Building Personal Intelligence:

A step towards truly personal AI

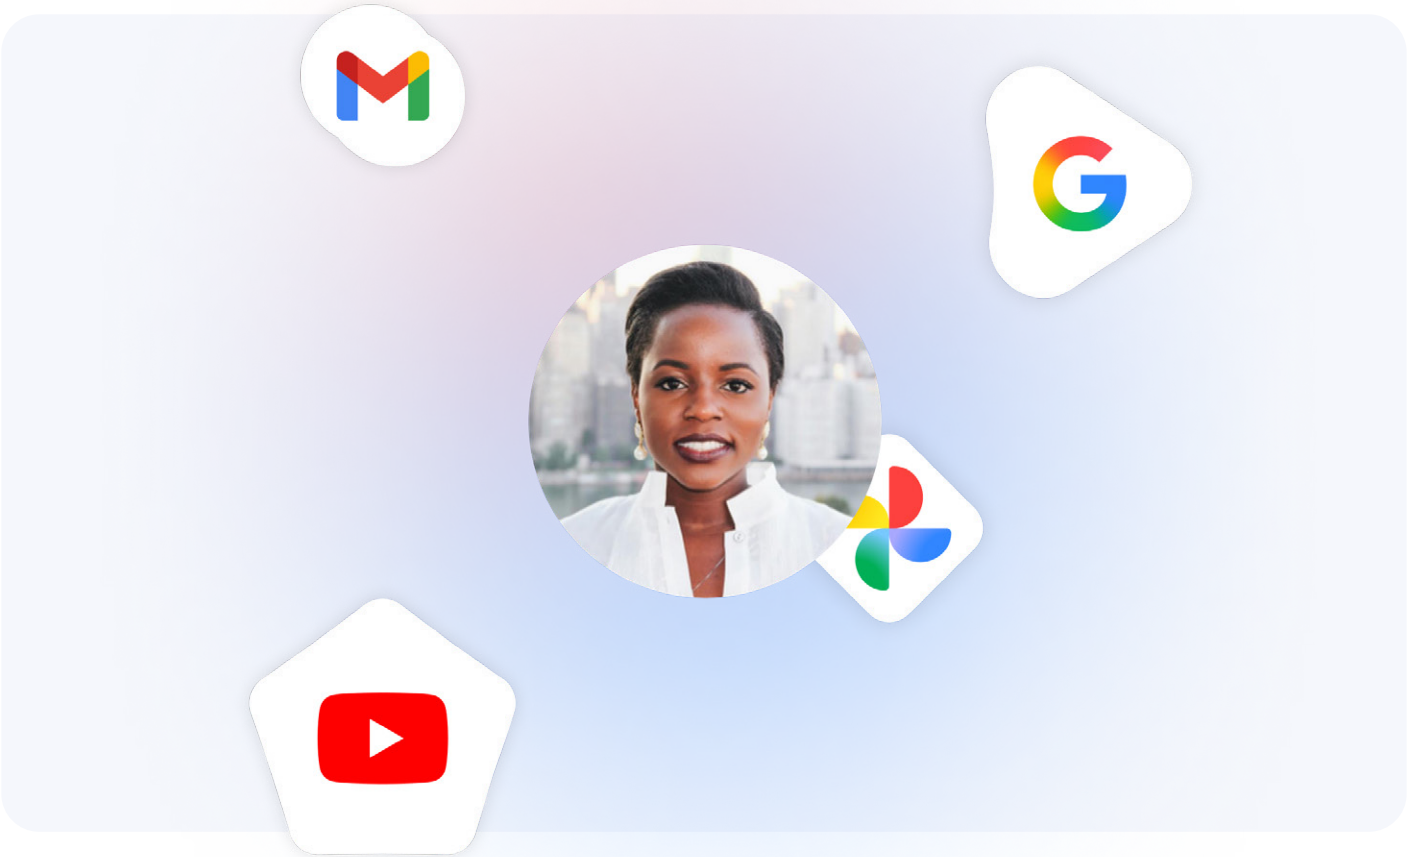

January 2026

Srinivasan (Cheenu) Venkatachary

8 min.

Making our products more helpful by making them personal

At Google, we've spent years refining our products to be not just helpful, but helpful to you.

We learned early on that personal context matters: when someone searches for "running shoes", they aren't usually looking for a list of generic bestsellers, but for the specific brands and styles they prefer. That insight — that your history and preferences should inform your results — enhanced Google Search, and in the years since, we've brought that same level of personalization to many of our products.

Interacting with your personal information has historically meant navigating independent product experiences: searching for your flight reservation required opening Gmail, while finding a memory meant searching or scrolling through Google Photos.

In the last few years, we began bridging these gaps with features that allowed for explicit information retrieval across products with your permission — like tagging @Gmail to find flight details in the Gemini app. We also introduced features that personalize your experience in a specific product — like using your past chats in the Gemini app to give you a more tailored response. But these still don't provide a fully personalized experience.

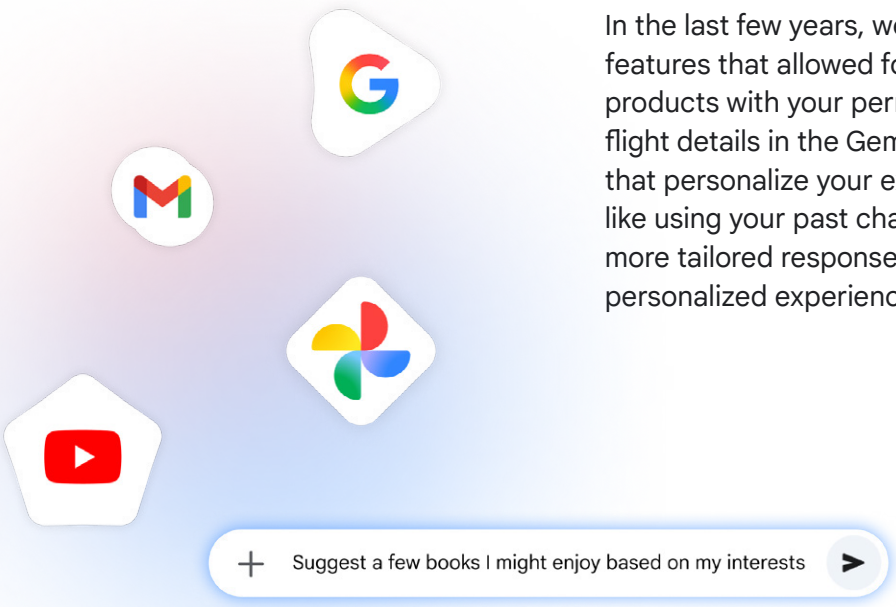

Introducing Personal Intelligence

Personal Intelligence marks the shift to AI that can truly understand your personal context.

It brings together the information you share and the activity within the products you use, while connecting the dots across your Google apps to to make AI uniquely helpful for you.

With Personal Intelligence, you can choose to connect certain Google apps to get more tailored responses — all while your information remains safeguarded.

Today, you can use a beta version of Personal Intelligence in the Gemini app. It securely connects information from apps like Gmail and Google Photos to make Gemini uniquely helpful. Personal Intelligence is coming soon to AI Mode in Search.

That means the Gemini app can help with more tailored discovery and complex planning — like helping you plan for spring break based on places you've already been or proactively finding the perfect set of tires for your specific car make and model.

If you turn this beta feature on, you control which apps to link, and each one supercharges the experience.

This is a foundational step toward moving beyond generic assistance to **AI that works for you**. It's still early days for this technology, and we're continuing to work through known technical issues and limitations. As we continue to test and learn, we are eager to hear your feedback.

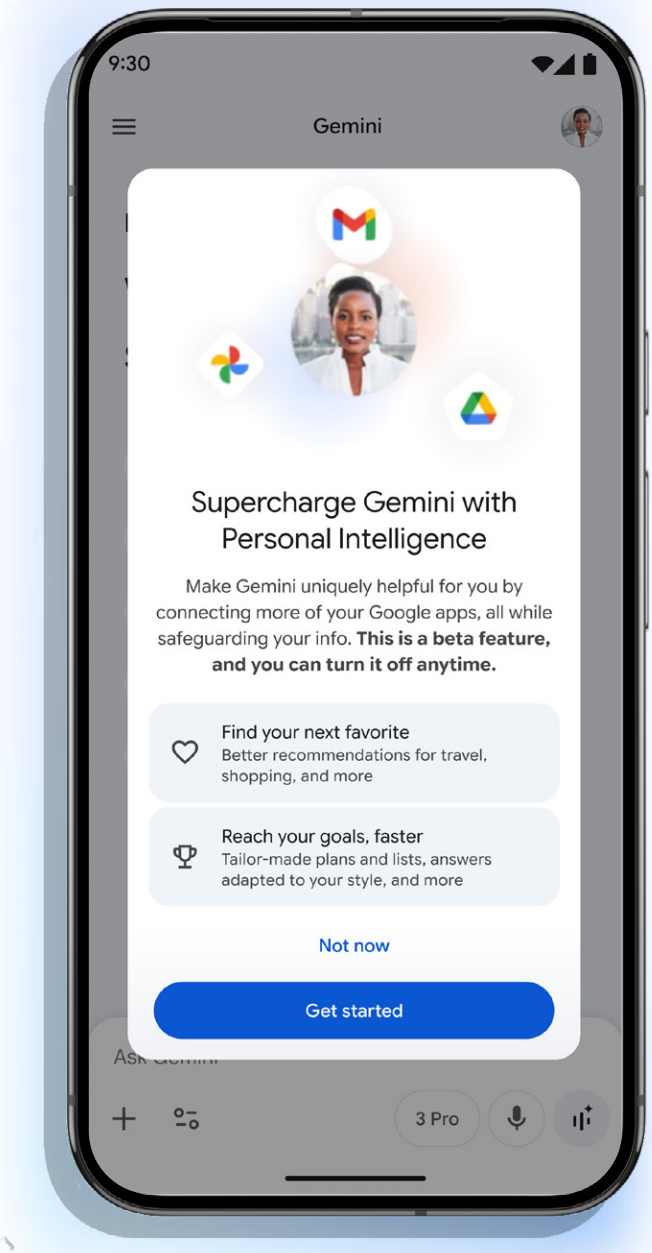

The Technical Challenge: Solving the Context Packing Problem

Personal Intelligence unlocks new utility by solving for the context packing problem: enabling our Gemini models to safely and accurately reason over disparate and vast amounts of personal data sources in real-time without compromising user privacy.

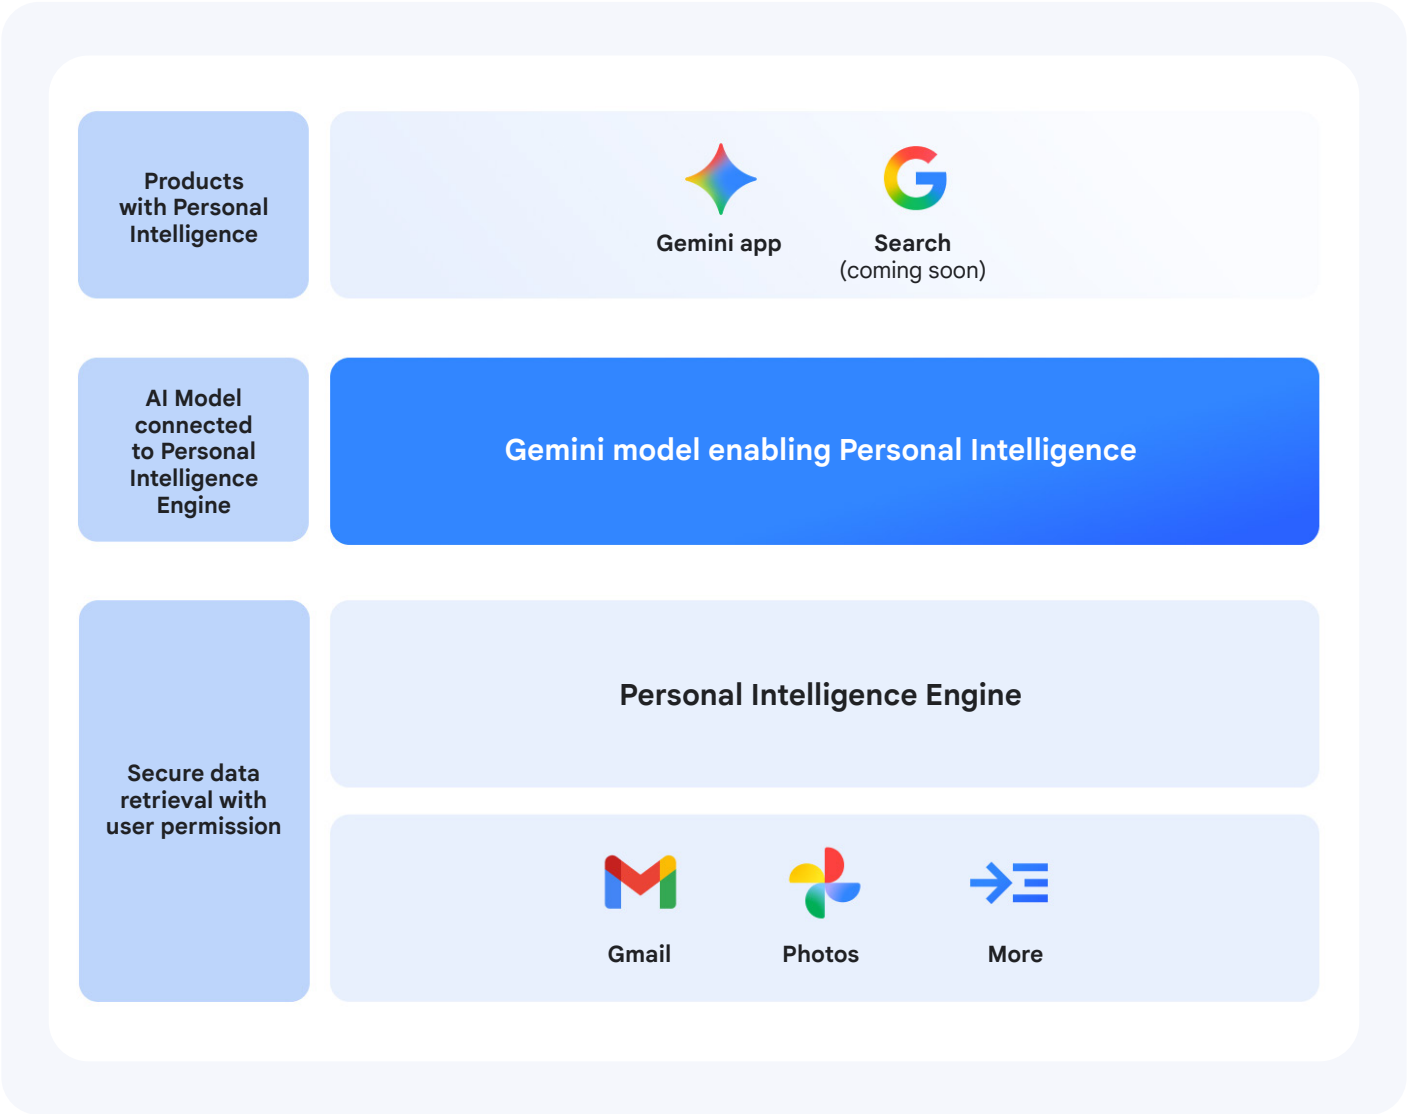

Advanced Gemini Models & A New Personal Intelligence Engine

To handle the complexity of retrieving data from multiple sources simultaneously and provide significantly more helpful responses to you, we built a new engine for Personal Intelligence.

Personal Intelligence has two core strengths: tool calls to retrieve specific details and reasoning across many complex sources. It often combines both approaches and can work across text, photos, and video to give you one-of-a-kind answers.

It enables Gemini models to securely retrieve relevant context, leveraging advances in tool use, dense retrieval, and long context capabilities to reason across your personal data from Google products in real-time.

Advanced Reasoning

Gemini 3, our most intelligent model series to date, is better at general understanding and deciphering more depth and nuance — capabilities that are critical for understanding complex personal context, such as mapping familial relationships or recognizing your specific aesthetic preferences.

Advanced Tool Use

The Gemini model also has significant advances in tool use capabilities which means it can understand your goal and retrieve more information related to your preferences from the Personal Intelligence engine. This retrieval also builds on the foundation of research we've done on search and dense retrieval, such as Gemini Embeddings.

Long Context

Gemini 3 has a 1 million token context window which enables reasoning across a vast amount of data. However, truly helpful personalization requires processing at a much larger scale, as a user's accumulated context across emails and photos alone often exceeds this window by orders of magnitude.

To bridge this gap, we utilize "context packing" — a technique that helps us dynamically identify and synthesize appropriate pieces of information into the working memory for the model.

How Personal Intelligence Works

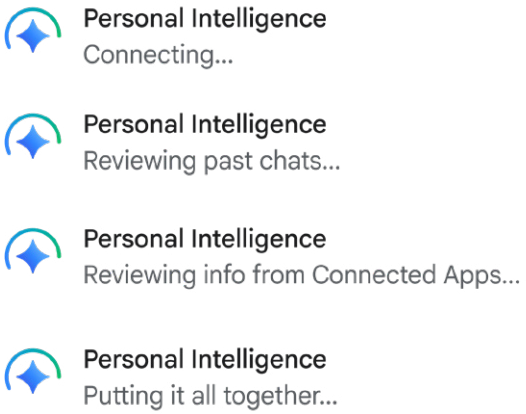

With this new engine and Gemini 3, we now have the capabilities required for true personalization.

For example, when you ask Gemini to "Plan a list of restaurants for my upcoming trip that are near my hotel", with Personal Intelligence, you won't just get generic top-rated spots. The Gemini model understands that this task requires synthesizing disparate personal details from across Google products — your hotel reservations, flight arrival time, past dining history, and the aspirational spots you've saved.

- The model agentically and securely executes searches for the personal information that's relevant to the response, looking for your trip in recent emails, but also other relevant information related to the query, like past restaurant reservations to understand what you love.
- Finally, it delivers a tailored set of personal recommendations close to your upcoming accommodations by making sense of your data like your photos and emails, as well as things like your past Gemini app chat conversations, Search queries and YouTube history.

This fundamentally shifts our architecture and our approach to personalization: we are moving towards a world where, with your permission, products like the Gemini app can securely access certain types of personal information as a continuous stream of context to inform every interaction – delivering real, tailored helpfulness.

It's critical to develop technology like Personal Intelligence responsibly. In accordance with our AI principles, we are focused on building this technology securely, while protecting user privacy and setting guardrails for sensitive topics.

For example, the model aims to avoid making proactive assumptions about sensitive data like your health, though it will discuss this data if you ask.

Privacy & Personal Intelligence

A key area of focus for us when building Personal Intelligence was protecting user privacy and securely connecting data sources across Google apps.

User controls by design

You can choose whether or not to turn these features on or off. In the Gemini app, you can manage all your preferences directly in your settings, like choosing which services — such as Google Workspace, Google Photos, YouTube, and Search — you want to connect as a part of Personal Intelligence in the "Connected Apps" settings. These connected app settings are off by default.

Securely connected data sources

We start by building on our best-in-class security infrastructure and implement additional industry-leading safeguards to ensure that this data remains protected even as it powers new AI experiences. For example, user data is encrypted at rest by default and protected in transit between our systems using Application Layer Transport Security (ALTS). We've also done work to increase resistance to prompt injections and have improved protection against misuse via cyberattacks.

Limited generative AI training

Our goal is to improve your experience while keeping your data secure and under your control. Built with privacy in mind, Gemini Apps don't train directly on your Gmail inbox or Google Photos library. To improve functionality over time, we train on info like prompts and responses in Gemini as well as summaries, excerpts and inferences used to help answer your prompts. To learn more, visit the Help Center.

Responsible Development

Like many emerging AI features, Personal Intelligence is still evolving. This technology may make mistakes — like misinterpreting context or making incorrect assumptions about your activity.

Known Technical Limitations & Issues

You can correct Gemini if it makes mistakes directly via a prompt (e.g. remember that I don't eat meat). We're also working to address known issues through rigorous internal testing and model tuning, but we know new challenges will arise. Your feedback is key to identifying them and making this technology as helpful as possible.

Here are some of the key technical challenges we're working to address:

- **Overpersonalization based on your interests**

  A known challenge is the tendency for the model to rely too heavily on a personalized inference where it's not appropriate — a phenomenon we call "tunnel vision".

  Example: You might be a big fan of coffee shops and the model understands that as part of Personal Intelligence. When you ask it to "plan a trip to Australia", it may inadvertently plan a trip where the itinerary is focused on coffee shops.

  Example: If you have an email in your inbox about your employment, it might start anchoring your responses around the fact that you're a software engineer.

  Example: If you ask "what kind of socks should I buy?", it could assume that because you have a marathon coming up you are only looking for athletic running socks.

- **Mistaking another person's preferences for your own**

  In testing we've also seen a challenge around the model conflating subjects — for instance, attributing a family member's interests to you. When you share a household account for services like YouTube, or do research or make a purchase for a friend or family member, the system may mistake others' preferences as your own.

  Example: Based on a receipt in your email, the model might think you enjoy listening to heavy metal and offer suggestions on concerts nearby, when you actually purchased the tickets as a birthday gift for your brother.

- **Incomplete information**

  There are some instances where you might not see all of your personal information if you ask. All the relevant information may not have been retrieved or our systems might make inaccurate inferences based on the information available.

  Example: If you ask for a summary of last month's activities, we might only have information for a fraction of them.

Known Technical Limitations & Issues

- **Mixing up timelines**

  Temporal relevance is a known issue for AI models generally and adding personal history adds additional complexity and makes it tricky to make relevant connections.

  Example: The model may mix up timing, noting that a graduate program application deadline from your email is in the past, when in fact it's still upcoming.

- **Misinterpreting relationships**

  The model also has challenges understanding and grasping the nuance of relationships and complex dynamics, sometimes misidentifying family roles.

  Example: It can misidentify a mother for a grandmother based on ambiguous text in emails or label a sibling as a friend.

- **Missing major life changes**

  The model won't always know when a major change in your life has occurred, such as a divorce or a death in the family.

  Example: The model might suggest an anniversary dinner reservation for a partner you are no longer with.

- **Incorrect assumptions**

  The model often assumes that a transaction record equals a completed action. It may assume you bought an item or attended an event based on a confirmation email, missing the subsequent context that you returned the item or cancelled the reservation.

  Example: The model could recommend a follow-up book in a series because you bought the first one, failing to realize you returned it the next day.

- **Overlooking corrections**

  If you correct the model about your personal information, it might be missed sometimes. This often happens with more ambiguous prompts.

  Example: You tell the model, "I don't usually eat steak," but it suggests a steakhouse recommendation again a week later even though you generally don't prefer steak.

Other Technical Challenges

- **Balancing speed and depth**

  Delivering true personalization requires more complex processing, so we're constantly navigating the delicate balance between latency and quality tradeoffs.

  To ensure the best user experience, the system distinguishes between general queries without personalization and complex personal requests and we'll continue to iterate on this as our models and technology evolve.

  For general questions without personalization, you'll get a faster response.

  With more personal questions, you may see a "thinking" indicator in the product which may say "Personalizing" or "Personal Intelligence". This visualizes the processing steps while the model securely retrieves and reasons over your personal information to provide a thorough answer.

  In the Gemini app, you can also choose the option "Answer now" to get a faster response if you're using the "Thinking" or "Pro" models.

- **Taking extra precaution with multimodal outputs**

  Making personal context more helpful isn't just about text; it involves going beyond text to provide rich, visual responses. This introduces additional challenges: generating multimodal outputs — like surfacing a specific photo or generating a visual summary — carry unique risks of inaccuracy or potential technical challenges. This is an area we're continuing to work on and test with trusted testers before we release more broadly.

The release of these features is just the first step in a broader research program designed to deliver a universal assistant that is helpful and truly personal.

To achieve this vision, we are continuing to do robust research in areas like securely integrating additional personal data and improving retrieval, long-context usage and model quality for personalization.

As we learn more from the initial release of these products, we look forward to continually improving the products and identifying further research opportunities.

Vision for the future

Personal Intelligence is more than a set of features. It serves as the technical foundation for the future of more personalized AI agents that can be even more helpful to all of us in our daily lives, and it's a significant step on our journey towards AGI.

This is just the beginning of the journey. We will continue to test, listen, and refine Personal Intelligence experiences across our products based on your feedback, so that we can make AI uniquely helpful for you.