Supplement: Failure-Conditioned Model-Order Expansion for Embodied Control

Seonhye Gu*
AI-Based Surgical Robot Innovation Lab

S1 Protocol Summary

The confirmatory experiment tests whether a structured target-motion residual that persists after zero-order parameter repair warrants activation of a prepared velocity state. All design choices in this supplement were fixed before the fresh candidate seeds were executed.

Candidate seeds were fixed as 300-339. Static eligibility was evaluated before dynamic treatment, and the first 10 eligible seeds in numeric order were selected: 300, 301, 302, 303, 304, 305, 307, 308, 310, and 311. Thirty-four of 40 candidates passed. Each selected seed was evaluated in a fresh Isaac process for every arm-condition cell, producing a complete 10-seed by 4-arm by 10-condition dynamic grid plus 20 static-retention cells. Dynamic outcomes did not influence seed inclusion.

S2 Exact Model Definitions

Let p_t be the observed target command and $H = 10$ the locked prediction horizon. The zero-order model is

$$\begin{aligned} \hat{p}_t &= \alpha p_t + (1-\alpha)\hat{p}_{t-1}, \\ \hat{p}_{t+H} &= \hat{p}_t. \end{aligned}$$

Episode 1 searches $\alpha \in \{0.25, 0.50, 0.75, 1.00\}$. The selected repair is $\alpha = 1.00$. This model can reduce current-position lag but cannot produce a nonzero open-loop displacement.

The order-one alternative adds

$$\begin{aligned} \hat{v}_t &= \beta\,(p_t - p_{t-1}) + (1-\beta)\hat{v}_{t-1}, \\ \hat{p}_{t+H} &= \hat{p}_t + H\hat{v}_t. \end{aligned}$$

after gate activation. The selected coefficients are $\alpha = \beta = 1.00$. Before activation, arm C returns the same position-only forecast as arm B. The oracle uses the true injected condition velocity and is never included in the primary contrast.

S3 Gate Definition And Controls

The gate uses an eight-step window and requires at least four deltas. It fires if all of the following are true:

*This work was independently conducted by the author as personal research while affiliated with the AI-Based Surgical Robot Innovation Lab. The affiliation is provided for identification purposes only and does not imply official laboratory output, institutional endorsement, or sponsorship.

Quantity	Locked threshold
Mean target speed	≥ 0.5 mm/step
Active-delta fraction	≥ 0.75
Directional consistency	≥ 0.90
Constant-velocity fit improvement	≥ 0.50

Directional consistency is the norm of the summed displacement divided by the sum of displacement norms. Constant-velocity fit improvement compares the RMSE of target deltas under zero velocity with the RMSE of delta differences under constant velocity.

Controls use the same origin and number of steps as the persistent-drift trace:

Control	Construction	Intended rejection
M0 static	Repeated fixed target	No motion evidence
M1 persistent drift	Fixed nonzero delta each step	Positive structural evidence
N1 observation noise	IID Gaussian noise, $\sigma = 0.2$ mm	High-frequency residual without direction
N2 single impulse	One eight-delta jump followed by rest	Transient event without persistence

S4 Confirmatory Drift Conditions

Table 1: Fresh Episode 2 condition set. Diagonal components were selected so that the vector norm matches the stated speed up to rounding.

ID	dx (mm/step)	dy (mm/step)	Delay	Duration
C01	+1.700	0.000	1	23
C02	+1.900	0.000	4	21
C03	-1.800	0.000	1	23
C04	-2.000	0.000	4	20
C05	0.000	+1.700	2	23
C06	0.000	-1.900	4	21
C07	+1.202	+1.202	1	23
C08	-1.273	+1.273	4	22
C09	+1.344	-1.344	2	21
C10	-1.414	-1.414	5	20

The set contains both signs of each axis and all four planar diagonal quadrants. It varies speed, delay, and duration and contains no exact condition from the excluded process-isolated pilot.

S5 Execution And Branch Replay

Each seed-arm-condition cell starts a new Isaac process. The procedure is:

1. Set the simulator seed and create the environment.
2. Run the static prefix until the target is within 20 mm for five stable steps, or stop at 200 prefix steps.
3. Record the reset-state and branch-state fingerprints.
4. Replay the shared Episode 1 evidence and fit the arm-specific model using only the locked candidates.
5. Restore the matched branch command and run one Episode 2 condition.

6. At every step, provide the model forecast to the shared controller, restore the true command for evaluation, and record target, end-effector, action, model state, gate state, and terminal information.
7. Close the Isaac process before starting the next cell.

Fresh-process execution was introduced after an excluded pilot demonstrated that sequential conditions inside one process could perturb later reset states. The confirmatory process-isolation rule therefore treats each cell as a separate simulator lifecycle rather than relying on nominal reset determinism.

S6 Validity Checklist

Table 2: Preregistered execution-validity fields and observed values.

Field	Requirement	Observed
Candidate selection	Fixed first-eligible order	Pass
Selected seed count	10	10
Dynamic cells	400 complete	400/400
Static-retention cells	20 complete	20/20
Maximum branch-start EE gap	≤ 1 mm	0 mm
Maximum branch-start command gap	≤ 0.001 mm	0 mm
Maximum branch-start target distance	≤ 21 mm	15.396 mm
Prefix readiness	All selected cells	Pass
$H = 10$ windows	Present in every dynamic cell	Pass
Drift exposure	Complete in every dynamic cell	Pass
Unexpected environment resets	0	0
Forbidden-region violations	0	0

No record was excluded after selection. The aggregate artifact sets `valid_confirmatory` to true.

S7 Full Confirmatory Outcomes

Table 3: Arm-level outcomes over the complete dynamic grid.

Arm	n	Success	Prediction error	Final distance	Forbidden	Resets
A frozen zero order	100	39%	20.081 mm	20.388 mm	0	0
B repaired zero order	100	76%	18.401 mm	18.905 mm	0	0
C gated constant velocity	100	100%	7.595 mm	5.601 mm	0	0
D oracle velocity	100	100%	0.000 mm	4.087 mm	0	0

Success is thresholded at 20 mm. If both B and C succeed, the pair is tied on the binary endpoint even when C has a smaller continuous final distance. This is why 24/100 pairs show a strict binary gain while 100/100 favor C on final distance.

S8 Condition-Level Effects

Every condition-level mean clears the favorable -5 mm threshold on both continuous endpoints. These rows are descriptive; confirmatory inference uses the crossed seed-condition bootstrap.

Table 4: Preregistered C-minus-B contrasts.

Endpoint	Contrast	Crossed-bootstrap 95% CI	Favorable pairs	Decision
$H = 10$ prediction error	-10.806 mm	[-11.360, -10.331] mm	100/100	H1 pass
Fixed-horizon final distance	-13.304 mm	[-13.599, -12.982] mm	100/100	H2 pass
20 mm resolution rate	+0.24	[+0.08, +0.45]	24/100 strict gains	Secondary

Table 5: C-minus-B effects and success counts by condition.

Condition	Prediction difference	Final-distance difference	B success	C success
C01	-11.769 mm	-12.865 mm	9/10	10/10
C02	-10.364 mm	-13.166 mm	9/10	10/10
C03	-12.462 mm	-13.807 mm	9/10	10/10
C04	-10.000 mm	-13.586 mm	5/10	10/10
C05	-10.461 mm	-12.801 mm	10/10	10/10
C06	-10.364 mm	-13.593 mm	5/10	10/10
C07	-11.768 mm	-12.716 mm	10/10	10/10
C08	-10.502 mm	-13.192 mm	9/10	10/10
C09	-10.367 mm	-13.427 mm	5/10	10/10
C10	-9.998 mm	-13.888 mm	5/10	10/10

S9 Retention And Gate Specificity

Table 6: Static retention and gate controls.

Analysis	Group	n	Rate	Wilson 95% CI	Mean final distance
Static retention	B repaired zero order	10	100%	-	1.501 mm
Static retention	C gated constant velocity	10	100%	-	1.501 mm
Gate	M1 persistent drift	100	100%	[96.30%, 100%]	-
Gate	M0 static	100	0%	[0%, 3.70%]	-
Gate	N1 observation noise	100	0%	[0%, 3.70%]	-
Gate	N2 single impulse	100	0%	[0%, 3.70%]	-

The B and C static trajectories are identical over their overlap. The gate does not activate on the static target, so C reduces exactly to repaired zero order.

S10 Statistical Procedure

For each continuous endpoint, form a 10 by 10 matrix of paired C-minus-B differences indexed by selected seed and fresh condition. For bootstrap replicate b, independently sample 10 seed indices and 10 condition indices with replacement and compute

$$\bar{\Delta}^{*(b)} = \frac{1}{100} \sum_{a=1}^{10} \sum_{c=1}^{10} \Delta_{I_a^{(b)}, J_c^{(b)}}.$$

The reported interval is the 2.5th and 97.5th percentile of 10,000 stored means. The RNG seed is

20260730. This preserves both crossed sampling factors rather than treating 100 cells as independent. The success-rate contrast is resampled identically but remains secondary. Wilson score intervals summarize individual gate rates.

S11 Reproducibility Environment

Item	Setting
Simulator	Isaac Sim 4.1
Framework	Isaac Lab 1.0.0 and ORBIT-Surgical
Task	`Isaac-Reach-Dual-STAR-IK-Rel-Play-v0`
Fabric	Disabled
Preregistration tag	`paper002-model-order-confirmatory-v1.0`
Bootstrap draws / seed	10,000 / 20260730

S12 Simulator Warnings And Scope

The ORBIT-Surgical assets emitted PhysX warnings for rigid bodies with negative mass or placeholder inertia tensors. Isaac Sim substituted approximate inertia. The warning was stable across arms and conditions and did not cause missing cells, unexpected resets, or branch mismatch. It nevertheless limits physical interpretation. The experiment supports a target-model and controller comparison inside the observed simulator stack; it does not validate force, tissue, mass-property, or hardware behavior.

No human participants, animals, patient data, or clinical procedures were involved. The surgical context names the robot benchmark. It is not evidence of clinical safety or efficacy.