



How AI Overviews in Search work

Recently, we launched AI Overviews on Search. In this document, we'll explain how AI Overviews work, how they can be helpful, and the work we've done to prioritize quality and safety in this experience.



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.

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 corroborate 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 only 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 can be 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 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.

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 to links from 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, hateful, violent, or contradictory of consensus on public interest topics, 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 were 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.

We're also [restricting](#) the types of election-related queries that trigger an AI Overview, out of an abundance of caution. And, we aim to not show AI Overviews for hard news topics, where freshness and factuality are particularly important.

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 [Search 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.



AI Overviews & more in Search Labs

Search Labs will continue as a testbed for innovation and experimentation in Search. We've learned a lot over the last year with AI experimentation in Search Labs to help inform how to roll out gen AI experiences more broadly in Search – from improving latency for responses, to increasing clicks to sites and driving ecosystem growth.

People will be able to opt-in to *AI Overviews & more* through Search Labs, where they'll get AI Overviews on more queries. People who are opted into *AI Overviews & more* will see the Labs icon by their AI Overview, indicating that they are in the Search Labs experience, which is experimental. Through *AI Overviews & more*, we'll continue to test new topics, models and query types to continue to learn more about how AI Overviews can be most helpful.

Overall, we'll keep improving when and how we show AI Overviews and strengthening our protections, and we're very grateful for ongoing feedback.