Google

# Genie 4
# Model card

Model Cards are intended to provide essential information on Gemini models, including known limitations, mitigation approaches, and safety performance. Model cards may be updated from time to time; for example, to include updated evaluations as the model is improved or revised. See the Google DeepMind site for a comprehensive list of model cards.

Published: November 2026

---

# Model Information

## Description

Genie 4 is a world generation system, which unlocks unprecedented realism and improves real-time interactive simulation. Genie 4 is multimodal, enabling audio — such as music, ambient noises, and interactive dialogue — to be simultaneously generated with the world.

Genie 4 incorporates and builds on capabilities from the Gemini family: Gemini Omni anchors the simulations in real-world appearance and behavior, while Gemini provides the complex reasoning needed to help maintain a simulation's coherence and continuity.

## Model dependencies

Genie 4 is built on top of Gemini Omni Flash and uses Gemini for reasoning and real-time steering.

## Inputs

Text strings (e.g., a prompt), images, and real-time user actions (e.g., keyboard and mouse navigation).

## Outputs

Real-time 720p video at up to 24 frames per second, with synchronized audio.

## Architecture

Genie 4 is a transformer-based system with native multimodal support for text, user action and image inputs for real-time interactive generation.

---

# Model Data

## Training Dataset

Genie 4 is trained on video, audio, image, and text data.

## Training Data Processing

Data filtering and preprocessing included techniques such as deduplication, safety filtering in-line with Google's commitment to advancing AI safely and responsibly, and quality filtering to mitigate risks and improve training data reliability and compliance. Once data is collected, it is cleaned and preprocessed to make it suitable for training.

---

# Implementation and Sustainability

## Hardware

Genie 4 is adapted from Gemini Omni Flash. For more information about the hardware for Genie 4, see the Gemini Omni Flash model card.

## Software

Genie 4 is adapted from Gemini Omni Flash. For more information about the software for Genie 4, see the Gemini Omni Flash model card.

---

# Distribution

Genie 4 is planned to be distributed via Project Genie.

---

# Evaluation

## Approach

Genie 4 was evaluated across a range of benchmarks, including human and automated evaluations, comparing performance against earlier world models including Genie 3. We focus on quality dimensions including visual quality and aesthetics, environmental realism, audio quality and synchronization, and responsiveness to real-time user control and steering via text prompts.

## Results

Genie 4 improved significantly compared to Genie 3 on photorealism, fluid movement, and real-time steering, while introducing native audio generation and interactive dialogue that adapts on the fly.

---

# Intended Usage and Limitations

## Benefit and Intended Usage

Understanding and simulating the world is a crucial aspect of building Artificial General Intelligence (AGI). Genie 4 advances our ability to create safe, infinitely variable simulations that can help with everything from developing training environments for robots and agents to unlocking the creation of new interactive experiences.

## Known Limitations

Genie 4 is part of an early, evolving area of research.

While it can maintain coherent environments over short sessions, visual details and object permanence can still shift or drift over time—especially as a simulation runs longer, after multiple complex prompt edits, or when navigating throughout the scene. Complex physical interactions, real-time audio accuracy, rendering legible text, spatial accuracy, and long-term consistency are active areas of ongoing research. We are sharing Genie 4 as a preview to learn alongside our community and safely advance the frontier of interactive simulation.

## Acceptable Usage

Genie 4 is adapted from Gemini Omni Flash. For more information about the acceptable usage for Genie 4, see the Gemini Omni Flash model card.

---

# Ethics and Content Safety

## Evaluation Approach

Genie 4 was developed in partnership with internal safety, security, and responsibility teams. A range of evaluations and red teaming activities were conducted to help improve the model safety and inform decision-making. These evaluations and activities align with Google's AI Principles and responsible AI approach, as well as Google's Generative AI policies (e.g. Gen AI Prohibited Use Policy and the Gemini API Additional Terms of Service). Evaluation types included but were not limited to:

1. **Training/Development Evaluations** including automated and human evaluations carried out continuously throughout and after the model's training, to monitor its progress and performance;
2. **Human Red Teaming** conducted by specialist teams who sit outside of the model development team, across the policies and desiderata, deliberately trying to spot weaknesses and ensure the model adheres to safety policies and desired outcomes;
3. **Ethics & Safety Reviews** conducted ahead of the model's release.

In addition to the evaluations above, system-level safety evaluations and reviews are run within the context of specific applications that models are deployed within.

## Frontier Safety Assessment

Genie 4 is part of the Gemini 3 series of models. To assess Genie 4 for frontier safety, we rely on our evaluations of Gemini 3.1 Pro and Gemini 3.7 Flash as outlined in our Frontier Safety Framework. We found that Gemini 3.1 Pro and Gemini 3.7 Flash did not reach any Tracked or Critical Capability Levels (T/CCLs). Our assessments have shown that Genie 4 does not have meaningful new capabilities or material increases in performance compared to Gemini 3.1 Pro or Gemini 3.7 Flash, therefore based on these results, we are confident that Genie 4 is also unlikely to reach any T/CCLs.

For more information on our Frontier Safety Assessment, read the Gemini 3.1 Pro Model Card and Gemini 3.7 Flash Model Card.

## Risks and Mitigations

Interactive world models can amplify creative possibilities, generate world simulations, and create immersive learning experiences, and can accelerate research in robotics, embodied agents and computer vision by providing controllable, realistic environments. Such capabilities require a proactive approach to safety.

We focused on two main content safety areas: (i) intentional adversarial misuse of the model; and (ii) unintentional model failure modes through benign use.

Safety and responsibility were built into Genie 4 throughout the training and deployment lifecycle, including pre-training, post-training, and product-level mitigations.

Mitigations include, but are not limited to:

- dataset filtering;
- conditional pre-training;
- supervised fine-tuning;
- safety policies and desiderata;
- product-level mitigations such as safety filtering and applying SynthID watermarking.