



omd



Gemini 3 Pro

Understanding implications and opportunities
from Google's latest model updates

Updated: November 28th, 2025



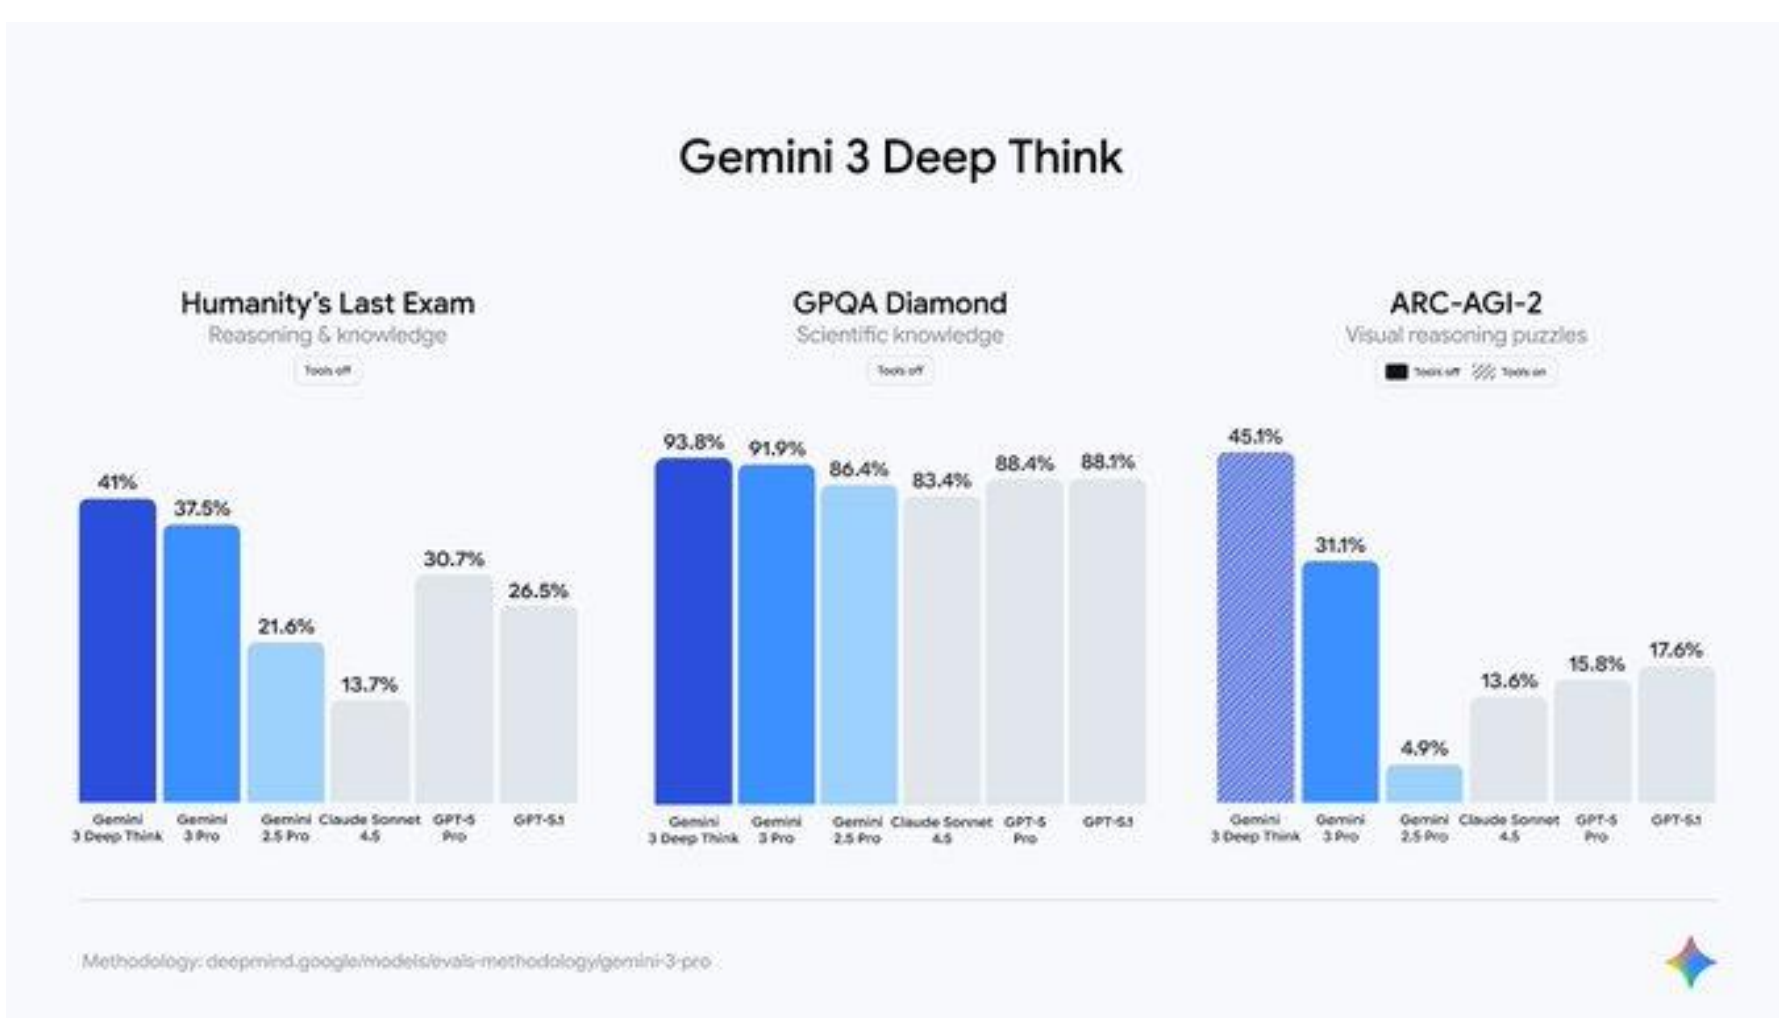
WE *Create* WHAT'S NEXT

Google Launches a Series of Leading-edge AI Models and Utilities

Gemini 3 Pro the New Leading Model

Over the past two weeks, the three frontier AI labs launched new models. OpenAI [launched GPT 5.1](#), and Anthropic launched [Claude 4.5 Opus](#).

The most significant releases came from Google. [Gemini 3 Pro](#), [Nano Banana Pro](#), [Antigravity](#) and [Notebook LM](#) were released or updated in the same week. Together, they deliver a fundamental leap in what AI can comprehend, create, and execute, as well as a [tangible thread](#) that connects Google's product lines.



Gemini 3 Pro is delivering breakthrough capabilities across [multiple benchmarks](#) - cognitive ability shown through [achieving 37% in Humanity's Last Exam](#), a 6-point improvement over GPT-5 Pro's previous record.

Gemini 3 Pro shows strong performance in [agentic scenarios](#) that require the AI to perform complex, multi-step tasks. Previously, the primary interaction mode involved a user prompting an AI for text or an image. Now, the AI can be tasked with completing complex projects that involve planning, using tools, and executing a sequence of actions.

Gemini 3 Pro enables [new apps, games, and user interfaces](#) to be easily created by anyone. We can expect to see an explosion of innovation as non-developers create new interface ideas.

Developers are empowered by AI-assisted Integrated Development Environments, such as [Google Antigravity](#). Capabilities at every part of the media value chain are growing at an accelerated rate.

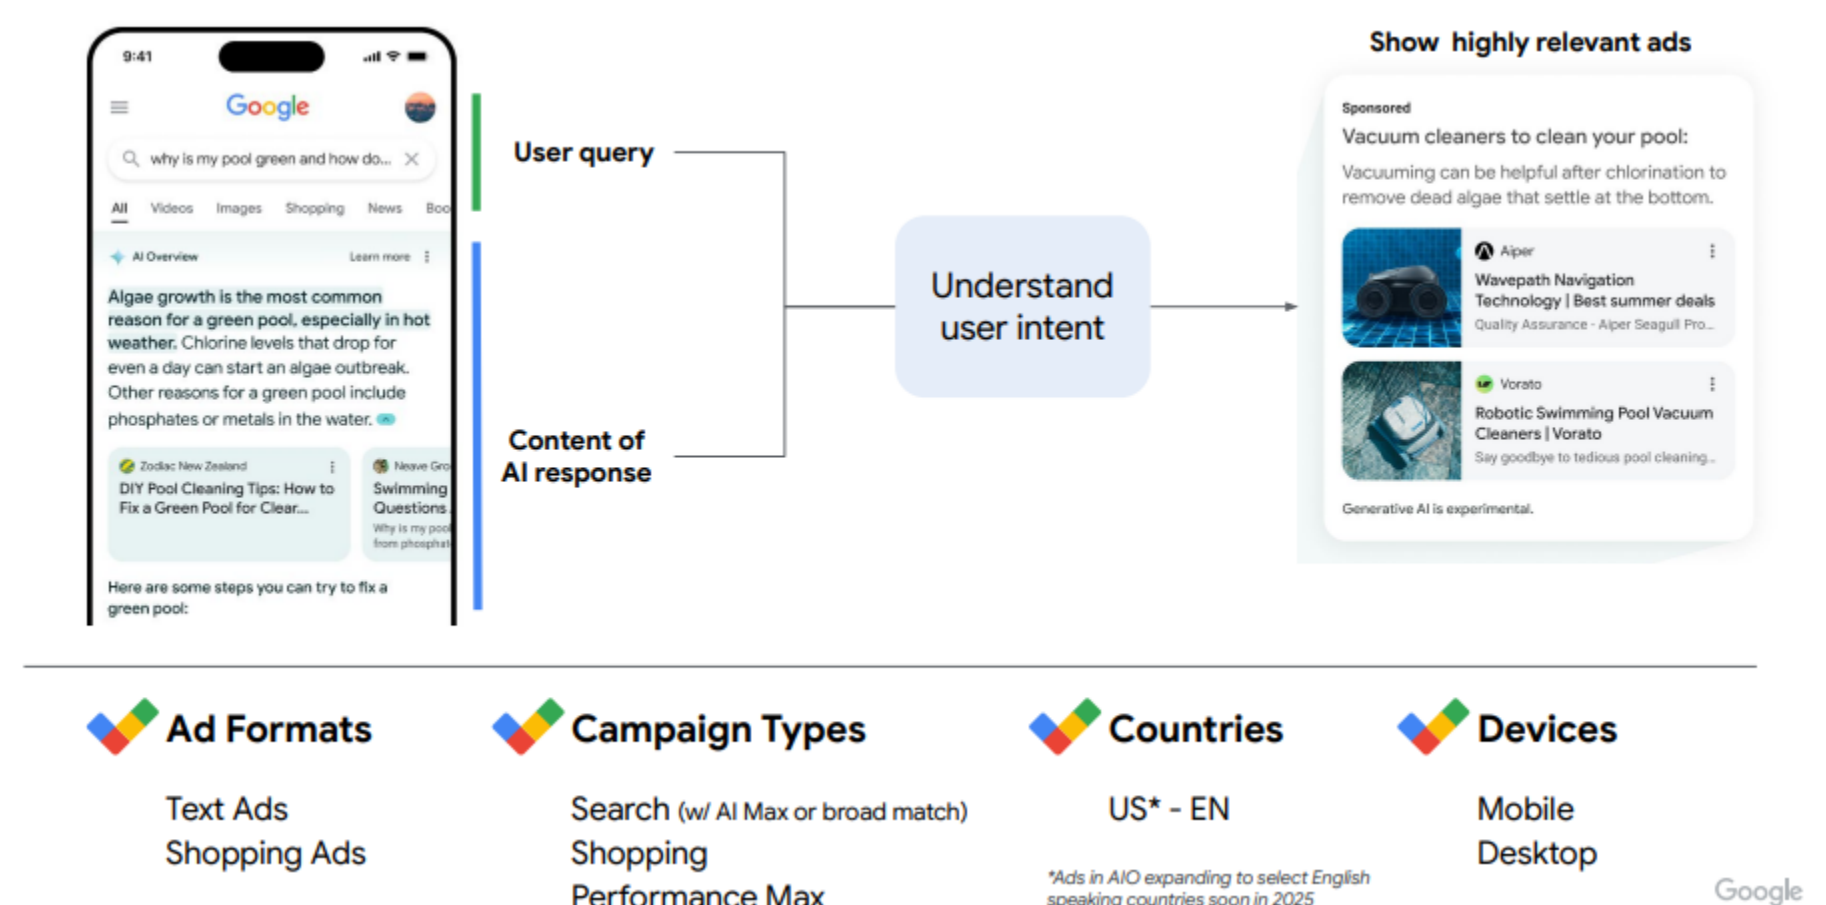
This is the start of a shift from brands defining the touchpoints they have with consumers [to dynamic AI-created touchpoints](#) that are built around specific needs by a wide variety of stakeholders. The role of the brand is increasingly to provide structured data to power those interfaces and to anticipate a wide variety of consumer intents that can be powered by them.

Ad Impact

Unlike many other new model releases, Google chose to release Gemini Pro 3 through both free and paid tiers of the Gemini app. This is in part due to their investment in [proprietary Tensor Processing Units](#), which are significantly more energy-efficient than GPUs and are now a [strategic growth](#) lever.

[Paid US users](#) now see Gemini 3 Pro in [Google AI Mode](#), meaning that consumers don't need to pull an AI interaction - it can be pushed via a search. This will be particularly useful in [complex search contexts](#), such as education or high-value purchase decisions.

Possibly not coincidentally, Google also [announced tests of Ads in AI mode](#) this week. Initially [launching in the US](#), they are currently below the AI Mode response but will be fully integrated within it by the end of the year. This [briefing document](#) details what formats are available for ads in both AI Overviews and AI Mode. [Text, shopping, local, and app ads](#), both above and below AI Overviews, are now available in over 200 countries worldwide.



Ads that are [integrated within AI Overviews](#) are now available in the US. AI Overviews will trigger when the query is complex, and an AI Mode Ad when the intent is marked as commercial.

With [2bn monthly AI Overviews](#) and [75m daily AI Mode](#) users, these ad formats are increasingly important as consumers learn to trust and value AI-generated responses. This drives an increased proportion of zero-click searches, where a significant portion of the purchase journey occurs [in stateful search environments](#). As commerce protocols mature, purchases will also occur in environments away from the brand's site.

Google is pushing for mass consumer adoption of AI services with capabilities and ideas in time for the [holiday season](#), including [holiday shopping](#) and [hosting](#). OpenAI is also focused on the giving season, launching [ChatGPT Shopping Research](#).

Information Distillation

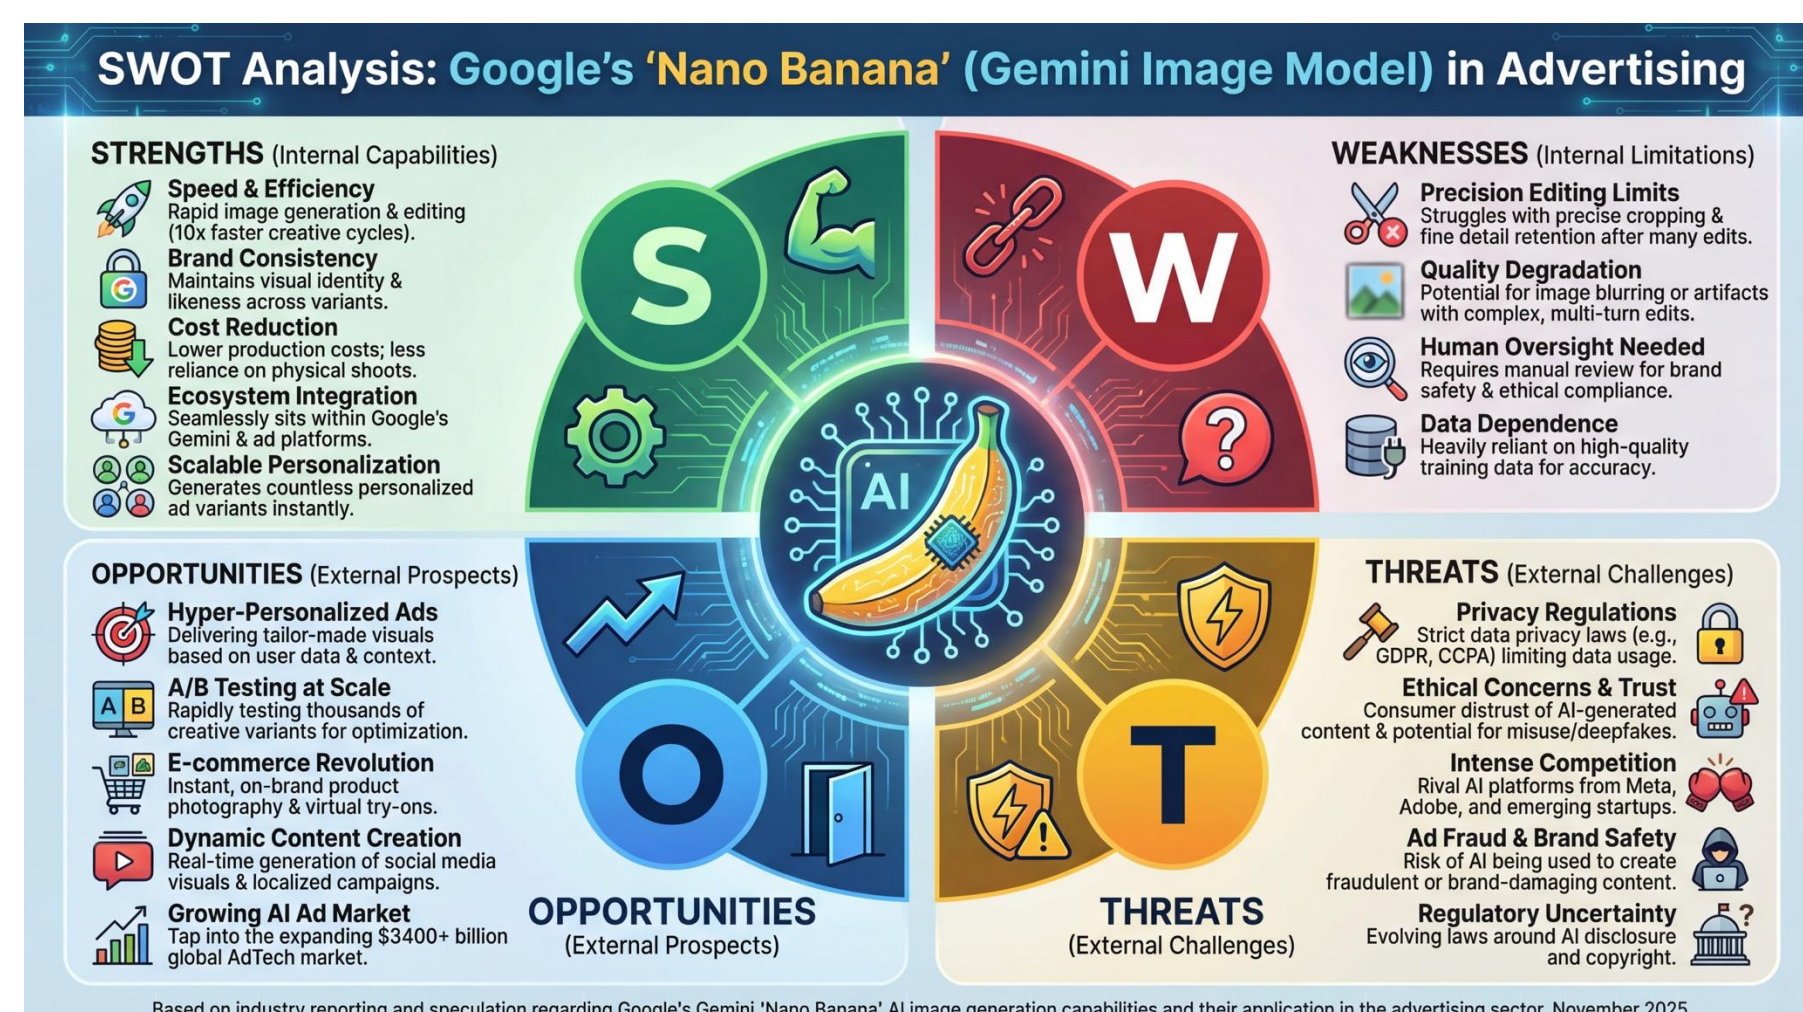
Officially named Gemini 3 Pro Image, [Nano Banana Pro](#) is the more widely used name of Gemini's image capabilities.

Built on the core Gemini 3 Pro architecture, it leverages the underlying model's intelligence for [visual creation and editing](#), and introduces several critical improvements aimed at [professional creative workflows](#), including higher-fidelity outputs up to 4K, precise local editing controls, and superior [text-in-image rendering](#).

The most significant upgrade is the introduction of [visual reasoning](#), meaning Nano Banana Pro can understand intent, perform research and curate information, for example:

- Specific [locations and times](#) in history.
- [3D renders from a blueprint](#).
- A [resume as a middle-earth](#) map.

The critical concept for businesses to understand is the ability to [effectively distill information](#). Entire books can be summarized as images, company reports refined into strategic infographics.



For example, a SWOT of Nano Banana in the advertising sector as a one-shot 4k image ([high-resolution here](#)). Brand maps can be researched, edited, and created quickly; it is a useful way to synthesize social listening data across multiple sentiments and sources.

In the same week as Nano Banana Pro launched, [Google Notebook LM saw significant upgrades](#) with implications for multiple business workflows. As well as the existing podcast and video summaries, [slide decks that synthesize information](#) from many sources can be created in minutes. For example, a [15-slide summary](#) of Emad Mostaque's 180-page e-book '[The Last Economy](#)'.


The challenging themes of AI's impact on employment and government policy can be entirely recontextualized by Gemini 3 Pro. For example, the Storybook Gem creates a [children's story](#) about a girl weaving a tapestry of skills through her life.

New Growth Opportunities

The launch of new models and utilities over recent weeks, each of which already have several hundred million users, highlights the urgency of [AI data imperatives](#).

Success in AI search ads will be based on the same principles of effective Generative Engine Optimization. This involves understanding **consumer** needs to generate trusted **content**, optimized effective **code**, in contexts that deliver **credibility**, for **commerce** outcomes. These 5Cs form the basis of Omnicom's GEO proposition AI Optix, the integration of paid, and GEO techniques will be a key differentiator.

We are starting to see a radically different architecture for the internet and by implication, the consumer economy. The distinct roles of publisher, platform, consumer and brand are dissolving. We need to imagine a [radically different relationship](#) between them.



This concept was explored by OpenAI this week in their guidance on how to [make a great ChatGPT App](#). The key advice is not to rebuild a brand website or app inside ChatGPT, but instead to 'see your app as capabilities the model can orchestrate'.

Key elements such as product availability, customization or price comparison can be packaged as discrete microservices that the model can embed in any consumer journey.

Personalization driven by real-time data, user permissions, and an intent layer will become a competitive advantage that delivers value beyond price.

If anything, AI [capabilities are exceeding predictions](#). Adoption is rapid as the utility is clear and switching costs are low. Just as AI streamlines existing workflows, new opportunities and capability requirements are coming into view quicker than ever.



Jean-Paul Edwards

Managing Director, Product
OMD Worldwide