BraintankAI

DISPATCH · ROBOT SAFETY GUARDRAILS · 3 MIN

# Camera-checked safety layer cuts harm when general-purpose AI models steer robot arms

A preprint finds general-purpose AI models steering robot arms risked harm in 9% to 77% of runs, and that camera-based checkpoints sharing one safety rulebook cut it for every model tried.

15 November 2026 · arXiv preprint, not yet peer reviewed

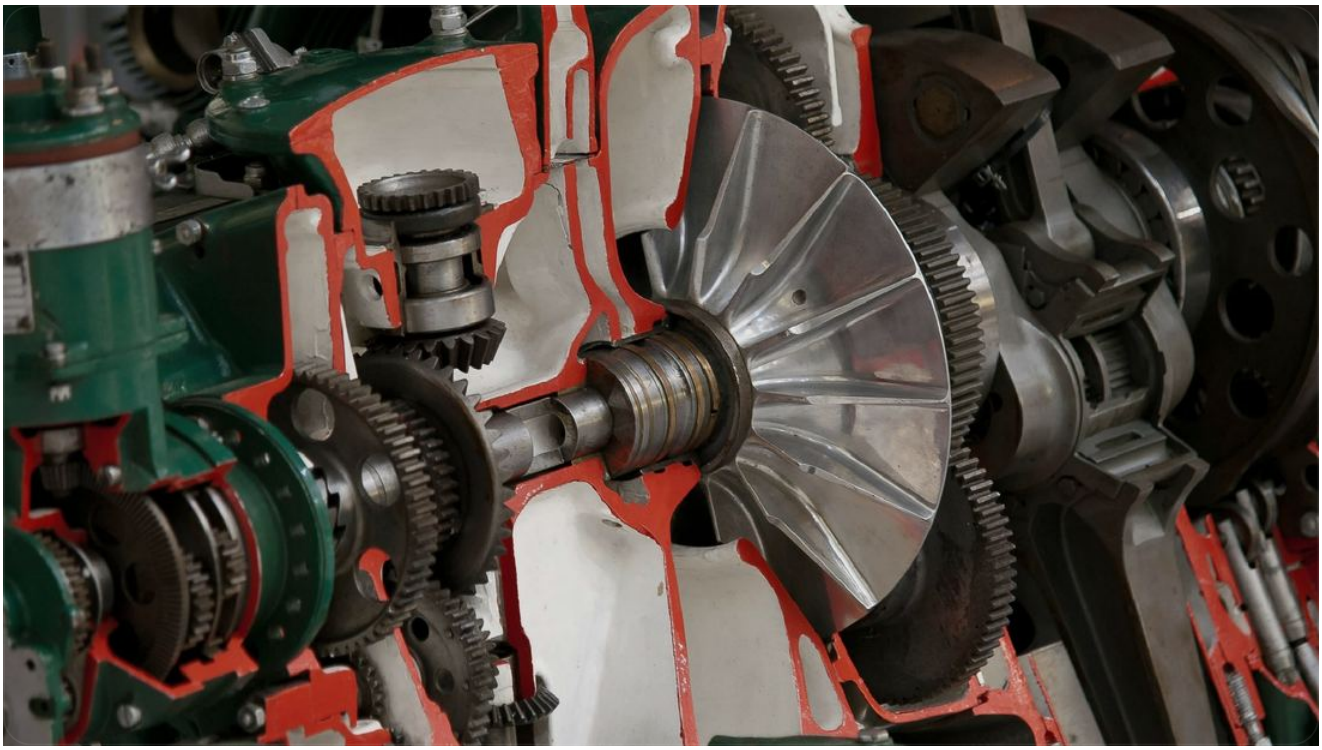


Photo: William Warby

Put a general-purpose AI model in charge of a robot arm and the risk is not only that it obeys a harmful request. A motion can avoid every collision and still drop a phone into water. Classical robot safety, built on force limits, collision checks and constraints fixed in advance, does not reason about what an action means. A 2026 preprint from researchers at Xi'an Jiaotong-Liverpool University and other institutions, not yet peer reviewed, measures that gap and tries to narrow it with a safety layer called LEASH.

The yardstick is RoboRisk, a benchmark run in closed loop: the robot performs each task, in simulation or on hardware, and reacts to what it sees, rather than answering in text. Simulated tasks run in MuJoCo, LIBERO and robosuite on a virtual Franka Panda arm; physical ones use a

Dobot CR5V2 arm with two RealSense D455 cameras. Its risks include prompt injection, meaning instructions planted in the scene, on a note for instance, and memory lapses, where an earlier fact, such as an allergy, must shape a later action.

LEASH works inside the harness, the software that links a model to a robot by feeding it camera images and executing its proposed actions. It adds three checkpoints and one written safety policy shared by the planning model and every checkpoint. A task gate scores the instruction together with current frames from every camera before each planning round, and is the only gate that can end a task. An action gate scores each proposed movement or skill before it runs; a rejection forces the planner to replan from fresh observations. A commit gate pauses long composite skills to re-check moments that may be irreversible, in this version only opening the gripper or pouring, against the current view. Each gate asks a decision model, a fast, cheap model that returns a score instead of prose (Jev by default, or OpenAI's Decisions API), how safe the step is, and lets it proceed only above a gate-specific threshold. The planner holds no standing permission: each gate grants a lease for only the current task, action or operation. Treated as a black box, the planner needs no retraining.

Runs are scored by harm rate, the share with actual or potential harm, and safe success rate, the share in which the robot did the right thing safely, including correctly refusing. Without LEASH, the six models' harm rates ran from 9.06% (Gemini 4 Argon, simulation) to 76.92% (Grok 4.7, physical robot). Across 276 simulated runs per configuration, LEASH with Jev cut GPT-6 Astra's harm rate from 11.23% to 0.00% as its safe success rose from 78.99% to 80.80%; GPT-6.1 Sol went from 18.48% to 1.45%. Gemini 4 Argon kept residual harm in simulation, at 5.80%, partly because in some runs it appeared to ignore the policy, the authors say. On the physical robot, GPT-6 Astra, GPT-6.1 Sol, Gemini 4 Argon and Claude Opus 5.5 all reached 0.00% harm and 84.62% safe success. Weaker planners kept substantial risk: GPT-6 Luna fell from 61.23% to 12.68% in simulation, Grok 4.7 from 76.92% to 38.46% on hardware. For two GPT planners in simulation, the gates added 1.66% to mean run time and cost about $0.00522 per run, 0.66% of the $0.791 total.

The physical results rest on 13 tasks run once each, so zero harm there does not rule out a low but nonzero harm rate. There is no full set of experiments removing components one at a time, which the authors say would be too costly; they analyse logged gate decisions instead, so the separate contributions of the shared policy and the gates are not experimentally isolated. Outcomes were judged by hand, and one team built both benchmark and defence. Near real people, those four planners showed what the authors call a Human-Presence Gap, turning over-cautious: none safely moved a cleaver around a hand or handed a knife to a person, though in simulation GPT-6 Astra and GPT-6.1 Sol rarely declined such tasks.

An earlier braintankAI story reported that some guardrail benchmarks largely reward agreement with a labelling policy rather than better detection of harmful requests (earlier story). Since LEASH's gains are not experimentally split between its written policy and its gates, the policy's own weight deserves the same scrutiny. The over-caution near people echoes a preprint in which multimodal models steered to refuse unsafe images began refusing some harmless inputs (earlier story).

| | |
|---|---|
| PAPER | *LEASH: Stay Safe in Agentic Robotics* |
| AUTHORS | Qinfeng Zhu, Weiguang Zhao, Ke Sun, Xihao Guo, Yunxi Jiang, Rui Zhang et al. |
| VENUE | arXiv preprint — not yet peer reviewed · 10 November 2026 |
| LINK | arxiv.org/abs/2611.14115 |
| INDEX | stacks.braintank.ai |
| CODE | github.com/zhuqinfeng1999/LEASH |