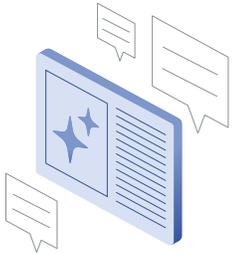


AI Overviews and the AI Mode experiment in Search





AI Overviews and the AI Mode experiment in Search

For decades, we’ve continued to reimagine and expand what Google Search can do. Now, our newest AI technologies are making it possible to create search experiences that feel much more effortless than ever before.

As we continue to apply AI to help people easily access information and explore the richness of the web, read on to learn more about how generative AI capabilities like AI Overviews and our new AI Mode experiment in Search work, how they can be helpful, and how we have prioritized quality and safety.

AI Overviews: The fundamentals

AI Overviews use generative AI to provide key information about a topic or question, with links to dig deeper and learn more on the web. More than a billion people around the world now use AI Overviews for help with their questions.

This application of generative AI has been specifically designed to be helpful for information journeys in Search, meaning it’s a different experience than interacting with an LLM-based chatbot. AI Overviews

use a customized Gemini model, which works in tandem with our existing Search systems – like our quality and ranking systems and the Google Knowledge Graph. AI Overviews are designed to carry out traditional “search” tasks, like identifying relevant, high-quality results from our index to support the information presented in the overview.

To ensure that AI Overviews are high quality, we’ve integrated our core web ranking systems into this experience, which are fundamentally designed to surface reliable and relevant information. AI Overviews are built to surface information that is backed up by top web results, and include links to web content that supports the information presented in the overview. People can use these links to dig deeper on a topic, because we know people come to Search to find a range of perspectives and to explore the open web.



How AI Overviews are helpful

Before rolling out AI Overviews in Search, we tested the experience extensively through Search Labs to help us understand where generative AI can be most helpful. Like all the Search features that we build, our goal is to help people find the information they're looking for quickly and reliably.

In particular, AI Overviews are helpful for more complex questions that might have previously taken multiple searches. AI Overviews help people get to the gist of a complicated topic or question more quickly, while also providing a jumping off point to explore links to learn more. AI Overviews were designed to show up on queries where they can add additional benefit beyond what people might already get on Search, and where we have high confidence in the quality of the responses. Meaning, they will show up in Search results when our systems determine that generative AI can be especially helpful – for example, when you want to quickly understand information from a range of sources.

With AI Overviews, people are visiting a greater diversity of websites for help with more complex questions. We've found that people who use AI Overviews actually use Search more and are more satisfied with their results. And when people click from search result pages with AI Overviews, these clicks are higher quality, where users are more likely to spend more time on the site, because we've done a better job of finding the right info and website for them.

Quality and safety protections

AI Overviews are rooted in our core Search quality and safety systems, which we have been refining for more than 20 years to deliver reliable and high-quality information. Some of the key ways we've worked to promote quality and safety in AI Overviews include:

- **Rigorous testing:** As we do with all Search features, we undertook a rigorous testing process before launch to maintain our high bar for quality with AI Overviews. We've utilized methods that we've long deployed on Search, like our Search quality rating program and side-by-side analyses. We've also engaged in a thorough adversarial red-teaming process designed to identify and address potential issues, including novel approaches focused on issues specific to generative AI. We are continually learning from each launch, and we are committed to improving our processes when needed.
- **Corroborating web results:** As we've mentioned, we designed AI Overviews to only show information that is supported by high-quality results from the web, meaning that AI Overviews generally don't "hallucinate" in the ways that other LLM experiences might. For "Your Money or Your Life" (YMYL) queries – where information quality is critically important – we have an even higher bar for showing supporting information from reliable and trustworthy sources. In addition, for many such queries, AI Overviews will inform people when it's important to seek out expert advice or to verify the information presented.
- **Safety Guardrails:** Similarly, we've incorporated our core safety systems, including technologies like SafeSearch, to prevent harmful, hateful, or otherwise explicit or shocking content from appearing in AI Overviews.
- **Policies:** AI Overviews adhere to our overall content policies for Search, as well some of our unique policies for Search features like featured snippets and autocomplete, which include careful considerations for content that may be explicit, graphic or cause real-world harm for example. We design our systems to automatically prevent policy violating content from appearing – but in the rare cases when such content does evade our systems, we take action.



- **Spam:** We utilize our core anti-spam protections to safeguard AI Overviews from low-quality content, and we've recently launched tailored updates to prevent spammy results from showing up in AI Overviews specifically. We constantly develop new techniques and implement updates to our ranking systems to protect against spam in AI Overviews and in Search more broadly, using systems like our AI-based spam-prevention system known as [SpamBrain](#).
- **Triggering:** AI Overviews are designed to show up on queries where they can add additional benefit beyond what people might already get on Search today, and where we have high confidence in the quality of the responses. Meaning, they will show up in Search results when our systems determine that generative AI can be especially helpful – for example, when you want to quickly understand information from a range of sources. As a result, we've worked to refine our systems with the goal of preventing AI Overviews from appearing on queries about highly sensitive, explicit or dangerous topics. For topics where up-to-the-minute accuracy is critical, such as current events or breaking stories, AI Overviews may not appear.
- **Addressing data voids:** AI Overviews are designed to show information supported by high quality web results, but for some searches (like uncommon or nonsensical searches), there might not be a lot of high quality information available on the web. These situations are often called "data voids," and can sometimes lead to lower quality information appearing. We've rolled out improvements designed to limit AI Overviews from appearing on these searches, and to prevent misleading, satirical, or otherwise unhelpful content from surfacing in AI Overviews.

Continued Innovation, Improvement, and Refinement

We are constantly measuring and have a detailed review process to assess the quality of AI Overviews across many different categories, including attributes like factuality, length, format and clarity of responses.

Since we began experimenting with this technology in [Labs](#), we've dramatically improved the quality of responses, with major gains in areas like factuality and freshness.

We approach improvements to AI Overviews in the same way that we approach Search more broadly: we learn from examples and feedback to identify areas of improvement. We come up with solutions that we believe could help not just one query, but a broad range of similar searches. People can [provide feedback](#) on the AI Overviews they see in Search, which helps us improve.



Introducing AI Mode, starting as a Labs experiment

As we've rolled out AI Overviews, we've heard from our power users that they want a way to get these AI-powered responses for even more of their searches. So we're beginning to test a new [AI Mode](#) as an opt-in experiment in Labs. This new experimental mode in Google Search expands on AI Overviews with more advanced reasoning, thinking and multimodal capabilities. We're starting this as a limited experiment in Labs, and we'll evolve the experience as we get feedback and learn what's most helpful.

Read on for an introduction to how this early experiment works, and for information about our core quality and safety protections.

How AI Mode works

Using a custom version of Gemini 2.0, this opt-in, experimental AI Mode is particularly helpful for queries where further exploration, reasoning, or comparisons are needed. With this new mode, people can ask nuanced questions that might have previously taken multiple searches – like exploring a new concept to comparing options and beyond – and get a comprehensive AI-powered response.

This new experimental AI Mode taps into real-time information from the web and from Google, like facts from our Knowledge Graph, info about the real-world, and shopping data for billions of products. It uses a “query fan-out” technique, issuing multiple related searches concurrently across subtopics and multiple data sources and then brings those results together to provide an easy-to-understand response. This approach helps you access more breadth and depth of information than a traditional search on Google.

People can also ask conversational follow-up questions to continue their search journeys. Thanks to the model's native ability to recognize relationships between consecutive queries, it will carry over context from previous questions to refine the user's intent and formulate a more precise, informed search.

And because AI Mode will be multimodal (meaning it will be able to process your question through voice, text or images), it will allow people to ask questions in a variety of different ways.

How to access:

This new mode is an opt-in experiment available in [Labs](#).

Once you're opted-in, there are different ways to access this new experience:

- Go to www.google.com, enter a question in the Search bar, and tap the “AI Mode” tab below the Search bar.
- Go directly to the AI Mode tab on Google Search at: google.com/aimode.
- In the Google app, tap the AI Mode icon below the Search bar on the home screen.



Making it easier to ask new questions and explore the web

At its core, AI Mode – like AI Overviews – has been designed to help people ask new questions and connect them with the breadth of content and perspectives across the web. People really value access to a range of information, so there are a number of ways we’re designing the experience to enhance user interaction with web content.

Just like AI Overviews, AI Mode prominently surfaces relevant links to help people find web pages and content they may not have discovered before. And as outlined above in “How AI Mode works,” the experiment uses a query fan out technique – issuing multiple related queries concurrently across subtopics and multiple data sources. While responses are being generated, our advanced models are now able to identify and access even more supporting web pages than was previously possible. This additional step allows us to display a wider and more diverse set of helpful web content associated with the response, enabling new opportunities for exploration. Because AI Mode is using more advanced models and novel techniques compared to AI Overviews, a different set of responses and links will likely appear for the same query.

As this is an early experiment, we’ll continue to update and refine the user interface and expand functionality, including new ways to display and help people engage with content across the web. For example, we’ll be making the experience even more visual and adding more types of rich content formats from a range of sources, like user-generated content.

We’re also training the models to intelligently determine when and how to link and best present information so it’s most useful and actionable. For example, teaching the model to decide when to include hyperlinks in the response if it’s likely that the user may want to take action or finish a task on a website (e.g. booking tickets). Or deciding when to prioritize visual information if the user’s question could benefit from an image or video (e.g. how-to queries).

Quality and safety protections

Like AI Overviews, AI Mode is rooted in our core Search ranking and safety systems and anti-spam protections, which we have been refining for more than 20 years to deliver reliable and helpful information. It also adheres to our overall content [policies](#) for Search, as well some of our policies for Search features.

AI Mode is designed to show information that is supported by high quality content from the web and other info sources. We’re also using novel approaches with the model’s reasoning capabilities to improve factuality. For example, in collaboration with Google Deepmind research teams, we use agentic reinforcement learning (RL) in our custom training to reward the model to generate statements it knows are more likely to be accurate (not hallucinated) and also backed up by inputs. The agentic RL also encourages the model to retrieve and add additional factual information to its claims from info sources like our Knowledge Graph.

As with any early-stage AI product, AI Mode won’t always get it right. That’s why we’re starting out as an experiment available only as an opt-in through Labs, and people have to click to access the mode. As we did with AI Overviews, before launching this new experimental AI Mode in Labs, we tested it extensively with trusted testers and conducted rigorous internal evaluations using methods we’ve been honing for decades in Search. We’re continuing to diligently test and identify areas for improvement.

Because the new mode will have some distinct capabilities (such as multimodality and allowing for follow-up questions), there are unique considerations for this experiment, including:

- **Triggering:** This experiment uses advanced reasoning on a wide range of queries. In some cases, our systems might determine that a question is better answered by links out to the web or existing Search features rather than a generative response. For example, over the years, our Search systems have gotten better at detecting if a query might indicate that



someone is in a vulnerable situation (such as self harm) or if someone is seeking explicit content. This experimental AI Mode leans on these Search systems, and aims to show a list of links and existing features, where appropriate. Out of an abundance of caution in this early testing phase, this mode also will generally respond with web search results if there is not high confidence in the helpfulness and quality of the AI-powered response. This is more likely to be the case for queries related to current events, where up-to-the-minute accuracy is important.

As was the case with AI Overviews, this experiment will answer more questions over time as we improve its capabilities and have higher confidence that responses will provide additional value for a query. In general, AI Mode will provide an AI-powered response for far more queries and more complex questions than AI Overviews do today, as it uses more advanced models.

- **Inaccuracies and Odd Responses:** AI Mode relies on Google Search’s deep understanding of web content – meaning that responses are backed up by high quality information – helping to mitigate against hallucinations and keep responses factual. However, this AI Mode experiment uses a large language model to help answer queries and it is possible that, in rare cases, it may sometimes confidently present information that is inaccurate, which is commonly known as “hallucination.” As with AI Overviews, in some cases this experiment may misinterpret web content or miss context, as can happen with any automated system in Search. And as with Search more broadly, you might potentially get odd responses from this new mode when you search for unusual things, as the experiment reflects the information that’s on the web that is relevant to the search.
- **Opinionated Responses:** We’ve designed this experimental AI Mode to present information as objectively and neutrally as possible based on what’s available from a wide range of sources across the web, and to avoid showing responses that could be interpreted as taking a position

on a disputed topic. Just as in search results, if the vast majority of information available on the web on a given topic represents a particular perspective or view, it’s possible that responses may reflect this imbalance. You can always ask another follow-up question to get more info, or click out of the mode to Search.

In addition, while we’ve designed this new mode to provide objective, neutral, and factual responses, this experiment may in some cases appear to take on a “persona,” if prompted by a user. While this is a known limitation of LLMs generally, we’ve designed protections that limit persona responses in this new mode.

- **False Equivalence Responses:** As previously noted, this experiment is designed to give objective, neutral responses based on information from the open web. Because it is designed to be objective and comprehensive, when asked to compare two topics, this experiment may occasionally provide responses that appear to equate or liken two topics or concepts that are not equivalent to one another.
- **Carrying Context Across Follow-up Questions:** When you ask follow-up questions, AI Mode is designed to carry over context from your previous question to provide more relevant results. For example, if you ask AI Mode to compare cats and dogs, and then ask “Which one lives longer?”, the answer will still be about cats and dogs. The systems are also designed to not carry over context if you do another search on a completely unrelated topic.

There may be some cases where this won’t work as intended: if you ask multiple follow-ups, context may not fully carry across. It’s also possible that context is carried across questions that the user meant to be separate intents, which can lead to results that feel less relevant or not what they were looking for. In cases where AI Mode is showing only web links, if previous query context is being applied, those results may not match what’s on the search results page if you did a standalone search with that same query.



- **Query Variety:** There are many ways to ask the same or similar questions. In some cases, this experiment may offer different responses to two queries that seem similar or are about the same topic, but are framed differently (for example, “Why is an iguana a good pet?” vs. “Why isn’t an iguana a good pet?”). This can occur in Search, as well, because our systems aim to provide results that are responsive and relevant to what a person is looking for. Our systems may also show a range of perspectives that are available on the web, where some differences of opinions may exist on a query.
- **Satire and humor:** We have made strong improvements in AI Overviews to detect satire and humor queries where outputs could be offensive, unsafe or misleading and we are extending these protections to this experiment. This AI Mode experiment aims to trigger on a wide range of queries and we’re continuing to make improvements to our detection mechanisms in this space, building on the protections we have put in place for AI Overviews.
- **Math:** The new experimental AI Mode uses Gemini 2.0’s advanced reasoning capabilities to help with complex questions like math equations. This is a known challenge area for large language models and while we have made significant improvements, errors can still occur. We’ll continue to invest in and improve on this in Labs.

Building the future of Search together

Part of our approach to experimenting with AI Mode is setting clear expectations on how it works, the protections we have and being mindful of the limitations that still exist. By making this mode first available through [Labs](#), we’re giving people the opportunity to explore and test this new technology while it’s still in an experimental phase.

We look forward to learning from user feedback, continuing to improve and iterate, and building the future of Search together.