This Week in Robotics

THE LINE

Three independent groups converged on the same answer to a question the field has been arguing about for two quarters: **the predicted future acts on a world-action policy through its presence and structure, not through its content.** Simple-WAM (preprint, Tsinghua Leap Lab) shows that leaving the future video tokens at *pure Gaussian noise* and running the video expert exactly once recovers almost all of the generalization that ten denoising steps buy — +26.4 points under environmental perturbation and +89.8 on video-conditioned task transfer over the latent paradigm, with the remaining nine passes worth at best +1.8. ReWAM (preprint, HUST + Horizon Robotics, **code released**) reaches 93.6% on RoboTwin 2.0 with no generative video pretraining at all, and shows raw DINO features beating Video-VAE latents by 4.9–9.2 points *despite* worse pixel reconstruction (SSIM 0.70 vs 0.91). V-JEPA Policy (preprint, **code released**) gets LIBERO-Plus 79.3 from 0.6B trainable parameters on a frozen V-JEPA 2.1 encoder. Read against W39's What Matters in WAM Design — corrupting future-latent *content* moved actions <1%, reversing their *temporal order* moved them 12.8–14.4% — this is the same finding arrived at from three directions, and it kills the Fast-WAM claim that inference-time future modeling is unnecessary while also killing the argument for paying for it.

Separately, one paper should change how you read every acceleration table you have ever seen: a seven-benchmark audit found 22 implementation bugs and 4 design limitations, and fixing them reverses method rankings. See §6.

Window: 2026-09-28 → 2026-10-04. Queue: 5 files, ~120 scored candidates, 30 at 4+. Coverage caveat carried from the harvester: cs.RO was enumerated completely only for the 2026-09-28 announcement; 09-29 through 10-02 are recall-limited (arXiv catchup 429'd and was abandoned, cs.LG and cs.CV were never enumerated). 10-03 and 10-04 had no announcement — weekend. Treat silence in any category as unmeasured, not empty.

1 Policy learning & VLAs

Simple-WAM — the future as an empty slot

arXiv:2609.34981 · preprint · Tsinghua Leap Lab + USTC + BIT · project page only, **no code stated**

Architecturally new:

Nothing is added. The intervention is (a) an attention mask that keeps future video tokens in the action expert's context, and (b) a scalar `p` that mixes the training flow-time schedule so 50% of samples draw τ = 1 (pure noise) instead of the usual continuous schedule. At inference the video expert runs one forward pass over tokens that are never denoised. Cost goes from `K·C_vid + K·C_act` (explicit) to `C_vid + K·C_act`, i.e. 74.7 ms/chunk against 286.9 for FastWAM-Joint and 62.0 for latent FastWAM.

What the eval actually covers:

A genuinely controlled three-axis protocol with matched backbone (Wan2.2-5B), matched data and matched budget — this is the paper's real contribution. Axis 1 is LIBERO-Plus's seven perturbation factors; axis 2 is retraining at 5 and 10 demos per task; axis 3 is four-fold cross-validation over LIBERO suites with the held-out suite's video available or not. Plus RoboTwin 2.0 and 4 real tasks × 30 trials on an AgileX Aloha. The headline table: 79.5 average under perturbation vs 67.7 explicit vs 53.8 latent; 10.1/73.6 on withheld tasks without/with action-free video vs 6.3/69.9 and 2.1/5.9.

The decisive number is in §4.3 and it is a pure inference-time sweep with *no retraining*: take the trained explicit model, run the video expert for only `t_denoise` of K steps. At `t_denoise = 1` the future tokens are $\varepsilon \sim N(0,I)$ — nothing has been denoised — and that single pass already beats the latent paradigm by 26.4 / 8.2 / 13.4 / 89.8 points across the four settings. The ablation in Tab. 3(b) is the one that makes the mechanism legible: substituting *learned* tokens or *zeros* for the noise collapses task generalization from 97.2 to 10.6 and 9.8. So it is not "the slot" alone — it has to be noise with the right statistics, which means the action expert learned to read a feature the video expert produces *from* noise, not a placeholder.

Reproducible with one arm and one GPU?

No for the result, yes for the mechanism. Training is a 5B video DiT plus a 1B action expert on A800s. But the two changes are an attention mask and one scalar in the noise sampler, both portable to whatever WAM or flow-matching VLA you already run. Latency was measured on a single 5090, which you may have.

ReWAM — the representation, not the generator

arXiv:2609.38163 · preprint · HUST + D-Robotics + Horizon Robotics + Fudan + XJTU · **code: https://github.com/hustvl/ReWAM**

Architecturally new:

Three pieces, and the third is the interesting one. *Feature Calibration* prepares frozen DINOv3-L features for diffusion — multi-layer aggregation over blocks {12,14,…,24} plus the final layer's spatial mean, running normalization, and a representation-aware noise shift $\alpha = \sqrt{(D/4096)}$. Worth 12.20 pp over raw DINO, decomposed as +3.92 normalization, +1.41 noise schedule, +6.87 multi-layer aggregation. A *Temporal Representation Bottleneck* flattens 2×2×2 spatiotemporal blocks across views into 128-d tokens. *Action-Grounded Representation Shaping* then routes **only action-loss gradients** into the bottleneck; world-loss gradients are stopped at both current and future representations. The gradient-routing ablation is the mechanism test: letting world-loss gradients reach the TRB through the current-state path costs 6.30 pp average, and through both paths costs more. Their reading — representation collapse or predictive shortcut — is plausible and untested.

Eval:

RoboTwin 2.0, 50 tasks × 100 episodes in clean and random; 93.6/93.6 vs Fast-WAM 91.9/91.8 and π0.5 82.7/76.8. RoboDojo across five capability categories. Real bimanual Piper, 3 tasks × 20 trials, 55/60 vs Fast-WAM 48/60. The number I would lead with in an interview is not the 93.6 — it is **600 hours of embodied pretraining against 10,000+ for π0.5 and ~20,000 for LingBot-VLA**, reaching RoboDojo average score 12.29 / SR 8.28% against π0.5's 11.41/6.91.

Where it's soft:

The Open instruction-following category is 0.25 against π0.5's 1.98 and HyVLA's 0.65, and the authors attribute the gap directly to the absence of generative video pretraining. That is an honest concession and it bounds the claim: representation-centric WAMs win on perturbation and long-horizon, lose on open-vocabulary instruction following. Precision also trails geometry-enhanced methods (X-VLA 18.32, Spatial Forcing 17.33, ReWAM 12.32). Main-table comparison is 5 epochs, ablations are 1 epoch — not the same budget. Single seed.

Reproducible?

Training is 3.58B params on 32 H20s — no. But the code is out, the DINOv3 backbone is frozen, and AGRS is a stop-gradient. The calibration recipe alone (normalize, shift the noise schedule, aggregate layers) is three lines and worth 12 points in their setting.

V-JEPA Policy — the one you can actually train

arXiv:2609.37250 · preprint · Tsinghua + SJTU + Fudan + USTC + WUSTL + TeleAI + γ-Robotics · **code: https://github.com/breez3young/VJEPA-Policy**

0.9B total, **0.6B trainable** on a frozen V-JEPA 2.1 encoder, no inherited video generator at all. A future-latent predictor's KV conditions a flow-matching action expert. LIBERO 97.3, LIBERO-Plus 79.3 against FastWAM 51.5 and ImageWAM 83.1; also RoboCasa-GR1. Pretrains the predictor on action-free DROID. This is the week's only WAM-class result at a scale a single workstation can finetune, and it arrives with code. **INFERENCE** \2192 taken together with ReWAM and Simple-WAM, the case for initializing from a video *generator* is now substantially weaker than it was a month ago — three groups got WAM-class generalization from frozen perceptual features or from noise.

Kinematic MeanFlow — one-step actions, measured end to end

arXiv:2610.00864 · preprint · Intel Labs China + Midea AI Research · code promised, URL not in abstract

Splits the velocity field at an intermediate point and estimates the two sub-intervals separately, constructing the intermediate state along the data-conditioned path to stop error propagation. On a finetuned GR00T-N1.6: **97.9% LIBERO at one step vs 97.3% at four**, with action-head latency down 67.5–74.4% and *end-to-end* down 30.3–54.9% (L40 eager: 149.6 → 67.4 ms). Reporting the end-to-end number separately from the head number is the right discipline and most acceleration papers do not do it. Edge-relevant: measured on Jetson Orin as well as L40.

Also worth your attention in this section. IronMind (preprint, XPeng) publishes an actual pretraining scaling curve for camera-space egocentric action — **55.0% OOD at 10k hours vs 11.7% at 5k vs 5.0% with none**, over six OOD tasks × 10 trials, with camera-space actions skipping body retargeting entirely; log-linear loss vs data, no code. RPG (preprint, Berkeley + Amazon FAR + MIT + UChicago) runs guided self-improvement with **no weight updates** — reconstructs practice tasks in sim, diagnoses failures from privileged simulator state, writes reusable symbolic skills; 28.6% after round 1 → 95.0% after round 15 over 22 tasks × 10 held-out inits, against ASPIRE 75.5 and a CodeAsPolicy/GPT-6-Astra-Pro agent at 60.0, plus 30/30 real over 3 tasks. Own task suite, robot platform not named, no code statement — hold it loosely. Causeway (preprint) is a training-free write into a frozen policy's action-stream representation that moves LIBERO-Goal instruction-switching from 3–26% to 47–65% over 71 instruction pairs on real xArm — directly adjacent to GAM's result from W39 and far cheaper, if it holds.

---

2 Manipulation, dexterity & sensing

UVTA — human tactile demonstrations, robot-only at deployment

arXiv:2609.34182 · preprint · SJTU + BIGAI + Sharpa + BIT · **code: github.com/uni-vta/UVTA**

Teleoperating a dexterous hand gives the operator almost no tactile feedback, so robot tactile demos do not scale. UVTA's answer is to collect tactile from *humans* with a mocap glove rig, build an aligned visual-tactile-action representation across human and robot data, and jointly predict future action and future tactile. **1,000 human demos + 150 robot demos per task over 5 contact-rich tasks. 70% average against the strongest baseline (T-Rex) at 29%. Without the tactile-prediction objective, 42%.** The scaling leg is the part that matters: 0% with no human data rising to 84% at 1,000 human demos, **with no visible saturation** at the maximum tested.

This is the cleanest instance yet of the pattern W39's thread named — supervision at training time, cheap sensing at deployment — and unlike ME-Dex (whose zero-tactile-input result used *simulator-generated* tactile) the tactile here is real human contact data. It does not close the thread's open question, because the robot still reads tactile at deployment; what it removes is the need to collect tactile *on the robot*.

TacGB — shear for free, bolted onto a sensor you already own

arXiv:2609.34006 · preprint · UC Berkeley (+ a high-school co-author, which is its own signal) · **no release stated**

A passive domed film laid over an existing normal-only pressure sensor. Tangential load tilts each dome and redistributes pressure across its footprint, so shear is encoded *mechanically* into the normal map. No added electronics, no force reconstruction, no taxel-to-dome alignment, no calibration — an end-to-end policy just consumes the richer map. **Up to +36 pp on insertion** across four imitation-learning tasks, with faster successful insertions and gentler contact on fragile placement. This is the highest ratio of mechanism-to-cost in the week's queue. It is also unreproduced, single-lab, with no fabrication files, no cost, and no stated release; the dome geometry is presumably the whole trick and it is not public.

Force as the action, not a feedback term

Four papers this week put force on the *output* side of the policy rather than the input side, which is a different claim from the contact-sensing thread and I am opening a thread for it (§9).

- **Wrench-ACT** (preprint, Siemens Research + UT Nuremberg): wrench as the *sole* action output into a pure force controller — not a hybrid position/force scheme. ACT backbone, five contact-rich tasks, matches or beats position baselines with gains scaling in how much deliberate force regulation each task needs. The ablation separates two things that are usually confounded: the wrench action space and **force-feedback bilateral teleop during data collection** contribute independently. That is the operationally important finding — you cannot learn force actions from position-teleop data. 1,000+ wrench-action demos promised on publication; no code stated.
- **FP2** (preprint, SJTU + UPenn + Flexiv + Noematrix): a lightweight downstream interface that splits task-level action generation from high-frequency force regulation on top of **frozen** robot foundation models — four RFM backbones, four real contact-rich tasks, novel-object generalization. No absolute numbers in the abstract, no code.
- **Impedance Cloning** (preprint): a particle filter recovers stiffness and equilibrium-point parameters from bilateral teleop **with no force sensor**, from one demonstration, on a CRANE-X7. Cheapest entry point here.
- **HACo** (preprint, HKU + BAAI + JHU): fingertip tactile plus joint torque, with the command/state discrepancy under compliance-regulated teleop used as the supervision signal and no online contact model. 83% vs 35% for the best baseline over a five-task real benchmark, 20 trials per task. No code.

The negative result you should keep

TACTIC (preprint): an identical-pipeline bakeoff of tactile encoders and conditioning schemes across **2,000+ real rollouts**, concluding there is **no universally optimal tactile representation**. No code and no encoder released, which is a shame for a paper whose value is the comparison. Together with ME-Dex's zeroed-tactile result from W39 and SpectRobot's blank-spectrogram probe, the tactile-literature picture is now: the representation choice is task-dependent, and the run-time contribution of the sensor is frequently near zero. Both of those are arguments for cheap sensing.

Shorter notes. DTWM (preprint, Yale + TAMU + UCLA + UW, repo `dextacwam`) conditions a video-diffusion world model on glove tactile and cuts hand-motion underestimation from 23% to 9% over 3 runs with a vision-matched control — and reports touch helping even when absent at inference, which is the same train-time/run-time split again. TaRL (preprint, NTU + Delta Electronics) regresses task progress from tactile deformation maps over successes *and* failures as an RL shaping reward: cube pickup 37→97% real, nut threading 34→56% sim, with cross-instance box-to-can transfer. Uni-VLaT (preprint) adapts VLAs to whole-body distributed skin with a tactile latent that also predicts future tactile, proprioception and visual state: 75% across five real tasks on two backbones, +43 pp. ReF-HIL (preprint, Tsinghua) is the week's online-correction entry — a value reference built from successful human experience plus a learned action fence that removes the imitation penalty inside it; 90% autonomy in 18–63 minutes on five real tasks, final 91.7–100%. No seeds stated, no code, so it does not clear the bar Res-HIL set last week. DITTO-X (preprint, Stanford + Columbia) is the data-rig item: an actuated exoskeleton that **back-drives the operator's fingers to the robot's state** before handover, rendering joint-level force and fingertip contact, hand-agnostic across Sharpa, Wuji and Inspire hands. No numbers in the abstract, no release stated. FlashDexRetarget (preprint, KAIST + Holiday Robotics) hits 90% over 50 human hand-object motions with ~100× less compute than RL baselines and claims scaling to 1,000 motions, with code and videos stated.

3 Platforms & embodiment

Demonstrated. Dyna Robotics' Taku + Dyna-2.1 (company tech report + uncut demo video, 2026-09-30) is the only platform item this week that cleared a real evidence bar: an **hour-long uncut autonomous laundry run**, published as uncut footage. Wheeled semi-humanoid, two 7-DoF arms, folding lower body, four steerable wheels; a three-layer stack of RL motion control, Dyna-2 manipulation and a VLM orchestrator with text memory, pretrained on a claimed million hours of human video, resuming after interruption. Length of uncut autonomy is a far better proxy for deployability than any success-rate table, and almost nobody publishes it. No code, weights, data, benchmark or pricing; "first of its kind" is self-assessed.

Nucleus (demo video, 10-01) released nearly two hours of factory footage with interventions left visible and **disclosed a 60% autonomy ratio** — rare and creditable, though the basis (time or tasks?) is unspecified and operator attention per robot is unknown.

Staged or vendor-asserted. Boston Dynamics' GR3 hands for Atlas (press release + sponsored content, 10-01): 13 DoF up from 7, four fingers and no pinky, a 4-DoF thumb plus three 3-DoF fingers, encapsulated direct-drive joints with one actuator type per joint, no exposed cables, **dense pressure tactile across fingertips and palm**, with sim-trained transfer claimed. Architecturally this is the most interesting hand of the week — direct-drive with dense pressure is a deliberate bet against tendon-plus-sparse-taxel — and there is not a single task success rate attached to it. Runway's Praxis-1 (company blog, no paper) claims the result that would matter most if true: after finetuning, **web video gives 16.1 cm placement error against 16.0 cm for teleoperated robot video** (±1 SEM over 93 evaluation pairs), with performance still improving through 10^2–10^3 hours of pretraining video, across Noble Machines bimanual, a Standard Bots RO1 6-DoF arm and an Ultra mobile base. They disclose four failure categories (rigid repeated objects, clutter with occlusion, transparent materials, deformables). **Open weights are pledged, not released** — "coming months," no license stated, no paper, single-lab self-reported. Flagged unreproduced. If the weights land, this is the most consequential item in the queue.

Priced hardware. Astribot T1 at IROS (press release, 10-01): **$18,000**, 23 DoF excluding end effectors, 1.55 m / 66 kg, cable-driven, 5 kg per arm, hot-swap battery, modular end effectors — and, unusually, a **documented SDK and API for joint, Cartesian and whole-body control**, with US orders open. Galbot ET1 (press release, 10-02) from 79,000 yuan across three compute tiers (RK3588/RDK X5 → Orin Nano 40 TOPS → NX 100 TOPS), 1.25 m / 32 kg, 26 DoF base — note that the **12-DoF dexterous hands are an optional add-on not in the headline price**, and there is no reliability or success-rate data. Neither vendor published a benchmark.

The counting problem got worse, not better. IDC puts H1 2026 humanoid shipments near 25,000 with AGIBOT leading (third-party count, no published methodology), against the IFR's ~7,000 humanoid *sales* worldwide in all of 2025. These are different quantities measured by different people with unmatched definitions, and neither reconciles with vendor production claims. Sources conflict; I am not picking one. Tesla's Texas Optimus plant is a drone-observer estimate of 40% steel completion from a single non-Tesla source — floorspace is not throughput, and it tells you nothing. Also this week: Agility and FORT signed an MoU on an offboard safety bridge for Digit 5 with **no standard named, no certification and no date**, and Singapore's HTX opened a humanoid training centre with a dated 2028 public-safety deployment target.

4 Open source & tooling

Highest friction reduction this month:

1. **The benchmark fixes from arXiv:2609.37771** — 22 bug repairs and revised settings across LIBERO, LIBERO-PRO, LIBERO-Plus, RoboTwin, RoboCasa, VLABench and RoboDojo. The paper states these are released and that PRs are submitted upstream, but **gives no repository URL in the text**, which is a problem for the one artifact this week that most people need. Until the PRs land, the *list* is still immediately actionable: pin MuJoCo to 3.3.2 rather than 3.8.1 (worth +26 pp on one LIBERO task on its own), read task language from BDDL rather than filenames, and reseed instruction sampling. Full detail in §6.

2. **Yeah Hand** (10-03, ROS Discourse): a 15-DoF tendon-driven hand from **five** Waveshare ST3215HS bus servos, 3D-printed TPU fingers and palm, ESP32 firmware over Bluetooth serial, and **per-finger torque sensing used for adaptive grasp with per-servo torque limits exposed as ROS 2 parameters**. Driver tested on Jazzy. ISO 9409-1-50-4-M6 flanges (UR, KUKA, FANUC, Techman, Yaskawa) plus Agilex. Hardware CERN-OHL-S-2.0, software GPL-3.0, CAD + firmware + driver all released at opsobot/yeah-hand and opsobot/yeah_hand_ros, build docs on Hackaday, kits on Tindie. The gap: **no evaluation of any kind** — no grasp success rate, no force accuracy, no cycle life, no BOM cost — single builder, no independent reproduction, and `ros2_control`/MoveIt integration solicited rather than built. Torque as a first-class per-finger signal at hobby cost is directly on your path; treat the hardware as real and every performance claim as absent rather than unfavorable.
3. **Three WAM codebases landed:** hustvl/ReWAM, breez3young/VJEPA-Policy, and uni-vta/UVTA. V-JEPA Policy is the only one trainable at your scale.
4. **easyminnn's tactile corpus** — six public tactile manipulation corpora converted into one loader format (FreeTacMan 5,627 / TitW 3,699 / OpenTouch 2,958 / EgoTouch 1,928 / exUMI 1,452 / VLA-Touch 380 episodes), spanning GelSight, taxel arrays and pressure gloves. LeRobot **v2.0, not v3.0**; mixed licenses; solo uploader; zero downloads; unreplicated. The value is the conversion work, and the mixed licensing means you check per-corpus before using any of it.
5. **FineART** (preprint, Scale AI + Hugging Face + USC + Stanford) released dataset, weights and code together: 40,543 episodes / 1,718 hours / 533,913 annotated subtasks / 151 tasks, with self-predicted-next-subtask mid-training taking spatial disambiguation 32→100% and unseen long-horizon 16→76%, and one-tenth data on a new robot. Bimanual, and the license is unstated — check before you build on it.

Release state of the stack, honestly reported. MuJoCo is **still 3.14.0** — the 10-04 sweep got a clean 404 on the 3.15.0 tag, which is a trustworthy negative, unlike the repo's `/releases` index page, which renders stale versions and dates. LeRobot reads as v0.6.1 (2026-08-03) but the index page served 2024 dates, so it is **UNVERIFIED, not flat** — and PyPI is unreachable from the harvester, so there is no cross-check. Nine weeks since the NVIDIA/Hugging Face deal and still no release; the v0.7.0 roadmap is an open community issue naming eight workstreams (runtime and inference, RL and reward models, hardware and sensing, data engine, eval and sim benchmarks among them) with in-flight PRs for an ONNX export guide and eval comparison. ROS 2 has **Lyrical and Jazzy sync rounds staged for tomorrow, 2026-10-05** — Lyrical, Jazzy. Isaac Lab, Genesis, ManiSkill and RoboCasa are **UNVERIFIED this week** — every release surface for them was blocked. Do not read any of that as "nothing shipped."

---

5 Money & moves

Two lines each; the second is the interpretation.

AMD agrees to buy World Labs for ~$8.2B in stock, Fei-Fei Li joining as chief scientist (press reports ~09-28/29, AMD IR confirmed, Bloomberg/TechCrunch/CNBC corroborated; no robotics product named). **INFERENCE** \2192 fourth consecutive week the acquirer is a compute or platform holder rather than an operator — NVIDIA/Hugging Face (W36), SoftBank/RAI (W38), Qualcomm/PickNik (W39), now this. The chip companies have decided the scarce asset is the world model and the software layer above it, because that is what pulls developers onto their silicon. The only part that touches your build is whether World Labs' asset-release policy survives.

Microagi acqui-hires a stealth team and names Animesh Garg Chief Research Officer (trade report 09-30; corroborated — a Georgia Tech professor plus seven researchers into one German team, terms undisclosed, framed as robot post-training for factory work). **INFERENCE** \2192 eight academics moving as a block says more about where value sits than any round this week, and the named function is **post-training**, not pretraining — consistent with base models commoditizing and the margin moving to adaptation.

Tangent Robotics raises $4.5M pre-seed (trade report 09-30). Columbia spinout — Ciocarlie, Piacenza, Kymissis — Fly Ventures and Toyota Ventures leading; optical contact sensing at the fingers, targeting assembly, threading, connector mating, gasket seating, snap fitting and gear meshing. No specs, nothing released. **INFERENCE** \2192 Ciocarlie's line that copying the human hand is "akin to building flying machines with flapping wings" is a funded bet against the anthropomorphic consensus

Boston Dynamics' GR3 and Unitree's Dex5-S just doubled down on. Note the task list is identical to every force-control paper in §2 — that is where near-term commercial manipulation money actually is.

Figure's multibillion-dollar compute commitment meets Nscale's data-centre buildout (press report 09-30; index headline, article slug unresolved by the harvester). **INFERENCE** \2192 when the binding constraint on a humanoid company's roadmap is someone else's concrete, the company has stopped believing the remaining gap is algorithmic.

T-Robotics leads a KRW 4.5B (~$3M) humanoid project for Korean display fabs (press release 09-29), alongside Singapore's HTX centre and its dated 2028 target (§3). **INFERENCE** \2192 national programs keep funding facilities, deployments and data libraries rather than models — Korea's MSIT deliverable in W39 was a physical-data library, Singapore built a training centre, this buys factory integration. Public money is going to the input side, consistently.

---

6 Deep read of the week

Faster and Better? Benchmark Bugs and Design Limitations Distort the Evaluation of Vision-Language-Action Acceleration

— arXiv:2609.37771v1, preprint, Shanghai Jiao Tong University (Chen, Zhou, Shi, Li, Feng, Gu), announced 2026-09-29.

Chosen over Simple-WAM and UVTA because it is the only paper this week that changes the *reliability of numbers you already believe*, and because its findings are actionable without reproducing anything.

Setup

The task is auditing, not manipulation. The entry point is an anomaly: **training-free** acceleration methods — which by construction only approximate the baseline's computation and add no capability — repeatedly score *higher* than the baseline. A success rate alone cannot distinguish "the approximation improved execution" from "the checker scored the altered behavior more generously." Embodiment and policy are the controlled variables: community-released π0.5 checkpoints throughout, on A800 and RTX 4090. The accelerated variants are DP-Cache-Fast (S=5, 2.56×), DP-Cache-Slow (S=2, 1.62×) and ProbeFlow (ε=0.008, 1.49×), with measured baseline latency 313.73 ms/chunk at batch 1 and horizon 32. Seven benchmarks audited: LIBERO, LIBERO-PRO, LIBERO-Plus, RoboTwin, RoboCasa, VLABench, RoboDojo. What is held out is nothing — this is not a generalization study; what is varied is the *benchmark's own implementation*.

Method

The genuinely new piece is a cheap observability substitute. Recording every rollout costs ~85% wall-clock overhead (160.4 s/episode with recording against 86.4 without, on A800). So instead they plot logged object and end-effector trajectories plus gripper states **against the checker's acceptance region**, mark the success-trigger time, and only render video where that view flags something. Everything else is inherited: manual code inspection, domain-expert verification, and a propagation step that inspects every task sharing an affected template, loader, checker or scene utility — including tasks with no anomalous gain, which is what turns anecdotes into a count.

The two worked examples define the failure taxonomy. *False baseline failure* (RoboTwin `place shoe`): the prompt says place the footwear on the mat; the checker silently demands a toe orientation the prompt never states, and the accelerated variant happens to reproduce the rule-based demonstrations' left-facing placement. Both shoes are on the mat; only one passes. *False acceleration success* (RoboTwin `put bottles dustbin`): the accelerated policy never grasps the bottle, knocks it off the table, the bottle rolls into the acceptance region, and the checker counts a success.

The 22 bugs partition into **task consistency (9)** — semantic omission (6) and semantic conflict (3); **initialization (6)** — prompt pollution (3) and invalid initial states (3); and **reproducibility (7)** — incomplete evaluation configuration (4), history-dependent scene construction (2), incomplete state restoration (1). The four design limitations are permissive checkers, threshold sensitivity, unrealistic object masses, and action smoothing that differs from real robots and therefore hides the jitter acceleration adds.

Results

The number that matters: on LIBERO-PRO's Spatial-Task, baseline success goes from 0.4% to 53.2% after one bug fix — last place to first. The bug is not subtle: the loader derived the prompt from a stale filename instead of the updated BDDL language field, so the policy was told to turn the stove *on* while the checker required it *off*. Every method scored ~0.4–0.6% before; every method scores ~52–53% after. The entire published task was measuring a loader defect.

The rest, with baselines:

- RoboTwin `rotate qrcode`: baseline 56 → 66, from 3 pp behind ProbeFlow to 6 ahead.
- RoboTwin `handover block`: 18 → 29. `blocks ranking size`: 21 → 27.
- VLABench `select mahjong`: baseline +17 pp against +6–10 for the accelerated methods; `select painting` +13.
- **Mass correction**, RoboTwin `place fan` (fan mass 10 g → 600 g): baseline 40 → 38 while DP-Cache-Fast falls 61 → 33 — **21 pp behind to 5 pp ahead**. Retraining at corrected mass takes the baseline to 53. `beat block hammer` (hammer 1 g → 600 g) collapses everything: 79/87/82/86 → 27/19/19/22.
- **LIBERO mass**: `Mug Placement in Microwave` (30.91 g → 300 g) 94/98/94/98 → 60/56/52/50.
- **Motion-aware scoring** (rank methods by pre-post-processing action variation, failures score zero): `scan object` moves the baseline from 5 pp behind ProbeFlow to 16.6 ahead; `turn switch` holds the baseline at 25 while DP-Cache-Slow drops 31 → 24.8.
- **MuJoCo version**: unpinned MuJoCo changes rendered policy inputs. Pinning 3.8.1 → 3.3.2 takes π0.5 on `pick up bowl` from 62% to 88%.
- **RoboCasa `ArrangeDrinkware`**: source and destination counters could resolve to the same object, so the task was *already successful at initialization*. After resampling, π0.5 goes from 20% to **0%**.

Where it's soft

- **No repository URL anywhere in the paper.** It says fixes are released and PRs submitted, and the whole value proposition is the artifact. That is the single biggest weakness.
- **Single-policy audit.** Everything is π0.5 and three training-free accelerators. Whether these checkers are equally unfair to a different policy class is untested, and some of the reversals plausibly turn on π0.5's particular behavior rather than on the fix.
- **No seeds, no intervals, anywhere.** Tables report bare percentages at n=100 episodes for RoboTwin and n=500 for LIBERO-PRO subsets. A 2-point reversal (`put bottles dustbin`: baseline 24/24 vs DP-Cache-Slow 28→22) is well inside the noise they never quantify. The large reversals are safe; the small ones are presented with the same confidence and should not be.
- **The mass corrections are unjustified.** Table 9 shows a hammer going 1 g → 600 g and a fan 10 g → 600 g with no measurement, no reference object, no sourcing. "More realistic" is asserted. A 600× change that collapses all four methods by 50+ points is a large intervention resting on the authors' judgment, and it is not the same kind of claim as a loader bug.
- **The motion-aware score is the authors' own invention and it is rank-based**, with failures assigned zero. That construction will favor a slower, smoother baseline almost by design. It may be the right metric for deployment; it is not a neutral referee, and reporting it beside success rate rather than as a replacement would have been more honest.
- **No ablation of the audit method itself** — no estimate of recall. They found 22 bugs; nobody knows the denominator, including them. Their own framing (propagate by failure mechanism) implies the count is a lower bound.
- Asymmetric compute is not an issue here (all methods share a checkpoint), which is one thing that is clean.

Steal list

1. **Plot trajectories against the acceptance region instead of recording video.** 85% overhead avoided, and it is strictly more diagnostic than a video for the question "did the checker fire for the right reason." This is a 50-line matplotlib utility and it belongs in your eval harness permanently, not just when something looks wrong.

2. **Make the anomaly the entry point.** A training-free approximation beating its own baseline is a *logical* impossibility signal, not a result. Any time a strictly-less-capable variant wins, audit the harness before you write it up. Generalize: any method that cannot in principle add capability but adds success is measuring your evaluation.
3. **Pin your simulator, and read task language from the source of truth.** Two of their highest- impact fixes are trivially portable: pin MuJoCo (26 pp on one task) and never derive a prompt from a filename. Audit your own loaders for the same pattern — it is the single most common bug class in their taxonomy (prompt pollution + semantic omission = 9 of 22).
4. **Verify success is not already satisfied at t=0, and that your seed actually controls the scene.** `ArrangeDrinkware` scored 20% on initialization alone. Their Advice #3 is the useful formulation: a shared seed does not establish that two runs instantiate the same task — test state restoration and action replay *separately* to find out which parts of an episode you can actually reproduce.

Cost to reproduce a scaled-down version

You do not need to reproduce this; you need to apply it. A worthwhile scaled-down version is a one-policy, one-benchmark audit of whatever you plan to cite:

- **Hardware:** none. Sim-only. One GPU.
- **Episodes:** their protocol is 100 episodes per task per condition. Pick the 5 tasks you would actually put in a table; two conditions (your policy, one cheap variant) × 5 tasks × 100 episodes = 1,000 rollouts. At their measured 86.4 s/episode on an A800 that is ~24 GPU-hours single-threaded, call it **8–12 GPU-hours** parallelized on one modern card, and under 4 if you drop to 50 episodes for the screening pass.
- **Engineering:** the trajectory-vs-acceptance-region plotter is an afternoon. The audit itself is human time, not compute — budget **2–3 days** of reading checker code for 5 tasks, which is the honest cost and the reason nobody does this.
- **GPU-hours for the re-training leg** (their corrected-mass retrain, which is the only part that needs training): one π0.5 finetune per corrected task. Skip it; the inference-only comparison carries the finding.
- **Total: ~12 GPU-hours, zero hardware, ~3 days of your attention, no new data.** This is the cheapest high-value piece of work in the queue, and it is work you would do anyway before citing a number in an interview.

---

7 Relevant to what you're building

`CONTEXT.md`'s `## Active build` section is still an unfilled placeholder — fourth week running — so there are no workstreams to map this week's items onto. Section skipped rather than invented.

---

8 Market signal

Boards checked 2026-10-04. (The harvester did not check boards on any sweep this week; these are read fresh today.)

- **Figure: 97 roles, 13 on AI–Helix** — identical totals to 2026-09-27, and identical composition (Modeling, Pretraining, Video Pretraining, Reinforcement Learning, Robot Learning, Perception, Generative AI, Data Infrastructure, Training Performance, Localization and Mapping, Android, Backend, XR). Plus a separate Reinforcement Learning Engineer – Whole Body Control under Controls.
- **OpenAI: 22 robotics roles**, against 23 on 09-27 and 21/26/27 from three sources on one day in W38 — so a 1-role delta is rendering noise. **Third consecutive week with zero titles containing policy, VLA, imitation or robot learning.** Composition still actuators, gears, motors, dynamometer and test infrastructure, thermal, PCB, firmware, commodity management. Two software exceptions worth noting: Inference Engineer, Robotics and Software Engineer, ML Systems & Training Architecture.

- **Physical Intelligence: 12 roles**, down from 13. New since 09-27: Robotics Software Engineer, Business Operations, Mechatronics Intern. Still **no policy-research titles**.

Comp datapoints, primary postings, base plus equity:

- **OpenAI, Machine Learning Engineer, Distributed Data Systems – Robotics**: **$380K–$500K + equity**, SF, hybrid 3 days. Unchanged for three weeks.
- **New this week, and the useful one: OpenAI, Software Engineer, Distributed Data Systems – Robotics**: **$230K–$385K + equity**, same team, same office policy. So the same function carries a ~$150K spread at the top of band between the "ML Engineer" and "Software Engineer" ladders. If you are negotiating into a data-systems-adjacent robotics role, the title is worth more than the work.
- **Figure, Helix AI Engineer, Modeling**: **$200,000–$400,000 base**, San Jose, 5 days in office. Unchanged. Description is explicitly model architectures for perception, reasoning and action, representation learning, and generalization/robustness.
- **Physical Intelligence** publishes exactly one rate: **Robot Operator at $25/hour + benefits**. Every engineering role withholds compensation.

INFERENCE \2192 nothing moved. Figure remains the only lab hiring policy and model people in volume and the only one publishing a band for them, and that band still tops out $100K below OpenAI's band for the data-pipeline engineer. Three weeks of identical composition at OpenAI is now a signal rather than a snapshot: whatever they are building, the hiring surface is a hardware program with an inference team attached.

Skills recurring in work that got attention this week:

- **Separating what a mechanism contributes from what it costs at inference.** Simple-WAM's `t_denoise` sweep, ReWAM's gradient-routing ablation, UVTA's tactile-prediction ablation (70→42), and K-MF reporting end-to-end latency separately from head latency are all the same discipline. The specific move — *change one thing at inference with no retraining and show the curve* — is the cheapest credibility you can buy and four papers did it this week.
- **Auditing an eval before trusting it.** This was differentiating in W39 and after arXiv:2609.37771 it is becoming table stakes to at least know which benchmarks are compromised. Being able to say "LIBERO-PRO's Spatial-Task was a loader bug and the real number is 53%, not 0.4%" is precisely the kind of thing that separates someone who read the paper from someone who read the leaderboard.
- **Force and torque on the output side of the policy.** Wrench-ACT, FP2, HACo, Impedance Cloning, plus W39's whole-hand QP force regulation. The EE-shaped skill from W39 has moved from "estimate contact without a sensor" to "command wrench as the action," and the data- collection corollary (you need force-feedback bilateral teleop to collect it) is a hardware problem, not a modeling one.
- **Mechanically encoding information instead of measuring it.** TacGB's domed film is the purest instance — shear without electronics, calibration or alignment. Same family as SpectRobot's off-surface mounting from W39.

Table stakes, updated. Released code at submission — this week ReWAM, V-JEPA Policy, UVTA, FineART, GraspTwin and Yeah Hand cleared it; Simple-WAM, TacGB, Wrench-ACT, FP2, HACo, TaRL, ReF-HIL and the benchmark audit itself did not. Latency reported end to end, not just at the head. An ID/OOD split. Multi-seed paired evaluation.

Still differentiating:

A controlled protocol with matched backbone, data *and* budget across the configurations you compare (Simple-WAM is the week's best instance, and it is the reason to read it even though the model is out of reach); an inference-only intervention sweep; publishing a pretraining *scaling curve* rather than a single data point (IronMind's 5.0 → 11.7 → 55.0% across 0/5k/10k hours); and disclosing uncut autonomy duration or an autonomy ratio rather than a success rate (Dyna, Nucleus). Seed bands on real hardware remain rare enough that Res-HIL from W39 is still the only clean example this quarter.

9 Threads

Updated `THREADS.md`; board as it now stands.

THREAD	STATUS	WHAT WOULD COUNT AS THE NEXT REAL UPDATE
World-action models as unified policy + simulator	**advancing (sharply)**	Three convergent results this week. Simple-WAM: fully-noised future tokens recover +26.4/+8.2/+13.4/+89.8 over the latent paradigm, nine further denoising passes worth ≤+1.8; learned or zero tokens collapse task generalization to ~10. ReWAM: 93.6% RoboTwin 2.0 with no generative video pretraining, DINO > Video-VAE despite worse SSIM, calibration worth 12.2 pp, AGRS routing worth 6.3 pp, 600 h vs 10,000+ h pretraining. V-JEPA Policy: LIBERO-Plus 79.3 from 0.6B trainable on a frozen encoder. All three independently confirm W39's What Matters temporal-order finding and contradict Fast-WAM's "training objective only" claim. Contested: Simple-WAM has no code and rests on a 5B backbone without embodied pretraining (authors say so); ReWAM's Open instruction-following is 0.25 vs π0.5's 1.98. **Next:** whether the noised-future result holds at scale or with embodied pretraining; a released WAM used for policy improvement by a group that did not build it (OpenWAM HF repos still untouched since 09-09, five weeks).
Benchmark validity in physical AI	**advancing (decisively)**	arXiv:2609.37771 audits seven benchmarks, finds 22 bugs + 4 design limitations, and shows a LIBERO-PRO task where the baseline goes 0.4% → 53.2% (last to first) because the loader read the prompt from a stale filename. Mass correction flips a 21-pp deficit to a 5-pp lead. Pinning MuJoCo 3.8.1 → 3.3.2 is worth 26 pp on one task. RoboCasa's ArrangeDrinkware scored 20% at initialization. Prior: REBOOT, RoboRecover, The Gaussian Is Enough, What Matters. Contested: the audit is single-policy (π0.5), has no seeds or intervals anywhere, and its 600× mass "corrections" are asserted rather than measured; no repo URL in the paper. **Next:** the fixes appearing as merged upstream PRs with a URL; and the still-unmet bar of a major-lab table with seed bands *and* paired evaluation together.
Force as a first-class action	**NEW — opened 2026-W40**	Four papers this week put force on the policy's output side rather than its input side. Wrench-ACT: wrench as sole action into a pure force controller, five contact-rich tasks, matches or beats position baselines, and the ablation shows the wrench action space and force-feedback bilateral teleop contribute *independently*. FP2: downstream force interface splitting task-level generation from high-frequency regulation on four frozen RFM backbones. Impedance Cloning: stiffness and equilibrium point recovered by particle filter from bilateral teleop with no force sensor, one demo, CRANE-X7. HACo: 83% vs 35%, command/state discrepancy as supervision. Plus W39's whole-hand QP force reallocation. Contested: none has code; Wrench-ACT's 1,000+ demos are promised on publication; FP2 publishes no absolute numbers. **Next:** any of these four releasing data or code; or a wrench-action policy trained on *position*-teleop data, which would falsify the collection-method claim.
Contact sensing without dedicated tactile hardware	**ADVANCING**	TACTIC's 2,000+ real rollouts find no universally optimal tactile representation (no encoder released). TacGB encodes shear *mechanically* into a normal-only sensor's pressure map — no electronics, no calibration, no alignment — for up to +36 pp on insertion. Prior: ME-Dex's zeroed input, SpectRobot's off-surface sensing, CoPRE's motor-current-only detection, Real-Time Force Regulation's SDF-plus-proprioception. Contested: TacGB is single-lab, unreproduced, with no fabrication files or cost. **Next (unchanged, five weeks):** an estimated- or predicted-force policy matching a real-sensor baseline on a task whose labels were *not* collected with that sensor.

THREAD	STATUS	WHAT WOULD COUNT AS THE NEXT REAL UPDATE
Paying at training time, not run time	**ADVANCING**	UVTA is the strongest instance yet and it uses *real* human tactile rather than simulator-generated: 1,000 human + 150 robot demos per task, 70% vs T-Rex's 29%, dropping the future-tactile prediction objective costs 28 points (70 → 42), and the human-data curve runs 0 → 84% with no saturation at 1,000. DTWM reports touch helping even when absent at inference. Contested: UVTA still reads tactile on the robot at deployment, so it removes the *collection* cost, not the sensor. ME-Dex's zero-tactile result is still only shown on simulator-generated tactile. **Next:** whether ME-Dex's zeroing survives with real sensor data; and a matched comparison of one-step heads (IMLE vs drifting vs consistency vs MeanFlow) on one backbone — K-MF adds a fourth candidate and still nobody has run the bakeoff.
Does VLA scale buy language-conditioned control?	**ADVANCING**	Causeway is a *training-free* write into a frozen policy's action-stream representation that takes LIBERO-Goal instruction switching from 3–26% to 47–65% over 71 instruction pairs, on real xArm — a far cheaper intervention than GAM's observation masking for the same failure mode. FRAM gets 92.2% LIBERO from 138.7M params against π0's 3.3B at 94.2%, wrist cameras only. Contested: Causeway has no code; FRAM is 138.7M, so the **sub-10M real-hardware multi-task bar is now six weeks unchanged**. The LIBERO-PRO Task column that GAM's claim rests on was itself a loader bug per §6 — GAM's 0.88 vs ≤0.11 needs rechecking against the fixed benchmark. **Next:** GAM or Causeway re-run on the *repaired* LIBERO-PRO; a sub-10M real-hardware multi-task policy against a published baseline.
Online correction as the last mile	**ADVANCING**	ReF-HIL (Tsinghua) builds a value reference from successful human experience plus a learned action fence that suppresses the imitation penalty inside it: 90% autonomy in 18–63 minutes on five real tasks, final 91.7–100%. Contested: no seeds stated and no code, so it does not clear the bar Res-HIL set with 3 seeds × 50 fixed inits. Res-HIL still has no code. EXPO-FT's released code is now **three weeks out with no third-party rerun**. **Next:** a third-party rerun of EXPO-FT, or code from Res-HIL or ReF-HIL.
Low-cost open hardware as training substrate	**ADVANCING**	Yeah Hand: 15 DoF from five bus servos, per-finger torque sensing exposed as ROS 2 parameters, CAD + firmware + Jazzy driver released under CERN-OHL-S-2.0 / GPL-3.0, with ISO 9409-1 flanges for UR/KUKA/FANUC/Techman/Yaskawa. The Antagonistic Tendon-Driven Hand (2609.36241, TAMU + Sogang) reports 17 DoF at 220 g and 200 mm with every DoF crossing neutral, 3.0% transmission error, 13.2 Hz joint bandwidth, 29 N fingertip force, 0.15 mm repeatability, Kapandji 8 and all 33 GRASP taxonomy types — **no design files**. Astribot T1 at $18,000 with a documented SDK; Galbot ET1 from 79,000 yuan with hands as a paid add-on. Contested: Yeah Hand has literally no evaluation and one builder; Cartesian Hand files still unreleased; still no third-party build of Aero Hand, BRIDGE or Peg-in-Bench reported by anyone. **Next:** a clearly independent group reporting success rates on released low-cost hardware with code. The Yeah Hand is now the cheapest candidate for someone to do that to.

THREAD	STATUS	WHAT WOULD COUNT AS THE NEXT REAL UPDATE
Claimed open releases vs actual artifacts	**narrowing holds (second week)**	Actually released: ReWAM (code), V-JEPA Policy (code), UVTA (code), FineART (dataset + weights + code, license unstated), Yeah Hand (hardware + firmware + driver), GraspTwin (code), easyminnn's six-corpus tactile conversion, FlashDexRetarget (code + videos stated). Still pending: Simple-WAM (project page), TacGB, Wrench-ACT (data on publication), FP2, HACo, TaRL, ReF-HIL, IronMind, RPG, the benchmark audit's own fixes (no URL). **Runway's Praxis-1 open weights are pledged for "coming months" with no license** — a new entry on the pending list and the most consequential one. StreamPI official weights are now **six weeks** pending. **Next:** whether the narrowing holds a third week; Praxis-1 weights; StreamPI weights.
Who owns the open robot-learning stack (merged with the former *Manipulation acquired rather than built*)	**ADVANCING**	AMD agrees to buy World Labs for ~$8.2B in stock with Fei-Fei Li as chief scientist — the **fourth consecutive week** in which the acquirer is a compute or platform holder rather than an operator (NVIDIA/Hugging Face W36, SoftBank/RAI W38, Qualcomm/PickNik W39). Microagi acqui-hires eight academics and names Animesh Garg CRO for robot *post-training*. LeRobot is nine weeks post-deal with no release and an unverifiable release page; v0.7.0 is an open roadmap issue. Contested: still no governance structure named in any of the four commitments — press quotes, not foundations. **The W35 in-quarter prediction (a logistics incumbent buying an end-effector or policy team) has ~3 weeks left and has not happened once; four platform/capital acquisitions have.** Kill the prediction on lapse and keep the consolidation framing. **Next:** any commitment written into a foundation or independent steering body; the first post-deal LeRobot release and whether GR00T or Isaac integrations get privileged treatment in it.
Simulation fidelity for transmission and contact	**stalled**	MuJoCo is still 3.14.0 — the 3.15.0 tag 404s, which is a trustworthy negative. Isaac Lab, Genesis, ManiSkill and RoboCasa are **UNVERIFIED** (release surfaces blocked). PneuTac (2609.38418, Oxford) is the one new item: MPM soft-body plus a deformable tactile membrane with 3DGS rendering and surrogate models in one simulator, three real tasks, no release stated. Still no sim-to-real result crediting MuJoCo's `ipc` flex contact or the discrete integrator; MuJoCable still has no repo. §6 adds a sharper point: **unpinned MuJoCo versions change rendered policy inputs enough to move success 26 pp**, so simulator fidelity is now also a reproducibility problem, not only a physics one. **Next:** a tendon-hand or underactuated-finger sim-to-real result attributing transfer to IPC flex contact; a MuJoCable repo.
Humanoid capital vs demonstrated capability	**ADVANCING**	Dyna's Taku publishes an **hour-long uncut autonomous** laundry run — the strongest evidence bar any humanoid-adjacent platform has cleared in this ledger — and Nucleus discloses a 60% autonomy ratio on ~2 hours of uncut factory footage. Against that, the counting problem got worse: IDC puts H1 2026 shipments near 25,000 while the IFR counted ~7,000 *sales* worldwide in all of 2025, with unmatched definitions. Boston Dynamics' 13-DoF GR3 hand with dense pressure tactile ships with zero task benchmarks. Tesla's Texas plant is a drone observer's 40% steel estimate from a single non-Tesla source. Contested: units built, shipped, sold and working remain four different quantities, and vendors report only the first two. **Next:** a humanoid company first-authoring a method with released code (still zero); or an independent count of humanoids in productive operation. Uncut-autonomy duration is now the best available proxy and two companies published it this week — track whether that becomes a norm.

Threads at 12 (cap). *Manipulation acquired rather than built* was merged into *Who owns the open robot-learning stack* — its W35 prediction is carried forward intact with its kill date, as W39 specified — and that slot went to *Force as a first-class action*.

10 Open question

Does the noised-future result survive embodied pretraining, or is it an artifact of starting from a video generator that was never shown a robot?

Simple-WAM's finding is clean and the mechanism test is good: a fully-noised future beats a dropped future by 26 points under perturbation, and substituting learned tokens or zeros for the noise collapses task generalization from 97.2 to 9.8. So the action expert is reading something real out of a Gaussian input. The authors' own closing sentence names the limit — their conclusions rest on a 5B backbone **without** embodied pretraining, and whether they hold at larger scale or with it remains open.

Here is why that limit might be load-bearing rather than pro forma. The feature the action expert reads at $\tau = 1$ cannot encode anything about *this* scene's future, because the input carries no information about it. What it can encode is whatever the video expert's weights say about futures *in general* given the clean current frame that sits beside the noise in the same sequence. That is a prior, not a prediction. A backbone with no embodied pretraining has a generic web-video prior, and a generic prior is exactly the thing you would expect to transfer robustly under perturbation — nothing in it is specific enough to break. Give the same backbone 10,000 hours of robot data and that prior sharpens into something scene- and embodiment-specific, at which point actually denoising it might start to pay, and the ordering in Fig. 2 could invert. ReWAM's one weak column is suggestive in the same direction: the model with no generative video pretraining is the one that falls apart on open instruction following (0.25 vs π0.5's 1.98). The two papers may be measuring the same boundary from opposite sides.

What would settle it:
The same `t_denoise` sweep — 1, 4, 6, 8, 10 — run on a WAM that *has* embodied pretraining, with everything else matched. Fast-WAM's codebase supports both paradigms and ReWAM published a 600-hour pretrained checkpoint protocol, so the comparison is a configuration change plus inference, not new training. If the single pass still captures the gain, the finding is about world-action models and you should change your inference loop today. If the curve flattens or reverses, the finding is about *un-pretrained* video priors and the field has been measuring its own initialization. That is one sweep, no retraining, and it is the highest-information experiment available this month at any budget.

ALSO THIS WEEK (4+ IN QUEUE, NOT COVERED ABOVE)

Nothing scoring 4 vanishes silently. All preprints unless noted; all abstract-level unless stated otherwise in the queue.

- InternW0-Delta — MoT WAM with "Causal Imprint" for predictive features without inference-time video rollout; claims the largest open corpus at 20K+ hours. Release promised, no numbers in abstract.
- RAPID — one visual human demo bootstraps verified reusable programs; 75.9±2.4% vs 14.6±2.5% for a CodeAsPolicy agent, real Franka + LIBERO-PRO. No code.
- Morphometric Imitation — contact-preserving hand retargeting plus residual RL; **89.3% zero-shot real over 300 trials, 30 objects**. Project page only.
- DA-GRD — sparse tactile probes with no re-vision; 84.7% lift vs 9.1% for stale AnyGrasp, 72.5% fewer probes. No code.
- ZeroBot — image-to-3D-mesh real2sim with a contact-biased action space; 87% real success, **119 s mean training**, zero demos, zero pretraining. Imperial + RAI Institute. No code stated.
- X-Reset — human hand-object states as an RL *reset distribution* rather than an imitation target; zero-shot sim-to-real, 22-DoF hand, three embodiments, 20 objects.
- DexRoam — tracker-free consumer-VR whole-body egocentric capture; GR00T 29→56%, π0.5 32→57%, halves required robot demos. No release stated.
- GraspTwin — single RGB-D digital twin, foundation-model affordance priors seeding Bayesian optimization; +33% task-oriented grasp success. **Code released.**

- Object-centric Tactile Interactive Perception — CMU; token learner selects tactile segments with contrastive text alignment; 41→92% seen, 19→67% unseen, three real tasks.
- SimpleICL — defines what a visual prompt is supposed to convey; minimalist in-context framework, no massive pretraining, low-cost collection protocol. Full open source promised (data + training pipeline). No numbers in abstract.
- Ego4WAM — fixed WAM backbone isolates alignment, duration, diversity and supervision in egocentric human data; video-only supervision works without action labels. Horizon Robotics, no code stated.
- GPT-6 Astra as embodied policies — frontier model as policy across six domains; **latency and token cost are the binding constraint, not decision quality**. (Answered in part by PyRUA-Lean, 2610.01939: 63.1→71.7 at equal LLM-call budget, 49% fewer calls, 65% fewer input tokens.)
- RoboCoach — imagined failures decide which subtask demos to buy; imagined-vs-deployed success correlates ρ=0.840.
- λ-0 / HumanVerse-500 — 500-hour wearable egocentric loco-manipulation dataset plus a three-stage whole-body VLA; code, models and data promised, robot unnamed in abstract.
- TacDyn-WAM — predicts implicit tactile *dynamics* rather than tactile pixels, multi-horizon deltas in one forward pass; UniVTAC 81.5, five real tasks 71.0–85.0. Directly contrasts with ME-Dex's zeroing result. No release stated.
- MikeHan517/robotwin-eef-lerobot — RoboTwin five-task end-effector conversion, LeRobot v3.0, MIT licence, 2,750 episodes / 856,023 frames. Bimanual, train split only, solo uploader.

Rate this issue 1-5 and tell me what was noise.