BETWEEN GLANCES: GEOMETRIC MEMORY FOR ADAPTIVE ACTION CHUNKING

Anonymous authors
Paper under double-blind review

ABSTRACT

For a robot acting between observations, each prediction must continue a movement and determine how far it should proceed before feedback returns. We introduce Geometric Replanning with Adaptive Chunk Execution (GRACE), a controller that connects these decisions through the geometry and stability of action generation. Curvature from the preceding denoising trajectory provides a geometric memory that weights directional guidance from the previous plan's remaining actions. Within the current guided generation, agreement among late action estimates determines the prefix to execute. Actual execution then positions the reference for the next observation, while the completed generation supplies its geometric memory. Direction, execution extent, and history thus evolve through one continuous control loop. Across four LIBERO suites, GRACE achieves the highest overall success at 97.60% while reducing replans by 36.6% relative to $\pi_{0.5}$. Component comparisons support the contributions of adaptive execution, geometric guidance, and directional weighting. An anonymous demonstration is available at `https://anonymous.4open.science/w/grace-demo-F4F2/`.

1 INTRODUCTION

Putting a mug in a microwave requires grasping it, carrying it inside, releasing it, and closing the door. During the transfer, a new observation reveals the current hand and object positions, while the preceding plan describes the intended motion toward placement. The next prediction must reconcile this updated scene with the movement already under way. Vision-language-action policies generate such movement in action chunks: they predict a sequence, execute a prefix, and replan from a new observation (Zhao et al., 2023; Brohan et al., 2023; Collaboration et al., 2023; Team et al., 2024; Chi et al., 2023; Black et al., 2024; Physical Intelligence et al., 2025). The prefix sets how long the robot acts before incorporating fresh visual feedback. Each replan therefore brings two connected decisions: *how should the new chunk continue the ongoing motion, and how much of it should be executed before observing again?*

A basic chunking controller generates each chunk without directional guidance from the previous plan and executes a fixed prefix regardless of the current prediction's stability. In Figure 1, the original $\pi_{0.5}$ controller grasps and transports objects but fails both illustrated tasks. To carry these operations through, we seek to coordinate motion continuation with the return of feedback. A useful connection lies in the previous chunk's unexecuted actions: they retain an intended continuation, and the executed prefix determines where it begins and how much remains. Execution length thus sets both when feedback returns and which planned actions remain to guide the next prediction.

We develop this connection through the process that forms each action chunk. Diffusion and flow policies refine action sequences through successive denoising revisions (Ho et al., 2020; Lipman et al., 2023; Chi et al., 2023). In cumulative motion coordinates, curvature records the directions in which the proposed path bends across denoising revisions. Retaining this geometry from the preceding generation provides a memory that can weight how the remaining plan guides a new prediction. As that prediction forms, agreement among its late estimates indicates how consistently each future action has emerged. The same generation can therefore use history to shape motion and provide the stability signal for choosing how far to execute.

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: **From local actions to complete tasks.** Matched LIBERO-10 sequences compare $\pi_{0.5}$ with the proposed controller on basket placement and microwave closure. Columns show steps 0, 130, 260, 390, and termination; completed episodes retain their terminal view and show the actual step.

We introduce GRACE, *Geometric Replanning with Adaptive Chunk Execution*, a controller that uses this generative information to connect motion continuation with execution timing. At each observation, the controller aligns the previous plan with actual execution to obtain the remaining reference. This reference specifies the intended continuation; geometric memory weights its directional influence while allowing motion magnitude to adapt to the fresh view. The stability of that same guided generation then determines the execution prefix. Agreement over more future actions supports a longer execution, while a shorter stable prefix brings feedback sooner. Execution positions the reference for the next observation, while the completed generation supplies the geometry for the following cycle.

Our contributions are threefold:

1. We connect motion direction and execution extent through the history that actual execution leaves for the next observation.
2. We introduce geometric memory from the preceding denoising trajectory within GRACE, guiding the generation whose stability determines the execution prefix. The memory retains the directional structure of how the preceding plan formed and complements the scalar stability summary used for execution.
3. We show that GRACE achieves the highest overall success across four LIBERO suites while using 36.6% fewer replans than $\pi_{0.5}$. Component comparisons support the contributions of adaptive execution, geometric memory, and directional weighting.

2 Related Work

Carrying motion across predictions. Action chunking predicts coordinated sequences for imitation and vision-language-action control (Zhao et al., 2023; Lazzati et al., 2026; Chi et al., 2023; Black et al., 2024; Physical Intelligence et al., 2025). Temporal aggregation smooths execution by combining overlapping predictions for the same action. Continuation guidance brings earlier plans into the new generation: SEAM uses aligned history, RTC preserves committed actions during asynchronous inference, and POTR guides from an action prior (Zhan et al., 2026; Black et al., 2025; Fang et al., 2026). FBFM combines action constraints with new state, while ChainVLA carries progress across queries through an execution state (Li et al., 2026; Huang et al., 2026). GRACE weights an aligned continuation using the preceding generation's denoising geometry. GRACE remembers the directional structure of the denoising process that formed the preceding plan, giving the next generation information about how that motion emerged.

Memory within generation. Temporal information can also enter through the starting point or intermediate state of generation. RTI-DP and variable-noise diffusion carry earlier actions or intermediate futures into new solves (Duan et al., 2025; Høeg et al., 2025); ActionCache retrieves intermediate actions for refinement (Oi et al., 2026). Learned temporal policies incorporate context across predictions (Jia et al., 2026; Miller & Jagersand, 2026; Jiang et al., 2025), while BID selects candidate chunks by temporal coherence (Liu et al., 2025). GRACE summarizes the curvature of the preceding denoising trajectory as a directional metric, retaining information about how the plan formed. Metric weighting and projection connect this construction to geometry-aware updates and proximal optimization (Amari, 1998; Rockafellar, 1976; Combettes & Pesquet, 2011) and give memory a direct role in shaping the current continuation.

Choosing when feedback returns. Execution length determines when a robot can respond to a fresh observation, echoing the receding-horizon view of feedback control and constrained predictive planning (Mayne et al., 2000; Betts, 2010). Concurrent work DVAC (Feng et al., 2026) also measures variation among late clean-action estimates and uses a rolling threshold to select a stable prefix; PACE (Nie et al., 2026) identifies low-speed transitions in the predicted motion as replanning boundaries. These signals connect the return of feedback to the movement being formed. We connect execution to motion continuation through geometric memory: memory shapes the generation whose stability selects the prefix, and actual execution aligns the remaining plan for the next observation.

3 GRACE: Geometric Replanning with Adaptive Chunk Execution

GRACE connects successive observations through the motion between them (Figure 2). Geometric memory helps each new prediction continue the ongoing movement, while agreement during generation determines how far the robot proceeds before looking again. Appendix A gives the complete mathematical construction.

3.1 The Continuation Cycle

As a robot carries an object toward a target, each new view enters a movement already under way. At replan n, GRACE predicts a chunk u_n of H actions and selects an execution prefix of length K_n:

$$u_n^{\langle K_n \rangle} := \big(u_n[j]\big)_{j=0}^{K_n-1}, \qquad 1 \le K_n \le H. \tag{1}$$

A longer prefix carries the movement further before feedback returns; a shorter one lets the next observation guide the continuation sooner.

3.2 Forming the Continuation with Memory

An action chunk takes shape through successive denoising revisions. At step ℓ, the policy refines the noisy sequence x_ℓ at time t_ℓ using velocity v_ℓ, conditioned on the current observation and task. With positive step size Δt, the candidate c_ℓ advances denoising and e_ℓ estimates the completed chunk:

$$c_\ell = x_\ell - \Delta t\, v_\ell, \qquad e_\ell = x_\ell - t_\ell v_\ell. \tag{2}$$

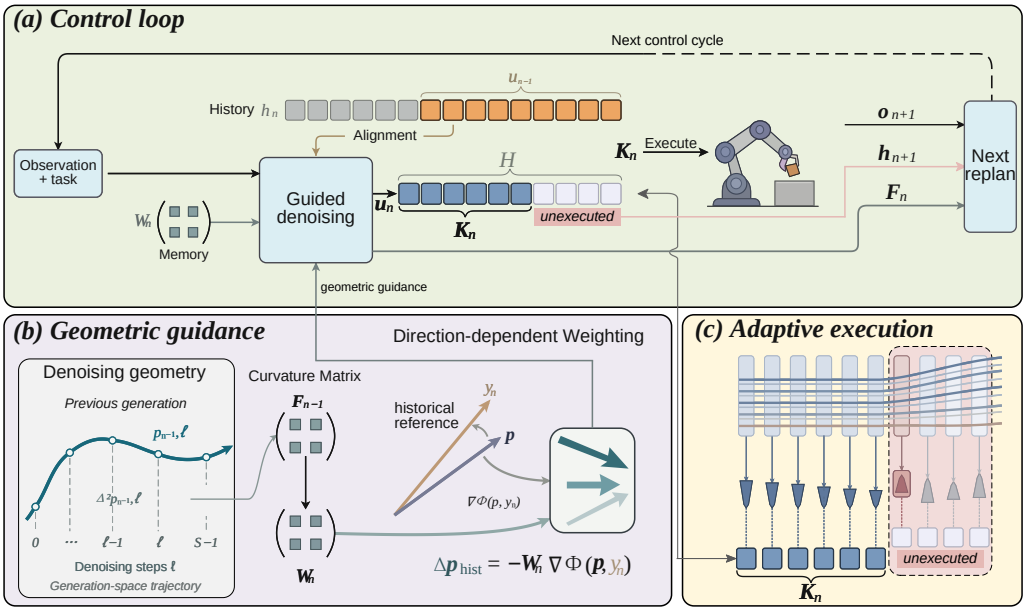

Figure 2: **Overview of GRACE.** (a) The continuation cycle. (b) Geometric memory guides denoising. (c) Agreement among late estimates from the same guided generation selects the execution prefix.

At action position j, let $c_\ell^{\text{mot}}[j] \in \mathbb{R}^{d_m}$ denote the candidate's normalized translation and rotation coordinates, where d_m is their dimension. For replan n, accumulating these coordinates over a geometric window of M actions gives

$$p_{n,\ell}[j] := \sum_{r=0}^{j} c_\ell^{\text{mot}}[r], \qquad p_{n,\ell} := \big(p_{n,\ell}[j]\big)_{j=0}^{M-1}. \tag{3}$$

The replan index n is implicit in the candidate c_ℓ.

The generated path itself moves as denoising revises the chunk. Second differences $\Delta^2 p_{n,\ell}$ across denoising nodes record how these revisions bend. For the nondegenerate case $E_n > 0$, the curvature matrix F_n collects these directions, normalized by the variation energy E_n:

$$F_n = \frac{1}{E_n} \sum_{\ell} \Delta^2 p_{n,\ell}\, \Delta^2 p_{n,\ell}^{\top}, \qquad E_n > 0. \tag{4}$$

This is curvature in the normalized action-generation space: it describes how the predicted cumulative action path changes across denoising nodes, distinct from curvature of the executed end-effector trajectory in the physical workspace. The resulting metric therefore records how the preceding plan formed, while the stability readout later measures agreement among the current plan's estimates. With I the identity, the preceding generation's generation-space curvature becomes the directional metric W_n:

$$W_n = F_{n-1}(I + F_{n-1})^{-1}, \qquad n \geq 1. \tag{5}$$

This transformation preserves the eigendirections of F_{n-1} and maps each eigenvalue $\lambda \geq 0$ to $\lambda/(1+\lambda)$. It gives larger weights to directions emphasized by the preceding denoising revisions while keeping every weight below one. The stored geometry therefore influences the distribution of the correction through a bounded directional operator.

The remaining plan traces a historical path y_n over the available continuation of length L_n. All guidance quantities use this shared support: in the following equation, p denotes the restricted candidate path $p_{n,\ell}^{(L_n)}$, W_n abbreviates its supported metric $W_n^{(L_n)}$, and $y = y_n$ (Appendix A.2). The

angular potential Φ measures directional disagreement, and geometric memory shapes a historical correction direction Δp_{hist} in path space:

$$\Phi(p, y) = \|p\| \, \|y\| - \langle p, y \rangle \,, \qquad \Delta p_{\text{hist}} = -W_n \nabla \Phi(p, y). \tag{6}$$

The potential is zero for positively collinear paths at any relative magnitude: a candidate can point along the historical continuation without reproducing its scale. Thus the observation-conditioned policy can adjust the amount of motion while history encourages a compatible direction. Applied throughout denoising, this guidance shapes both the final chunk u_n and the evolving predictions used to choose its execution length. Appendix A.3 gives the complete correction.

3.3 Execution and the Next Continuation

The same memory-guided generation also reveals how far its continuation has settled. When its late predictions agree over a longer prefix, the robot can carry that movement further before observing again.

At action position j, $e_\ell[j] \in \mathbb{R}^{d_s}$ denotes the normalized motion and gripper components of the estimate e_ℓ, with d_s their combined dimension. Curvature uses the translation and rotation coordinates, while this stability readout also includes the gripper coordinate. With $\bar{e}[j]$ their mean over the final Q denoising steps, the stability score is

$$V_n[j] = \mathbb{E}_\ell \left[\|e_\ell[j] - \bar{e}[j]\|_2^2 \right] . \tag{7}$$

Smaller scores mean closer agreement. We choose K_n to end the prefix before the first score above a threshold η_n drawn from recent generations (Appendix A.4).

Actual execution positions the remaining plan for the next observation, while F_n supplies its geometry. Each executed prefix thus sets both when the robot looks again and which planned motion remains to guide it.

4 Experiments

Our evaluation follows the control loop from outcome to mechanism. We first establish task completion and feedback efficiency, then isolate the roles of adaptive execution and geometric memory, and finally trace how a generated continuation becomes the next executed prefix. Across four LIBERO suites, the experiments connect aggregate performance to the motion and geometry that produce it.

Table 1: **Task completion across four LIBERO suites.** Success (%) is mean $\pm$ sample SD across repeated evaluation runs. Each suite contains 2,500 episodes per method; Overall averages the four suites within each run before computing the mean and SD (10,000 episodes per row). Bold marks the highest mean in each column, breaking ties by the lowest SD.

Method	Spatial	Object	Goal	Long	Overall
$\pi_{0.5}$ (Physical Intelligence et al., 2025)	98.04 ± 0.78	97.84 ± 0.62	97.48 ± 0.46	92.52 ± 0.84	96.47 ± 0.35
SEAM (Zhan et al., 2026)	98.40 ± 0.73	97.48 ± 0.64	97.88 ± 0.77	92.80 ± 1.03	96.64 ± 0.54
RTC (Black et al., 2025)	98.04 ± 0.65	98.68 ± 0.18	98.08 ± 0.46	92.92 ± 0.76	96.93 ± 0.21
POTR (Fang et al., 2026)	98.48 ± 0.61	98.60 ± 0.51	97.76 ± 0.52	93.76 ± 1.01	97.15 ± 0.54
BID (Liu et al., 2025)	$\mathbf{98.84 \pm 0.52}$	98.48 ± 0.59	97.12 ± 0.61	93.48 ± 0.79	96.98 ± 0.43
DVAC (Feng et al., 2026)	98.76 ± 0.65	98.96 ± 0.41	97.64 ± 0.65	94.24 ± 1.24	97.40 ± 0.33
PACE (Nie et al., 2026)	98.72 ± 0.52	98.84 ± 0.61	$\mathbf{98.08 \pm 0.44}$	93.72 ± 1.10	97.34 ± 0.36
GRACE	98.52 ± 0.69	$\mathbf{99.12 \pm 0.23}$	97.88 ± 0.48	$\mathbf{94.88 \pm 0.44}$	$\mathbf{97.60 \pm 0.17}$

4.1 Evaluation Setting

We evaluate on LIBERO-Spatial, LIBERO-Object, LIBERO-Goal, and LIBERO-Long (LIBERO-10), covering 40 manipulation tasks (Liu et al., 2023). All methods use the same pretrained $\pi_{0.5}$-LIBERO policy (Physical Intelligence et al., 2025) and the same official initial-state coverage. The action-step horizons are 220, 280, 300, and 520, respectively. We report mean success and its

sample standard deviation across repeated evaluation runs. Overall success first averages the four suites within each run, then computes the mean and standard deviation across runs. Mean replans includes the initial policy call and averages over all outcomes; it describes how often the controller incorporates a new observation. Appendix B gives the statistical definitions and detailed results.

We compare GRACE with $\pi_{0.5}$, SEAM (Zhan et al., 2026), RTC (Black et al., 2025), POTR (Fang et al., 2026), BID (Liu et al., 2025), DVAC (Feng et al., 2026), and PACE (Nie et al., 2026) (Table 1).

4.2 Task Completion Across Observations

GRACE achieves the highest overall mean success among the eight evaluated controllers (Table 1), reaching 97.60% with a run-to-run spread of 0.17 percentage points. It also returns to observation substantially less often than $\pi_{0.5}$. Figure 3 presents these results through the success–feedback frontier, paired completions, and cumulative progress: GRACE occupies the high-success, low-feedback region, wins more paired episodes than each comparator, and maintains its advantage as long-horizon tasks unfold. The task-level results show that the reduction in replans holds for every task relative to the base policy (Appendix B.3). Mean action counts remain close at the aggregate level, indicating that the lower replan count comes from longer intervals between observations over a comparable amount of executed motion.

The suite and task breakdowns show where this control-loop advantage matters. GRACE performs strongly across Spatial, Object, Goal, and Long, with the clearest relative gain on longer manipulation sequences where several object relations must be carried through one episode. At the task level, it matches or exceeds the base policy on most tasks, including the majority of Long tasks. Carrying a directional continuation between observations therefore supports both broad performance and the extended interactions for which feedback timing is most consequential. In the matched episodes in Figure 4, GRACE completes all four illustrated tasks, whereas the compared controllers leave the task objectives unmet. Appendix B.2 reports the complete task matrix.

4.3 Contributions Within the Control Loop

Table 2 compares three component substitutions across the same four suites, with 10,000 episodes per configuration. *w/o adaptive execution* retains geometry and executes five actions per replan. *w/o geometric memory* retains adaptive execution and uses the base policy field. *Isotropic previous memory* retains guidance with an equal-trace isotropic metric on the available history. The complete controller is the reference configuration.

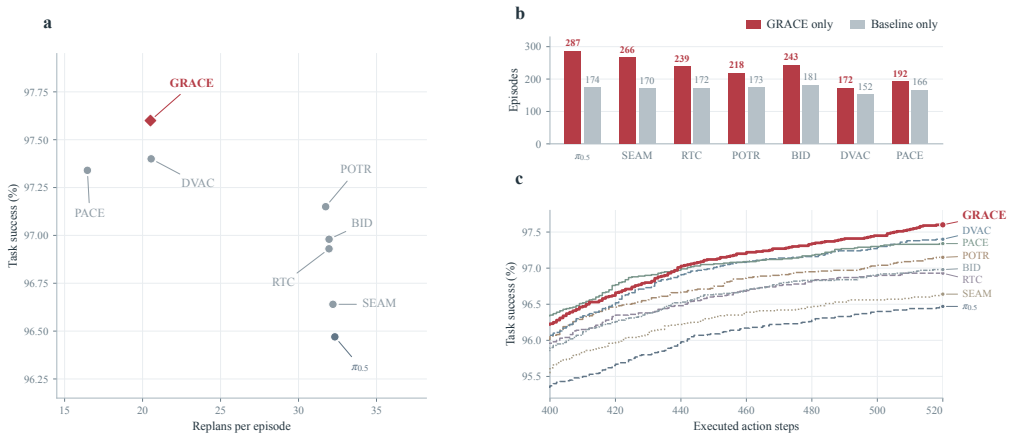

Figure 3: **Task completion and feedback efficiency.** (a) Success and mean replans across controllers. (b) Paired completion outcomes against the baselines. (c) Cumulative success as execution proceeds.

Table 2: **Components of the control loop across four suites.** Success (%) is mean ± sample SD across repeated evaluation runs. Each suite contains 2,500 episodes per configuration; Overall averages the four suites within each run before computing the mean and SD (10,000 episodes per row). Bold marks the highest mean in each column, breaking ties by the lowest SD. The complete controller is the reference configuration.

Configuration	Spatial	Object	Goal	Long	Overall
w/o adaptive execution	98.48 ± 0.44	98.20 ± 0.68	97.52 ± 0.41	92.80 ± 1.14	96.75 ± 0.51
w/o geometric memory	98.40 ± 0.73	98.84 ± 0.22	97.24 ± 0.91	94.04 ± 0.46	97.13 ± 0.27
Isotropic previous memory	98.28 ± 0.84	99.00 ± 0.24	97.12 ± 1.17	93.56 ± 1.07	96.99 ± 0.62
GRACE	$\mathbf{98.52 \pm 0.69}$	$\mathbf{99.12 \pm 0.23}$	$\mathbf{97.88 \pm 0.48}$	$\mathbf{94.88 \pm 0.44}$	$\mathbf{97.60 \pm 0.17}$

Panels (a–b) of Figure 5 relate the selected execution prefix to late-generation stability, showing that GRACE extends coherent continuations and brings feedback forward when the later proposal becomes less settled. Panel (c) compares cross-replan continuation revision across memory variants.

Adaptive execution determines how far the guided continuation becomes motion. With geometry retained, it raises overall success by 0.85 percentage points over fixed execution, with the clearest gain on Long. Choosing the prefix therefore coordinates feedback timing with the continuation available for the next replan.

The selected prefix determines where the next observation enters the task, and the actions that remain after execution become the historical continuation for the next generation. Stable predictions therefore support a longer motion interval, while an earlier observation refreshes the continuation before the later part of the proposal is committed.



Figure 4: **Carrying connected interactions through completion.** Matched LIBERO-10 episodes compare four controllers across transfers, spatial arrangements, and closure. Paths are overlaid on the final scenes with a shared progress scale.

Panel (c) isolates the complementary role of geometric memory with adaptive execution retained. Directional memory shifts the cross-replan revision distribution toward smaller changes, and the table shows corresponding gains in overall success over both the no-memory and isotropic variants. Each variant generates its own actions and carries the resulting history forward, allowing its guidance choices to affect subsequent candidates and execution decisions. The isotropic substitution preserves the total metric weight but removes its directional allocation. Its lower success, together with the revision distributions, supports retaining the geometry of the preceding solve when shaping the next continuation.

Geometric memory determines the content of that continuation. The curvature-derived metric retains the directional structure of the preceding denoising revisions, so the next candidate can inherit the plan's orientation while adapting its magnitude to the new observation. The lower revision distribution in Panel (c) reflects this carry-over across replans, and the success ordering in Table 2 shows how directional structure complements the execution schedule. Adaptive execution decides when the loop returns to feedback; geometric memory shapes what the next interval continues.

The complete controller achieves the highest mean success in each suite of Table 2, extending the overall ordering across spatial relations, object identities, goal conditions, and longer task sequences. These results support the complementary roles of adaptive execution and directional memory across the different task demands represented in the benchmark.

4.4 From Memory to the Next Observation

Figure 6(a–c) shows the historical continuation, current candidate, and GRACE-returned path in a common normalized coordinate system, with the selected prefix directly below each scene. The returned paths move toward the remembered direction while preserving the current observation-conditioned proposal.

The historical path supplies the direction, the current candidate supplies the observation-conditioned motion, and the returned path combines them before the stability readout selects the executed prefix. The same prefix then determines which portion of the plan remains available at the next observation.

Panel (d) measures the within-generation directional effect: the directional metric maintains a larger alignment gain than its equal-trace isotropic replacement throughout denoising, with a 5.18× directional-to-isotropic gain ratio at the final update. Panel (e) reports cumulative closed-loop task success, with GRACE leading the component variants through the late execution horizon.

In the paired intervention in Panel (d), both branches start from the same observation-conditioned candidate and historical reference, use the same continuation support, and apply the same correction and projection procedure. Matching the metric trace preserves the total directional weight while changing its distribution. The alignment comparison therefore isolates how remembered geometry allocates guidance within a denoising update.

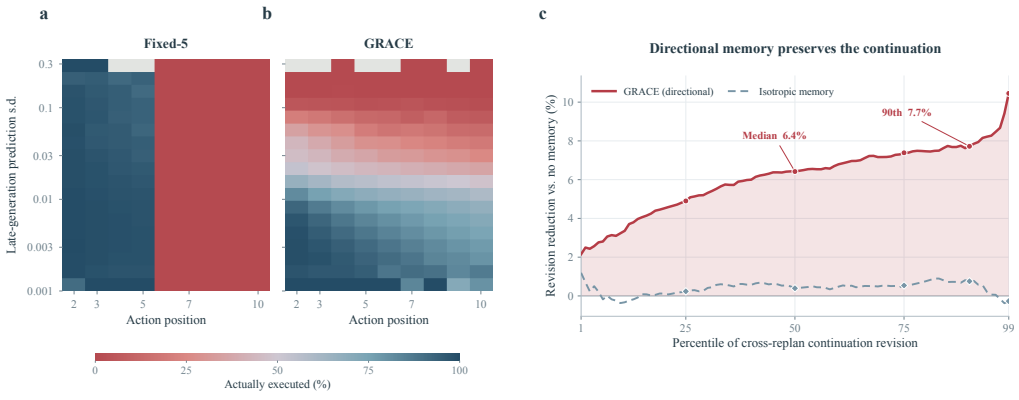

Figure 5: **Adaptive execution and continuation continuity.** (a–b) Execution positions alongside late-generation prediction stability. (c) Cross-replan continuation revision for the memory variants.

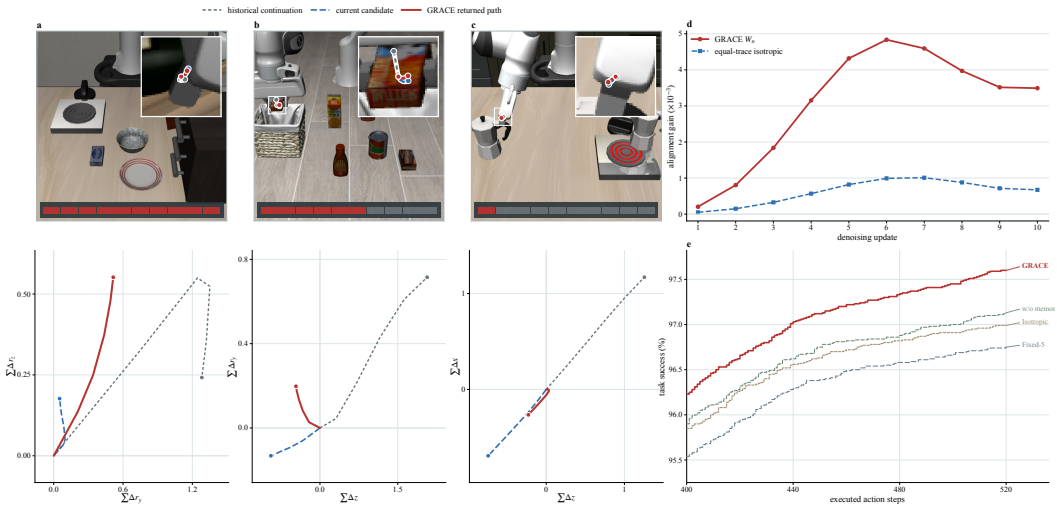

Figure 6: **From geometric memory to task completion.** (a–c) Scene, continuation, and selected prefix for three calls. (d) Directional alignment during denoising. (e) Cumulative success for the component variants and GRACE.

Reference alignment after execution. Across the evaluation, the selected prefix and actual execution jointly determine the next reference, while the completed generation supplies the geometry for the following replan. A longer executed prefix advances the reference farther along the plan and leaves a shorter remaining continuation; an earlier return to observation preserves more of the prior plan for guidance. The controller therefore updates the reference's content together with the timing of feedback. Geometric memory provides the directional weighting for that newly aligned reference, linking execution decisions back to the next guided generation. After full execution, the next call has no unexecuted reference, so its historical correction vanishes. The policy still forms an observation-conditioned chunk and records fresh curvature and a stability profile. When that chunk leaves an unexecuted continuation, the following call can again combine its aligned reference with the newly recorded geometry. The first call of an episode follows the same observation-driven generation with empty history. Appendix B.4 reports the resulting execution profiles and history availability by task.

5 Conclusion

A complete task unfolds through many local predictions. GRACE connects the decisions inside each prediction to the observation that follows: geometric memory carries the direction of the preceding generation, denoising agreement selects the execution prefix, and actual execution positions the next reference. The component results separate these roles while showing how they work together. Adaptive execution sets the timing of feedback, and directional memory preserves the motion content that feedback should continue.

Across four LIBERO suites, GRACE achieves the highest overall success at 97.60%, with 36.6% fewer replans per episode than $\pi_{0.5}$. The mechanism evidence follows the same path from local to global: directional weighting improves alignment during generation, continuation revision becomes more coherent across replans, and the complete controller maintains the strongest task-level progress. Each glance renews the prediction; memory carries direction into the movement that follows.

Reproducibility statement

Appendix A provides the complete mathematical formulation and controller parameters. Appendix B documents the evaluation protocol, metric definitions, component comparisons, task-level results, and mechanism-analysis procedures.

AI use statement

We used generative AI tools to aid or polish writing. The authors take responsibility for the final paper.

References

Shun-ichi Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.

John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. SIAM, 2010.

Kevin Black, Noah Brown, Danny Driess, Alejandro Escontrela, Adam Equall, et al. Pi0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

Kevin Black, Manuel Y. Galliker, and Sergey Levine. Real-time execution of action chunking flow policies. In *Advances in Neural Information Processing Systems*, 2025.

Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.

Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin CM Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems*, 2023. doi: 10.15607/RSS.2023.XIX.026.

Open X-Embodiment Collaboration et al. Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864*, 2023.

Patrick L. Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pp. 185–212. Springer, 2011.

Yufei Duan, Hang Yin, and Danica Kragic. Real-time iteration scheme for diffusion policy. *arXiv preprint arXiv:2508.05396*, 2025.

Kai Fang, Hailong Pei, and Xuemin Chi. Smoother action chunking flow policy via prior-corrected orthogonal trust-region guidance. *arXiv preprint arXiv:2605.24433*, 2026.

Xiangdong Feng, Yuxuan Cheng, Chen Shi, Boyao Han, Yuxuan Yan, Yitong Hong, Zhuotao Tian, and Li Jiang. Denoising tells when to replan: Denoising-variance adaptive chunking for flow-based robot policies. *arXiv preprint arXiv:2606.03847*, 2026.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.

Sigmund H. Høeg, Yilun Du, and Olav Egeland. Fast policy synthesis with variable noise diffusion models. In *IEEE International Conference on Robotics and Automation*, 2025.

Yuzhi Huang, Weijue Bu, Ziyi Xiong, Jie Wu, Fanding Huang, Jingyan Jiang, and Zhi Wang. Chain-VLA: Chaining vision-language-action queries through a unified execution state for long-horizon manipulation. *arXiv preprint arXiv:2608.02326*, 2026.

Jindou Jia, Gen Li, Xiangyu Chen, Tuo An, Yuxuan Hu, Jingliang Li, Xinying Guo, and Jianfei Yang. Action-to-action flow matching. *arXiv preprint arXiv:2602.07322*, 2026.

Sunshine Jiang, Xiaolin Fang, Nicholas Roy, Tomás Lozano-Pérez, Leslie Pack Kaelbling, and Siddharth Ancha. Streaming flow policy: Simplifying diffusion/flow-matching policies by treating action trajectories as flow trajectories. In *Conference on Robot Learning*, 2025.

Filippo Lazzati, Kyle Stachowicz, William Chen, Alberto Maria Metelli, Andrew Wagenmaker, and Sergey Levine. Why does action chunking improve behavioral cloning performance in robotic control? *arXiv preprint arXiv:2608.02547*, 2026.

Peize Li, Ruimeng Zhang, Ru Zhang, Cong Huang, Kai Chen, and Shanghang Zhang. FBFM: A training-free asynchronous feedback mechanism for flow-matching in world-action models execution. *arXiv preprint arXiv:2607.29235*, 2026.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *arXiv preprint arXiv:2306.03310*, 2023.

Yuejiang Liu, Jubayer Ibn Hamid, Annie Xie, Yoonho Lee, Maximilian Du, and Chelsea Finn. Bidirectional decoding: Improving action chunking via guided test-time sampling. In *International Conference on Learning Representations*, 2025.

David Q. Mayne, James B. Rawlings, Christopher V. Rao, and Patrick O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

Dylan Miller and Martin Jagersand. Temporal policy: History-initialized action generation for robotic learning from demonstration. *arXiv preprint arXiv:2607.29482*, 2026.

Junnan Nie, Jiayi Li, Jiachen Zhang, Junyi Lao, Chenghao Liu, Tianle Zhang, Liang Lin, and Songfang Huang. PACE: Phase-aware chunk execution for robot policies with action chunking. *arXiv preprint arXiv:2606.00537*, 2026.

Ryuji Oi, Hikari Otsuka, Kosuke Matsushima, Yuki Ichikawa, Masato Motomura, Tatsuya Kaneko, and Daichi Fujiki. ActionCache: Training-free acceleration for vision-language-action models with action caching and refinement. *arXiv preprint arXiv:2607.06370*, 2026.

Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael 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. *arXiv preprint arXiv:2504.16054*, 2025.

R. Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.

Octo Model Team et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.

Dijia Zhan, Xuemiao Xu, Jinyi Li, and Jie Tang. Seam: Smooth execution of action-chunked motion for vision-language-action policies. *arXiv preprint arXiv:2607.04609*, 2026.

Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Robotics: Science and Systems*, 2023.

A THE MATHEMATICAL CONSTRUCTION OF GRACE

Indices n, ℓ, and j denote replans, denoising nodes, and action positions, respectively. Action positions use zero-based indexing.

A.1 ALIGNED CONTINUATIONS AND MOTION COORDINATES

Let u_{n-1} be the previous chunk and k_{n-1}^{act} its actual executed length. For maximum reference window M, the available history is

$$L_n = \min\{M, H - k_{n-1}^{\mathrm{act}}\}, \qquad h_n[j] = u_{n-1}[k_{n-1}^{\mathrm{act}} + j], \quad 0 \le j < L_n. \tag{8}$$

Each episode begins with an empty reference, zero geometry, and empty stability history, so the first replan has $L_0 = 0$. Within a fixed replan, we abbreviate the reference length as L. We represent h_n over M action positions by setting its remaining entries to zero. All geometric quantities use the motion components of the policy's normalized action coordinates.

Let d_m be the motion dimension and $a^{\mathrm{mot}}[j]$ denote the translation and rotation components of action $a[j]$ in the policy's normalized coordinates. Define the cumulative map $\mathcal{T}_M$ by

$$[\mathcal{T}_M(a)]_j = \sum_{r=0}^{j} a^{\mathrm{mot}}[r], \qquad 0 \le j < M. \tag{9}$$

Its image is a path vector in $\mathbb{R}^{Md_m}$. In Eq. 3, $p_{n,\ell} = \mathcal{T}_M(c_{n,\ell})$, with the M cumulative points concatenated into one vector and n implicit in the main-text candidate c_ℓ. Let P_L retain the first L motion blocks and set the others to zero. The map used for a reference of length L is $\mathcal{T}_L = P_L\mathcal{T}_M$. Its adjoint satisfies

$$[\mathcal{T}_L^*(q)]_j^{\mathrm{mot}} = \sum_{r=j}^{L-1} q_r \quad (j < L), \qquad [\mathcal{T}_L^*(q)]_j = 0 \quad (j \ge L), \tag{10}$$

and inserts zeros in the remaining action coordinates. The historical path is $y_n = \mathcal{T}_{L_n}(h_n)$.

A.2 CURVATURE MEMORY AND DIRECTIONAL WEIGHTS

At node ℓ of replan n, let $c_{n,\ell}$ be the Euler candidate before that node's geometric correction; it inherits corrections from earlier nodes. Its cumulative path over the fixed window is $p_{n,\ell} = \mathcal{T}_M(c_{n,\ell})$. Define

$$\Delta p_{n,\ell} = p_{n,\ell} - p_{n,\ell-1}, \tag{11}$$

$$\Delta^2 p_{n,\ell} = p_{n,\ell} - 2p_{n,\ell-1} + p_{n,\ell-2}. \tag{12}$$

For S denoising nodes, use positive step size $\Delta t = 1/S$ and times $t_\ell = 1 - \ell\Delta t$, with $0 \le \ell < S$. The variation energy E_n and curvature contribution F_n are

$$E_n = (\Delta t)^2 \sum_{\ell=1}^{S-1} \|\Delta p_{n,\ell}\|_2^2, \qquad F_n = \begin{cases} \dfrac{\sum_{\ell=2}^{S-1} \Delta^2 p_{n,\ell} (\Delta^2 p_{n,\ell})^\top}{E_n}, & E_n > 0, \\ 0, & E_n = 0. \end{cases} \tag{13}$$

The fixed window M defines the curvature representation at every replan, including calls with $L_n = 0$. A call with empty history still produces an action chunk, curvature, and a stability profile. After replan n, F_n replaces the geometric state for the next replan. With I the identity, the metric uses the immediately preceding generation's curvature:

$$W_0 = 0, \qquad W_n = F_{n-1}(I + F_{n-1})^{-1} \quad (n \ge 1). \tag{14}$$

For an available history of length L, the current candidate path p and metric are restricted to the continuation as

$$p^{(L)} = P_L p, \qquad W_n^{(L)} = P_L W_n P_L. \tag{15}$$

The metric is constructed from the full single-generation curvature representation before its directional weights are restricted to the current historical support. Restoring the indices gives $p_{n,\ell}^{(L_n)} = P_{L_n} p_{n,\ell}$. In the main-text guidance formula, Eq. 6, p and W_n denote the supported quantities $p_{n,\ell}^{(L_n)}$ and $W_n^{(L_n)}$, respectively.

Spectral properties. Each F_n is positive semidefinite. For $n \geq 1$, an eigenvalue $\lambda \geq 0$ of F_{n-1} gives the eigenvalue $\lambda/(1+\lambda) \in [0, 1)$ of W_n. The transformation preserves eigenvectors; weights increase with the input eigenvalue, with diminishing increments. Moreover, for $z = P_L z$,

$$0 \leq z^\top W_n^{(L)} z \leq \lambda_{\max}(W_n)\|z\|_2^2 < \|z\|_2^2 \qquad (z \neq 0). \tag{16}$$

The normalization in Eq. 13 makes F_n invariant to a common nonzero scaling of the action coordinates. Orthogonal changes of motion coordinates act by congruence on F_n and W_n.

A.3 Angular Guidance and Projection

Let o_n collect the observation and task instruction at replan n, and let v_θ denote the policy velocity field. At denoising state x and time t, evaluate $v = v_\theta(x, t \mid o_n)$. With positive step size Δt, the candidate c, clean estimate e, and next time τ are

$$c = x - \Delta t\, v, \qquad e = x - tv, \qquad \tau = t - \Delta t. \tag{17}$$

Write $L = L_n$, $p = \mathcal{T}_M(c)$, $p^{(L)} = P_L p = \mathcal{T}_L(c)$, and $y = y_n$. The angular potential is

$$\Phi(p, y) = \|p\|\,\|y\| - \langle p, y\rangle, \qquad \Phi_\tau(p, y) = (1-\tau)\Phi(p, y). \tag{18}$$

For positively collinear paths, Φ equals zero at every magnitude. Gradients act on the first argument. For $p^{(L)} \neq 0$, the scheduled gradient g on the aligned path is

$$g = (1-\tau)\left(\frac{\|y\|}{\|p^{(L)}\|}p^{(L)} - y\right). \tag{19}$$

At $p^{(L)} = 0$, the selected subgradient is $g = -(1-\tau)y$. With angular strength α, the action-space correction for $L > 0$ is

$$\begin{aligned}\delta &= -\frac{2\alpha(1-\tau)}{L}\,\mathcal{T}_L^*\Big(W_n^{(L)} g\Big) \\ &= \frac{2\alpha(1-\tau)^2}{L}\,\mathcal{T}_L^*(\Delta p_{\text{hist}}).\end{aligned} \tag{20}$$

The second form follows from $g = (1-\tau)\nabla\Phi(p^{(L)}, y)$ and the main-text direction Δp_{hist}. The two factors of $1-\tau$ arise from the scheduled potential and the action correction. For $L = 0$, $\delta = 0$. Both the gradient and the output of the metric lie in the current support, as expressed by Eq. 15.

The half-space of aggregate historical alignment is

$$\mathcal{K}_L(y) = \{a : \langle \mathcal{T}_L(a), y\rangle \geq 0\}. \tag{21}$$

Setting $r = \mathcal{T}_L^*(y)$ gives its Euclidean projection:

$$\Pi_{\mathcal{K}_L(y)}(a) = \begin{cases} a + \dfrac{\max\{0, -\langle a, r\rangle\}}{\|r\|_2^2} r, & \|r\| > 0, \\ a, & \|r\| = 0. \end{cases} \tag{22}$$

The complete denoising update is

$$e^+ = \Pi_{\mathcal{K}_L(y)}\left(e + \frac{t}{\Delta t}\delta\right), \qquad x^+ = c + \frac{\Delta t}{t}(e^+ - e). \tag{23}$$

For $L = 0$, the historical path is zero, the projection is the identity, and $x^+ = c$. The reference and metric remain fixed during the solve. With the node index restored, x^+ gives the corrected state $x_{\ell+1}$ from which the next clean prediction is evaluated. At each node, the execution readout retains the estimate e from Eq. 17 before that node's correction.

Projection and angular properties. The adjoint identity $\langle \mathcal{T}_L(a), y\rangle = \langle a, r\rangle$ makes Eq. 22 the solution of

$$\min_w \ \tfrac{1}{2}\|w - a\|_2^2 \quad \text{subject to} \quad \langle w, r\rangle \geq 0. \tag{24}$$

For an input outside the half-space, the active-constraint solution is $w = a - \langle a, r\rangle\, r/\|r\|^2$; inputs inside it remain fixed. The correction has support on the selected motion coordinates. Cauchy–Schwarz gives $\Phi(p, y) \geq 0$, with equality for positively collinear nonzero paths. At $p = 0$, $-y$ belongs to the subdifferential of Φ, establishing the gradient branch above.

A.4 Stability Readout and Execution Length

Let d_s count the normalized motion and gripper coordinates, and let $e_{n,\ell}[j,d]$ denote the clean estimate at action position j and coordinate $d \in \{1,\dots,d_s\}$. For the final Q denoising nodes $\mathcal{I}_Q = \{S-Q,\dots,S-1\}$, define

$$\bar{e}_n[j,d] = \frac{1}{Q}\sum_{\ell\in\mathcal{I}_Q} e_{n,\ell}[j,d], \tag{25}$$

$$V_n[j] = \sum_{d=1}^{d_s}\frac{1}{Q}\sum_{\ell\in\mathcal{I}_Q}\left(e_{n,\ell}[j,d]-\bar{e}_n[j,d]\right)^2. \tag{26}$$

The estimates are collected from the current guided trajectory: each node after the first starts from the preceding node's geometrically corrected state.

The main-text estimate is $e_\ell[j] = (e_{n,\ell}[j,d])_{d=1}^{d_s}$, and the expectation in Eq. 7 is uniform over $\mathcal{I}_Q$.

Let $\mathcal{B}_n$ concatenate the score profiles from the preceding B calls, using all available calls during startup. The first call uses $\mathcal{B}_0 = \{V_0[j] : 0 \le j < H\}$. For threshold multiplier β, use the population standard deviation to set

$$\eta_n = \mathrm{mean}(\mathcal{B}_n) + \beta\,\mathrm{std}(\mathcal{B}_n). \tag{27}$$

Writing $J_n = \{j : V_n[j] > \eta_n\}$, the selected execution length is

$$K_n = \begin{cases} H, & J_n = \varnothing, \\ \max\{1, \min J_n\}, & J_n \neq \varnothing. \end{cases} \tag{28}$$

For a crossing at position $j > 0$, the first j actions form the selected prefix. A crossing at $j = 0$ still executes one action.

After choosing K_n, the current profile enters the rolling buffer and F_n becomes the geometric memory for the next replan. The controller then executes the selected prefix. Actual execution determines the aligned reference in Eq. 8 at the following call.

A.5 Parameterization

The evaluated controller uses the following parameterization of the construction:

Quantity	Value
Prediction horizon H	10
Maximum geometric window M	5
Denoising nodes S	10
Motion and stability dimensions (d_m, d_s)	$(6, 7)$
Angular strength α	0.1
Final nodes for stability Q	5
Recent score profiles B	5
Threshold multiplier β	2

B Evaluation Protocol and Detailed Results

B.1 Shared Evaluation and Metric Definitions

All methods and component variants use the same pretrained $\pi_{0.5}$-LIBERO checkpoint and official initial-state coverage. Spatial, Object, Goal, and Long each contain ten tasks, with 250 episodes per task, 2,500 per suite, and 10,000 per configuration. Success follows the benchmark's completion condition within 220, 280, 300, and 520 action steps for the four suites, respectively.

For an aggregate containing N episodes, let y_i denote success, r_i the number of policy calls including the initial call, and a_i the executed action steps. We report

$$\mathrm{SR} = \frac{100}{N}\sum_{i=1}^{N} y_i, \quad \mathrm{SR@60} = \frac{100}{N}\sum_{i=1}^{N} y_i \mathbf{1}[r_i \le 60], \quad \overline{r} = \frac{1}{N}\sum_{i=1}^{N} r_i, \quad \overline{a} = \frac{1}{N}\sum_{i=1}^{N} a_i. \tag{29}$$

Success differences are in percentage points. SR@60 includes all episodes in its denominator; replans and action steps average successes and failures. Table 3 reports these metrics for the main comparison and component substitutions, with the complete controller as their shared reference.

Table 3: **Task completion and feedback across four suites.** Every row contains 10,000 episodes under the common evaluation protocol. SR@60 is the percentage of all episodes successful within 60 policy calls, including the first call. Replans and action steps are means over all episodes. The GRACE row also serves as the complete-controller ablation reference.

Method / configuration	Successes	SR@60 (%)	Replans	Steps
Main comparison				
$\pi_{0.5}$	9,647	93.42	32.33	159.75
SEAM	9,664	93.52	32.21	159.14
RTC	9,693	94.25	31.95	157.85
POTR	9,715	94.35	31.74	156.78
BID	9,698	93.91	31.96	157.90
DVAC	9,740	97.23	20.55	157.47
PACE	9,734	97.33	16.46	156.42
GRACE	9,760	97.37	20.50	157.19
Component substitutions				
w/o adaptive execution	9,675	93.63	32.30	159.62
w/o geometric memory	9,713	96.91	20.62	157.95
Isotropic previous memory	9,699	96.87	20.57	157.84

Component substitutions. The configuration *w/o adaptive execution* sets $K_n = 5$ and retains the preceding generation's geometry and guidance. The configuration *w/o geometric memory* retains the stability readout and uses the base denoising field, with angular correction and historical projection disabled. *Isotropic previous memory* retains both operations and preserves the total metric weight on the available history. For $L > 0$, it uses

$$W_{n,\mathrm{iso}}^{(L)} = \frac{\mathrm{tr}(W_n^{(L)})}{Ld_m} P_L. \tag{30}$$

Here P_L selects the aligned support and $W_n^{(L)}$ is its effective metric (Appendices A.1 and A.2). Each configuration generates its own actions and history; the isotropic metric is computed from that configuration's preceding generation.

B.2 Success Across Tasks

Tables 5 and 6 report success on all 40 tasks, with equal task weights in the overall comparison. Identifiers T0–T9 follow the benchmark order within each suite; Table 4 lists the Long task instructions.

Table 4: **LIBERO-10 task descriptions.** Task identifiers follow the benchmark order.

Task	Instruction
T0	Put both the alphabet soup and the tomato sauce in the basket.
T1	Put both the cream cheese box and the butter in the basket.
T2	Turn on the stove and put the moka pot on it.
T3	Put the black bowl in the bottom drawer of the cabinet and close it.
T4	Put the white mug on the left plate and put the yellow and white mug on the right plate.
T5	Pick up the book and place it in the back compartment of the caddy.
T6	Put the white mug on the plate and put the chocolate pudding to the right of the plate.
T7	Put both the alphabet soup and the cream cheese box in the basket.
T8	Put both moka pots on the stove.
T9	Put the yellow and white mug in the microwave and close it.

Table 5: **Success by task (%): Spatial and Object.** Each cell aggregates the common official initial-state coverage. Task identifiers follow the order within each suite; Long instructions appear in Table 4.

Task	$\pi_{0.5}$	SEAM	RTC	POTR	BID	DVAC	PACE	GRACE
Spatial								
T0	99.60	100.00	100.00	100.00	100.00	100.00	99.20	100.00
T1	100.00	100.00	99.20	99.60	100.00	100.00	100.00	98.40
T2	99.60	100.00	99.60	99.60	100.00	99.20	100.00	99.60
T3	95.20	96.40	95.20	96.80	98.80	97.60	98.80	98.80
T4	94.40	94.80	95.60	95.20	97.20	96.40	96.00	95.60
T5	97.20	97.20	96.80	97.20	97.20	97.20	97.20	97.60
T6	99.60	100.00	99.60	99.60	99.20	99.60	99.60	99.60
T7	98.00	99.60	98.80	98.80	98.80	98.80	99.20	98.40
T8	99.20	98.80	98.80	98.80	98.40	99.60	99.60	99.60
T9	97.60	97.20	96.80	99.20	98.80	99.20	97.60	97.60
Object								
T0	95.60	92.40	99.60	98.00	98.00	99.60	98.80	98.80
T1	100.00	99.20	100.00	99.60	99.60	100.00	99.60	99.60
T2	96.40	97.20	98.00	97.20	97.20	100.00	99.20	99.20
T3	96.80	97.60	97.20	98.00	96.80	97.60	98.00	96.80
T4	96.80	98.40	98.80	98.00	99.20	96.00	98.00	98.00
T5	98.00	98.00	96.80	97.60	98.40	98.40	98.00	99.60
T6	98.00	97.60	99.20	99.60	99.20	98.40	99.20	100.00
T7	99.20	99.20	99.60	99.60	99.20	100.00	99.60	99.60
T8	99.20	96.80	100.00	100.00	99.60	100.00	99.60	99.60
T9	98.40	98.40	97.60	98.40	97.60	99.60	98.40	100.00

Table 6: **Success by task (%): Goal and Long.** Each cell aggregates the common official initial-state coverage. Task identifiers follow the order within each suite; Long instructions appear in Table 4.

Task	$\pi_{0.5}$	SEAM	RTC	POTR	BID	DVAC	PACE	GRACE
Goal								
T0	96.40	98.40	97.60	98.80	96.80	98.00	98.80	98.80
T1	99.60	99.20	99.20	98.00	99.20	98.80	99.60	98.80
T2	97.60	97.20	97.20	98.40	98.00	97.20	98.00	98.40
T3	92.80	92.00	94.40	92.40	90.40	93.60	94.00	94.00
T4	98.00	98.80	99.60	98.40	98.00	98.00	98.00	96.40
T5	96.80	97.60	97.60	97.60	97.60	98.40	98.00	99.20
T6	98.40	98.00	98.00	97.60	95.20	96.00	97.60	98.80
T7	100.00	100.00	100.00	100.00	100.00	100.00	99.60	99.60
T8	100.00	99.60	99.60	99.60	99.20	100.00	99.20	98.00
T9	95.20	98.00	97.60	96.80	96.80	96.40	98.00	96.80
Long								
T0	95.20	95.20	94.40	94.00	92.80	96.80	96.80	97.60
T1	100.00	99.60	99.20	100.00	99.20	99.60	99.60	99.60
T2	93.60	96.40	98.00	97.20	96.00	95.60	94.00	96.40
T3	98.40	98.00	96.80	97.20	97.20	96.40	96.80	97.20
T4	96.80	97.60	98.00	97.60	97.20	97.60	94.80	96.80
T5	97.60	96.80	97.60	97.20	99.20	100.00	98.00	98.80
T6	92.00	94.00	92.80	93.60	91.20	94.00	92.40	96.80
T7	100.00	97.60	100.00	99.60	99.20	98.80	98.80	100.00
T8	56.00	60.40	58.40	67.20	68.40	70.80	72.40	71.60
T9	95.60	92.40	94.00	94.00	94.40	92.80	93.60	94.00

B.3 TASK-LEVEL OUTCOMES AND EFFICIENCY

Tables 7 and 8 compare task-level success, SR@60, mean replans, and mean action steps for $\pi_{0.5}$ and GRACE.

Table 7: **Detailed task metrics: Spatial and Object.** Each task aggregates 250 episodes under the shared evaluation coverage. Success and SR@60 are episode percentages; Δ is GRACE minus $\pi_{0.5}$ success in percentage points. Replans and action steps are per-episode means over successes and failures.

	Success (%)			SR@60 (%)		Replans		Action steps	
Task	$\pi_{0.5}$	GRACE	Δ	$\pi_{0.5}$	GRACE	$\pi_{0.5}$	GRACE	$\pi_{0.5}$	GRACE
Spatial									
T0	99.60	100.00	+0.40	99.60	100.00	16.75	10.90	81.70	79.11
T1	100.00	98.40	-1.60	100.00	98.40	22.15	14.63	108.65	109.82
T2	99.60	99.60	+0.00	99.60	99.60	20.32	12.84	99.67	99.00
T3	95.20	98.80	+3.60	95.20	98.80	18.42	12.09	89.98	89.12
T4	94.40	95.60	+1.20	94.40	95.60	26.69	17.55	131.56	130.74
T5	97.20	97.60	+0.40	97.20	97.60	20.00	13.29	97.94	96.85
T6	99.60	99.60	+0.00	99.60	99.60	22.03	13.98	108.30	107.23
T7	98.00	98.40	+0.40	98.00	98.40	25.24	15.55	124.26	121.40
T8	99.20	99.60	+0.40	99.20	99.60	20.34	13.19	99.72	98.71
T9	97.60	97.60	+0.00	97.60	97.60	24.89	17.06	122.54	123.48
Object									
T0	95.60	98.80	+3.20	95.60	98.80	31.64	19.11	156.32	146.01
T1	100.00	99.60	-0.40	100.00	99.60	27.02	17.18	133.30	131.07
T2	96.40	99.20	+2.80	96.40	99.20	25.96	16.44	127.83	125.41
T3	96.80	96.80	+0.00	96.80	96.80	29.11	18.96	143.53	144.94
T4	96.80	98.00	+1.20	96.80	98.00	30.38	19.54	149.73	147.13
T5	98.00	99.60	+1.60	98.00	99.60	28.83	18.27	142.51	138.09
T6	98.00	100.00	+2.00	98.00	100.00	32.10	20.38	158.81	154.96
T7	99.20	99.60	+0.40	99.20	99.60	26.97	17.44	132.82	134.01
T8	99.20	99.60	+0.40	99.20	99.60	31.66	19.88	156.17	149.90
T9	98.40	100.00	+1.60	98.40	100.00	25.96	16.51	127.66	128.13

Table 8: **Detailed task metrics: Goal and Long.** Each task aggregates 250 episodes under the shared evaluation coverage. Success and SR@60 are episode percentages; Δ is GRACE minus $\pi_{0.5}$ success in percentage points. Replans and action steps are per-episode means over successes and failures.

	Success (%)			SR@60 (%)		Replans		Action steps	
Task	$\pi_{0.5}$	GRACE	Δ	$\pi_{0.5}$	GRACE	$\pi_{0.5}$	GRACE	$\pi_{0.5}$	GRACE
Goal									
T0	96.40	98.80	+2.40	96.40	98.80	27.10	16.63	133.57	124.21
T1	99.60	98.80	-0.80	99.60	98.80	18.07	11.94	88.26	91.17
T2	97.60	98.40	+0.80	97.60	98.40	19.68	13.00	96.68	95.87
T3	92.80	94.00	+1.20	92.80	94.00	39.75	26.98	197.12	199.39
T4	98.00	96.40	-1.60	98.00	96.40	18.64	12.21	91.20	93.14
T5	96.80	99.20	+2.40	96.80	99.20	28.10	18.21	138.55	134.71
T6	98.40	98.80	+0.40	98.40	98.80	19.40	13.13	95.16	93.74
T7	100.00	99.60	-0.40	100.00	99.60	15.61	10.65	76.04	76.45
T8	100.00	98.00	-2.00	100.00	98.00	15.73	11.17	76.61	81.42
T9	95.20	96.80	+1.60	95.20	96.80	30.11	19.33	148.62	146.89
Long									
T0	95.20	97.60	+2.40	80.00	95.60	58.55	36.86	290.98	292.38
T1	100.00	99.60	-0.40	96.80	99.60	50.84	32.71	252.66	251.93
T2	93.60	96.40	+2.80	88.00	96.40	54.36	32.60	270.01	257.20
T3	98.40	97.20	-1.20	97.60	97.20	46.01	29.40	228.27	233.10
T4	96.80	96.80	+0.00	94.40	96.80	49.91	31.27	247.68	243.78
T5	97.60	98.80	+1.20	95.20	98.80	39.28	24.20	194.51	188.38
T6	92.00	96.80	+4.80	86.00	96.40	51.06	30.71	253.36	241.70
T7	100.00	100.00	+0.00	88.80	100.00	54.06	33.98	268.44	262.69
T8	56.00	71.60	+15.60	0.00	65.60	92.87	55.30	463.21	447.56
T9	95.60	94.00	-1.60	76.40	93.20	57.54	35.14	285.88	276.90

B.4 Execution and Available History

Tables 9 and 10 report execution profiles by task. The controller selects K_n actions; k_n^{act} actions are executed before the next replan or episode termination. Prefix statistics weight calls equally.

History is available when $L_n > 0$. The first call has $L_0 = 0$; subsequent calls satisfy $L_{n+1} = \min\{5, 10 - k_n^{\text{act}}\}$ under the evaluated parameterization. Nonempty history occurs in 32.18% of calls.

Table 9: **Execution profiles by task: Spatial and Object.** Each task contains 250 GRACE episodes. Replans is the total number of calls. Selected and Actual give mean prefix lengths; Full is the percentage selecting ten actions; History is the percentage with nonempty aligned history. All prefix statistics weight calls equally.

Task	Replans	Selected	Actual	Full (%)	History (%)
Spatial					
T0	2,726	7.55	7.25	60.31	36.68
T1	3,658	7.69	7.51	64.16	32.26
T2	3,211	7.95	7.71	66.61	30.05
T3	3,022	7.64	7.37	64.20	32.86
T4	4,387	7.61	7.45	63.00	34.40
T5	3,323	7.50	7.29	60.52	35.09
T6	3,494	7.84	7.67	66.54	30.05
T7	3,887	8.01	7.81	68.48	28.07
T8	3,297	7.74	7.48	63.69	33.24
T9	4,265	7.43	7.24	61.64	35.52
Object					
T0	4,777	7.81	7.64	66.28	33.22
T1	4,294	7.83	7.63	67.75	32.16
T2	4,110	7.83	7.63	67.45	30.61
T3	4,740	7.83	7.64	66.79	32.51
T4	4,885	7.69	7.53	66.61	32.65
T5	4,567	7.71	7.56	65.49	33.72
T6	5,095	7.80	7.60	66.75	33.15
T7	4,360	7.86	7.68	67.04	30.69
T8	4,971	7.75	7.54	66.39	33.53
T9	4,127	7.94	7.76	68.82	28.45

Table 10: **Execution profiles by task: Goal and Long.** Each task contains 250 GRACE episodes. Replans is the total number of calls. Selected and Actual give mean prefix lengths; Full is the percentage selecting ten actions; History is the percentage with nonempty aligned history. All prefix statistics weight calls equally.

Task	Replans	Selected	Actual	Full (%)	History (%)
Goal					
T0	4,157	7.73	7.47	66.61	32.48
T1	2,985	7.86	7.64	64.92	31.86
T2	3,249	7.57	7.38	63.22	32.53
T3	6,745	7.58	7.39	63.75	35.70
T4	3,053	7.89	7.63	64.85	31.05
T5	4,552	7.65	7.40	63.31	35.43
T6	3,282	7.38	7.14	61.55	38.21
T7	2,662	7.58	7.18	62.89	36.81
T8	2,793	7.59	7.29	62.26	34.80
T9	4,832	7.79	7.60	66.45	32.70
Long					
T0	9,216	8.02	7.93	68.68	31.28
T1	8,177	7.81	7.70	66.97	33.01
T2	8,150	8.01	7.89	68.93	29.85
T3	7,350	8.06	7.93	69.63	29.46
T4	7,818	7.90	7.80	67.59	31.08
T5	6,050	7.94	7.78	67.29	32.02
T6	7,678	8.03	7.87	69.41	30.42
T7	8,495	7.82	7.73	66.25	33.67
T8	13,824	8.17	8.09	70.49	28.91
T9	8,784	7.99	7.88	67.53	31.89

B.5 Matched Scenes and Motion Paths

Figures 1 and 4 compare controllers from matched initial scenes and camera views. In Figure 1, GRACE completes basket placement and microwave closure at steps 497 and 277, respectively; later display times retain the terminal frame. In Figure 4, its completion steps are 497, 451, 221, and 261 in row order.

Trajectory overlays project the end-effector position before execution and after each action onto the shared camera view. Color denotes the fraction of executed actions within each episode. Circles and stars mark the initial and terminal positions; intermediate markers and arrows indicate progress. Each row uses a common crop, each background is the controller's terminal scene, and completion labels follow the benchmark criterion.

B.6 Mechanism Analysis Protocol

The quantitative panels of Figure 6 use five evaluation runs across the four LIBERO suites. Panels (a–c) illustrate selected successful calls in normalized motion coordinates, with strips marking execution prefixes. Panel (e) reports closed-loop component results (Appendix B.1).

Paired geometric updates. For each aligned call and denoising update, both branches share the candidate c_ℓ, clean prediction e_ℓ, historical reference, and continuation support. The directional branch uses W_n from the preceding generation's curvature F_{n-1}; the other uses the equal-trace isotropic metric in Eq. 30. Both apply angular correction and projection; paired results are pooled across the four suites.

Let $\widetilde{p}^{\mathrm{dir}}$ and $\widetilde{p}^{\mathrm{iso}}$ denote the resulting cumulative paths on the aligned support, and y the historical path. Alignment gains are measured relative to the same candidate path. Their paired difference at replan n and update ℓ is

$$\cos(p, y) = \frac{\langle p, y \rangle}{\|p\| \, \|y\|}, \qquad D_{n,\ell} = \cos(\widetilde{p}^{\mathrm{dir}}, y) - \cos(\widetilde{p}^{\mathrm{iso}}, y). \tag{31}$$