Paper review · Visionary

Calibrated noise injection for VLA fine-tuning

DART, Laskey, Lee, Fox, Dragan and Goldberg, CoRL 2017 (arXiv:1703.09327).

Pavel Bushuyeu · November 2026

Does DART's calibration still matter for VLAs?

What DART does

It adds noise to the expert's executed actions and keeps the expert's clean action as the label.

It sizes the noise to the learner's error. Collection stays off-policy, so the expert keeps control.

What changed since

Zhang et al. (2025): one-shot isotropic noise with clean labels is on par with DART, for small state-based policies.

MimicGen adds noise but stores the noisy action as the label.

DiscoDemo (2026) ran fixed-noise DART for π0.5 in simulation. No paper sizes the noise to a VLA's own error.

Hypotheses

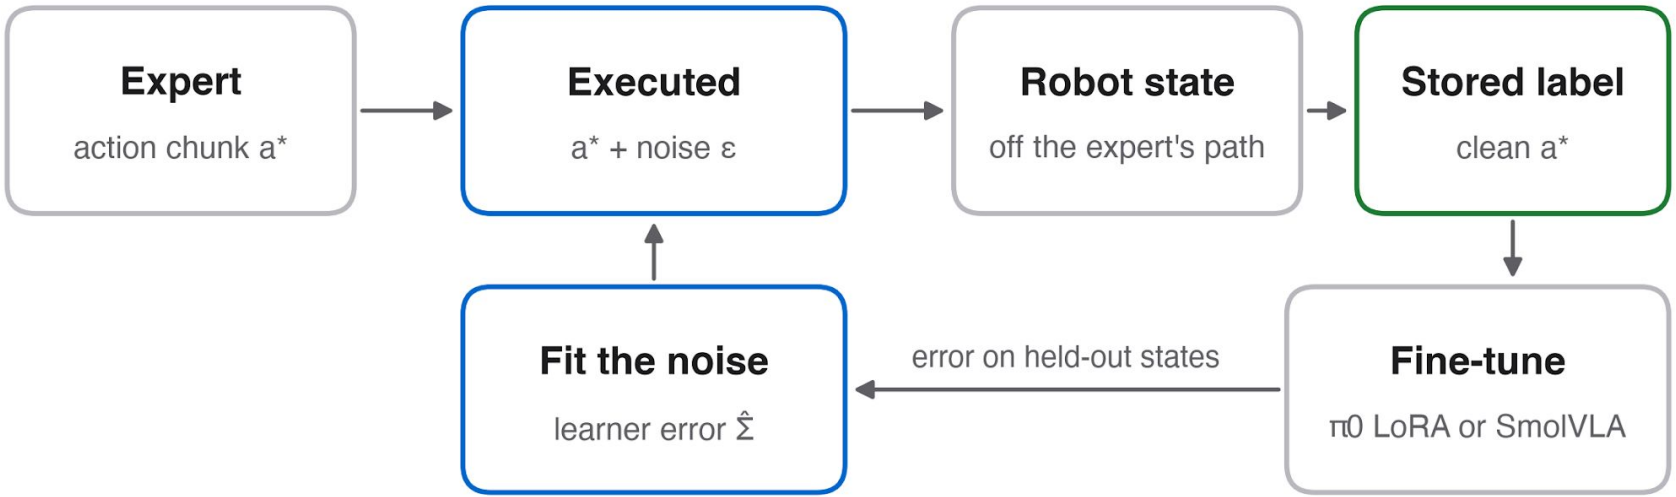
H1 For π0 LoRA fine-tuning, noise sized to the learner's error beats isotropic noise of the same total size by at least 5 points of success on 2 of 3 tasks.

H2 Noise correlated across the action chunk beats noise drawn independently at each step.

H3 At the same noise, clean labels beat MimicGen's noisy labels.

If H1 fails, the cheaper isotropic recipe is enough for VLAs too. Either outcome changes how fine-tuning data is collected.

One collection loop, six arms



Arms: clean · MimicGen default (σ 0.05, noisy labels) · the same noise with clean labels · isotropic sweep · DART per step · DART across the chunk.

Design

Tasks	MimicGen robosuite Square, Threading and Coffee (D1).
Expert	The MimicGen generator, patched to log the clean action.
Policy	π0 with LoRA (openpi); SmolVLA for the noise sweeps.
Round 0	200 clean demos; fit the learner's error covariance and its correlation across the chunk on held-out expert states.
Round 1	800 more successful demos per arm.
Metrics	Success with Wilson intervals, DART's shift gap, the expert's success under noise, recovery after a push.

Compute and risks

Compute (estimate)

150–250 GPU-hours for collection and LoRA fine-tunes, on one 32 GB GPU plus a university cluster.

One arm is timed first.

Risks

The pretrained VLA may already recover: π0 credits its diverse pretraining data. A null result still tells data collectors to skip calibration.

The waypoint expert corrects arm drift, not object drift.

Noise lowers the expert's success: keep successful demos, report the yield.

A noise-on mode for VLA teleoperation

1. Fit the noise from the first fine-tune's errors on held-out demos.
2. During leader-follower teleoperation (ALOHA, DROID, LeRobot), add short calibrated pushes to the follower arm.
3. Store the operator's command as the label.
4. Keep the pushes short, so the operator corrects the follower and does not build the push into the command.

Why it can work

Evidence

DART on a real robot: grasping in clutter went from 49% (behavior cloning) to 79% success, averaged over 4 human supervisors.

Twice that noise hurt the humans (72%).

Prior work

BDI and DIPA inject disturbances into operator actions on real robots.

Data Inoculation (2026) steers teleoperation toward a learner's likely errors with a learned model, for a small flow policy.

None collects VLA fine-tuning data with noise sized from the VLA's error.

What this builds on

Work	What it does	What this adds
DART, 2017	Noise sized to the learner's error, clean labels, small policies.	A pretrained VLA, action chunks.
Zhang et al., 2025	One-shot isotropic noise matches DART for small policies.	The test at VLA scale.
Policy-Calibrated DAgger, 2026	Offline noise statistics from a diffusion policy.	Noise during collection.
MimicGen, 2023	Noise in synthetic demos, noisy labels.	Clean labels, sized noise.
DiscoDemo, 2026	Fixed-noise DART for π0.5 in simulation.	Noise sized to the VLA.
Data Inoculation, 2026	Learned perturbations of teleoperation.	A VLA, no learned model.
π*0.6 (Recap), 2025	Corrections, greedy exploration.	Designed perturbations.

Sources

1. Laskey et al. DART: Noise Injection for Robust Imitation Learning, CoRL 2017.
2. Zhang et al. Action Chunking and Exploratory Data Collection Yield Exponential Improvements in Behavior Cloning, 2025.
3. Wang and Kantor. Policy-Calibrated DAgger, 2026.
4. Mandlekar et al. MimicGen, 2023.
5. Park et al. DiscoDemo, 2026. Li, Zhou, Xu and Levine. Robotic Data Inoculation, 2026.
6. Black et al. π0, 2024. Physical Intelligence. π*0.6 (Recap), 2025.
7. Oh et al. Bayesian Disturbance Injection, 2021. Tahara et al. DIPA, 2023.