

How developers build AI for good with Gemma 4

24.8.2026 - Glenn Cameron | Google Italia

We're announcing the winners of the Gemma 4 Good Challenge, which asked developers to build AI solutions to challenges facing people around the world.

Making the most of AI models requires brilliant engineering. That's why today we're excited to announce the winners of the Gemma 4 Good Challenge, a Kaggle competition that asked participants to bring impactful AI solutions to the world.

While Gemma 4 models allow developers to innovate without limits, deploying them in resource-constrained environments represents a technical challenge. Developers used key technologies, including LiteRT, Cactus, Ollama, llama.cpp, and Unsloth, to create highly performant solutions for everyday hardware. Here are the winners:

First Place

GEM-4 is a physical robotic assistant and control pipeline built to physically assist elderly and disabled individuals with daily living tasks. Using a Gemma 4 31B model to label video training clips, a fine-tuned, lightweight Gemma 4 E2B controller translates visual observations and language instructions into real-world physical movements. Judges lauded GEM-4 for its innovative closed-loop data engine and Vision-Language-Action (VLA) architecture, creating an embodied system that moves AI reasoning directly into hands-on physical help.

Second Place

Trido enabled a voice-controlled digital whiteboard built to assist educators, particularly those with physical disabilities, in managing interactive classroom instruction. Using real-time speech processing and dual-path audio input, a locally hosted Gemma 4 E2B model dynamically generates mind maps, quizzes, and visual widgets on canvas. Judges praised its "deep user empathy and robust offline fallback loop", creating a collaborative educational dashboard that can run on budget laptops.

Third Place

PenguinAgent is an offline-first research application and data analysis dashboard built to support wildlife ecologists conducting field expeditions in remote environments. Using Segment Anything (SAM 3) and SigLip 2 tracking inputs, a locally hosted Gemma 4 26B model ingests sequential video frames and kinematic telemetry to analyze thermodynamic huddling behavior and query local ornithological papers. Judges commended its thorough technical implementation and RAG-grounded scientific reporting, creating a local-first analysis tool that ensures complete scientific data privacy at the edge of the world.

Fourth Place

DEMENTOR is an ambient edge companion and hardware sensor network built to support individuals living with dementia. Using offline camera feeds and haptic feedback, Gemma 4 E2B helps caregivers monitor behavioral patterns and assists with daily memory retrieval. Judges highlighted its low-level hardware integration and edge orchestration, creating a system that

prioritizes user safety and absolute data privacy for sensitive health metrics.

Global Resilience Prize: Acuífero·Vigía

Acuífero·Vigía is an offline-first flood early warning system deployed in South America. The system fuses local telemetry, volunteer audio descriptions, and river camera feeds to generate real-time flood predictions and localized safety instructions. Judges appreciated the project's technical depth, specifically highlighting auditable reasoning trails for accountability and the decision to use Gemma 4 E2B to prevent WebGPU buffer overflows on Raspberry Pi 5 units.

AI Trust & Safety Prize: TrueVoice

TrueVoice tackles the rise of voice-cloning fraud, where bad actors can clone a loved one's voice to mislead unsuspecting victims. The project leverages Gemma 4 E4B's audio capabilities to identify voice cloning scams by analyzing vocal anomalies, emotional tones, and micro-timbres. Evaluated as a solution to a high-stakes problem, TrueVoice addresses a critical and growing threat in AI-driven fraud by deploying detection models directly to the edge, building a critical safety layer that actively intercepts scams before they can cause emotional or financial harm.

Digital Equity & Inclusion Prize: Gem-Care

Gem-Care focuses on resolving the digital exclusion of millions of people with non-normative speech in an increasingly voice-activated world. While modern Automatic Speech Recognition systems achieve high accuracy for standard speech patterns, they frequently fail for individuals with dysarthria or speech disabilities. Gem-Care fine-tuned Gemma 4 E2B to perform multimodal clinical-context adaptation, translating and reconstructing non-normative speech inputs. Judges celebrated the achievement of a Word-Error-Rate (WER) of 19.0% compared to the base model's 32.7% WER.

Future of Education Prize: CodeBuddy

CodeBuddy helps bridge the digital divide for 28 million students in Indonesia by providing Socratic programming tutoring. Since many schools lack internet-connected computers, students write Python code in physical paper notebooks, snap a picture, and let Gemma 4 E4B transcribe, compile, and debug their code entirely offline. The judges highlighted the teacher dashboard and praised the project as a "masterclass in offline edge AI design."

Health & Sciences Prize: Project Rosie

Project Rosie developed an agentic oncology assistant that democratizes access to personalized neoantigen mRNA cancer vaccines for canine oncology. It uses the Gemma 4 26B model to guide veterinarians through the clinical interpretation and design of personalized neoantigen mRNA cancer vaccines for dogs, while using deterministic templates for generating critical manufacturing specifications—an architectural decision praised by judges for its disciplined, non-hallucinatory design.

LiteRT Prize: Gilbeot (🇮🇩)

Gilbeot is an on-device walking navigation assistant for elderly users. Recognizing that seniors struggle with traditional spatial mapping, Gilbeot uses local visual analysis to translate machine coordinates into human descriptions, such as turning left past a specific landmark sign. The project demonstrates exceptional performance tuning on Android devices, deploying Gemma 4 E2B via Google AI Edge's LiteRT-LM framework to achieve fast multimodal loops without cellular service,

which judges noted as a highly novel on-device walking guidance architecture.

llama.cpp Prize: PreVillage

PreVillage is a voice-first WhatsApp navigator that helps rural citizens in Nepal access government services. Using Gemma 4 E4B, it bypasses complex online government portals by using Romanized-Nepali ASR and a local RAG database to map out the undocumented routes, counters, and forms inside administrative offices. The judges hailed PreVillage as a “masterclass in user empathy and robust AI engineering,” praising its use of a self-healing RAG loop and llama.cpp execution on a Raspberry Pi 5 to serve responsive conversational routing at 7.5 tokens per second.

Ollama Prize: PathOS

PathOS is an offline histopathology screening assistant designed to bring same-day, zero-connectivity cancer detection to under-resourced rural clinics. Powered by a custom-trained Gemma 4 E2B model optimized for low-resource CPUs, the pipeline was praised by judges for its advanced reinforcement learning (GRPO) alignment to prevent clinical hedging and its complete on-device execution pipeline.

Unsloth Prize: MindBridge

MindBridge is a mental health screening tool designed to help community health workers conduct vital psychiatric evaluations in regional Hindi dialects. Built by fine-tuning Gemma 4 E2B’s native audio layers on a bespoke conversational dataset, the system was lauded by judges for its robust, TestFlight-ready mobile build, detailed on-device benchmarks, and authentic local-idiom translations.

Cactus Prize: KawanIsyarat

KawanIsyarat is an offline two-way communication assistant designed to connect deaf and hard of hearing BISINDO (Indonesian Sign Language) users. Deployed on mid-range Android smartphones, the system uses Cactus to run the Gemma 4 E2B model locally, translating camera-captured sign landmarks into natural sentences while simplifying spoken audio. The judges praised its memory-aware model routing—which dynamically loads a Whisper model for speech-to-text and unloads it to preserve VRAM for Gemma—and latency optimizations that disabled thinking mode to make real-time performance viable.

Congratulations to all participants

More than 1,600 entries were submitted to the Gemma 4 Good Challenge over the course of 6 weeks. That is over 1,600 projects aimed at helping people, making it a win for the global community. These submissions demonstrate a vision for AI that is private-by-design and practical for the real world.

As Gemma celebrates 1 billion downloads, we extend a massive thank you and congratulations to all the participants who contributed code, datasets, and models to the Gemmaverse. We can’t wait to see how these solutions mature and bring a positive impact for the world.

<https://blog.google/innovation-and-ai/technology/developers-tools/winning-entries-gemma-4-good-challenge>