A new way to search with generative AI

An overview of SGE

Author: Srinivasan Venkatachary and collaborators

We're excited to bring cutting-edge, generative AI front and center in Google Search. For decades, AI has been helping Google Search under the hood, allowing us to reimagine how people interact with and discover information, improve quality and relevance, and support a healthy open web. In fact, one of our first applications of machine learning in a Google product was our early spelling correction system back in 2001 — over two decades ago — which helped people get relevant results faster, regardless of spelling mistakes or typos.

In recent years, breakthroughs in AI have dramatically improved Search. In 2019, we brought Bidirectional Encoder Representations from Transformers (BERT) to Search ranking, resulting in a huge step change in search quality. Rather than aiming to understand words individually, BERT helped Search understand words in the context that they were used, empowering people to ask longer, more conversational queries and get connected with more relevant, helpful results.

We've since applied even more powerful large language models (LLMs) to Search, like our Multitask Unified Model (MUM) — a model that is 1,000 times more powerful than BERT, trained across 75 different languages and many different tasks simultaneously, and is multimodal. MUM has been deployed in dozens of features across Search to improve quality and help us understand and organize information in new ways. For example, we applied MUM to help people find related topics in videos, even when the topics aren't explicitly mentioned.

But we've only scratched the surface of what's possible with generative AI. These models will enhance Google Search and revolutionize how people engage with information. We're experimenting with what's possible through our new program Search Labs, starting with an experiment called **SGE (Search Generative Experience)**.

Here we outline how we're approaching SGE — what it is, how it works and the steps we've taken in developing it responsibly. Our approach to SGE will evolve as we learn from ongoing research, experience and user feedback in Labs. To sign up for Search Labs, available in English in the U.S. to start, visit labs.google.com/search.

What is SGE?

SGE is an early step in transforming the Search experience with generative AI. When using SGE, people will notice their search results page with familiar web results, organized in a new way to help them get more from a single search.

With SGE, people will be able to:

- Ask entirely new types of questions that you never thought Search could answer

- Quickly get the lay of the land on a topic, with links to relevant results to explore further

- Ask follow-up questions naturally in a new conversational mode

- And get more done easily, like generating creative ideas and drafts right in Search

SGE is rooted in the foundations of Search, so it will continue to connect people to the richness and vibrancy of content on the web, and strive for the highest bar for information quality.

How SGE works

AI-powered snapshots

When appropriate, SGE will show an AI-powered snapshot to help people quickly get an overview, with factors to consider and a helpful synthesis of relevant insights and information.

These snapshots serve as a jumping-off point from which people can explore a wide range of content and perspectives on the web. SGE will show links to resources that corroborate the information in the snapshot, so people can check the information themselves and explore further. This allows people to dig deeper and discover a diverse range of content, from publishers, creators, retailers, businesses, and more, and use the information they find to advance their tasks.

Conversational mode

People can tap to "ask a follow up" or on one of the suggested next steps under the snapshot. This will launch the new conversational mode, where they can naturally ask Google more about the topic they're exploring. Beyond the succinct snapshots of information generated by AI, people will be able to go deeper with additional links to resources to explore.

Conversational mode is especially useful for follow-up questions, as well as more complex or evolving information journeys. It uses AI to understand when a person is searching for something that is related to a previous question. It carries over context from previous questions to reformulate the query to better reflect the intent. In conversational mode, people will see their web links below SGE change throughout the course of the conversation so they can easily explore the most relevant content from across the web.

Vertical experiences

SGE will also be helpful for information journeys related to verticals, like shopping or local searches, as they often have multiple angles or dimensions to explore. In shopping, for example, SGE helps uncover insights so people can make considered, complex purchase decisions faster and easier.

For product searches, SGE will generate a snapshot of noteworthy factors to consider and a range of product options. SGE also provides product descriptions that include relevant, up-to-date reviews, ratings, prices and product images. This up-to-date product information is possible because SGE is built on Google's Shopping Graph, the world's most comprehensive dataset of constantly-changing products, sellers, brands, reviews and inventory.

Similarly, SGE will provide context about local places, using AI-powered insights that make it easy to compare and explore options.

Advertisements

As Search applies the power of generative AI, Search ads will continue to play a critical role. They serve as additional sources of useful information while helping people discover millions of businesses online.

With SGE, Search ads will continue to appear in dedicated ad slots throughout the page. In this new experience, advertisers will continue to have the opportunity to reach potential customers along their search journeys. We'll test and evolve the ads experience as we learn more.

As always, we're committed to transparency and making ads distinguishable from organic search results. When Search ads do appear, they will continue to feature our industry-leading clear and transparent ad labels with the "Sponsored" label in bold black text.

Creativity

These new generative AI capabilities can help people continue their journeys in more creative ways, going beyond just finding information to making use of it. We believe this has useful applications for people to carry out more creative steps of an information journey, like going from searching for a new ebike to helping write the perfect social post to show it off.

While SGE is adept at both informational and creative applications, users will notice constraints on creative uses to start, as we've intentionally placed a greater emphasis on safety and quality. Over time, we will continue to expand creative capabilities as quality improves.

User experience

Everything we do at Google is grounded in extensive user research. We've been evolving the Search user interface (UI) for many years to be more useful and accessible. To bring the power of AI to Search in a user-friendly way, we've built SGE as an integrated experience, applying what we've learned about user behavior. The AI-powered snapshot has easy-to-access resources and recognizable UI for links that allow people to further explore across both desktop and

mobile to bring the power of generative AI right into Google Search. Further, we help users make the transition into conversational mode through thoughtfully crafted call-outs and highlighted states that show the user how to use this new paradigm. For example, in conversational mode, a user can see how the AI stitches together the context of the query and its follow-ups to reformulate it.

Color also plays a big role in helping people clearly understand SGE is a new way to interact with search. For instance, the color container of the AI-powered snapshot will dynamically change, and this usage of color will evolve over the next few months to better reflect specific journey types and the query intent itself.

3

Applying generative AI responsibly

We are rolling out SGE thoughtfully as an experiment, and in accordance with our AI Principles. We took extensive steps and a careful, considered approach to develop this experience responsibly, leaning on protections and approaches that we've honed for years in Search. The following are some of our key considerations. We will continue to evolve and improve our approach over time, and we will engage with industry experts, policymakers, civil and human rights leaders, content creators and more as part of that process.

Training the model

SGE is currently powered by a variety of LLMs, including an advanced version of MUM and PaLM2. By using a variety of models across the feature set, we are able to further optimize and fine-tune the models to meet the unique needs of users and help them throughout their information journeys.

Many people have now engaged directly with LLMs, including through experiments like Bard. While SGE also applies LLMs, it has been purposefully trained to carry out tasks

specific to Search, including identifying high-quality web results that corroborate the information presented in the output. These models are used in tandem with our core ranking systems to deliver helpful and reliable results.

By constraining SGE to these specific tasks, including corroboration, we're able to significantly mitigate some of the known limitations of LLMs, like hallucination or inaccuracies. We further mitigate these challenges by using our existing Search quality systems and our ability to identify and rank high-quality, reliable information. More on that below.

Human input and evaluation

Human input and evaluation are important in developing products responsibly. Among other quality checks, we also utilize human input when training SGE's outputs, focusing on attributes like length, format and clarity.

Search has long applied robust processes for human input and evaluation to train and improve our ranking systems, and we've applied many of those well-tested principles and learnings to how we train and evaluate SGE, including:

- Raters: We work with independent Search Quality Raters to help us measure the quality of outputs and the results displayed. These ratings do not directly impact SGE's output, but are used to train the LLMs and improve the experience overall.

- Focused analysis: Following our process for significant launches in Search today, we analyze results across multiple broad, representative query sets, as well as conduct more focused studies to confirm responses meet our quality thresholds. In particular, we focus on topic areas that may be more susceptible to known quality risks, or that are more complex and nuanced. This includes classes of queries that may be at higher risk for safety or inclusion issues and aims to ensure our protections and responses are effective across those domains.

- **Red-teaming:** We conduct adversarial testing of these systems to identify areas where the systems aren't performing as intended. This helps identify bias issues, safety concerns and other areas where we can improve the product.

We are continuing to evolve and improve these approaches with new research-backed techniques.

Search quality systems

We know that people come to Search for trusted information – including to verify what they've heard elsewhere – and we hold ourselves to a high standard for reliable, helpful and high-quality information. That's why we built a customized integration of generative AI in Search that is rooted in our core Search ranking and quality systems, which we have been honing for decades. We've been bringing AI advancements to Search for many years, and developed a careful, rigorous evaluation process to ensure that any update maintains the high bar we set for delivering reliable results.

Just as our ranking systems are designed not to unexpectedly shock or offend people with potentially harmful, hateful, or explicit content, SGE is designed not to show such content in its responses.

We hold SGE to an even higher standard when it comes to generating responses about certain queries where information quality is critically important. On Search, we refer to these as "Your Money or Your Life" (YMYL) topics – such as finance, health, or civic information – areas where people want an even greater degree of confidence in the results. Just as we do on Search, for YMYL topics, SGE places even more emphasis on producing informative responses that

are corroborated by reliable sources. We've also trained the model to include disclaimers in its output, where appropriate. For example, on health-related queries where we do show a response, the disclaimer emphasizes that people should not rely on the information for medical advice, and they should work with medical professionals for individualized care.

There are some topics for which SGE is designed to not generate a response. On some topics, there might simply be a lack of quality or reliable information available on the open web. For these areas – sometimes called "data voids" or "information gaps" – where our systems have a lower confidence in our responses, SGE aims to not generate an AI-powered snapshot. SGE is also designed not to generate snapshots for explicit or dangerous topics, or for queries that indicate a vulnerable situation – for example, on self-harm queries, where our systems will instead automatically surface trusted hotline resources at the top of Search.

Policies

Our automated systems work to prevent policy-violating content from appearing in SGE. SGE also aligns with some of our unique policies for 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.

While our systems generally work well in catching policy-violating content before it appears, SGE is an experimental feature, and violating content might surface. In such cases, we may take action under our policies to prevent this violative content from appearing again. We also use these examples to inform future improvements to the models. Given that SGE is an experimental feature in Search Labs, we will continue to examine and appropriately evolve our policies to make the experience more helpful over time.

Factuality v. fluidity

Another intentional choice we made pertains to the fluidity of responses in SGE, both in the AI-powered snapshots and conversational mode.

We have found that giving the models leeway to create fluid, human-sounding responses results in a higher likelihood of inaccuracies (see limitations below) in the output. At the same time, when responses are fluid and conversational in nature, we have found that human evaluators are more likely to trust the responses and less likely to catch errors.

Given the trust people put in Search, we were intentional in constraining conversationality. What this means, for example, is that people might not find conversational mode in SGE to be a free-flowing creative brainstorm partner — and instead find it to be more factual with pointers to relevant resources.

Balancing the fluidity of the experience with information quality is important, and SGE will improve over time as we iterate on this balance.

Lack of persona

At times, LLMs have the potential to generate responses that seem to reflect opinions or emotions, since they have trained on language that people use to reflect the human experience. We intentionally trained SGE to refrain from reflecting a persona. SGE is not designed to respond in the first person, for example, and we finetuned the model to provide objective, neutral responses that are corroborated with web results.

Known limitations

While we've built a range of protections into SGE, there are known limitations of both LLMs and this experience in its initial, experimental form. The following are some of the loss patterns that we observed during evaluations and adversarial testing, and other limitations we expect in SGE. In many cases, we have already made improvements with model updates and additional fine-tuning, and we expect to make further progress as SGE develops.

- **Misinterpretation during corroboration:** We have seen some instances where SGE has appropriately identified information to corroborate its snapshot, but with slight misinterpretations of language that change the meaning of the output.

- **Hallucination:** Like all LLM-based experiences, SGE may sometimes misrepresent facts or inaccurately identify insights.

- **Bias:** Because SGE is trained to corroborate responses with high-quality resources – which are thus highly ranked results – it may show a snapshot that reflects a narrower range of perspectives than what is available on the web, thereby reflecting biases in those results. This can create the impression that the

model has learned this bias. However, it is more likely that SGE is providing outputs that reflect biases that are contained within top results. This is a phenomenon that occurs in search results today, as well. For example, authoritative organizations and media entities often do not add a qualifier of "male" when writing about male sports, and generic queries about that sport may thus bias towards male players or teams, even if information about female players or teams is an equally or perhaps even more accurate response.

- Opinionated content implying persona: While SGE is designed to reflect a neutral, objective tone in its generative output, there may be instances in which the output reflects opinions that exist on the web, giving the impression of the model displaying a persona.

- Duplication or contradiction with existing Search features: Because SGE is integrated into Search alongside other results and features on the search results page, it's possible for the output of SGE to appear to be in contradiction with other information in those results. For example, people might see a featured snippet result that highlights the perspective of a single source, while SGE represents a synthesized perspective corroborated in a range of results.

Building the future of Search together

Part of our approach to rolling out SGE responsibly is setting clear expectations with our users, and being mindful of the limitations that still exist, even as we make progress in quality and safety. By making SGE first available through Search Labs, we're giving people the opportunity to interact with this new technology while also being transparent that this is still in an experimental phase.

While it's still early days for SGE, we're actively improving the experience and committed to gains in quality and capabilities. We look forward to learning from feedback from users in Labs, continuing to improve and iterate and building the future of Search together.

Google