HANGZHOU · Ant Group (inclusionAI)

Two thousand hours of smartphone video of human hands, packaged for training robots

Open-AoE gathers phone-recorded footage of people handling objects, annotates hand poses and action steps, and supplies tools to turn it into robot training data.

15 November 2026 · arXiv preprint, not yet peer reviewed

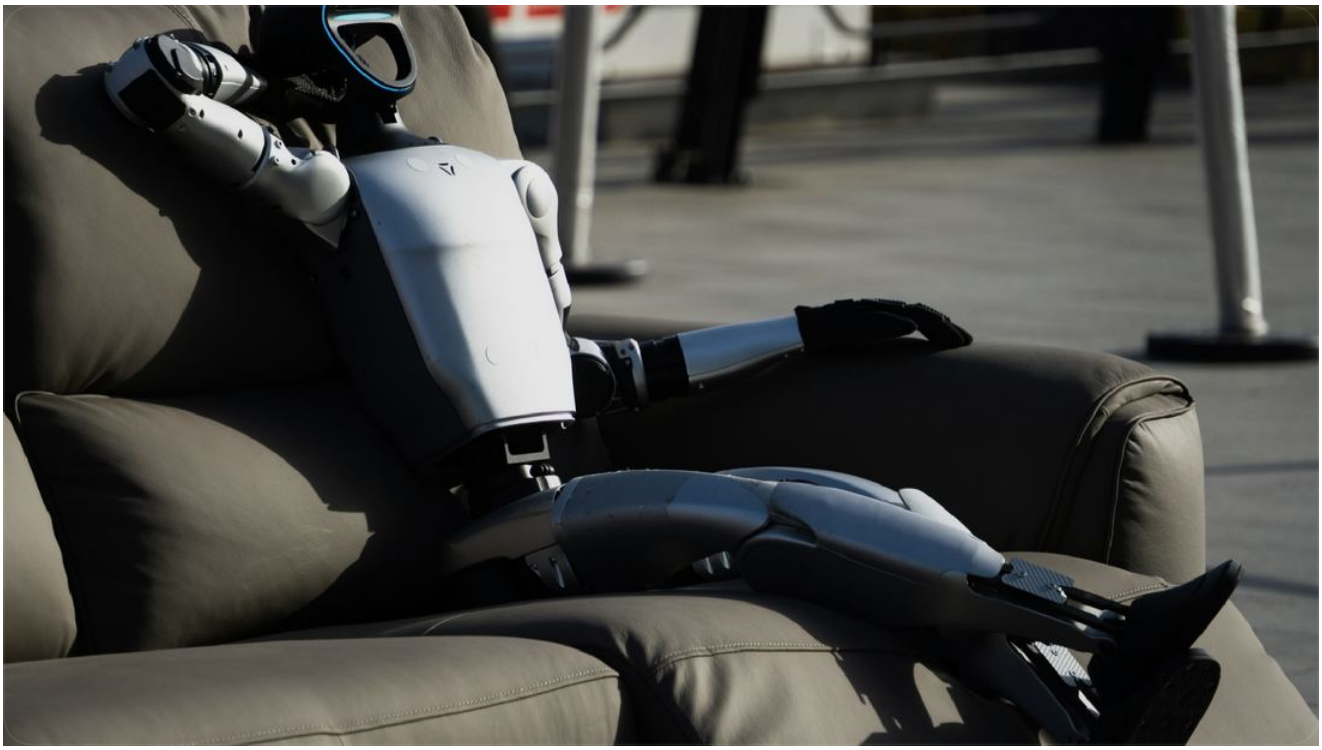


Photo: Yang Plasticine

Robots that learn to handle objects need large amounts of examples, and one of the cheapest sources is video of people doing ordinary tasks. A team led by Ant Group has released Open-AoE, about 2,000 hours of egocentric video, meaning footage filmed from the person's own point of view, here by a smartphone worn while they work. The video comes from more than 500 contributors using over 400 phone models in more than 400 scenes. Each clip is annotated with 3D hand poses, the camera's path through space and short text descriptions of each step. The release also includes open software for converting all of this into inputs for several kinds of robot-learning model. The work, described in a 2026 preprint on arXiv that has not yet been peer reviewed, is a dataset and infrastructure release. It does not show that robots trained on the data perform better.

The authors start from a practical gap. High-quality recordings of human manipulation usually depend on specialised headsets or controlled capture setups, which limits who can contribute. Large collections of everyday video are easier to gather but generally lack the details a robot needs: where the hands are, how the camera moved and when one action ends and the next begins. Tools for processing such footage exist, but each tends to be built for one dataset. The authors frame their goal as releasing data that can be used directly for training, not just videos.

Capture runs on ordinary phones. Lightweight on-device vision models start and stop recording according to whether hands are visible, give voice prompts when the hands drift out of frame or the phone is badly worn, and upload only valid clips. In the cloud, faulty, shaky or hand-less footage is discarded, and faces and sensitive content are pixelated. A vision-language model, an AI system that reads images and text together, then checks each clip and tags it by domain, scene, task and key objects. Reconstruction software estimates the camera's six-degree-of-freedom trajectory (three numbers for position, three for rotation) and fits MANO, a standard 3D model of the human hand, to both hands. A final optimisation refines hands and camera together so they share one consistent coordinate frame. Clips are split into atomic actions, short steps such as picking up a cup, with human correction.

The toolchain can retarget human hand motion onto several commercial robot hands and humanoids, meaning it converts the motion into joint movements a particular machine can execute. It can also produce robotised video, in which the human arm is painted out and a simulated robot is rendered performing the same motion. Further adapters reshape the data for existing vision-language-action models, which map images and instructions to robot commands, and for world models, which predict how a scene will change.

The evaluation compares Open-AoE with three other datasets, OpenEgo, EgoDex and EgoXtreme, using roughly 100-hour samples, or EgoXtreme's whole, smaller release. Across 3,000 resampled trials using image embeddings, Open-AoE scored highest on average on all six measures of visual diversity, though its lead on the coverage and entropy measures was small, particularly over OpenEgo. In a 100-hour audit, action labels covered 99.99% of video time, with 32,407 distinct descriptions and 175 action verbs. About 98% of entries had both hands validly tracked, and the footage yielded around 1,760 four-second training windows per hour, 97.8% of the theoretical maximum. When a separate vision-language model, Idefics2, rated how well annotations matched images on a five-point scale, Open-AoE averaged 4.583, against about 3.0 for OpenEgo, 2.9 for EgoDex and 2.0 for EgoXtreme.

The authors call that last result an audit under a shared judge rather than a controlled test of annotation accuracy. The datasets use very different labelling styles, and Open-AoE's free-text English labels were themselves model-generated, then reviewed by humans. Some of the high completeness figures also partly reflect the authors' own filtering gates. The accuracy of the reconstructed hands and camera paths is not checked against ground truth, and the detailed statistics come from a 100-hour sample of the full 2,000. Most importantly, no model or robot is trained on the data. The authors describe the benefit of having many different phone cameras as a hypothesis still to be tested through controlled experiments, and whether this collection makes better training material remains an open question.

PAPER	*Open-AoE: An Open Egocentric Manipulation Dataset and Toolchain for Embodied Learning*
AUTHORS	Zishuo Li, Bowen Yang, Changtao Miao, Kai Zhu, Hao Chen, Qingze Guan, Zhengxing Wu, Wanke Zhan, Yang Sun, Zhiyi Huang, Zitong Shan, Zhenchao Jin et al.
VENUE	arXiv preprint — not yet peer reviewed · 15 July 2026
LINK	arxiv.org/abs/2607.14183
INDEX	stacks.braintank.ai
CODE	github.com/ant-research/Open-AoE