



Gemini Robotics

ER 2

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: July 2026

Model Information

Description

Gemini Robotics ER (Embodied Reasoning) 2 is a Vision-Language-Model that enhances Gemini's spatial, temporal, and physical reasoning capabilities.

Model dependencies

Gemini Robotics ER (Embodied Reasoning) 2 is based on Gemini 3.5 Flash.

Inputs

Text / Image / Video / Audio (interleaving possible). Context window of up to 128k.

Outputs

Text, with up to 64K token output.

Architecture

Gemini Robotics ER 2 is based on the Gemini 3.5 Flash. For more information about the model architecture, see the Gemini 3.5 Flash [model card](#).

Model Data

Training Dataset

Gemini Robotics ER 2 was trained on Gemini 3.5 training datasets and additional datasets representing various embodied reasoning tasks. For more information about the training data, see the Gemini 3.5 Flash [model card](#).

Training Data Processing

For more information about the training data processing, see the Gemini 3.5 Flash [model card](#).

Implementation and Sustainability

Hardware

Gemini Robotics ER 2 was trained using Google's Tensor Processing Units (TPUs). See details in the Gemini 3.5 Flash model card.

The efficiencies gained through the use of TPUs are aligned with Google's [commitment to operate sustainably](#).

Software

Software JAX (Bradbury et al., 2018), ML Pathways (Dean, 2021).

Distribution

Gemini Robotics ER 2 is distributed in the following channels; respective links shared in line:

- [Google AI Studio](#)
- [Gemini API](#)
- [Gemini Enterprise Agent Platform \(private preview\)](#)

Our models are available to downstream providers via an application program interface (API) and subject to relevant terms of use. There is no required hardware or software to use the model. For AI Studio and Gemini API, see the [Gemini API Additional Terms of Service](#); for Gemini Enterprise Agent Platform, see [Google Cloud Platform Terms of Service](#). For more information, see [Gemini Model API instructions](#) and [Gemini API quickstart](#).

Evaluation

Approach

See Figure 1-4 in the [release post](#) for Gemini Robotics ER 2 evaluation procedures and Figure 5 for Gemini Robotics Safety evaluation procedures.

Results

See Figure 1-4 in the [release post](#) for Gemini Robotics ER 2 evaluation results and Figure 5 for Gemini Robotics Safety evaluation results.

Intended Usage and Limitations

Benefit and Intended Usage

Gemini Robotics ER 2 represents an advancement to our reasoning-first Gemini Robotics ER models that enables robots to understand their environments with precision. This model specializes in reasoning critical for robotics, including visual, spatial, and temporal understanding, tool orchestration and success detection for physical agents.

Known Limitations

For more information about the known limitations for Gemini Robotics ER 2, see the Gemini 3.5 Flash [model card](#).

Acceptable Usage

For more information about the acceptable usage for Gemini Robotics ER 2, see the Gemini 3.5 Flash [model card](#). Additionally, users are required to use discretion before using the Robotics Models in a production, commercial, or public environments, and to not use the Robotics Models for safety-critical applications or work, such as in the following settings: (i) healthcare, (ii)

transportation, or (iii) other areas where safety protocols are vital, and a malfunction could reasonably foreseeably lead to death, personal injury, or property damage.

Ethics and Safety

Ethical Considerations & Risks

Previous impact assessment and risk analysis work as discussed in ([Gemini Robotics 1.5](#)) and references therein remain relevant to Gemini Robotics ER 2. See Section 6 for information on responsible development and safety mitigations.

Evaluation

Gemini Robotics ER 2 performs similarly to [Gemini 3.5 Flash](#) in terms of compliance with Gemini safety policies which prevent our AI models from generating specific types of harmful content. It achieves significant gains on Safety Instruction Following, Human Safety monitoring and other [semantic safety](#) benchmarks, see [the safety section](#) in the release post. For new agentic safety and uncertainty resolution evaluations, please see our [safety technical report](#).