Sentinel Vision Navigator

Giving eyes to blind and partially sighted people through live camera understanding, fluent voice guidance, Android handoffs, and opt-in native Accessibility Service control.

- AMD Developer Hackathon
- Vision and Multimodal AI
- AI Agents
- Akash + Qwen
- MI300X Training
- Hugging Face

The Problem

Blind and partially sighted people may be alone without knowing what is directly around them: chairs, stairs, curbs, doors, signs, people approaching, vehicles, or blocked paths.

- Maps solve long-range routing, not immediate room safety.
- OCR reads text, but does not guide movement.
- General assistants do not understand the physical environment.
- WHO reports at least 2.2B people with near or distance vision impairment.

Our Solution

A voice-first AI mobility assistant that turns a phone camera into practical environmental awareness.

- Ask: guide me, what is ahead, read that sign, is someone approaching, call my mother.
- Camera-first local guidance: ignore GPS for room-scale movement.
- One calm physical action at a time: stop, turn, step, pan, wait.

Product Experience

Setup, ask, move safely.

- First-run connector setup: camera, speech, contacts, location, email/WhatsApp handoffs, SMS, calls, maps, and Sentinel Vision Control.
- Natural voice commands and text fallback.
- 3D camera HUD for detected objects, hazards, connector status, and debugging.

3D Environment HUD

A Terminator-style visual layer sits over the camera for demos, partial sight, caregivers, and debugging.

- Perspective grid, scan line, reticle, object chips, hazard markers.
- Shows door, chair, stairs, car, person, sign, obstacle, path.
- Voice remains primary; visuals support testing and presentation.

AI Architecture

Multimodal model + RAG + agent router + OS/Accessibility handoffs.

- Akash ML API with Qwen/Qwen3.6-35B-A3B for live app inference.
- Vision RAG for indoor rooms, sidewalks, stairs, curbs, signs, and object finding.
- Agent router for navigation, calls, SMS, WhatsApp/email handoffs, maps, battery, time, apps, Accessibility Service actions, and conversation.

Why AMD Hackathon

The project combines high-impact accessibility with multimodal AI, agentic workflows, and AMD GPU training.

- Primary fit: Vision and Multimodal AI.
- Secondary fit: AI Agents and Agentic Workflows.
- Fine-Tuning on AMD GPUs: SentinelBrain-14B-MoE-v0.1 is a proof-of-concept MI300X/ROCm training artifact.
- Hugging Face Space and public model progress pages support the submission.

Build Stack

Web + Android + AMD training + Apple roadmap.

- Web: HTML/CSS/JS, camera, speech, HTTPS delivery, local command routing.
- Android: Expo SDK 54, React Native, EAS APK, Android connectors, native Accessibility Service.
- Apple: planned Expo/EAS iOS path when credentials and platform work are ready.
- AI: Akash/Qwen live inference, SentinelBrain MI300X training artifact, Hugging Face pages.

Benefits

Independence, safety, and access on everyday phones.

- Know if someone approaches or if an obstacle blocks the path.
- Read signs, labels, screens, and door numbers.
- Use calls, messages, WhatsApp/email handoffs, maps, nearby search, ride handoff, and contacts by voice.
- Lower barrier than specialized wearable hardware.

SentinelBrain Training

A proof-of-concept competition training artifact, not the finished production app model.

- AMD Developer Cloud + MI300X + ROCm/PyTorch supported 14B+ BF16 experiments.
- Training architecture included sharded datasets, long-context SFT, watchdogs, checkpoint recovery, and Hugging Face publication.
- The planned 8,544-step 6K SFT pass needs roughly 45-60 hours on one MI300X-class GPU; production quality needs more training and benchmark gates.

Roadmap

From hackathon prototype to production assistive platform.

- Keep HF Space and SentinelBrain model page current.
- Continue SentinelBrain training and evaluation beyond proof of concept.
- Add OAuth connectors for Gmail/Outlook/calendar and emergency contacts.
- Add depth, optical flow, on-device fallback, caregiver dashboard, and field testing.

Submission

Public web demo, Android APK, whitepaper, pitch deck, Hugging Face Space target, and SentinelBrain training artifact.

- Demo: amdvision.qubitpage.com
- APK: includes opt-in Sentinel Vision Control Accessibility Service.
- Model/progress: SentinelBrain-14B-MoE-v0.1
- Goal: practical AI eyes for blind users.