

August 5, 2025 Models


Genie 3: A new frontier for world models

Jack Parker-Holder and Shlomi Fruchter

Try Project Genie

Share

Learn more



Today we are announcing Genie 3, a general purpose world model that can generate an unprecedented diversity of interactive environments.

Given a text prompt, Genie 3 can generate dynamic worlds that you can navigate in real time at 24 frames per second, retaining consistency for a few minutes at a resolution of 720p.

DeepMind



Towards world simulation

At Google DeepMind, we have been pioneering research in simulated environments for over a decade, from training agents to [master real-time strategy games](#) to developing simulated environments for [open-ended learning](#) and [robotics](#). This work motivated our development of world models, which are AI systems that can use their understanding of the world to simulate aspects of it, enabling agents to predict both how an environment will evolve and how their actions will affect it.

World models are also a key stepping stone on the path to AGI, since they make it possible to train AI agents in an unlimited curriculum of rich simulation environments. Last year we introduced the first foundation world models with [Genie 1](#) and [Genie 2](#), which could generate new environments for agents. We have also continued to push the state of the art in video generation with our models Veo 2 and Veo 3, which exhibit a deep understanding of intuitive physics.

Each of these models marks progress along different capabilities of world simulation. Genie 3 is our first world model to allow interaction in real-time, while also improving consistency and realism compared to Genie 2.


	GameNGen	Genie 2	Veo	Genie 3
Resolution	320p	360p	720p to 4K	720p
Domain	Game-specific	3D Environments	General	General
Control	Game-specific	Limited keyboard / mouse actions	Video-level description*	Navigation; Promptable world events
Interaction Horizon	A few seconds	10-20 seconds	8 seconds	Multiple minutes
Interaction Latency	Real time	Not real time	N/A	Real time

*Additional controls such as references, style and camera are available.

DeepMind



Genie 3 can generate a consistent and interactive world over a longer horizon



- Capabilities
- Embodied agent research
- Limitations
- Responsibility
- Next steps

Genie 3's capabilities include:

The following are recordings of real time interactions from Genie 3.

Modelling physical properties of the world

Experience natural phenomena like water and lighting, and complex environmental interactions.



DeepMind




Prompt: The video shows a first person perspective of someone navigating difficult terrain in the middle of a volcanic area. This is a real worl... +

< >

Simulating the natural world

Generate vibrant ecosystems, from animal behaviors to intricate plant life.



Prompt: Running by the shores of a glacial lake, exploring branching paths through the forest, crossing flowing mountain streams. Set amidst... +

< >

Modelling animation and fiction

Tap into imagination, creating fantastical scenarios and expressive animated characters.

DeepMind



Prompt: A vibrant 3D style, an adorable, fluffy creature bounding across a vibrant rainbow bridge in a fantastical landscape. The creature is s... +

< >

Exploring locations and historical settings

Transcend geographical and temporal boundaries to explore places and past eras.



Prompt: A real world mountainous environment in the Alps. The landscape features steep, rocky cliffs and narrow gorges filled with loose scr... +

< >

Pushing the frontier of real-time capabilities

DeepMind

information from a minute ago. To achieve real-time interactivity, this computation must happen multiple times per second in response to new user inputs as they arrive.




Environmental consistency over a long horizon

In order for AI generated worlds to be immersive, they have to stay physically consistent over long horizons. However, generating an environment auto-regressively is generally a harder technical problem than generating an entire video, since inaccuracies tend to accumulate over time. Despite the challenge, Genie 3 environments remain largely consistent for several minutes, with visual memory extending as far back as one minute ago.



Prompt: POV action camera of a tan house being painted by a first person agent with a paint roller

< >



The trees to the left of the building remain consistent throughout the interaction, even as they go in and out of view.

Genie 3's consistency is an emergent capability. Other methods such as NeRFs and Gaussian Splatting also allow consistent navigable 3D environments, but depend on the provision of an explicit 3D representation. By contrast, worlds generated by Genie 3 are far more dynamic and rich because they're created frame by frame based on the world description and actions by the user.

DeepMind



Prompt: First-person view drone video. High speed flight into and along a narrow canyon in Iceland with a river at the bottom and moss on the rocks, golden hour, realworld

Promptable world events


In addition to navigational inputs, Genie 3 also enables a more expressive form of text-based interaction, which we refer to as *promptable world events*.

Promptable world events make it possible to change the generated world, like altering weather conditions or introducing new objects and characters, enhancing the experience from navigation controls.

This ability also increases the breadth of counterfactual, or “what if” scenarios, that can be used by agents learning from experience to handle unexpected situations.

Choose a world setting. Then, pick an event, and see Genie 3 create it.

DeepMind



What event should happen next?

Fueling embodied agent research

To test the compatibility of Genie 3 created worlds for future agent training, we generated worlds for a recent version of our [SIMA agent](#), our generalist agent for 3D virtual settings. In each world we instructed the agent to pursue a set of distinct goals, which it aims to achieve by sending navigation actions to Genie 3. Like any other environment, Genie 3 is not aware of the agent's goal, instead it simulates the future based on the agent's actions.

Choose a world setting. Then, pick a goal you'd like an agent to achieve and watch how it accomplishes it.


DeepMind

Bakery

Garden

Market


Museum



Goal: Approach the industrial mixer Go to the cooling racks Walk to the glass case

Choose a goal

Since Genie 3 is able to maintain consistency, it is now possible to execute a longer sequence of actions, achieving more complex goals. We expect this technology to play a critical role as we push toward AGI, and agents play a greater role in the world.



DeepMind

Limitations

While Genie 3 pushes the boundaries of what world models can accomplish, it's important to acknowledge its current limitations:

- **Limited action space.** Although promptable world events allow for a wide range of environmental interventions, they are not necessarily performed by the agent itself. The range of actions agents can perform directly is currently constrained.
- **Interaction and simulation of other agents.** Accurately modeling complex interactions between multiple independent agents in shared environments is still an ongoing research challenge.
- **Accurate representation of real-world locations.** Genie 3 is currently unable to simulate real-world locations with perfect geographic accuracy.
- **Text rendering.** Clear and legible text is often only generated when provided in the input world description.
- **Limited interaction duration.** The model can currently support a few minutes of continuous interaction, rather than extended hours.

Responsibility

We believe foundational technologies require a deep commitment to responsibility from the very beginning. The technical innovations in Genie 3, particularly its open-ended and real-time capabilities, introduce new challenges for safety and responsibility. To address these unique risks while aiming to maximize the benefits, we have worked closely with our Responsible Development & Innovation Team.

At Google DeepMind, we're dedicated to developing our best-in-class models in a way that amplifies human creativity, while limiting unintended impacts. As we continue to explore the potential applications for Genie, we are announcing Genie 3 as a limited research preview, providing early access to a small cohort of academics and creators. This approach allows us to gather crucial feedback and interdisciplinary perspectives as we explore this new frontier and continue to build our understanding of risks and their appropriate mitigations. We look forward to working further with the community to develop this technology in a responsible way.

Next steps

We believe Genie 3 is a significant moment for world models, where they will begin to have an impact on many areas of both AI research and generative media. To that end, we're exploring how we can make Genie 3 available to additional testers in the future.

Genie 3 could create new opportunities for education and training, helping students learn and experts gain experience. Not only can it provide a vast space to train agents like robots and autonomous systems, Genie 3 can also make it possible to evaluate agents' performance, and explore their weaknesses.

At every step, we're exploring the implications of our work and developing it for the benefit of humanity, safely and responsibly.

Please cite using the following BibTeX

[Download BibTeX](#)

Acknowledgments

Genie 3 was made possible due to key research and engineering contributions from Phil Ball, Jakob Bauer, Frank Belletti, Bethanie Brownfield, Ariel Ephrat, Shlomi Fruchter, Agrim Gupta, Kristian Holsheimer, Aleks Holynski, Jiri Hron, Christos Kaplanis, Marjorie Limont, Matt McGill, Yanko Oliveira, Jack Parker-Holder, Frank Perbet, Guy Scully, Jeremy Shar, Stephen Spencer, Omer Tov, Ruben Villegas, Emma Wang and Jessica Yung.

We thank Andrew Audibert, Cip Baetu, Jordi Berbel, David Bridson, Jake Bruce, Gavin Buttimore, Sarah Chakera, Bilva Chandra, Paul Collins, Alex Cullum, Bogdan Damoc, Vibha Dasagi, Maxime Gazeau, Charles Gbadamosi, Shan Han, Woohyun

DeepMind

Dimple Vijaykumar, Luyu Wang, Piers Wingfield, Nat Wong, Keyang Xu, Christopher Yew, Nick Young and Vadim Zubov for their invaluable partnership in developing and refining key components of this project.

Thanks to Tim Rocktäschel, Satinder Singh, Adrian Bolton, Inbar Mosseri, Aäron van den Oord, Douglas Eck, Dumitru Erhan, Raia Hadsell, Zoubin Ghahramani, Koray Kavukcuoglu and Demis Hassabis for their insightful guidance and support throughout the research process.

Feature video was produced by Suz Chambers, Matthew Carey, Alex Chen, Andrew Rhee, JR Schmidt, Scotch Johnson, Heysu Oh, Kaloyan Kolev, Arden Schager, Sam Lawton, Hana Tanimura, Zach Velasco, Ben Wiley, and Dev Valladares. Including samples generated by Signe Nørly, Eleni Shaw, Andeep Toor, Gregory Shaw, and Irina Blok.

We thank Frederic Besse, Tim Harley and the rest of the SIMA team for access to a recent version of their agent.

Finally, we extend our gratitude to Mohammad Babaeizadeh, Gabe Barth-Maron, Parker Beak, Jenny Brennan, Tim Brooks, Max Cant, Harris Chan, Jeff Clune, Kaspar Dagaard, Dumitru Erhan, Ashley Feden, Simon Green, Nik Hemmings, Michael Huber, Jony Hudson, Dirichi Ike-Njoku, Hernan Moraldo, Bonnie Li, Simon Osindero, Georg Ostrovski, Ryan Poplin, Alex Rizkowsky, Giles Ruscoe, Ana Salazar, Guy Simmons, Jeff Stanway, Metin Toksoz-Exley, Xinchun Yan, Petko Yotov, Mingda Zhang and Martin Zlocha for their insights and support.

Related posts

Genie 3

[Learn more >](#)


DeepMind



How to create effective prompts with Genie

[Learn more >](#)

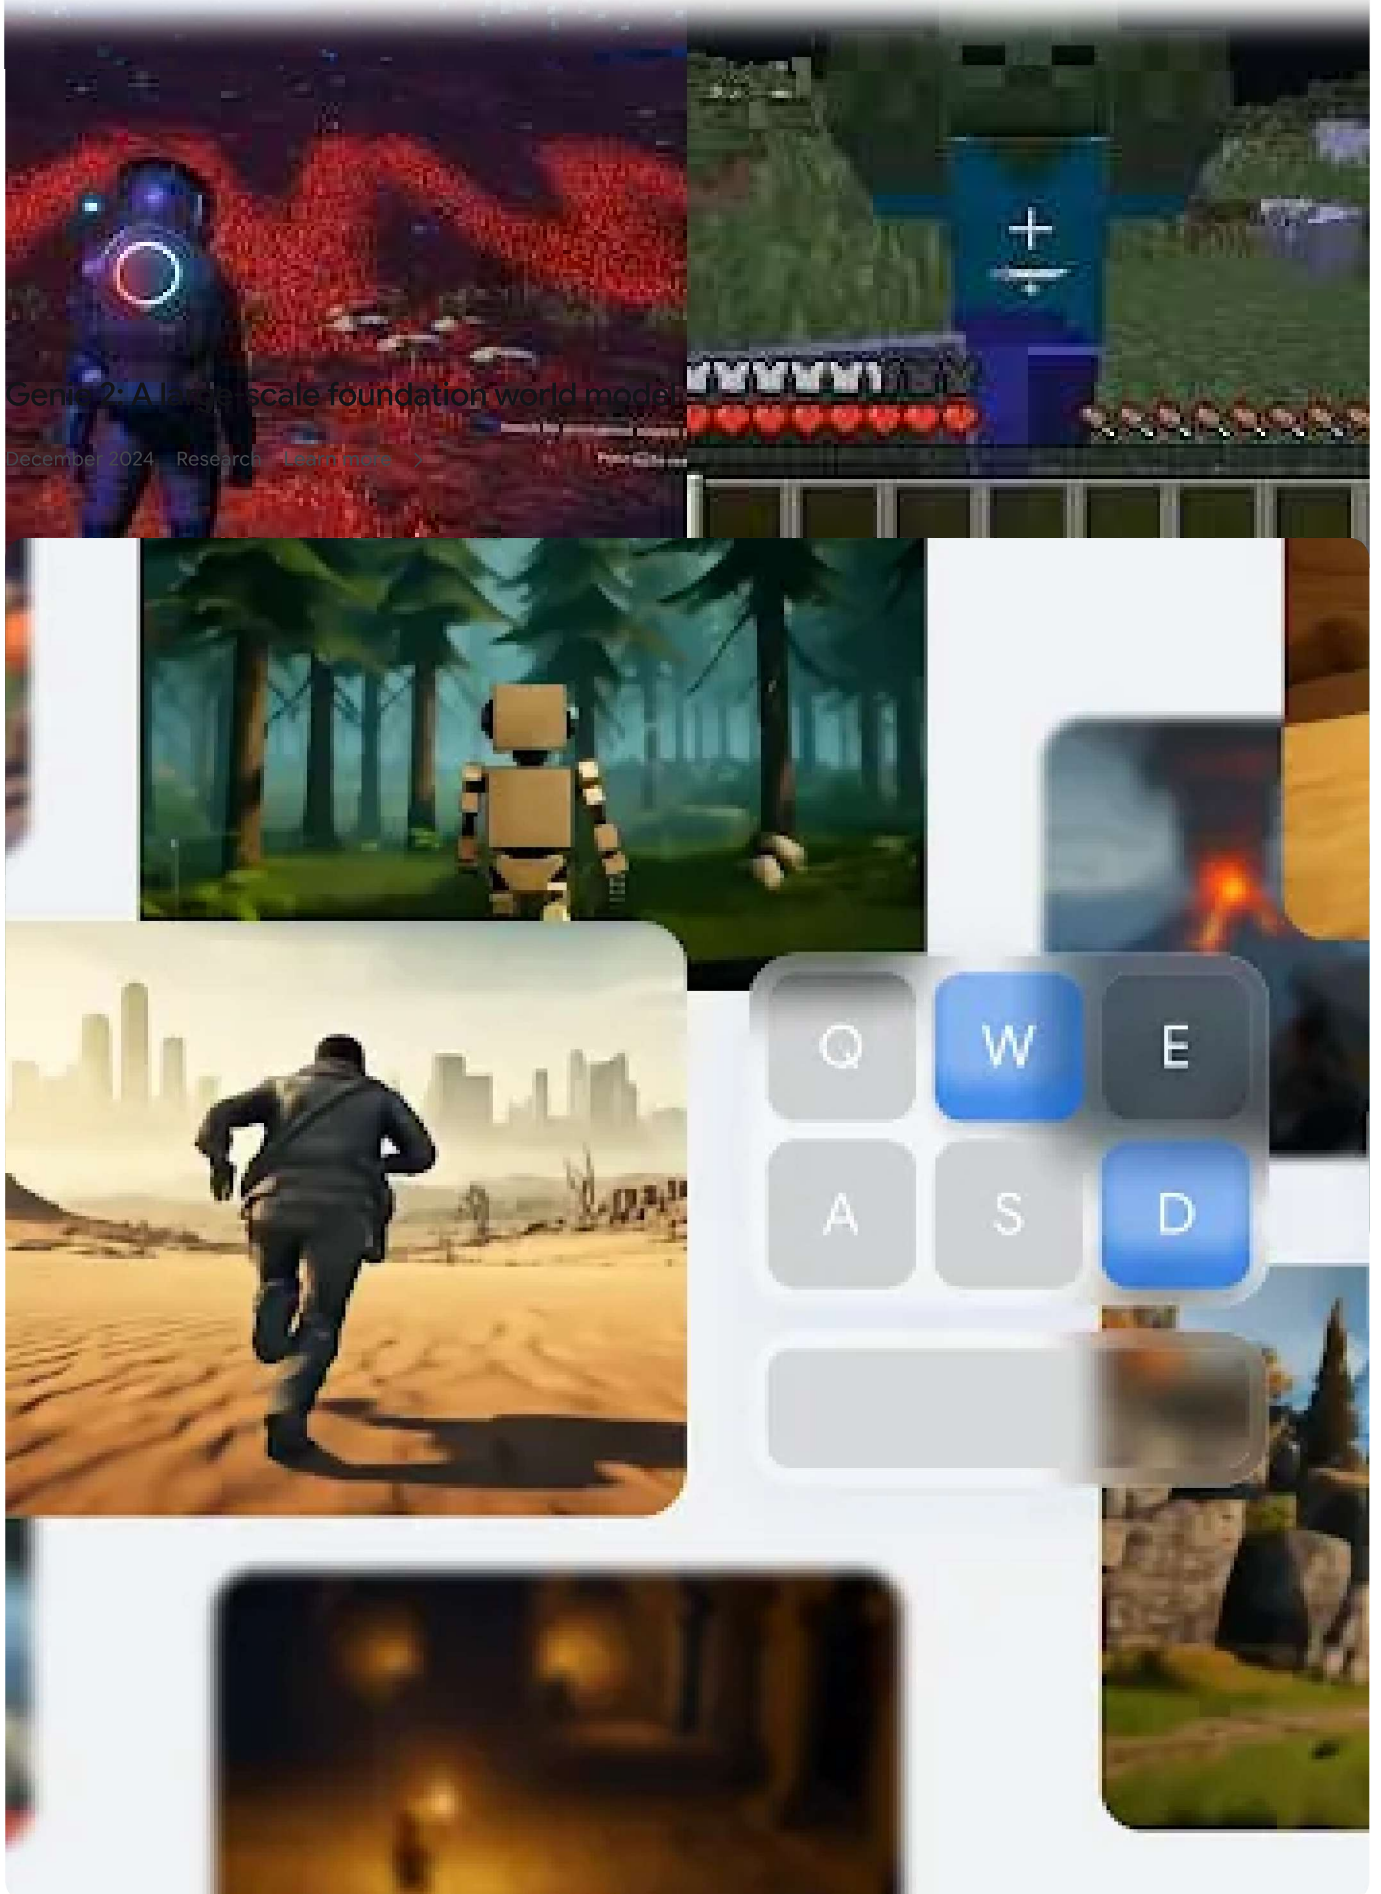
DeepMind



SIMA 2: An Agent that Plays, Reasons, and Learns With You in Virtual 3D Worlds

November 2025 Research [Learn more >](#)

DeepMind



DeepMind



Follow us



Sign up for updates on our latest innovations

I accept Google's Terms and Conditions and acknowledge that my information will be used in accordance with [Google's Privacy Policy](#).

Get the latest updates

Sign up

Build AI responsibly to benefit humanity

Models

Research

DeepMind

↳ Gemini Audio

◇ Gemma

🌐 Genie

📺 Lyria

📺 Veo

Responsibility

Science

🕒 AlphaFold

🕒 AlphaGenome

📷 WeatherNext

🌍 AlphaEarth

Products

◆ Gemini app

<> Google AI Studio

⚖️ Google Antigravity

Learn more

About

News

Careers

National Partnerships for AI

The Podcast

Google

About Google

Google products

Privacy

Terms

