TEMPORAL CAUSAL OBSERVABILITY: SEPARATING PERSISTENT CAUSAL STATE FROM PERCEPTUAL ACCESS

Anonymous authors
Paper under double-blind review

ABSTRACT

Temporal causal representation learning (CRL) can recover latent causal factors and their dynamics, yet still fail when perceptual access to those factors changes over time: a change in observation need not imply a change in the underlying causal state. We formalize this distinction as Temporal Causal Observability, separating causal-state evolution from time-varying factor-wise perceptual access and studying controlled interventions that alter access while preserving the causal trajectory. For a frozen temporal CRL model, we show that synchronized reference evidence can evaluate the expected benefit of a fixed measurement correction without observing the true state, under a conditional orthogonality assumption. We then establish an identifiability boundary: evaluating the benefit of a fixed correction does not in general identify when that correction should be applied, because the same observation change may arise from either state dynamics or changing perceptual access. This motivates Regime-Marginalized Elimination (RME), which marginalizes latent access regimes instead of making hard observability decisions. Across multiple temporal CRL systems, reduced perceptual access causes substantial degradation despite strong full-observation recovery. On frozen CITRIS representations, recursive RME initialized from a single clean encoding reduces hidden-state error by 22–74% relative to robust median fusion and by 7–34% relative to recursive hard regime decisions under the same evidence and uncertainty-propagation procedure, across all eight representations. A matched structural ablation further shows that recursive maintenance benefits from aligning access regimes with the causal-factor partition learned by the CRL model. In one-step updates, RME also outperforms hard regime decisions on BISCUIT. These results identify temporal causal observability as a distinct post-identification problem: learning a causal state and its dynamics does not determine whether current perceptual evidence should update that state.

1 INTRODUCTION

Temporal CRL recovers causal variables from high-dimensional observations and learns how they evolve (Lippe et al., 2022; 2023a; Yao et al., 2022). Consider a cyclist who passes behind a car and reappears (Figure 1a). A temporal CRL model may represent the cyclist's position and its dynamics well. The cyclist keeps moving throughout; only the camera's perceptual access to the cyclist changes. Identification leaves open a question that arises at every frame: *given a learned causal state and its dynamics, when should a new encoding change the state we currently believe in?*

The learned dynamics predict the state, D_t, and the encoder reads the current frame, M_t. When the two disagree, the disagreement has two possible causes: a genuine innovation of the causal state, or a change of perceptual access. A faithful encoder reports the pixels in both cases. Factor identification therefore does not settle causal-state maintenance.

We call this post-identification problem **Temporal Causal Observability (TCO)**. TCO separates the persistent causal state C_t from factor-wise perceptual access O_t. We study it by intervening on access rather than on the world: for one factor at a time, we replace the image content that carries the factor with a plausible alternative, while the causal trajectory and all transition inputs stay fixed.

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.

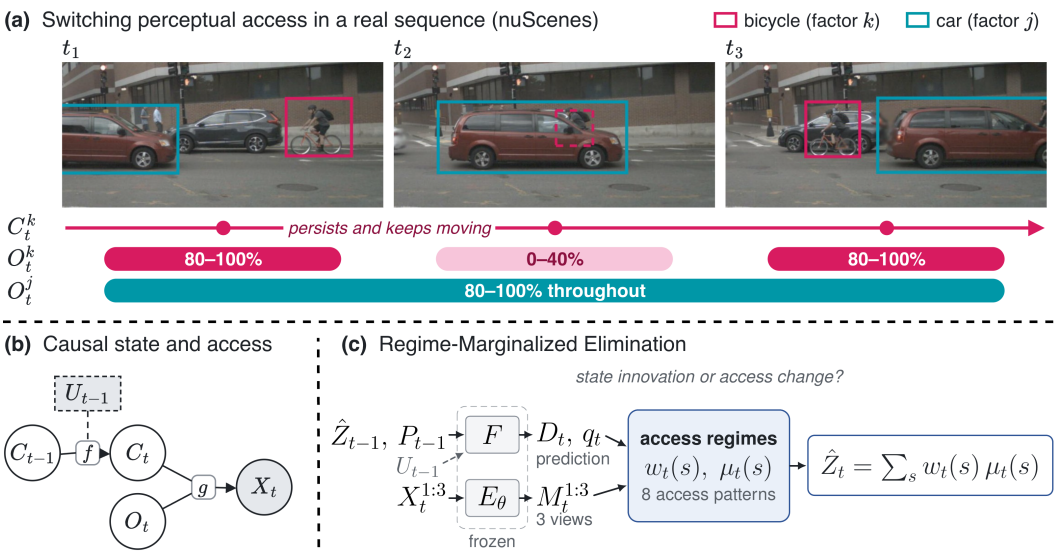

Figure 1: **State persists while access changes.** (a) In a nuScenes sequence, a bicycle passes behind a car and reappears (Caesar et al., 2020); percentages are the dataset's visibility bins. The bicycle's causal state C_t^k keeps evolving while its perceptual access O_t^k drops. (b) State and access enter the observation through separate mechanisms. (c) RME propagates the previous estimate and its variance through the frozen transition F to obtain the prediction (D_t, q_t), weights the 2^3 latent access patterns s of the current block by $w_t(s)$ using three synchronized encodings $M_t^{1:3}$, and averages the regime-conditioned updates $\mu_t(s)$.

Each factor is seen through three synchronized views whose access fails independently, so a hidden factor usually remains visible in some view: the task is to maintain the state from a mix of valid and replaced evidence.

The paper follows one chain of questions, and each answer raises the next.

1. *Is the representation at fault?* No. Encoders pass replaced evidence through almost exactly as a one-line calculation predicts (Section 3.1, Table 1), and under distribution-matched replacement no per-frame model can detect the change of access (Proposition 1).
2. *Can synchronized evidence decide the update?* A synchronized reference identifies the expected benefit of a fixed correction when its error is orthogonal to the correction direction (Proposition 2). Replaced references violate this condition, so access validity must itself be inferred.
3. *How can the causal state be maintained?* RME marginalizes causal-state updates over latent access regimes (Section 4). With a representation-calibrated scale and propagated uncertainty, it maintains the state recursively and outperforms committing to a single regime (Sections 5.2–5.3).

Across eight frozen CITRIS representations, recursive RME consistently outperforms memoryless fusion and matched hard regime decisions; the same ordering holds in one-step BISCUIT updates. A matched repartitioning control further shows that maintenance depends on alignment with the learned causal-factor partition. These results make TCO a testable post-identification problem rather than a failure of factor recovery itself.

Our contributions are: (i) TCO, a post-learning intervention framework that exposes the distinction between persistent causal state and perceptual access, with a closed-form manipulation check; (ii) an information boundary for post-CRL state updating: single-frame non-identifiability of access under distribution-matched replacement, and identification of correction benefit by a synchronized reference under a conditional orthogonality condition that replaced references violate; and (iii) a minimal regime-marginalized causal-state updater showing how a frozen causal representation can be maintained when validity-informative evidence is available, together with matched ablations showing the roles of representation-specific scale, recursive state uncertainty, and alignment between latent access regimes and the learned causal-factor partition.

2 Related Work

Temporal causal representation learning. CRL recovers latent variables that match the causal structure of a system (Schölkopf et al., 2021). In time series, identifiability comes from observed intervention targets (Lippe et al., 2022; 2023b), binary interactions (Lippe et al., 2023a), temporal structure and nonstationarity (Yao et al., 2022; Song et al., 2023), or sparse mechanisms (Lachapelle et al., 2022). These methods determine which latent variables and mechanisms can be identified. TCO starts after identification: given the learned factors and dynamics, when should new evidence change the state?

Partial observability. Multi-view CRL identifies what is shared across partial views (Yao et al., 2024), sparsity principles handle instance-dependent partial observability (Xu et al., 2024), and CaRiNG recovers latent processes under non-invertible generation from temporal context (Chen et al., 2024). These works recover information from incomplete observations during learning. TCO holds the trajectory fixed and changes access after learning, including replacements that look like valid evidence.

State estimation with unreliable measurements. Filtering separates dynamics from measurement noise (Särkkä, 2013). Robust filters down-weight outlying innovations (Agamennoni et al., 2012; Fang et al., 2018), probabilistic data association averages over which measurement originates from the target (Bar-Shalom & Tse, 1975), and interacting multiple models average over switching regimes (Blom & Bar-Shalom, 1988). Learned state-space models filter in learned feature spaces (Fraccaro et al., 2017; Becker et al., 2019), and trackers learn object permanence through occlusion (Tokmakov et al., 2021). RME uses the regime-averaging principle as it stands; we do not propose it as a new filtering method. Our contribution lies elsewhere: identifying the post-identification TCO problem in temporal CRL, characterizing its information boundary, and testing what it takes to maintain a frozen learned causal state under changing access. Our reference score is related to learning from noisy targets (Lehtinen et al., 2018; Batson & Royer, 2019).

3 Temporal Causal Observability

Temporal causal representation learning aims to recover causal factors that persist and evolve over time. However, whether a factor exists and evolves is different from whether the current observation provides reliable evidence about it. We call this distinction **Temporal Causal Observability (TCO)**.

Causal state and perceptual access. Let $C_t = (C_t^1, \dots, C_t^K)$ denote the causal state at time t. Each factor evolves according to

$$C_t^k = f_k\Big(C_{t-1}^{\mathrm{Pa}(k)}, U_{t-1}, \epsilon_t^k\Big), \tag{1}$$

where U_{t-1} is an optional observed input and is omitted for autonomous dynamics. We introduce factor-wise perceptual access $O_t = (O_t^1, \dots, O_t^K)$ through the observation mechanism

$$X_t = g(C_t, O_t, \eta_t). \tag{2}$$

The two mechanisms play different roles: Eq. 1 determines how the state evolves, whereas Eq. 2 determines how that state is exposed to perception. Hence a factor may persist and evolve while its current access changes. An intervention $\mathrm{do}(O_t^k = 0)$ changes only the observation mechanism, it leaves the structural equation for C_t^k, the underlying causal trajectory, and the transition inputs unchanged.

Prediction and current encoding. We consider a pretrained and frozen temporal CRL model with encoder E_θ and learned transition F. The dynamics predict the current state from the previous estimate, while the encoder reads the current observation:

$$D_t = F(\hat{Z}_{t-1}, U_{t-1}), \qquad M_t = E_\theta(X_t). \tag{3}$$

Here D_t is the prediction from the learned causal dynamics and M_t is the current encoding. A disagreement $D_t \neq M_t$ is ambiguous: it may reflect a genuine innovation of the causal state or a change in perceptual access. The current encoding is faithful to the observation in either case.

3.1 What the Current Observation Can Reveal

Our interventions replace the carrier of an inaccessible factor by a plausible donor value and render the resulting observation through the same mechanism. For a standardized factor and a least-squares probe with full-access fit R^2_{full}, if a fraction p of observations retains access and the replacement is independent of the target, then

$$R^2_{\text{changed}} = (2p-1)R^2_{\text{full}}. \tag{4}$$

The derivation is given in Appendix A.1. This relation serves as a manipulation check: when the observed degradation follows Eq. 4, the encoder is faithfully passing the replaced evidence through rather than failing to represent the factor.

Proposition 1 (Single-frame non-identifiability)**.** *Suppose the replacement satisfies*

$$\tilde{C}^k_t \sim P(C^k_t \mid C^{-k}_t), \qquad \tilde{C}^k_t \perp C^k_t \mid C^{-k}_t,$$

and is rendered through the same observation mechanism with the same exogenous-noise law. Then X_t has the same distribution under $O^k_t = 1$ and $O^k_t = 0$. Consequently, any detector of O^k_t that observes only X_t has ROC AUC $1/2$, and

$$p(O^k_t = 1 \mid X_t) = p(O^k_t = 1). \tag{5}$$

The proof is given in Appendix A.2. Thus, under this construction, access validity cannot be recovered by improving the current-frame encoder alone; additional evidence must come from outside the current observation.

3.2 What a Synchronized Reference Can Reveal

Consider a correction rule fixed in advance, $U_t = u(D_t, M_t)$, and the target $Z^\star_t$ that represents the persistent causal state in the learned coordinates. The benefit of correcting is

$$\Delta_t = \|D_t - Z^\star_t\|^2_Q - \|U_t - Z^\star_t\|^2_Q, \qquad \|a\|^2_Q = a^\top Q a. \tag{6}$$

The target is unobserved. A synchronized reference $W_t = Z^\star_t + \xi_t$, such as another view of the same instant, yields the score $\Gamma_t = \|D_t - W_t\|^2_Q - \|U_t - W_t\|^2_Q$, and expanding both squares gives

$$\Gamma_t - \Delta_t = 2\langle U_t - D_t, \xi_t \rangle_Q. \tag{7}$$

Proposition 2 (Reference identification)**.** *Let the correction rule be fixed before scoring. If $\mathbb{E}[\langle U_t - D_t, \xi_t\rangle_Q \mid \mathcal{H}_t] = 0$ for conditioning information $\mathcal{H}_t$, then $\mathbb{E}[\Gamma_t \mid \mathcal{H}_t] = \mathbb{E}[\Delta_t \mid \mathcal{H}_t]$.*

Thus a synchronized reference need not be noise free; its error must be orthogonal, in expectation, to the correction direction. A replaced reference generally violates this condition, because its error is no longer centered around the maintained causal state. The learned dynamics cannot serve as the reference either: setting $W_t = D_t$ makes every non-trivial correction appear harmful.

Propositions 1–2 define the information boundary of TCO. The current frame need not reveal whether its evidence still refers to the maintained factor, while synchronized evidence evaluates a fixed correction only when the reference itself is suitable. Hence evaluating correction benefit does not, in general, identify which correction should be applied to an individual sample: access validity remains latent and must be inferred jointly with the causal state.

4 Maintaining Learned Causal State with RME

The previous section characterizes what the available evidence can and cannot identify. RME addresses the remaining update problem by treating perceptual access as a latent regime. Rather than selecting one access hypothesis and committing to its state update, it maintains the competing hypotheses and marginalizes their causal-state estimates.

Latent access regimes. Assume V synchronized encodings $M_t^1, \ldots, M_t^V$ of the same factor. An access pattern

$$s = (s_1, \ldots, s_V) \in \{0,1\}^V \tag{8}$$

indicates which encodings currently carry that factor. Access is assumed to persist within a block of L frames, whose boundary is known, and the prior over patterns is

$$\pi(s) = \prod_{v=1}^{V} p^{s_v}(1-p)^{1-s_v}, \tag{9}$$

where p is the nominal access probability. With the three synchronized encodings used in our main setting, there are $2^3 = 8$ possible regimes.

For each learned coordinate, we use a simple channel model. The dynamics prediction defines

$$Z_t \mid D_t \sim \mathcal{N}(D_t, q_t), \tag{10}$$

while an encoding is modeled as

$$M_t^v \mid Z_t, s_v \sim \begin{cases} \mathcal{N}(Z_t, r), & s_v = 1, \\ \mathcal{N}(m_0, \sigma_0^2 + r), & s_v = 0, \end{cases} \tag{11}$$

where m_0 and σ_0^2 are the training moments of the learned coordinate. TCO defines perceptual access at the causal-factor level: the coordinates assigned by the frozen CRL model to the same causal factor C_t^k therefore share the corresponding access variable O_t^k, and hence one latent access regime. This ties the update to the factorization supplied by the learned causal representation rather than treating latent coordinates as exchangeable.

Regime-marginalized update. Let $\mathcal{I}_t$ denote the observed prefix of the current block. Integrating out the latent state in Eqs. 10–11 gives a likelihood $\ell_u(s)$ for each regime. Its posterior weight is

$$w_t(s) \propto \pi(s) \prod_{u \in \mathcal{I}_t} \ell_u(s). \tag{12}$$

Conditioned on pattern s, the causal-state estimate combines the dynamics prediction with exactly those encodings that are valid under that pattern:

$$\mu_t(s) = \frac{D_t/q_t + \sum_{v:s_v=1} M_t^v/r}{1/q_t + |s|/r}, \qquad |s| = \sum_v s_v. \tag{13}$$

If no encoding is valid, the update retains D_t. RME then averages over the latent access regimes,

$$\boxed{\hat{Z}_t = \sum_s w_t(s)\mu_t(s)} \tag{14}$$

instead of committing to the most probable pattern. The same regime weights give the posterior variance of the maintained state by the law of total variance; the full expression is given in Appendix A.5.

Importantly, the reference score in Section 3.2 does not determine the regime weights in Eq. 12. The reference analysis characterizes when a fixed correction can be evaluated, whereas RME directly infers a posterior over latent access patterns from the synchronized measurements and the learned dynamics.

Learned coordinates and recursive maintenance. Two additional quantities matter when applying Eq. 14 to a frozen causal representation. First, learned coordinates have no natural physical scale. We therefore estimate the measurement variance r from the frozen encoder on training data and the transition residual variance from held-out training trajectories; no test state or access label enters this calibration. Second, under recursive maintenance, $\hat{Z}_{t-1}$ is itself uncertain. Its variance is therefore propagated through the frozen transition to obtain the predictive uncertainty q_t, rather than treating the previous estimate as exact.

We consider two protocols. In a one-step update, the dynamics prediction starts from the clean encoding of the preceding frame and isolates a single update. In recursive maintenance, the model receives one clean encoding at $t = 0$ and subsequently feeds back its own maintained causal-state belief. Both protocols use only observations available up to the current time. Calibration, uncertainty propagation, the integrated likelihood, and the complete recursive algorithm are detailed in Appendix A.5.

Table 1: **Changing access breaks recovery; evidence beyond a single current observation is required.** $p = 0.8$. (A) Probe R^2 under full and changed access with IND or COH replacement, and the prediction of Eq. 4. CaRiNG is scored on frames whose target is hidden (effective $p = 0$). (B) Aggregate R^2 on CITRIS–Voronoi when a model receives more information. The GRU is trained on state labels; it is an information-budget diagnostic, not a matched baseline for RME.

A. Faithful encoders pass donors through

System	Full	IND	COH	$(2p-1)R^2_{\text{full}}$
CITRIS–Voronoi	.985	.557	.595	.591
BISCUIT–Arm-Voronoi	.985	.534	.593	.591
CaRiNG–NG	.802	$-.802$	$-.805$	$-.802$
DMS, native	.599	.363	–	.359

B. What additional information helps?

Model	IND	COH
Frozen encoder	.557	.595
Access-matched retraining	.548	.585
Supervised GRU, one view	.691	.679

5 EXPERIMENTS

Setup. We freeze temporal CRL backbones trained on full observations: CITRIS on interventional Pong (three representation seeds) and on Voronoi (five), and BISCUIT on robotic-arm Voronoi (five independently generated and trained systems). CaRiNG and a mechanism-sparsity backbone (DMS) mark the boundary. Each test sequence is rendered as three synchronized views, which are replicated renderings of one scene with separate replacement masks. For every factor and view, access is withheld on 20% of four-frame blocks ($p = 0.8$), with masks drawn separately for each view. Withheld carriers are replaced in one of two ways: IND independently resamples the donor value at each frame, and COH copies a temporally coherent segment from another trajectory. We report factor-macro any-hidden error E_{H} on frames where at least one view misses the factor, all-valid error E_{V} on frames where all three views carry it, and aggregate R^2, all in standardized factor units through probes fitted on training data. Of the any-hidden factor-frames, 78–79% keep two valid views, 19–20% one, and 1.5–1.7% none, so E_{H} mainly measures maintenance from mixed valid and replaced evidence. A representation or system is the unit of replication, and intervals are paired bootstraps over these units. Within each seed, all compared methods share the same frozen encodings, transitions, masks, and donors (Appendix E).

5.1 CHANGING PERCEPTUAL ACCESS BREAKS CAUSAL-STATE RECOVERY

Table 1A shows the same pattern on four backbones. Recovery that is nearly perfect under full access falls to the value that Eq. 4 predicts: $0.6 \times 0.985 = 0.591$ on both Voronoi systems, 0.359 on DMS, and -0.802 on CaRiNG frames whose target is hidden. The encoders are faithful; they report the donor. Table 1B asks what additional information helps. Retraining the encoder on partially observed data leaves recovery at 0.548, consistent with Proposition 1. A supervised GRU that sees the history of one view reaches 0.691: temporal history recovers part of the lost information. Because it is trained on state labels, the GRU serves only to diagnose the value of that information, not as a baseline for RME.

The value of synchronized evidence is tested directly. We score the fixed correction $U_t = M_t^v$ with a reference that averages the two other views and compare how well the score ranks the realized benefit Δ_t against a shuffled reference. With two valid references, the gain in Spearman correlation is 0.55–0.60 (Figure 2a); with one valid reference it shrinks to 0.14–0.20; with none it turns negative. Spearman correlation measures ranking rather than unbiasedness, but the pattern is the one expected when replaced references violate the condition of Proposition 2: synchronized evidence helps where it is valid, so its validity must be inferred together with the state.

5.2 RME MAINTAINS STATE UNDER AMBIGUOUS EVIDENCE

Table 2 is the main test of causal-state maintenance. Each method receives one clean encoding at $t = 0$ and afterwards only the incoming views and its own causal-state belief, whose mean and variance it propagates through the frozen transition with the same Monte Carlo draws. The methods share the evidence, the scale, the access prior, and the propagation, and differ only in the decision rule: average over access regimes, commit to the most probable regime, or keep the views whose posterior validity is at least 0.5. Recursive RME has the lowest hidden error in every column. It

Table 2: **Recursive causal-state maintenance: same evidence, different decisions.** Every method starts from one clean encoding at $t = 0$ and then feeds back its own mean and variance through the frozen transition; all share the calibrated scale, the evidence of the current block prefix, and the same Monte Carlo draws; $p = 0.8$; means over three Pong and five Voronoi representations; lower is better. Block MAP commits to the most probable access regime; hard validity keeps the views whose posterior validity is at least 0.5; RME averages over access regimes.

	Pong, IND		Pong, COH		Voronoi, IND		Voronoi, COH	
Recursive update	E_{H}	E_{V}	E_{H}	E_{V}	E_{H}	E_{V}	E_{H}	E_{V}
Median of views (memoryless)	.3770	**.2409**	.3797	**.2409**	.1615	.0119	.1540	**.0119**
Block MAP regime	.3104	.2572	.3247	.2607	.0639	.0177	.0764	.0200
Hard validity at 0.5	.3107	.2565	.3214	.2589	.0577	.0169	.0715	.0187
RME (average regimes)	**.2901**	.2527	**.2972**	.2530	**.0424**	**.0117**	**.0514**	.0124

lowers hidden error below the memoryless median by 0.083–0.119 (22–74%) and below the two hard decisions by 0.015–0.028 (7–34%), in every one of the eight representations and for both replacement types (paired intervals in Appendix C.3). Figure 2b shows that the advantage over the median holds at every frame of a hidden block, and on Voronoi it grows as the block accumulates evidence.

All-valid frames show a second difference. Under recursion, hard decisions also raise all-valid error, by 1.5–3% on Pong and by 44–61% on Voronoi relative to RME. RME matches the median on Voronoi and costs 0.012 (5%) on Pong, the backbone with the weakest full-observation probes (R^2 from 0.69 to 0.88). A validity-stratified analysis (Appendix C.4) locates the gain in ambiguous cases: when only one view is valid, RME substantially reduces error in every representation, whereas the median is slightly better when two views agree. Thus RME's largest advantage occurs when valid evidence is the minority.

Table 3 repeats the comparison in single updates from the clean previous encoding, with classical robust updates and a second backbone, BISCUIT. RME again has the lowest hidden error in every column, and hard validity decisions lose to it on all 13 representations and systems for both replacement types. The gap is much smaller in one step: with IND replacement, hard decisions exceed RME's hidden error by 0.004 on Pong and 0.005–0.008 on Voronoi, against 0.020–0.021 and 0.015–0.022 under recursion, and on CITRIS they match RME on all-valid frames. The cost of committing to one regime grows about three- to fivefold when the state is maintained recursively, consistent with each wrong commitment being carried forward in the maintained state. A χ^2 gate and a local Student-t update judge each encoding by its own innovation, whereas RME also asks whether the encodings agree; both lag on hidden frames, although the Student-t update reaches the lowest all-valid error on Pong and BISCUIT.

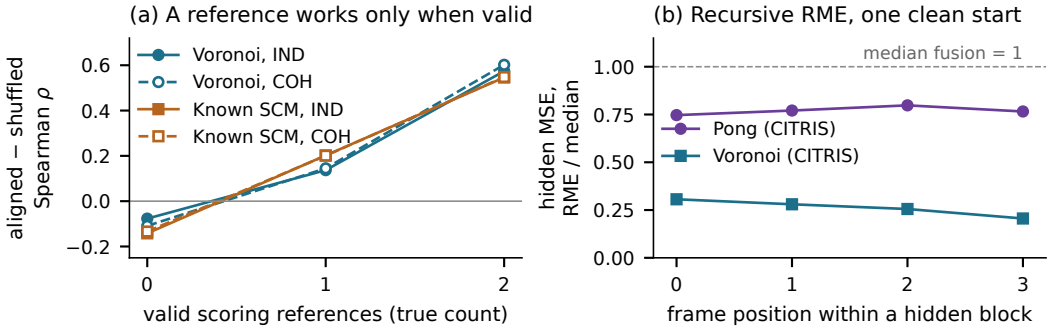

Figure 2: **References need validity; recursive RME keeps its advantage across hidden blocks.** (a) Spearman correlation between the synchronized score and the realized correction benefit, minus that of a shuffled reference, grouped by how many of the two scoring references are valid. (b) Recursive RME with calibrated scale and propagated uncertainty: any-hidden error relative to median fusion at each frame of a hidden block (IND replacement).

Table 3: **One-step updates with the same evidence.** One update from the clean previous encoding with the calibrated scale; IND replacement, $p = 0.8$. BISCUIT uses its native action-conditioned transition. Lower is better. COH replacement is in Appendix C.5.

	CITRIS–Pong		CITRIS–Voronoi		BISCUIT–Arm-Voronoi	
One-step update	E_{H}	E_{V}	E_{H}	E_{V}	E_{H}	E_{V}
RME (average regimes)	**.2570**	.2433	**.0326**	.0116	**.0424**	.0146
Block MAP regime	.2613	.2437	.0409	.0116	.0484	.0153
Hard validity at 0.5	.2610	.2437	.0378	.0116	.0474	.0151
Frame MAP regime	.2793	.2461	.0645	.0116	.0630	.0137
χ^2 innovation gate, 99%	.2913	.2565	.1010	.0147	.0967	.0199
Local Student-t, $\nu = 4$	.2961	**.2418**	.0991	.0116	.0750	**.0119**

Table 4: **Requirements for maintaining a learned causal state.** Lower is better; means over three Pong and five Voronoi representations; $p = 0.8$. (A) One update from the clean previous encoding, IND replacement; access validity is ignored, inferred, or read from the true access labels, and the scale is a universal constant or calibrated to the representation. (B) Recursive RME from one clean encoding; point estimate: the previous estimate is fed back as if exact; propagated: its uncertainty is pushed through the learned transition. COH in panel A gives the same ordering (Appendix C.1).

A. One update

			CITRIS–Pong		CITRIS–Voronoi	
Update	Validity	Scale r	E_{V}	E_{H}	E_{V}	E_{H}
Median of views	–	–	.2409	.3770	.01189	.1614
Precision-weighted, all views	ignored	calibrated	.2404	.4828	.01158	.3036
RME	inferred	fixed 0.49	.2615	.2891	.04599	.1003
RME	inferred	calibrated	.2433	**.2570**	.01159	**.0326**
Validity-informed update	given	calibrated	.2404	.2431	.01158	.0149

B. Recursive maintenance

		Pong, IND		Pong, COH		Voronoi, IND		Voronoi, COH	
Scale r	Previous state	E_{H}	E_{V}	E_{H}	E_{V}	E_{H}	E_{V}	E_{H}	E_{V}
fixed 0.49	point estimate	.4011	.3312	.4128	.3396	.2070	.1046	.2332	.1122
fixed 0.49	propagated	.3723	.2933	.3811	.2956	.1384	.0453	.1577	.0466
calibrated	point estimate	.3184	.2730	.3174	.2691	.0605	.0209	.0789	.0292
calibrated	propagated	**.2901**	.2527	**.2972**	.2530	**.0424**	.0117	**.0514**	.0124

5.3 What causal-state maintenance requires

RME itself is the regime marginalization of Eq. 14. The ablations ask what is required to apply this update reliably in learned causal coordinates. Table 4A isolates one update from the clean previous encoding. An update that ignores access validity trusts every donor: with the calibrated scale, hidden error reaches 0.304 on Voronoi, almost twice the median's 0.161. Inferring validity is not enough on its own either, because a universal fixed scale misweights prediction and encoding in coordinates with arbitrary units: with $r = 0.49$, RME rejects donors yet gives the prediction too much weight, and on frames where every view is valid it pays 0.046 against the median's 0.012. With the representation-specific scale, RME nearly matches the median on all-valid frames while reducing hidden error by 32% on Pong and 80% on Voronoi. Calibration details and validity diagnostics are in Appendix B.3.

Recursive maintenance adds one more requirement, because the previous causal state is estimated rather than observed. In Table 4B, the calibrated scale and uncertainty propagation each reduce recursive hidden-state error; together they give the lowest mean hidden-state error in every column, and with the calibrated scale, propagation lowers it by a further 0.018–0.028 in every seed. With a fixed scale and point-estimate feedback, recursive RME loses to the memoryless median in every column.

Causal-factor alignment. TCO assigns access factor-wise, so RME shares a regime across coordinates of the same learned causal factor. Holding observations, dynamics, calibration, and inference fixed, the learned partition reduces hidden-state error by approximately 30% relative to size-matched repartitions on CITRIS–Voronoi. It wins across all five representations under both replacement policies (Appendix C.11).

Evidence boundary. On native single-view DMS, RME is worse than the dynamics prior in all five seeds because the available evidence does not separate state change from access change. Three synchronized views restore a modest gain ($0.746 \rightarrow 0.726$; Appendix C.9). RME can combine informative evidence, but it cannot create information that is absent.

6 Discussion

What TCO adds to CRL. Temporal CRL is judged by the causal variables and dynamics it identifies under full observation. TCO adds a second requirement: a learned causal state must remain coherent when perceptual access changes without a corresponding intervention on the underlying causal state. The experiments show that post-CRL state maintenance requires more than recovering identifiable factors and dynamics. The update must also respect the scale of the learned representation, and it must distinguish uncertainty about the maintained causal state from uncertainty about whether current perceptual evidence still refers to that state. Recursion sharpens this distinction: collapsing the uncertainty about access into a hard decision costs little in a single update, but the error then enters the maintained state itself. Each link of the chain in Section 1 is visible in the data: retraining leaves the loss in place, a reference fails without valid references, an update that ignores access validity trusts donors, point-estimate recursion drifts, and recursive hard decisions fall behind regime averaging.

Limits of the current constructive setting. RME's channel model treats views as independent, but our views are replicated renderings: on frames where every factor is valid the three encodings are identical, and otherwise residual correlations average 0.65 on Voronoi and 0.76 on Pong. The likelihood of Eq. 11 thus acts as a scoring rule; its weights still rank access well (macro AUROC 0.91–0.996). The propagated variance stays nearly flat while the one-step prediction error grows with the time since the clean start (Appendix C.2), so the maintained belief is not yet a calibrated long-horizon uncertainty. The setting also supplies synchronized views, a nominal access probability, and known block boundaries, which carry real information (Appendix D); and our donors replace evidence, whereas physical occlusion also removes and distorts it. Full occlusion is tested only indirectly: 1.5–1.7% of hidden factor-frames have no valid view (22–52 factor-blocks per representation), and there RME is compared with the median of replaced views, not with the learned dynamics alone. The grouping ablation isolates one way in which learned causal structure enters maintenance: breaking factor co-membership while holding the observations, transition, calibration, and inference fixed increases hidden-state error by approximately 42–44% on CITRIS–Voronoi. This does not establish that the learned partition is the unique causal factorization or isolate which properties of the transition must themselves be causal. Inferring boundaries online, single-view maintenance, long-horizon calibration, and physical occlusion remain open.

7 Conclusion

Recovering causal factors and maintaining their states are different problems: CRL can identify a causal state and its dynamics without identifying whether current evidence should update that state. When perceptual access changes, a faithful encoder reports the wrong thing, and when the replacement is plausible, no single frame can reveal it. Treating factor-wise perceptual access as a latent regime, aligning that regime with the learned causal factors, marginalizing the corresponding state updates, and respecting the scale and uncertainty of the learned representation lets a frozen temporal CRL model maintain its causal state from a single clean start. We suggest judging temporal causal representations not only by the factors they recover, but also by how their beliefs respond when evidence changes.

AI USE STATEMENT

Generative AI tools were used to assist with research ideation and critique, theoretical discussion, feedback on experimental design, software implementation and debugging, analysis checking and result interpretation, literature organization, figure and table preparation, and manuscript drafting and editing. The authors made the research decisions, evaluated the experimental evidence, reviewed AI-assisted claims and artifacts, and take full responsibility for the content of this paper.

REPRODUCIBILITY STATEMENT

The appendix and anonymous supplementary material document the benchmark construction, causal sanity checks, model and data provenance, configurations, seed protocol, information available to each estimator, metric definitions, paired statistical comparisons, and the code required to regenerate the reported tables and figures from saved per-seed artifacts.

REFERENCES

Gabriel Agamennoni, Juan I. Nieto, and Eduardo M. Nebot. Approximate inference in state-space models with heavy-tailed noise. *IEEE Transactions on Signal Processing*, 60(10):5024–5037, 2012.

Yaakov Bar-Shalom and Edison Tse. Tracking in a cluttered environment with probabilistic data association. *Automatica*, 11(5):451–460, 1975.

Joshua Batson and Loic Royer. Noise2Self: Blind denoising by self-supervision. In *ICML*, pp. 524–533, 2019. URL https://proceedings.mlr.press/v97/batson19a.html.

Philipp Becker, Harit Pandya, Gregor Gebhardt, Cheng Zhao, C. James Taylor, and Gerhard Neumann. Recurrent Kalman networks: Factorized inference in high-dimensional deep feature spaces. In *ICML*, 2019.

Henk A. P. Blom and Yaakov Bar-Shalom. The interacting multiple model algorithm for systems with Markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, 1988. doi: 10.1109/9.1299.

Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11621–11631, 2020.

Guangyi Chen, Yifan Shen, Zhenhao Chen, Xiangchen Song, Yuewen Sun, Weiran Yao, Xiao Liu, and Kun Zhang. CaRiNG: Learning temporal causal representation under non-invertible generation process. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 7236–7259, 2024.

Huazhen Fang, Mulugeta Haile, and Yebin Wang. Robustifying the Kalman filter against measurement outliers: An innovation saturation mechanism. In *IEEE Conference on Decision and Control*, 2018. doi: 10.1109/CDC.2018.8619140. URL https://www.merl.com/publications/TR2018-173.

Marco Fraccaro, Simon Kamronn, Ulrich Paquet, and Ole Winther. A disentangled recognition and nonlinear dynamics model for unsupervised learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Sébastien Lachapelle, Pau Rodriguez, Yash Sharma, Katie E. Everett, Rémi Le Priol, Alexandre Lacoste, and Simon Lacoste-Julien. Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. In *Proceedings of the First Conference on Causal Learning and Reasoning*, volume 177 of *Proceedings of Machine Learning Research*, pp. 428–484. PMLR, 2022. URL https://proceedings.mlr.press/v177/lachapelle22a.html.

Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In *ICML*, pp. 2965–2974, 2018. URL https://proceedings.mlr.press/v80/lehtinen18a.html.

Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, and Efstratios Gavves. CITRIS: Causal identifiability from temporal intervened sequences. In *ICML*, 2022.

Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, and Efstratios Gavves. BISCUIT: Causal representation learning from binary interactions. In *UAI*, 2023a.

Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, and Efstratios Gavves. Causal representation learning for instantaneous and temporal effects in interactive systems. In *ICLR*, 2023b.

Simo Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013. URL https://users.aalto.fi/~ssarkka/pub/cup_book_online_20131111.pdf.

Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. doi: 10.1109/JPROC.2021.3058954.

Xiangchen Song, Weiran Yao, Yewen Fan, Xinshuai Dong, Guangyi Chen, Juan Carlos Niebles, Eric P. Xing, and Kun Zhang. Temporally disentangled representation learning under unknown nonstationarity. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to track with object permanence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.

Danru Xu, Dingling Yao, Sébastien Lachapelle, Perouz Taslakian, Julius von Kügelgen, Francesco Locatello, and Sara Magliacane. A sparsity principle for partially observable causal representation learning. In *ICML*, 2024.

Dingling Yao, Danru Xu, Sébastien Lachapelle, Sara Magliacane, Perouz Taslakian, Georg Martius, Julius von Kügelgen, and Francesco Locatello. Multi-view causal representation learning with partial observability. In *ICLR*, 2024.

Weiran Yao, Guangyi Chen, and Kun Zhang. Temporally disentangled representation learning. *Advances in Neural Information Processing Systems*, 35:26492–26503, 2022.

A Theory

All statements concern a fixed learned representation and the measurement model stated with each result.

A.1 Faithful encoding under replacement

Let the target factor C be standardized, let $\hat{C} = h(X)$ be a probe fitted by least squares, and let R^2_{full} be its fit on full observations. For a least-squares fit, $\mathrm{Var}(\hat{C}) = \mathrm{Cov}(\hat{C}, C) = R^2_{\text{full}}$, so $\mathbb{E}(\hat{C} - C)^2 = 1 - R^2_{\text{full}}$ on accessible frames. On a replaced frame the encoder reports the donor, $\hat{C}' = h(X')$ with X' rendered from $\tilde{C}$. If $\tilde{C}$ has the law of C and is independent of it, then $\mathrm{Var}(\hat{C}') = R^2_{\text{full}}$ and $\mathrm{Cov}(\hat{C}', C) = 0$, hence $\mathbb{E}(\hat{C}' - C)^2 = 1 + R^2_{\text{full}}$. Mixing a fraction p of accessible frames with $1 - p$ replaced frames gives

$$1 - R^2_{\text{changed}} = p(1 - R^2_{\text{full}}) + (1 - p)(1 + R^2_{\text{full}}) \quad \Longrightarrow \quad R^2_{\text{changed}} = (2p - 1)\, R^2_{\text{full}}.$$

Rendering nonlinearities, cross-factor perturbations of the encoding, and probe error move the observed values slightly away from this prediction; Table 13 lists the deviations.

A.2 Proof of Proposition 1

Write C_t^{-k} for the other factors. Under $O_t^k = 1$, $X_t = g\big((C_t^k, C_t^{-k}), \eta_t\big)$; under $O_t^k = 0$, $X_t = g\big((\tilde{C}_t^k, C_t^{-k}), \eta_t\big)$. Because $\tilde{C}_t^k$ is drawn from the law of C_t^k given C_t^{-k}, the pairs $(\tilde{C}_t^k, C_t^{-k})$ and (C_t^k, C_t^{-k}) have the same joint law, and η_t is independent of both. Hence X_t has the same law under both values of O_t^k. Any score $\phi(X_t)$ then has identical class-conditional distributions, so its ROC curve is the diagonal and its AUC is $1/2$. By Bayes' rule, $P(O_t^k = 1 \mid X_t) = P(O_t^k = 1)$. □

A.3 Reference identification

Let $D, U, Z^\star, W \in \mathbb{R}^d$, let $Q \succeq 0$ be symmetric, and let the quadratic losses be integrable. The correction U is fixed before scoring. Write $W = Z^\star + \xi$ and $a = U - D$. Expanding both scores,

$$\Gamma = D^\top Q D - U^\top Q U + 2(U - D)^\top Q W,$$
$$\Delta = D^\top Q D - U^\top Q U + 2(U - D)^\top Q Z^\star,$$

so $\Gamma - \Delta = 2a^\top Q\xi$, and $\mathbb{E}[a^\top Q\xi \mid \mathcal{H}] = 0$ implies $\mathbb{E}[\Gamma \mid \mathcal{H}] = \mathbb{E}[\Delta \mid \mathcal{H}]$. This projected orthogonality is weaker than independence. With $\mathcal{G} = \sigma(Z^\star, D, U, Q)$, the stronger condition $\mathbb{E}[\xi \mid \mathcal{G}] = 0$ gives $\mathbb{E}[\Gamma \mid \mathcal{G}] = \Delta$.

Validity and orthogonality. A replaced reference $\tilde{W}$ drawn independently of $(Z^\star, D, U)$ has $\mathbb{E}[\langle U - D, \xi\rangle_Q \mid \mathcal{G}] = \langle U - D, \mathbb{E}\tilde{W} - Z^\star\rangle_Q$, which is nonzero unless $U = D$ or $\mathbb{E}\tilde{W} = Z^\star$, so it violates the condition in general. A valid reference $W = Z^\star + \varepsilon$ satisfies the stronger condition $\mathbb{E}[\xi \mid \mathcal{G}] = 0$ only if its error is conditionally centered given $(Z^\star, D, U)$; an error shared with the candidate, such as a common encoding error across replicated renderings, breaks it. Access validity therefore governs the condition in our benchmarks without being equivalent to it, and the rank diagnostics of Figure 2a measure ranking quality, not unbiasedness.

Finite references. Suppose m references have conditionally independent projected errors $Y_j = a^\top Q\xi_j$ given $\mathcal{G}$, each sub-Gaussian with $\mathbb{E}[\exp(\lambda Y_j) \mid \mathcal{G}] \le \exp(\lambda^2\sigma_a^2/2)$. The averaged score satisfies $\bar{\Gamma} - \Delta = 2m^{-1}\sum_j Y_j$, a sub-Gaussian variable with variance proxy $4\sigma_a^2/m$, so

$$\Pr\big(|\bar{\Gamma} - \Delta| \ge u \mid \mathcal{G}\big) \le 2\exp\left(-\frac{mu^2}{8\sigma_a^2}\right).$$

A guarantee on the sign of a single update additionally requires a margin $|\Delta| > u$.

Two references. Direct expansion gives $(D - W_1)^\top Q(D - W_2) - (U - W_1)^\top Q(U - W_2) = (\Gamma_1 + \Gamma_2)/2$. Removing the reference-reference term does not relax the orthogonality condition.

A.4 When a score cannot act as a gate

Shared anchor. Let $D \sim \mathcal{N}(0, a)$, $e \sim \mathcal{N}(0, c)$ and $N_v, N_w \sim \mathcal{N}(0, b)$ be mutually independent, with $Z^\star = D + e$ and two replaced measurements $M_v = N_v$, $M_w = N_w$. Each donor is independent of the target, yet the residuals $r_v = N_v - D$ and $r_w = N_w - D$ share the anchor and satisfy $\mathrm{Cov}(r_v, r_w) = a$. For the candidate $U = M_v$ scored with $W = M_w$, the score is $S = 2r_v r_w - r_v^2$ and the true benefit is $\Delta = 2r_v e - r_v^2$. Gaussian fourth moments give $\mathrm{Cov}(r_v r_w, r_v^2) = 2a(a + b)$ and $\mathrm{Var}(r_v^2) = 2(a + b)^2$, while every term involving e vanishes. Therefore

$$\mathrm{Cov}(S, \Delta) = -4a(a + b) + 2(a + b)^2 = 2(a + b)(b - a),$$

which is negative whenever $a > b$: agreement between two invalid views then predicts harm, not benefit.

Decision ambiguity. Take identical evidence $D = 0$, $M = m \neq 0$ in two worlds: $Z^\star = 0$ with an invalid M, and $Z^\star = m$ with a valid M. An estimator that sees only this evidence returns the same x in both worlds, with worst-case squared error $\max\{x^2, (x - m)^2\} \ge m^2/4$, attained at $x = m/2$; Jensen's inequality extends the bound to randomized estimators. Additional channels, such as synchronized measurements, are what separate the two worlds.

Selection. Two candidates with zero true benefit and independent Rademacher scores Γ_1, Γ_2 have unbiased scores, yet $\mathbb{E}[\max(\Gamma_1, \Gamma_2)] = 1/2$. Scores stay unbiased only for candidates fixed before scoring, or when an independent reference evaluates the selected candidate.

A.5 RME: LIKELIHOOD, VARIANCE, AND ALGORITHM

For pattern s with valid set $J_s = \{v : s_v = 1\}$, the valid measurements share one state. Integrating $Z_t \sim \mathcal{N}(D_t, q_t)$ gives the coordinate likelihood

$$\ell_t(s) = \mathcal{N}\big(M_{t,J_s};\, D_t\mathbf{1},\; rI + q_t\mathbf{1}\mathbf{1}^\top\big) \prod_{v \notin J_s} \mathcal{N}\big(M_t^v;\, m_0,\, \sigma_0^2 + r\big)\,,$$

with the valid factor equal to one when J_s is empty. The rank-one term carries the shared state uncertainty; the determinant lemma and the Sherman–Morrison formula evaluate it in $O(|J_s|)$. Coordinate likelihoods multiply within a factor group, which shares one pattern. For each factor and block, the implementation (i) initializes log weights with $\log \pi(s)$; (ii) adds the integrated log likelihood of each new frame in the block prefix; (iii) normalizes by log-sum-exp and returns $\hat{Z}_t$ and P_t from Eqs. 12–15. With three views there are eight patterns, and the cost per block is proportional to $2^V LV d_g$ for a group of d_g coordinates. No score threshold and no access label enter the weights.

Uncertainty of the maintained state. The weights of Eq. 12 also give the posterior variance of the maintained causal state by the law of total variance,

$$P_t = \sum_s w_t(s) \left[(1/q_t + |s|/r)^{-1} + \big(\mu_t(s) - \hat{Z}_t\big)^2\right], \tag{15}$$

which the recursive update propagates to the next step.

Recursive algorithm. Algorithm 1 states the recursive update for one factor group. The anchored protocol replaces step 2 by the prediction from the clean previous encoding, $D_t = F(M_{t-1}^{\text{clean}}, U_{t-1})$, with $q_t = q_{\text{res}}$.

Algorithm 1 Recursive RME for one factor group

Require: frozen encoder E_θ and transition F; calibrated r and residual variance q_{res}; nominal p; block length L; clean initial encoding $\hat{Z}_0$ with $P_0 = r$

1: **for** $t = 1, 2, \ldots$ **do**
2: Draw 32 antithetic pairs $z^{(i)} \sim \mathcal{N}(\hat{Z}_{t-1}, P_{t-1})$; set $D_t \leftarrow \text{mean}_i\, F(z^{(i)}, U_{t-1})$ and $q_t \leftarrow \text{Var}_i\, F(z^{(i)}, U_{t-1}) + q_{\text{res}}$
3: Encode the synchronized views: $M_t^v \leftarrow E_\theta(X_t^v)$ for $v = 1, \ldots, V$
4: **if** t starts a block **then** set $\log w(s) \leftarrow \log \pi(s)$ for all $s \in \{0,1\}^V$
5: For every pattern s: $\log w(s) \leftarrow \log w(s) + \log \ell_t(s)$ and compute $\mu_t(s)$
6: Normalize w; output $\hat{Z}_t = \sum_s w(s)\mu_t(s)$ and P_t from Eq. 15
7: **end for**

Posterior mean. For conditioning information Y and $\mu = \mathbb{E}[Z \mid Y]$ under the channel model, every Y-measurable a satisfies $\mathbb{E}[\|a - Z\|_Q^2 \mid Y] = \|a - \mu\|_Q^2 + \text{tr}\{Q\,\text{Cov}(Z \mid Y)\}$. The average over patterns therefore minimizes conditional quadratic risk under the model, whereas the most probable pattern need not. This statement concerns one update given its prediction (D_t, q_t). The recursive algorithm is an assumed-density approximation (Särkkä, 2013): after every frame it collapses the regime mixture to its mean and diagonal variance, and the regime weights of later frames in the block are computed from these collapsed predictions. Exact inference under the model would carry a regime-conditioned state posterior for each access pattern through the block; the collapse discards the dependence between the state history and the regime. The regime weights are therefore weights of this approximate model, and the optimality statement does not extend to the recursive algorithm as a whole. An invertible affine reparameterization $z' = Az + b$ with $Q' = A^{-\top} Q A^{-1}$ preserves all quadratic comparisons.

Table 5: Benchmarks, their role, and the unit of replication.

System	Role	Unit
CITRIS–Pong	visual coordinates, fitted transition	3 representation seeds
CITRIS–Voronoi	four factors, fitted transition	5 representation seeds
BISCUIT–Arm-Voronoi	six factors, native action prior	5 independent systems
CaRiNG–NG	non-invertible observation	5 training seeds
DMS	single-view and repeated-view interfaces	5 seeds, fixed generator
Known-coordinate SCM	inference without representation learning	5 simulation seeds

B BENCHMARKS, CALIBRATION, AND INFORMATION ACCESS

B.1 ACCESS INTERVENTIONS AND UNITS

A carrier replacement changes the rendered evidence for one factor while leaving the stored causal trajectory and transition inputs unchanged. IND replacement independently resamples the donor value at each frame from the factor's empirical pool, whereas COH replacement copies a temporally coherent segment from another trajectory. We avoid calling IND replacement i.i.d., because the empirical donor pool can retain cross-factor dependence.

Each view receives its own replacement mask with marginal access probability $p = 0.8$, held constant within four-frame blocks. Pong masks hide a fixed 20% of blocks per factor and view, whereas Voronoi masks threshold separately seeded block-level uniform variables. Methods compared within a representation share the same underlying states, masks, and donors.

Figure 3 illustrates representative realizations of this intervention principle. In every case, the causal state continues to follow its original trajectory while only the factor-specific observation carrier is replaced during $O_t^k = 0$.

BISCUIT uses the same carrier-level Voronoi intervention principle, while DMS applies the analogous access manipulation through its native or repeated-view measurement interface. Table 5 summarizes the role and unit of replication of each benchmark.

B.2 BACKBONES AND QUALIFICATION

CITRIS. Encoders and factor assignments are frozen. The three Pong representations (seeds 42–44) reach full-observation probe R^2 of 0.877, 0.698 and 0.694 against a qualification threshold of 0.65; the angle's sine and cosine count as one of five factors. Two additional Pong trainings did not qualify: seed 45 assigns no latent to one factor, and seed 46 fails the separation gate. Pong uses at most 20,000 training and 5,000 test examples. Voronoi uses four factors, temporal edge probability 0.5, no instantaneous edges, single-target interventions and two flow layers; trajectories have 30,000, 3,000 and 3,000 steps for training, validation and test. CITRIS uses eight latents, hidden width 32, batch size 512, learning rate 10^{-3}, 200 epochs, and a 100-epoch causal encoder. Transitions are 96-tree ExtraTrees regressors fitted on training encodings; probes are ridge regressions with regularization 1.

BISCUIT. Arm-Voronoi has six factors and two-dimensional actions with unknown interaction targets. Five independently generated systems (150,000 examples, temporal edge probability 0.4, two flow layers) are trained with 12 latents, hidden width 32, batch size 256, learning rate 10^{-3} and 100 epochs. Qualification requires probe $R^2 \geq 0.95$; the systems reach grouped R^2 0.9888 ± 0.0038, held-out probe R^2 0.9849 ± 0.0028 and interaction F1 0.9929 ± 0.0007. RME uses the native action-conditioned prior mean and variance.

CaRiNG. Five seeds reach MCC 0.9066 ± 0.0100 under the official protocol. The target-hidden R^2 is 0.802 ± 0.022 under full access and -0.802 ± 0.056 and -0.805 ± 0.119 under the two donor policies. Training on a 40,000-example partial-access mixture gives -0.607 ± 0.067 and -0.614 ± 0.074; a supervised GRU gives 0.679 ± 0.019 and 0.575 ± 0.020.

DMS. The mechanism-sparsity backbone uses five training seeds on one fixed generator with 20-dimensional observations and two-frame trajectories. The native interface offers one measurement of

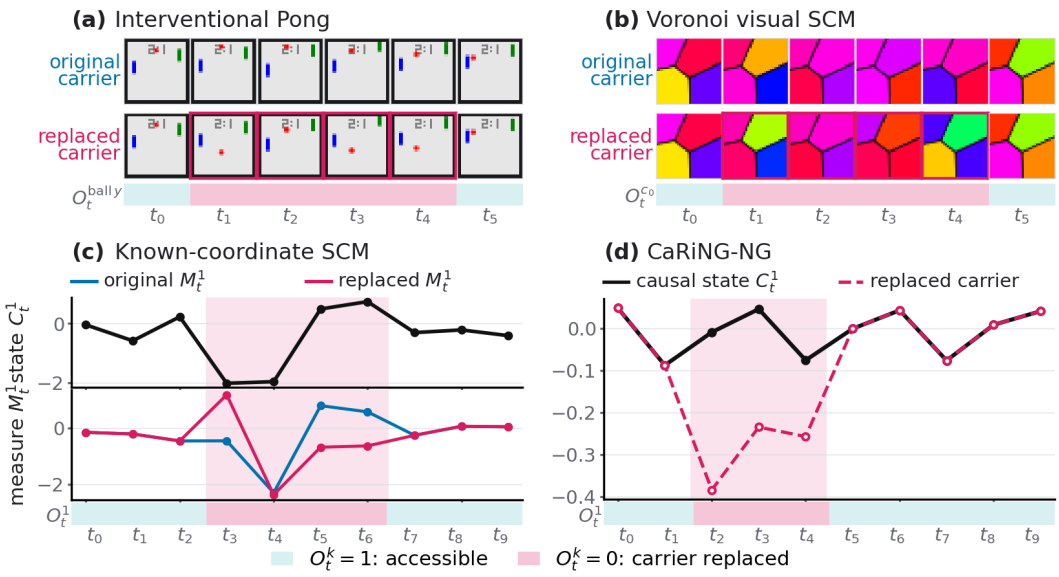

Figure 3: **Representative factor-wise perceptual-access interventions.** TCO interventions change the observation carrier of a selected causal factor while leaving its underlying causal trajectory and transition inputs unchanged. Shaded intervals denote $O_t^k = 0$; outside them, $O_t^k = 1$. Panels (a)–(b) show carrier replacement in visual systems, panel (c) separates the unchanged causal state from its replaced measurement in known causal coordinates, and panel (d) shows the same state–carrier distinction for CaRiNG-NG.

the second frame and a full-observation first frame. The repeated interface renders three synchronized views of the second frame through the fixed decoder with independent renderer noise; RME uses $L = 1$. The selected configuration uses the true latent dimension (10) and the true temporal edge count (55) as structural hyperparameters.

Known-coordinate SCM. Three factors follow $C_t^1 = 0.75C_{t-1}^1 + I_t^1 + \epsilon_t^1$, $C_t^2 = 0.55C_{t-1}^2 + 0.35\tanh(C_{t-1}^1) + I_t^2 + \epsilon_t^2$ and $C_t^3 = 0.50C_{t-1}^3 + 0.30\sin(C_{t-1}^2) + I_t^3 + \epsilon_t^3$, with innovation standard deviation 0.7, interventions with probability 0.08 and scale 0.8, and measurement noise 0.7. Seeds 542–546 each provide 80 episodes of 192 frames.

B.3 Scale calibration and uncertainty propagation

Measurement scale. For each CITRIS representation, 4,096 training frames at evenly spaced indices are encoded 16 times by sampling the frozen encoder posterior and passing each sample through the trained flow. Samples are normalized with the training latent mean and standard deviation used by RME, and $r_j = \max\{\text{median}_t \operatorname{Var}_s[z_j \mid x_t], 10^{-4}\}$. The estimate ranges from 2.7×10^{-4} to 1.4×10^2 across Pong coordinates and from 4.3×10^{-3} to 6.1×10^2 across Voronoi coordinates. BISCUIT uses the same posterior-sampling principle with 16 samples on 4,096 training frames. The invalid channel uses the training latent moments plus r.

Process scale. The transition residual variance q is estimated from five contiguous out-of-fold partitions of the training trajectory, conditioned on the intervention regime when at least 32 examples are available, and floored at 10^{-4}.

Propagation. At each step, 32 standard normal vectors and their negatives define 64 particles $\hat{Z}_{t-1} + \epsilon \odot P_{t-1}^{1/2}$ for every propagating method. All particles pass through the frozen transition in one call with the known intervention input; D_t is the particle mean and q_t is the particle variance plus the residual variance. Random numbers are shared across the compared conditions of a seed,

and the recursion starts from $P_0 = r$. Fixed-precision and validity-informed updates propagate their analytical posterior variance; the median is memoryless.

B.4 Information available to each method

Table 6: Information available to each method. An anchor is a clean encoding of the previous frame. No method receives future frames.

Method	Measurements	Clean anchor	Propagated var.	Access labels
Median / mean of views	current	–	–	no
Dynamics only	none	anchored or t=0	–	no
Precision-weighted, all views	current	anchored or t=0	recursive	no
RME update	block prefix	anchored or t=0	recursive	no
MAP, hard, χ^2, t (one step)	current or prefix	anchored	–	no
MAP, hard (recursive)	block prefix	t=0	recursive	no
Validity-informed	valid current views	anchored or t=0	recursive	yes
Supervised GRU	1 or 3 views, history	–	–	no (state labels)

B.5 Metrics and uncertainty

For factor k, let $H_{tk} = 1$ if at least one view misses the factor at time t and $V_{tk} = 1 - H_{tk}$. The any-hidden error E_{H} averages squared error over $H_{tk} = 1$ and the all-valid error E_{V} over $V_{tk} = 1$, both in standardized factor coordinates through probes fitted on training data, then averaged over factors. Any-hidden frames may still contain two valid views. Intervals resample whole representations or systems with replacement, pair methods within each unit, and use percentile endpoints; frames are never treated as independent units.

C Additional results

C.1 Anchored update

Table 7: One-step update from the clean previous encoding, both replacement types. Means over three Pong and five Voronoi representations; E_{V} / E_{H}.

Update	Pong, IND	Pong, COH	Voronoi, IND	Voronoi, COH
Median of views	.2409 / .3770	.2409 / .3798	.01189 / .1614	.01189 / .1540
RME update, $r = 0.49$	.2615 / .2891	.2615 / .2895	.04599 / .1003	.04582 / .1022
RME update, calibrated	.2433 / .2570	.2432 / .2588	.01159 / .03262	.01162 / .03434
Validity-informed, calibrated	.2404 / .2431	.2402 / .2435	.01158 / .01492	.01159 / .01487

Calibration lowers the paired hidden error of the RME update by 0.0321 on Pong IND (95% interval $[0.0285, 0.0372]$, 3/3 seeds) and by 0.0677 on Voronoi IND ($[0.0647, 0.0702]$, 5/5), and its all-valid error by 0.0182 and 0.0344. Aggregate R^2 rises by 0.026 and 0.061. With the fixed scale, the fixed-precision and validity-informed updates already pay an all-valid error of 0.0434 on Voronoi IND, against 0.0119 for the median; with the calibrated scale both reach 0.0116. The macro AUROC of the RME validity posterior rises from 0.857 to 0.913 on Pong IND and from 0.897 to 0.996 on Voronoi IND (COH: 0.806 to 0.907 and 0.880 to 0.994). A sweep over scalar multiples of 0.49, fixed before execution, never reaches the coordinate-wise estimate: its smallest value, 0.49×0.125, gives $E_{\mathrm{V}}/E_{\mathrm{H}} = 0.0149/0.0502$ on Voronoi IND and $0.2458/0.2628$ on Pong IND.

C.2 Recursive maintenance

Against the median, recursive RME with the calibrated scale and propagated uncertainty reduces hidden error by 0.0868 on Pong IND ($[0.0799, 0.1002]$, 3/3), 0.0825 on Pong COH ($[0.0775, 0.0911]$, 3/3), 0.1191 on Voronoi IND ($[0.1103, 0.1287]$, 5/5) and 0.1026 on Voronoi COH ($[0.0987, 0.1062]$, 5/5). With the calibrated scale, propagation reduces hidden error by 0.0283, 0.0202, 0.0181 and

Table 8: Recursive 2×2 ablation: aggregate R^2 / E_{H} / E_{V}. Point estimate: the previous estimate is fed back as if exact. The median row is memoryless.

Update	Pong, IND	Pong, COH	Voronoi, IND	Voronoi, COH
$r = 0.49$, point estimate	.629 / .401 / .331	.619 / .413 / .340	.815 / .207 / .105	.796 / .233 / .112
$r = 0.49$, propagated	.664 / .372 / .293	.658 / .381 / .296	.891 / .138 / .045	.879 / .158 / .047
calibrated, point estimate	.701 / .318 / .273	.703 / .317 / .269	.952 / .061 / .021	.936 / .079 / .029
calibrated, propagated	.726 / .290 / .253	.722 / .297 / .253	.968 / .042 / .012	.962 / .051 / .012
Median of views	– / .377 / .241	– / .380 / .241	– / .162 / .012	– / .154 / .012

0.0275 in the same order, winning in every seed. The all-valid difference to the median is $+0.0119$ on Pong IND ($[0.0037, 0.0201]$), $+0.0120$ on Pong COH ($[0.0055, 0.0194]$), -0.0002 on Voronoi IND ($[-0.0004, -0.0001]$, 5/5) and $+0.0005$ on Voronoi COH ($[-0.0001, 0.0014]$).

Error by block position and by sequence age. With IND replacement, hidden error of recursive RME at positions 0–3 of a hidden block is 0.286, 0.289, 0.296 and 0.291 on Pong (median 0.383, 0.375, 0.371, 0.380) and 0.0536, 0.0465, 0.0378 and 0.0321 on Voronoi (median 0.175, 0.166, 0.148, 0.156). Using the full-observation encoding as a proxy target, the mean propagated variance on Pong IND is 0.316 at ages 1–8 and 0.311 beyond age 512, while the squared one-step prediction error rises from 0.245 to 0.496; on Voronoi the variance moves from 0.427 to 0.430 while the error rises from 0.304 to 0.429. The propagated variance therefore tracks local uncertainty but not the slow growth of error over long horizons.

C.3 Recursive same-evidence controls

The recursive controls use exactly the evidence of recursive RME: the same frozen encodings, masks, donors, calibrated scale, residual variance, access prior, block-prefix likelihood, and Monte Carlo draws. Each control feeds back its own mean and diagonal variance through the frozen transition, with the same particles as RME; at the first step all three methods have identical predictive means and variances, which every run checks. Block MAP takes the regime-conditioned update $\mu_t(s^\star)$ and conditional variance of the most probable pattern $s^\star$ of the current block prefix; hard validity thresholds RME's marginal validity probabilities at 0.5 and applies a precision-weighted update to the retained views, with its analytical posterior variance. No test state or access label enters any update. The RME latent predictions in this comparison coincide exactly with those of the recursive ablation (maximum absolute difference 0 over all 16 seed–replacement runs), so Tables 2 and 4B share one RME row.

Table 9: Recursive same-evidence controls: hidden-error difference, control minus RME, with 95% intervals from 20,000 resamples of whole representations. Positive values favor RME, which wins in 3/3 Pong and 5/5 Voronoi representations in every cell.

Benchmark, replacement	Block MAP − RME	Hard 0.5 − RME
Pong, IND	$+.0203$ $[.0180, .0245]$	$+.0206$ $[.0146, .0247]$
Pong, COH	$+.0275$ $[.0192, .0347]$	$+.0242$ $[.0220, .0286]$
Voronoi, IND	$+.0215$ $[.0135, .0305]$	$+.0153$ $[.0109, .0197]$
Voronoi, COH	$+.0250$ $[.0143, .0357]$	$+.0201$ $[.0139, .0271]$

RME also has lower all-valid error than either control in every paired representation: control-minus-RME differences range from $+.0037$ to $+.0077$ on Pong and from $+.0053$ to $+.0076$ on Voronoi, with positive intervals in every cell. RME's aggregate R^2 is $.7255/.7219$ (Pong IND/COH) and $.9679/.9623$ (Voronoi IND/COH). With three Pong representations, the whole-representation intervals are coarse; the per-representation win counts carry the Pong evidence.

C.4 Recursive errors by number of valid views

Access labels enter only this evaluation. RME and median predictions are those of Table 2.

Table 10: Recursive RME versus the memoryless median by number of valid views. Share of any-hidden factor-frames, factor-macro MSE, and the paired difference (median minus RME) with 95% whole-representation intervals.

Benchmark	Valid	Share	Median	RME	Median − RME
Pong, IND	2	79.2%	.2401	.2655	$-.0254\ [-.0329, -.0205]$
	1	19.1%	.8644	.3658	$+.4986\ [.4403, .6007]$
	0	1.7%	1.3363	.6247	$+.7117\ [.6120, .8361]$
Pong, COH	2	79.2%	.2408	.2677	$-.0269\ [-.0410, -.0196]$
	1	19.1%	.8825	.3984	$+.4841\ [.4262, .5971]$
	0	1.7%	1.2618	.5434	$+.7183\ [.6447, .7860]$
Voronoi, IND	2	78.4%	.0124	.0129	$-.0005\ [-.0015, .0002]$
	1	20.0%	.6641	.1233	$+.5409\ [.5016, .5802]$
	0	1.5%	1.2037	.5063	$+.6975\ [.4255, .8858]$
Voronoi, COH	2	78.4%	.0123	.0170	$-.0047\ [-.0066, -.0028]$
	1	20.0%	.6294	.1527	$+.4767\ [.4458, .5076]$
	0	1.5%	1.2336	.5427	$+.6909\ [.5483, .8294]$

RME wins every representation with one or no valid view. With two valid views, the median wins every Pong and every Voronoi COH representation; Voronoi IND is mixed (2/5). The no-valid stratum (52 and 22 factor-blocks per representation) compares RME with a median of replaced views, not with the learned dynamics alone.

C.5 One-step same-evidence controls

Table 11: Same-evidence controls on BISCUIT–Arm-Voronoi, aggregate R^2 / E_{H} / E_{V}; five systems, one-step updates.

Update	IND	COH
RME update (average regimes)	.9626 / .04239 / .01458	.9642 / .04054 / .01457
Block MAP regime	.9583 / .04844 / .01533	.9602 / .04602 / .01535
Hard validity at 0.5	.9590 / .04742 / .01514	.9607 / .04537 / .01514
Frame MAP regime	.9510 / .06301 / .01372	.9548 / .05783 / .01370
χ^2 innovation gate, 99%	.9272 / .09669 / .01994	.9258 / .09904 / .01988
Local Student-t, $\nu = 4$	.9458 / .07497 / .01188	.9420 / .08066 / .01164

The BISCUIT controls were specified after the CITRIS comparison had been run. On BISCUIT, the hard validity rule has higher hidden error than RME by 0.0050 with IND ($[0.0038, 0.0062]$) and 0.0048 with COH ($[0.0037, 0.0060]$), and block MAP by 0.0061 ($[0.0046, 0.0076]$) and 0.0055 ($[0.0043, 0.0067]$); RME wins all five systems in each case. On CITRIS, the hard rule exceeds RME's hidden error by 0.0041 and 0.0050 on Pong IND and COH and by 0.0052 and 0.0058 on Voronoi IND and COH, with no seed in which it wins; block MAP exceeds it by 0.0043 and 0.0051 on Pong and by 0.0083 and 0.0086 on Voronoi. All controls use the calibrated scale and fixed hyperparameters: a 99% χ^2 gate with factor-group degrees of freedom, and five iterations of Student-t reweighting with $\nu = 4$.

C.6 Supervised GRU information budget

Table 12: Supervised GRU on CITRIS–Voronoi, aggregate R^2 / E_{H} / E_{V}; five systems, $p = 0.8$. Diagnostic only; not a matched RME baseline.

Input	IND	COH
One view	.6911 / .4534 / .0752	.6793 / .4724 / .0766
Three views	.8522 / .2327 / .0214	.8534 / .2311 / .0210

Both GRUs are trained on rendered training trajectories with IND access masks and causal-state supervision, using the same recipe: 20 epochs, hidden width 64, sequence length 32, learning rate

10^{-3}. They measure how much information one or three views carry for a supervised learner; they are not an information-matched comparison with RME. Three views raise aggregate R^2 by 0.161 with IND ($[0.158, 0.165]$) and 0.174 with COH ($[0.164, 0.184]$), 5/5 systems each. The three-view model has 18,308 parameters, the one-view model 15,236.

C.7 Channel dependence and the manipulation check

On every frame where all factors are valid in all views, the three frozen encodings are identical (176 frames per Pong seed and policy, 212–280 per Voronoi seed and policy). When the target factor is valid in all views but other factors may be replaced, pairwise residual correlations average 0.76 on Pong and 0.65 on Voronoi, and the participation rank of the three-view residual covariance averages 1.50 and 1.68. These diagnostics use test access labels only to form strata; they enter neither calibration nor inference.

Table 13: Observed changed R^2 against Eq. 4.

System, replacement	Effective p	Observed	Prediction	Difference
CITRIS–Voronoi, IND	0.8	.5573	.5911	$-.0338$
CITRIS–Voronoi, COH	0.8	.5952	.5911	$+.0041$
BISCUIT–Arm-Voronoi, IND	0.8	.5344	.5909	$-.0565$
BISCUIT–Arm-Voronoi, COH	0.8	.5927	.5909	$+.0018$
CaRiNG–NG, hidden target, IND	0	$-.8016$	$-.8021$	$+.0005$
CaRiNG–NG, hidden target, COH	0	$-.8046$	$-.8021$	$-.0025$

COH donors are not independent across frames, so their close agreement with the formula reflects the per-frame marginal rather than the independence premise.

C.8 Reference diagnostics

Table 14: Spearman correlation between the score of the fixed correction $U_t = M_t^v$ and its realized benefit, $p = 0.8$. The aligned reference averages the two other views; the residual score is $-\|M_t^v - D_t\|_Q^2$. Mean and sample SD over seeds.

System	Replacement	Aligned	Shuffled	Residual
CITRIS–Pong	IND	$0.236_{.024}$	$-0.104_{.008}$	$0.154_{.037}$
CITRIS–Pong	COH	$0.237_{.027}$	$-0.105_{.002}$	$0.153_{.036}$
CITRIS–Voronoi	IND	$0.417_{.037}$	$0.019_{.023}$	$-0.067_{.044}$
CITRIS–Voronoi	COH	$0.417_{.039}$	$-0.001_{.020}$	$-0.096_{.044}$
Known-coordinate SCM	IND	$0.558_{.002}$	$0.110_{.003}$	$0.170_{.005}$
Known-coordinate SCM	COH	$0.559_{.002}$	$0.111_{.003}$	$0.170_{.005}$

Figure 2a stratifies the aligned-minus-shuffled correlation by the true number of valid scoring references, which is used only after inference. For Voronoi IND/COH and the SCM IND/COH, the values are $-0.077/-0.109/-0.142/-0.134$ with no valid reference, $0.137/0.145/0.202/0.201$ with one, and $0.574/0.602/0.546/0.548$ with two.

C.9 DMS transfer

With one view, RME's hidden error exceeds that of the dynamics prior in all five seeds. With three views it falls from 0.746 to 0.726, while all-valid error rises from 0.632 for mean fusion to 0.698. On an entirely clean re-rendering the corresponding all-valid errors are 0.7007 and 0.7010.

C.10 Known-coordinate SCM

With known coordinates and correct scales, the ordering of Table 3 appears without representation learning: averaging over regimes beats the most probable regime, which beats hard elimination and residual selection, and block evidence adds to the per-frame update.

Table 15: DMS at $p = 0.8$, five seeds, $L = 1$: aggregate R^2 / E_{H} / E_{V} within the partially accessible stream.

Update	Native, one view	Repeated, three views
Paired full render	.599 / .701 / .701	.600 / .700 / .701
Dynamics only	.572 / .747 / .747	.572 / .746 / .748
Current measurement or view mean	.363 / 2.615 / .714	.506 / 1.092 / .632
Coordinate median	–	.540 / .948 / .661
Fixed precision	.373 / 2.544 / .712	.507 / 1.088 / .632
RME, frame	.582 / .793 / .716	.594 / .726 / .698
Validity-informed	.589 / .747 / .712	.631 / .665 / .632

Table 16: Update rules in known causal coordinates, $p = 0.8$, IND replacement, five seeds: aggregate R^2 / E_{H} / E_{V}. Both fusion rows include the dynamics prediction.

Update	R^2 / E_{H} / E_{V}
Dynamics only	.495 / .488 / .493
Precision fusion of prior and all views	.747 / .382 / .122
Median of views fused with prior	.747 / .317 / .179
Smallest-residual view	.604 / .443 / .328
Hard elimination	.733 / .361 / .165
MAP regime	.771 / .312 / .138
RME update, one frame	.791 / .280 / .130
RME update, block prefix	.800 / .264 / .128
RME update, full block (uses future frames)	.810 / .245 / .127
Validity-informed	.843 / .184 / .122

C.11 Causal-factor partition ablation

In TCO, perceptual access is indexed by causal factor: changing O_t^k changes how C_t^k is exposed to perception without changing the structural equation that evolves C_t^k. A learned CRL factor may occupy multiple latent coordinates, so RME assigns those coordinates a shared access regime. E51 asks whether this factor-wise alignment matters after identification.

We rerun recursive RME on the five frozen CITRIS–Voronoi representations while changing only the partition of latent coordinates into factor groups. For each representation, eight deterministic non-trivial repartitions preserve the complete assigned-coordinate pool and the original group-size multiset while breaking the original within-factor co-membership. Whole-group permutations are excluded because they only relabel factors. The encoder, transition, observations, masks, donors, calibration, access prior, block boundaries, and uncertainty-propagation draws are unchanged.

The eight shuffled results are averaged within each representation before computing paired differences, so the statistical unit remains the five frozen representations rather than the forty individual repartitions.

The learned causal-factor partition has lower hidden-state error than the mean matched repartition in all five representations under both replacement mechanisms. Relative to the shuffled-partition mean, it lowers E_H by 30.3% under IND and 29.4% under COH. The E_V comparison has the same 5/5 ordering, so this gain is not obtained by sacrificing all-valid frames relative to the matched partition control. The same ordering holds when only one of the three views remains valid, the common regime in which valid evidence is a minority: E_H is .1233 versus .1589 under IND and .1527 versus .1876 under COH, with the learned partition winning all five representations in both cases. The zero-valid stratum is sparse and is not used for the primary claim. This is a post-identification structural ablation, not a causal-discovery test. It shows that causal-state maintenance benefits from aligning factor-wise access with the factorization supplied by the frozen CRL model; it does not establish that this factorization is unique or recoverable from arbitrary observations.

Table 17: **Breaking causal-factor alignment degrades recursive maintenance on CITRIS–Voronoi.** Shuffled values are within-representation means over eight pre-specified non-trivial repartitions that preserve the assigned coordinate pool and group sizes. Positive ΔE_H means that breaking the learned factor alignment increases hidden-state error.

	Any-hidden error E_H			All-valid error E_V	
Replacement	Learned	Shuffled	ΔE_H [95% CI]	Learned	Shuffled
IND	.04240	.06086	$+.01846$ $[.01621, .02070]$	.01165	.01320
COH	.05139	.07281	$+.02142$ $[.01845, .02390]$	.01244	.01364

D PRELIMINARY EXECUTIONS

Preliminary executions used the fixed scale $r = 0.49$ and point-estimate feedback. Their RME results are superseded by the paired comparisons of Tables 2–4 and are not used for any claim. Four exhibits come from these executions because they do not involve the measurement scale: the recovery measurements of Table 1A and the retraining row of Table 1B, the reference diagnostics (Figure 2a and Appendix C.8), the DMS results (Appendix C.9), and the known-coordinate SCM (Appendix C.10). One preliminary result supports the Discussion: on BISCUIT, an RME update that used full-block evidence with block boundaries shifted by half a block performed worse than the per-frame update despite seeing future frames (hidden error 0.067 against 0.053).

E REPRODUCIBILITY AND PROVENANCE

The confirmatory paired experiments underlying Tables 2–4 were specified before execution. Each run records hashes of its source and configuration at start and end, aborts if either changes, and archives the executed source, checkpoint hashes, calibration indices, random seeds, and raw predictions. Within a representation, compared conditions use the same frozen checkpoint, cached encodings, transition, masks, donors, and, where applicable, matched Monte Carlo draws. Tree-based transitions can amplify small hardware-dependent differences between encodings, so paired comparisons are computed from a shared cache rather than from independently regenerated inputs.

The valid-view analysis in Appendix C.4 is a post-hoc descriptive evaluation of archived predictions rather than a new inference experiment. Its stratification protocol was fixed before inspecting stratum-specific errors, and true access labels are used only for evaluation. The causal-factor partition ablation in Appendix C.11 was registered before E51 inference. For each CITRIS–Voronoi representation, all eight deterministic, non-trivial, size-matched repartitions and the within-representation aggregation rule were fixed in advance. The learned-partition condition reproduces the corresponding E50 recursive RME latent predictions exactly (maximum absolute difference zero). No test causal state or true access label is available to inference in either experiment.

The supplement contains experiment configurations, executed-source snapshots, per-representation summaries, aggregation scripts, and unit tests covering the regime likelihood, mixture variance, recursive uncertainty propagation, one-step and recursive controls, causal-factor repartition checks, and information-access constraints. These checks include that future measurements cannot change past estimates, access labels reach only evaluation or the explicit validity-informed control, shuffled factor partitions preserve the assigned-coordinate pool and group-size multiset, and the learned-partition E51 condition reproduces E50. Training datasets and checkpoint binaries are regenerated from the official generators and training scripts described in Appendix B. Table 18 maps each reported exhibit to the supplement experiment or analysis that produces it.

Table 18: Reported exhibits and the supplement experiments or analyses that produce them.

Exhibit	Source experiment / analysis
Table 2, Table 9	`e50_{pong,voronoi}_corrected_recursive_controls`
Table 4B, Figure 2b, Table 8	`e43t_{pong,voronoi}_recursive_uncertainty`
Table 3 (CITRIS)	`e44_{pong,voronoi}_matched_controls`
Table 3 (BISCUIT), Table 11	`e47_biscuit_matched_controls`
Table 4A, Table 7	`e42s_{pong,voronoi}_paired_calibration`
Table 10	`e49_valid_view_strata` (post-hoc descriptive analysis of archived predictions)
Table 17	`e51_voronoi_causal_grouping_ablation`
Table 12	`e46_voronoi_fair_gru`
Table 13, channel-dependence diagnostics	`e45_replacement_r2_sanity`
Table 1A, Figure 2a, Tables 14–16	Preliminary executions described in Appendix D