CLOSED-LOOP ADAPTATION OF FROZEN VISION-LANGUAGE-ACTION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Execution errors can prevent vision-language-action (VLA) policies from completing tasks they solve under nominal robot dynamics. Correcting these errors changes both motion and the observations driving subsequent policy commands. We present an execution interface, calibrated from fault-free motion, that estimates one additive offset per modeled command coordinate and applies bounded correction with frozen policy weights. To analyze the resulting closed-loop system, we introduce crossed command-stream replay, which separates direct correction from changed-command effects while retaining observer feedback. Geometric and finite-horizon analyses relate physical repair to response direction, correction authority, motion-error persistence, and information delay. Our simulation study spans π_0, $\pi_{0.5}$, OpenVLA-OFT, and GR00T N1.5/N1.7 in four robot interfaces. Across four LIBERO benchmark suites, $\pi_{0.5}$ success under additive command faults rises from 58/220 to 121/220 over two sampler seeds on the same 110 task/state cases held out from calibration and method selection; healthy success changes from 209/220 to 213/220. With four previously untested rotation-offset vectors in forty new task/state cases, adaptation increases success from 48/160 to 125/160; healthy success is 37/40 with and without adaptation. The averaging gain selected for development is 0.16 versus 0.08 in the initial-state holdout. Panda-arm calibration transfers to OpenVLA-OFT and GR00T N1.7. With commands fixed, correction reduces median endpoint deviation from 13.7 to 7.8 mm. Replay shows that changed commands can reduce or increase reference-relative translation path-deviation cost. Separate predictor interventions distinguish identification accuracy from task recovery.

1 INTRODUCTION

Vision-language-action (VLA) policies translate images and instructions into robot commands (Black et al., 2025b;a; Kim et al., 2025b; NVIDIA et al., 2025). Their success also depends on how the robot executes those commands. A policy that completes a task under nominal dynamics can fail when an execution fault changes the resulting motion. We address this problem by adapting the execution interface while keeping the pretrained policy fixed.

Our interface uses commanded actions and measured robot motion to estimate additive command offsets (Figure 1). A predictor calibrated from fault-free motion estimates the response to the commands sent to the robot. An observer updates the offset estimate from the difference between predicted and measured motion, then supplies a bounded correction to subsequent commands. This adaptation requires neither policy-weight updates nor access to internal VLA representations.

Correction changes robot motion and can therefore change the observations and subsequent commands of the VLA. These policy responses can reinforce or offset the direct benefit of correction. To distinguish the two effects, we introduce *crossed command-stream replay*. We record command sequences from executions with adaptation off and on, then replay each sequence from the same simulator checkpoint under both adapter settings. Each enabled observer updates from its own replay measurements. Comparing adapter settings under the same sequence measures direct correction; comparing sequences under the same adapter setting measures the effect of changed commands. We assess these effects through deviation from a common reference trajectory and evaluate task success separately.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

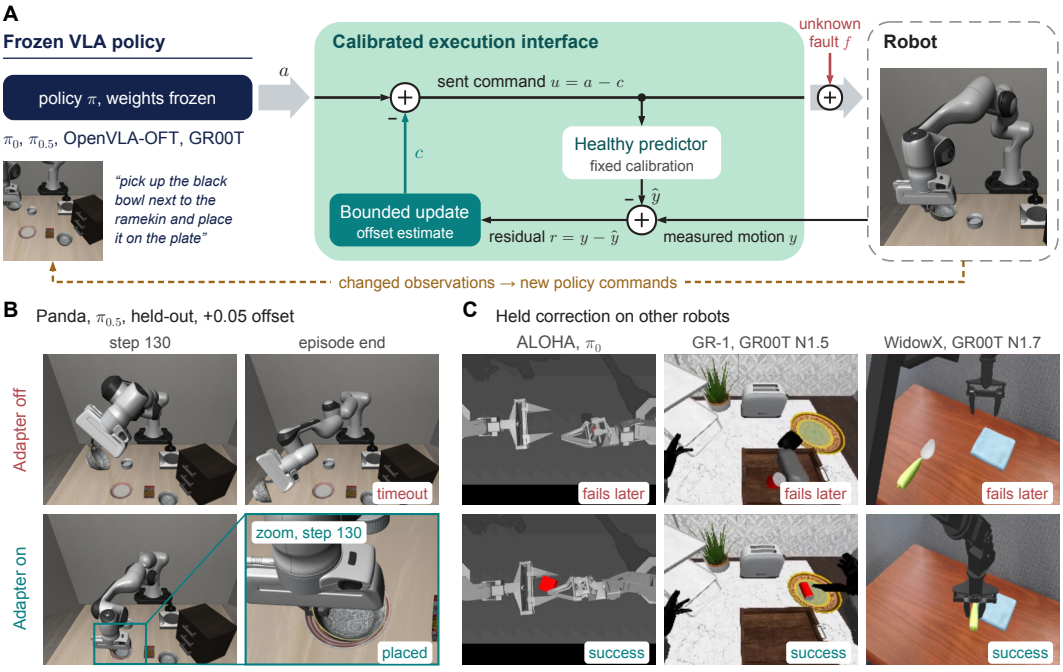


Figure 1: **Frozen policy, adaptive execution.** (A) Policy command a minus bounded correction c gives $u = a - c$ before unknown fault f. Measured motion y and healthy prediction $\hat{y}$ give residual $r = y - \hat{y}$; dashed feedback carries policy observations. (B) A held-out Panda/$\pi_{0.5}$ episode under a six-coordinate $+.05$ offset, re-rendered from its recorded commands; off and on share the sampler seed. On places the bowl at step 130; off times out at step 229. (C) Held correction on ALOHA, GR-1 and WidowX under separate protocols: uncorrected and corrected runs at the same step; the uncorrected runs later fail (provenance: Appendix E).

We evaluate five VLA variants in simulation on Panda, ALOHA (Zhao et al., 2023), GR-1 (NVIDIA, n.d.), and WidowX (Walke et al., 2023), including GR00T N1.5 and N1.7 (NVIDIA, 2025; 2026b) (Table 1). Panda calibration is reused across compatible policy backbones; other robots use separate calibration. Panda adapts during task execution; faulted evaluations on the other robots identify offsets in preceding episodes and then hold the correction fixed.

Our contributions are:

1. **Calibrated execution adaptation across frozen policies.** On held-out initial states across two sampler seeds, adaptation raises faulted $\pi_{0.5}$ task success from $58/220$ to $121/220$. A separate evaluation of four previously untested rotation-offset vectors raises success from $48/160$ to $125/160$.
2. **Controlled analysis of policy responses.** Crossed command-stream replay separates direct correction, changed-command effects, and their interaction while retaining the observer's execution feedback.
3. **Control-theoretic limits on physical recovery.** Conditional analysis and experiments examine how correction bounds, information delay, accumulated motion error, and predictor mismatch affect recovery, distinguishing offset estimation from physical and task outcomes.

2 Related Work

Adaptation around learned policies. Adaptive augmentation (Cheng et al., 2022; Huang et al., 2023) and residual policies (Silver et al., 2018; Johannink et al., 2019) already adapt pretrained behavior. Meta-learned adaptive control uses visual foundation-model features to adapt online to terrain and actuation changes (Lupu et al., 2025). Adaptive Policy Execution (APEX) adapts references beneath reacting policies, including $\pi_{0.5}$ on LIBERO, and is tested on hardware (Zhao

et al., 2026). The Torque Adaptation Module (TAM) learns torque residuals in randomized, robot-specific simulation (Son et al., 2026). Our comparisons include multiple observers and a reconstructed structured adapter. A faithful APEX test requires its reference reconstruction and discrete action-space update; we could not verify the latter from the available specification and found no released reference implementation. Our APEX-style comparison matches command support and cap, but does not reproduce published APEX (Appendix C.7).

VLA adaptation and robustness. The Joint-level Physical-fault Aware Residual Calibrator (J-PARC) uses fault-conditioned supervision (Jo et al., 2026); Asynchronous Action Chunk Correction (A2C2) learns from demonstrations (Sendai et al., 2025). Policy fine-tuning (Kim et al., 2025a; Peng et al., 2026) changes the learned policy. LIBERO-Plus studies scene, sensing and instruction shifts (Fei et al., 2025); our evidence principally concerns execution faults.

Identification, control, and attribution. Prediction accuracy and control usefulness differ (Gevers, 1993; Forssell & Ljung, 1999). Our observer uses standard components (Shim, 2020; Isermann, 2005; Ioannou & Sun, 1996), with bounds from established tools (Slotine & Li, 1989; Forni & Sepulchre, 2014; Geiselhart & Wirth, 2016). Causal and geometric identification study useful physical changes (Taheri et al., 2025; 2026). Li audits commands supplied to a frozen locomotion policy against fixed and frequency-matched interventions (Li, 2026). Our intervention crosses VLA-generated output streams with a downstream observer, rerunning its updates from each replay's evolving measurements. The protocol thus preserves downstream feedback while imposing the policy's recorded output. These controlled contrasts do not identify natural mediation effects (Pearl, 2001; Ge et al., 2023).

3 A CALIBRATED EXECUTION INTERFACE

3.1 POLICY, EXECUTION, AND ADAPTATION STATE

At execution step $k \geq 0$, a frozen policy π_ϕ issues the nominal command $a_k = \pi_\phi(h_k, \ell; \omega_k) \in \mathbb{R}^d$. Here ϕ are fixed weights, h_k contains observations and internal history, ℓ is the instruction, ω_k is sampling randomness, and d counts modeled joint-target or Cartesian coordinates. Action chunks are command sequences; unmodeled outputs pass through. The offset estimate $\hat{f}_k \in \mathbb{R}^d$ produces correction c_k through a diagonal binary mask $D_m \in \mathbb{R}^{d\times d}$:

$$c_k = D_m \hat{f}_k, \qquad u_k = a_k - c_k, \qquad x_{k+1} = \mathcal{F}_0(x_k, u_k + f_k). \tag{1}$$

The sent command is u_k, the additive execution fault is $f_k \in \mathbb{R}^d$, and $\mathcal{F}_0$ maps the robot–controller–object state x_k and command to the next state. The observer receives motion, not f_k. Position-zero errors can appear as joint-position offsets (ROBOTIS, n.d.); a wrong action-decoder mean at fixed scale gives an additive command error (Physical Intelligence, n.d.). This equivalence depends on the interface. Friction, torque and locking instead alter a physical-parameter vector θ_k in the dynamics map $\mathcal{F}(x_k, u_k, \theta_k)$ and need not admit an exact command-equivalent fault.

Measured motion $y_{k+1} \in \mathbb{R}^d$ comprises normalized end-effector displacement and relative rotation on Panda/WidowX, or joint position on ALOHA/GR1. Each interface has its own units and calibration. Measurements update the next command; history persists across replans. Unlike OpenVLA-OFT and GR00T, $\pi_{0.5}$-LIBERO does not condition directly on robot state (Appendix B.1).

3.2 CALIBRATED EXECUTION PREDICTION

A finite impulse response (FIR) predictor h_W maps the sent-command history $U_k = (u_k, \ldots, u_{k-K})$ to predicted healthy motion $\hat{y}_{k+1}$. The integer $K \geq 0$ is its maximum lag; j indexes lags:

$$\hat{y}_{k+1} = h_W(U_k) = \sum_{j=0}^{K} W_j u_{k-j} + b_h, \qquad r_{k+1} = y_{k+1} - \hat{y}_{k+1}. \tag{2}$$

Healthy data fit diagonal lag matrices $W_j \in \mathbb{R}^{d\times d}$ and intercept $b_h \in \mathbb{R}^d$; r_{k+1} is the motion-prediction residual. Physical dynamics may still be coupled. The calibrated *sensitivity* $M \in \mathbb{R}^{d\times d}$ maps command offsets to motion; its Moore–Penrose pseudoinverse $M^\dagger$ gives the offset target

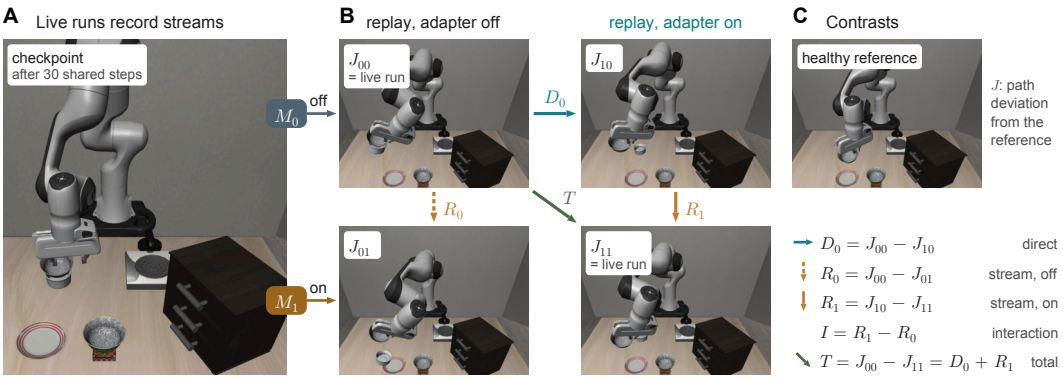


Figure 2: **Crossed command-stream replay.** (A) Two live runs of one source share 30 policy steps; from this checkpoint, adaptation off/on records streams M_0, M_1. (B) Each stream is replayed from the checkpoint with adapter setting $a \in \{0, 1\}$; frames show the end of the 50-step window, and diagonal cells reproduce the live runs. (C) $J(a, M)$ is healthy-reference path-deviation cost with imposed stream M; cost indices select adapter and stream. D_0 is direct benefit, R_0, R_1 stream effects without/with adaptation, I their interaction and T total benefit; positive values mean lower cost. Enabled observers use their own measurements; task outcomes retain all assigned sources.

$z_{k+1} = M^{\dagger} r_{k+1} - b$, with healthy centering vector $b \in \mathbb{R}^d$ (zero in reference Panda runs). Finite probes measure a windowed response, not necessarily steady-state gain. Policy transfer requires a shared physical interface and calibrated command-history domain (Appendix B.2).

3.3 Bounded observer updates

For coordinate $i = 1, \ldots, d$, let $c_i^{\max} \geq 0$ be its correction cap. Euclidean projection $\Pi_{\mathcal{C}}$ onto $\mathcal{C} = \{v \in \mathbb{R}^d : |v_i| \leq c_i^{\max}\}$ clips each coordinate. The matrix $J_g \in \mathbb{R}^{d_g \times d}$ selects d_g residual coordinates for gating. Write $\|\cdot\|_2$ for the Euclidean norm and $\mathbf{1}[\cdot]$ for the indicator of a true condition. For magnitude $n \geq 0$, set $s(n) = \mathbf{1}[n \geq \delta]/[1 + (n/\rho)^2]$, with deadzone threshold $\delta \geq 0$ and attenuation scale $\rho > 0$. The normalized-target (NT) rule, with averaging gain $0 < \gamma \leq 1$, is

$$\hat{f}_{k+1}^{\mathrm{NT}} = \Pi_{\mathcal{C}}\big[(1-\gamma)\hat{f}_k + \gamma s(\|J_g r_{k+1}\|_2) z_{k+1}\big]. \tag{3}$$

NT smooths the attenuated target and decays toward zero when $s = 0$. Projection bounds correction, not healthy task performance. The calibrated disturbance observer (DOB), with averaging gain $0 < \alpha \leq 1$, removes NT's gate and attenuation: $\hat{f}_{k+1} = \Pi_{\mathcal{C}}[(1-\alpha)\hat{f}_k + \alpha z_{k+1}]$. Appendix B.1 specifies innovation-based and robot variants.

Development fixes masks and gains. Panda updates from zero within an episode; faulted evaluations on the other robots identify then hold across reset episodes. Scheduled studies enable adaptation at fault onset; hidden-onset studies enable it earlier. A separate gain-compensation variant estimates a command multiplier g_i through $g_i - 1$ and sends $a_i/\hat{g}_i$, where $\hat{g}_i$ estimates g_i. Proposition 6 bounds NT error under attenuation, model mismatch and fault variation.

Adaptation dimension and computation. The adapted state scales with command dimension, independently of VLA size; VLA weights require no updates or backpropagation. A recorded Panda collection measures median adapter computation of 0.174 ms per step versus 895 ms per policy call (Appendix B.1).

4 Four mechanisms of closed-loop execution repair

We analyze command-stream response, correction authority, accumulated displacement and observer bias. Their consequences concern physical execution; task recovery is measured separately.

Controlled trajectories and their cost. Fix an integer horizon $N \geq 1$, initial physical state, controller histories and adapter initialization, fault schedule, exogenous randomness and healthy

reference trajectory. For n measured coordinates, let $e_{i,k} \in \mathbb{R}^n$ be branch i's deviation from the reference at step $k = 1, \ldots, N$, stacked as $E_i \in \mathbb{R}^{nN}$. Positive-definite step weights $Q_k \in \mathbb{R}^{n\times n}$ form the block-diagonal matrix $Q = \mathrm{diag}(Q_1, \ldots, Q_N)$. Define cost J_i and benefit $\mathcal{B}_N$ for unadapted/adapted branches $0/A$, using the weighted norm of any vector $v \in \mathbb{R}^{nN}$:

$$J_i = \|E_i\|_Q^2, \qquad \mathcal{B}_N = J_0 - J_A, \qquad \|v\|_Q^2 = v^\top Q v. \tag{4}$$

Positive $\mathcal{B}_N$ denotes reduced physical cost. In Figure 2's within-source 2×2 design, M_0, M_1 are live unadapted/adapted command streams and J_{am} is replay cost with adapter setting $a \in \{0, 1\}$ and stream index $m \in \{0, 1\}$. The fixed-stream adapter contrast includes the observer's dynamic response to its own measurements.

Define direct benefit D_0, stream replacement R_0 without adaptation, cost interaction I, and total benefit T by

$$\begin{aligned} D_0 &= J_{00} - J_{10}, \qquad R_0 = J_{00} - J_{01}, \\ I &= J_{10} + J_{01} - J_{11} - J_{00}, \qquad T = J_{00} - J_{11} = D_0 + R_0 + I. \end{aligned} \tag{5}$$

The additional stream effect under adaptation is $R_1 = J_{10} - J_{11}$. It includes changes in subsequent observer behavior.

Proposition 1 (Replay decomposition and response geometry)**.** *Let $q_0 = \|E_{00}\|_Q$, $q = \|E_{10}\|_Q$ and $\Delta = E_{11} - E_{10}$. The response magnitude is $r_\Delta = \|\Delta\|_Q$ and $A = E_{10}^\top Q \Delta$ its alignment with the remaining deviation. Then*

$$R_1 = -2A - r_\Delta^2, \qquad T = D_0 + R_1 = D_0 - 2A - r_\Delta^2. \tag{6}$$

For $q r_\Delta > 0$, define the response ratio $g_{\mathrm{resp}} = r_\Delta / q$ and weighted angle $\theta \in [0, \pi]$ by $\cos\theta = A/(q r_\Delta)$. Then $R_1 > 0$ exactly when $\cos\theta < -g_{\mathrm{resp}}/2$; $g_{\mathrm{resp}} > 2$ forces $R_1 < 0$. For any response-magnitude bound $\beta \geq r_\Delta$,

$$T \geq D_0 - 2q\beta - \beta^2 = q_0^2 - (q + \beta)^2. \tag{7}$$

This bound is sharp over the full norm ball: strict benefit for every displacement in that ball holds if and only if $\beta < q_0 - q$.

The ratio classifies realized R_1, not total benefit T or stability. It is reference-relative and includes plant and observer responses. When the reference is the healthy J_{00} trajectory, $T \leq 0$, but not necessarily $R_1 \leq 0$. Quadratic interaction need not imply trajectory nonadditivity (Appendix A).

Mask and cap limit command cancellation. For a command fault $f \in \mathbb{R}^d$, candidate estimate $h \in \mathcal{C}$, and the mask and box in Section 3, let I_{cmd} denote the identity on the d command coordinates. The smallest remaining error is

$$\min_{h\in\mathcal{C}} \|f - D_m h\|_2^2 = \|(I_{\mathrm{cmd}} - D_m) f\|_2^2 + \|D_m(f - \Pi_{\mathcal{C}} f)\|_2^2. \tag{8}$$

The two terms are uncorrected coordinates and cap excess. This exact distance bounds command cancellation, not task success (Lemma 1).

Theorem 1 (Memory–authority–delay limit)**.** *Let $F, C \geq 0$ be the known fault magnitude and correction cap, $\lambda \geq 0$ the memory factor and $g > 0$ the command-to-motion gain. For an unknown constant fault $f \in \{-F, F\}$, scalar deviation p_k and correction $c_k \in [-C, C]$, suppose $p_0 = 0$ and $p_{k+1} = \lambda p_k + g(f - c_k)$ for $k = 0, \ldots, N-1$. Suppose deterministic correction policies receive identical information under either fault sign before an integer $\tau \in \{0, \ldots, N\}$, and the exact sign thereafter. With no further constraints, define $S_m(\lambda) = \sum_{j=0}^{m-1} \lambda^j$ for integers $m \geq 1$, $S_0 = 0$, and $[x]_+ = \max(x, 0)$ for real x. The smallest worst-case terminal displacement is*

$$\inf_{\text{admissible corrections}} \max_{f=\pm F} |p_N| = g[F S_N(\lambda) - C S_{N-\tau}(\lambda)]_+. \tag{9}$$

For an integrator ($\lambda = 1$) with $N > \tau$, zero worst-case terminal error is ideally attainable exactly when $C \geq FN/(N - \tau)$ (Corollary 1). Positive delay and $F > 0$ thus require authority beyond fault cancellation, not merely $C \geq F$. With $C \geq F$, zero correction before the delay followed by $c_k = f$ leaves $|p_N| = gF\tau$; state feedback with authority beyond F can remove it. These ideal scalar limits do not guarantee deployed-observer recovery (Appendix A.7).

Table 1: **Execution repair across robots, benchmarks and VLAs.** Successes, Off → corrected, with robot-specific calibration. Faults are additive command errors except the listed gain condition. Each arrow compares Off and correction within its stated protocol; identification episodes are excluded. †: healthy GR1 identifies for one episode and then holds; healthy WidowX adapts online. These differ from the corresponding three-episode faulted identification protocols. "Not reported": no matched healthy result. OpenVLA-OFT Object lists rotation/translation/gain faults. MuJoCo (Todorov et al., 2012) runs LIBERO, gym_aloha (Cadène et al., 2024) and RoboCasa-GR1 (Nasiriany et al., 2024; NVIDIA, n.d.); SAPIEN (Xiang et al., 2020) runs SimplerEnv (Li et al., 2025). Rows are not pooled.

Robot / benchmark	VLA	Task	Faulted	Healthy controls
Online adaptation during the evaluated episode				
Panda / LIBERO	$\pi_{0.5}$	Four suites	$58/220 \rightarrow 121/220$	$209/220 \rightarrow 213/220$
	OpenVLA-OFT	Spatial	Table 3	Table 3
		Object	$0 \rightarrow (10, 15, 19)/20$	Not reported
	GR00T N1.7	Spatial	Table 3	Table 3
Identify for one episode, then hold: same healthy/faulted schedule				
ALOHA / gym_aloha	π_0	Transfer cube	$0/39 \rightarrow 15/39$	$17/39 \rightarrow 13/39$
Identify for three faulted episodes, then hold: healthy protocols differ†				
GR1 / RoboCasa-GR1	GR00T N1.5	Plate → plate	$1/27 \rightarrow 17/27$	$18/29 \rightarrow 18/29$ †
		Tray → plate	$1/27 \rightarrow 13/27$	$14/29 \rightarrow 15/29$ †
WidowX / SimplerEnv	GR00T N1.7	Spoon → towel	$0/20 \rightarrow 14/20$	$15/20 \rightarrow 16/20$ †

Predictor mismatch shifts the observer's target. Let $G_j \in \mathbb{R}^{d\times d}$ map a lag-j command to motion for $j = 0, \ldots, K$. The sums $G = \sum_{j=0}^{K} G_j$ and $W = \sum_{j=0}^{K} W_j$ are constant-input, or direct-current (DC), gain maps. For an exact FIR response, constant nominal commands and fault, exact healthy centering, invertible M, zero estimator bias and inactive projection, DOB or active full-estimate innovation (19) at a stationary estimate $h_* \in \mathbb{R}^d$ requires

$$[M + (G - W)D_m]h_* = Gf. \tag{10}$$

An invertible bracket determines a unique candidate, not its stability. In the scalar corrected case $D_m = 1$, $M = G = g \neq 0$, $W = w$ and $f \neq 0$, constant NT attenuation $s \in [0, 1]$ gives $h_*/f = s/[1 + s(1 - w/g)]$ when the denominator is nonzero. Thus matching $w = g$ removes this mismatch bias at $s = 1$, while NT attenuation can remain. These are conditional stationary identities, not measured probe identities or recovery guarantees (Proposition 2).

5 Experiments

Coverage and protocol. Table 1 spans five VLA variants and four robot interfaces. LIBERO (Liu et al., 2023) supplies forty tasks in Spatial, Object, Goal, and LIBERO-10. Healthy $\pi_{0.5}$ Spatial motion calibrates Panda. The initial-state holdout and unseen-fault study use cases excluded from calibration and method selection. Historical reference/disjoint-calibration results retain tuning overlap (Appendix B.2). A *key* identifies task, state, sampler seed, and fault condition across arms; a *source* is a live rollout, and an *evaluation group* shares a protocol. Rescues/regressions change failure/success to success/failure. Historical sampling is unpinned. Unless stated otherwise, means weight tasks equally and marginal 95% intervals resample tasks without multiplicity adjustment. Registration denotes dated internal protocols archived with the submission; amendments and reused data are specified by study.

5.1 Calibrated feedback repairs execution across backbones and interfaces

With a six-coordinate Panda offset of $+.05$ normalized action units and bounded rotation correction, the holdout and its second-seed replication raise success from $58/220$ to $121/220$: 72 rescues, nine regressions, $+28.6$ percentage points $[16.2, 41.5]$. Both blocks use fixed NT settings and disjoint calibration; all four suite point estimates improve. Healthy success is $209/220 \rightarrow 213/220$ (nine rescues, five regressions). Both seeds reuse 110 task/state cases; all 880 episodes are retained (Table 2; Appendix B.2). On the Panda holdout, repair remains partial: faulted success is 55%, versus 95%

Table 2: **Repair on held-out initial states and historical evaluations.** Frozen $\pi_{0.5}$ versus rotation-only NT under the six-coordinate $+.05$ offset; n counts paired cases per condition and $+/-$ denotes rescues/regressions. Holdout A/B use pinned seeds 102001/102002 on the same 110 task/state cases, with fixed settings and disjoint calibration. B was registered after A, before B collection. Historical rows retain selection overlap and unpinned sampling; they are not pooled with the holdout.

		Faulted			Healthy		
Evaluation	n	Off	NT	$+/-$	Off	NT	$+/-$
Holdout: block A	110	29	62	37/4	107	108	3/2
block B	110	29	59	35/5	102	105	6/3
Holdout pooled	220	58	121	72/9	209	213	9/5
Historical: disjoint	120	34	68	37/3	118	118	2/2
reference	120	28	78	51/1	116	117	4/3

Table 3: **Panda calibration transfers to two further backbones.** Spatial successes, frozen $\rightarrow$ corrected, reusing $\pi_{0.5}$ calibration. Offsets use normalized units; gain uses a separate compensator. Healthy controls cover only the original twenty keys. The forty-key evaluation starts six-coordinate fault and rotation adaptation at step 30, without matched healthy-adapter arms. "Not evaluated": untested condition (Appendices B.3 and C.4).

Condition	OpenVLA-OFT	GR00T N1.7
Rotation offset .10	$0/20 \rightarrow 11/20$	$0/20 \rightarrow 14/20$
Translation offset .15	$0/20 \rightarrow 14/20$	$2/20 \rightarrow 15/20$
Multiplicative gain .20	$0/20 \rightarrow 17/20$	$0/20 \rightarrow 19/20$
Healthy, rotation adapter	$20/20 \rightarrow 19/20$	$18/20 \rightarrow 18/20$
		$4/40 \rightarrow 11/40$ (NT)
Mid-episode, uniform offset .10	Not evaluated	$4/40 \rightarrow 10/40$ (DOB)

for healthy Off. Rotation-only support avoids direct translation updates; six-channel correction adds faulted successes but also healthy regressions (Section 5.3).

ALOHA's identify-and-hold protocol repairs $0/39 \rightarrow 15/39$ faulted tasks. Identifying instead on a healthy first episode and holding that estimate gives $17/39 \rightarrow 13/39$ (four rescues, eight regressions). This healthy cost occurs without holding a fault-derived estimate; unpinned sampling limits attribution (Appendix B.3).

Repair transfers to previously untested command faults. Four withheld rotation-offset vectors on forty new cases give NT $48/160 \rightarrow 125/160$, with eighty rescues and three regressions (Table 4); all four NT–Off intervals are positive. Each correcting arm selects among four development candidates before testing unseen faults. NT selects $\gamma = .16$, versus the holdout's .08. The nominal-error integral, which integrates deviation from predicted nominal motion, exceeds NT by 8.1 points $[1.9, 14.4]$; healthy rankings remain unresolved (Appendix C.6).

Comparison with a structured reconstruction. With matched support, cap and four candidates each, NT exceeds the APEX-style reconstruction by 18.2 faulted-success points $[8.2, 28.4]$ (healthy Off/NT/reconstruction: $211/215/166$ of 220). The reconstruction selects the largest tested rate multiplier (1), leaving its optimum unresolved. Information, calibration, and updates differ from published APEX (Appendix C.7, Table 15).

5.2 The policy response changes the benefit of physical correction

For three-dimensional end-effector positions p_k and healthy-reference positions p_k^{ref}, translation cost is $J = .05 \sum_{k=1}^{50} \|p_k - p_k^{\mathrm{ref}}\|_2^2$: fifty samples at $.05$ s intervals. Angular cost uses relative-rotation error. Campaign primaries distinguish interaction, stream replacement, and total benefit (Appendix C.2). Physical costs use fidelity-qualified replays; task outcomes retain all keys. Imperfect source reproducibility limits attribution solely to replanning.

Table 4: **Execution repair on previously untested fault vectors.** Four rotation-supported faults on forty unused task/state cases. Faulted gains versus Off use 95% task-cluster intervals; +/− denotes healthy rescues/regressions. Static selects zero. This evaluation is separate from Table 2.

Adapter	Faulted /160	Gain (points)	Healthy /40	+/−
Off	48	—	37	—
Static (zero)	47	$-.6\ [-1.9, 0]$	37	0/0
NT	125	$48.1\ [40.0, 56.2]$	37	1/1
DOB	116	$42.5\ [35.0, 50.0]$	35	1/3
Nominal-error integral	138	$56.2\ [48.8, 63.8]$	36	2/3
L1-style	130	$51.2\ [43.8, 58.8]$	37	2/2

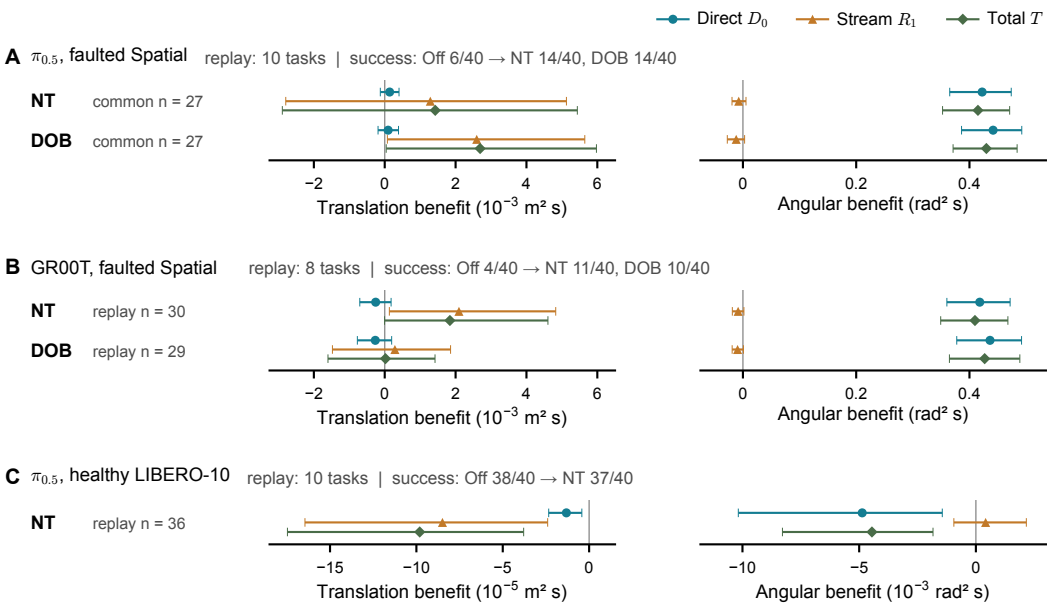


Figure 3: **Direct correction, changed commands and task recovery.** Equal-task effects with 95% task-bootstrap intervals; positive reduces healthy-reference deviation. Replay crosses command streams with the adapter off/on; n counts retained source cases. Metrics have separate units/scales. Task totals use forty keys; replay uses 27 common $\pi_{0.5}$, 30/29 GR00T NT/DOB and 36 healthy sources. Marginal intervals do not compare observers.

Under faults, correction has coordinate-specific benefit. Rotation-only NT correction has positive angular direct benefit on all thirty retained faulted sources in each backbone. GR00T's registered angular total benefit is $.410\ [.349, .468]$ rad² s. Translation R_1 is $2.10\ [.13, 4.83] \times 10^{-3}$ m² s on GR00T NT, contributing a benefit on average; its total translation benefit is unresolved. The $\pi_{0.5}$ NT translation contrasts are unresolved, while matched DOB replay has positive mean stream and total benefits (Figure 3). NT–DOB total-benefit differences remain unresolved. Cost interaction does not establish trajectory nonadditivity (Appendix C.10). Physical benefit and task success disagree in both directions: a $\pi_{0.5}$ regression has positive T, a rescue negative T (Appendix C.11).

Healthy stream costs depend on the reference. Relative to each paired healthy-Off path, $J_{00} = 0$: direct correction adds cost $-D_0 = J_{10}$. Added stream cost $-R_1 = J_{11} - J_{10}$ exceeds direct cost on average in all four evaluation groups; their intervals for $S = D_0 - R_1 = J_{11} - 2J_{10}$ are positive (Table 10). Using other healthy-Off paths of the same task/state in shared observation coordinates, all five $\pi_{0.5}$ NT intervals instead include zero (three positive means). GR00T retains four negative means, one resolved; integral correction retains three, two resolved. Departure from a paired path therefore does not establish worse typical healthy execution (Appendix C.11).

The response persists across adapters and without live state input. Each adapter gives healthy $R_1 < 0$ on all 33 valid integral/L1 sources (Appendix C.5). Holding GR00T's proprioceptive input fixed with live images leaves mean R_1 negative with a resolved interval: live proprioceptive changes are unnecessary for this cost. Visual mediation is unquantified (Appendix C.8).

Response magnitude and direction have distinct roles. The overshoot threshold $r_\Delta > 2q$ holds on $27/34$ Spatial, $20/36$ LIBERO-10 and $18/36$ GR00T healthy sources, versus $1/27$ NT and $0/27$ DOB in faulted $\pi_{0.5}$ replay. Magnitude-only bounds miss realized benefits: none of 113 strong-fault source–matrix records pass, although 66 have positive T. Records reuse keys; prospective forecasts remain nondecisive (Appendices C.9 and D.3).

5.3 Interventions test the limits of execution repair

A predicted identification change need not improve recovery. The predictor pilot meets its registered rotation-y median-estimate threshold in both suites (Appendix D.1). In the separate constraint confirmation (E2), lower Spatial observation bias yields no resolved task gain; LIBERO-10's point estimates move in the opposite directions (Table 5).

Table 5: **Identification and task recovery are distinct endpoints.** Sixty paired cases per suite; rotation-y constrained versus unconstrained predictor. Positive error reduction means lower absolute late-window observation bias. Marginal 95% intervals include zero for both task effects and LIBERO-10's error change (Appendix D.1).

Suite	Observation-error reduction (%)	Task change (percentage points)
Spatial	$+35.5\ [21.9, 47.6]$	$-5.0\ [-11.7, 0]$
LIBERO-10	$-33.4\ [-135.6, 39.8]$	$+11.7\ [0, 25]$

Additional authority permits more recovery, with healthy costs. Changing online correction from rotation-only to six channels raises faulted LIBERO-10 success $25/60 \to 41/60$; this comparison also changes normalization and calibration. Healthy Off versus six-channel correction gives $56/60 \to 53/60$ (three rescues, six regressions); rotation-only scores $52/60$. The known-fault oracle isolates support: rotation-to-six-channel success rises $36/60 \to 57/60$. Rotation-only recovery shows that task completion can occur without directly correcting every faulted coordinate (Appendix D.2).

Memory and delay affect physical and task outcomes differently. On nineteen fixed-command sources, NT reduces median endpoint deviation $13.74 \to 7.80$ mm. Delay raises cost on $18/19$ sources; halving the cap on $14/19$. Known-fault correction leaves median endpoint deviations 0, 2.74, 5.72 mm with immediate correction, ten-step delay and half cap, respectively. Fitted persistence $\hat{\lambda}$ is 1.00 on Panda, $.66$ on ALOHA. Delayed/immediate task success is $13/40$ versus $14/40$: unresolved -2.5 points $[-10, 5]$ (Appendix D.3).

Repair does not require an onset signal in the tested condition. With adaptation from step 10 and a six-coordinate $+.10$ fault at hidden step 30, 50 or 70, registered Spatial Off/NT/DOB success is $13/25/24$ of forty (healthy: $40/40/38$). NT gains 30 points $[12.5, 47.5]$ (fourteen rescues, two regressions). Scheduled activation gives $6/40 \to 14/40$ (Off→NT) on the same inspected states but another seed and schedule; results are not pooled (Appendix D.5).

6 Discussion and limitations

Calibrated feedback repairs withheld command faults; crossed replay measures how direct correction combines with changed commands. Control analysis identifies limits from authority, memory and predictor mismatch. Predicting the changed-command response before deployment remains open.

Simulation coverage is limited; replay uses at most ten task clusters, with imperfect source duplication and exclusions. Joint locking scores $0/20$ in both arms; friction-shift gain is unresolved ($+7.5$ points $[-5, 21.25]$).

AI USE STATEMENT

Generative AI tools assisted with code development, literature discovery, evidence checking, experimental design, interpretation of results, and manuscript preparation, including writing up mathematical arguments and proofs developed by the authors. The authors independently checked cited sources, tested the implementation, and take full responsibility for the final manuscript and reported results.

REPRODUCIBILITY STATEMENT

The accompanying `submission_package.zip` contains manuscript sources, figures, registrations, verification reports and an extended evidence document. Its CPU-only artifact reconstructs translation costs, task-cluster uncertainty, retention and task counts for all assigned cases from saved trajectories and outcomes: run the evidence commands in `artifact/README.md` from the packaged manuscript directory. This recomputes saved replay measurements; it does not run a simulator or VLA inference. Source identities and metadata transformations are recorded. The evidence manifest identifies the repository's full raw artifacts; Appendices A–E give proofs, protocols, supporting results and image provenance.

REFERENCES

Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 17–40. PMLR, 2025a. URL `https://proceedings.mlr.press/v305/black25a.html`.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, Laura Smith, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A Vision-Language-Action Flow Model for General Robot Control. In *Proceedings of Robotics: Science and Systems*, 2025b. doi: 10.15607/RSS.2025.XXI.010. URL `https://www.roboticsproceedings.org/rss21/p010.html`.

Rémi Cadène, Quentin Gallouédec, Alexander Soare, and Simon Alibert. gym-aloha: A gym environment for ALOHA, 2024. URL `https://github.com/huggingface/gym-aloha`. Software; accessed September 2026.

Yikun Cheng, Pan Zhao, Fanxin Wang, Daniel J. Block, and Naira Hovakimyan. Improving the robustness of reinforcement learning policies with $\mathcal{L}_1$ adaptive control. *IEEE Robotics and Automation Letters*, 7(3):6574–6581, 2022. doi: 10.1109/lra.2022.3169309. URL `http://dx.doi.org/10.1109/LRA.2022.3169309`.

Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, Jinlan Fu, Jingjing Gong, and Xipeng Qiu. LIBERO-Plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025. URL `https://arxiv.org/abs/2510.13626`.

Fulvio Forni and Rodolphe Sepulchre. A differential Lyapunov framework for contraction analysis. *IEEE Transactions on Automatic Control*, 59(3):614–628, 2014. doi: 10.1109/TAC.2013.2285771. URL `https://arxiv.org/abs/1208.2943`.

Urban Forssell and Lennart Ljung. Closed-loop identification revisited. *Automatica*, 35(7):1215–1241, 1999. doi: 10.1016/S0005-1098(99)00022-9. URL `https://doi.org/10.1016/S0005-1098(99)00022-9`.

Lin Ge, Jitao Wang, Chengchun Shi, Zhenke Wu, and Rui Song. A reinforcement learning framework for dynamic mediation analysis. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 11050–11097. PMLR, 2023. URL https://proceedings.mlr.press/v202/ge23a.html.

Roman Geiselhart and Fabian R. Wirth. Relaxed ISS small-gain theorems for discrete-time systems. *SIAM Journal on Control and Optimization*, 54(2):423–449, 2016. doi: 10.1137/14097286X. URL https://arxiv.org/abs/1406.3224.

Michel Gevers. Towards a joint design of identification and control? In Harry L. Trentelman and Jan C. Willems (eds.), *Essays on Control: Perspectives in the Theory and its Applications*, pp. 111–151. Birkhäuser, 1993. URL https://perso.uclouvain.be/michel.gevers/PublisMig/ECC93.pdf.

Kevin Huang, Rwik Rana, Alexander Spitzer, Guanya Shi, and Byron Boots. DATT: Deep adaptive trajectory tracking for quadrotor control. In Jie Tan, Marc Toussaint, and Kourosh Darvish (eds.), *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pp. 326–340. PMLR, 2023. URL https://proceedings.mlr.press/v229/huang23a.html.

Petros A. Ioannou and Jing Sun. *Robust Adaptive Control*. Prentice-Hall, 1996.

Rolf Isermann. Model-based fault-detection and diagnosis – status and applications. *Annual Reviews in Control*, 29(1):71–85, 2005. doi: 10.1016/j.arcontrol.2004.12.002. URL http://dx.doi.org/10.1016/j.arcontrol.2004.12.002.

Minsoo Jo, Taeju Kwon, Junha Chun, Youngjoon Jeong, and Taesup Kim. Uncovering vulnerability of vision-language-action models under joint-level physical faults. *arXiv preprint arXiv:2606.10501*, 2026. URL https://arxiv.org/abs/2606.10501.

Tobias Johannink, Shikhar Bahl, Ashvin Nair, Jianlan Luo, Avinash Kumar, Matthias Loskyll, Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. Residual reinforcement learning for robot control. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 6023–6029. IEEE, 2019. doi: 10.1109/icra.2019.8794127. URL http://dx.doi.org/10.1109/ICRA.2019.8794127.

Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-Tuning Vision-Language-Action Models: Optimizing Speed and Success. In *Proceedings of Robotics: Science and Systems*, 2025a. doi: 10.15607/RSS.2025.XXI.017. URL https://www.roboticsproceedings.org/rss21/p017.html.

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. OpenVLA: An open-source vision-language-action model. In Pulkit Agrawal, Oliver Kroemer, and Wolfram Burgard (eds.), *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 2679–2713. PMLR, 2025b. URL https://proceedings.mlr.press/v270/kim25c.html.

Xuanlin Li, Kyle Hsu, Jiayuan Gu, Oier Mees, Karl Pertsch, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 3705–3728. PMLR, 2025. URL https://proceedings.mlr.press/v270/li25c.html.

Zongtan Li. When is a learned command adapter worth it? closed-loop identification and counterfactual auditing of frozen locomotion policies, 2026. URL https://arxiv.org/abs/2607.21867.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *Advances in Neural Information*

Processing Systems, volume 36, pp. 44776–44791, 2023. doi: 10.52202/075280-1939. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/8c3c666820ea055a77726d66fc7d447f-Abstract-Datasets_and_Benchmarks.html.

Elena Sorina Lupu, Fengze Xie, James Alan Preiss, Jedidiah Alindogan, Matthew Anderson, and Soon-Jo Chung. MAGIC-VFM: Meta-learning adaptation for ground interaction control with visual foundation models. *IEEE Transactions on Robotics*, 41:180–199, 2025. doi: 10.1109/TRO.2024.3475212.

Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. RoboCasa: Large-scale simulation of household tasks for generalist robots. In *Proceedings of Robotics: Science and Systems*, 2024. doi: 10.15607/RSS.2024.XX.050. URL https://www.roboticsproceedings.org/rss20/p050.html.

NVIDIA. GR00T N1.5: Open foundation model for generalist humanoid robots, 2025. URL https://research.nvidia.com/labs/gear/gr00t-n1_5/.

NVIDIA. GR00T-N1.7-SimplerEnv-Bridge model card, 2026a. URL https://huggingface.co/nvidia/GR00T-N1.7-SimplerEnv-Bridge. Accessed September 2026.

NVIDIA. Isaac GR00T N1.7 release, 2026b. URL https://github.com/NVIDIA/Isaac-GR00T/releases/tag/n1.7-release.

NVIDIA. RoboCasa GR-1 Tabletop Tasks, n.d. URL https://github.com/robocasa/robocasa-gr1-tabletop-tasks/tree/4840e671596f93ca03651524b9f72ffb1aadfeff. Simulation software; accessed September 2026.

NVIDIA, Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi "Jim" Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, Joel Jang, Zhenyu Jiang, Jan Kautz, Kaushil Kundalia, Lawrence Lao, Zhiqi Li, Zongyu Lin, Kevin Lin, Guilin Liu, Edith Llontop, Loic Magne, Ajay Mandlekar, Avnish Narayan, Soroush Nasiriany, Scott Reed, You Liang Tan, Guanzhi Wang, Zu Wang, Jing Wang, Qi Wang, Jiannan Xiang, Yuqi Xie, Yinzhen Xu, Zhenjia Xu, Seonghyeon Ye, Zhiding Yu, Ao Zhang, Hao Zhang, Yizhou Zhao, Ruijie Zheng, and Yuke Zhu. GR00T N1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025. URL https://arxiv.org/abs/2503.14734.

Judea Pearl. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pp. 411–420, 2001. URL https://arxiv.org/abs/1301.2300.

Yuanfang Peng, Jingjing Fu, Chuheng Zhang, Li Zhao, Jiang Bian, Mingyu Liu, Ling Zhang, Jun Zhang, and Rui Wang. What to ignore, what to react: Visually robust RL fine-tuning of VLA models. *arXiv preprint arXiv:2605.13105*, 2026. URL https://arxiv.org/abs/2605.13105.

Physical Intelligence. OpenPI: Action normalization and unnormalization transforms, n.d. URL https://github.com/Physical-Intelligence/openpi/blob/main/src/openpi/transforms.py. Source code, accessed September 2026.

ROBOTIS. DYNAMIXEL XM430-W350 e-manual: Homing offset, n.d. URL https://emanual.robotis.com/docs/en/dxl/x/xm430-w350/. Accessed September 2026.

Kohei Sendai, Maxime Alvarez, Tatsuya Matsushima, Yutaka Matsuo, and Yusuke Iwasawa. Leave no observation behind: Real-time correction for VLA action chunks, 2025. URL https://arxiv.org/abs/2509.23224.

Hyungbo Shim. Disturbance observers. In *Encyclopedia of Systems and Control*, pp. 1–8. Springer London, 2020. ISBN 9781447151029. doi: 10.1007/978-1-4471-5102-9_100068-1. URL http://dx.doi.org/10.1007/978-1-4471-5102-9_100068-1.

Tom Silver, Kelsey Allen, Josh Tenenbaum, and Leslie Kaelbling. Residual policy learning. *arXiv preprint arXiv:1812.06298*, 2018. doi: 10.48550/ARXIV.1812.06298. URL https://arxiv.org/abs/1812.06298.

Jean-Jacques E. Slotine and Weiping Li. Composite adaptive control of robot manipulators. *Automatica*, 25(4):509–519, 1989. doi: 10.1016/0005-1098(89)90094-0. URL http://dx.doi.org/10.1016/0005-1098(89)90094-0.

Dongwon Son, Florian Shkurti, Jason Lee, Naman Shah, Beomjoon Kim, and Dieter Fox. TAM: Torque adaptation module for robust motion transfer in manipulation, 2026. URL https://arxiv.org/abs/2606.06218.

Mahdi Taheri, Soon-Jo Chung, and Fred Y. Hadaegh. Closing the loop inside neural networks: Causality-guided layer adaptation for fault recovery control. *IEEE Control Systems Letters*, 9: 3137–3142, 2025. doi: 10.1109/LCSYS.2026.3652898. URL https://doi.org/10.1109/LCSYS.2026.3652898.

Mahdi Taheri, Haeyoon Han, Soon-Jo Chung, and Fred Y. Hadaegh. Geometric fault identification via mirror descent learning, 2026. URL https://arxiv.org/abs/2605.17103.

Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, 2012. doi: 10.1109/IROS.2012.6386109. URL https://homes.cs.washington.edu/~todorov/papers/TodorovIROS12.pdf.

Homer Rich Walke, Kevin Black, Tony Z. Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, Abraham Lee, Kuan Fang, Chelsea Finn, and Sergey Levine. BridgeData V2: A dataset for robot learning at scale. In *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pp. 1723–1736. PMLR, 2023. URL https://proceedings.mlr.press/v229/walke23a.html.

Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A SimulAted part-based interactive ENvironment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11097–11107, 2020. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Xiang_SAPIEN_A_SimulAted_Part-Based_Interactive_ENvironment_CVPR_2020_paper.html.

Mengfei Zhao, Chenxi Jiang, Tuo An, Jindou Jia, and Jianfei Yang. APEX: Adaptive policy execution for precise manipulation, 2026. URL https://arxiv.org/abs/2606.16504.

Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Proceedings of Robotics: Science and Systems*, 2023. doi: 10.15607/RSS.2023.XIX.016. URL https://www.roboticsproceedings.org/rss19/p016.html.

A DEFINITIONS AND PROOFS FOR THE EXECUTION-REPAIR FRAMEWORK

Command streams are time-indexed nominal outputs, including unchanged coordinates; *physical cost* is squared reference-trajectory deviation, not task success. Write $\|\cdot\|_2$ for Euclidean norm. For vectors u, v and compatible symmetric positive-definite Q, define $\langle u, v\rangle_Q = u^\top Q v$ and $\|u\|_Q = \sqrt{\langle u, u\rangle_Q}$.

A.1 RESIDUALS CONTAIN MODEL ERROR ON CORRECTED COMMANDS

Let $a_k, c_k, f_k \in \mathbb{R}^d$ denote nominal command, subtracted correction and additive fault at step k; the sent command is $u_k = a_k - c_k$. Let $b_g \in \mathbb{R}^d$ be a fixed response offset and $v_{k+1} \in \mathbb{R}^d$ collect omitted dynamics and approximation error. For memory length $K \geq 0$ and fixed $G_j \in \mathbb{R}^{d\times d}$, $j = 0, \ldots, K$, assume the measured response $y_{k+1} \in \mathbb{R}^d$ satisfies, on a stated domain,

$$y_{k+1} = \sum_{j=0}^{K} G_j(u_{k-j} + f_{k-j}) + b_g + v_{k+1}.$$

Pre-horizon commands and faults are specified. With fixed matrices W_j, offset b_h and the same history, the predictor is $\hat{y}_{k+1} = \sum_{j=0}^{K} W_j u_{k-j} + b_h$. For residual-model remainder $e_{k+1} = b_g - b_h + v_{k+1}$, subtraction yields

$$r_{k+1} := y_{k+1} - \hat{y}_{k+1} = \sum_{j=0}^{K} G_j f_{k-j} + \sum_{j=0}^{K} (G_j - W_j)(a_{k-j} - c_{k-j}) + e_{k+1}. \tag{11}$$

Model error depends on corrected commands, estimates and policy responses; the remainder need not be bounded or independent. Exact prediction still measures fault history $\sum_j G_j f_{k-j}$, not generally the current fault.

A.2 CONDITIONAL STATIONARY OBSERVER TARGETS

Proposition 2 (DC mismatch and observer stationarity)**.** *Use the DC gains G, W of Section 4, invertible sensitivity $M \in \mathbb{R}^{d\times d}$ and binary correction mask D_m. At constant nominal command $\bar{a} \in \mathbb{R}^d$ and fault $f \in \mathbb{R}^d$, assume $v_{k+1} = 0$ at every step, $b_h = b_g + (G - W)\bar{a}$, zero estimator bias, constant estimate history $h_* \in \mathbb{R}^d$ and inactive projection. DOB stationarity, or full-estimate innovation* (19) *stationarity with strictly active update gain, requires* (10)*. NT with constant attenuation $s \in [0, 1]$ instead requires*

$$[M + s(G - W)D_m]h_* = sGf. \tag{12}$$

For scalar $D_m = 1$, $M = G = g \neq 0$, $W = w$, $f \neq 0$, and nonzero denominators, stationary estimates $h_*^{\mathrm{DOB}}, h_*^{\mathrm{NT}}$ satisfy

$$\frac{h_*^{\mathrm{DOB}}}{f} = \frac{1}{2 - w/g}, \qquad \frac{h_*^{\mathrm{NT}}}{f} = \frac{s}{1 + s(1 - w/g)}.$$

Proof. For constant estimate h, the centered residual is $r(h) = Gf - (G - W)D_m h$. Stationarity sets $Mh = r(h)$ for DOB or active innovation, and $Mh = sr(h)$ for NT. Substitution and scalar division give the identities. □

An invertible bracket gives a unique candidate subject to its cap, gate and self-consistent NT attenuation. Closed innovation gates permit other held estimates. Invertibility does not imply stability: memoryless scalar DOB convergence requires $|1 - \alpha(2 - w/g)| < 1$; constant-$s$ NT requires $|1 - \gamma[1 + s(1 - w/g)]| < 1$; FIR dynamics require a history-state analysis. At $w = g$, NT still gives $h_*/f = s$. Finite-window probes need not measure the DC map G assumed here.

A.3 CONTROLLED REPLAY, RESPONSE GEOMETRY AND A SHARP TOLERANCE

Proof of Proposition 1. Since $D_0 = q_0^2 - q^2$, expanding $J_{11} = q^2 + 2A + r_\Delta^2$ proves (6). For $qr_\Delta > 0$, substitution of $A = qr_\Delta \cos\theta$ gives $R_1/q^2 = -g_{\mathrm{resp}}(g_{\mathrm{resp}} + 2\cos\theta)$. Cauchy–Schwarz

gives $\cos\theta \geq -1$, proving the sign conditions; at $g_{\text{resp}} = 2$, $R_1 \leq 0$ with equality only at $\cos\theta = -1$. If $q = 0$, $R_1 = -r_\Delta^2$; if $r_\Delta = 0$, $R_1 = 0$, without an angle.

The triangle inequality gives $\|E_{10} + \Delta\|_Q \leq q + \beta$ whenever $\|\Delta\|_Q \leq \beta$, proving (7). Equality holds at $\Delta = (\beta/q)E_{10}$ for $q > 0$, or any norm-β vector for $q = 0$, proving the exact infimum. Since $q_0, q, \beta \geq 0$, strict benefit throughout the ball is equivalent to $q_0 > q + \beta$, or $\beta < q_0 - q$. □

Faithful replay requires matched histories and randomness, with its observer rerun (Section 4). For fixed Euclidean trajectories P_{am} and common reference P^{ref}, write $E_{am} = P_{am} - P^{\text{ref}}$. With fixed weights, $\Delta = P_{11} - P_{10}$ and r_Δ are reference-independent, whereas q, A, g_{resp} are reference-relative. Exposing the reference as $R_1(P^{\text{ref}})$, any same-dimensional shift w gives $R_1(P^{\text{ref}} + w) = R_1(P^{\text{ref}}) + 2\langle\Delta, w\rangle_Q$. The shifted reference need not be healthy. Choosing $P^{\text{ref}} = P_{00}$ forces $T \leq 0$ but permits either sign of R_1.

The full ball can contain unattainable responses. Realized r_Δ is retrospective; prospective decisions require an independently justified β.

A.4 Quadratic interaction does not imply nonadditive trajectories

Proposition 3 (Cost interaction and trajectory nonadditivity)**.** *Let $P_{am} \in \mathbb{R}^{3N}$ stack translation trajectories, $P^{\text{ref}} \in \mathbb{R}^{3N}$ be their common reference, and $Q \in \mathbb{R}^{3N\times 3N}$ be fixed and positive definite. For costs $J_{am} = \|P_{am} - P^{\text{ref}}\|_Q^2$, define the reference deviation e, direct displacement u, stream displacement v and trajectory contrast Z:*

$$e = P_{00} - P^{\text{ref}}, \quad u = P_{10} - P_{00}, \quad v = P_{01} - P_{00}, \quad Z = P_{11} - P_{10} - P_{01} + P_{00}.$$

For $I = J_{10} + J_{01} - J_{11} - J_{00}$,

$$I = -2\langle u, v\rangle_Q - 2\langle e + u + v, Z\rangle_Q - \|Z\|_Q^2. \tag{13}$$

Z is reference-independent.

Proof. The deviations are e, $e + u$, $e + v$, $e + u + v + Z$; expanding their squared norms cancels all other terms. Z contains no reference. □

Additivity, $Z = 0$, means $P_{11} - P_{00} = (P_{10} - P_{00}) + (P_{01} - P_{00})$. For example, additive scalar endpoints $(P_{00}, P_{10}, P_{01}, P_{11}) = (0, 1, 1, 2)$ with zero reference and unit weight have $Z = 0$ but $I = -2$. Conversely, $(0, 1, 1, \sqrt{2})$ gives $I = 0$ and $Z \neq 0$. Z includes physical and observer responses without separating them. Reference independence concerns Euclidean translation, not arbitrary rotation coordinates.

A.5 A finite-horizon bound that includes the policy's response

Compare unadapted live execution 0, adapted replay R of M_0 and adapted live execution L, with common initial histories, matched sampling and faithful own-stream replay. For vectorized histories $h_{i,t}$ sufficient for policy/output bounds and nominal commands $a_{i,k}$, define history discrepancies r_t, d_t and command discrepancies v_k in fixed norms:

$$r_t = \|h_{L,t} - h_{R,t}\|, \quad d_t = \|h_{R,t} - h_{0,t}\| \ (1 \leq t \leq N), \qquad v_k = \|a_{L,k} - a_{0,k}\| \ (0 \leq k < N).$$

Stack them in time order as $r, d, v \in \mathbb{R}^N$; r_t is a history distance, distinct from residual (11). Assume nonnegative execution/policy coefficients $g_{t,k}^{\text{exec}}, \ell_{k,j}^{\text{pol}}$ and omitted-effect bounds ϵ_t, η_k satisfy, on a declared domain,

$$r_t \leq \sum_{k=0}^{t-1} g_{t,k}^{\text{exec}} v_k + \epsilon_t, \qquad v_k \leq \sum_{j=1}^{k} \ell_{k,j}^{\text{pol}}(r_j + d_j) + \eta_k, \tag{14}$$

Stack errors $\epsilon, \eta \in \mathbb{R}^N$ in the respective time orders of r, v. Define comparison matrices $\mathcal{G}, \mathcal{L} \in \mathbb{R}^{N\times N}$ by $\mathcal{G}_{t,k+1} = g_{t,k}^{\text{exec}}$ for $0 \leq k < t \leq N$ and $\mathcal{L}_{k+1,j} = \ell_{k,j}^{\text{pol}}$ for $1 \leq j \leq k < N$, with other entries zero. Then $\mathcal{K} = \mathcal{G}\mathcal{L}$ is strictly lower triangular. Set the combined error vector $\varepsilon = \epsilon + \mathcal{G}\eta$.

Proposition 4 (Finite-horizon policy-response bound)**.** *Assume* (14)*, and let* $\bar{d} \in \mathbb{R}^N$ *bound* d*:* $0 \preceq d \preceq \bar{d}$*, where* $\preceq$ *denotes componentwise order. Then*

$$r \preceq \bar{r} := \sum_{j=0}^{N-1} \mathcal{K}^j(\mathcal{K}\bar{d} + \varepsilon). \tag{15}$$

Local premises apply while histories remain in domain; a predicted history enclosure contained there suffices through N*.*

Proof. Substitution gives $r \preceq \mathcal{K}r + \mathcal{K}\bar{d} + \varepsilon$. Iterate this order-preserving inequality N times and use $\mathcal{K}^N = 0$. Local validity follows by induction within the enclosure. □

No contraction is required or asymptotic stability asserted. If output sensitivity bounds $w_t \geq 0$ satisfy $\|e_{L,t} - e_{R,t}\|_{Q_t} \leq w_t r_t$, then

$$\|E_{11} - E_{10}\|_Q \leq \beta := \left(\sum_{t=1}^{N} w_t^2 \bar{r}_t^2 \right)^{1/2}.$$

Squaring and summing proves this inequality; Proposition 1 bounds benefit. Prospective use requires jointly valid gains, errors, domains and D_0, q bounds, unqualified here.

A.6 Exact command-space authority

Lemma 1 (Mask and cap distance)**.** *For the command-space identity* I_{cmd}*, diagonal binary mask* D_m*, centered box* $\mathcal{C} = \{h : |h_i| \leq c_i^{\max}\}$*,* $c_i^{\max} \geq 0$*, and any* f*,* (8) *holds, attained at* $h = \Pi_{\mathcal{C}} f$*. Exact cancellation is possible precisely when* $(I_{\mathrm{cmd}} - D_m)f = 0$ *and* $D_m f \in \mathcal{C}$*.*

Proof. On unselected coordinates the error is f_i; on selected coordinates its smallest magnitude is $(|f_i| - c_i^{\max})_+$. Squaring and summing proves the formula and its zero set. □

For a linear physical map G on the declared command coordinates, exact motion cancellation instead requires $Gf \in G(D_m\mathcal{C})$; task recovery has no such equivalence.

A.7 Exact scalar limit from memory, authority and information delay

Proof of Theorem 1. Write $S_j = S_j(\lambda)$ for the geometric sums. Unrolling yields $p_N = g[fS_N - \sum_{k=0}^{N-1} \lambda^{N-1-k} c_k]$. Before τ, identical information forces identical corrections for both signs. Denote their common weighted sum by $b = \sum_{k<\tau} \lambda^{N-1-k} c_k$. The remaining weighted correction can take every value in $[-D, D]$, where $D = CS_{N-\tau}$, since its nonnegative weights sum to $S_{N-\tau}$ and the component intervals add. For $A = FS_N \geq 0$, the optimal worst-case remainder at fixed b is

$$g \max\{[|A - b| - D]_+, [|-A - b| - D]_+\} = g[A + |b| - D]_+ \geq g[A - D]_+.$$

Here $\max\{|A - b|, |-A - b|\} = A + |b|$. Zero correction before τ gives $b = 0$. If $N > \tau$, constant post-delay correction with the fault sign and magnitude $\min\{C, A/S_{N-\tau}\}$ attains the lower bound. If $N = \tau$, no such correction is available and zero correction attains gA. This includes $F = 0$ and $C = 0$. □

At $\lambda = 1$ and $C = F$, the optimum is $gF\tau$. At $0 \leq \lambda < 1$ and $C = F$, it is $gF\lambda^{N-\tau} S_\tau$. The unknown sign belongs to the additive fault; control direction is fixed by $g > 0$. Zero terminal displacement implies neither zero accumulated path cost nor task success.

Corollary 1 (Authority required for a terminal tolerance)**.** *Under Theorem 1, let* $N > \tau$ *and let* $\epsilon \geq 0$ *be an allowed terminal displacement. An admissible deterministic correction policy can guarantee* $|p_N| \leq \epsilon$ *for both fault signs if and only if*

$$C \geq \frac{[FS_N(\lambda) - \epsilon/g]_+}{S_{N-\tau}(\lambda)}. \tag{16}$$

In particular, for $\lambda = 1$ *and* $\epsilon = 0$*, this is* $C \geq FN/(N - \tau)$*.*

Proof. Since $N > \tau$ and $\lambda \geq 0$, $S_{N-\tau}(\lambda) \geq 1$. Attainment in (9) makes ϵ feasible exactly when $FS_N - CS_{N-\tau} \leq \epsilon/g$. Solving for $C \geq 0$ gives (16). For $\lambda = 1$, $S_m = m$. □

The threshold assumes exact scalar dynamics, known fault magnitude, exact sign from τ and otherwise unconstrained correction. It bounds terminal displacement, not deployed-observer performance, intermediate safety or task success. With $N = \tau$, no post-information control remains: feasibility requires exactly $gFS_N \leq \epsilon$, regardless of C.

Proposition 5 (Scalar DOB performance under exact fault information)**.** *Under Theorem 1, let $f = \sigma F$, $\sigma \in \{-1, 1\}$, with F known. An idealized DOB with gain $0 < \alpha \leq 1$ initializes $c_{-1} = 0$, applies $c_k = 0$ for $k < \tau$, and, given the exact target before command $k \geq \tau$, uses projection (clipping) to $[-C, C]$:*

$$c_k = \Pi_{[-C,C]}[(1-\alpha)c_{k-1} + \alpha f].$$

For $m = N - \tau$ remaining commands, define correction magnitudes $b_j = \min\{C, F[1-(1-\alpha)^{j+1}]\}$ for $j = 0, \ldots, m-1$, and $U_\alpha = \sum_{j=0}^{m-1} \lambda^{m-1-j} b_j$ (zero if $m = 0$). Its worst-case terminal displacement E_N^α and gap to the ideal $E_N^\star$ from (9) *are*

$$\begin{aligned} E_N^\alpha &= g(FS_N - U_\alpha), \\ E_N^\alpha - E_N^\star &= g[\min\{FS_N, CS_m\} - U_\alpha] \geq 0. \end{aligned} \tag{17}$$

Proof. Induction gives $c_{\tau+j} = \sigma b_j$: the unclipped average approaches f monotonically, and a reached cap remains active. Unrolling p_N gives the first identity, since $0 \leq b_j \leq F$. Moreover, $U_\alpha \leq \min\{CS_m, FS_m\} \leq \min\{CS_m, FS_N\}$. Subtracting the ideal value proves the gap, including $m = 0$, $F = 0$ or $C = 0$. □

With $\alpha = 1$ and $C \leq F$, the gap is zero. Here τ is first exact information, not correction-threshold latency: measurements after command k first affect command $k + 1$.

Cancelling increments versus recovering displacement. For $\lambda = 1$, the positive fault $f = F > 0$, $\tau > 0$ and $C \geq F$, zero correction before τ followed by $c_k = F$ gives $p_N = gF\tau$ for $N \geq \tau$. For $m \geq 1$ additional steps, summation shows that a continuation returns $p_{\tau+m}$ to zero exactly when $\sum_{k=\tau}^{\tau+m-1}(c_k - F) = F\tau$. Thus $C = F$ cannot recover the offset, whereas $C > F$ permits recovery in $m_{\min} = \lceil F\tau/(C-F) \rceil$ steps, rounding up: use the cap, adjust the last step, then cancel F. With $C < F$, even cancellation is infeasible.

With known F, observable reference error p_k and feedback gain $\kappa > 0$, $c_k = F + \kappa p_k$ gives $p_{k+1} = (1 - g\kappa)p_k$, converging for $0 < g\kappa < 2$ while the cap remains inactive; $C > F + \kappa|p_\tau|$ suffices. For $F = g = 1$, $\tau = 1$, $C = 2$ and $\kappa = 1/2$, $p_1 = 1$ then halves each step with $c_k \leq 1.5$. This uses additional reference-state feedback without integrating position error. For $0 \leq \lambda < 1$, exact cancellation dissipates the offset; a constant uncancelled command error f_{res} instead leaves equilibrium displacement $p_* = gf_{\text{res}}/(1-\lambda)$.

A.8 NT ESTIMATION RELATIVE TO ITS CAP

Proposition 6 (NT error relative to the capped fault)**.** *Let $\mathcal{C} = \{v \in \mathbb{R}^d : |v_i| \leq c_i^{\max}\}$, with fixed $c_i^{\max} \geq 0$, and let $\Pi_\mathcal{C}$ be Euclidean projection onto this box. For any fault sequence f_k, define observation error $w_{k+1} = z_{k+1} - f_k$. Let $s_k \in [0, 1]$ denote the actual NT attenuation in* (3)*, including zero when its gate closes. For gain $0 < \gamma \leq 1$, suppose $\hat{f}_{k+1} = \Pi_\mathcal{C}[(1-\gamma)\hat{f}_k + \gamma s_k z_{k+1}]$. Define capped targets $\vartheta_k = \Pi_\mathcal{C}(f_k)$ and $\vartheta_k^s = \Pi_\mathcal{C}(s_k f_k)$, and their estimation error $e_k^\mathcal{C} = \|\hat{f}_k - \vartheta_k\|_2$. Then*

$$e_{k+1}^\mathcal{C} \leq (1-\gamma)e_k^\mathcal{C} + \gamma s_k \|w_{k+1}\|_2 + \gamma\|\vartheta_k - \vartheta_k^s\|_2 + \|\vartheta_{k+1} - \vartheta_k\|_2, \tag{18}$$

and $\|\vartheta_{k+1} - \vartheta_k\|_2 \leq \|f_{k+1} - f_k\|_2$.

Proof. Coordinatewise clipping gives $\Pi_\mathcal{C}[(1-\gamma)\vartheta_k + \gamma s_k f_k] = (1-\gamma)\vartheta_k + \gamma\vartheta_k^s$. If $f_{k,i}$ is unsaturated, so is $s_k f_{k,i}$. Otherwise $\vartheta_{k,i}$ is the signed endpoint: both sides clip there if $s_k f_{k,i}$ is outside, and interpolate inside otherwise, including zero caps. Projection nonexpansiveness bounds the update's distance to this reference by $(1-\gamma)e_k^\mathcal{C} + \gamma s_k\|w_{k+1}\|_2$. Its displacement from ϑ_k is

$\gamma(\vartheta_k^s - \vartheta_k)$. The triangle inequality and change to ϑ_{k+1} give the bound; nonexpansiveness on f_k, f_{k+1} proves the final inequality. □

The bound separates observation error, attenuation bias and fault variation beyond the cap. Applied correction $D_m\hat{f}_k$ additionally depends on the mask. The remainder w_{k+1} can depend on robot, policy and earlier corrections; its definition implies neither independence nor a uniform bound. Bounded estimation error does not guarantee better execution or task success. A small selected residual $\|J_g r_{k+1}\|_2$ need not bound the full innovation $\|z_{k+1} - \hat{f}_k\|_2$: J_g may have a nontrivial kernel, and the sensitivity map can amplify components absent from the gate.

B Experimental interfaces, calibration and task protocols

B.1 Interfaces and adaptation state

Simulated tasks use benchmark success criteria and episode limits. Panda uses LIBERO's operational-space pose controller at 20 Hz. One action unit requests .05 m or .5 rad, rather than guaranteeing that displacement; a $+.05$ offset adds a 2.5 mm/.025 rad target per affected coordinate. Measured increments use these scales and relative rotation in the three-dimensional rotation group $\mathrm{SO}(3)$; the gripper is unchanged. $\pi_{0.5}$/GR00T execute five/eight actions between replans. The observer updates each step, preserving command history across replans.

Policy inputs and observer measurements. The $\pi_{0.5}$-LIBERO configuration disables continuous state tokens, state tokenization in the prompt and state-relative action conversion. It conditions on images and instruction despite receiving state from the client; the observer separately measures motion. OpenVLA-OFT enables a proprioceptive projector; GR00T supplies state with external/wrist images. Packaged implementation snapshots establish these pathways. Holding GR00T's state input constant (Appendix C.8) tests effect persistence, not each observation channel's contribution.

ALOHA accepts twelve arm angles and two normalized gripper targets at 50 Hz; GR1 accepts twenty-nine joint targets at 20 Hz. Their predictors observe joint position, with ten/sixteen actions between replans and 300/720-step limits. WidowX uses SimplerEnv Cartesian increments, eight-action replanning and a 150-step limit in its held-estimate evaluation. Its predictor observes position and wrapped-angle increments; correction affects translation only.

Table 6: **Adapted state and interface-specific constants.** Stored/updated/applied count scalar coordinates eligible during identification; held evaluations make no estimator updates. The gate threshold δ, normalization scale ρ and coordinatewise cap use each interface's calibrated units. All use $\gamma = .08$; the unseen-fault study (Appendix C.6) selects $\gamma = .16$ with cap .15.

Interface	Stored/updated/applied	δ	ρ	Cap	Evaluation
Panda, headline	6/6/3	.008	.15	.15	Online from zero
Panda, recent command faults	6/6/3	.008	.15	.05	Online from zero
ALOHA	14/14/6	.002	.4	.08	Identify 1, then hold
GR1	29/29/7	.013	.11	.2	Identify 3, then hold
WidowX	6/3/3	.001	.009	.03	Identify 3, then hold

Writing superscript I for innovation, the alternative to NT is

$$\hat{f}^{\mathrm{I}}_{k+1} = \Pi_{\mathcal{C}}\left[\hat{f}_k + \gamma s(\|J_g\xi_{k+1}\|_2)M^{\dagger}\xi_{k+1}\right], \qquad \xi_{k+1} = r_{k+1} - M(\hat{f}_k + b). \tag{19}$$

For invertible M, the unprojected increment is $\gamma s(\|J_g\xi_{k+1}\|_2)(z_{k+1} - \hat{f}_k)$. A closed gate makes NT decay toward zero and innovation hold. Both estimate command-equivalent offsets, not full physical state; neither integrates a physical tracking error. Panda applies rotational coordinates; ALOHA applies left-arm joints 0–5, GR1 right-arm joints 7–13, and WidowX translation coordinates 0–2. Panda and ALOHA use all measured channels in the residual norm. GR1 and WidowX use corrected channels and form innovation from the applied estimate, $\xi_m = r - MD_m\hat{f}$, with zero bias. GR1's increment is $\gamma s(\|J_g\xi_m\|_2)M^{\dagger}\xi_m$; WidowX additionally masks that increment by D_m and masks its NT increment. Controllers therefore differ across robots.

Panda stores 42 predictor coefficients, six intercepts, 36 sensitivity entries, a pseudoinverse and 42 command-history values besides its six adapted estimates. No VLA weights change; this is not a speedup measurement. Gain compensation can update six gains and six recursive-least-squares (RLS) covariances; historical files leave some RLS/least-mean-squares launch variants unresolved.

Recorded computation time. Twenty healthy Spatial episodes with disjoint calibration give adapter median/99th-percentile $0.174/0.266$ ms (2,073 steps) and policy-call median 894.668 ms (425 five-action replans). Adapter timing covers correction, motion processing, FIR and estimation; policy timing includes inference transport. Timing is not Off-subtracted; exact hardware/load are unrecorded beyond a Turing graphics processor. Ten simulator warmup steps are excluded, without filtering inference warmup. Real-time 20 Hz execution remains unestablished.

B.2 Calibration, evaluation splits and uncertainty

Reference Panda calibration uses three healthy Spatial episodes (285 steps). Per-channel ridge fits use current and six lagged commands with an intercept, penalty 10^{-2} including the intercept, and episode-bounded windows. Online history starts at zero, excluding simulator warmup commands. The sensitivity matrix is measured with controlled offset probes on replayed commands. The reference probe uses initial state 45; two Spatial evaluation keys also overlap the predictor log. The disjoint alternative uses ten healthy episodes starting at state 25 and sensitivity probes there, separate from evaluation states 45–48. Both retain the development-selected correction mask and gains. Four law-constant settings were tried on the headline fault before fixation. Development overlaps fifteen of twenty historical Spatial keys (tasks 0–9/state 45 and tasks 0–4/state 46); faulted comparisons also selected the correction mask. Thus neither historical evaluation holds out method selection, despite disjoint predictor/probe calibration. Later evaluations specify development-state reuse.

Other robots are recalibrated: eight healthy trajectories for ALOHA, six/ten for GR1 plate-to-plate/tray-to-plate, and ten for WidowX. ALOHA/GR1 use ordinary least squares; GR1's held-pose sensitivity is approximately .99 times the identity on the arm block. WidowX uses its predictor response after an inaccurate held-pose z probe. Calibration is local: one Panda configuration change alters translation-z/rotation-y sensitivity by approximately $107\%/-11\%$.

Historical sampling is unpinned; matching prefixes cannot establish future randomness identity. Task analyses retain all keys; only physical replay excludes sources.

Historical headline uncertainty uses 20,000 bootstrap resamples of forty (suite, task) clusters, retaining their states, without suite stratification. Net faulted-gain intervals are $[36, 65]$ episodes for reference calibration and $[20, 49]$ for disjoint calibration. Recent intervals generally use 10,000 whole-task draws and registered seeds. Intervals are marginal 95%, not simultaneous; McNemar tests are descriptive under task dependence. Unresolved contrasts establish neither equivalence nor population noninferiority.

Registered state holdout and second-seed replication. The registered method fixes disjoint calibration, NT $(\gamma, \delta, \rho) = (.08, .008, .15)$, rotation support, cap .15, five-action replanning, and uniform $+.05$ fault/adaptation from the first policy step, with the estimate reset to zero. A precollection repository/ledger scan selects Spatial state 4; Object 0/1; Goal 0–3; LIBERO-10 4–7. The 10/20/40/40 allocation covers 110 task/state pairs across forty tasks. Blocks A/B use sampler seeds 102001/102002; each suite's four arms share one process within a block. B was registered after A's outcome and before B's collection, with every method setting unchanged. Reusing the states gives 220 task/state/seed evaluations. All 880 episodes complete without exclusions or recorded retries.

Paired differences weight evaluations equally. Intervals use 10,000 unstratified resamples of forty (suite, task) clusters, retaining all states and both seeds within each sampled cluster (bootstrap seed 20260926; bootstrap seeds throughout are eight-digit integers chosen for uniqueness and denote neither collection nor registration dates). Faulted gains are $+30.0$ points $[17.3, 43.3]$ in A, $+27.3$ $[13.6, 41.5]$ in B and $+28.6$ $[16.2, 41.5]$ pooled (Table 2). The pooled healthy gain, $+1.8$ points $[-1.6, 5.5]$, does not establish noninferiority. Healthy (task, state) regressions are Goal $(3, 0)$ and LIBERO-10 $(8, 6)$ in A; Goal $(9, 2)$ and LIBERO-10 $(8, 5), (9, 4)$ in B. Faulted regressions are Goal $(5, 0)$ and Object $(0, 0), (7, 1), (8, 0)$ in A; Goal $(5, 0), (5, 1), (5, 2)$ and Object $(8, 0), (9, 1)$ in B.

Table 7: Registered suite outcomes: Off → NT (rescues/regressions), with n task/state pairs per block.

Suite	n	Faulted A	Faulted B	Healthy A	Healthy B
Spatial	10	7 → 10 (3/0)	4 → 10 (6/0)	10 → 10 (0/0)	10 → 10 (0/0)
Object	20	7 → 13 (9/3)	7 → 14 (9/2)	19 → 20 (1/0)	19 → 20 (1/0)
Goal	40	15 → 27 (13/1)	18 → 26 (11/3)	40 → 39 (0/1)	39 → 39 (1/1)
LIBERO-10	40	0 → 12 (12/0)	0 → 9 (9/0)	38 → 39 (2/1)	34 → 36 (4/2)

Registered minimum-gain and healthy-net criteria pass in both blocks; all suite point gains are positive (Table 7). A's equal 30-point suite gains contradict the predicted smallest gain on LIBERO-10. B's Goal interval crosses zero ($[-10, 47.5]$ points); Object's touches zero ($[0, 65]$). The packaged artifact reproduces outcomes, intervals and configurations.

Historical calibration evaluations. Disjoint/reference faulted suite totals are Spatial $9/20 \rightarrow 18/20$ / $8/20 \rightarrow 18/20$, Object $7/20 \rightarrow 15/20$ / $5/20 \rightarrow 16/20$, Goal $18/40 \rightarrow 28/40$ / $15/40 \rightarrow 29/40$, and LIBERO-10 $0/40 \rightarrow 7/40$ / $0/40 \rightarrow 15/40$. Reference healthy success is $19/20 \rightarrow 20/20$ on Spatial, $20/20 \rightarrow 20/20$ on Object, $39/40 \rightarrow 40/40$ on Goal and $38/40 \rightarrow 37/40$ on LIBERO-10: total $116/120 \rightarrow 117/120$, with four rescues and three regressions. The lost LIBERO-10 (task, state) keys are $(9, 45), (0, 47), (9, 48)$. Disjoint healthy success is $20/20 \rightarrow 20/20$, $20/20 \rightarrow 20/20$, $40/40 \rightarrow 39/40$ and $38/40 \rightarrow 39/40$: total $118/120 \rightarrow 118/120$, with two rescues and two regressions, at Goal $(5, 48)$ and LIBERO-10 $(9, 47)$. All indices are zero-based. Different frozen outcomes prevent interpreting the two calibration evaluations as a paired calibration intervention.

B.3 Transfer and identification before held execution

Panda transfer reuses the predictor and sensitivity with each backbone's replanning schedule. GR00T N1.7 uses the LIBERO checkpoint on Panda and the SimplerEnv-Bridge checkpoint (NVIDIA, 2026a) on WidowX, post-trained on BridgeData V2 (Walke et al., 2023). OpenVLA-OFT Object gives $0/20 \rightarrow 10/20, 15/20, 19/20$ for rotation, translation and gain faults. Healthy Spatial rotation controls give $20/20 \rightarrow 19/20$ for OpenVLA-OFT and $18/20 \rightarrow 18/20$ for GR00T (one rescue, one regression). Gain compensation is distinct.

Other robots identify on preceding faulted episodes, then reset and hold estimates. Every arm excludes identification episodes. ALOHA's .02 rad left-arm result removes one episode from forty: healthy Off/faulted Off/corrected are $17/0/15$ of 39. GR1 removes three from thirty, giving $20/1/17$ of 27 for plate-to-plate and $15/1/13$ for tray-to-plate under .10 rad right-arm offsets. Its held vector is the median of estimates over reach steps 50–200 in the three identification episodes. WidowX removes three from 23, giving $14/0/14$ of twenty under a .005 translation offset; held estimates have median $(.0050, .0051, .0049)$. These protocols do not establish single-episode online repair across robots.

Healthy controls qualify the transfer claims. ALOHA identifies in its first healthy episode, then holds that estimate: $17/39 \rightarrow 13/39$, four rescues and eight regressions. GR1's separate healthy controls identify for one episode, then hold for 29: plate-to-plate gives $18/29 \rightarrow 18/29$ (six rescues, six regressions), and tray-to-plate $14/29 \rightarrow 15/29$ (seven, six). WidowX's healthy control instead updates online from zero and gives $15/20 \rightarrow 16/20$ (three, two). GR1/WidowX healthy protocols thus differ from their three-episode faulted identification.

With a matched identified vector reset before each episode, hold/innovation/NT score $15/14/0$ of forty ALOHA tasks and $22/17/1$ of thirty GR1 plate-to-plate tasks. Hold versus innovation has discordances $4/3$ and $8/3$; hold versus NT has $15/0$ and $21/0$. NT attenuation can erase useful initialized estimates. WidowX's continuous, faster-update and held variants score $6/20, 12/20, 14/20$ on different episode allocations, so they are not a matched adaptation-speed ablation.

Base OpenVLA feasibility screen. Base OpenVLA-7B (Kim et al., 2025b) on SimplerEnv Google-robot move-near gives healthy $2/10$ and translation-x predictor coefficient of determination $R^2 = .57$. Its healthy/faulted screens omit adaptive repair and differ from OpenVLA-OFT (Table 1).

B.4 Observer controls and scheduled fault recovery

With a fixed policy-sampling schedule, FIR ($K = 6$) versus static ($K = 0$) prediction gives $10/80$ versus $9/80$ on LIBERO-10 states 42–49 (five versus four discordances). Calibrated feedback versus raw integral correction gives $18/20$ versus $3/20$ on Spatial states 45–46 under a .15 translation offset (fifteen versus zero discordances). The latter integrates $y - u$ with gain .005, without the calibrated predictor/map; it tests that package, not FIR memory alone. Frozen command traces match over shared recorded steps within 10^{-6} on 61/80 and 17/20 pairs, respectively, despite equal frozen success outcomes. Among observers, Kalman uses covariance-weighted residual updates; composite adds fitted-reference tracking, and calibrated integral accumulates z.

Table 8: **ALOHA observer confirmation.** Each observer uses thirty independent confirmation seeds with one healthy and seven fault conditions. Balanced success weights healthy success and mean faulted success equally. Shared Off/known-fault oracle controls are not candidate observers.

Observer	Healthy /30	Faulted /210	Balanced
NT	16	98	50.0%
RLS	17	63	43.3%
DOB	14	79	42.1%
Kalman	12	71	36.9%
Composite	12	67	36.0%
Calibrated integral	11	57	31.9%
Off / oracle	15 / 13	48 / 103	36.4% / 46.2%

The ALOHA study uses 2,608 search, 896 validation and 1,920 confirmation rollouts with separate fitting/search/validation/confirmation seed sets. No observer contrast survives the reported multiplicity correction. Composite minus Kalman is $-.95$ points $[-9.52, 7.62]$ under 50,000 whole-seed resamples; the whole-seed sign-flip p-value is .86261. A separate twenty-episode Panda rotation-.10 comparison gives NT/DOB/RLS/Kalman/integral/oracle $19/17/17/18/15/20$ successes. Neither study establishes an observer ranking.

Scheduled stronger fault. Development on Spatial states 15/17, seed 84001 selects uniform $+.10$ (Off $5/20$, NT $7/20$) after a pre-evaluation amendment adds two stronger faults; three initial candidates missed the registered Off-success threshold. Evaluation uses states 18/19/21/22, seed 85001, the historical calibration, rotation cap .05 and Table 6's constants. Fault and adaptation begin at policy step 30 after ten simulator warmup steps; the first update affects command 31. Episodes end on success or after 220 policy steps. All forty keys reach activation. Off/NT give $6/40 \to 14/40$ under fault (nine rescues, one regression at $(2, 19)$), or $+20$ points $[10, 30]$, with healthy $40/40$ in both arms. The 10,000 whole-task bootstrap uses seed 20260916 and the registered contrast order. Scheduled enablement does not test unknown-onset detection.

Matched DOB comparison. Six new forty-key arms reuse these inspected keys, sharing one process per healthy/faulted Off–NT–DOB trio. DOB uses the projected smoother (Section 3.3) without NT's gate/attenuation. Development gains $.02, .08, .20$ give healthy/faulted totals $20/6, 20/8, 19/9$ of twenty each; maximizing combined success and breaking ties by fewer healthy losses selects .08. Observers had unequal gain-selection budgets. Evaluation Off/NT/DOB are $6/14/14$ of forty under fault and $39/40/40$ under health. NT–DOB is 0 points $[-7.5, 7.5]$, with one success unique to each: NT at $(5, 18)$ and DOB at $(2, 19)$. DOB–Off is $+20$ points $[10, 32.5]$ (eight rescues, no loss); NT–Off is $+20$ $[10, 30]$ (nine, one). Each adapter rescues healthy $(9, 19)$ without losing a key. NT/DOB pre-fault prefixes match on 40/40 keys; each matches faulted Off on 36/40. All keys remain in the task analysis. Equal totals leave the observer comparison unresolved; they do not establish equivalence.

C Crossed replay: protocol, replication and interpretation

C.1 What is replayed and what is measured?

Sources are live rollouts identified by task, state and sampler seed; seven-dimensional streams include the gripper. Crossed cells (Figure 2) share checkpoint, fault and reference; observers update from

replay measurements. The reference executes the same key's healthy-Off stream without fault or correction. For adapter setting $a \in \{0,1\}$ and stream index $m \in \{0,1\}$, p_k^{am} and p_k^{ref} are the measured and reference positions at step k. Translation cost is

$$J_{am} = .05 \sum_{k=1}^{50} \|p_k^{am} - p_k^{\mathrm{ref}}\|_2^2 \quad [\mathrm{m}^2\,\mathrm{s}]. \tag{20}$$

Angular cost replaces the position difference by the relative axis–angle vector, whose norm is the rotation angle in $[0, \pi]$, and has units $\mathrm{rad}^2\,\mathrm{s}$; subscript r marks angular contrasts. Positive D_0, R_1 or T means reduced cost.

Timing and controller state. Ten warm-up steps and thirty policy commands precede the checkpoint (environment step 40/policy step 30); fifty subsequent commands span $2.5\,\mathrm{s}$. FIR history holds the seven most recent sent commands; estimates start at zero. Each step corrects using the current estimate, measures response, then updates. Sources use five-action $\pi_{0.5}$ chunks and eight-action GR00T chunks. At step 30, GR00T is inside its step-24 chunk, first permits nonzero correction at step 31 and next replans at step 32. The registered delayed archival control suppresses updates and correction for ten window steps. The stronger fault is a uniform $+.10$ six-coordinate offset, whereas the archival fault is $+.05$ on rotation-y. Both use rotation correction, cap $.05$, gain $.08$, and the fixed Panda calibration; NT uses deadzone $.008$ and normalization scale $.15$.

Eligibility and fidelity. Eligibility requires eighty commands per arm and exact action/position/joint agreement throughout the thirty-step preactivation prefix. Each key supplies seven continuations: reference, four cells and fresh-reset diagonal repeats. Fidelity compares diagonals with sources and fresh resets with checkpoint restoration. Tolerances are $10^{-8}\,\mathrm{m}$ for position, 10^{-8} rad for joints and diagonal orientation, and 10^{-8} in command units for correction. All four cells must complete the fifty-step window. Success flags neither stop continuation nor add trials. Excluded keys are not replaced; task denominators retain all assigned keys.

C.2 Evaluation groups, source reproducibility and inference

Original and matched stronger-fault $\pi_{0.5}$ collections share keys but use separately collected streams; their costs are not pooled (Table 9). GR00T uses inspected Spatial states.

Table 9: **Assigned keys and physical retention.** All allocations cover ten tasks. "Eligible" applies source length and prefix rules; "valid" additionally applies replay fidelity. A full Off duplicate repeats a live episode's key, checking its entire saved trajectory. A dash denotes an unallocated duplicate, not perfect matching.

Evaluation	States; sampler seed	Assigned	Eligible	Valid/tasks	Off duplicate
Spatial archive, fault	13,14; 83001	20	18	15/9	—
Spatial archive, healthy	13,14; 83001	20	18	16/9	—
Spatial replication	23,24,29,31; 86001	40	35	34/9	36/40
LIBERO-10, healthy	30–33; 87001	40	36	36/10	31/40
GR00T, healthy	23,24,29,31; 89001	40	36	36/10	35/40
$\pi_{0.5}$ strong, original	18,19,21,22; 85001	40	32	30/10	—
$\pi_{0.5}$ strong, matched	18,19,21,22; 85001	40	29	27/10	—
GR00T strong, NT / DOB	23,24,29,31; 98001	40	34	30/8; 29/8	—

What matching establishes. Off–NT prefixes match on $39/40$, $36/40$ and $38/40$ keys in the independent Spatial, LIBERO-10 and healthy GR00T collections. Their full Off duplicates match less often (Table 9). Separate-process Spatial attempts lacked matching prefixes and are excluded. A GR00T repeat changes process and inference device together: all prefixes but only $35/40$ complete trajectories match, with gaps up to $8.76\,\mathrm{mm}$. Thus seeds/resets do not ensure identical sources or isolate hardware effects. Attribution concerns realized stream replacement. Checks cover recorded robot fields, not every object state; off-diagonals lack independent fresh repeats.

Statistical units and registered endpoints. Equal-task means average retained states within tasks, then tasks. Task-cluster bootstraps resample whole tasks with their states and paired cells (10,000 draws; pointwise 95% intervals). Sensitivities weight sources equally or omit one task. Bootstrap seeds are 20260916 for archival/independent Spatial, 20260918 for stronger $\pi_{0.5}$/LIBERO-10, 20260920 for healthy GR00T, and 20260923 for faulted GR00T; these values do not establish collection dates. The archival primary is interaction I; healthy replications and original stronger $\pi_{0.5}$ register R_1; matched observers register paired translation $T_{\mathrm{NT}} - T_{\mathrm{DOB}}$; faulted GR00T registers angular NT T_r. Translation decisions use margin 10^{-5} m^2 s where specified, equivalent to constant 2 mm error over 2.5 s. Secondary intervals are unadjusted for multiplicity.

C.3 Healthy response: replicated cost, distinct task outcomes

Healthy-Off reference implies $J_{00} = 0$ and $D_0 = -J_{10} \leq 0$, but not R_1's sign; its mean is negative in all four initial evaluation groups.

$$S = D_0 - R_1 = J_{11} - 2J_{10} \tag{21}$$

tests whether added stream cost exceeds direct cost. This analysis is retrospective for Spatial and registered for LIBERO-10/GR00T. Every evaluation group's interval is positive; sources can disagree (Table 10).

Table 10: **Healthy translation response.** Equal-task means and 95% task-bootstrap intervals, in 10^{-4} m^2 s. Negative R_1 means added cost; positive S means that increment exceeds direct cost. Task counts use the full allocation, independently of replay retention.

Evaluation	R_1 [interval]	S [interval]	Task Off → NT
Spatial archive	$-3.64\,[-5.58, -1.70]$	$3.40\,[1.53, 5.30]$	$20/20 \to 19/20$
Spatial replication	$-1.71\,[-2.52, -.984]$	$1.52\,[.874, 2.25]$	$40/40 \to 39/40$
LIBERO-10	$-.849\,[-1.645, -.240]$	$.718\,[.107, 1.532]$	$38/40 \to 37/40$
GR00T	$-1.811\,[-4.239, -.357]$	$1.638\,[.285, 3.988]$	$38/40 \to 39/40$

All sixteen archival and 34 independent-Spatial sources have negative R_1; LIBERO-10 has 31 negative/five positive and GR00T 35 negative/one positive. Every retained task mean is negative. Restricting to exact full Off duplicates preserves negative R_1 intervals: independent Spatial 33 sources/nine tasks, $-1.67\,[-2.50, -.926] \times 10^{-4}$; LIBERO-10 31/eight, $-1.045\,[-1.972, -.342] \times 10^{-4}$; GR00T 32/ten, $-1.843\,[-4.254, -.381] \times 10^{-4}$ m^2 s. The Spatial restriction is retrospective; the other two are registered sensitivities. The main results retain all eligible sources, including discordant task outcomes. Independent Spatial endpoint medians are 2.68 versus 12.30 mm with original/adapted streams, both with correction enabled. The ordering is metric-specific: healthy GR00T's mean angular stream cost is 1.193×10^{-3} rad^2 s, smaller than direct angular cost 2.536×10^{-3}; LIBERO-10 angular R_1 remains unresolved.

C.4 Faulted response and observer comparisons

The archival rotation-y fault is task-benign ($20/20$ in all source arms). Its complete signed contrasts remain in the extended evidence; the stronger fault below tests task-relevant repair.

The stronger six-coordinate fault has separate original and matched $\pi_{0.5}$ source collections, each with Off/NT task counts $6/40 \to 14/40$ (Appendix B.4). Original translation $D_0 = .0518\,[-.1724, .3036] \times 10^{-3}$ m^2 s and all translation contrasts remain unresolved, while angular direct benefit is $.423\,[.366, .474]$ rad^2 s, positive on all thirty valid sources. The matched evaluation retains the same 27 sources for both observers. Table 11 keeps these populations separate.

For matched $\pi_{0.5}$ observers, paired translation $T_{\mathrm{NT}} - T_{\mathrm{DOB}}$ is $-1.267\,[-3.537, .0621] \times 10^{-3}$ m^2 s and remains unresolved. Secondary differences are $R_{1,\mathrm{NT}} - R_{1,\mathrm{DOB}} = -1.310\,[-3.562, -.00751] \times 10^{-3}$ and $I_{\mathrm{NT}} - I_{\mathrm{DOB}} = -.1265\,[-.2604, -.0149] \times 10^{-3}$. The latter exceeds the negative practical margin; the former excludes zero but not that margin. These secondaries do not replace the primary or establish universal observer superiority.

Table 11: **Faulted physical response across two backbones.** Equal-task means and 95% task-bootstrap intervals; n counts retained sources. Translation R_1 uses 10^{-3} m^2 s; angular T_r uses rad^2 s. All angular totals are positive on every retained source. They are secondary except for the registered GR00T NT angular primary. Different marginal intervals do not establish an observer difference.

Evaluation	n	Translation R_1	Angular T_r
$\pi_{0.5}$ original NT	30	$1.204\ [-2.667, 4.530]$	$.417\ [.355, .472]$
$\pi_{0.5}$ matched NT	27	$1.287\ [-2.793, 5.128]$	$.415\ [.353, .471]$
$\pi_{0.5}$ matched DOB	27	$2.596\ [.075, 5.647]$	$.430\ [.371, .484]$
GR00T NT	30	$2.095\ [.131, 4.826]$	$.410\ [.349, .468]$
GR00T DOB	29	$.287\ [-1.476, 1.858]$	$.427\ [.365, .489]$

Completed GR00T stronger-fault replication. Off/NT/DOB score $4/11/10$ out of forty. NT–Off has eight rescues/one regression, $+17.5$ percentage points $[-2.5, 40]$; DOB–Off has six/zero, $+15$ points $[0, 32.5]$; NT–DOB has three/two, $+2.5$ points $[-10, 15]$. NT passes its registered five-net-gain minimum but fails interval exclusion. Healthy Off is $39/40$; no healthy adaptive arm was allocated here. Off–NT prefixes match $39/40$, Off–DOB $40/40$ and healthy-Off/fault-Off $40/40$. The NT mismatch at task 7/state 29 fails in both source arms.

The angular primary passes: NT $T_r = .410\,[.349, .468]$ rad^2 s. Its direct angular benefit is $.418\,[.360, .472]$; translation total $T = 1.839\,[-.00516, 4.606] \times 10^{-3}$ m^2 s remains unresolved. On 29 common valid keys, paired translation $T_{\mathrm{NT}} - T_{\mathrm{DOB}}$ is $1.884\,[-.202, 5.736] \times 10^{-3}$, and paired R_1 and I are also unresolved. Differing marginal significance does not establish an observer difference. The generic scorer's `--primary T` mislabeled translation T in metadata but also computed angular T_r. Preserved outputs and independent trajectory reconstruction support the registered angular endpoint.

Retention matters: four task-0 healthy sources and task 3/state 31 are too short, and task 7/state 29 fails prefix matching. Of 34 replayed keys, NT excludes task 3 at states 23/24/29 and task 7/state 24 for fidelity; DOB also excludes task 8/state 29. Both physical analyses therefore omit tasks 0 and 3. The only NT task regression, task 3/state 24, remains in the forty-key task analysis but has no qualified replay effect.

C.5 Adapter families separate recovery from healthy response cost

A registered $\pi_{0.5}$ Spatial comparison uses inspected states 18/19/21/22, sampler seed 100001, and separate healthy/faulted processes. Each process contains all comparator arms. Fault and adaptation start at policy step 30; rotation support and cap .05 are shared, with the same Panda calibration for the calibrated arms. With integral gain k_i, the uncalibrated rule is $\hat{f}_{k+1} = \Pi_{\mathcal{C}}[\hat{f}_k + k_i(y_{k+1} - u_k)]$. The L1-style comparator uses intermediate target σ_k and filter gains κ, γ to filter the calibrated target twice: $\sigma_{k+1} = (1-\kappa)\sigma_k + \kappa z_{k+1}$, then $\hat{f}_{k+1} = \Pi_{\mathcal{C}}[(1-\gamma)\hat{f}_k + \gamma\sigma_{k+1}]$, from zero, without gate/attenuation. At $\kappa = 1$ it reduces to the calibrated DOB; this implementation does not inherit a general L1 stability guarantee or reproduce another published adapter.

Development uses twenty inspected keys (states 15/17, ten tasks), separate from evaluation, seed 100002. Maximizing faulted successes over three integral gains/four L1 gain pairs (ties favor smaller gains) selects $k_i = .10$ and $(\kappa, \gamma) = (.3, .16)$. NT retains $(\gamma, \delta, \rho) = (.08, .008, .15)$. L1–NT gives five rescues/no regressions, $+12.5$ points $[5, 20]$; exact McNemar $p = .0625$ is descriptive. NT–integral is unresolved: $+2.5$ points $[-7.5, 15]$, four/three discordances. Unequal tuning prevents a general adapter ranking. Integral's four healthy losses remain in the forty-key task denominator; their net interval is $[-22.5, 0]$ points. They occur at task/state $(5, 19), (5, 22), (7, 22), (9, 21)$. The separate faulted-process healthy reference scores $39/40$, distinct from the healthy comparator block's $40/40$ Off arm.

Every adapter has negative healthy R_1 on all 33 retained sources. Paired $R_{1,\mathrm{NT}} - R_{1,\mathrm{integral}} = 1.394\,[.347, 2.789] \times 10^{-3}$ m^2 s: the 7.0 ratio of mean costs reflects a resolved difference. Healthy NT–L1 R_1 and both faulted paired T contrasts remain unresolved. Retention (assigned/eligible/valid) is $40/35/33$ healthy, $40/30/28$ faulted NT/integral and $40/29/27$ L1; fidelity is qualified only on retained matrices. Comparisons reuse keys and are not independent; intervals use 10,000 task-bootstrap

Table 12: **Registered adapter comparison.** Successes out of forty; healthy translation R_1 in 10^{-4} m^2 s and faulted total T in 10^{-3} m^2 s, equal-task means with pointwise 95% intervals. Healthy matrices share 33 retained sources/ten tasks; faulted NT/integral retain 28/ten and L1-style 27/nine. Adapter comparisons use common keys.

Adapter	Healthy	Faulted	Healthy R_1	Faulted T
Off	40	7	—	—
NT	40	15	-2.31 $[-4.81, -.799]$	2.04 $[-.256, 4.62]$
Integral	36	14	-16.25 $[-32.21, -4.44]$	2.75 $[.366, 5.51]$
L1-style	40	20	-2.78 $[-6.14, -.796]$	2.19 $[-.162, 4.77]$

draws, seed 20260925; this seed value does not denote a collection or registration date. Registrations and the portable `artifact/new_results/` retain development selection, exclusions, paired contrasts, and the complete signed outputs.

C.6 Previously untested faults with equally budgeted adapters

The registered study retains all 880 development and 1,440 evaluation episodes, with no recorded technical failures or configuration mismatches. Forty new task/state cases span forty $(\text{suite}, \text{task})$ clusters, one per task in each LIBERO suite, with sampler seed 104002. Evaluation initial-state indices are 0/2/4/8 for Spatial/Object/Goal/LIBERO-10; development uses Spatial states 15/17 and seed 104001. The exposure inventory excludes prior calibration, development, live evaluation and replay cases. "Unseen" refers to this project; foundation-model pretraining exposure is unknown.

Faults and selection. All four new vectors lie within the rotation correction mask. The development fault is the historical six-coordinate $+.05$ offset; it is retained as the seen anchor. The new vectors change sign/support combinations or magnitude (Table 13), with cap .15. One configuration per arm is selected from four candidates using at most two healthy regressions out of twenty, then faulted success, healthy success, correction energy and fixed candidate order. NT selects $\gamma = .16$, DOB .04, nominal-error integral $k_i = .05$, and L1-style $\kappa = .6$, $\gamma = .16$. These settings and fault support differ from the earlier holdout; success rates are not pooled or interpreted as a gain-tuning intervention.

Table 13: **Withheld vectors and NT repair.** Translation fault coordinates are zero. Listed rotation offsets are normalized command units; all effects are NT–Off percentage points with marginal 95% task-bootstrap intervals. Each row evaluates the same forty task/state cases.

Condition	Rotation offset	Off	NT	Gain (points)
U1	$(-.05, -.05, -.05)$	14	34	50.0 $[32.5, 67.5]$
U2	$(+.05, -.05, +.05)$	8	35	67.5 $[52.5, 82.5]$
U3	$(+.14, +.14, +.14)$	1	19	45.0 $[27.5, 62.5]$
U4	$(0, -.05, +.05)$	25	37	30.0 $[15.0, 45.0]$

U1/U2/U4 are new sign-and-support combinations; U3 is a withheld magnitude in an existing family, close to the correction cap. It does not isolate correction authority as the cause of its lower recovery.

The fixed correction candidates $s = 0, .025, .05, .10$ act negatively on the three rotation coordinates. Their development faulted successes are $9/18/18/6$ and healthy successes $20/17/5/0$ out of twenty. Only zero meets the healthy screen, so the evaluation's static arm is a zero control. This is a result for that grid and selection rule, not a general impossibility for fixed correction.

Nominal-error integral. Let $A_k = (a_k, \ldots, a_{k-K})$ denote nominal-command history and let $k_i > 0$ be the per-step integration gain. Using the healthy predictor and centering of Section 3.2, this arm updates

$$e_{k+1} = M^{\dagger}[y_{k+1} - h_W(A_k)] - b, \qquad \hat{f}_{k+1} = \Pi_{\mathcal{C}}(\hat{f}_k + k_i e_{k+1}), \qquad c_{k+1} = D_m \hat{f}_{k+1}.$$

The projected state supplies anti-windup; the update is ungated. This differs from the earlier integral in Appendix C.5. In the matched memoryless scalar limit with $b = 0$, $e = f - c$ and this recursion

equals a DOB with the same gain. That limiting equivalence is not a stability result for the FIR implementation or the VLA closed loop.

All outcomes and uncertainty. The primary averages NT–Off over the 40×4 unseen case-condition grid. Ten thousand bootstrap draws resample forty task clusters, keeping all four conditions together (seed 20260925): $+48.1$ points $[40.0, 56.2]$, eighty rescues and three regressions. Leave-one-task-out estimates range 46.8–50.0 points. Spatial/Object/Goal/LIBERO-10 each contribute forty case-conditions, with Off–NT counts $15 \to 37$, $15 \to 34$, $17 \to 33$, $1 \to 21$. The four conditions do not supply independent task replications.

Table 14: **Complete six-arm outcomes.** Successes out of forty per cell; healthy and anchor rows are excluded from the unseen-fault primary. Integral denotes the nominal-error integral.

Condition	Off	Static	NT	DOB	Integral	L1-style
Healthy	37	37	37	35	36	37
Seen anchor	11	11	26	22	25	27
U1	14	13	34	35	37	34
U2	8	8	35	33	39	34
U3	1	1	19	14	25	25
U4	25	25	37	34	37	37

Integral exceeds NT by 8.1 points $[1.9, 14.4]$, with twenty integral-only and seven NT-only successes. NT–DOB is $+5.6$ points $[0, 11.2]$; NT–L1 is -3.1 points $[-8.8, 2.5]$. Healthy differences versus Off are NT $0\,[-7.5, 7.5]$, DOB $-5\,[-15, 5]$, integral $-2.5\,[-12.5, 7.5]$, and L1 $0\,[-10, 10]$ points. Marginal, unadjusted intervals do not establish healthy equivalence or a healthy-performance ranking. The integral's higher faulted success is not an identified tradeoff with healthy performance. All regressions and suite partitions are retained in `artifact/unseen_faults/`.

Reproducibility and scope. One of 240 Off/zero-static case-condition pairs differs in task outcome (Goal task 1, state 4, U1). Thus zero correction does not give exact empirical replication here; $1/240$ is an observed discrepancy rate, not a universal noise floor or resolution bound. Arm order is cyclically rotated, which does not balance every relative arm position within the unseen contrasts. The dated protocol records this limitation and partial totals viewed during a scorer dry run. Core conditions, selection and estimands were fixed before evaluation, with disclosed amendments. These results establish task effects under the specified command faults, not new replay effects or a validation of the physical-recovery bounds.

C.7 A RECONSTRUCTED STRUCTURED ADAPTER UNDER MATCHED COMMAND CONSTRAINTS

The completed $\pi_{0.5}$ comparison retains all 1,320 episodes: 110 task/state cases, two sampler seeds (102001/102002), three arms and healthy/faulted conditions. These cases and seeds were used in the earlier holdout; this is a matched benchmark comparison, not additional untouched-state evidence. The fault adds $+.05$ to all six normalized arm commands from the first policy action. Both adapters are enabled then, after ten warmup steps, correct rotation coordinates only, and share cap .15 and five-action replanning.

Implementation and selection. The APEX-style arm reconstructs the published tracking-error, velocity and acceleration command terms (Zhao et al., 2026). The available specification does not determine their discrete coefficient updates or their mapping to our normalized operational-space interface; the published joint-space implementation was unavailable. The reconstruction uses backward differences, an inferred update and sliding variable, and conditional integration against outward coefficient growth at output saturation. It has no healthy plant calibration; NT retains the fixed healthy predictor. Consequently, this is neither a faithful APEX reproduction nor an ablation of the update law alone.

Both methods score four candidates on twenty development cases, retaining candidates with at most two healthy regressions and selecting by faulted success, then healthy success, then smaller gain. NT gains .02/.04/.08/.16 yield faulted $17/19/20/19$ and healthy $20/19/20/20$; structured-adapter rate multipliers .001/.01/.1/1 yield faulted $12/9/10/14$ and healthy $20/20/19/18$. NT

Table 15: **Completed comparison under matched correction constraints.** Successes per 220 matched keys; $+/-$ counts rescues/regressions versus Off. The 1,320 episodes form a separate collection on the holdout's 110 task/state cases and two seeds: faulted Off/NT totals are $61/116$, versus $58/121$ in Table 2. Collections are not pooled. The comparator is a documented reconstruction, not published APEX.

	Faulted		Healthy	
Adapter	Success	$+/-$	Success	$+/-$
Off	61	—	211	—
NT	116	64/9	215	7/3
APEX-style	76	31/16	166	3/48

selects .08; the reconstruction selects multiplier 1, the largest tested value; the finite grid does not establish the reconstruction's optimal performance. Matching numerical published gains does not establish matching effective gains across interfaces. The scored budget is 400 episodes, including each method's Off controls; forty superseded Off episodes and an aborted technical diagnostic are additionally archived. Candidate counts and selection rules match, not total development histories.

Paired results and uncertainty. NT–reconstruction faulted success is $+18.2$ points $[8.2, 28.4]$, with 56 NT-only and sixteen reconstruction-only successes. Effects weight keys equally; intervals resample forty $(\text{suite}, \text{task})$ clusters, keeping states and seeds together (10,000 percentile draws, seed 20260925). NT–Off is $+25.0$ points $[14.3, 36.0]$; reconstruction–Off is $+6.8$ points $[-.5, 14.4]$. The latter's faulted improvement is unresolved. Healthy NT–Off is $+1.8$ points $[-1.4, 5.1]$, with seven rescues and three regressions; this is not a noninferiority result. The reconstruction loses 20.5 healthy points $[10.2, 31.8]$, with three rescues and 48 regressions.

Table 16: **All seed and suite outcomes for the structured-adapter comparison.** Successes in Off/NT/APEX-style order; n counts keys per arm and condition. Seed and suite rows partition the same 220 keys in different ways.

Partition	n	Faulted	Healthy
Seed 102001	110	30/55/39	107/107/82
Seed 102002	110	31/61/37	104/108/84
Spatial	20	12/20/16	20/20/17
Object	40	13/27/21	38/40/34
Goal	80	35/52/33	79/79/64
LIBERO-10	80	1/17/6	74/76/51

Both seed blocks favor NT over the reconstruction; leaving out any one task gives $+15.6$ to $+19.9$ points. No evaluation keys are excluded. The initial registration followed the first development Off collection but preceded scored candidate outcomes. The anti-windup amendment preceded scored evaluation; the second-seed amendment preceded collection of either block. The dated protocol, retrospective chronology correction, reconstruction choices and portable outcome verification are in `artifact/apex_style_comparison/`.

C.8 GR00T RESPONSE WITH PROPRIOCEPTIVE INPUT HELD FIXED

A registered healthy Spatial intervention uses states 23/24/29/31 and seed 101001. In the stale-state process, each episode's eight-dimensional state input is held at its first post-reset inference value; external/wrist images remain live. Off and NT share a process within each condition; live and stale conditions use separate processes and healthy references. The 1,246 episode inferences comprise eighty initial-state calls and 1,166 substitutions; two unseeded warmups are excluded. All eighty stale episode acknowledgments enable the intervention.

Both conditions score Off $38/40$, NT $37/40$; stale NT has one rescue/two regressions, -2.5 points $[-10, 5]$. Live retains $40 \to 37 \to 36$ sources; stale retains $40 \to 36 \to 36$, each covering ten tasks. Live/stale mean R_1 are -3.41 $[-7.14, -1.02]$ and -2.86 $[-5.63, -1.01]$, in 10^{-4} m^2 s; 34/36

and 35/36 source effects and every task mean are negative. Corresponding S is 3.22 [.84, 6.90] and 2.71 [.85, 5.39] in the same units. Cost persists without live proprioceptive changes, consistent with visual feedback. Different processes, references and populations preclude a mediated-fraction or equivalent-behavior interpretation. Faulted repair was not tested.

C.9 Why response magnitude and direction both matter

Let $P_{am}, P^{\text{ref}} \in \mathbb{R}^{50\times 3}$ hold positions; vec stacks their coordinates. Use $E_{am} = \sqrt{.05}\,\text{vec}(P_{am} - P^{\text{ref}}) \in \mathbb{R}^{150}$ with identity weight Q in Proposition 1. Its identity $R_1 = -2A - r_\Delta^2$ separates directional contribution from magnitude penalty; for $qr_\Delta > 0$, $R_1 > 0$ requires $\cos\theta < -r_\Delta/(2q)$. Figure 4 displays this condition on observed trajectories, all with $q, r_\Delta > 0$.

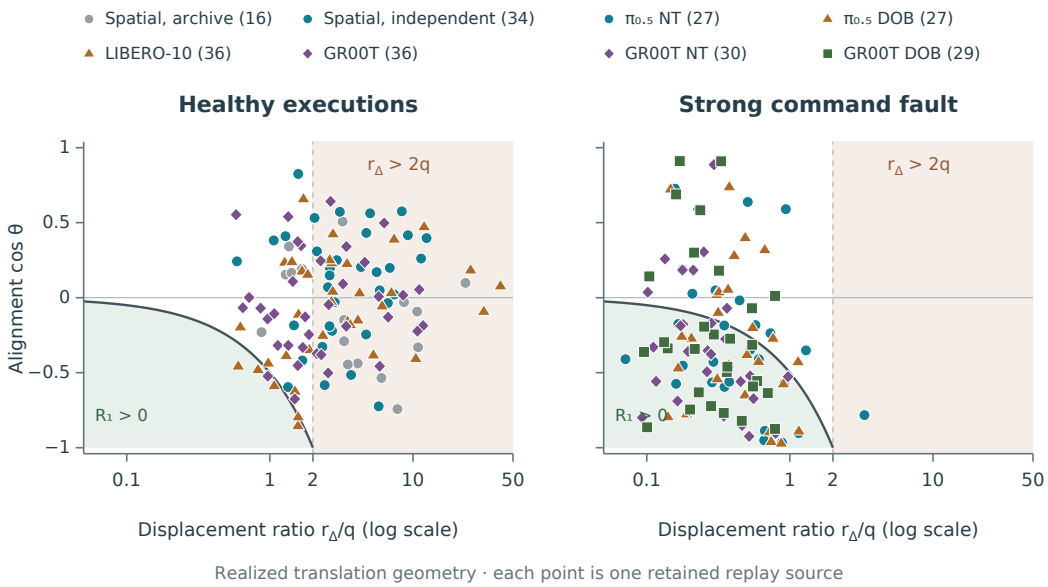


Figure 4: **Retrospective geometry of the stream effect.** Each point is a retained source; below the boundary, stream replacement reduces translation cost. Healthy and faulted panels use the same axes. Matched $\pi_{0.5}$ observers share 27 sources; GR00T has 30 NT and 29 DOB sources. This displays an identity on observed trajectories; it neither predicts prospectively nor independently validates that identity.

In healthy independent Spatial, LIBERO-10 and GR00T, $r_\Delta > 2q$ on $27/34$, $20/36$ and $18/36$ sources; only $12/34$, $20/36$ and $21/36$ responses point toward the reference. In the matched strong-fault $\pi_{0.5}$ NT/DOB evaluations, only $1/27$ and $0/27$ exceed $2q$, while $21/27$ and $19/27$ point toward the reference. GR00T faulted NT/DOB have no such magnitude exceedances and $22/30$ and $21/29$ reference-directed responses. Table 17 reports both terms' uncertainty. Joint bootstrap draws preserve their difference R_1; marginal interval endpoints cannot be subtracted.

The realized norm-only sufficient condition $r_\Delta < q_0 - q$ passes on $0/113$ source–matrix records in these four strong-fault matrices, although $T > 0$ on $15/27$, $15/27$, $20/30$ and $16/29$ sources, totaling $66/113$. Repeated keys preclude 113 independent trials. Each original smaller-fault matrix passes $1/15$. Ignoring direction is conservative; this does not qualify a prospective bound on unseen policy responses.

C.10 Cost interaction is not trajectory nonadditivity

For position arrays define $e = P_{00} - P^{\text{ref}}$, $u = P_{10} - P_{00}$, $v = P_{01} - P_{00}$ and the reference-independent displacement $Z = P_{11} - P_{10} - P_{01} + P_{00}$. For any such arrays U, V, use $\langle U, V\rangle_Q = .05\sum_{k=1}^{50} U_k^\top V_k$:

$$I = -2\langle u, v\rangle_Q - 2\langle e + u + v, Z\rangle_Q - \|Z\|_Q^2.$$

Table 17: **Directional contribution and magnitude penalty.** Equal-task means and exploratory 95% task-bootstrap intervals. Healthy rows use 10^{-4} m^2 s and faulted rows 10^{-3} m^2 s. Their difference gives the R_1 results above. These retrospective terms are unregistered.

Evaluation	$-2A$	r_Δ^2
Healthy Spatial archive	.253 [−.142, .715]	3.89 [1.83, 6.03]
Healthy Spatial replication	−.239 [−.521, −.020]	1.47 [.910, 2.08]
Healthy LIBERO-10	.107 [−.024, .296]	.957 [.378, 1.71]
Healthy GR00T	.017 [−.169, .185]	1.83 [.372, 4.14]
Matched strong $\pi_{0.5}$ NT	7.19 [3.00, 12.30]	5.91 [2.31, 10.34]
Matched strong $\pi_{0.5}$ DOB	6.34 [1.75, 11.81]	3.74 [1.52, 6.27]
Strong GR00T NT	3.50 [.818, 7.15]	1.41 [.656, 2.37]
Strong GR00T DOB	1.35 [−.737, 3.39]	1.07 [.650, 1.66]

Thus $Z = 0$ does not imply $I = 0$: additive displacements can interact in a quadratic cost. Conversely, nonzero Z need not produce a resolved cost interaction. Retrospective matched strong-fault $\pi_{0.5}$ estimates give mean root-mean-square nonadditivity $\sqrt{\sum_{k=1}^{50} \|Z_k\|_2^2/50}$ of 3.39 mm for NT and 3.34 mm for DOB. Their paired difference is .049 [−.389, .558] mm, leaving its magnitude unresolved despite the secondary difference in I. The combined Z-dependent contribution to that difference likewise has interval $[-3.015, .023] \times 10^{-4}$ m^2 s. This Euclidean-position analysis does not transfer unchanged to global rotation subtraction or isolate observer recurrence, contacts and plant dynamics from one another.

C.11 Reference choice and task-level counterexamples

Holding trajectories fixed, shifting reference-position array r by array w gives

$$R_1(r + w) = R_1(r) + 2\langle P_{11} - P_{10}, w\rangle_Q, \qquad S(r + w) = S(r) - 2\langle P_{00} - 2P_{10} + P_{11}, w\rangle_Q.$$

Displacement is unchanged; alignment is reference-relative. A registered sweep uses healthy-Off paths of the same task/initial state under other seeds, plus one same-seed/other-process comparison. Windows align by policy step (30–79), without spatial alignment or time warping; post-prefix physical states can differ. Cells use observed `position`, common to telemetry/replay, rather than primary `ee_pos` (maximum saved discrepancy .63 mm); recomputed declared-reference results differ slightly.

Table 18: **Measured healthy-reference sensitivity.** All costs use 10^{-4} m^2 s and equal-task means. Paired R_1 is retrospectively recomputed on the alternative's same n retained sources; alternative 95% intervals preserve the registered 10,000-draw schedule, seed 20260927. J_{00} measures departure of unadapted execution from the alternative. An asterisk denotes the same sampler seed in a separate process. References/matrices are reused; rows are not independent replications.

Matrix (seed)	Reference seed	n	Paired R_1	Alternative R_1 [interval]	J_{00}
$\pi_{0.5}$ NT (100001)	85001	33	−2.304	−1.064 [−3.257, 1.158]	24.36
	99001	33	−2.304	2.781 [−1.943, 10.209]	22.03
	100001*	32	−2.325	.288 [−1.434, 2.876]	2.29
$\pi_{0.5}$ NT (86001)	93001	34	−1.701	−.472 [−2.656, 1.463]	27.96
	93002	34	−1.701	1.641 [−.474, 4.553]	26.79
$\pi_{0.5}$ L1 (100001)	85001	33	−2.777	−2.535 [−5.471, −.076]	24.36
	99001	33	−2.777	1.926 [−3.024, 9.474]	22.03
	100001*	32	−2.795	−.236 [−1.900, 1.784]	2.29
$\pi_{0.5}$ integral (100001)	85001	33	−16.191	−13.508 [−23.897, −5.215]	24.36
	99001	33	−16.191	−3.925 [−14.196, 10.301]	22.03
	100001*	32	−16.142	−9.884 [−20.588, −3.173]	2.29
GR00T NT (89001)	98001	34	−1.971	−1.106 [−3.301, .219]	3.99
	101001	35	−1.798	−1.454 [−3.146, −.116]	12.32
GR00T NT (101001)	98001	35	−3.762	−2.643 [−7.424, 1.153]	13.38
	89001	36	−3.397	−4.262 [−10.345, .212]	12.32

Cross-seed $\pi_{0.5}$ variability is $J_{00} = 2.20$–2.80×10^{-3} m^2 s, compared with paired NT/L1 stream costs $.17$–$.28 \times 10^{-3}$. The calibrated adapters' sign changes limit a typical-path degradation claim. GR00T and integral point estimates stay negative, but most GR00T intervals include zero. Finite references do not characterize successful-path distributions; healthy-Off episodes can fail. `artifact/new_reference_sweep/` reproduces every signed D_0, R_1, T, S, matched population and exclusion, distinguishing equal-task J_{00} from archived source-weighted summaries. Analytical perturbation envelopes for the primary coordinate remain valid: both initial Spatial R_1, S intervals survive 2.5 mm perturbations, while LIBERO-10's S interval includes zero at 1 mm. These local bounds and the measured alternative paths answer different sensitivity questions.

Discordant task outcomes. Rescues/regressions are adapter-only/comparator-only successes. Table 19 includes opposing physical signs and unavailable replays: cost reduction is neither necessary nor sufficient for task rescue. Costs span fifty translation steps; task success covers the episode and can depend on other coordinates.

Table 19: **Task outcomes can oppose reference-path costs.** Representative counterexamples and every healthy NT rescue/regression in the four initial evaluation groups. Values are source-level NT R_1, T in 10^{-3} m^2 s; positive values reduce reference cost. A dash marks unavailable replay, never zero. Retrospective selection exposes conflicting signs.

Evaluation	Task/state	Outcome	R_1	T
Healthy Spatial archive	8/14	regression	$-.0492$	$-.0634$
Healthy Spatial replication	4/31	regression	$-.2622$	$-.3970$
Healthy LIBERO-10	4/31	rescue	$-.0150$	$-.0154$
	6/33	regression	$-.0142$	$-.0144$
	8/30	regression	—	—
Healthy GR00T	2/29	rescue	.000696	$-.00845$
Strong $\pi_{0.5}$, original	2/19	regression	.8034	.0814
	7/21	rescue	-6.8019	-6.6345
	2/22	rescue	$-.5624$	-1.1123
	3/21	rescue	—	—
Strong GR00T	2/29	rescue	-5.2385	-5.7175
	2/31	rescue	2.0575	$-.0556$
	9/24	rescue	$-.5228$	.2743
	3/24	regression	—	—

The artifact retains all outcomes, contrasts and NT–DOB discordances. Source costs are not paired observer differences.

Reproduction and scope. The evidence index and CPU-only artifacts retain registrations, exclusions and signed outcomes; offline analyses add no independent samples or prospective certificate. Earlier strong-$\pi_{0.5}$ executable provenance and LIBERO-10 timestamps are incomplete; saved arithmetic is checked. Extended material retains historical collection details, secondary endpoints and diagnostics.

D IDENTIFICATION, AUTHORITY, MEMORY AND TIMING: SUPPORTING EXPERIMENTS

Registrations and extended evidence retain omitted historical configurations.

D.1 IMPROVED IDENTIFICATION NEED NOT IMPROVE TASK RECOVERY

Theory-guided pilot. The mismatch diagnosis motivated a pre-evaluation registered prediction: constraining predictor rotation gains to the probe would raise the median late rotation-y estimate to at least 70% of the fault. NT reached 81.5% on Spatial and 70.5% on LIBERO-10, versus historical $41\%/45\%$. Unpinned task counts were 17 versus 18 of twenty and 18 versus 15 of forty; healthy Spatial changed $20 \to 19$. This supports the predicted direction and threshold, not exact equilibrium or resolved task benefit.

E2 fits unconstrained (U) and rotation-y tap-sum-constrained (.254; C) FIR predictors with shared data, ridge penalty, free-intercept procedure and sensitivity: thirty healthy episodes for fitting, ten for qualification. Only one of three checkpoints passes response checks; steady-state gain remains unqualified.

Each suite crosses ten tasks, two states and three sampler schedules (sixty keys); seven arms give 840 episodes. Both innovation observers use corrected-channel normalization, cap .15 and uniform fault $+.05$. Writing subscript r_y for rotation about y, observation error averages absolute last-fifty-step episode *signed mean* errors in $(M^{-1}r)_{r_y}$ relative to the fault, not instantaneous errors or pooled bias. Intervals use whole-task resampling (10,000 draws; exact bootstrap distribution for binary contrasts; on the knife-edge Spatial C–U contrast the archived 10,000-draw scorer reports $[-13.3, 0]$).

Table 20: **Controlled predictor outcomes.** Successes out of sixty; parentheses give healthy counts. The oracle cancels rotation using the known fault. Only the response constraint changes between U/C.

Suite	Off	U	C	Rotation oracle
Spatial	30 (60)	58 (58)	55 (60)	54
LIBERO-10	0 (56)	18 (53)	25 (52)	36

Spatial observation bias falls $.02678 \to .01727$, or 35.5% $[21.9, 47.6]\%$, but C–U task success changes -5 points $[-11.7, 0]$ (one gain, four losses). LIBERO-10 bias rises $.02520 \to .03361$: relative reduction -33.4% $[-135.6, 39.8]\%$, with uncertain sign. C–U task gain is $+11.7$ points $[0, 25]$ (twelve gains, five losses). Healthy LIBERO-10 U/C lose four/five Off successes and rescue one each: -5 $[-11.7, 0]$/-6.7 $[-16.7, 0]$ points. C misses its five-point healthy-loss target. The mean signed applied-correction fractions rise $.529 \to .941$ on Spatial and $.490 \to .889$ on LIBERO-10. These average the scalar ratios c_{r_y}/f_{r_y}; remaining-disturbance fractions are $1 - c_{r_y}/f_{r_y}$. Neither is an absolute estimation-error measure.

Predictor-memory control. Eighty LIBERO-10 keys use NT with pinned sampling and historical calibration. FIR, an autoregressive model with exogenous inputs (ARX) using one output lag, and response-constrained FIR score $10/80, 8/80, 14/80$, with episodewise last-fifty-step rotation-y means averaging $26.7\%, 33.3\%, 61.3\%$ of the .05 fault. Constrained minus FIR gives $+5$ points $[-6.25, 17.5]$ (eight gains, four losses). ARX also rescales sensitivity; no healthy arm was collected. Constrained FIR's episode median is 71.5%, but its mean misses the 70% target. The pilot's 70.5% is a median from different, unpinned episodes, not comparable to that mean. Task improvement remains unresolved.

D.2 Correction authority and healthy regressions

Six-channel correction reuses E2's sixty LIBERO-10 keys and innovation constants. Online comparisons change support, normalization and calibration; only the oracle isolates support.

Table 21: **Additional authority and its healthy cost.** Successes out of sixty matched task/state/sampler keys. Faulted/healthy Off are 0/56. U/C denote unconstrained/constrained predictors; the oracle uses the known fault.

Correction	Fault U	Fault C	Oracle	Healthy U	Healthy C
Rotation	18	25	36	53	52
Six channels	35	41	57	47	53
Six channels with gate	—	32	—	—	54

Write U6/C6 for six-channel U/C. C6 versus rotation-only C gives $+26.7$ points $[8.3, 46.7]$ (22 gains, six losses); C6 versus U6 gives $+10$ $[-1.7, 21.7]$. The six-channel oracle exceeds C6 by 26.7 $[15, 38.3]$ (sixteen gains, no loss), and its $36/60 \to 57/60$ measures available authority. Healthy U6/C6 lose ten/six and rescue one/three Off successes: -15 $[-26.7, -3.3]$/-5 $[-16.7, 3.3]$ points, both exceeding three losses.

From twenty healthy estimate-only episodes, per-channel late-estimate means b_i and across-episode standard deviations s_i define a gate opening channel i when $|\hat{f}_i - b_i| > 3\max(s_i, .002)$. This gives no instantaneous three-sigma guarantee. Gated C6 gives healthy $54/60$ (one rescue, three losses), but faulted $32/60$, missing the recovery target $36/60$. Relative to ungated C6 the faulted change is -15 points $[-31.7, 1.7]$.

Physical faults and allocation. Weighted allocation regularizes correction across all inputs; masked allocation selects coordinates after applying the sensitivity pseudoinverse. Both support and regularization change. Panda joint-5 torque Off/masked/weighted scores are $36/15/41$ of sixty across three suites; weighted gains/losses are $29/3$ against masked and $14/9$ against Off. Object joint-3 totals are $8/16/2$ of twenty, with weighted gains/losses $1/15$ against masked. ALOHA faulted totals are $60/70/66$ of 240, with healthy $11/8/6$ of twenty. Allocation benefit is fault-dependent. Separate Panda joint-3 friction increase gives $0/40 \rightarrow 13/40$; locking gives $0/20$ in both arms, without isolating a rank-loss mechanism.

D.3 Physical memory, delay and the limits of prediction

State-restored, same-command continuations give local fits $\hat{\lambda} \simeq 1$ on Panda and .663 on ALOHA. Panda's near-unit median perturbation-gap ratio and expanding steps leave strict contraction unqualified.

Length/fidelity checks retain seventeen/nineteen of twenty qualification/test sources each. Fifty-command branches follow thirty-command prefixes, with restore agreement within 10^{-8} rad, shared calibration and rotation mask, and zero initialization. Reference rotation-y fault/cap are $+.05/.05$; interventions delay updates ten steps, halve the cap or reverse fault sign. Updates affect the next command; the known-fault reference cancels at its assigned onset.

Table 22: **Measured physical benefit and delay/cap costs.** Equal-source paired means with task-bootstrap 95% intervals on nineteen test sources, in 10^{-5} m^2 s over 2.5 s. The first row is Off minus corrected cost; subsequent rows are increases relative to each controller's reference condition.

Contrast	NT	Innovation	Known-fault reference
Reference benefit	8.79 [4.69, 12.72]	8.99 [4.83, 13.08]	14.67 [10.26, 19.48]
Delay penalty	2.69 [.95, 4.38]	2.83 [1.10, 4.74]	4.00 [2.18, 6.88]
Half-cap penalty	1.05 [.18, 1.91]	1.46 [.53, 2.50]	4.83 [3.28, 6.87]

Median endpoint deviation falls from 13.74 mm Off to 7.80 mm with reference NT. Delay raises NT and innovation cost on 18/19 sources; halving the cap raises it on 14/19. Known-fault median endpoint errors are zero immediately, 2.74 mm delayed and 5.72 mm at half cap. The capped reference retains median .250 of Off energy, below the registered half-energy expectation: halved displacement quarters squared cost. Known-fault correction does not dominate both laws on every delayed/capped source. Healthy NT/innovation mean costs are $1.22\,[.50, 2.10]$ and $1.44\,[.58, 2.48] \times 10^{-5}$ m^2 s.

Prospective qualification. A twenty-tap development-fitted forecast uses sampled healthy residuals. Neither constant nor error-dependent qualified radius yields decisive intervals on 228 test comparisons, distinct from the retrospective $0/113$ norm-budget check (Appendix C.9). Whole-source coverage is 16/19 and 19/19; all constant-radius misses involve reversed signs. Point predictions always favor adaptation: 88.2% NT/86.8% innovation sign accuracy equals the always-benefit baseline. No prospective rule or impossibility of prediction follows.

ALOHA comparison. Twenty ALOHA streams split 6/4/10 for response fitting/qualification/testing with two fifty-step checkpoints per test source and a .02 rad left-arm fault. Predictor fitting uses all logs; its data are not disjoint. Off/NT/innovation/held median endpoints are $.04915/.00313/.00197/.00240$ rad; the empirical response envelope covers 72/100 nonhealthy branches. Interfaces, durations and metrics differ from Panda, confounding robot effects.

Task outcomes with delayed activation. On forty scheduled stronger-fault keys, delayed NT updates/corrects at steps 40/41 versus immediate 30/31, scoring $13/40$ versus $14/40$ and archived Off

$6/40$. Delayed minus immediate is -2.5 points $[-10, 5]$ (two gains, three losses); delayed minus Off is $+17.5$ $[7.5, 30]$ (seven, zero). The arms share pre-fault prefixes and seed 85001, without delayed healthy controls. Delay's task cost remains unresolved; a benign-fault bridge scores $20/20$ in all faulted arms (healthy NT: $19/20$).

D.4 A SCHEDULED PHYSICAL DYNAMICS SHIFT

Joint-3 friction loss increases at step 30 and resets per episode. Translation correction uses cap .30 and fixed calibration. Development selects added friction 20 by the registered Off-success threshold: Off/NT score $10/20 \to 16/20$ (seven rescues, one regression), an exploratory $+30$ points $[5, 55]$.

Table 23: **Completed friction-shift evaluation.** Entries are Off/NT/DOB successes. Both blocks reuse forty inspected task/state keys; Block B was registered after observing A. Intervals cluster whole tasks, retaining both seeds in the combined analysis.

Block	Faulted	Healthy	NT minus Off (points)
A, seed 93001	23/28/31 of 40	40/40/38 of 40	12.5 $[0, 27.5]$
B, seed 93002	23/24/31 of 40	40/40/39 of 40	2.5 $[-10, 15]$
A+B	46/52/62 of 80	80/80/77 of 80	7.5 $[-5, 21.25]$

Block A NT gives seven rescues/two regressions; combined eleven/five fail the amended ten-net-gain and positive-interval criteria. Combined DOB minus Off is $+20$ points $[-6.25, 50]$ (22 rescues, six regressions); NT minus DOB is -12.5 $[-32.5, 3.75]$ (four, fourteen). Observer superiority remains unresolved. Healthy NT loses none; DOB loses $(4, 23)$, $(9, 31)$ in A and $(5, 23)$ in B, without rescues. States are reused and no dynamics-fault crossed replay is qualified.

D.5 ADAPTATION BEFORE A HIDDEN FAULT ONSET

NT/DOB starts at step 10 before hidden onset 30/50/70, without signaling or resetting, with fixed rotation support, cap .05, calibration and gain .08. The 180-episode screen uses uniform $\pm.10$ or rotation-y $\pm.02$ faults and thirty-step pulses.

Table 24: **Hidden-onset task outcomes.** Complete assigned outcomes are retained, including success before the scheduled fault. The confirmation was registered before the screen, conditional on its gate; physical states were previously inspected.

Allocation	Condition	Off	NT	DOB
Screen, twenty keys	Healthy	20	19	19
	Step	15	16	15
	Thirty-step pulse	19	20	20
Confirmation, forty keys	Healthy	40	40	38
	Uniform $+.10$ step	13	25	24

Screen NT step and both pulse gains are $+5$ points $[0, 15]$; healthy losses give -5 $[-15, 0]$. Censored pulse responses preclude unconditional washout bounds.

Confirmation uses uniform $+.10$ six-coordinate steps, seed 99001, allocating 13/14/13 keys to onsets 30/50/70. NT–Off is $+30$ points $[12.5, 47.5]$ (fourteen rescues, two regressions at $(6, 22)$, $(7, 19)$); DOB–Off is $+27.5$ $[12.5, 45]$ (twelve, one at $(6, 22)$). NT–DOB is $+2.5$ $[-10, 15]$ (three, two). Healthy DOB loses $(4, 19)$, $(4, 22)$, giving -5 $[-15, 0]$; NT loses none.

Latency measures first post-onset correction-norm crossing of $.5\sqrt{3}(.05)$, not detection: medians six NT/five DOB steps on all 39 exposed episodes. One DOB source exceeds threshold before onset; one key succeeds before onset but remains in the denominator. Healthy NT/DOB mean correction norms are $.01146/.01333$ command units, measuring activity rather than physical deviation. Different seeds/schedules prevent ranking against scheduled $6/40 \to 14/40$; adapting prefixes prevent reusing its crossed replay.

D.6 STOPPED MIRROR-DESCENT DEVELOPMENT STUDY

The registered screen stopped before live-policy confirmation. Five six-setting families reuse 36 healthy LIBERO-10 and thirty faulted Spatial sources (two fail reproduction), minimizing healthy translation cost subject to positive angular direct benefit and faulted angular cost ≤ 1.10 times NT at gain .08. Selected healthy-translation/faulted-angular changes are $-33.35\%/+6.45\%$ (mirror), $-34.70\%/+7.67\%$ (curvature-scaled Euclidean), $-30.31\%/+9.64\%$ (fixed-rate diagonal quadratic), and $+14.72\%/-12.97\%$ (DOB). Neither mirror-specific superiority nor task preservation follows. On a unit integrator, the tested mirror law leaves displacement 16.02 versus NT's 6.84.

E PROVENANCE OF THE ROBOT ILLUSTRATIONS

Figure 1B is one of the 72 initial-state holdout keys where Off fails and Online succeeds (LIBERO-Spatial task 1, initial state 4, sampler seed 102002; six-coordinate $+.05$ fault); both runs use the same sampler seed. Frames are re-rendered by replaying the logged executed commands, and the replayed end-effector paths match the logs within 10^{-6} m. The agent-view camera is widened to 72° for the illustration; the policy observed $45°$. Online places the bowl at step 130 (zoom); Off times out at step 229. Of the nine such Spatial keys, only this one replays exactly in all three recorded arms (Off, Online, healthy) on our rendering machine.

Figure 1C splits single frames of recorded side-by-side videos into uncorrected and identify-then-hold halves at the same step: ALOHA/π_0 transfer-cube ($+.02$ rad left arm, one identification episode, frame 272/10.88 s), GR1/GR00T N1.5 tray-to-plate ($+.10$ rad right arm, three episodes, frame 149/7.45 s), and WidowX/GR00T N1.7 spoon-to-towel ($+.005$ translation, three episodes, frame 65/6.5 s). The uncorrected runs later fail (ALOHA at step 299; GR1 and WidowX time out at steps 719 and 149). Pairs share task and episode seed but not policy samples, so they illustrate recorded outcomes rather than controlled contrasts. Protocols differ across robots.

Figure 2 shows source (task 5, initial state 18, sampler seed 85001) of the stronger-fault $\pi_{0.5}$ NT collection, re-rendered by rerunning the replay driver; it was selected among sources that reproduce within the driver's tolerances on our rendering machine, and every branch matches the archived replay within 1.2×10^{-13} m.

Selection is illustrative: ALOHA's candidate count is unknown; GR1 retains four successes from seven; WidowX is a selected success. Images add no trials. Panel C keeps the recorded pixels with external labels; provenance files record processing.