

Gemini Deep Research

Save hours of work with Deep Research as your personal research assistant. Now with the ability to draw context from your Gmail, Drive and even Chat in addition to the web, and transform reports into interactive content in Canvas.

Try Deep Research today



Gemini

Get up to speed on just about anything with Deep Research, an agentic feature in Gemini that can automatically browse hundreds of websites and even your Gmail, Drive and Chat on your behalf, think through its findings and create insightful multi-page reports in minutes.

With the Gemini 3 model, Deep Research is even better at all stages of research, from planning to delivering even more insightful and detailed reports.

Planning

Deep Research transforms your prompt into a personalised multi-point research plan

Searching

Deep Research autonomously searches and deeply browses the web – and your Gmail, Drive and Chat if you so choose – to find relevant, up-to-date information

Reasoning

Deep Research shows its thoughts as it reasons over information gathered iteratively and thinks before making its next move

Reporting

Deep Research provides comprehensive custom research reports with more detail and insights, generated in minutes and available as an Audio Overview, saving you hours of time

Gemini

Gemini Deep Research is designed to tackle your complex research tasks by breaking them down, exploring sources like websites and your Workspace content if you choose it. Once Deep Research has found answers, it will synthesise its findings into comprehensive results.

You can also upload your own files to Deep Research and make your reports even more immersive by turning them into interactive content, quizzes, Audio Overviews and more in Canvas.

Competitive analysis

Building a competitor report that cross-references public web data with your internal strategy memos, feature-comparison spreadsheets and team chats about a rival product.

Due diligence

Topic understanding

Diving deep into subjects by comparing and contrasting key concepts, identifying relationships between ideas and explaining underlying principles.

Product comparison

Evaluating different models of an appliance based on features, performance, price and customer reviews.

It's a step towards more agentic AI that can move beyond simple question answering to become a true collaborative partner capable of sophisticated thinking and execution.

How to access Deep Research

Try Deep Research today

- On desktop
- On mobile
- In 150 countries
- In 45+ languages
- Available to Google Workspace users

Just select Deep Research in the prompt bar to get started and let Gemini do the research for you.

How we built the first Deep Research

The day after we pioneered the Deep Research product category on Gemini in December 2024, we gathered some of the team behind the product for a discussion.



Watch on

An agentic system

To build Deep Research, we developed a new planning system that enables the Gemini app to work through complex problems. For Deep Research, we trained Gemini models to be capable of:

- **Breaking down the problem:** When presented with a complex user query, the system first formulates a detailed research plan, breaking the problem into a series of smaller, manageable sub-tasks. You're in control of the plan: Gemini presents it to you, and you can refine it to make sure that it's focused on the right areas.
- **Research:** The model oversees the execution of this plan, and intelligently determines which sub-tasks can be tackled simultaneously and which need to be done sequentially. The model can use tools like search and web browsing to fetch information and reason over it. At each step, the model reasons over

Gemini

- **Synthesis:** Once the model determines that enough information has been gathered, it synthesises the findings into a comprehensive report. In building the report, Gemini critically evaluates the information, identifies key themes and inconsistencies, and structures the report in a logical and informative way, even performing multiple passes of self-critique to enhance clarity and detail.

New category, new problems, new solutions

In building Deep Research, we had to work through three significant technical challenges:

Multi-step planning

Research tasks require multiple steps of iterative planning. At each step, the model has to ground itself on all information gathered so far, then identify missing information and discrepancies that it wants to explore – all while trading off comprehensiveness with compute time and user-waiting time. Training the model to be effective at long multi-step planning in a data-efficient manner enabled us to make Deep Research function in an open domain setting across all topics.

Long-running inference

A typical Deep Research task involves many model calls over several minutes. This creates a challenge for building agents: It has to be built so that a single failure doesn't mean having to restart the task from the beginning.

Gemini

for graceful error recovery without restarting the entire task. This system is truly asynchronous: you can hop to a different app or quite literally turn off your computer after starting a Deep Research project and the next time that you visit Gemini, you'll get notified when your research is done.

Context management

Over the course of a research session, Gemini can process hundreds of pages of content. To maintain continuity and enable follow-up questions, we use Gemini's industry-leading 1 million token context window complemented with a RAG setup. This effectively allows the system to 'remember' everything it has learned during that chat session, making it smarter the longer that you interact with it.

Evolving with new models

When Deep Research launched in December, it was powered by Gemini 1.5 Pro. With the introduction of Gemini 2.0 Flash Thinking (experimental), we were able to dramatically improve both the quality and serving efficiency of this product. With thinking models, Gemini takes more time to plan out its approach before it makes its next steps. This innate characteristic of self-reflection and planning makes it a great fit for these kinds of long-running agentic tasks. What we see is that Gemini is now even better at all stages of research and delivers more detailed reports. At the same time, the compute efficiency of the Flash model allows us to expand access to Deep Research to far more users. We're really excited about developing Flash and Thinking models in general, and expect Deep Research to keep getting better and better.

And with our most capable model, Gemini 3, Deep Research is even better at all stages of research, delivering even more insightful and detailed reports

Gemini

capabilities by giving you more control over what it can browse and giving it sources beyond the open web.

We are excited to see how people use Deep Research, and these real-world experiences will inform how we continue to build and improve Deep Research. Ultimately, our goal is a truly agentic and universally helpful AI assistant.

Agentic Gemini



Reason



Search



Browse

Gemini's **new agentic AI system** brings together the **best of Gemini, Google Search and web technologies** to continuously search, browse and think through information in a continuous reasoning loop for more comprehensive results.

Check responses. Compatibility and availability vary. 18+.



What Gemini can do

Gemini Live

Image generation

Gemini

Deep Research

Personal Intelligence

Canvas

Storybook

Apps

Gems

Long context

Students

Import your memory

About Gemini

Overview

Our approach

Policy guidelines

Gemini Drops

Latest news

Subscriptions

Release notes

Gemini for macOS

YouTube

Gemini

[Instagram](#)

[GitHub](#)

[Discord](#)

Google

[Privacy](#)

[Terms](#)



South Africa