.
- This feature is available on Business and Enterprise plans.

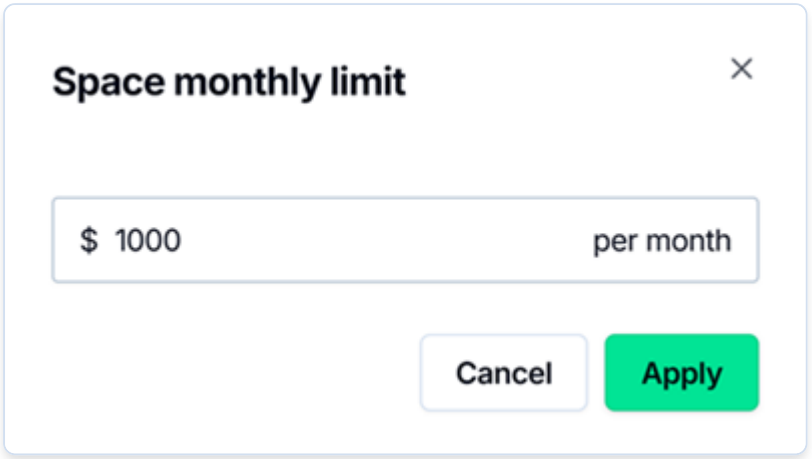
Set a space monthly limit

A space monthly limit restricts the total amount that can be spent on extra-usage credits across all users in your space during each billing cycle. When the limit is reached, all extra-usage is blocked until you raise the limit or the limit resets at the start of the next billing cycle.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Limits**.
3. In the **Extra usage space monthly limit** section, click **Set space monthly limit**.



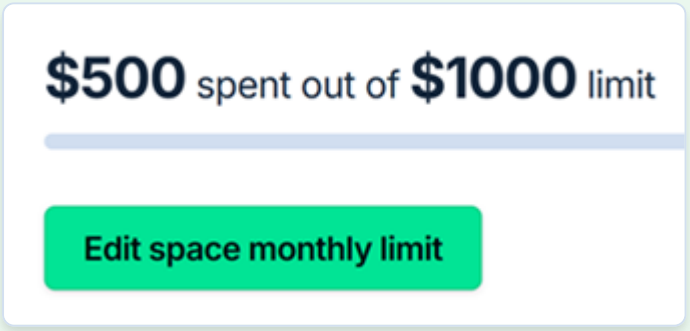
4 Enter the amount, and click **Apply**.



The space monthly limit immediately applies to all users in the space.

EDIT OR REMOVE

- To edit the space monthly limit, click **Edit space monthly limit**, edit the amount, and click **Apply**.
- To remove the space monthly limit, click **Edit space monthly limit** and click **Remove**.



Set user monthly limits

User monthly limits restrict the amount that an individual user can spend on extra-usage credits during each billing cycle. When a user reaches their limit, that user's extra-usage is blocked until the limit is raised or the limit resets at the start of the next billing cycle. Other users in the space can continue using the add-ons normally.

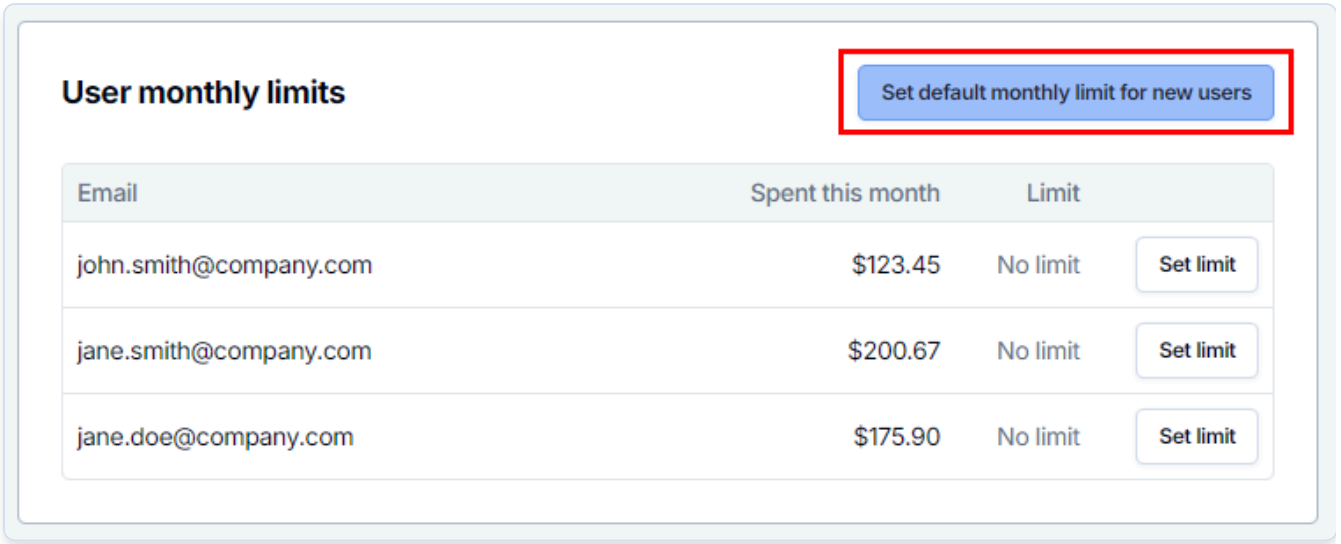
Set a default limit for future users

You can set a default monthly limit that applies to all new users who join the space. You can adjust the monthly limit for individual users after they join.

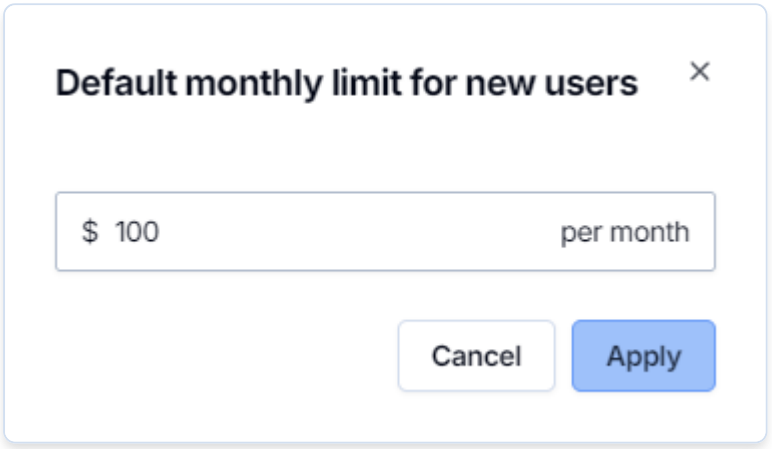
1 Sign in to the GPT for Work dashboard with the account that you use for your space.

2 In the sidebar, select **Limits**.

3 In the **User monthly limits** section, click **Set default monthly limit for new users**.



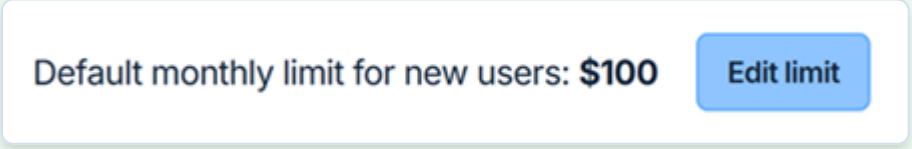
4 Enter the amount, and click **Apply**.



The default monthly limit automatically applies to all users who join the space after this point.

EDIT OR REMOVE

- To edit the default monthly limit, click **Edit limit**, edit the amount, and click **Apply**.
- To remove the default monthly limit, click **Edit limit** and click **Remove**.

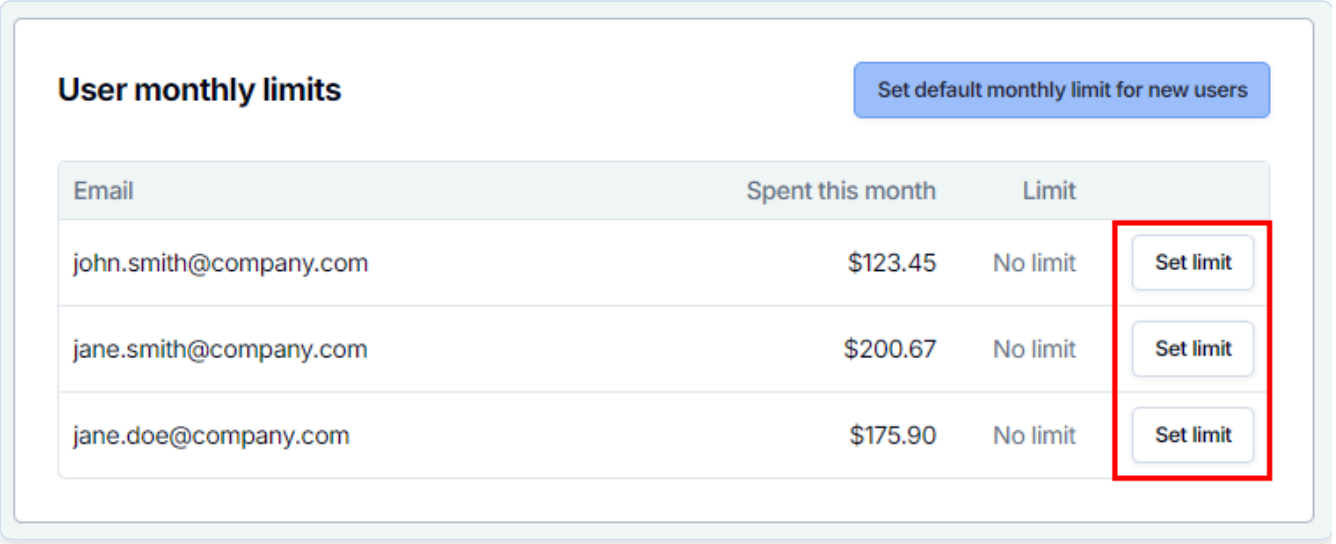


Set limits for current users

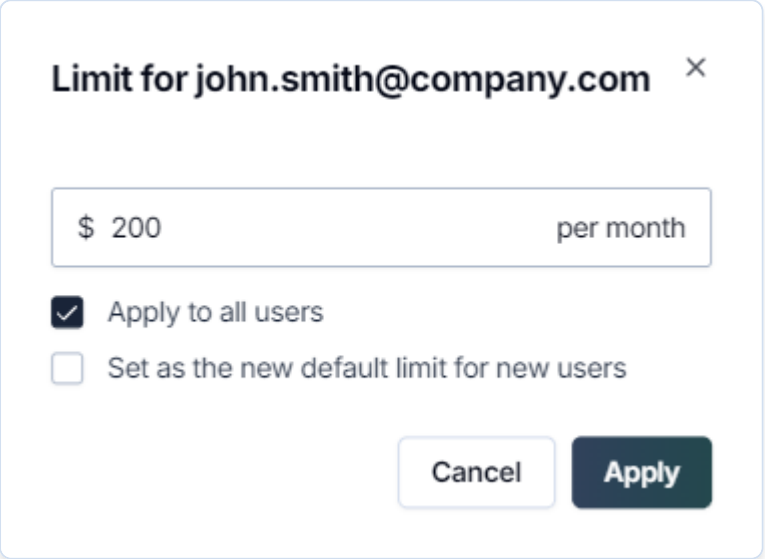
You can set monthly limits for individual users in your space, or apply the same limit to all current users in one go.

1 Sign in to the GPT for Work dashboard with the account that you use for your space.

2. In the sidebar, select **Limits**.
3. In the **User monthly limits** section, find the user and click **Set limit**.



4. Enter the amount.
5. *(Optional)* Select **Apply to all users** to apply the same limit to all current users in the space.

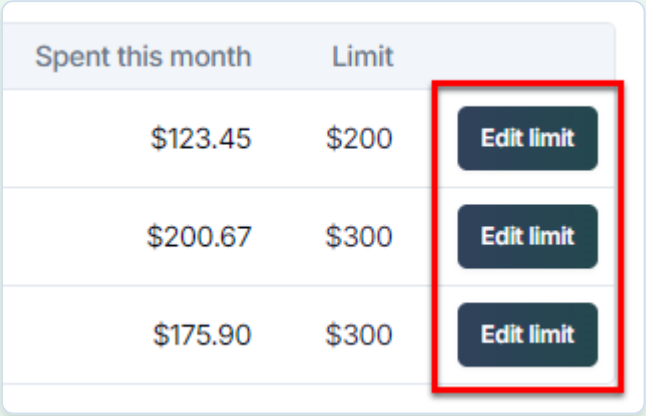


6. Click **Apply**.

The user monthly limit immediately applies.

EDIT OR REMOVE

- To edit monthly limits:
  - Single user: Click **Edit limit** for the user, edit the amount, and click **Apply**.
  - All users: Click **Edit limit** for any user, edit the amount, select **Apply to all users**, and click **Apply**.
- To remove monthly limits:
  - Single user: Click **Edit limit** for the user and click **Remove**.
  - All users: Click **Edit limit** for any user, select **Apply to all users**, and click **Remove**.



What's next

- Buy extra-usage credits

Monitor your usage

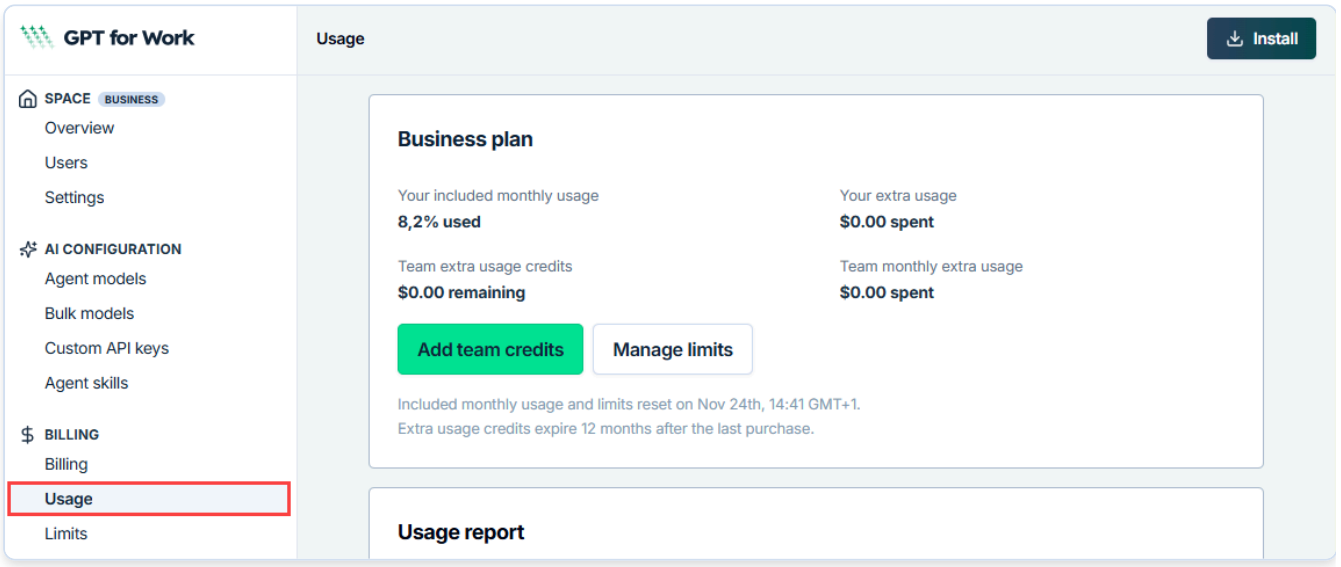
With subscription pricing, each seat comes with its own included usage that resets at the start of each billing cycle, and your space can buy extra-usage credits that are shared across the whole team. Track your usage and remaining extra-usage credits, and generate detailed usage reports, from the GPT for Work dashboard.

Track your usage

Prerequisites

You are the owner or an admin of the space.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Usage**.



The page shows the following figures:

Figure	Description
Your included monthly usage	How much of your seat's included usage you have consumed, as a percentage.
Team extra usage credits	How much is left in the extra-usage credits shared across your space. Available on Business and Enterprise plans.
Your extra usage	How much you have spent from the extra-usage credits shared across your space during the current billing cycle.

Figure	Description
Team monthly extra usage	How much your whole space has spent from the shared extra-usage credits during the current billing cycle. Available on Business and Enterprise plans.

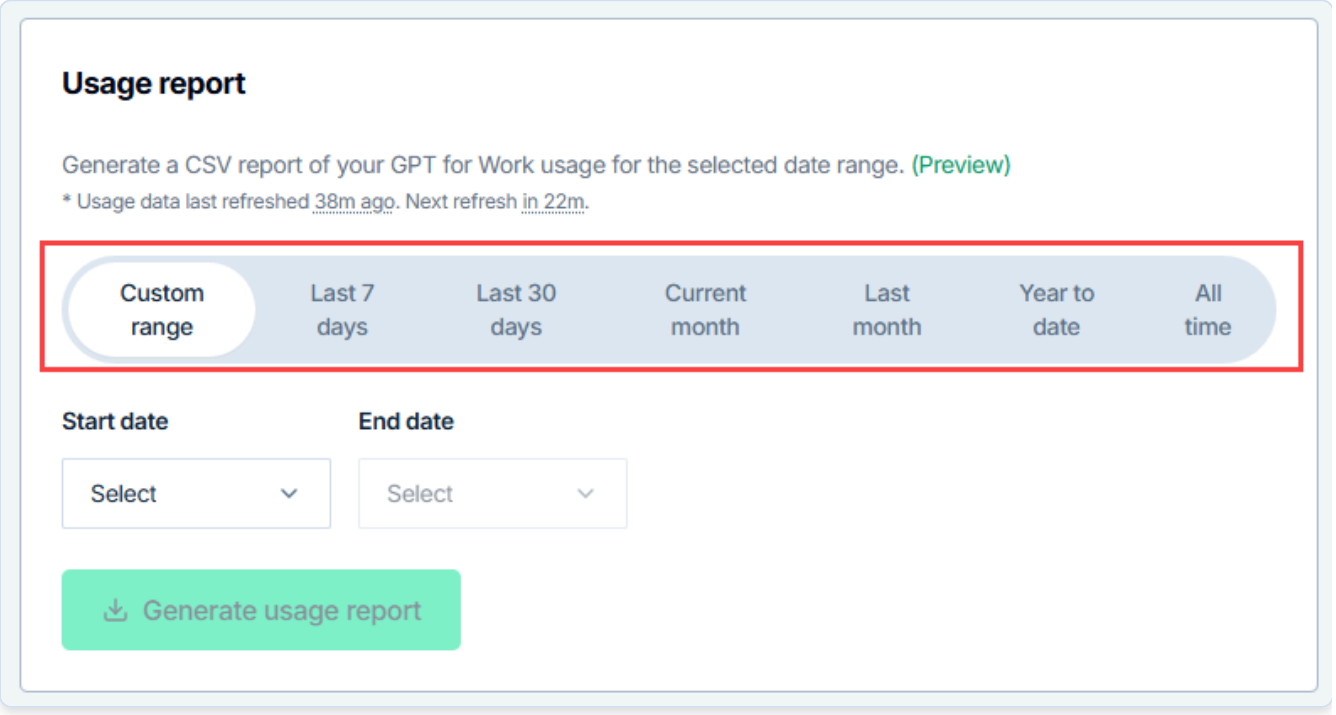
Generate a usage report

As a space owner or admin, you can generate and download CSV-format usage reports for your space. You can select a predefined time range or define a custom time range for a report. The reports provide detailed per-user information about feature and model usage, token consumption, and credit consumption.

Prerequisites

You are the owner or an admin of the space.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Usage**.
3. Select the time range for which you want to generate the report.

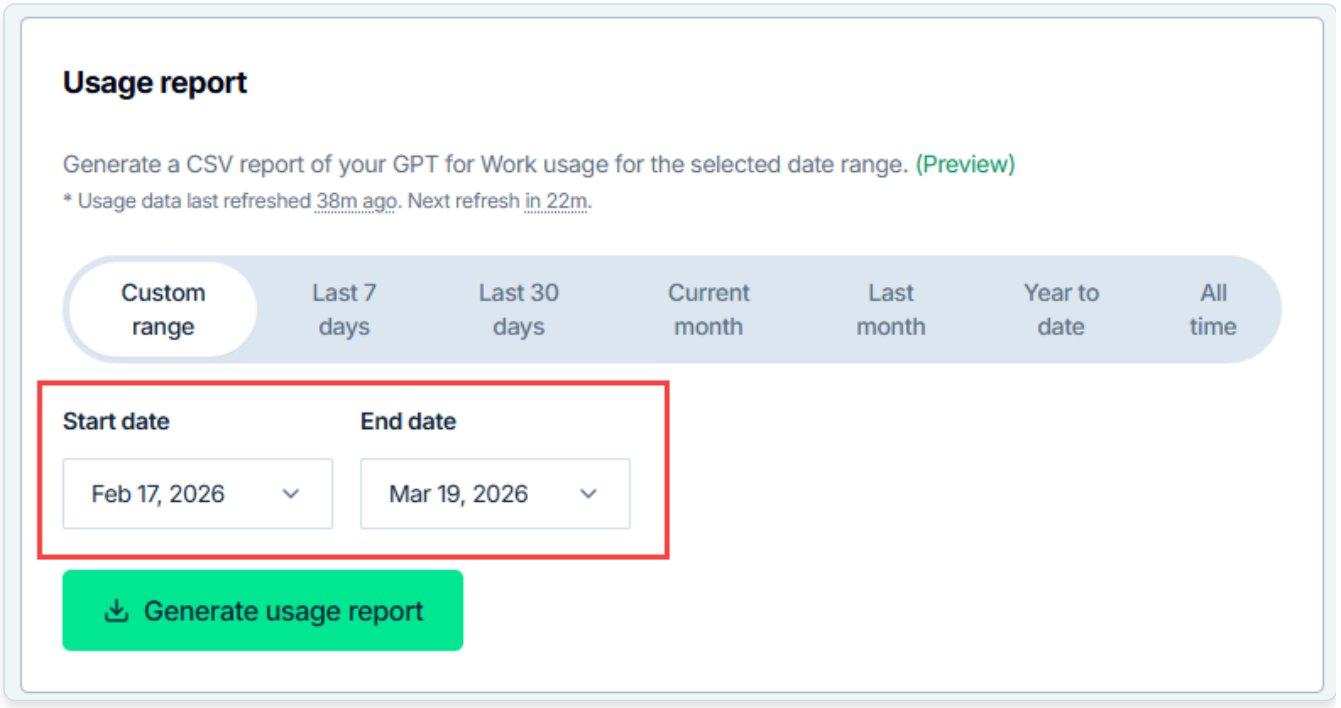


If you want to define a custom time range:

1. Select **Custom range**.
2. Select **Start date**.

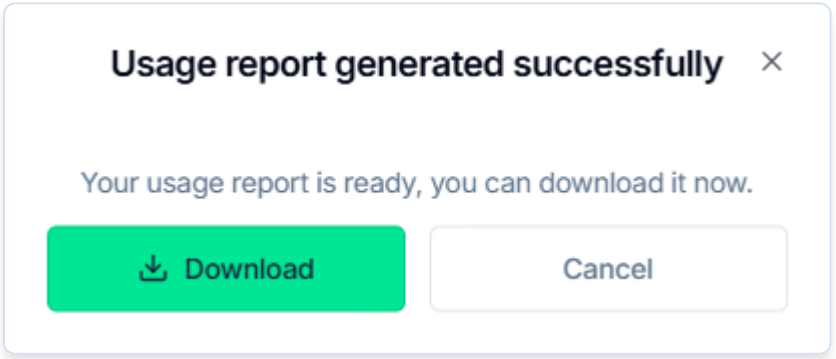
3 Select **End date**.

The range is inclusive, so the report will include usage data from both the start date and end date.



4 Click **Generate usage report**. The dashboard generates the report for the selected time range.

5 Click **Download**. Your browser downloads the CSV file to your computer.



You've generated and downloaded the usage report. You can now view it on your computer.

	A	B	C	D	E	F	G	H	I	J	K
1	calendar_day	space_id	user_email	document_id	feature	gpt_function	tool_name	service_tier	model	nb_executions	total_tokens
2	19.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	5	55471
3	19.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	18	207040
4	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	31	583003
5	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
6	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
7	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
8	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
9	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
10	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
11	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
12	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
13	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	9	218501
14	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
15	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
16	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
17	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
18	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
19	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
20	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
21	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
22	11.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	19	263140
23	11.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	19	272706
24	11.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	function	GPT		priority	gpt-4.1	14	2076
25	10.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	8	117248
26	10.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	14	260680
27	10.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	4	53985
28	05.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	function	GPT		priority	gpt-5.2	3	2964
29	23.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	27	391609
30	20.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	8	118653
31	19.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	21	297865
32	18.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	27	390988

What's next

- Buy extra-usage credits.
- Manage users in your team.

Pay-as-you-go billing (legacy)

This section applies to pay-as-you-go spaces. If you're on a subscription plan, see Subscription billing instead.

Buy credits

Buy credits to use the GPT for Work add-ons.

Set spending limits

Set monthly spending limits for your space and for individual users from the GPT for Work dashboard.

Monitor your usage

Track the credits your pay-as-you-go space has left, and generate usage reports.

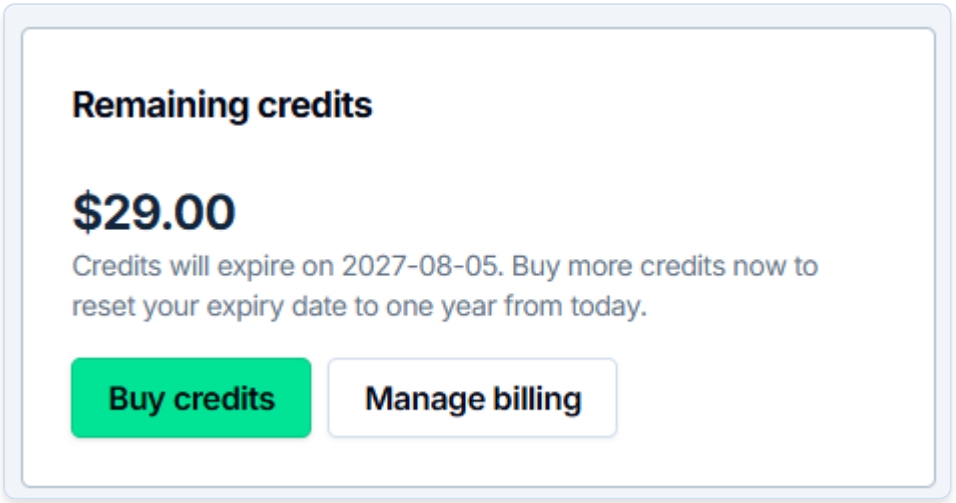
Download invoices

Download invoices and receipts for credit pack purchases from the GPT for Work dashboard.

Buy credits

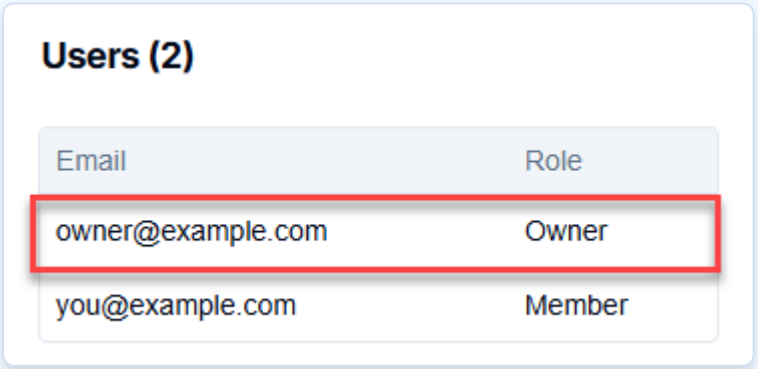
You can buy credits at any time to use the GPT for Work Google add-on (GPT for Sheets and Docs) or the GPT for Work Microsoft add-in (GPT for Excel Word). Credits expire one year after the last purchase.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the **Remaining credits** section, click **Buy credits**.

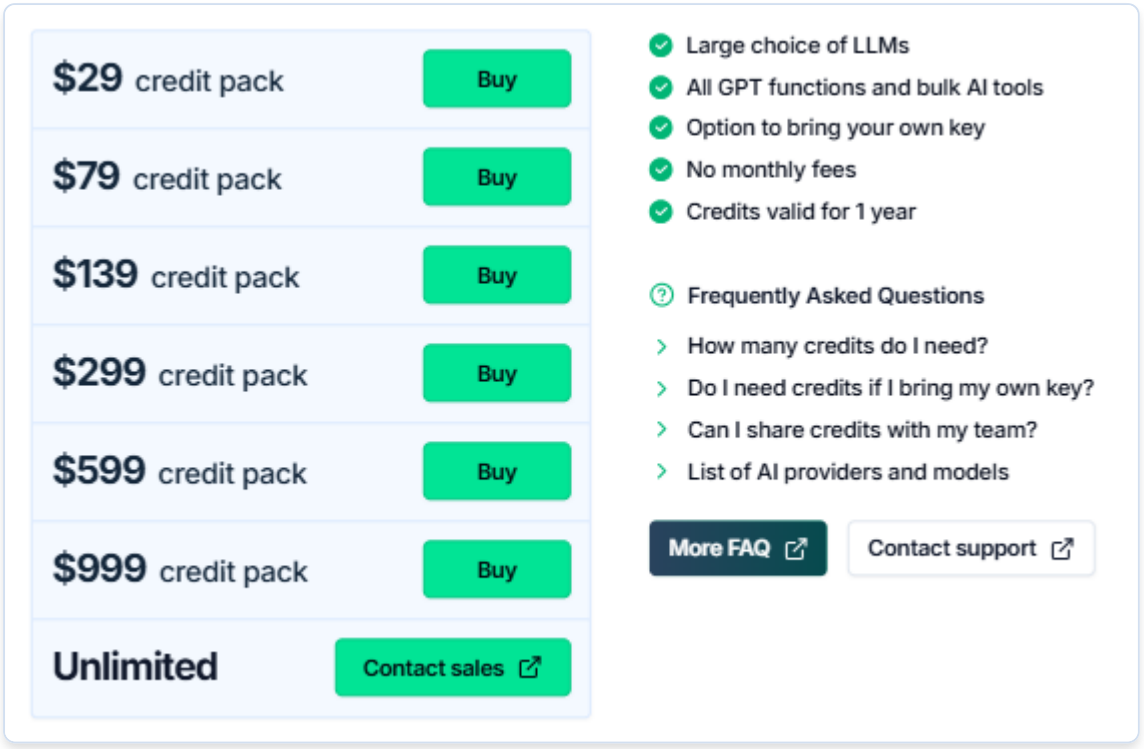


INFO

If this option is not available, you can ask the owner of the space to buy credits.



3 Click **Buy** for the credit pack that best suits your needs.



4 Enter your billing information and click **Go to checkout**.

5 Enter your payment details and click **Pay**.

Once your payment is processed, your credits increase by the purchased amount.

> **TIP**
>
> To estimate your credit usage based on common use cases, try our estimation tools.

What's next

- Download an invoice or a receipt for a credit pack purchase.

Set spending limits

Set monthly spending limits to control how much can be spent in your space each month. You can set a limit for the entire space and individual limits per user.

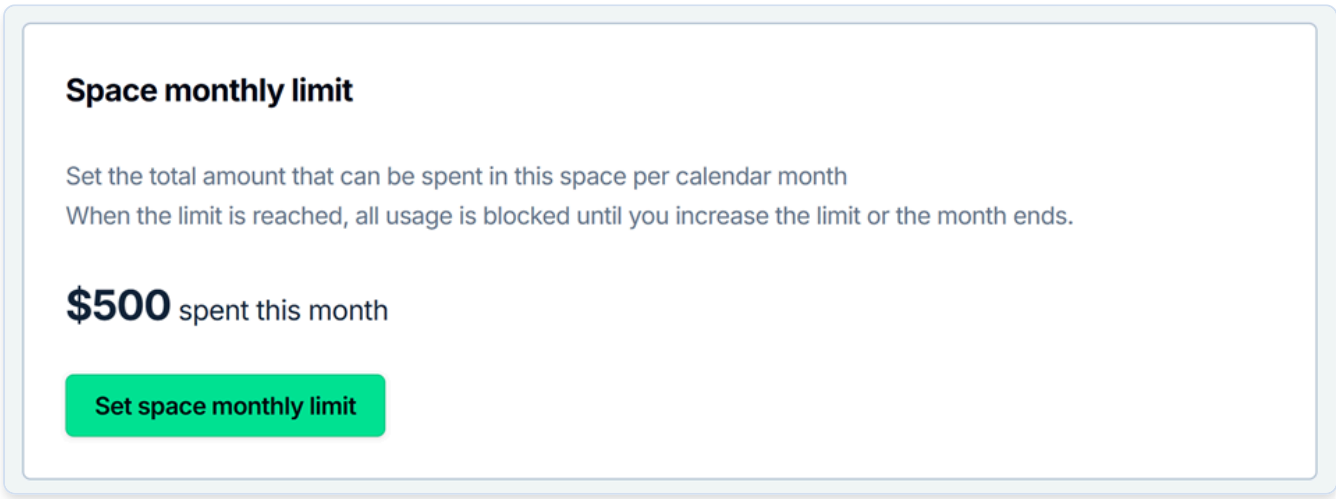
Prerequisites

Only the space owner or a space admin can set spending limits for a space.

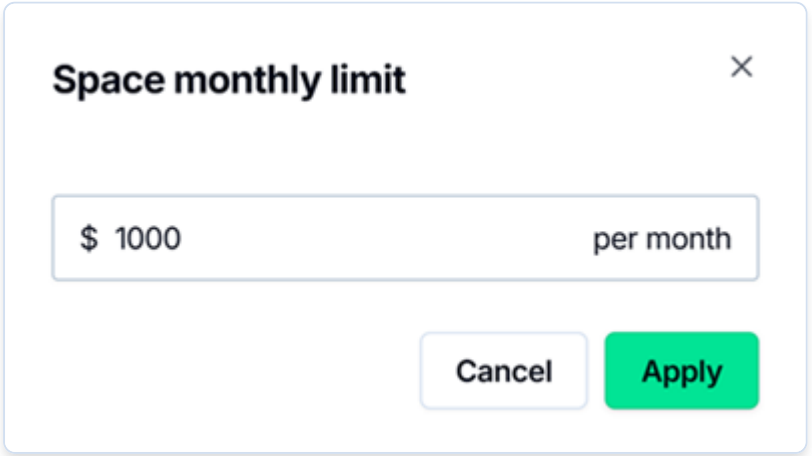
Set a space monthly limit

A space monthly limit restricts the total amount that can be spent across all users in your space per calendar month. When the limit is reached, all usage is blocked until you raise the limit or the limit resets at the end of the calendar month.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Limits**.
3. In the **Space monthly limit** section, click **Set space monthly limit**.



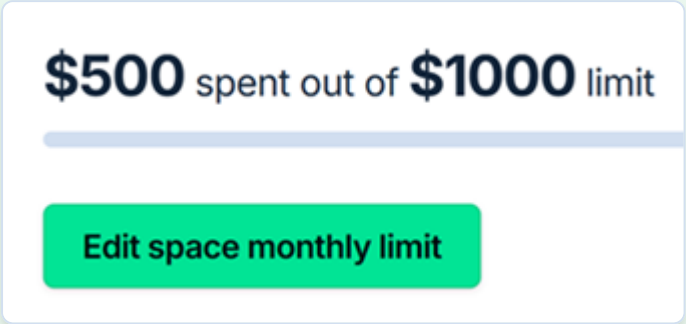
4. Enter the amount, and click **Apply**.



The space monthly limit immediately applies to all users in the space.

EDIT OR REMOVE

- To edit the space monthly limit, click **Edit space monthly limit**, edit the amount, and click **Apply**.
- To remove the space monthly limit, click **Edit space monthly limit** and click **Remove**.



Set user monthly limits

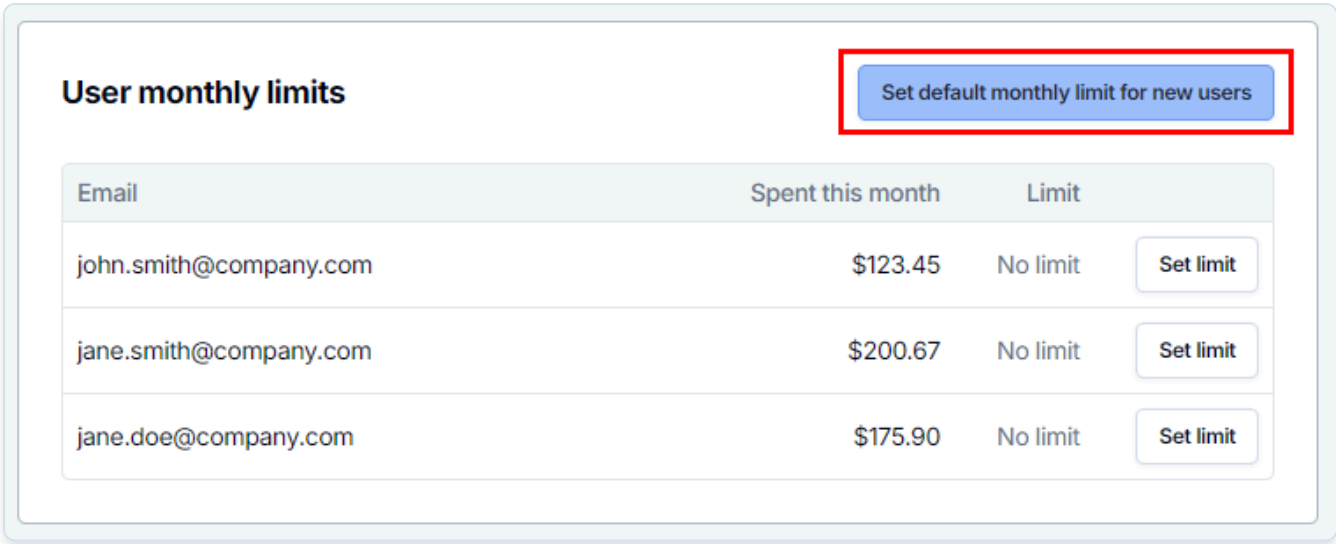
User monthly limits restrict the amount that an individual user can spend per calendar month. When a user reaches their limit, that user's usage is blocked until the limit is raised or the limit resets at the end of the calendar month. Other users in the space can continue using the add-ons normally.

Set a default limit for future users

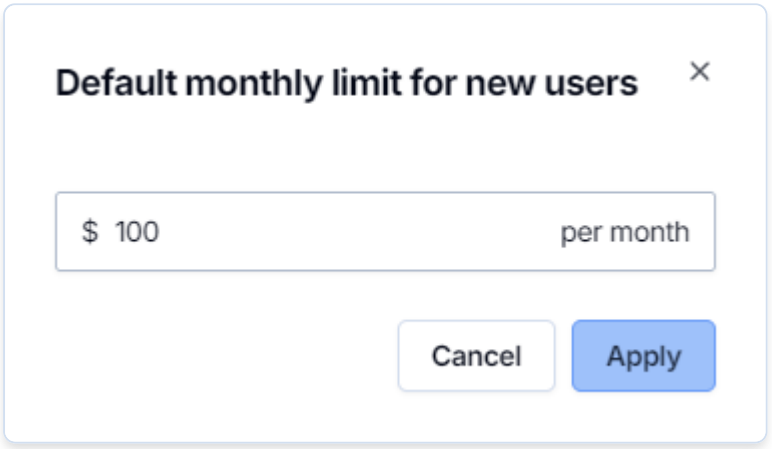
You can set a default monthly limit that applies to all new users who join the space. You can adjust the monthly limit for individual users after they join.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Limits**.

3 In the **User monthly limits** section, click **Set default monthly limit for new users**.



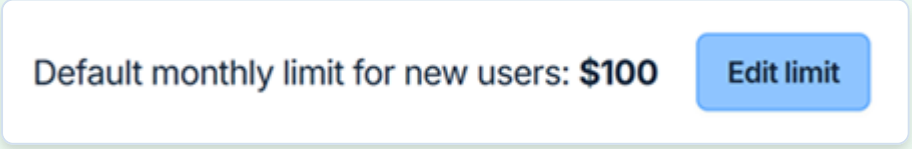
4 Enter the amount, and click **Apply**.



The default monthly limit automatically applies to all users who join the space after this point.

EDIT OR REMOVE

- To edit the default monthly limit, click **Edit limit**, edit the amount, and click **Apply**.
- To remove the default monthly limit, click **Edit limit** and click **Remove**.

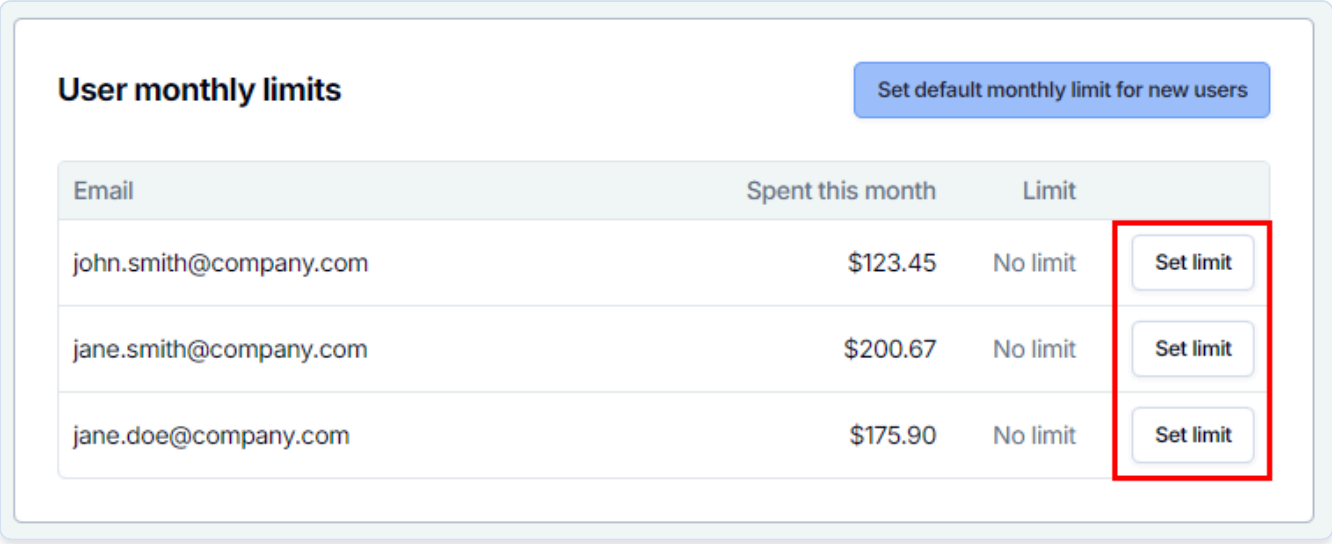


Set limits for current users

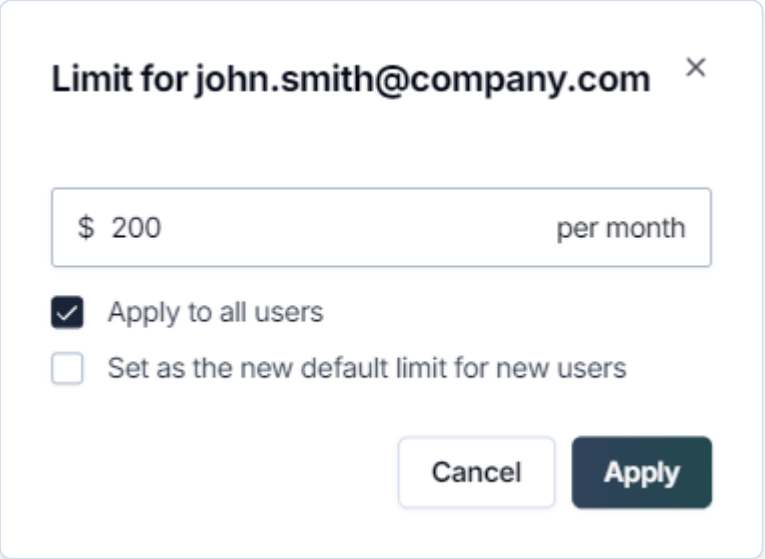
You can set monthly limits for individual users in your space. You can also apply the same limit to all current users in the space in one go.

1 Sign in to the GPT for Work dashboard with the account that you use for your space.

2. In the sidebar, select **Limits**.
3. In the **User monthly limits** section, find the user and click **Set limit**.



4. Enter the amount.
5. (Optional) Select **Apply to all users** to apply the same limit to all current users in the space.

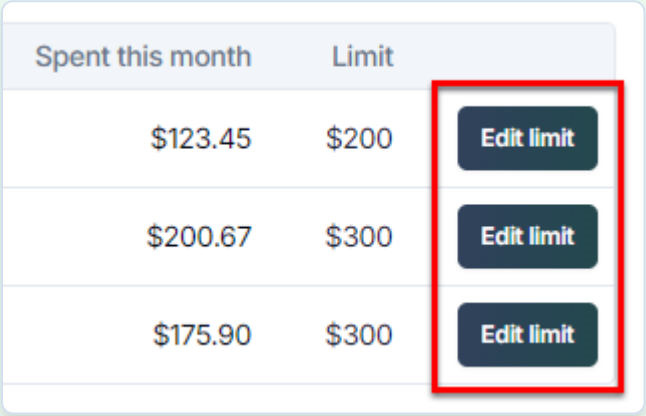


6. Click **Apply**.

The user monthly limit immediately applies.

EDIT OR REMOVE

- To edit monthly limits:
  - Single user: Click **Edit limit** for the user, edit the amount, and click **Apply**.
  - All users: Click **Edit limit** for any user, edit the amount, select **Apply to all users**, and click **Apply**.
- To remove monthly limits:
  - Single user: Click **Edit limit** for the user and click **Remove**.
  - All users: Click **Edit limit** for any user, select **Apply to all users**, and click **Remove**.



What's next

- Manage users in your team.
- Buy more credits.

Monitor your usage

With legacy pay-as-you-go pricing, your space consumes the credits in the credit packs it has bought. Track how much is left on the **Billing** page, and generate detailed usage reports from the GPT for Work dashboard.

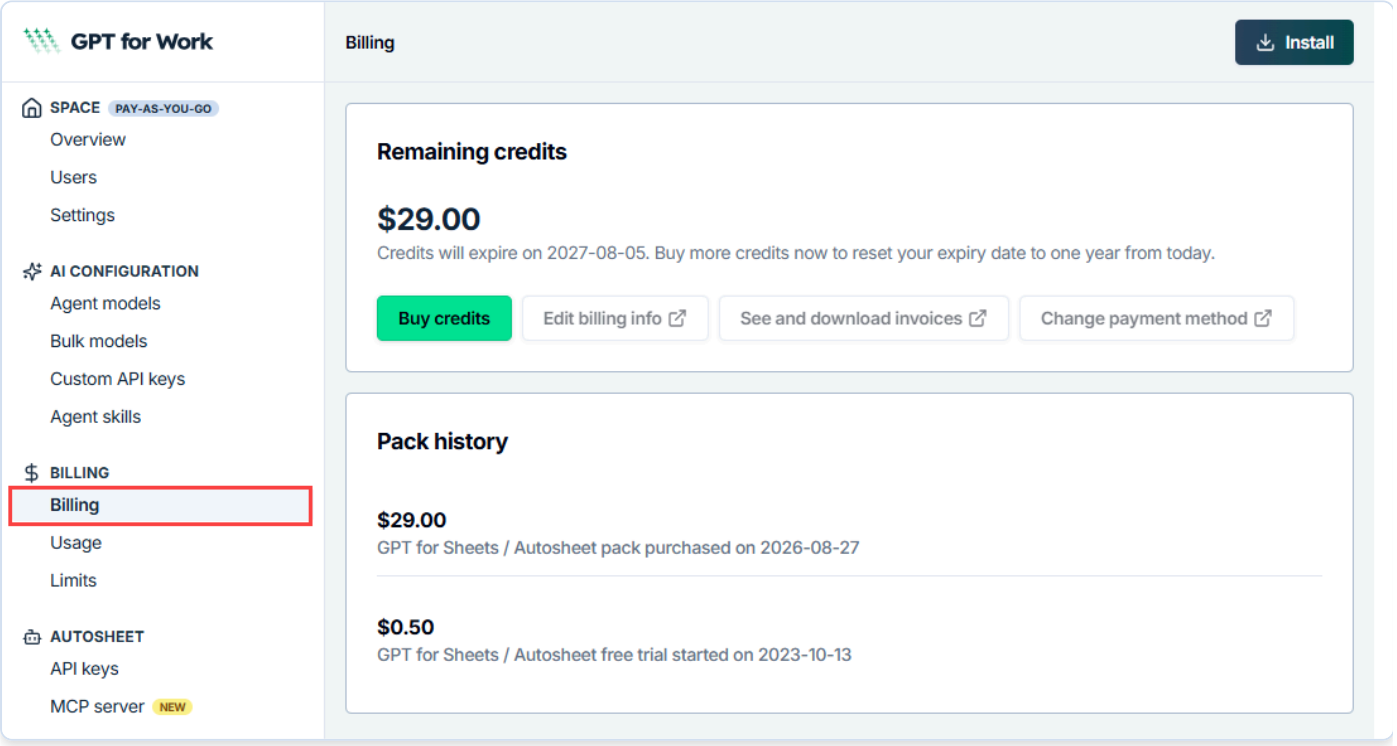
Track your usage

Prerequisites

You are the owner or an admin of the space.

1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.

The **Remaining credits** section shows how many credits you have left, and when they expire.



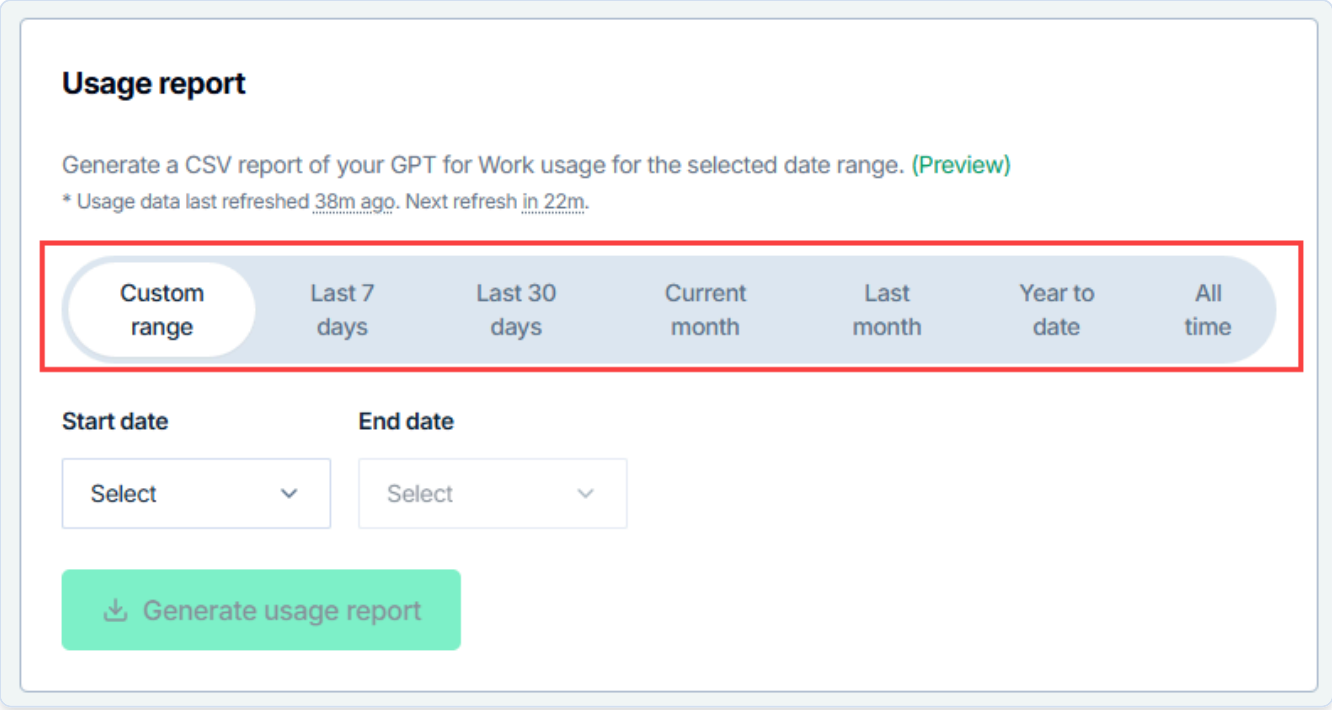
Generate a usage report

As a space owner or admin, you can generate and download CSV-format usage reports for your space. You can select a predefined time range or define a custom time range for a report. The reports provide detailed per-user information about feature and model usage, token consumption, and credit consumption.

Prerequisites

You are the owner or an admin of the space.

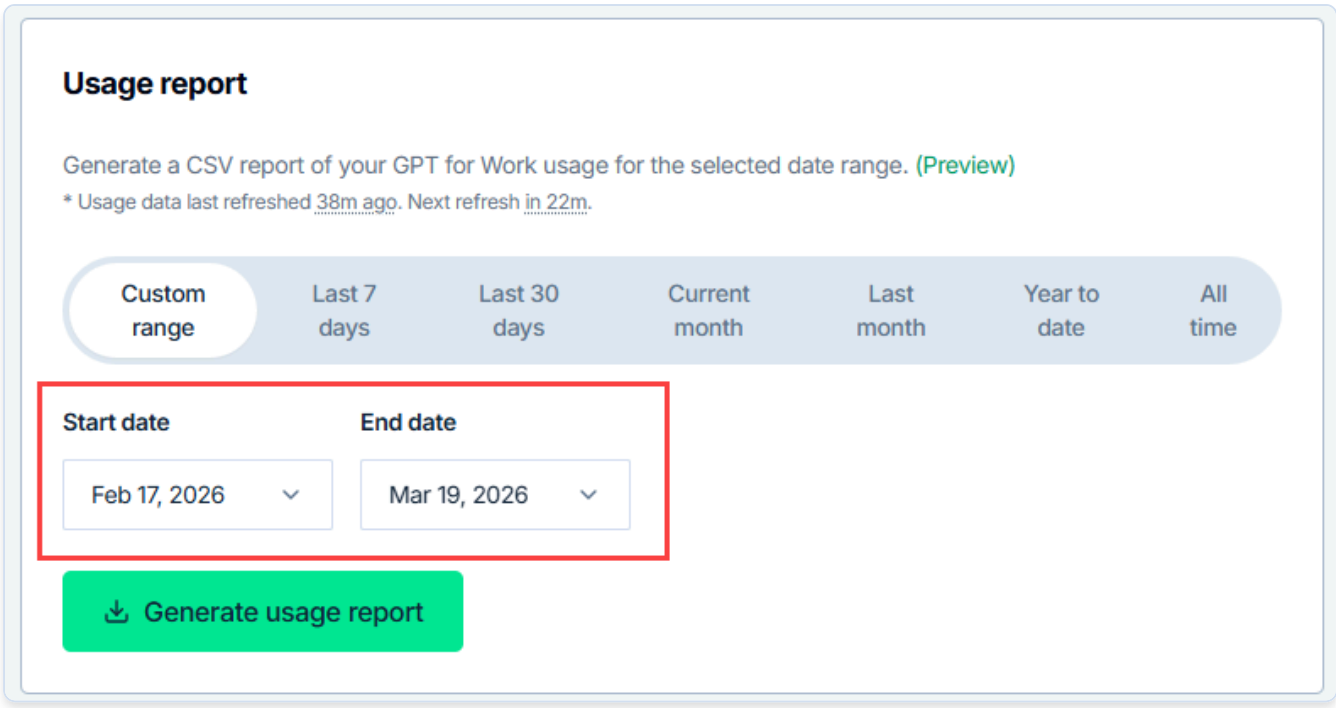
1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Usage**.
3. Select the time range for which you want to generate the report.



If you want to define a custom time range:

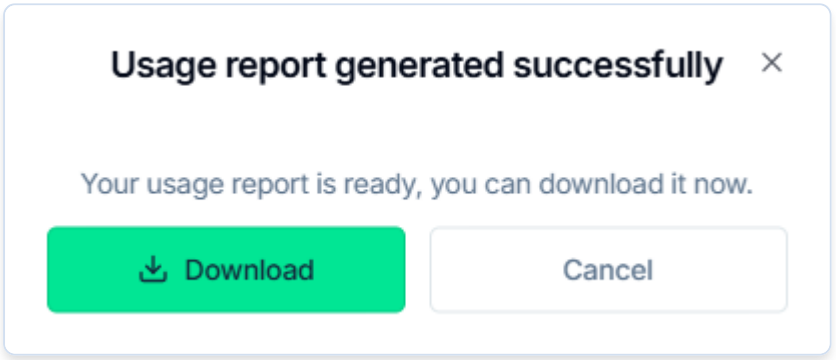
1. Select **Custom range**.
2. Select **Start date**.
3. Select **End date**.

The range is inclusive, so the report will include usage data from both the start date and end date.



4 Click **Generate usage report**. The dashboard generates the report for the selected time range.

5 Click **Download**. Your browser downloads the CSV file to your computer.



You've generated and downloaded the usage report. You can now view it on your computer.

	A	B	C	D	E	F	G	H	I	J	K
1	calendar_day	space_id	user_email	document_id	feature	gpt_function	tool_name	service_tier	model	nb_executions	total_tokens
2	19.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	5	55471
3	19.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	18	207040
4	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	31	583003
5	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
6	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
7	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
8	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
9	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
10	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
11	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
12	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	3	36614
13	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	9	218501
14	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
15	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
16	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
17	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
18	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
19	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
20	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
21	18.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.4	3	31470
22	11.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	19	263140
23	11.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	19	272706
24	11.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	function	GPT		priority	gpt-4.1	14	2076
25	10.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	8	117248
26	10.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop				claude-opus-4-6	14	260680
27	10.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	agent_loop			priority	gpt-5.1	4	53985
28	05.03.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH	function	GPT		priority	gpt-5.2	3	2964
29	23.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	27	391609
30	20.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	8	118653
31	19.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	21	297865
32	18.02.2026	xxxxxxxx-xxxx-xxxx-xxxx-	john.smith@company.co	Zr9KfQ2mXvL8-TeWp3aH					gpt-5.1	27	390988

What's next

- Buy credits.
- Manage users in your team.
- Download invoices and receipts.

Download credit pack invoices and receipts

Download invoices and receipts for past credit pack purchases for the GPT for Work Google add-on (GPT for Sheets and Docs) or the GPT for Work Microsoft add-in (GPT for Excel Word).

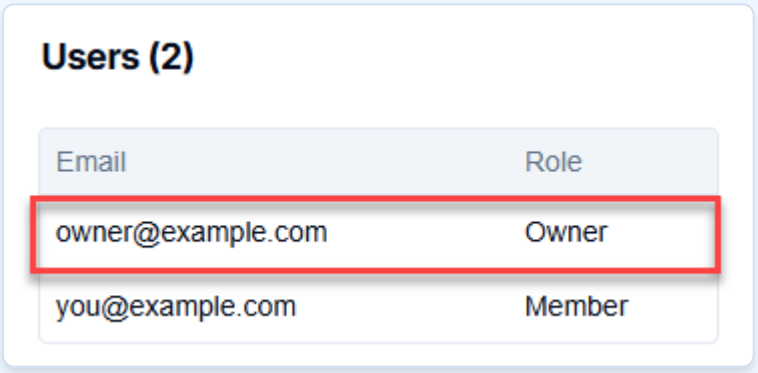
Prerequisites

You have purchased a credit pack.

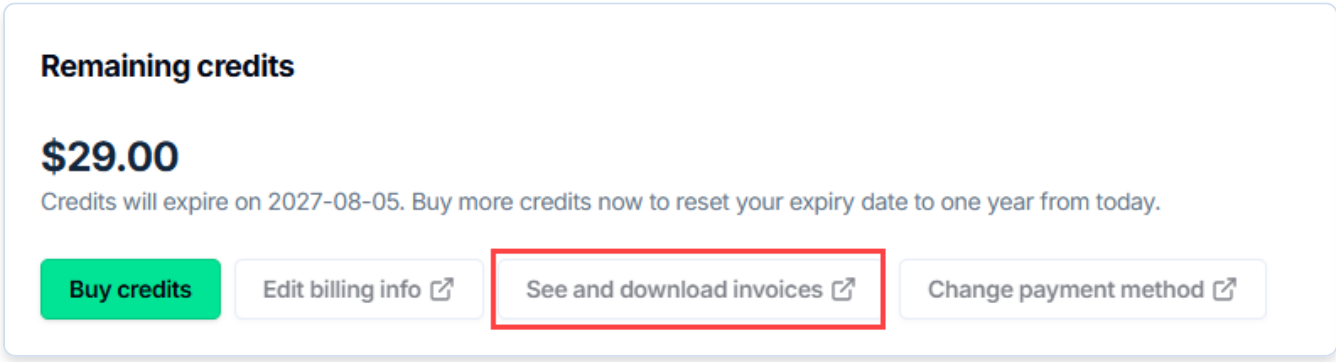
1. Sign in to the GPT for Work dashboard with the account that you use for your space.
2. In the sidebar, select **Billing**.

   INFO

   If this option is not available, you can ask the owner of the space for an invoice or a receipt.

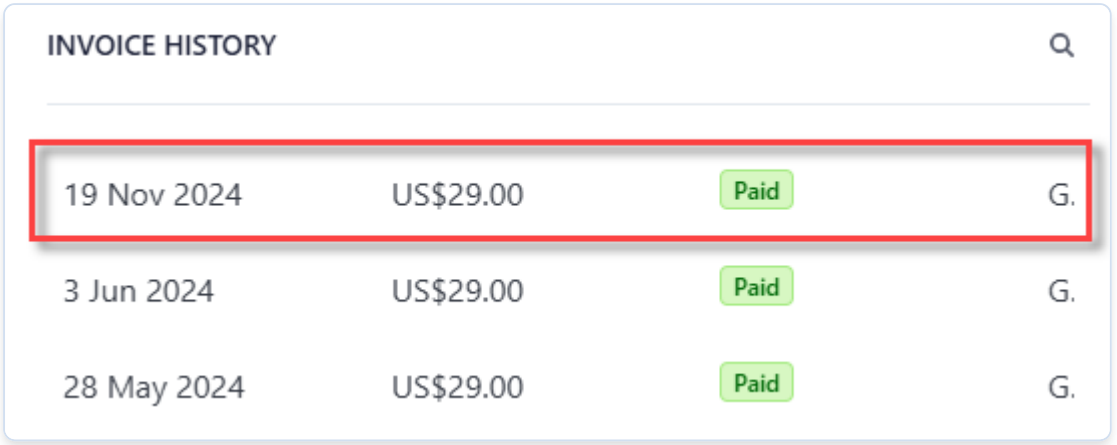
   

3. Click **See and download invoices**.

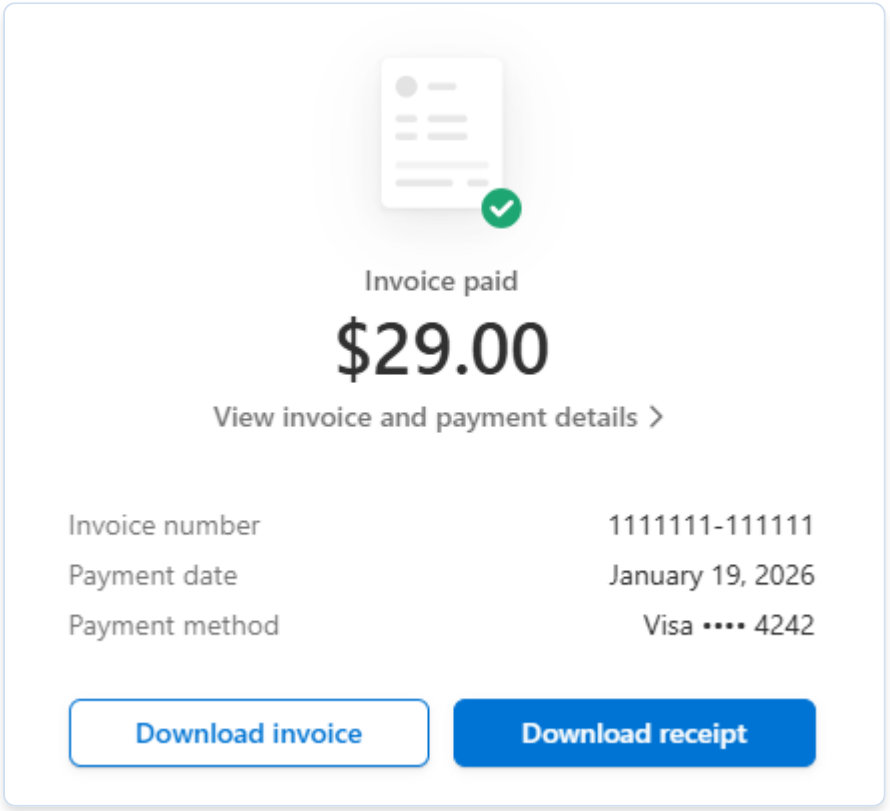
   

4. In the **Invoice history** section, click on the purchase whose invoice or receipt you want to

download.



5 Click **Download invoice** or **Download receipt**.



What's next

- Buy credits

Resources

Explore release notes, supported AI models, security and privacy, pricing, and user guides for GPT for Work.

Release notes

Stay informed about the latest GPT for Work product updates on Sheets, Excel, Word, and Docs: New models, new features, improvements, and fixes.

Models

Explore the AI providers and models supported by GPT for Work, and learn how rate limits and tokens affect your usage.

Security and privacy FAQ

Learn how Talarian, the maker of GPT for Work, secures your data, protects your privacy, and complies with international standards and regulations such as ISO 27001 and GDPR.

Pricing

Explore GPT for Work subscription plans, estimate your costs, and find information about legacy pay-as-you-go pricing.

Guides

Explore Agent use cases, familiarize yourself with basic AI concepts, and learn how to write better prompts and scrape the web from Google Sheets.

Models

Explore the AI providers and models supported by GPT for Work, and learn how rate limits and tokens affect your usage.

AI providers and models

Explore the AI providers and models supported by GPT for Work on Microsoft Excel and Word, and Google Sheets and Docs.

Recommended models

Find the recommended built-in models for GPT for Sheets and GPT for Excel, and choose the one that best fits your use case and budget.

Rate limits

Learn how GPT for Work relies on AI services that enforce rate limits.

Tokens

Learn what tokens are, how they work, and how they affect your credit usage in GPT for Work add-ons.

AI providers and models supported by GPT for Work

GPT for Work supports two types of models:

- **Agent models** are used by the Agent to handle chat messages, create formulas, and execute complex tasks by breaking them down into smaller tasks and delegating these to bulk models.
- **Bulk models** handle bulk tasks and web search tasks. The Agent uses bulk models when needed, and you use bulk models with bulk tools and GPT functions. The Agent can only use the bulk models you set in the GPT for Work dashboard.

> **INFO**
>
> For information about the prices for each model, see How credits are consumed.

Agent models

Built-in and provider-connected Agent models

The following table lists the **Agent models** available in GPT for Work, as built-in and provider-connected.

Provider	Model	Built-in	Provider-connected 🔑
Anthropic	Opus 5.5	Yes	Yes
Anthropic	Sonnet 5.5	Yes	Yes
Anthropic	Haiku 5.5	Yes	Yes
Anthropic	Fable 5.1	Yes	Yes
Anthropic	Fable 5	No	Yes
Anthropic	Opus 5	Yes	Yes
Anthropic	Opus 4.8	Yes	Yes
Anthropic	Opus 4.7	No	Yes
Anthropic	Opus 4.6	Yes	Yes
OpenAI	GPT-6 Astra	Yes	Yes

Provider	Model	Built-in	Provider-connected 🔑
OpenAI	GPT-6.1 Sol	Yes	Yes
OpenAI	GPT-6 Sol	Yes	Yes
OpenAI	GPT-6 Luna	Yes	Yes
OpenAI	GPT-5.6 Sol	Yes	Yes
OpenAI	GPT-5.6 Luna	Yes	Yes
OpenAI	GPT-5.5	No	Yes
OpenAI	GPT-5.4	Yes	Yes
OpenAI	GPT-5.1	No	Yes

Endpoint-connected Agent models

In GPT for Excel, you can use models from OpenAI-compatible API endpoints as Agent models. The endpoint must be compatible with OpenAI's Responses API and be publicly accessible over the internet. Not all endpoints and models can run the Agent reliably: some endpoints implement the Responses API only partially, and smaller models may fail to complete Agent tasks. For setup instructions, see Connect to OpenAI-compatible API endpoints from Excel.

Models available through dedicated API endpoints do not consume credits, but require credits to be available.

Bulk models

The following tables list the **bulk models** available in GPT for Work, as built-in and connected.

Built-in bulk models

You can use models from OpenAI, Perplexity, Google, Anthropic, and Parallel.

Provider	Model	Reasoning	Vision	Web search	URL fetch
Search...	Search...	All	All	All	All
OpenAI	gpt-6.1-sol	Yes	Yes	No	No
OpenAI	gpt-6-luna	Yes	Yes	No	No

Provider	Model	Reasoning	Vision	Web search	URL fetch
OpenAI	gpt-5.6-sol	Yes	Yes	No	No
OpenAI	gpt-5.6-terra	Yes	Yes	No	No
OpenAI	gpt-5.6-luna	Yes	Yes	No	No
OpenAI	gpt-5.5	Yes	Yes	No	No
OpenAI	gpt-5.4	Yes	Yes	No	No
OpenAI	gpt-5.4-mini	Yes	Yes	No	No
OpenAI	gpt-5.4-nano	Yes	Yes	No	No
OpenAI	gpt-5.2	Yes	Yes	No	No
OpenAI	gpt-5.1	Yes	Yes	No	No
OpenAI	gpt-5-mini	Yes	Yes	No	No
OpenAI	gpt-5-nano	Yes	Yes	No	No
OpenAI	gpt-4.1	No	Yes	No	No
OpenAI	gpt-4.1-mini	No	Yes	No	No
OpenAI	gpt-4.1-nano	No	Yes	No	No
OpenAI	gpt-5-search-api	No	No	Yes	No
OpenAI	gpt-4o	No	Yes	No	No
OpenAI	gpt-4o-mini	No	No	No	No
OpenAI	gpt-5	Yes	Yes	No	No
OpenAI	gpt-4-turbo	No	No	No	No
OpenAI	o1	Yes	Yes	No	No
OpenAI	o3-mini	Yes	No	No	No
OpenAI	o3	Yes	Yes	No	No
OpenAI	o4-mini	Yes	Yes	No	No
OpenAI	gpt-3.5-turbo	No	No	No	No

Provider	Model	Reasoning	Vision	Web search	URL fetch
Perplexity	perplexity-fast	Yes	No	Yes	No
Perplexity	perplexity-low	Yes	No	Yes	No
Google	gemini-2.5-flash	Yes	No	Optional	Yes
Anthropic	claude-5.5-opus	Yes	Yes	No	No
Anthropic	claude-5.5-sonnet	Yes	Yes	No	No
Anthropic	claude-5.5-haiku	Yes	Yes	No	No
Anthropic	claude-5-opus	Yes	Yes	No	No
Anthropic	claude-5-sonnet	Yes	Yes	No	No
Anthropic	claude-4.8-opus	Yes	Yes	No	No
Anthropic	claude-4.7-opus	Yes	Yes	No	No
Anthropic	claude-4.6-opus	Yes	Yes	No	No
Anthropic	claude-4.6-sonnet	Yes	Yes	No	No
Anthropic	claude-4.5-opus	Yes	Yes	No	No
Anthropic	claude-4.5-sonnet	No	Yes	No	No
Anthropic	claude-4.5-haiku	No	Yes	No	No
Parallel	speed	No	No	Yes	No

> **INFO**
>
> GPT for Docs ONLY supports OpenAI and Anthropic models.

Provider-connected bulk models

You can use models from OpenAI, Perplexity, Google, Anthropic, DeepSeek, Mistral, and xAI.

Provider	Model	Reasoning	Vision	Web search	URL fetch
Search...	Search...	All	All	All	All
OpenAI	gpt-6.1-sol	Yes	Yes	No	No
OpenAI	gpt-6-astra	Yes	Yes	No	No
OpenAI	gpt-6-luna	Yes	Yes	No	No
OpenAI	gpt-5.6-sol	Yes	Yes	No	No
OpenAI	gpt-5.6-terra	Yes	Yes	No	No
OpenAI	gpt-5.6-luna	Yes	Yes	No	No
OpenAI	gpt-5.5	Yes	Yes	No	No
OpenAI	gpt-5.4	Yes	Yes	No	No
OpenAI	gpt-5.4-mini	Yes	Yes	No	No
OpenAI	gpt-5.4-nano	Yes	Yes	No	No
OpenAI	gpt-5.2	Yes	Yes	No	No
OpenAI	gpt-5.1	Yes	Yes	No	No
OpenAI	gpt-5-mini	Yes	Yes	No	No
OpenAI	gpt-5-nano	Yes	Yes	No	No
OpenAI	gpt-4.1	No	Yes	No	No
OpenAI	gpt-4.1-mini	No	Yes	No	No
OpenAI	gpt-4.1-nano	No	Yes	No	No
OpenAI	gpt-5-search-api	No	No	Yes	No

Provider	Model	Reasoning	Vision	Web search	URL fetch
OpenAI	gpt-4o	No	Yes	No	No
OpenAI	gpt-4o-mini	No	No	No	No
OpenAI	gpt-5	Yes	Yes	No	No
OpenAI	gpt-4-turbo	No	No	No	No
OpenAI	o1	Yes	Yes	No	No
OpenAI	o3-mini	Yes	No	No	No
OpenAI	o3	Yes	Yes	No	No
OpenAI	o4-mini	Yes	Yes	No	No
OpenAI	gpt-3.5-turbo	No	No	No	No
OpenAI	gpt-3.5-turbo (fine-tuned)	No	No	No	No
OpenAI	gpt-4o (fine-tuned)	No	No	No	No
OpenAI	gpt-4o-mini (fine-tuned)	No	No	No	No
OpenAI	gpt-4.1 (fine-tuned)	No	Yes	No	No
OpenAI	assistant	No	No	No	No
Perplexity	perplexity-fast	Yes	No	Yes	No
Perplexity	perplexity-low	Yes	No	Yes	No
Perplexity	perplexity-medium	Yes	No	Yes	No
Google	gemini-3.5-flash	Yes	No	Optional	Yes
Google	gemini-3.1-flash-lite	Yes	No	Optional	Yes
Google	gemini-3-flash-preview	Yes	No	Optional	Yes

Provider	Model	Reasoning	Vision	Web search	URL fetch
Google	gemini-3.1-pro-preview	Yes	No	Optional	Yes
Google	gemini-2.5-flash	Yes	No	Optional	Yes
Google	gemini-2.5-flash-lite	Yes	No	Optional	Yes
Google	gemini-2.5-pro	Yes	No	Optional	Yes
xAI	grok-4-1-fast	Yes	Yes	No	No
xAI	grok-4-1-fast-non-reasoning	No	Yes	No	No
Anthropic	claude-5.5-opus	Yes	Yes	No	No
Anthropic	claude-5.5-sonnet	Yes	Yes	No	No
Anthropic	claude-5.5-haiku	Yes	Yes	No	No
Anthropic	claude-fable-5.1	Yes	Yes	No	No
Anthropic	claude-fable-5	Yes	Yes	No	No
Anthropic	claude-5-opus	Yes	Yes	No	No
Anthropic	claude-5-sonnet	Yes	Yes	No	No
Anthropic	claude-4.8-opus	Yes	Yes	No	No
Anthropic	claude-4.7-opus	Yes	Yes	No	No
Anthropic	claude-4.6-opus	Yes	Yes	No	No
Anthropic	claude-4.6-sonnet	Yes	Yes	No	No

Provider	Model	Reasoning	Vision	Web search	URL fetch
Anthropic	claude-4.5-opus	Yes	Yes	No	No
Anthropic	claude-4.5-sonnet	No	Yes	No	No
Anthropic	claude-4.1-opus	No	Yes	No	No
Anthropic	claude-4.5-haiku	No	Yes	No	No
DeepSeek	deepseek-v4-pro	Yes	No	No	No
DeepSeek	deepseek-v4-flash	Yes	No	No	No
Mistral	mistral-large	No	No	No	No
Mistral	mistral-medium	No	Yes	No	No
Mistral	mistral-small	No	No	No	No
Mistral	ministral-8b	No	No	No	No
Mistral	ministral-3b	No	No	No	No

INFO

GPT for Docs ONLY supports OpenAI and Anthropic models.

Endpoint-connected bulk models

You can use any OpenAI-compatible API endpoint with GPT for Excel and Word. For bulk models, the endpoint must be compatible with OpenAI's Chat Completions API. You can connect to two main types of services:

- **Cloud-based AI platforms** provide access to models over the internet with no software installation or setup required on your part. Popular examples include Anyscale, Azure AI, Fireworks AI, and Together AI. The available models vary from platform to platform.

- **Local LLM servers** run on a local machine, such as your own computer or another computer on a local network. Popular examples include LM Studio, LocalAI, and Open WebUI. The available models depend on what's installed on the server you're using.

For setup instructions, see Connect to OpenAI-compatible API endpoints from Excel or Word.

Models available through dedicated API endpoints do not consume credits, but require credits to be available.

What's next

- Recommended models
- Model prices
- Manage available models
- Estimate your costs

Recommended models

Find the recommended built-in models for GPT for Sheets and GPT for Excel, and choose the one that best fits your use case and budget. For more information about the full range of supported models, see AI providers and models supported by GPT for Work.

Agent models

Best	Balanced	Cheapest
Opus 5.5 (Medium)	GPT-6 Luna (High)	GPT-5.4 (Low)

For the benchmark results behind these recommendations, see Agent model benchmark.

Examples

- For simple or straightforward tasks such as creating formulas or fixing errors, use **GPT-6 Luna** at **High** reasoning effort.
- For high-level tasks such as financial modeling, overall formatting, or visual tasks, use **Opus 5.5** at **Medium** reasoning effort.

Bulk models

The same bulk models are recommended whether the Agent delegates a bulk task, or you run bulk tools or GPT functions yourself.

Use case	Best	Balanced	Cheapest
Content generation, translation, categorization, extraction, scoring	gpt-5.6-sol	gpt-5.4-mini	gpt-5.4-nano
Enrichment from the web	gpt-5-search-api	perplexity-low	perplexity-fast

Agent model benchmark

To compare Agent models, we run an internal benchmark of roughly 200 everyday spreadsheet tasks, such as writing a formula, inserting a column, creating a chart, or formatting a header. Each

model runs every task once with the same spreadsheets and prompts, and each result is scored as a pass or a fail. Newly released models are always evaluated in real-world use before we recommend them, even when they score well in the benchmark.

We update the benchmark regularly. The latest run is from September 2026 and tested several model configurations in GPT for Excel, six of which this page compares:

- **Opus 5.5** at **Medium** reasoning effort achieved the best overall performance and ranked first or tied for first in six of the eight task categories.
- **GPT-6 Luna** at **High** reasoning effort offered the best balance of score and cost. It matched the score of Opus 4.8 while consuming about half as many credits per task as Opus 5.5.

> **NOTE**
>
> The benchmark ran in GPT for Excel only. The Agent uses different tools in Google Sheets, so results in Sheets may differ.

Overall scores

The following figure shows the share of tasks each model completed successfully.



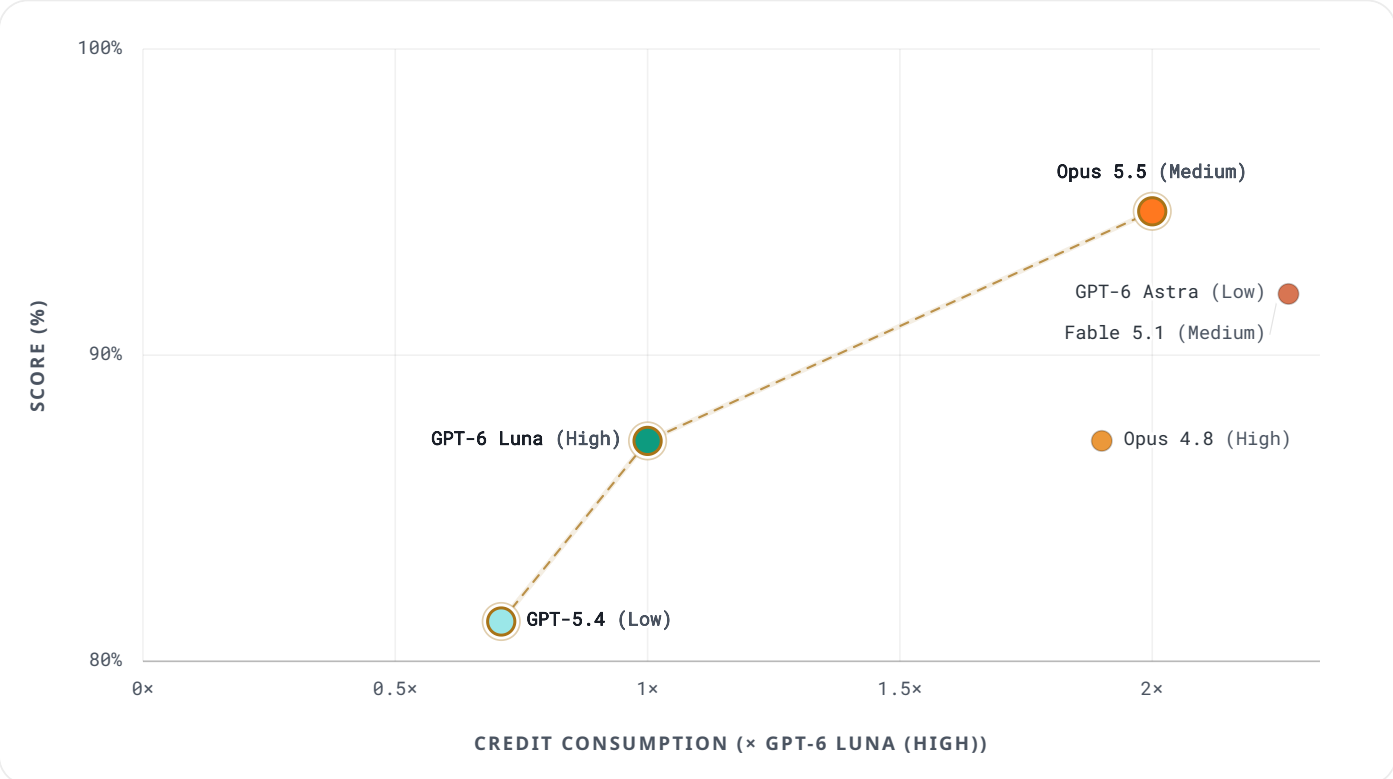
Click here for the chart data as a table ▼

Model	Reasoning effort	Score
Opus 5.5	Medium	94.7%

Model	Reasoning effort	Score
GPT-6 Astra	Low	92.0%
Fable 5.1	Medium	92.0%
Opus 4.8	High	87.2%
GPT-6 Luna	High	87.2%
GPT-5.4	Low	81.3%

Scores vs. credit consumption

The following figure shows each model's score against its credit consumption per task, relative to GPT-6 Luna at High reasoning effort. Credit consumption covers the Agent model only, not the bulk models the Agent delegates to, and assumes built-in models on a subscription. For example, Opus 5.5 consumes about twice as many credits per task as GPT-6 Luna. The dashed line connects the models that offer the best trade-off: Every model off the line is matched or outscored by a model that consumes fewer credits.

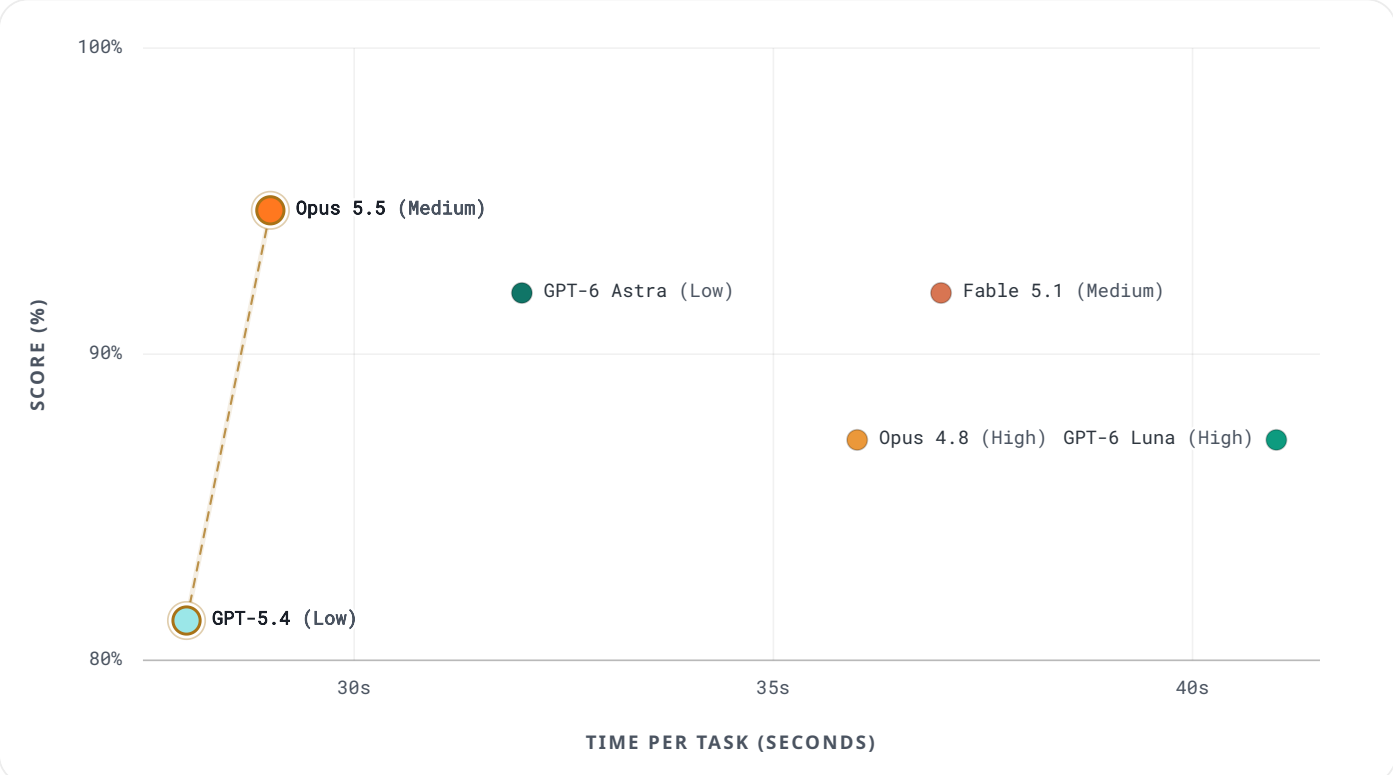


Click here for the chart data as a table ▼

Model	Reasoning effort	Score	Credit consumption relative to GPT-6 Luna (High)
Opus 5.5	Medium	94.7%	2×
GPT-6 Astra	Low	92.0%	2.27×
Fable 5.1	Medium	92.0%	2.27×
Opus 4.8	High	87.2%	1.9×
GPT-6 Luna	High	87.2%	1×
GPT-5.4	Low	81.3%	0.71×

Scores vs. speed

The following figure shows each model's score against the average time it took to complete a task. Shorter times are better. The dashed line connects the models that offer the best trade-off: Every model off the line is matched or outscored by a faster model.



Click here for the chart data as a table ▼

Model	Reasoning effort	Score	Average time per task
Opus 5.5	Medium	94.7%	29s
GPT-6 Astra	Low	92.0%	32s
Fable 5.1	Medium	92.0%	37s
Opus 4.8	High	87.2%	36s
GPT-6 Luna	High	87.2%	41s
GPT-5.4	Low	81.3%	28s

Rankings by task category

The following table shows how the models rank in each task category. Models with the same score in a category share a rank.

Task category	● Opus 5.5 Medium	● GPT-6 Astra Low	● Fable 5.1 Medium	● Opus 4.8 High	● GPT-6 Luna High	● GPT-5.4 Low
Bulk content generation	**1st**	2nd	2nd	4th	4th	3rd
Bulk data processing	2nd	**1st**	2nd	2nd	2nd	3rd
Charts	**1st**	2nd	3rd	4th	3rd	5th
Data operations	**1st**	2nd	**1st**	2nd	3rd	4th
Formatting	**1st**	2nd	2nd	3rd	4th	3rd
Formulas	2nd	2nd	**1st**	2nd	**1st**	3rd
Pivot tables	**1st**	**1st**	**1st**	**1st**	**1st**	**1st**
Spreadsheet manipulation	**1st**	2nd	3rd	3rd	4th	5th

Learn more

- AI providers and models supported by GPT for Work
- LiveBench (a good model benchmark)

Rate limits

GPT for Work extensions rely on services that enforce rate limits to ensure smooth and fair access to everyone.

If you choose not to set up API keys, you will share rate limits with other GPT for Work users. By setting API keys, you benefit from your own rate limits, which you can manage from these links:

- For OpenAI models
- For Anthropic models
- For Mistral AI models
- For Gemini models

Google rate limits:

If you use the GPT for Sheets and Docs add-on, you are also subject to the rate limits enforced by Google, based on your type of account:

	External service calls daily limits (URL fetch)
gmail.com	20k / day / gmail.com account
Google Workspace	100k / day / Google Workspace account

> **INFO**
>
> **Google's rate limits apply to all add-ons and scripts**, not just ours. If you use multiple scripts or add-ons, your usage counts towards the same rate limits. GPT for Excel and Word are not affected by Google's rate limits.

Tokens

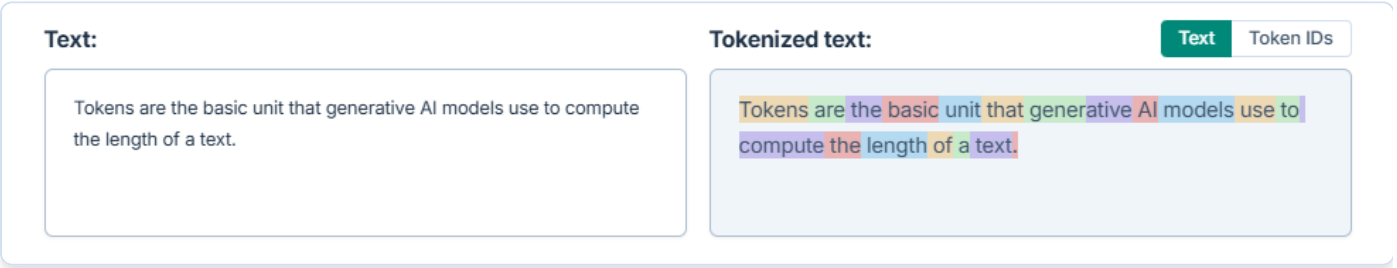
Tokens are the basic unit for measuring AI model usage. They directly affect your credit consumption.

What is a token?

A token is a small unit of text that a model processes. Tokens can be:

- Whole words (like "hello")
- Parts of a word (like "mor" and "ning")
- Punctuation marks (like "." or ",")
- Symbols (like "$" or "@")
- Spaces

The process of splitting text into tokens is called **tokenization**. The way text is split into tokens depends on the model's **tokenizer**, which is specific to each AI provider. To see how text is split into tokens, use our tokenizer tool.



Token counts for languages

Token counts vary by language because the way text is split into tokens differs from one language to another.

- **English:** 1 word ≈ 1.3 tokens
- **French/Spanish:** 1 word ≈ 1.4 tokens
- **Chinese:** 1 word ≈ 1.9 tokens

These figures are based on Beyond Fertility: Analyzing STRR as a Metric for Multilingual Tokenization Evaluation and OpenAI documentation.

> **INFO**
>
> These are rough estimates. Actual token counts vary based on the specific words, punctuation, and formatting in your text. To see how text is split into tokens, use our tokenizer tool.

What counts as tokens in GPT for Work?

GPT for Work uses the following types of tokens:

- **Input tokens:** The text you send to a model, including your prompt, cell values, global instructions, and any internal instructions used by the add-on. Images can also be used as input and are converted to tokens.
- **Output tokens:** The text the model generates in response, such as translations, summaries, or extracted data.

Reasoning models are trained to think before they answer, producing an internal chain of thought before responding to a prompt. They generate two types of output tokens:

- **Reasoning tokens** make up the model's internal chain of thought. These tokens are typically not visible in the output from most providers.
- **Completion tokens** make up the model's final visible response.

For more information, see How credits are consumed.

What's next

- How credits are consumed
- Model prices
- Estimate your usage

Security and privacy FAQ

In this article, we (Talarian, maker of GPT for Work) aim to answer the most frequent questions about the security and privacy measures we implement and our compliance with global regulations like the GDPR. For more details, check the Talarian security portal.

Compliance

Is your organization ISO 27001 certified?

Yes, we are certified under the ISO 27001 standard. The certificate is available through the Talarian security portal.

Do you comply with GDPR and other privacy regulations?

Yes, we comply with GDPR. As a data controller, we process personal data in a fair, transparent and secure way in accordance with our Privacy Policy. As a data processor, we process personal data according to our Data Processing Agreement.

Security

Do you conduct independent third-party security testing?

Yes, we rely on annual third-party security assessments. The resulting reports are available through the Talarian security portal. Also, our infrastructure is hosted on Google Cloud Platform, which undergoes its own independent security audits.

We are certified under the CASA Tier 3 certification for Google, ensuring GPT for Work meets Google's standards for security and privacy. Additionally, we have completed the Microsoft publisher attestation for the Microsoft 365 App Compliance Program.

Is my data encrypted at rest and in transit?

Yes. We rely on Google's encryption-at-rest mechanism (which is based on the AES algorithm with a key size of 256 bits). Google Cloud Platform's security infrastructure uses a common cryptographic library called Tink, which includes the FIPS 140-2 validated module (named BoringCrypto) to implement encryption consistently across Google Cloud.

For data in transit, we exclusively use TLS 1.2 and TLS 1.3 to ensure secure transmission. The following data is also encrypted when in use: API keys, bulk tools configuration history (when GPT

for Work is used with Google Sheets and Excel spreadsheets) and GPT functions cache (when GPT for Work is used with Google Sheets and cache is enabled).

What measures are in place to prevent unauthorized access to customer data?

We restrict and monitor access to production data on an as-needed basis. We use a robust access control policy with quarterly admin rights reviews. We enforce stringent password security policies and enforce Multi-factor Authentication (MFA) for all our personnel.

How do you ensure secured user identification?

We offer Single Sign-On (SSO) access to GPT for Work with Microsoft 365 (for GPT for Excel Word) or Google Workspace (for GPT for Sheets and Docs). User identification security levels are therefore those provided by either Microsoft or Google.

How do you handle software and hardware vulnerabilities?

With respect to the devices used to develop our application: hardware and software provided to our employees are controlled under a centralized Mobile Device Management (MDM) solution. This allows us to control our fleet and its usage. In this respect, critical software (such as antivirus and vulnerability management solutions) are deployed on personnel's devices and strictly controlled via our MDM solution.

With respect to the hosting of our application: we rely on the infrastructure provided by Google Cloud Platform (GCP) and on its vulnerability and patch management programs.

Finally, the GPT for Work application is subject to our software development lifecycle (SDLC) process where security by design is key, from the product initiative stage until the product delivery stage. Through our continuous integration and continuous delivery (CI/CD) process, every new piece of application code is subject to state-of-the-art code security scanning (which includes detecting public vulnerabilities and ensuring compliance with frameworks such as OWASP Top10).

How do you ensure data center physical security?

As our application is hosted within data centers provided by Google, the physical security of our infrastructure is managed by Google Cloud Platform. Additional information can be found in Google Cloud Platform's security documentation available at: https://www.google.com/about/datacenters/data-security/.

Do you have a formal incident response plan?

Yes, as part of our ISO 27001 compliance framework, Talarian has security incident management procedures in place. The customer notification process in the event of a data incident involving personal identifiable information is set forth in our Data Processing Agreement.

Data & Privacy

Do the GPT for Work website and documentation use cookies?

The GPT for Work website and documentation site do not set first-party cookies.

On the documentation site, embedded YouTube videos load only after you choose to play them, using YouTube's privacy-enhanced mode. Once the player loads, YouTube may set third-party cookies. For more information, see How Google uses information from sites or apps that use our services.

The external support service may use necessary security cookies to protect the service and prevent abuse. These cookies are set by the support service, not by the GPT for Work website or documentation site.

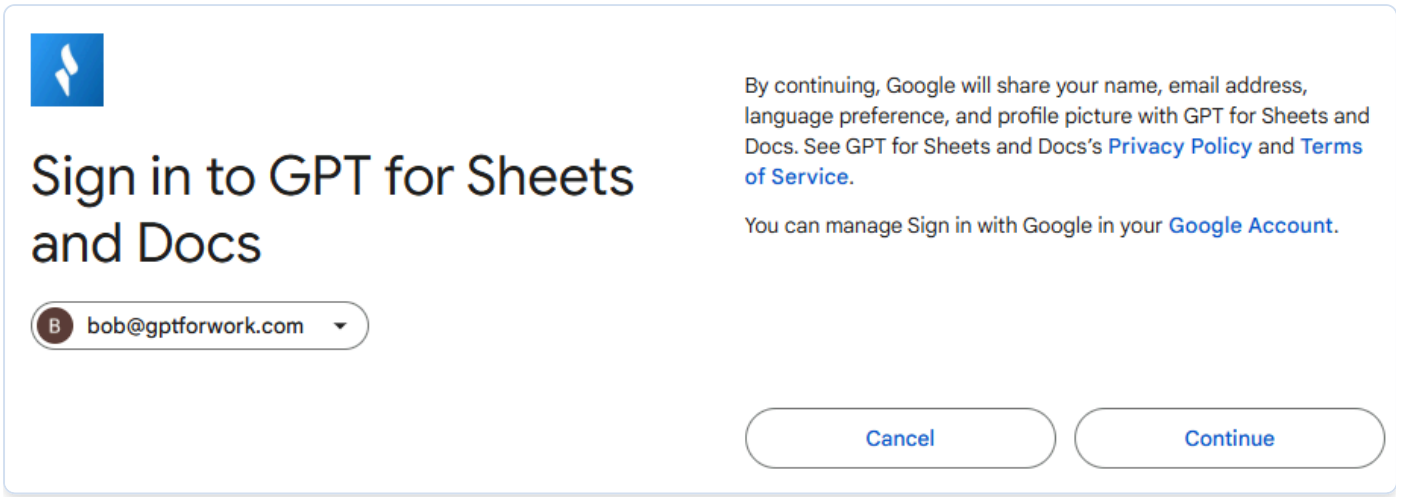
What permissions are needed to use GPT for Work?

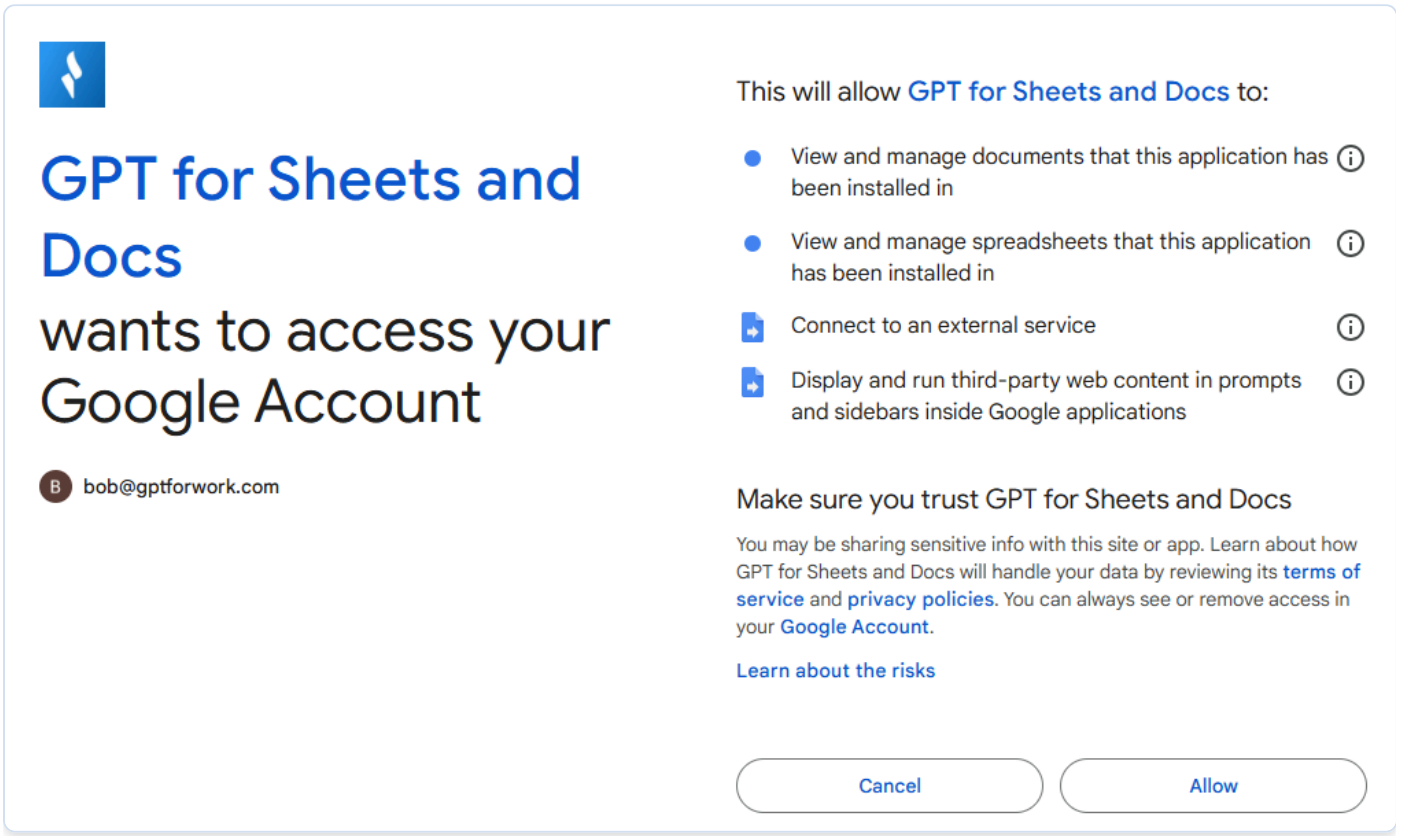
GPT for Work includes add-ons for Google Sheets and Google Docs (GPT for Sheets and Docs) as well as add-ins for Microsoft Excel and Microsoft Word (GPT for Excel Word). GPT for Work requires permissions to interact with these Google and Microsoft applications. Only the authorization scopes necessary for GPT for Work's performance are requested.

Google permissions for GPT for Sheets and Docs

Authorization scope

When you install GPT for Sheets and Docs from the marketplace, you are asked to allow the following:



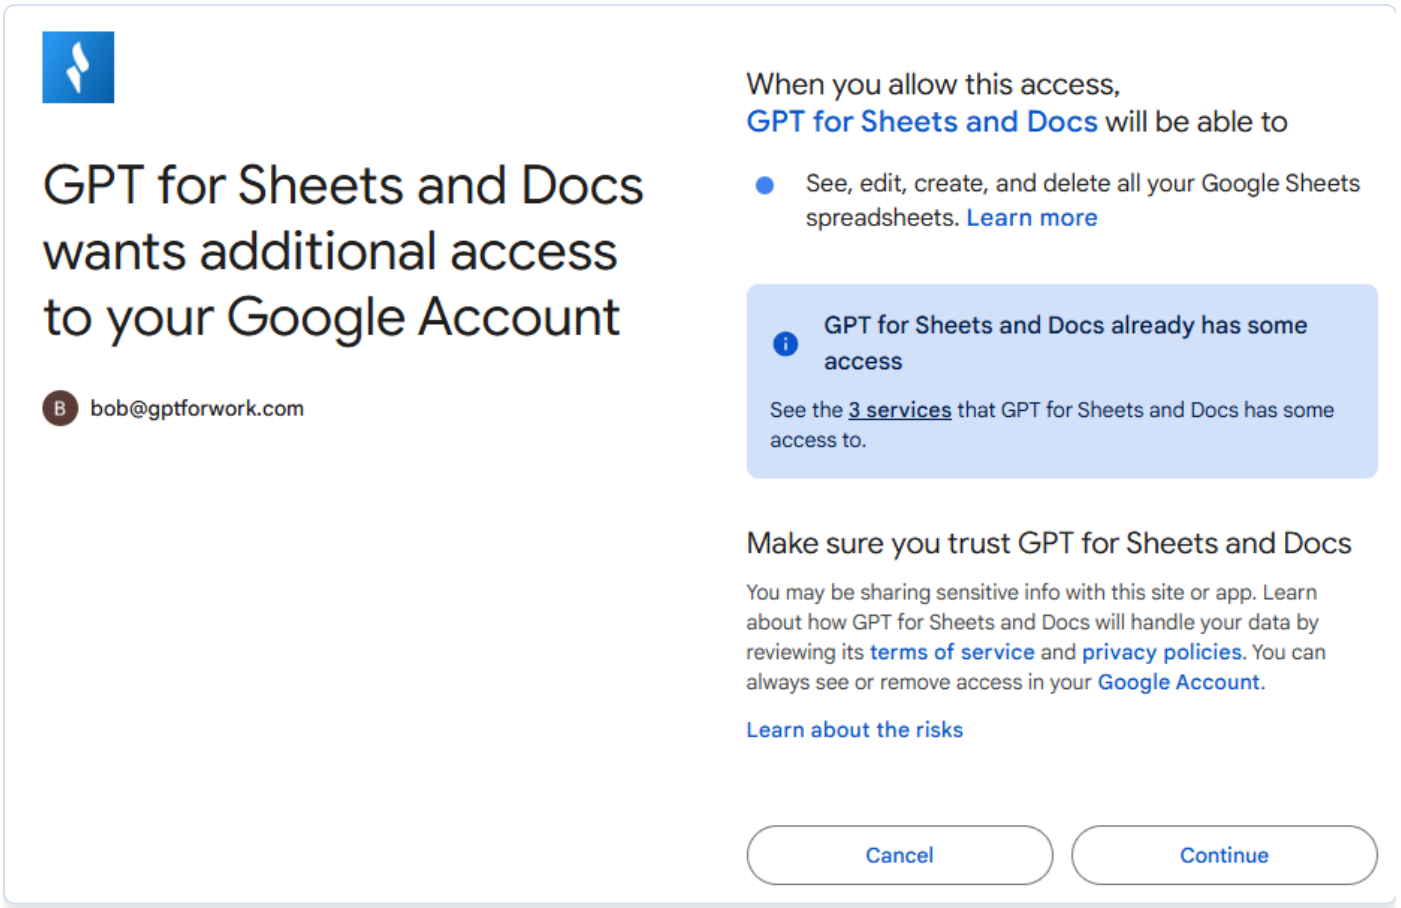


INFO

By clicking 'Allow' or 'Continue', you indicate your acceptance of the Terms of Service and Privacy Policy of GPT for Work.

Authorization scopes	Why does GPT for Work need it
See your primary Google account email address	Required to identify the user and fetch their associated metadata.
See your personal info, including any personal info you've made publicly available	
View and manage documents that this application has been installed in	Allows GPT for Work to fetch the user-selected Google Docs content in accordance with requests launched by the user on the application. Also allows GPT for Work to insert data in Google Docs in response to a user's request.
View and manage spreadsheets that this application has been installed in	Allows GPT for Work to fetch the user-selected Google Sheets data in accordance with requests launched by the user on the application. Also allows GPT for Work to insert data in Google Sheets in response to a user's request.
Connect to an external service	Required to allow GPT for Work add-on to send/retrieve information from its backend. Even if this authorization is described as an 'external service', it still stays within the Google environment because GPT for Work is hosted on Google Cloud Platform.
Display and run third-party web content in prompts and sidebars inside Google applications	Required to display GPT for Work sidebars inside Google Sheets and Google Docs.

When you enable the 'Bulk Tools' functionality, you are asked to allow the following:



Authorization scopes	Why does GPT for Work need it
See, edit, create, and delete all your Google Sheets spreadsheets	Allows GPT for Work to write results from AI providers into the user's spreadsheets in response to an asynchronous user request (through bulk tool runs). There is unfortunately no way to limit that scope to specific spreadsheets with the way that the Google Sheets API works. However, GPT for Work never reads from or writes to spreadsheets you didn't launch the add-on in.

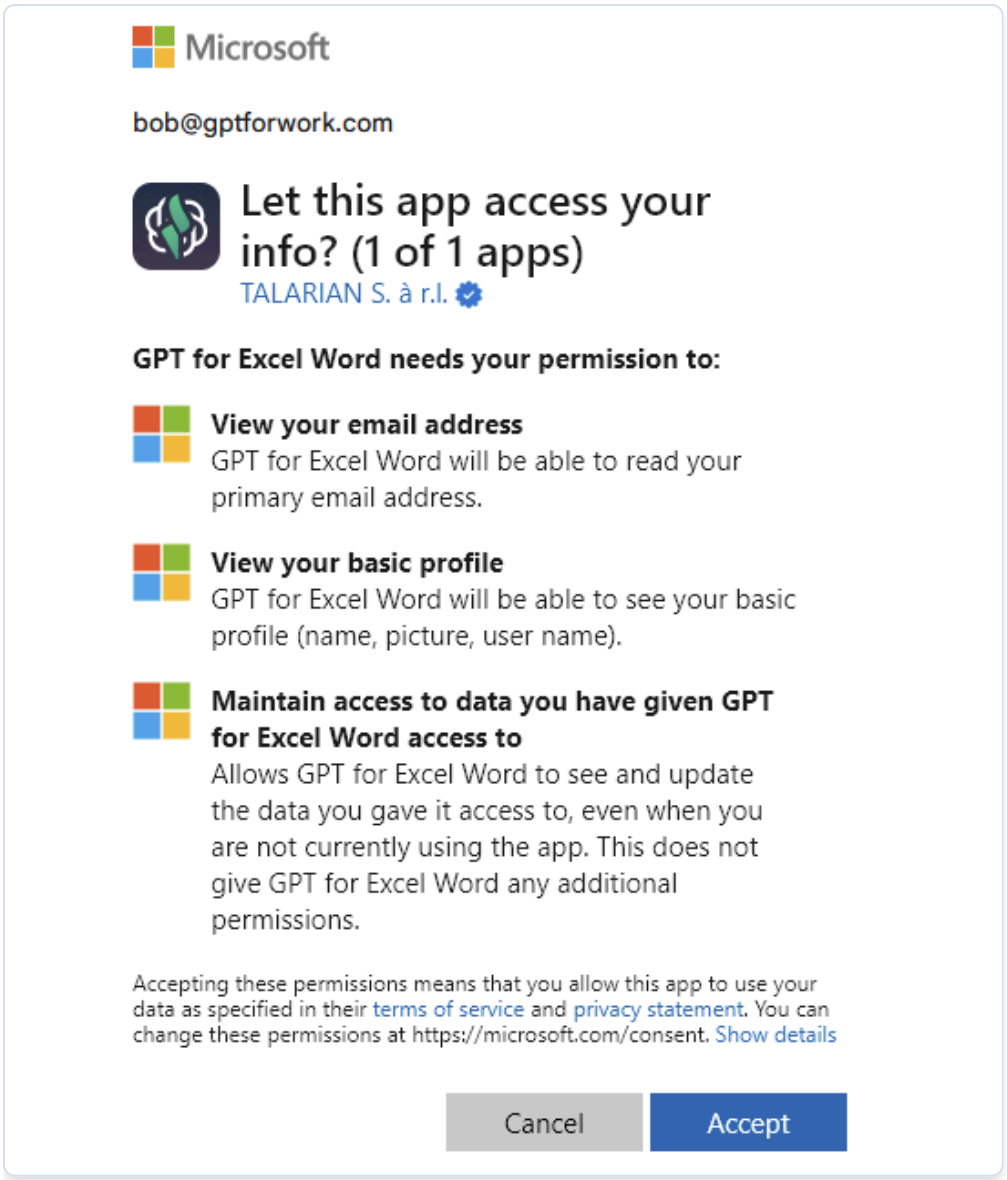
How to review or revoke permissions

At any point in time, you can review the permissions you have granted for GPT for Sheets and Docs. See more details here.

Microsoft Permissions for GPT for Excel Word

Authorization scope

When you install GPT for Excel Word from the marketplace, you are asked to allow the following:



INFO

By clicking **Accept**, you indicate your acceptance of the Terms of Service and Privacy Policy of GPT for Work.

Authorization scopes	Why does GPT for Work need it
User's email address	Required to identify the user and fetch their associated metadata.
User basic profile information (including name, surname, preferred username)	
OpenID (user's unique identifier and token)	

Authorization scopes	Why does GPT for Work need it
offline_access	Required to keep the user signed-in. The offline_access scope allows GPT for Work to receive refresh tokens from the Microsoft identity platform token endpoint.

How to review or revoke permissions

At any point in time, you can review the permissions you have granted for GPT for Excel Word. See more details here.

How is my data processed?

Here are some details as to the customer data used by GPT for Work, and which of this data is being stored on GPT for Work's infrastructure (i.e., Google Cloud Platform).

User identification

This includes the following stored data: username, account email address, basic profile information, GoogleID or MicrosoftID (as applicable). No other user identification data is stored.

Documents, Inputs, Outputs and Configurations

First, let's introduce three key concepts used in this paragraph:

- **Documents** mean the spreadsheets (Google Sheets or Microsoft Excel) and documents (Google Docs or Microsoft Word) in which the GPT for Work add-on is used;
- **Inputs** are the information you submit through GPT for Work. This includes:
  - the prompts and custom instructions you give (e.g. generating text, summarizing data or analyzing content). When using the 'Bulk tools' functionality in Google Sheets or Excel, you can configure these prompts and instructions, and these configurations are saved in your history for easy reuse (**Configurations**); and
  - the portion of the Documents submitted along with these prompts and instructions (e.g. relevant cells of your spreadsheets and parts of your documents) (**Submitted Document Portions**). To avoid any doubt, Configurations do not contain any Submitted Document Portions; and
- **Outputs** are the results generated by AI providers based on your Inputs.

Here is how GPT for Work uses your Documents, Inputs, Outputs and Configurations:

Documents

- Except when you are using the 'Agent' functionality, GPT for Work does not read or store any data from your Documents that you do not use as an Input in your prompt.
- When using the 'Agent' functionality, GPT for Work automatically selects the relevant parts of the Document that are necessary to respond effectively to your request. As of July 21, 2025, the 'Agent' feature is available in Beta mode to a select number of users only in Google Sheets and operates exclusively via Talarian's own API key(s) (OpenAI's model).
- GPT for Work sends your Inputs to the relevant AI provider's API or custom endpoint and writes the resulting Outputs in your Document.

Inputs & Outputs

- If you use a custom endpoint (as opposed to a model with API keys), such as on-premise inference servers (e.g., Ollama) or an Azure endpoint, the Inputs submitted and Outputs received do not go through our infrastructure, as the GPT for Work add-ons directly call the set endpoint. The use of custom endpoints therefore constitutes the most private option currently available.
- If you use a model with an API key, GPT for Work does not log any of the Inputs submitted or Outputs received. By using an API key, your AI provider usage (e.g. data retention, rate limits, cost, moderation) is bound by the terms you negotiated with your AI provider.
- If you use a model without an API key, we log all the Inputs and Outputs for support purposes for a duration of 30 days. We also store them for product improvement purposes (excluding the training of any language model) for no longer than 1 year. Outputs and Inputs may be stored for up to 5 years and can only be accessed on an *ad hoc* and as-needed basis by the authorized personnel.
- In Google Sheets only, and only if cache is enabled, Inputs and Outputs of executions made through GPT functions are cached encrypted for a short period of time (30 days) to prevent unexpected costs and loss of data due to Google Sheets automated recalculations. This caching occurs regardless of whether you use a model with or without an API key.

Configurations

- When you use the 'Bulk tools' functionality in Google Sheets or Excel, your bulk tools Configurations are securely stored in encrypted form for a limited period (90 days) to support the history feature.
- Examples of data that may be contained in Configurations include your desired translation glossary (if you are using the 'translate' tool) or your tailored classification instructions (if you are using the 'classify/categorize' tool). To avoid any doubt, Configurations do not contain any Submitted Document Portions.

Settings

These include the name of the model and the AI provider, the temperature (i.e., the setting that controls the creativity of a model's response), and any custom instruction you decide to insert to

provide context and shape how the chosen model responds.

Metadata relating to your Documents

GPT for Work stores the following additional data from your Document:

- Google Sheets: spreadsheets ID (identifier of the spreadsheet), sheet IDs (identifier of the sheet within a specific spreadsheet), sheet name; and
- Google Docs: Google Doc ID.

Usage metadata

These are data events corresponding to the different actions of GPT for Work users (e.g., user clicking on the 'Run' button).

API keys and custom endpoints

If you set up API keys for AI providers or use custom endpoints, they will be stored encrypted in transit, at rest and in use on Google Cloud Platform.

When setting up a team, space owners and admins can choose whether team members are allowed to use their own API keys and, in GPT for Excel Word, their own custom endpoints. In other words, space owners and admins have complete control over which AI providers and models are available to their space members.

Subscription data

This includes the following stored data: current plan details, transaction details (amount, invoice ID and payment date), details of the invoiced person or company (customer email and customer ID). We do not store payment card credentials on Talarian's infrastructure and always rely on Stripe, a third-party PCI-DSS-compliant payment processor, for payment card processing.

Is my data used by GPT for Work to train models?

No. Whether you use a model with or without an API key, Talarian does not use your Inputs and Outputs to train language models. Also, according to their terms as of July 21, 2025, the AI provider APIs we rely on (when using GPT for Work through Talarian's own API key(s)) currently do not use your Inputs and Outputs to train their models.

To check the AI provider's models available with or without an API key, please see our list of supported models.

Who owns Inputs and Outputs?

As between you and Talarian, you retain all rights in your Inputs and Outputs. More details are available in GPT for Work terms of service.

What do you do with my API Key and custom endpoints, when I use them within GPT for Work add-ons?

If you set up API keys for AI providers or use custom endpoints, they will be stored encrypted in transit, at rest and in use on Google Cloud Platform.

Do you store payment card information?

We don't store payment card credentials. Payments are processed through Stripe (https://stripe.com/en-fr). They are a PCI Service Provider Level 1 organization. Using Stripe means we don't need to store your payment card credentials.

You can read more about security at Stripe here: https://stripe.com/docs/security/stripe.

Where is my data located?

The Google Cloud Platform datacenters on which your data is stored are located in the United States - us-central1 region (Iowa). More details on Google Cloud Platform's region are available here.

Do you sell any of my data?

No, we never sell data.

Do you use third-party service providers?

Yes. GPT for Work relies on or interacts with several service providers in order to operate. They fall within 2 different categories:

Talarian's third-party service providers

These are service providers with which we hold a contractual relationship. They help us with various services including cloud hosting, generative AI provision or support ticket management.

Whenever we share data with these service providers, we make sure that they use it in compliance with data protection legislation, and that the personal data processing they carry out is covered by a specific data processing contract. They are of two sorts:

Talarian's subprocessors

They process personal data for which we are a processor, following our instructions. This processing is carried out as part of the services you have subscribed to. You will find details of our subprocessors in our Data Processing Agreement.

Service provider	Purpose
Google LLC or any of its affiliates	Data hosting, Generative AI *
Zendesk, Inc.	Support
OpenAI OpCo, LLC or any of its affiliates	Generative AI *
Anthropic, PBC or any of its affiliates	Generative AI *
Perplexity, Inc.	Generative AI *
Functional Software, Inc. (d/b/a Sentry)	Error Monitoring

* Google LLC, OpenAI OpCo, LLC, Anthropic, PBC, and Perplexity, Inc. are engaged for Generative AI purposes only to the extent (i) you are using GPT for Work through Talarian's own API key(s) and (ii) choose any such third-party service provider as AI provider.

Talarian's processors

They process data for which we act as the controller, and therefore act as processors.

Service provider	Purpose
SalesForce, Inc.	CRM
Sendinblue, SAS (d/b/a Brevo)	Transactional and marketing email
Slack Technologies Limited	Internal communication tool
Google Ireland Limited (Google Workspace)	Collaboration tool
Microsoft Corporation (Microsoft Clarity)	Front end analytics
Retool, Inc.	Admin dashboards for customer support

Other third-party service providers

These are third-party service providers with which you hold directly a contractual relationship. These can be of various types, notably:

- *AI providers that you use by creating and setting API keys or setting up custom API endpoints.* The full list of these AI providers are set forth in sections **Provider-connected models** and **Endpoint-connected models** on the page AI providers and models supported by GPT for Work.
- *Payment processor.* Payments are processed through Stripe (https://stripe.com/en-fr), which is a PCI Service Provider Level 1 organization. By purchasing GPT for Work you accept to be bound by the Stripe terms and conditions (https://stripe.com/en-fr/legal/consumer).

Can I have my data deleted?

Yes, you can.

As long as your GPT for Work account is active, we retain your data in accordance with our Privacy policy and Data Processing Agreement. Upon deletion of your GPT for Work account, we may archive your data only for a specified period of time in accordance with our internal data retention policy, and in any case no longer than legally permitted.

To request all your data to be deleted following deletion of your GPT for Work account, submit your request via our contact form and write '*Data Deletion request*' as the subject line.

Deleting your data will interrupt GPT for Work's functionality. Signing up again with the same Google or Microsoft account will not restore the deleted data.

How can I send you a privacy complaint?

For privacy complaints, you can reach us at: legal@talarian.io.

Guides

Explore Agent use cases, familiarize yourself with basic AI concepts, and learn how to write better prompts and scrape the web from Google Sheets.

Agent use cases

Learn about the different spreadsheet use cases that you can tackle with the Agent in GPT for Sheets and GPT for Excel.

Concepts

Familiarize yourself with the basic concepts related to GPT for Work, AI, and large language models (LLM).

How to write a good prompt

Write a good prompt to get the best results from our add-ons.

Scrape the web

Retrieve information from specific web pages.

Agent use cases for Sheets and Excel

Learn about the different spreadsheet use cases that you can tackle with the Agent in GPT for Sheets and GPT for Excel. You can explore the use cases by business application or product capability.

Business applications

Use the Agent in GPT for Sheets or GPT for Excel to handle product management, SEO, paid ads, outbound sales, market analysis, and customer feedback.

Product capabilities

Use the Agent in GPT for Sheets or GPT for Excel to handle content generation, translation, web search, data preparation, data analysis, formulas, and categorization.

Business applications for the Agent in Sheets and Excel

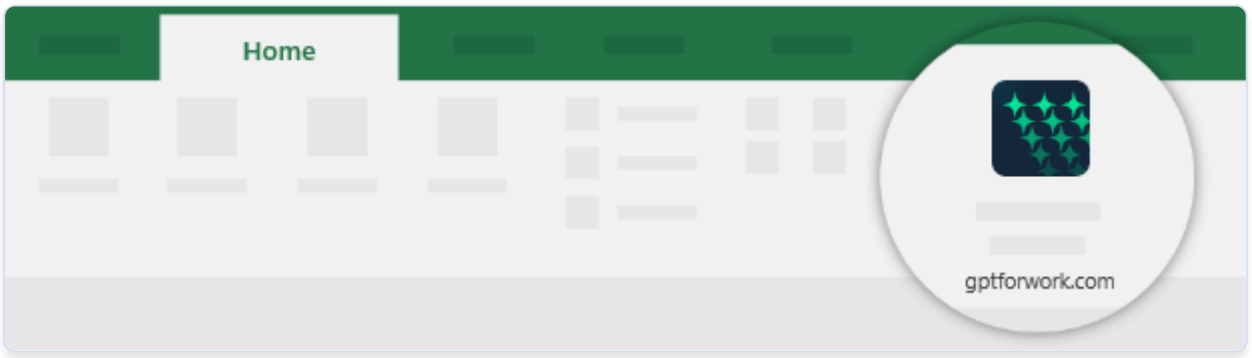
The Agent in GPT for Sheets and GPT for Excel allows you to tackle common business tasks in bulk with conversational AI. This page provides practical examples of how to use the Agent across different business applications.

How to use the Agent examples

To try out the use cases in GPT for Sheets or GPT for Excel, use the built-in examples from the add-on sidebar.

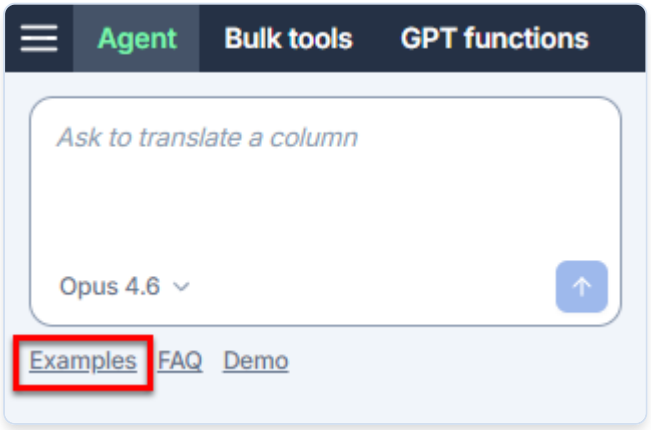
MS Excel

1. Create a new blank workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.

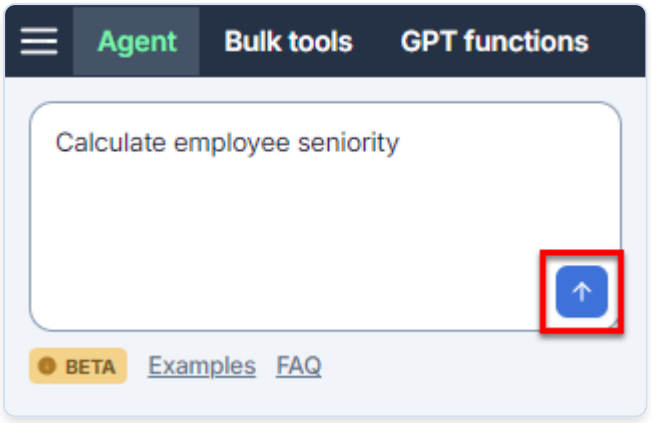


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, click **Examples** and select the example you want to try out.



4 Click the send button to start the example.



The following video shows the entire process.

Product management

E-commerce

Use the Agent to generate product descriptions, product listings, and other content for your e-commerce website.

In this example, the Agent is used to generate a title and description for the products in column A, then categorize them and extract the main features.

MS Excel

Watch video on YouTube

Data prep (inventory)

Use the Agent to prepare your data for analysis or reporting. This can include cleaning, transforming, or formatting your data to make it more useful for your purposes.

MS Excel

Watch video on YouTube

SEO

Translate to multiple languages

Use the Agent to translate spreadsheet content to multiple languages in bulk. You can provide specific instructions for the translation, and even use a glossary stored in another sheet to have the translation follow your brand's terminology.

In this example, the Agent is used to translate the content of column A to the languages provided in row 1.

MS Excel

Watch video on YouTube

SEO (h1, metas, FAQs)

Use the Agent to generate SEO-friendly content for your website. Bulk-generate meta titles and descriptions, headers, FAQs, image alt texts, and more.

In this example, the Agent is used to generate titles and descriptions for car spare parts, then generate FAQs about them.

MS Excel

Watch video on YouTube

Paid ads

Google Shopping (titles & categories)

Use the Agent to generate optimized content for your Google Shopping listings or for other product catalogs or paid ads.

In this example, the Agent is used to improve product titles and assign categories based on the official Google Shopping guidelines.

MS Excel

Watch video on YouTube

Outbound sales

ICP fitness

Use the Agent to filter your leads and find companies and contacts that match your ideal customer profile, based your own criteria.

In this example, the Agent is used to analyze the leads based on their industry fit, role fit, location, and company size.

MS Excel

Watch video on YouTube

Market analysis

Market research

Use the Agent to generate market overviews or company profiles in bulk by fetching information from the web.

In this example, the Agent is used to research the companies in column A, and find their website, CEO name, and other information.

MS Excel

Watch video on YouTube

Customer feedback

Feedback analysis

Use the Agent to categorize, filter, tag and analyze customer feedback and survey responses. This can include sentiment analysis, key topic extraction, brand and product mention identification, and more.

In this example, the Agent is used to translate support tickets to a common language, classify them based on sentiment, and define their relevance to product quality.

MS Excel

Watch video on YouTube

What's next?

- Use the Agent in GPT for Sheets or GPT for Excel.
- Learn about the Agent's product capabilities.

Product capabilities of the Agent in Sheets and Excel

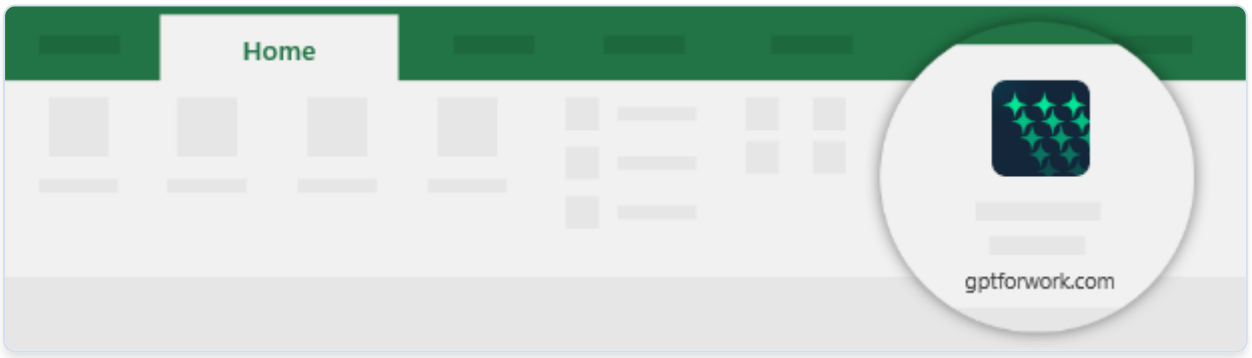
The Agent in GPT for Sheets and GPT for Excel allows you to tackle common tasks in bulk with conversational AI. This page provides practical examples of how to make the most out of the Agent's various capabilities.

How to use the Agent examples

To try out the use cases in GPT for Sheets or GPT for Excel, use the built-in examples from the add-on sidebar.

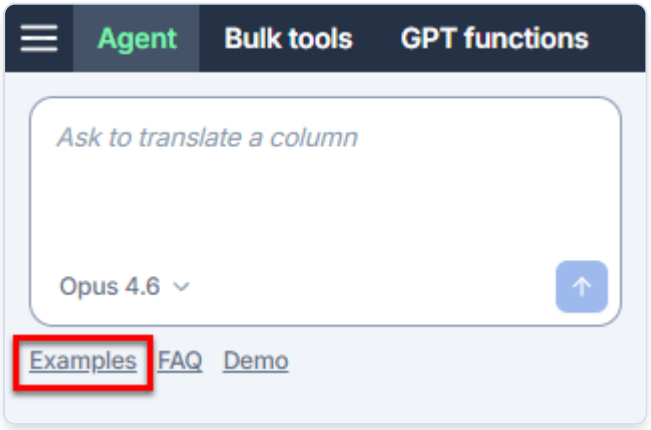
MS Excel

1. Create a new blank workbook.
2. On the ribbon, select **Home > GPT for Excel Word**.

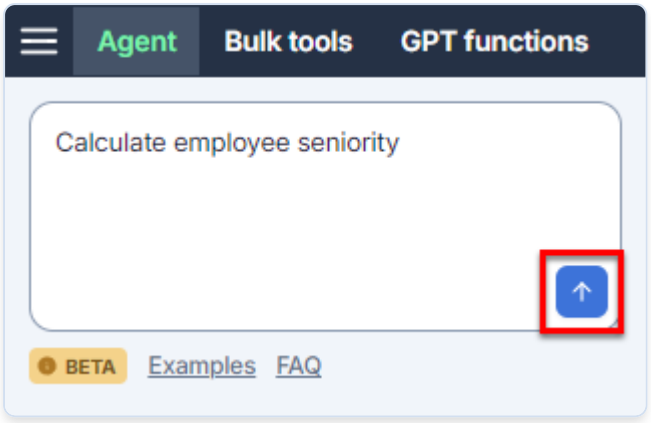


> **INFO**
>
> If **GPT for Excel Word** is not visible on the ribbon, or if it's grayed out, select **Home > Add-ins > My Add-ins > GPT for Excel Word**. Learn more.

3. In the sidebar, click **Examples** and select the example you want to try out.



4 Click the send button to start the example.



The following video shows the entire process.

Content generation

SEO (h1, metas, FAQs)

Use the Agent to generate SEO-friendly content for your website. Bulk-generate meta titles and descriptions, headers, FAQs, image alt texts, and more.

In this example, the Agent is used to generate titles and descriptions for car spare parts, then generate FAQs about them.

MS Excel

Watch video on YouTube

Translation

Translate to multiple languages

Use the Agent to translate spreadsheet content to multiple languages in bulk. You can provide specific instructions for the translation, and even use a glossary stored in another sheet to have the translation follow your brand's terminology.

In this example, the Agent is used to translate the content of column A to the languages provided in row 1.

MS Excel

Watch video on YouTube

Feedback analysis

Use the Agent to categorize, filter, tag and analyze customer feedback and survey responses. This can include sentiment analysis, key topic extraction, brand and product mention identification, and more.

In this example, the Agent is used to translate support tickets to a common language, classify them based on sentiment, and define their relevance to product quality.

MS Excel

Watch video on YouTube

Web search

Market research

Use the Agent to generate market overviews or company profiles in bulk by fetching information from the web.

In this example, the Agent is used to research the companies in column A, and find their website, CEO name, and other information.

MS Excel

Watch video on YouTube

Data preparation

Data prep (inventory)

Use the Agent to prepare your data for analysis or reporting. This can include cleaning, transforming, or formatting your data to make it more useful for your purposes.

MS Excel

Watch video on YouTube

Data analysis

ICP fitness

Use the Agent to filter your leads and find companies and contacts that match your ideal customer profile, based your own criteria.

In this example, the Agent is used to analyze the leads based on their industry fit, role fit, location, and company size.

MS Excel

Watch video on YouTube

Feedback analysis

Use the Agent to categorize, filter, tag and analyze customer feedback and survey responses. This can include sentiment analysis, key topic extraction, brand and product mention identification, and more.

In this example, the Agent is used to translate support tickets to a common language, classify them based on sentiment, and define their relevance to product quality.

MS Excel

Watch video on YouTube

Formulas

Generate and autofill formulas

Use the Agent to generate, troubleshoot, and explain spreadsheet formulas without writing any formula code yourself. Simply describe in your own words what calculation or logic you need, and the Agent creates the appropriate formulas. This is particularly useful for complex nested formulas or functions you're unfamiliar with.

In this example, the Agent is used to generate a formula to calculate the seniority of employees in years and months.

MS Excel

Watch video on YouTube

Categorization

E-commerce

Use the Agent to generate product descriptions, product listings, and other content for your e-commerce website.

In this example, the Agent is used to generate a title and description for the products in column A, then categorize them and extract the main features.

MS Excel

Watch video on YouTube

Google Shopping (titles & categories)

Use the Agent to generate optimized content for your Google Shopping listings or for other product catalogs or paid ads.

In this example, the Agent is used to improve product titles and assign categories based on the official Google Shopping guidelines.

MS Excel

Watch video on YouTube

What's next?

- Use the Agent in GPT for Sheets or GPT for Excel.
- Learn about the Agent's business applications.

Concepts

Learn about the key concepts related to GPT for Work.

AI concepts

AI provider

Company, platform, or server that provides models to GPT for Work add-ons. GPT for Work integrates with multiple AI providers, including OpenAI, Anthropic, and Google. See the complete list of supported models.

API endpoint

URL address that gives access to one or more models served by an AI provider. An API endpoint acts as the "gateway" through which GPT for Work sends your prompts, model settings, and other parameters, and from which it receives the model's responses. API endpoints give you access to more models and afford you more control and privacy, and let you benefit from the terms of service you have with an AI service, such as cost and rate limits.

You set up API endpoints for your space in the GPT for Work dashboard. All users in the space automatically have access to any and all models available through the endpoints, but only the space owner and admins can set and manage the endpoints. Learn more.

API key

Unique identifier that grants access to models served by AI providers. API keys give you access to more models and afford you more control and privacy, and let you benefit from the terms of service you have with an AI service, such as cost and rate limits.

You set API keys for your space in the GPT for Work dashboard. All users in the space automatically have access to any and all models associated with the keys, but only the space owner and admins can set and manage the keys. Billing for a key reflects its usage by all users. Learn more.

Formula (spreadsheet)

A spreadsheet formula is an expression you enter into a spreadsheet cell to calculate a value, process text, retrieve data, or perform some other task that returns a result. Formulas can contain functions, cell references, operators, as well as numbers, text, and other values. A

formula always starts with the equal sign `=`, which tells the spreadsheet to evaluate the expression instead of treating the entry as plain text.

The following simple formula consists of single function, `SUM`, which adds values:

```
=SUM(A1:A5)
```

The following more complex formula, which consists of two functions, one nested inside the other, runs the inner `GPT` function only if the cell `A1` is not blank, which is determined by the outer `IF` function. If `A1` is not blank, the formula returns the result from the `GPT` function; if `A1` is blank, the formula returns an empty value.

```
=IF(A1="", "", GPT("Write a tagline for", A1))
```

Function (spreadsheet)

A spreadsheet function is a predefined formula that performs a specific task in a spreadsheet, such as calculating values, processing text, or retrieving data. Functions always return a result. Functions typically take one or more parameters as input, which tell them how exactly to perform their task and generate their result. For example, the `SUM` function adds values:

```
=SUM(A1:A5)
```

Functions can be divided into native functions, which are built into the spreadsheet software, such as `SUM` in Excel and Sheets, and custom functions, which come from user scripts or third-party software, such the `GPT` function in GPT for Excel and GPT for Sheets.

You can use a function on its own (a simple formula) or combine it with other functions to create more complex formulas, optionally with nesting and conditional logic.

Global instructions

Specific guidelines provided to the model to define its role, expertise, or how it should respond to prompts. Global instructions are set in the sidebar of the GPT for Work add-ons, and remain active in the document or spreadsheet until you change them.

Model

Large language model (LLM) used in a GPT for Work add-on.

The available models depend on whether you are using only built-in models or also connected models. See the complete list of supported models.

Cloud-hosted model

Model served over the internet by a cloud-based AI platform. This is in contrast to models served by an LLM server running on a local machine. Compared to local models, cloud-hosted models generally run on more powerful hardware, but require online access and involve sending data over the internet to an external service. For using cloud-hosted models, GPT for Work supports multiple AI providers, both model creators, such as OpenAI, Anthropic, and Google, and model hosts, such as Anyscale, Fireworks AI, Together AI, or any platform with an OpenAI-compatible API endpoint.

Content fetching model

Some web search models can access the full content from specific URLs. The full content is added to the context of the models, so they can use it to answer questions about the content or extract information from it.

Local model

Model served by an LLM server running on a local machine, such as your own computer or another computer on a local network. This is in contrast to models served over the internet by a cloud-based AI platform. Compared to cloud-hosted models, local models offer enhanced data privacy and offline access. For using local models, GPT for Work supports any LLM server with an OpenAI-compatible API.

Reasoning model

Reasoning models are trained to think before they answer, producing an internal chain of thought before responding to a prompt. They generate two types of output tokens:

- **Reasoning tokens** make up the model's internal chain of thought. These tokens are typically not visible in the output from most providers.
- **Completion tokens** make up the model's final visible response.

Vision model

Vision models can process images as input. The following features support vision models:

- **Custom prompt** bulk tool in GPT for Sheets and GPT for Excel
- **Prompt images (Vision)** bulk tool in GPT for Sheets and GPT for Excel
- **GPT_VISION** function in GPT for Sheets and GPT for Excel

For information about supported and recommended vision models, see:

- Use images in bulk tool prompts in Sheets / Use images in bulk tool prompts in Excel
- Use images in GPT function prompts in Sheets / Use images in GPT function prompts in Excel

Web search model

All web search models can go beyond their built-in knowledge by accessing information from search engines, like Google or Bing. This allows models, which have compressed historical knowledge, to explore topics in greater depth, with more current data or with a specific focus.

Prompt

Input text given to the model to generate a response or perform a task. Prompts can be included in the sidebar of the GPT for Work add-ons, or passed as arguments in GPT functions. For guidance, see How to write a good prompt.

Rate limit

Restriction on the number of tokens or on the number of requests that can be processed within a certain time frame (minute, day, month). AI providers set rate limits on their models. Google also enforces rate limits on the daily number of requests. Learn more.

Token

A token is a unit of text, such as a word, subword, punctuation mark, or symbol. Models split all input and output text into tokens for processing. Model providers use tokens as the basic unit for measuring model usage. Input tokens from your prompt and output tokens from the model's response contribute to your credit consumption. For more information, see Tokens and How credits are consumed.

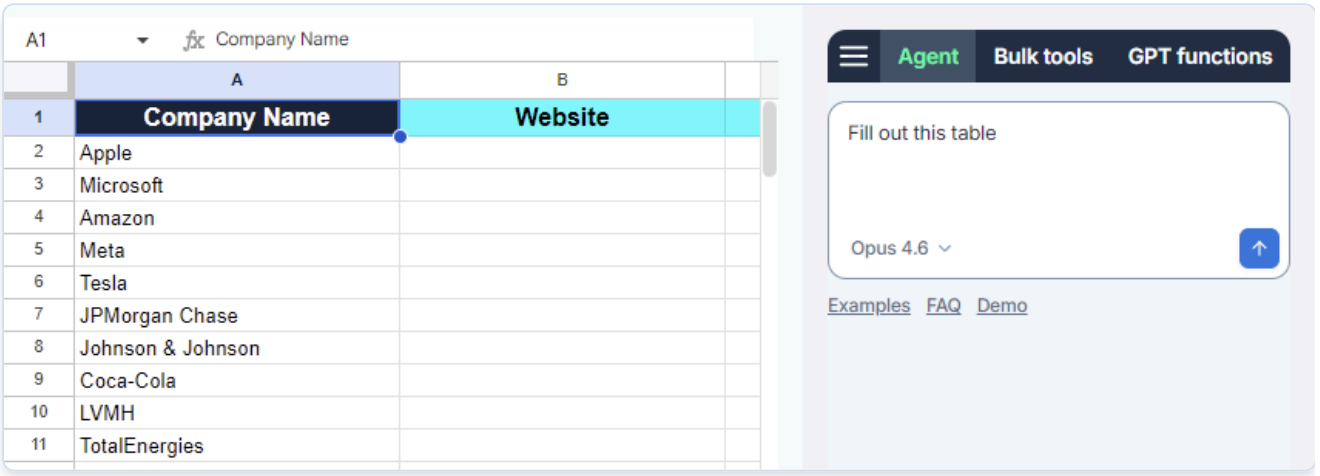
Web search context

Web search models do not fetch full page content. Instead, they read search result snippets, like you would see on a Google search results page. The snippets are added to the **context** of the models, so they can use the information to generate a response. The larger the context, the more information a model can retrieve from each web source, producing richer and more detailed responses. For some models, you can select the **context size** in the model settings.

GPT for Work concepts

Agent

The **Agent** is the chat interface for spreadsheet work in GPT for Excel and GPT for Sheets. Describe what you need in plain language — from everyday spreadsheet tasks to row-by-row bulk processing across large datasets — and the Agent reads your data, determines the required steps, and executes them directly in your spreadsheet.



For more information, see Agent in GPT for Sheets and Agent in GPT for Excel.

Agent model

Model used by the Agent to handle chat messages, thinking, and spreadsheet tasks that do not require bulk processing with AI. For a list of supported Agent models, see AI providers and models supported by GPT for Work.

Built-in model

Model that is available in a GPT for Work add-on out of the box. This is in contrast to a connected model, which you make available yourself by setting an API key (for a provider-connected model) or setting up an API endpoint (for an endpoint-connected model). Built-in models let you get started quickly and are available on both subscription and pay-as-you-go pricing.

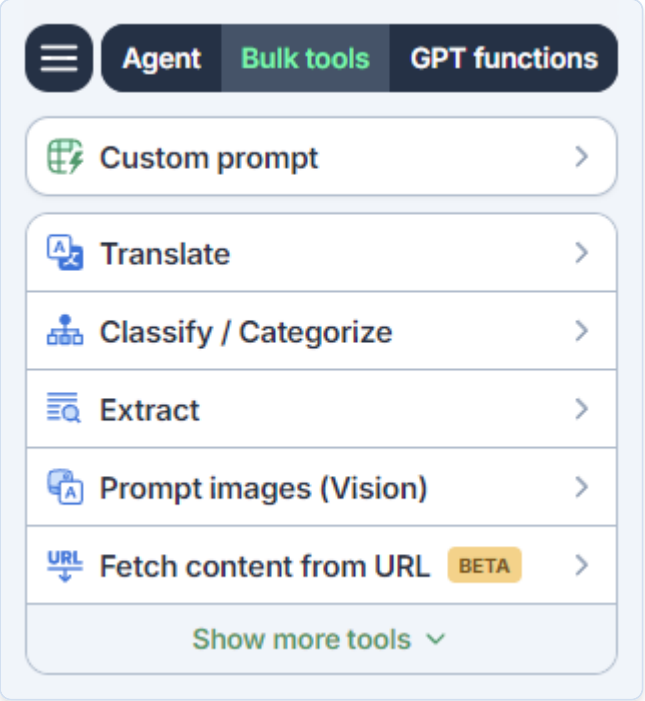
Bulk model

Models used in bulk tools (manually or Agent-run) and GPT functions to perform bulk processing with AI. For a list of supported bulk models, see AI providers and models supported by GPT for Work.

Bulk tools

Bulk tools allow you to run prompts on specific spreadsheet columns at once without writing any formulas. You configure and run bulk tools from the add-on sidebar. You can also easily reuse past prompts and configurations with the **History** feature.

GPT for Work provides dedicated bulk tools for selected use cases, such as translation, classification/categorization, and extraction, as well as a generic bulk tool for fully custom prompts.



For more information, see Bulk tools in GPT for Sheets and Bulk tools in GPT for Excel.

Connected model

Model that you make available in a GPT for Work add-on yourself. This is in contrast to a built-in model, which is available out of the box.

You have two types of connected models:

- **Endpoint-connected models** are available through an API endpoint that you set up for a cloud-based AI platform or local LLM server.
- **Provider-connected models** are available with an API key you set for an AI provider.

API keys and API endpoints give you access to more models and afford you more control and privacy, and let you benefit from the terms of service you have with an AI service, such as cost and rate limits.

Endpoint-connected model

Connected model that you make available in a GPT for Work add-on by setting up an API endpoint for a cloud-based AI platform or local LLM server.

API endpoints give you access to more models and afford you more control and privacy, and let you benefit from the terms of service you have with an AI service, such as cost and rate limits.

Provider-connected model

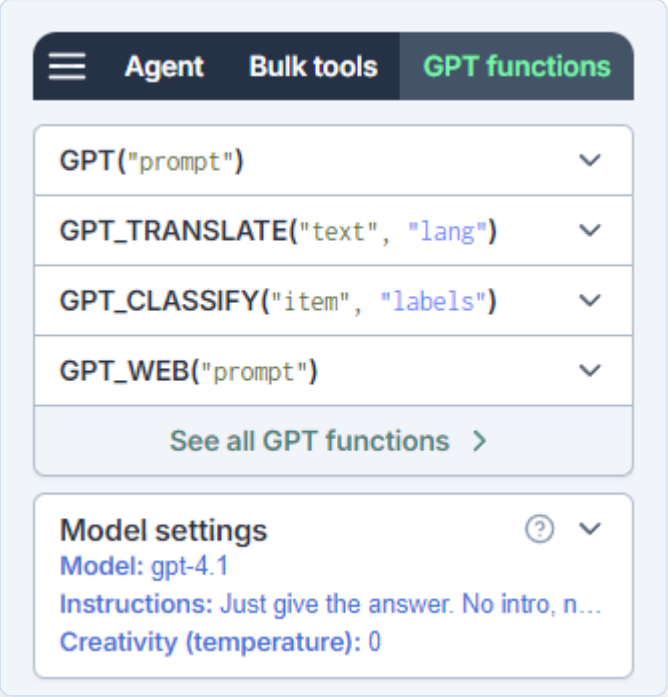
Connected model that you make available in a GPT for Work add-on by setting an API key for an AI provider.

API keys give you access to more models and afford you more control and privacy, and let you benefit from the terms of service you have with an AI service, such as cost and rate limits.

GPT functions

GPT functions are custom spreadsheet functions that allow you to prompt AI from inside spreadsheet cells. GPT functions work exactly like native functions in that you can use them on their own or combine them with other functions when creating formulas.

GPT for Work provides dedicated GPT functions for selected use cases, such as translation, classification, and web search, as well as a generic GPT function for all other use cases. You can define function-specific parameters to further refine how the AI generates responses. You can find a list of all available functions with usage examples in the add-on sidebar.

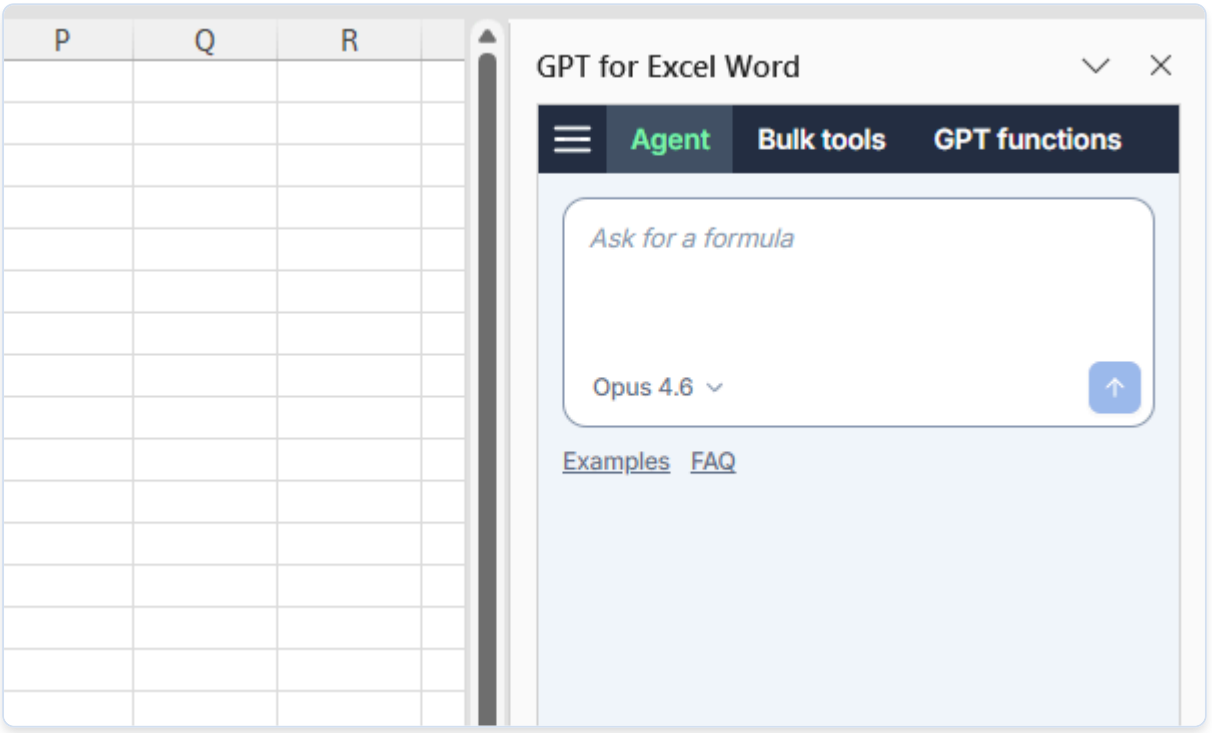


For more information, see All GPT functions in GPT for Sheets and All GPT functions in GPT for Excel.

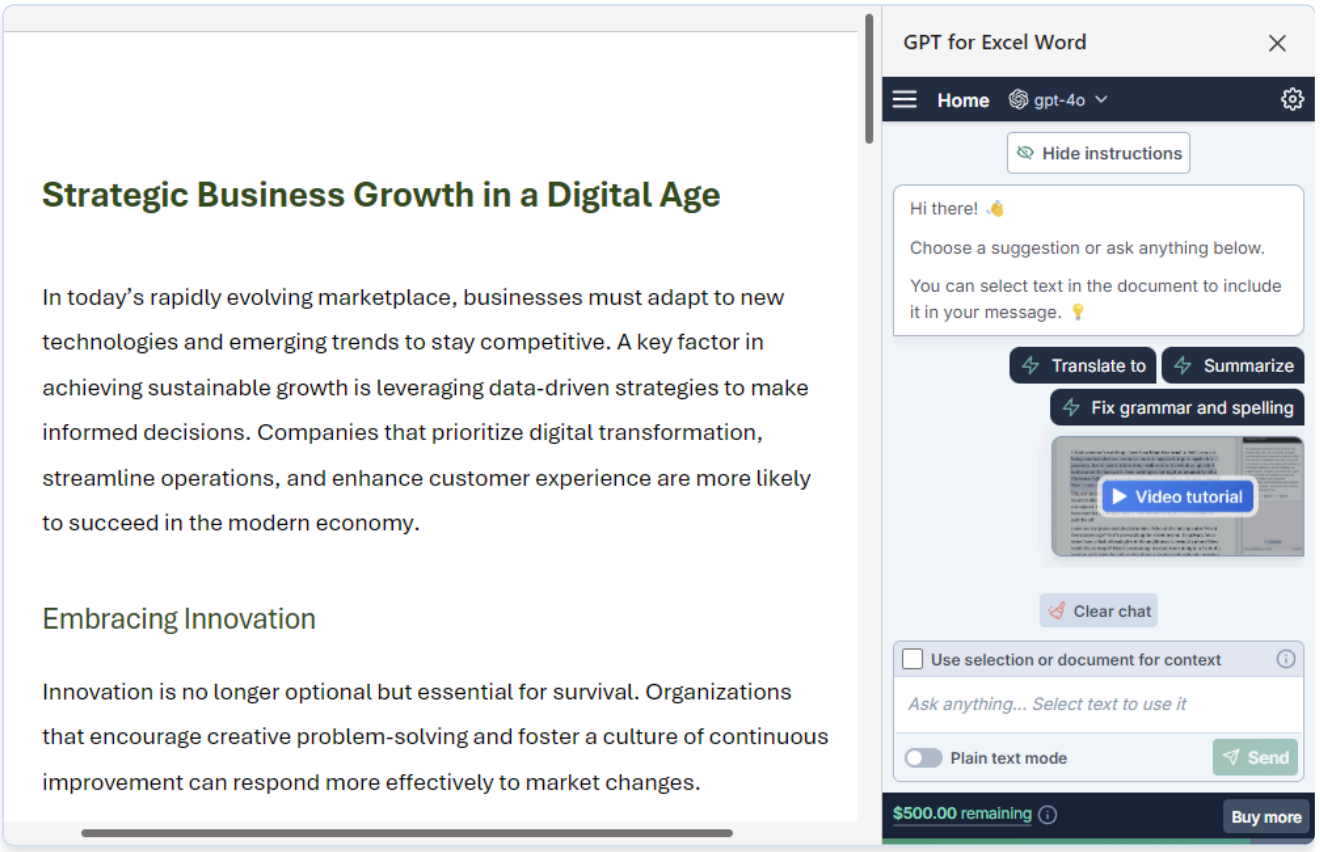
Sidebar

Panel in GPT for Work add-ons where users can interact with the AI, input prompts, and adjust settings. When you launch the GPT for Work add-ons, the sidebar opens on the right side of the screen.

MS Excel



MS Word



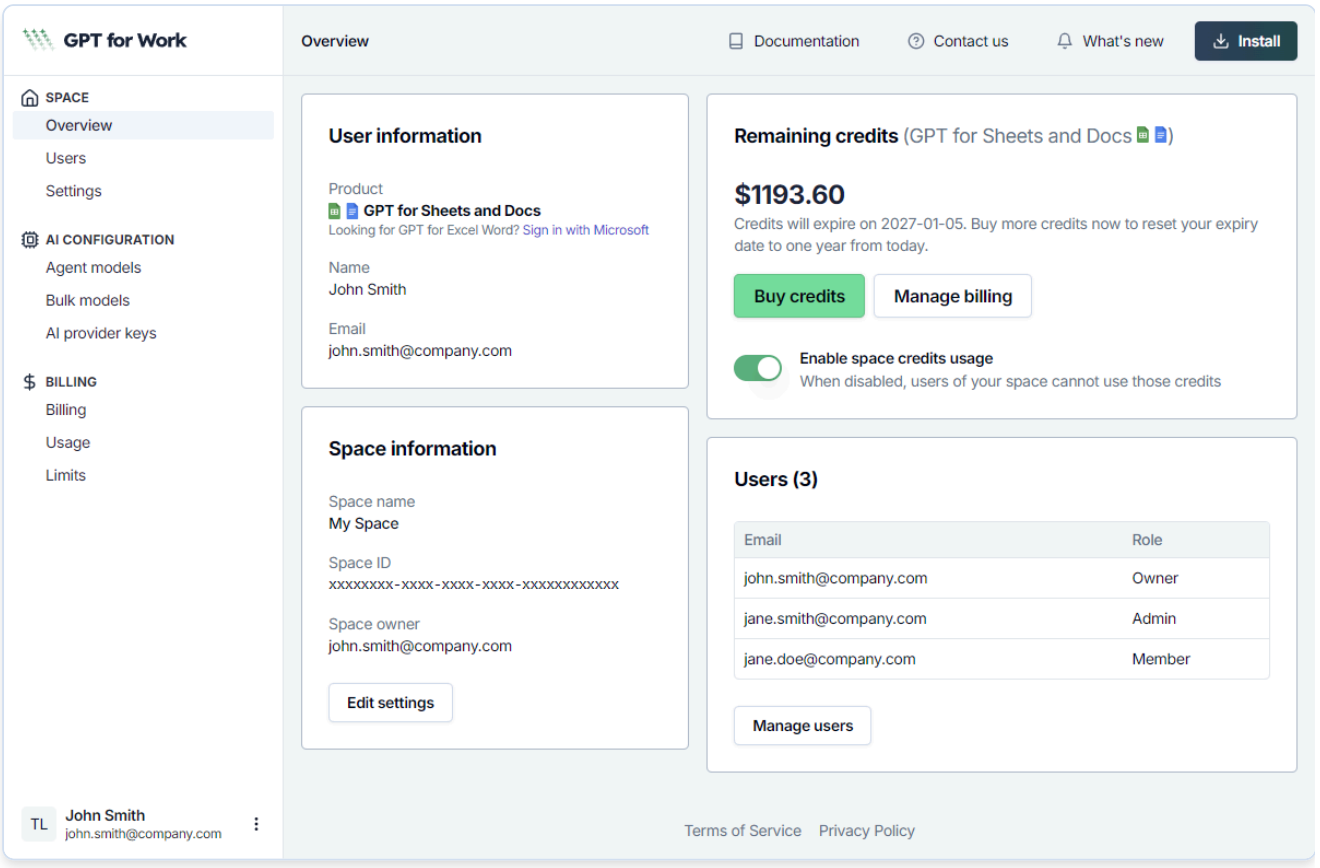
Space

A space is the administrative hub for GPT for Work. Available through the GPT for Work dashboard, it centralizes purchases, billing, usage controls, and API keys for an individual user or a team.

If you sign in to both GPT for Sheets and Docs and GPT for Excel and Word, you will have two separate spaces, even if you sign in with the same account on both platforms. These spaces are independent: credits, users, and API keys do not carry over between platforms.

You can be active in only one space at a time. If you join another user's space, you lose access to your own space's credits, API keys, and API endpoints.

In the screenshot below, you can see the overview of a GPT for Sheets and Docs space with three space users.


GPT for Work
Overview
Documentation
Contact us
What's new
Install
SPACE
Overview
Users
Settings
AI CONFIGURATION
Agent models
Bulk models
AI provider keys
BILLING
Billing
Usage
Limits
TL
John Smith
john.smith@company.com
User information
Product
GPT for Sheets and Docs
Looking for GPT for Excel Word? Sign in with Microsoft
Name
John Smith
Email
john.smith@company.com
Space information
Space name
My Space
Space ID
xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
Space owner
john.smith@company.com
Edit settings
Remaining credits (GPT for Sheets and Docs)
$1193.60
Credits will expire on 2027-01-05. Buy more credits now to reset your expiry date to one year from today.
Buy credits
Manage billing
Enable space credits usage
When disabled, users of your space cannot use those credits
Users (3)
Email
Role
john.smith@company.com
Owner
jane.smith@company.com
Admin
jane.doe@company.com
Member
Manage users
Terms of Service
Privacy Policy

How to write a good prompt

Write a good prompt to get the best results from our add-ons. Be very specific about what you want and prompt exactly like you would brief a person to do the job, by providing clear instructions and ample context.

While models can sometimes process multi-lingual prompts or generate a response in a different language from the prompt language, it is highly recommended to write your prompt in the language that you expect your response to be in.

Label	Description	Examples
Task	What you want to achieve **Note:** This should be composed of at least a verb and a noun. You can replace "Task" by "Question" if it makes more sense.	`Task: write a cold email` `Task: write a blogpost` `Task: write 5 taglines` `Task: summarize in 3 paragraphs` `Task: translate from English to Spanish` `Task: create a tweet thread of 10 tweets` `Question: How can I find an SEO expert?`
Topic	What your text is about **Note:** Don't hesitate to be highly specific.	`Topic: sell a shampoo` `Topic: an ice-cream shop` `Topic: food in Ancient Rome`

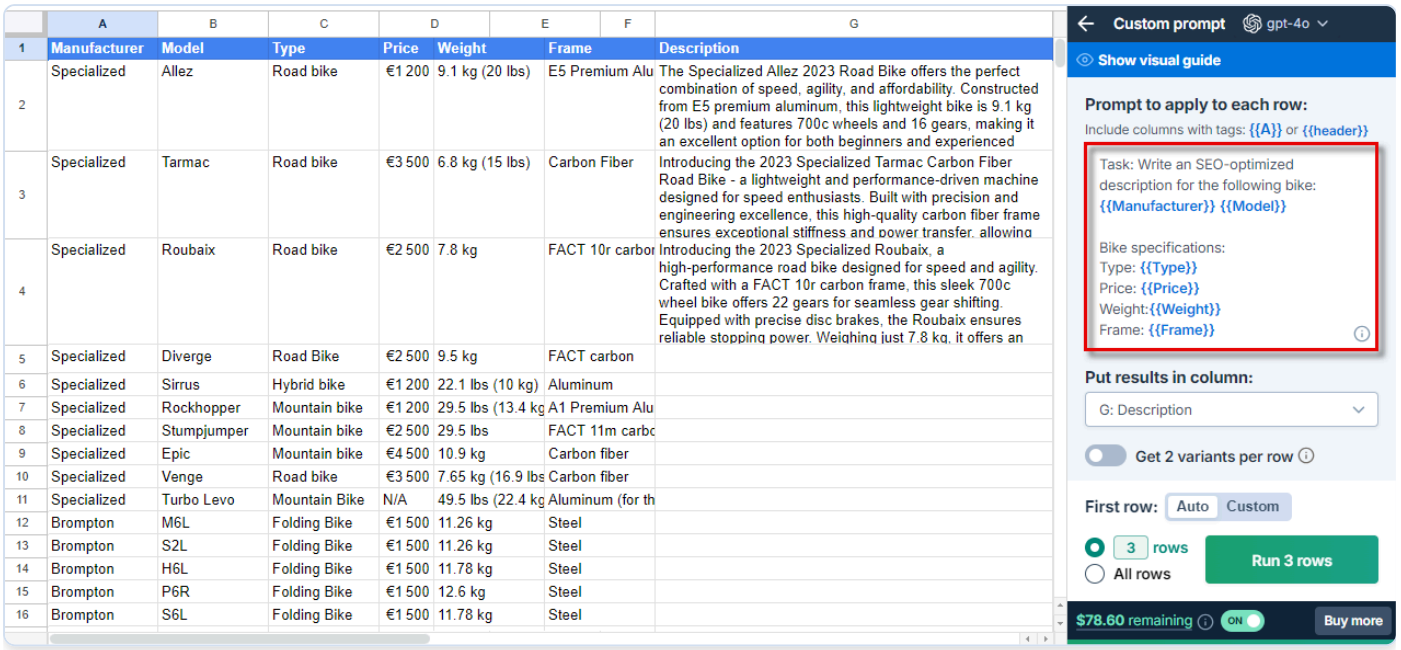
Label	Description	Examples
Role	What character the AI has to impersonate	`Role: Act as a teacher. Write in very basic terms, and illustrate with examples adapted to your audience. Then provide links to explore the topic further.` `Role: Act as a tech reviewer. You write in-depth reviews including pros, cons, features, and benchmarks.` `Role: Act as a technical writer. Your writing is very structured and hierarchical. Instructions are provided step by step so the reader can apply them without thinking.`
Style	What type of language you want to use	`Style: casual` `Style: formal` `Style: business` `Style: creative` `Style: academic`
Tone	How you want your text to sound	`Tone: joyful` `Tone: angry` `Tone: funny` `Tone: serious` `Tone: excited` `Tone: sad`

Label	Description	Examples
Audience	Who this text is for **Note:** This enables the model to select appropriate words and the right level of abstraction.	`Audience: 5-year old` `Audience: teenager who likes rock music` `Audience: my boss` `Audience: topic expert`
Length	How long the output is expected to be (any unit) **Note:** Language models have limited numerical accuracy, so it's good to be overly precise if you have a hard limit.	`Length: 4 paragraphs` `Length: 250 words` `Length: strictly under 30 characters including spaces`
Format	What structure or syntax your output must comply with	`Format: HTML` `Format: table` `Format: markdown`
	Any other parameter or instruction that is useful to you	`Country: X` `Product: door bell` `Relate recipient job title to the user selling proposition` `Sign as XYZ` `Start with xxx` `Include references to X, Y and Z`

It is then extremely easy to write a good prompt that you can then use with our add-ons.

Spreadsheets

This is what your prompt should look like in a spreadsheet with GPT for Sheets or GPT for Excel Custom prompt bulk tool:

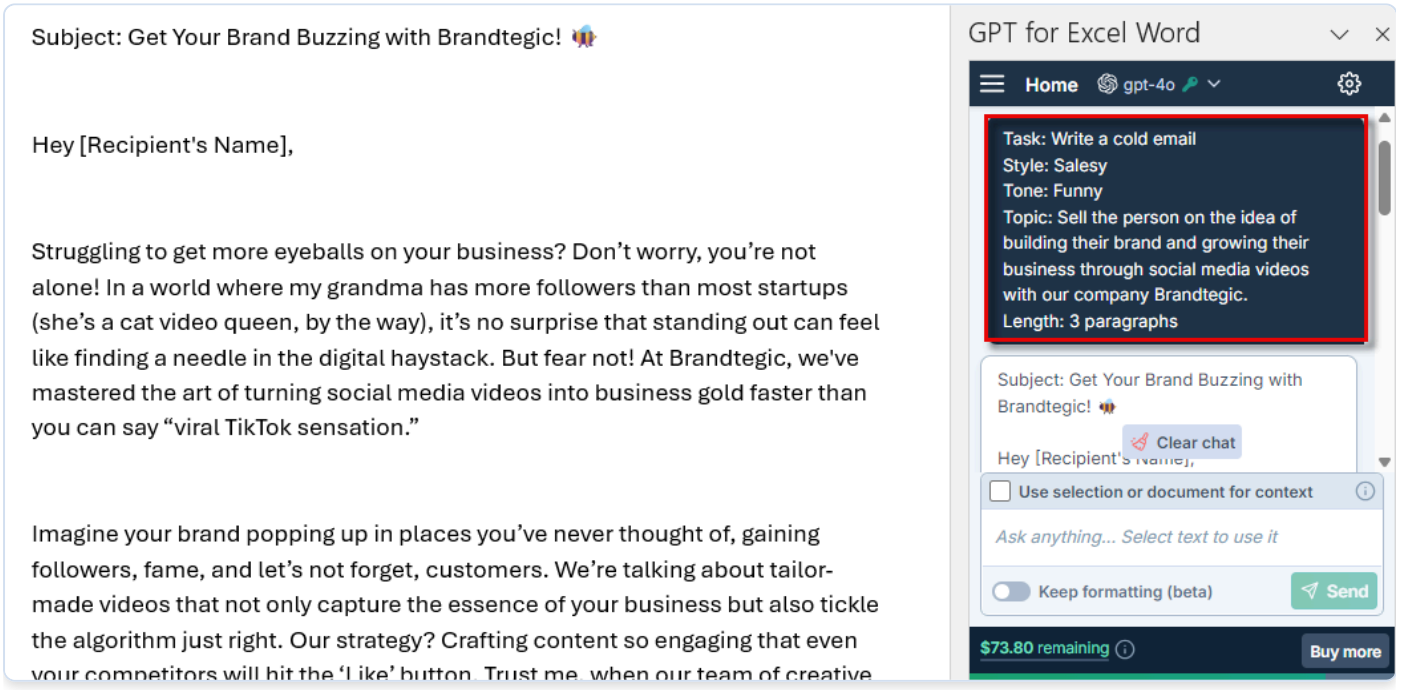


	A	B	C	D	E	F	G
1	Manufacturer	Model	Type	Price	Weight	Frame	Description
2	Specialized	Allez	Road bike	€1 200	9.1 kg (20 lbs)	E5 Premium Alu	The Specialized Allez 2023 Road Bike offers the perfect combination of speed, agility, and affordability. Constructed from E5 premium aluminum, this lightweight bike is 9.1 kg (20 lbs) and features 700c wheels and 16 gears, making it an excellent option for both beginners and experienced
3	Specialized	Tarmac	Road bike	€3 500	6.8 kg (15 lbs)	Carbon Fiber	Introducing the 2023 Specialized Tarmac Carbon Fiber Road Bike - a lightweight and performance-driven machine designed for speed enthusiasts. Built with precision and engineering excellence, this high-quality carbon fiber frame ensures exceptional stiffness and power transfer, allowing
4	Specialized	Roubaix	Road bike	€2 500	7.8 kg	FACT 10r carbon	Introducing the 2023 Specialized Roubaix, a high-performance road bike designed for speed and agility. Crafted with a FACT 10r carbon frame, this sleek 700c wheel bike offers 22 gears for seamless gear shifting. Equipped with precise disc brakes, the Roubaix ensures reliable stopping power. Weighing just 7.8 kg, it offers an
5	Specialized	Diverge	Road Bike	€2 500	9.5 kg	FACT carbon	
6	Specialized	Sirrus	Hybrid bike	€1 200	22.1 lbs (10 kg)	Aluminum	
7	Specialized	Rockhopper	Mountain bike	€1 200	29.5 lbs (13.4 kg	A1 Premium Alu	
8	Specialized	Stumpjumper	Mountain bike	€2 500	29.5 lbs	FACT 11m carbo	
9	Specialized	Epic	Mountain bike	€4 500	10.9 kg	Carbon fiber	
10	Specialized	Venge	Road bike	€3 500	7.65 kg (16.9 lbs	Carbon fiber	
11	Specialized	Turbo Levo	Mountain Bike	N/A	49.5 lbs (22.4 kg	Aluminum (for th	
12	Brompton	M6L	Folding Bike	€1 500	11.26 kg	Steel	
13	Brompton	S2L	Folding Bike	€1 500	11.26 kg	Steel	
14	Brompton	H6L	Folding Bike	€1 500	11.78 kg	Steel	
15	Brompton	P6R	Folding Bike	€1 500	12.6 kg	Steel	
16	Brompton	S6L	Folding Bike	€1 500	11.78 kg	Steel	

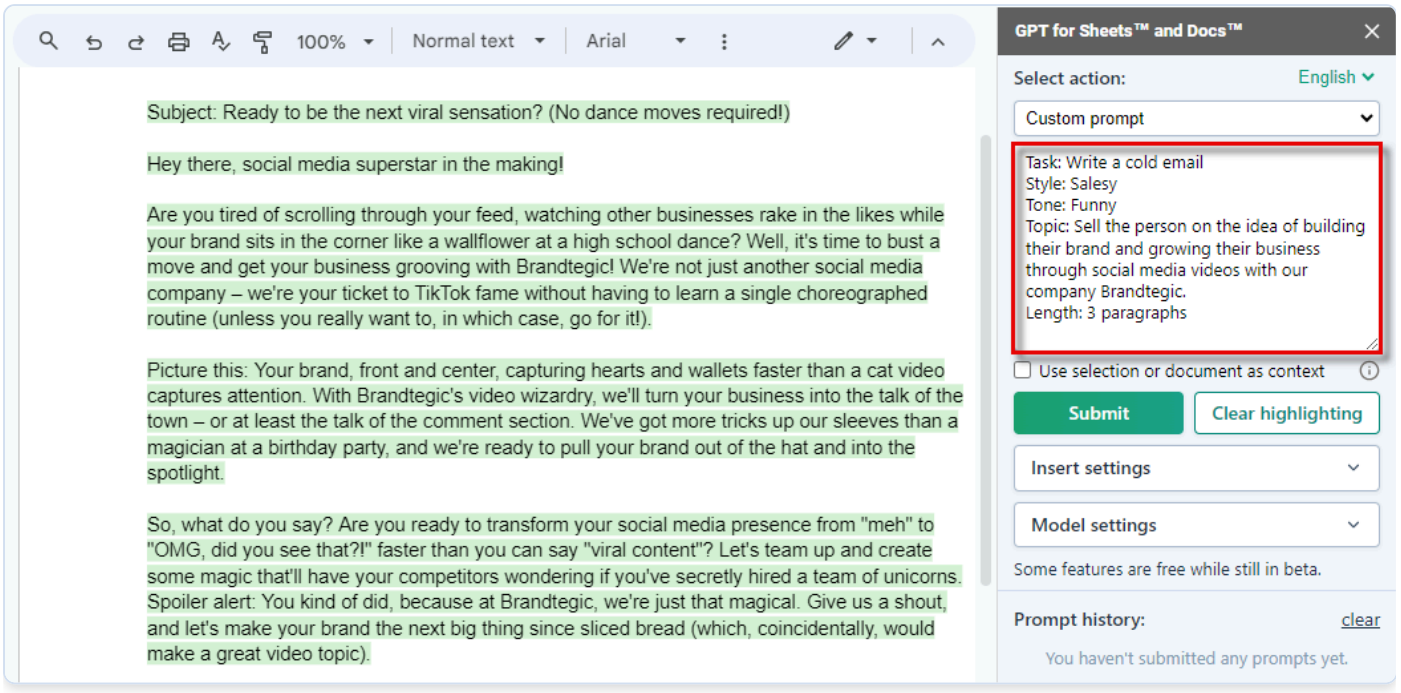
Learn more about GPT for Sheets Custom prompt bulk tool or GPT for Excel Custom prompt bulk tool.

Documents

This is what a prompt should look like in a document with GPT for Word:



This is what a prompt should look like in a document with GPT for Docs:



You can copy this prompt from the text below:

```
Task: Write a cold email
Style: Salesy
Tone: Funny
Topic: Sell the person on the idea of building their brand and growing their
business through social media videos with our company Brandtegic.
Length: 3 paragraphs
```

Scrape the web from Google Sheets

Retrieve information from specific web pages directly into your Google Sheets spreadsheet using our web scraping script. The script imports text into your cells while automatically removing HTML and formatting.

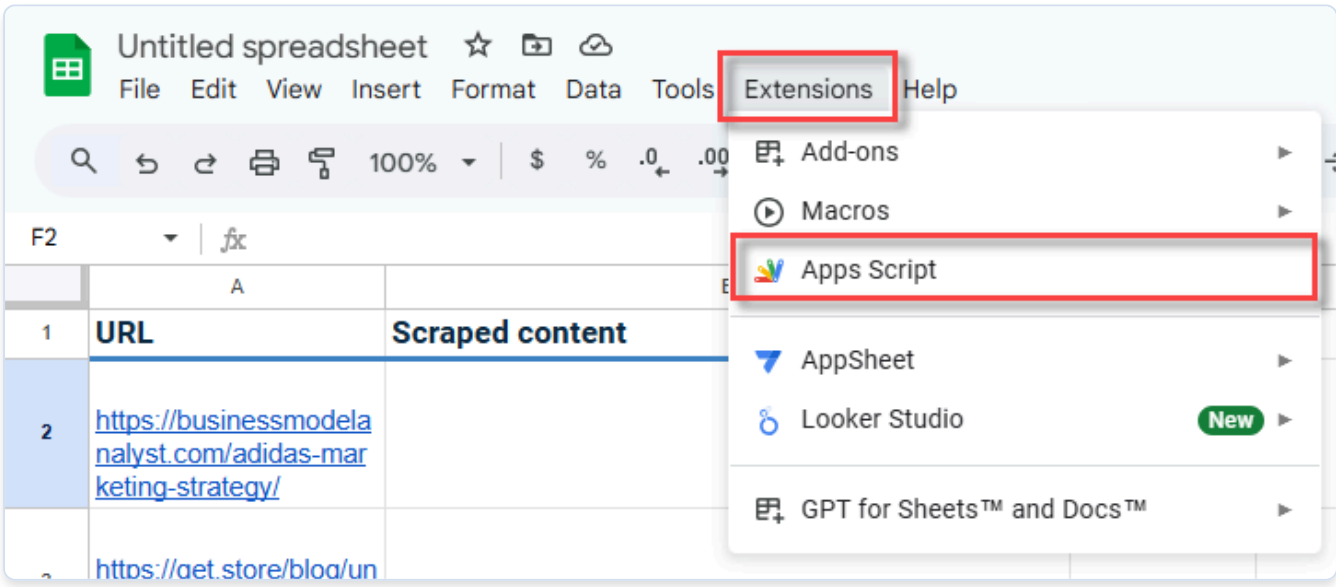
> **TIP**
>
> You can also use the Fetch content from URL bulk tool.

1. Install the script in your spreadsheet.
2. Launch scraping:
   - On multiple URLs at once using the **Scrape** menu
   - On a single URL using the GPT_SCRAPE function

Install the script in your spreadsheet

Install the web scraping script in your spreadsheet to enable the import of web page content:

1. From the **Extensions** menu in your spreadsheet, click **Apps Script**.



   A new project opens on the *Apps Script* page.

2. Copy and paste the script into the script editor.

```
/**
* @OnlyCurrentDoc
*/

const SHEETS_CELL_LIMIT = 50000;
```

```
function onOpen() {
  SpreadsheetApp
  .getUi()
  .createMenu("Scrape")
  .addItem("Scrape URLs in column A", "ScrapeAB")
  .addToUi()
}

/**
  * @param {"https://en.wikipedia.org/wiki/ChatGPT"} url The url to scrape
  * @customfunction
  */
function GPT_SCRAPE(url) {
  const res = UrlFetchApp.fetch(url);
  if (res.getResponseCode() >= 300) {
    return `Error: response code ${res.getResponseCode()}`;
  }
  const html = res.getContentText();
  const text = cleanup_html(html);
  // html.substring(0,SHEETS_CELL_LIMIT),
  return text.substring(0,SHEETS_CELL_LIMIT);
}

function cleanup_html(html) {
  // remove tags
  const tag_free = html
    .replace(/<style([\s\S]*?)<\/style>/gi, "")
    .replace(/<script([\s\S]*?)<\/script>/gi, "")
    .replace(/<[^>]+>/g, "");

  // replace html entities
  let decoded = tag_free;
  try {
    const xml = XmlService.parse('<d>' + tag_free + '</d>');
    decoded = xml.getRootElement().getText();
  } catch (e) {
    // unable to decode some HTML entity
  }

  // remove repeated spaces
  const clean = decoded
    .replace(/\n\s*\n/gi, "\n")
    .replace(/ /gi, " ")
    .trim();

  return clean;
}
```

```
function ScrapeAB() {
  const range = SpreadsheetApp.getActiveSpreadsheet().getRange("A:B");
  const values = range.getValues().slice(1)

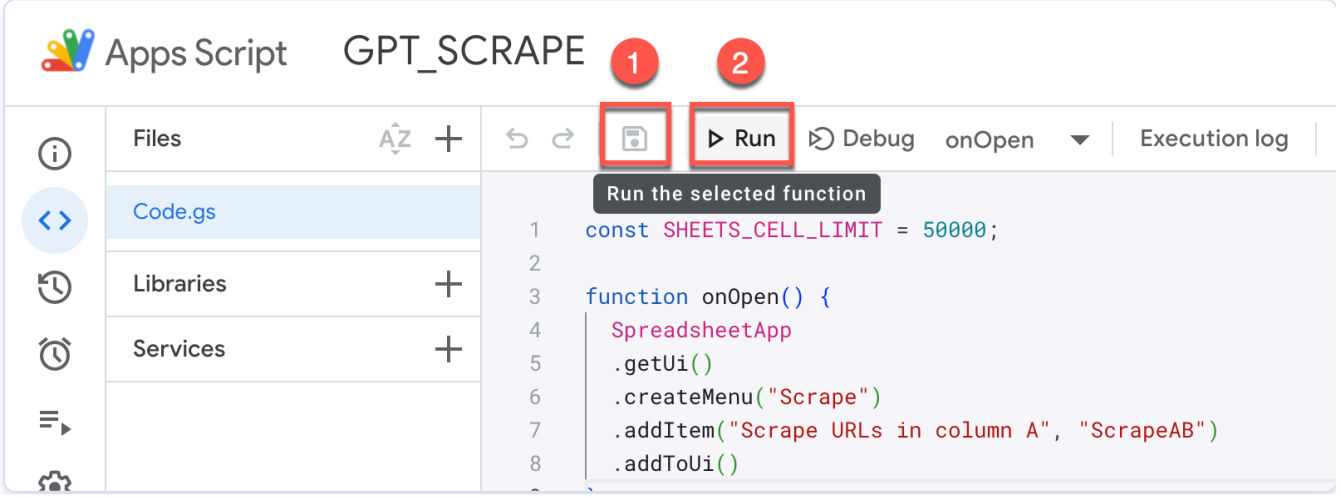
  let urls = 0
  let done = 0

  for (const row of values) {
    const url = row[0];
    const res = row[1];
    if (!url) continue;
    urls++;
    if (!res) continue;
    done++;
  }

  if (urls === done) {
    SpreadsheetApp.getUi().alert(`Found ${urls} urls, ${done} already
scraped. Nothing to be done.`)
    return;
  }
  SpreadsheetApp.getUi().alert(`Found ${urls} urls, ${done} already
scraped. Scraping ${urls - done}...`)

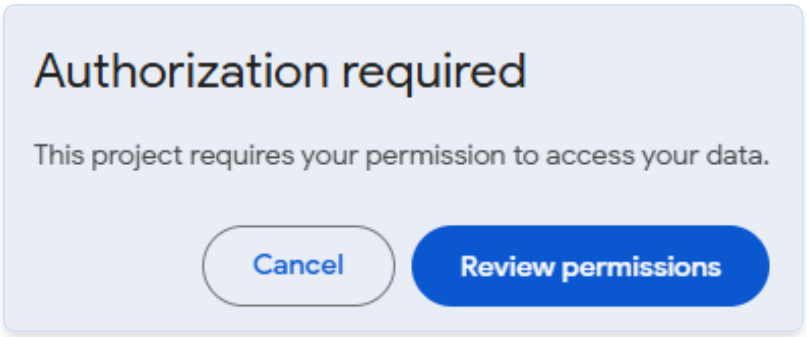
  for (let irow = 0; irow < values.length; irow++) {
    row = values[irow];
    const url = row[0];
    const res = row[1];
    if (!url || res) continue;
    const text = GPT_SCRAPE(url)
    range.getCell(irow+2, 2).setValue(text);
  }
}
```

3 Save the project and click **Run**.



An *Authorization required* pop-up appears.

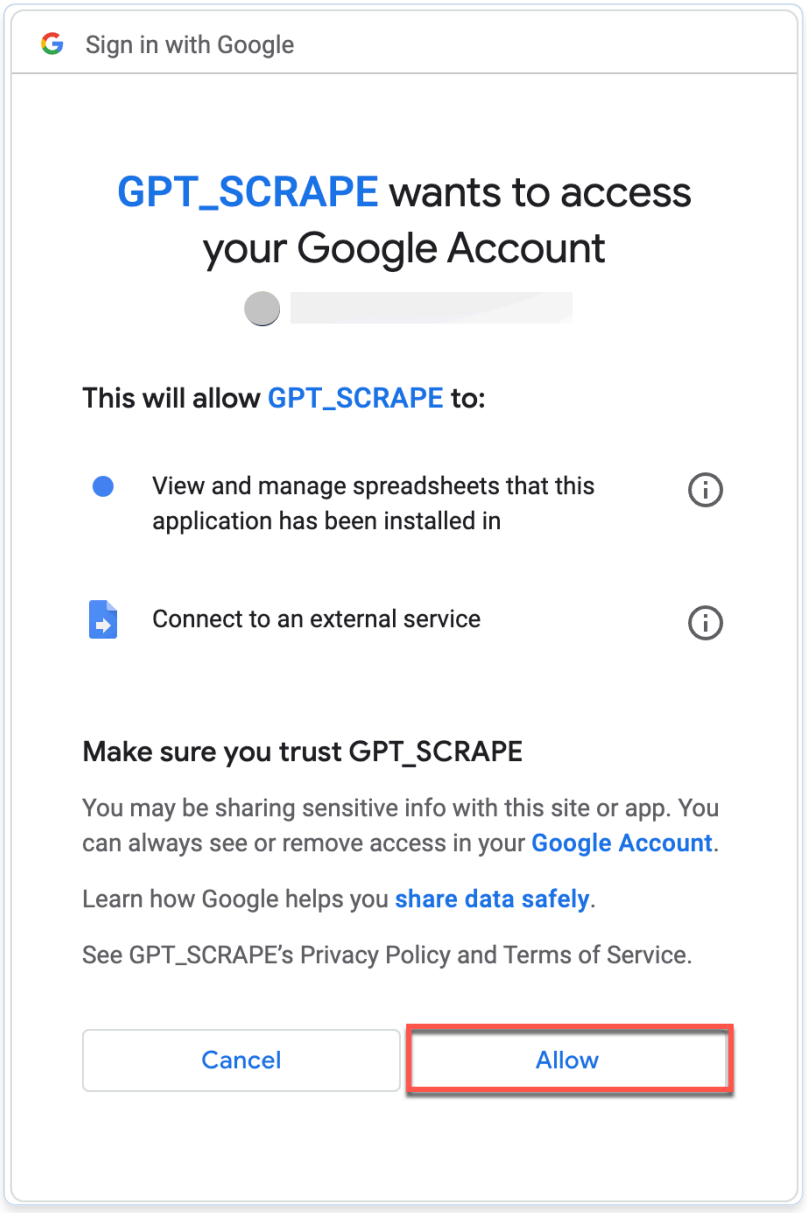
4 Click **Review permissions** to accept the necessary access rights for the script.



You are prompted to select a Google account.

5 Choose the account you are using for the spreadsheet. This ensures the script works correctly within your document.

6 Click **Allow** to grant these permissions.



INFO

The following permissions are necessary for the script to run:

- **View and manage spreadsheets that this application has been installed in**: Required for the script to specifically interact with the spreadsheets where it's installed, whether it's adding, modifying, or removing data.
- **Connect to an external service**: Required for the script to browse web pages and import the content to your Google Sheets.

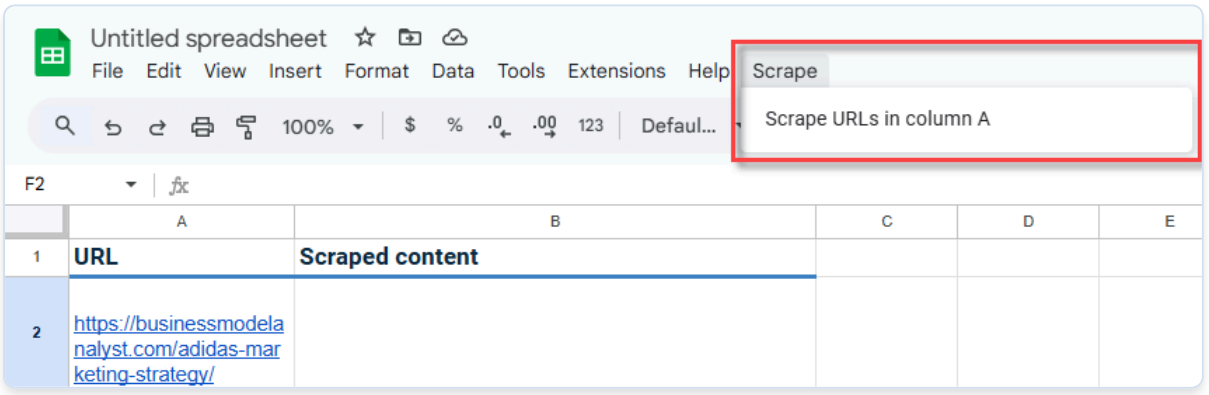
You can now scrape multiple URLs at once in your spreadsheet or scrape a single URL using GPT_SCRAPE.

Scrape multiple URLs with the Scrape menu

Prerequisites

- You have installed the web scraping script in your spreadsheet.
- You have entered the URLs in column A of your spreadsheet, starting from the second row.

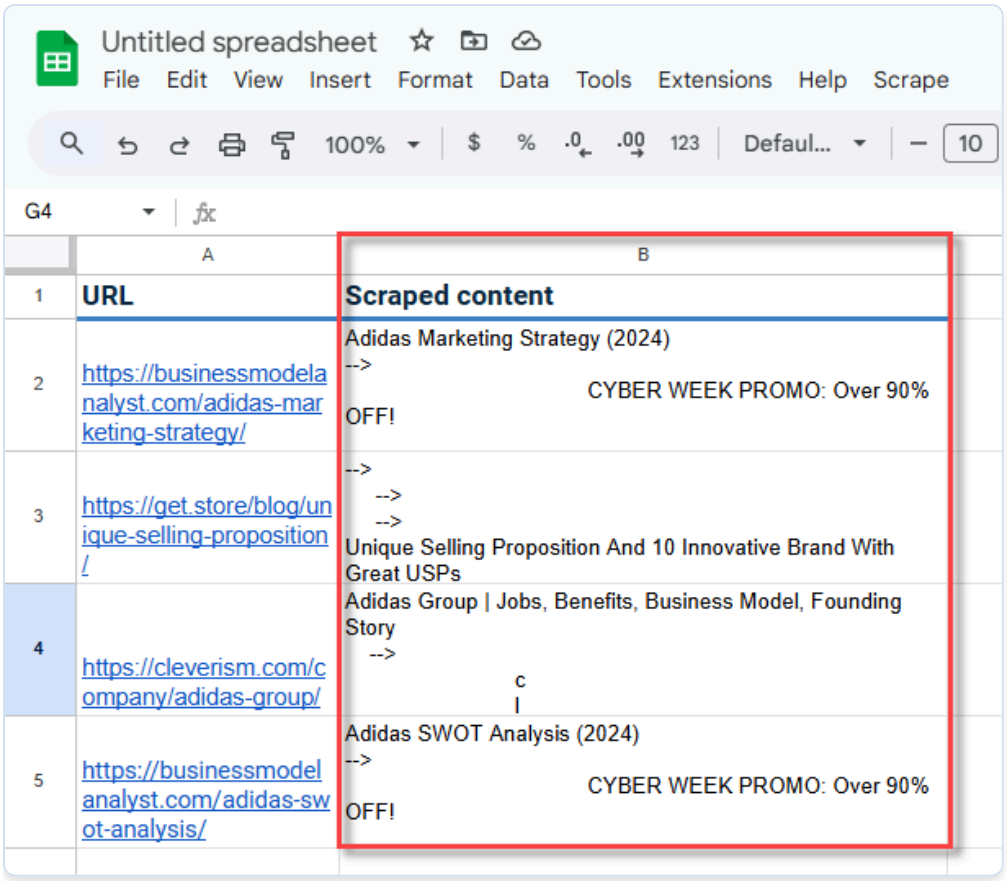
1. From the **Scrape** menu, select **Scrape URLs in column A**.



   A confirmation pop-up appears.

2. Click **OK** to launch the script.

This script processes the web page URLs in column A, and write the scraped content into column B, ready for you to edit directly. It removes complex codes and styles, leaving only text, and includes the page's title and main body text.

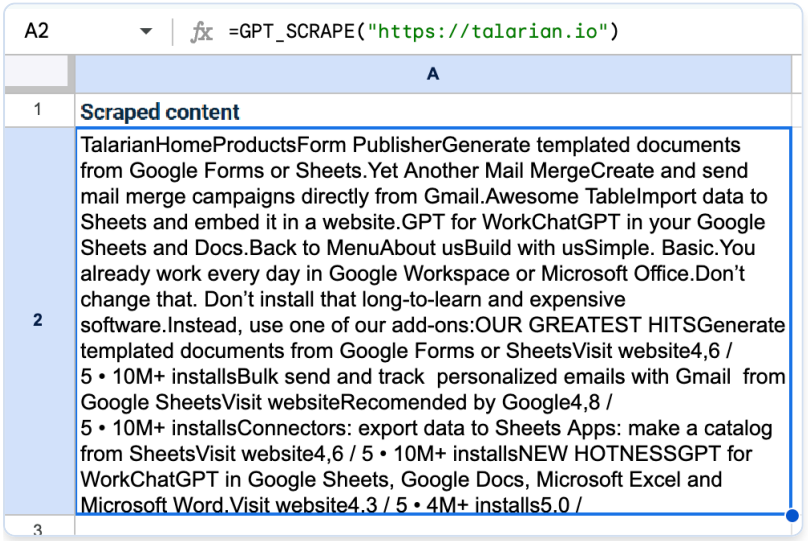


Scrape a single URL with GPT_SCRAPE

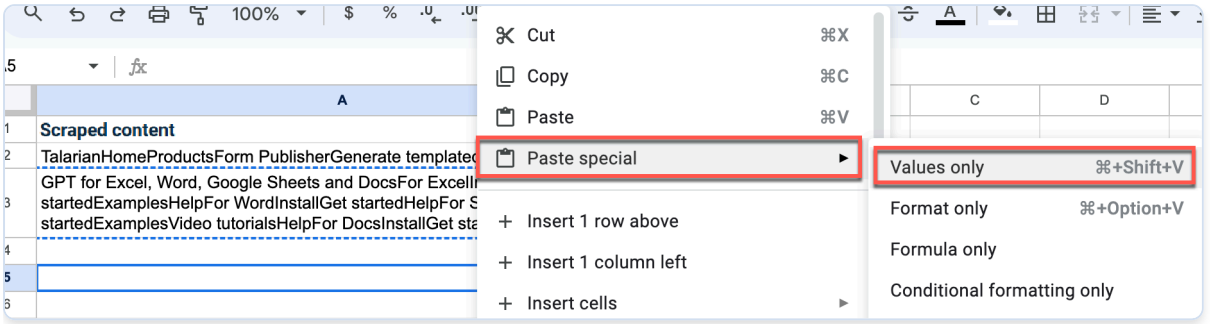
Prerequisites

You have installed the web scraping script in your spreadsheet.

Type a GPT_SCRAPE formula into a cell, for example `=GPT_SCRAPE("https://talarian.io")`, and press **Enter** to generate the result:



This function browses content from the specified web page and imports it into the cell. To edit a cell's content, copy the cell and use **Paste special**, and **Values only** for pasting.



INFO

Using the GPT_SCRAPE function does not affect your credit usage.

Limitations

Note the following limitations of the web scraping script:

- Caution is advised when scraping content from web pages to adhere to website policies and prevent server overload. Spacing out requests can prevent issues such as IP blocking or service restrictions.
- Each import is capped at 50,000 characters of text. Content beyond this limit will be truncated.
- Web pages heavily reliant on JavaScript might not be fully captured by this script.
- This script is optimized for HTML content. It may not perform optimally with APIs that return structured data, such as JSON.

What's next

- Use bulk tools on the scraped content.