IDENTIFIABLE HIERARCHICAL FACTOR DISENTANGLEMENT IN DUAL NON-IID TIME SERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Many nonstationary time series contain two coupled latent scales: a persistent sequence-level environment and recurrent mechanism states within each sequence. Flat switching representations conflate these sources of variation, obscuring the hierarchy, state cardinalities, and cross-environment state correspondence. We formalize this setting as dual non-IID time series and study when the persistent–recurrent hierarchy is identifiable from unlabeled sequences. Under identifiable within-sequence switching and translated invariant transition signatures, we recover the environment partition, both unknown cardinalities, and a shared indexing of recurrent states across environments. With sufficient transition variation, dynamic coordinates are identifiable up to a component permutation and componentwise transformations within contexts. Cross-context anchoring yields a global chart, and unique minimal target sufficiency further identifies the target-relevant subset. We also characterize representative recovery objects and ambiguities at the corresponding identification boundaries. These recovery stages motivate Hierarchical Dual Non-IID Component Estimation (HiDICE), a hierarchical variational state-space model with persistent–recurrent inference and additive signature sharing. Controlled experiments reproduce the predicted hierarchy recovery and stagewise behavior at the identification boundaries. Among evaluated label-free methods, HiDICE performs best under shortcut reversal on HPR-MNIST. It also achieves the lowest mean leave-one-subject-out heart rate error on PPG-DaLiA.

1 INTRODUCTION

Across a collection of nonstationary sequences, latent variation may operate at two latent scales as depicted in Figure 1(a): a sequence-level environment persists throughout a sequence, while mechanism states switch and recur within it (Linderman et al., 2017; Balsells-Rodas et al., 2024). We refer to this persistent–recurrent regime as dual non-IID time series. The two scales can be coupled, with recurrent mechanisms exhibiting environment-dependent frequencies or transition laws (Huang et al., 2020; Günther et al., 2023). A natural approach absorbs both scales into a flat switching state, $F_{n,t} = (E_n, R_{n,t})$. Figure 1(b) illustrates the resulting ambiguity. Although expressive, such a representation collapses the hierarchy: even if the flat switching states are recovered, their decomposition into persistent environments and recurrent mechanisms remains unidentified. The obstacle is therefore not expressive capacity, but identifiable hierarchical factorization.

Our key observation is that this hidden hierarchy leaves an observable geometric signature that resolves this factorization ambiguity. Once switching states are identifiable within each sequence up to permutation, a fixed observable functional of their state-conditional transition kernels yields an unordered signature set $\mathcal{Q}_n = b_{E_n} + \mathcal{W}$. Each sequence thus carries an environment-specific translation of a shared recurrent template. Set equality identifies the environment partition, while unique translations align recurrent states across environments. Thus, a latent factorization problem becomes an observable set alignment problem, captured geometrically in Figure 1(c).

Aligned recurrent states further enable stagewise recovery of the dynamic representation: from context-specific coordinates to a globally aligned chart, and, when target information is available, to the target-relevant subset.

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.

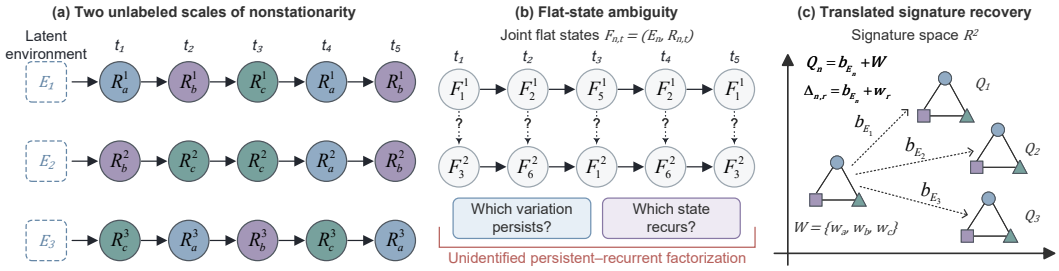

Figure 1: Persistent–recurrent ambiguity and recovery from translated signatures. (a) Persistent environments and recurrent mechanism states form two latent scales of nonstationarity. (b) Flat joint states leave their factorization and cross-environment correspondence unidentified. (c) Translated signature sets identify environments and align recurrent states.

Existing label-free work identifies within-sequence switching states or unknown domains (Hälvä & Hyvärinen, 2020; Balsells-Rodas et al., 2024; Song et al., 2023; 2024), while hierarchical temporal models organize time-indexed latents (Li et al., 2025b). None jointly identifies a sequence-fixed persistent environment and shared recurrent states when both cardinalities are unknown. Invariant methods instead target stable prediction rather than this latent organization (Peters et al., 2016; Krueger et al., 2021; Creager et al., 2021). Appendix B gives a broader comparison.

Motivated by this recovery argument, we introduce Hierarchical Dual Non-IID Component Estimation (HiDICE), a hierarchical variational state-space model that mirrors these recovery stages. We evaluate it through controlled recovery, shortcut reversal on Hierarchical Persistent–Recurrent MNIST (HPR-MNIST), and subject–activity transfer on PPG-DaLiA. Our contributions are:

- We establish identifiability of the persistent–recurrent hierarchy in the dual non-IID setting: from unlabeled sequences, the environment partition, both unknown cardinalities, and a shared cross-environment indexing of recurrent states are recoverable.
- We characterize a stagewise recovery chain for dynamic factors and target refinement, including representative recovery objects and ambiguities at its boundaries. Under progressively stronger conditions, recovery advances from context-specific coordinates to a global chart and then to the unique minimal target-relevant subset.
- We develop HiDICE, a hierarchical variational state-space model that translates the population identification results into finite-sample inductive biases, and evaluate it through controlled recovery, shortcut reversal, and held-out heart rate estimation.

2 Identifying the Persistent–Recurrent Hierarchy

2.1 Problem Setup

Dual non-IID setting. We observe N variable-length sequences $\mathcal{D} = \{\boldsymbol{X}_n\}_{n=1}^N$, where $\boldsymbol{X}_n = (X_{n,1}, \ldots, X_{n,T_n})$ and $X_{n,t} \in \mathbb{R}^{d_x}$. We use *dual non-IID* to denote a persistent–recurrent latent hierarchy with two discrete sources of variation operating at different temporal scales. The persistent environment $E_n \in [K_E]$ remains fixed throughout sequence n, while the recurrent mechanism states $R_{n,t} \in [K_R]$ switch and recur within the sequence. Their dynamics may depend on the persistent environment. Each sequence additionally contains an environment-conditioned static context S_n capturing sequence-specific static variation and continuous dynamic coordinates $Z_{n,t}$. The latent and observation dynamics satisfy

$$p(S_n \mid E_n), \quad p(R_{n,t} \mid R_{n,t-1}, E_n), \quad p(Z_{n,t} \mid H^Z_{n,t}, E_n, R_{n,t}), \quad X_{n,t} = F(Z_{n,t}, S_n) + \varepsilon_{n,t}. \tag{1}$$

Here $H^Z_{n,t}$ denotes the prescribed latent lag history, and h^Z denotes one of its realizations. Only the sequences $\mathcal{D}$ are observed; the latent assignments, change points, and cardinalities (K_E, K_R) are unknown.

Identification convention. We work in a population regime with fixed N and every $T_n \to \infty$. Unless stated otherwise, identification claims concern covered environments, contexts, and support components, and equalities are understood up to null sets. We assume absent or known independently deconvolvable observation noise. Appendix C gives the formal population, observational-equivalence, and coverage conventions.

2.2 Within-Sequence Switching Recovery

Each observed sequence law first identifies its recurrent switching representation up to a sequence-specific permutation.

Let $H_t^X = h_t^X(X_{1:t-1})$ denote the prescribed finite observed history. For a covered pair (e, s) and recurrent state r, let

$$\mu_{e,r,s}(B \mid h^X) := P\big(X_t \in B \mid H_t^X = h^X, E = e, S = s, R_t = r\big), \qquad B \in \mathrm{Bor}(\mathbb{R}^{d_x}), \quad (2)$$

denote a version of the state-conditional observed transition kernel. For each fixed W, let $P_{e,s}^{(W)} := \mathcal{L}(X_{u:u+W-1} \mid E = e, S = s)$. By conditional stationarity, this law is independent of u. The switching representation expresses it as a finite positive mixture over recurrent state paths having positive conditional probability.

Condition A (Identifiable within-sequence switching)**.** Within each sequence, the stationary window law has a unique minimal mixture over distinguishable positive-probability recurrent state paths. State and transition coverage, separated state-conditional kernels, and a nondegenerate minimal switching representation then identify the recurrent state quotient and its kernels up to a sequence-specific permutation. Conditional stationarity and ergodicity recover the corresponding population window laws from a long observed sequence as in Eq. (3).

Condition A recovers an unordered quotient of recurrent states within each sequence while allowing its labels to vary across sequences. For each fixed window length W and Borel set $A \subseteq \mathbb{R}^{W d_x}$, stationarity and conditional ergodicity give

$$\widehat{P}_{n,W,T}(A) := \frac{1}{T - W + 1} \sum_{t=1}^{T-W+1} \mathbf{1}\{X_{n,t:t+W-1} \in A\} \xrightarrow[T\to\infty]{\text{a.s.}} P_{E_n,S_n}^{(W)}(A). \quad (3)$$

Thus each sequence reveals stationary finite-window laws, not numerical (E_n, S_n) values or an observed state path. Translated signatures address the remaining cross-sequence ambiguity.

2.3 Translated Signature Hierarchy Recovery

Assumption 1 (Additive invariant kernel signatures)**.** A fixed observable functional Γ satisfies

$$\Delta_{n,r} := \Gamma(\mu_{E_n,r,S_n}) = b_{E_n} + w_r, \qquad b_e, w_r \in \mathbb{R}^{d_\Delta}. \quad (4)$$

The functional Γ is invariant to S_n, and the maps $e \mapsto b_e$ and $r \mapsto w_r$ are injective. Hence every sequence produces a translated copy of the common row set $\mathcal{W} = \{w_1, \ldots, w_{K_R}\}$:

$$\mathcal{Q}_n = \{\Delta_{n,r} : r = 1, \ldots, K_R\} = b_{E_n} + \mathcal{W}. \quad (5)$$

The decomposition has the unavoidable global translation gauge $b'_e = b_e + c$ and $w'_r = w_r - c$. The absolute origins of b_e and w_r are therefore not identified; the translated row sets, environment partition, and correspondence between recurrent states are gauge invariant.

Observable signature example. Suppose $X_t = (D_t, Q_t)$, where Q_t carries static nuisance information and the conditional law of D_t on a common history grid has natural parameter $\eta_{e,r}(h) = \eta_0(h) + Bb_e + L_R w_r$. If Γ discards Q_t and returns the identified natural parameter contrast of D_t relative to $\eta_0(h)$, then $\Gamma(\mu_{e,r,s}) = \bar{b}_e + \bar{w}_r$, where $\bar{b}_e = Bb_e$ and $\bar{w}_r = L_R w_r$, independently of s. Thus Γ is observable and uses no latent labels.

Theorem 1 (Persistent–recurrent hierarchy identification)**.** *Under Condition A, Assumption 1, and environment coverage (each $e \in [K_E]$ occurs in at least one observed sequence), observed sequence laws identify the environment partition, (K_E, K_R), environment-specific recurrent transitions, state-conditional kernels, and one shared cross-environment recurrent state indexing, up to separate environment and state permutations. The recurrent state law may depend on the environment through $A^{(e)}$ and its stationary distribution.*

Proof sketch. Condition A identifies each sequence's state-conditional kernels up to permutation, so Assumption 1 yields the unordered row $\mathcal{Q}_n = b_{E_n} + \mathcal{W}$. A finite nonempty set has no nonzero translation symmetry; together with injectivity of $e \mapsto b_e$ and environment coverage, row equality identifies the environment partition and K_E. Unique translations to a reference row then identify K_R and globally align the recurrent states. Full details are in Appendices C.2 and D.

3 From Aligned States to Dynamic Factors

Theorem 1 establishes a shared cross-environment indexing of recurrent states. We next use this alignment to identify the continuous dynamic representation.

3.1 Dynamic Coordinate Identification

The transition ratios of these aligned states share the observation map and static context, so the unknown Jacobian cancels. For an identified pair $(r_q, r_{0,q})$, let $\Lambda_{n,q}$ denote the resulting log Radon–Nikodym ratio:

$$\Lambda_{n,q}(f_{S_n}(z) \mid h_q^X) = \log \frac{p_Z(z \mid h_q^Z, E_n, r_q)}{p_Z(z \mid h_q^Z, E_n, r_{0,q})}. \tag{6}$$

Conditional innovation factorization makes the latent contrast additive:

$$\rho_q(z \mid h_q^Z) := \log \frac{p_Z(z \mid h_q^Z, E_n, r_q)}{p_Z(z \mid h_q^Z, E_n, r_{0,q})} = \sum_{j=1}^{d_Z} \rho_{qj}(z_j \mid h_q^Z). \tag{7}$$

Its first and second coordinate derivatives define

$$v_q := \big((\partial_{z_j}\rho_{qj})_{j=1}^{d_Z}, (\partial_{z_j}^2\rho_{qj})_{j=1}^{d_Z}\big), \qquad M_\rho := \begin{bmatrix} v_1^\top & \cdots & v_m^\top \end{bmatrix}^\top. \tag{8}$$

Condition B (Coordinate-separating transition variation)**.** On the covered joint support, $F : (z, s) \mapsto F(z, s)$ is injective and bimeasurable onto its image, and every covered slice $f_s : z \mapsto F(z, s)$ is a C^2 embedding.

(B1) Factorized and separating contrasts. Conditional innovations factorize across coordinates on the covered support. Let Ω_q be the common support of contrast q and $\Omega_\star := \bigcap_{q=1}^m \Omega_q$ their common contrast support. There are $m \geq 2d_Z$ identified log-transition contrasts satisfying

$$\operatorname{rank}\big(M_\rho(z, h_{1:m}^Z)\big) = 2d_Z \quad \text{for a.e. } z \in \Omega_\star. \tag{9}$$

(B2) Connected contrast support. Within each covered context, $\Omega_\star$ has full probability under the dynamic law and is nonempty, open, and path-connected. For every $j \in [d_Z]$, each nonempty coordinate fiber $\Omega_{\star,-j}(u) := \{z_{-j} : (u, z_{-j}) \in \Omega_\star\}$ is connected.

(B3) Dynamic indecomposability. For each covered static context, any measurable function of the dynamic state invariant across all covered transitions is almost surely constant, with the same value across covered environments. Thus a nontrivial component of Z cannot be reassigned to the time-invariant context S. This exclusion holds uniformly for every admissible representation compatible with the observed X-law.

Together with the observation map regularity above, these clauses rule out coordinate mixing and static–dynamic reassignment: factorization and full-rank variation eliminate local cross-coordinate derivative terms, connected support prevents incompatible chart choices across support pieces, and dynamic indecomposability prevents a time-invariant function of Z from being reassigned to S.

Condition B permits a different chart T_s in each covered static context, consisting of a coordinate permutation followed by componentwise C^2 diffeomorphisms. A global chart requires unique alignment on every overlapping pair and a connected overlap graph.

Condition C (Cross-context anchoring)**.** For any context-specific charts compatible with the observed X-law, equality of the pushed-forward transition families on a positive-overlap context pair uniquely determines their relative chart. Formally, each positive-overlap edge satisfies

$$\Big[(T_s)_{\#}p_Z(\cdot \mid T_s^{-1}(\widetilde{h}), e, r) = (T_{s'})_{\#}p_Z(\cdot \mid T_{s'}^{-1}(\widetilde{h}), e, r),\ \lambda_{ss'}\text{-a.e. on } \mathcal{A}_{ss'}\Big] \Longrightarrow T_s = T_{s'} \text{ a.e.}, \tag{10}$$

The context graph has the covered static contexts as vertices and an edge (s, s') whenever $\lambda_{ss'}(\mathcal{A}_{ss'}) > 0$. We assume that this graph is connected. Thus anchoring is an additional edgewise-uniqueness and graph connectivity condition, not a consequence of within-context recovery.

Theorem 2 (Dynamic coordinate identification)**.** *Under Theorem 1, the population regime, and Condition B, the static context is identified by a bimeasurable bijection. Within each context, Z_t is identified by a coordinate permutation and componentwise C^2 diffeomorphisms. Under Condition C, one permutation and one set of componentwise diffeomorphisms apply globally across the connected context graph.*

This target-agnostic equivalence may relabel any later target subset. Explicitly, any X-equivalent admissible parameterization must satisfy

$$\widetilde{S} = \kappa(S), \qquad \widetilde{Z}_{t,j} = a_{S,j}\big(Z_{t,\pi_S(j)}\big)\,, \tag{11}$$

where κ is bimeasurable, π_S is a coordinate permutation, and every $a_{S,j}$ is a C^2 diffeomorphism. Condition C removes the S dependence, producing a single global chart.

Proof sketch. Theorem 1 supplies aligned state contrasts, and Eq. (6) cancels the observation Jacobian. Differentiating the resulting additive contrast identities and using the full-rank condition rules out cross-coordinate mixing, leaving only a permutation and componentwise C^2 diffeomorphisms; Condition C propagates this chart across the connected context overlap graph. Full details are in Appendices C.3–E.

3.2 Target-Dependent Refinement

Target-dependent refinement. Under cross-context anchoring, Theorem 2 yields a global target-agnostic chart. We then impose a unique-minimality condition on target sufficiency.

Assumption 2 (Unique minimal target sufficiency)**.** The population joint law of (X, Y) is observed, and there is a unique inclusion-minimal set $J_C \subseteq [d_Z]$ satisfying $Y \perp\!\!\!\perp Z_{-J_C} \mid Z_{J_C}$.

For this target-dependent refinement, define $(C, V) := (Z_{J_C}, Z_{-J_C})$.

Corollary 1 (Target-relevant subset)**.** *Among parameterizations with the same population joint law of (X, Y), under the global conclusion of Theorem 2 and Assumption 2, J_C and $(C, V) = (Z_{J_C}, Z_{-J_C})$ are identified up to the global coordinate equivalence; under $\widetilde{Z}_j = a_j(Z_{\pi(j)})$, the set becomes $\pi^{-1}(J_C)$. This is a representational identification result and does not by itself identify a causal effect or graph.*

Componentwise bijections preserve the sigma-field generated by each coordinate, while the global permutation only relabels coordinates. Hence target sufficiency and unique inclusion minimality are preserved under the same global relabeling.

Target-causal interpretation. If, in addition, $Y = g_Y(C, \varepsilon_Y)$ with $\varepsilon_Y \perp\!\!\!\perp (E, R, S, Z)$, then C admits a target-causal interpretation; this assumption is not required for the corollary.

Representative recovery boundaries. Representative recovery objects and ambiguities include local or unseparated switching classes, mixed within-context charts or coordinate-basis ambiguities, context-specific charts, and multiple sufficient target subsets. Formal target transfer and representative boundary constructions are collected in Appendices E and F.

4 The HiDICE Framework

HiDICE is a hierarchical variational state-space model guided by our stagewise identification analysis. As illustrated in Figure 2, it combines persistent–recurrent inference, alignment through additive transition signatures, coordinate-factorized dynamics, and target-subset refinement when Y

is observed. All components are optimized jointly under a structured variational objective, and the effective environment and recurrent-state cardinalities are estimated without hierarchy labels.

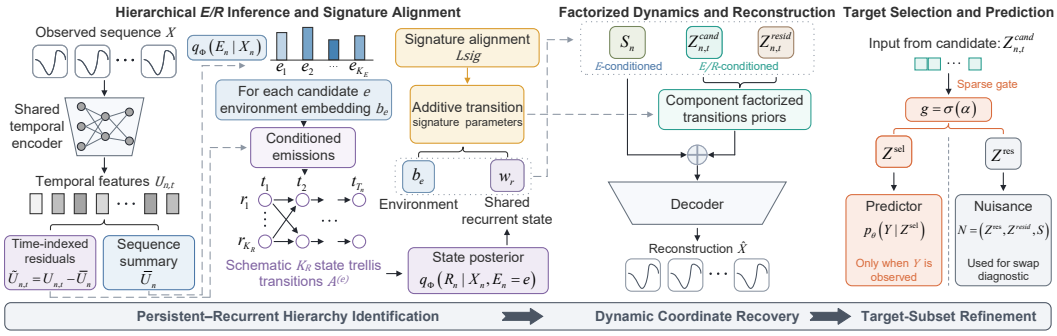

Figure 2: Theory-guided overview of HiDICE. HiDICE translates the stagewise identification chain into finite-sample inductive biases through hierarchical inference of persistent environments and recurrent mechanism states, alignment via additive transition signatures, and coordinate-factorized latent dynamics. When Y is observed, a sparse gate selects a subset of candidate dynamic coordinates for target prediction. The lower strip summarizes the corresponding theoretical recovery stages.

Hierarchical inference and state alignment. To represent uncertainty at the two latent scales, a sequence summary parameterizes $q_\phi(E_n \mid \boldsymbol{X}_n)$, while time-indexed residuals parameterize a chain of recurrent states conditioned on each candidate environment:

$$q_\phi(\boldsymbol{r} \mid \boldsymbol{X}_n, e) \propto \pi^{(e)}_{r_1} \exp\left(\ell^{(e)}_{n,1,r_1}\right) \times \prod_{t=2}^{T_n} A^{(e)}_{r_{t-1},r_t} \exp\left(\ell^{(e)}_{n,t,r_t}\right). \tag{12}$$

Forward–backward marginalizes path uncertainty. Conditional encoders also infer a static branch S_n, an overcomplete candidate dynamic branch $Z^{\text{cand}}_{n,t}$, and a residual dynamic branch $Z^{\text{resid}}_{n,t}$. The two dynamic branches are not assigned a priori to the population partition (Z_{J_C}, Z_{-J_C}).

To mirror the additive signature structure in Assumption 1, define $\mathcal{B} := \{\text{cand}, \text{resid}\}$. HiDICE uses shared learned estimator embeddings $\widetilde{b}_e$ and $\widetilde{w}_r$, distinct from the population signature components b_e and w_r. These embeddings enter each branch's transition parameters additively:

$$\eta^{\beta}_{e,r,j} = \eta^{\beta}_{0,j} + B^{\beta}_j \widetilde{b}_e + L^{\beta}_{R,j} \widetilde{w}_r, \qquad \beta \in \mathcal{B}, \tag{13}$$

so both dynamic branches retain E/R-conditioned transition structure. Figure 2 suppresses the tildes for readability. The fitted state-conditional transition distributions yield signatures $\widehat{\Delta}_{e,r} := \text{concat}(\widehat{\boldsymbol{m}}_{e,r}, \widehat{\boldsymbol{\ell}}_{e,r})$ from conditional-mean and log-marginal-variance features. These signatures are regularized toward the translated-row geometry:

$$\mathcal{L}_{\text{sig}} = (K_E K_R)^{-1} \sum_{e,r} \|\widehat{\Delta}_{e,r} - B_\Delta \widetilde{b}_e - L_{R,\Delta} \widetilde{w}_r\|_2^2. \tag{14}$$

Thus $\mathcal{L}_{\text{sig}}$ biases the fitted signatures toward the translated-row geometry induced by the shared environment and recurrent embeddings.

Factorized dynamics and target refinement. Given the aligned E/R states, HiDICE models each dynamic branch with coordinate-factorized transition families:

$$p_\theta(Z^{\beta}_t \mid H^{Z^{\beta}}_t, E, R) = \prod_j p_\theta(Z^{\beta}_{t,j} \mid H^{Z^{\beta}}_t; \eta^{\beta}_{E,R,j}). \tag{15}$$

Coordinate factorization instantiates the innovation separation bias in Condition B. Because the transition parameters are S-independent, sharing them across inferred static contexts supplies the cross-context alignment bias motivated by Condition C. The decoder uses all branches. When Y is observed, sparse gates $g_j = \sigma(\alpha_j)$ define $Z^{\text{sel}} = g \odot Z^{\text{cand}}$. Only Z^{sel} enters the target head, while $\mathcal{L}_{\text{min}} := \sum_j g_j$ encourages a minimal selection corresponding to the target-relevant subset.

Training and model order. Let $q_{\phi,n}$ denote HiDICE's structured posterior and $p^{\mathrm{lat}}_{\theta,n}$ its hierarchical prior; both include E_n, S_n, $\boldsymbol{R}_n$, and the candidate and residual dynamic paths. Let $\mathcal{O}_n \subseteq [T_n]$ index observed targets and be empty when Y is unavailable. The sequence ELBO is

$$\mathcal{E}_n = \mathbb{E}_{q_{\phi,n}}\left[\sum_{t=1}^{T_n} \log p_\theta(X_{n,t} \mid Z^{\mathrm{cand}}_{n,t}, Z^{\mathrm{resid}}_{n,t}, S_n) + \sum_{t\in\mathcal{O}_n} \log p_\theta(Y_{n,t} \mid Z^{\mathrm{sel}}_{n,t})\right] - \mathrm{KL}\big(q_{\phi,n} \| p^{\mathrm{lat}}_{\theta,n}\big)\,. \tag{16}$$

The reconstruction term uses Z^{cand}, Z^{resid}, and S, whereas the target term uses only Z^{sel}. The KL couples the structured posterior to the hierarchical prior. Training minimizes

$$\mathcal{J} = -\sum_n \mathcal{E}_n + \lambda_{\mathrm{sig}}\mathcal{L}_{\mathrm{sig}} + \lambda_{\mathrm{min}}\mathcal{L}_{\mathrm{min}} + \sum_{a\in\mathcal{A}_{\mathrm{reg}}} \lambda_a \mathcal{L}_a, \tag{17}$$

where $\mathcal{A}_{\mathrm{reg}}$ collects the auxiliary regularizers detailed in Appendix G.1. Candidate (K_E, K_R) capacities are selected by held-out ELBO with a prespecified standard error rule. Given the selected capacity pair, we define the effective latent-state cardinalities by thresholding the posterior occupancies:

$$\widehat{K}^{\mathrm{eff}}_E = \sum_e \mathbf{1}\{\widehat{\pi}_e \geq \delta\}, \qquad \widehat{K}^{\mathrm{eff}}_R = \sum_r \mathbf{1}\left\{\sum_e \widehat{\pi}_e \widehat{\rho}_{e,r} \geq \delta\right\}. \tag{18}$$

Here, $\widehat{\pi}_e$ is the posterior occupancy of environment e, $\widehat{\rho}_{e,r}$ is the posterior occupancy of recurrent state r conditional on environment e, and $\delta > 0$ is a prespecified threshold.

5 Experiments

5.1 Experimental Setup

We evaluate HiDICE along three theory-guided axes: controlled tests of hierarchy recovery and stagewise behavior; shortcut-reversal robustness on HPR-MNIST, constructed from MNIST (LeCun et al., 1998); and held-out heart rate prediction on PPG-DaLiA (Reiss et al., 2019). HPR-MNIST uses a target–color shortcut whose reliability depends jointly on the persistent environment and recurrent mechanism state and reverses at test time.

In all experimental tables, superscripts ^{I}, ^{L}, and ^{P}denote an in-house baseline defined for this study, label-assisted training, and an evaluation-only post-hoc hierarchy map, respectively. We report mean $\pm$ s.d. over five matched seeds. Primary endpoints are hierarchy and factor recovery, reversal accuracy, and LOSO MAE. Boundary tests report evaluation objects and metrics tailored to each construction. Sections G.2–G.3 specify the shared metrics and reporting rules. Appendix H provides the dataset-specific protocols and complete results.

5.2 Hierarchy Recovery and Identification Boundaries

The reference controlled SCM has $K_E = 3$ and $K_R = 4$, both withheld from the estimator, and covers every positive-probability (e, r) configuration. We probe the predicted recovery chain before relaxing signature separation, contrast rank, cross-context anchoring, and unique target minimality.

Reference recovery. Table 1 tests whether a label-free estimator can jointly recover both latent cardinalities, the persistent environment partition, cross-environment recurrent-state correspondence, and dynamic coordinates. HiDICE alone recovers both cardinalities in all five matched seeds and achieves the highest mean values for E ARI, R F1, and Z $\mathrm{MCC_S}$. Together, these results indicate that HiDICE's advantage spans successive hierarchy-to-coordinate recovery stages, a pattern consistent with preserving the persistent–recurrent factorization rather than collapsing the hierarchy into a flat switching representation.

Identification boundaries. These condition-targeted boundary interventions complement the component ablations by probing signature separation, transition variation, and cross-context alignment.

Table 1: Controlled SCM hierarchy and factor recovery.

Method	Exact ↑	E ARI ↑	R F1 ↑	Z MCC$_S$ ↑
Flat AR-HMM^{I,P}	2/5	0.595 ± 0.051	0.644 ± 0.042	0.664 ± 0.043
rSLDS (2017)	–	–	0.688 ± 0.051	0.706 ± 0.047
HM-NLICA (2020)^{P}	3/5	0.682 ± 0.054	0.792 ± 0.045	0.787 ± 0.035
CHiLD (2025b)	3/5	0.767 ± 0.047	0.759 ± 0.041	0.779 ± 0.038
CtrlNS (2024)	–	–	0.786 ± 0.036	0.812 ± 0.024
HiDICE (ours)	**5/5**	$\mathbf{0.886 \pm 0.046}$	$\mathbf{0.852 \pm 0.033}$	$\mathbf{0.862 \pm 0.027}$

Notes. Exact is the number of seeds recovering both (K_E, K_R); R F1 uses a single global alignment; – indicates that the method does not define the corresponding output.

Figure 3 summarizes recovery results along these three paths. At the signature-aliasing endpoint, sequence-local recurrent state F1 remains 0.776 ± 0.036, whereas globally aligned F1 is 0.601 ± 0.052. Along the rank path, at $\lambda = 0$, the generated σ_8 vanishes and its fitted counterpart is 0.028 ± 0.009, while the other seven fitted singular values remain bounded away from zero. At this endpoint, the trajectory subspace error is 0.176 ± 0.041 (Appendix G.3). Without anchoring, within-context MCC$_S$ remains 0.844 ± 0.033, whereas global MCC$_S$ is 0.593 ± 0.071.

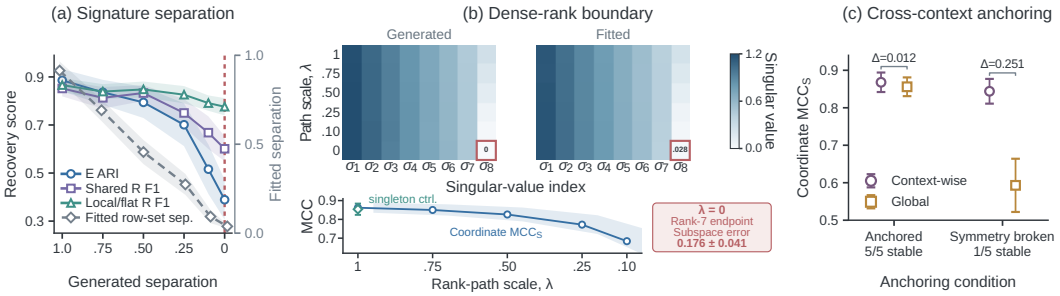

Figure 3: Recovery along three identification-boundary paths. (a) Under signature aliasing, sequence-local state recovery remains identifiable, whereas environment recovery and global state alignment do not. (b) At $\lambda = 0$, the generated contrast matrix has rank seven, and the endpoint is evaluated using trajectory subspace error. (c) Without cross-context anchoring, within-context coordinate recovery remains stable but global alignment degrades. Bands and error bars show mean $\pm$ one s.d.

5.3 Generalization under Distribution Shifts

Beyond controlled recovery, we test whether the hierarchy recovered by HiDICE supports generalization under two complementary distribution shifts. Table 2 reports shortcut reversal on HPR-MNIST and participant transfer on PPG-DaLiA.

Table 2: Generalization under distribution shifts. (a) HPR-MNIST shortcut reversal. (b) PPG-DaLiA heart rate MAE (BPM) under unseen participant–activity pairs, leave-one-subject-out (LOSO), and worst-activity evaluation.

(a) HPR-MNIST

Method	R F1 ↑	Leak. ↓	Rev. ↑
Temporal ERM^{I}	–	–	.371±.067
COGS (2026)	–	–	.704±.051
Flat AR-HMM^{I,P}	.698±.051	.203±.044	.619±.074
REx (2021)^{L}	–	–	.812±.036
HiDICE (ours)	**.831±.044**	**.091±.032**	**.788±.049**

(b) PPG-DaLiA

Method	Pair ↓	LOSO ↓	Worst-act. ↓
SAB-HRT (2025)	5.106±.144	6.217±.182	11.624±.315
Deep PPG (2019)	5.979±.172	7.412±.247	13.283±.372
PULSE (2026)	3.852±.127	4.785±.145	8.918±.281
Temporal ERM^{I}	4.031±.168	4.938±.163	9.181±.343
REx (2021)^{L}	3.794±.133	4.661±.154	8.634±.296
HiDICE (ours)	**3.687±.143**	**4.513±.126**	**8.346±.339**

Notes. R F1 uses a single global alignment. Leakage is the above-chance Δ accuracy of a probe using (Z^{resid}, S). REx^{L} receives the observed (E, R) groups. Boldface marks the best label-free result.

HPR-MNIST shortcut reversal. HPR-MNIST flips only the target–color shortcut reliability at test time while holding the digit and state processes and support fixed. Reversal accuracy therefore measures shortcut robustness. HiDICE achieves the highest label-free reversal accuracy (0.788 ± 0.049), together with higher globally aligned R F1 and lower target leakage than the flat AR-HMM. Label-assisted REx receives the true (E, R) groups and reaches 0.812 ± 0.036. The joint improvement in reversal accuracy, recurrent-state recovery, and leakage provides evidence that shortcut robustness arises from hierarchical factor separation rather than stronger prediction alone.

PPG-DaLiA participant transfer. Pair evaluation tests unseen participant–activity compositions, LOSO tests held-out participants, and worst-activity MAE summarizes activity-level robustness. LOSO is our primary participant-generalization setting and excludes the test participant from normalization, model selection, early stopping, and calibration. HiDICE achieves the lowest mean MAE among label-free methods on all three endpoints. Its LOSO MAE is 4.513 ± 0.126 BPM, compared with 4.785 ± 0.145 for PULSE. Relative to label-assisted REx, HiDICE's margin widens from 0.107 BPM on pair evaluation to 0.148 BPM under LOSO and 0.288 BPM on worst-activity evaluation. This widening margin indicates that HiDICE's advantage becomes more pronounced when evaluation emphasizes held-out participants or difficult activities.

5.4 Ablations and Diagnostic Evidence

Figure 4 shows stage-aligned ablation effects. Removing the persistent branch lowers E ARI by 0.485, while removing the recurrent branch lowers R F1 by 0.304. Removing signature decomposition lowers both metrics by 0.239 and 0.150, respectively, consistent with its role in hierarchy separation and cross-environment state alignment. Removing the factor split increases target leakage by 0.155 and reduces reversal accuracy by 0.137, supporting contributions to factor allocation and shortcut robustness. Removing target selection reduces controlled SCM J_C F1 by 0.284, whereas PPG-DaLiA LOSO MAE increases by 0.127 BPM. This contrast indicates that predictive error alone need not track target subset recovery.

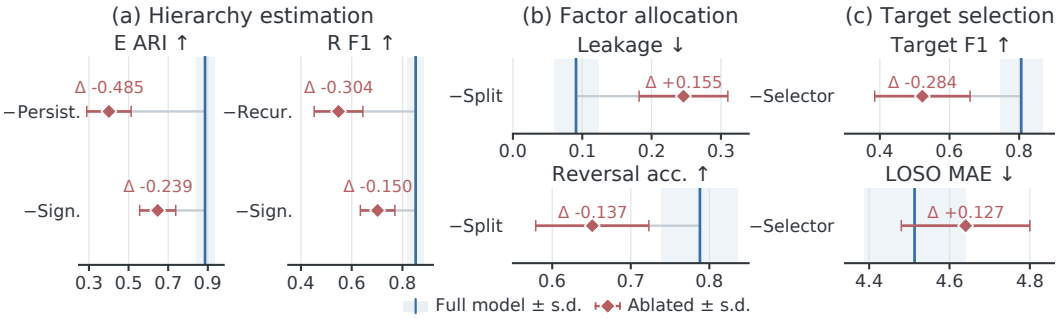

Figure 4: Component ablations for (a) hierarchy recovery, (b) factor allocation, and (c) target selection. Means ± one s.d.; annotations report $\Delta = \text{Ablated} - \text{Full}$.

6 Conclusion

We introduced dual non-IID time series and showed that observable translated transition signatures identify the persistent–recurrent hierarchy, enabling dynamic coordinate recovery. Cross-context anchoring yields a global chart, while unique minimal target sufficiency further identifies the minimal target-relevant subset. The guarantees are population-level and apply to covered environments, contexts, and support components under the stated identification conditions. HiDICE translates these conditions into finite-sample inductive biases, and the controlled diagnostics probe the predicted recovery stages rather than certify identifiability. Overall, the real-data experiments support predictive transfer under the evaluated shifts, while ground-truth latent mechanism recovery remains subject to the stronger assumptions specified above. An important direction for future work is to determine when target observations and persistent–recurrent mechanism variation jointly identify target-causal factors, potentially under weaker requirements than full dynamic-coordinate recovery.

AI USE STATEMENT

Generative AI tools were used for language editing, manuscript restructuring, and LaTeX revision. No generative AI tools were used for substantive research tasks, including the development of research ideas, theoretical results, experimental design, data generation or analysis, or determination of results. The authors reviewed all AI-assisted revisions for technical correctness and consistency with the underlying work and take responsibility for the final content.

ETHICS STATEMENT

No new human-subject data were collected. PPG-DaLiA is used under its public protocol, and learned states are not clinical categories without validation.

REPRODUCIBILITY STATEMENT

Appendices C–F provide assumptions, proofs, and counterexamples; Appendices G–H provide the training objective, diagnostics, experimental protocols, and full results. Code and experiment scripts will be made publicly available upon publication.

REFERENCES

Carles Balsells-Rodas, Yixin Wang, and Yingzhen Li. On the identifiability of switching dynamical systems. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 2639–2672. PMLR, 2024.

Carles Balsells-Rodas, Xavier Sumba, Tanmayee Narendra, Ruibo Tu, Gabriele Schweikert, Hedvig Kjellström, and Yingzhen Li. Causal discovery from conditionally stationary time series. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 2715–2741. PMLR, 2025.

Manuel Brenner, Elias Weber, Georgia Koppe, and Daniel Durstewitz. Learning interpretable hierarchical dynamical systems models from time series data. In *International Conference on Learning Representations*, 2025.

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. PMLR, 2024.

Elliot Creager, Joern-Henrik Jacobsen, and Richard S. Zemel. Environment inference for invariant learning. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 2189–2200. PMLR, 2021.

Sabrina C. Crepaldi, Jiabin Wang, Fumiya Matsumoto, Hiroki Takeuchi, Tatsuhiko Watanabe, and Yoshiharu Yamamoto. A stability- and aggregation-based method for heart rate estimation using photoplethysmographic signals during physical activity. *Sensors*, 25(14):4315, 2025.

Rodrigo González Laiz, Tobias Schmidt, and Steffen Schneider. Self-supervised contrastive learning performs non-linear system identification. In *International Conference on Learning Representations*, 2025.

Christian Gumbsch, Noor Sajid, Georg Martius, and Martin V. Butz. Learning hierarchical world models with adaptive temporal abstractions from discrete latent dynamics. In *International Conference on Learning Representations*, 2024.

Wiebke Günther, Urmi Ninad, and Jakob Runge. Causal discovery for time series from multiple datasets with latent contexts. In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 766–776. PMLR, 2023.

Wiebke Günther, Oana-Iuliana Popescu, Martin Rabel, Urmi Ninad, Andreas Gerhardus, and Jakob Runge. Causal discovery with endogenous context variables. In *Advances in Neural Information Processing Systems*, volume 37, pp. 36243–36284, 2024.

Hermanni Hälvä and Aapo Hyvärinen. Hidden Markov nonlinear ICA: Unsupervised learning from nonstationary time series. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence*, volume 124 of *Proceedings of Machine Learning Research*, pp. 939–948. PMLR, 2020.

Zhuo He, Shuang Li, Wenze Song, Longhui Yuan, Jian Liang, Han Li, and Kun Gai. Learning time-aware causal representation for model generalization in evolving domains. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 22523–22547. PMLR, 2025.

Çağlar Hızlı, Çağatay Yıldız, Matthias Bethge, S. T. John, and Pekka Marttinen. Identifying latent state transitions in non-linear dynamical systems. In *International Conference on Learning Representations*, 2025.

Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020.

Aapo Hyvärinen, Hiroaki Sasaki, and Richard E. Turner. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 859–868. PMLR, 2019.

Panagiotis Kasnesis, Lazaros Toumanidis, Daniele Jahier Pagliari, and Alessio Burrello. Replacing attention with modality-wise convolution for energy-efficient PPG-based heart rate estimation using knowledge distillation. *IEEE Journal of Biomedical and Health Informatics*, 30(1):339–352, 2026.

Ilyes Khemakhem, Diederik P. Kingma, Ricardo P. Monti, and Aapo Hyvärinen. Variational autoencoders and nonlinear ICA: A unifying framework. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2207–2217. PMLR, 2020.

David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (REx). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5815–5826. PMLR, 2021.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Haoyang Li, Ziwei Zhang, Xin Wang, and Wenwu Zhu. Invariant node representation learning under distribution shifts with multiple latent environments. *ACM Transactions on Information Systems*, 42(1):26:1–26:30, 2024.

Yuke Li, Yujia Zheng, Guangyi Chen, Kun Zhang, and Heng Huang. Identification of intermittent temporal latent process. In *International Conference on Learning Representations*, 2025a.

Zijian Li, Minghao Fu, Junxian Huang, Yifan Shen, Ruichu Cai, Yuewen Sun, Guangyi Chen, and Kun Zhang. Towards identifiability of hierarchical temporal causal representation learning. In *Advances in Neural Information Processing Systems*, volume 38, pp. 22318–22357, 2025b.

Scott W. Linderman, Matthew J. Johnson, Andrew C. Miller, Ryan P. Adams, David M. Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 914–922. PMLR, 2017.

Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, and Stratis Gavves. CITRIS: Causal identifiability from temporal intervened sequences. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 13557–13603. PMLR, 2022.

Sarah Mameche, Lénaïg Cornanguer, Urmi Ninad, and Jilles Vreeken. SPACETIME: Causal discovery from non-stationary time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 19405–19413, 2025.

Hiroshi Morioka and Aapo Hyvärinen. Causal representation learning made identifiable by grouping of observational variables. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 36249–36293. PMLR, 2024.

Khalid Oublal, Saïd Ladjal, David Benhaiem, Emmanuel Le Borgne, and François Roueff. Disentangling time series representations via contrastive independence-of-support on l-variational inference. In *International Conference on Learning Representations*, 2024.

Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: Identification and confidence intervals. *Journal of the Royal Statistical Society: Series B*, 78(5):947–1012, 2016.

Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. Deep PPG: Large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19(14):3079, 2019.

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.

Sagar Shrestha and Xiao Fu. Content-style learning from unaligned domains: Identifiability under unknown latent dimensions. In *International Conference on Learning Representations*, 2025.

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, pp. 8092–8113, 2023.

Xiangchen Song, Zijian Li, Guangyi Chen, Yujia Zheng, Yewen Fan, Xinshuai Dong, and Kun Zhang. Causal temporal representation learning with nonstationary sparse transition. In *Advances in Neural Information Processing Systems*, volume 37, pp. 77098–77131, 2024.

Xinxin Song, Yuxiao Cheng, Tingxiong Xiao, and Jinli Suo. COGS: A causal representation learning framework for out-of-distribution generalization in time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 25572–25580, 2026.

Julius von Kügelgen, Michel Besserve, Wendong Liang, Luigi Gresele, Armin Kekić, Elias Bareinboim, David M. Blei, and Bernhard Schölkopf. Nonparametric identifiability of causal representations from unknown interventions. In *Advances in Neural Information Processing Systems*, volume 36, pp. 48603–48638, 2023.

Binghui Xie, Yongqiang Chen, Jiaqi Wang, Kaiwen Zhou, Bo Han, Wei Meng, and James Cheng. Enhancing evolving domain generalization through dynamic latent representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16040–16048, 2024.

Danru Xu, Dingling Yao, Sebastien Lachapelle, Perouz Taslakian, Julius von Kügelgen, Francesco Locatello, and Sara Magliacane. A sparsity principle for partially observable causal representation learning. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 55389–55433. PMLR, 2024.

Zirui Yan, Emre Acartürk, and Ali Tajer. Reward-oriented causal representation learning. In *Advances in Neural Information Processing Systems*, volume 38, pp. 46294–46346, 2025.

Dingling Yao, Caroline Muller, and Francesco Locatello. Marrying causal representation learning with dynamical systems for science. In *Advances in Neural Information Processing Systems*, volume 37, pp. 71705–71736, 2024.

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

Weiran Yao, Yuewen Sun, Alex Ho, Changyin Sun, and Kun Zhang. Learning temporally causal latent processes from general temporal data. In *International Conference on Learning Representations*, 2022b.

Qiuhao Zeng, Wei Wang, Fan Zhou, Gezheng Xu, Ruizhi Pu, Changjian Shui, Christian Gagné, Shichun Yang, Charles X. Ling, and Boyu Wang. Generalizing across temporal domains with Koopman operators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16651–16659, 2024.

Kun Zhang, Shaoan Xie, Ignavier Ng, and Yujia Zheng. Causal representation learning from multiple distributions: A general setting. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 60057–60075. PMLR, 2024.

Jonathan Yuyang Zhou and Yao Xie. Nonlinear sequence embedding by monotone variational inequality. In *International Conference on Learning Representations*, 2025.

APPENDIX

TABLE OF CONTENTS

A Core Notation **15**

B Related Work **15**

C Formal Assumptions and Observational Equivalence **16**

C.1 Population regime and observational equivalence 16

C.2 Switching-state identification 17

C.3 Coordinate-recovery conditions 19

C.4 Cross-context alignment and target conditions 20

D Proofs of Hierarchical State Identification **20**

D.1 Quotient of recurrent states within a sequence 21

D.2 Global–local hierarchy separation 22

E Proofs of Coordinate and Target Identification **23**

E.1 Static context preservation 23

E.2 Observable density ratio contrasts 24

E.3 Within-context coordinate recovery 24

E.4 Cross-context alignment 26

E.5 Target-minimality transfer 27

F Identification Boundaries and Counterexamples **27**

F.1 State hierarchy boundaries 27

F.2 Coordinate-recovery boundaries 28

F.3 Cross-context and target boundaries 29

G HiDICE Estimation and Shared Evaluation Protocol **29**

G.1 Optimization and inference 29

G.2 Reproducibility, model selection, and computational reporting 31

G.3 Metrics and alignment 31

H Experimental Protocols and Complete Results **32**

H.1 Baselines and implementation parity 32

H.2 Controlled SCM: protocol and results 33

H.3 HPR-MNIST: protocol and results 35

H.4 PPG-DaLiA: protocol and results 38

H.5 Cross-dataset ablations and diagnostics 39

A CORE NOTATION

Hierarchy Identification	
$E_n,\ K_E$	Persistent environment and environment cardinality.
$R_{n,t},\ K_R$	Recurrent mechanism state and recurrent-state cardinality.
$\mu_{e,r,s}$	State-conditional observed transition kernel.
$\Gamma,\ \Delta_{n,r},\ b_e,\ w_r$	Signature functional, state signature, environment shift, and recurrent template point.
$\mathcal{W},\ \mathcal{Q}_n$	Shared recurrent template and its environment-translated signature set.
Coordinate Identification	
$S_n,\ Z_{n,t}$	Static context and continuous dynamic coordinates.
$F,\ f_s$	Observation map and its context-specific dynamic slice.
$\rho_q,\ M_\rho,\ \Omega^\star$	Transition contrast, derivative contrast matrix, and common contrast support.
$T_s,\ \kappa,\ \pi_s,\ a_{s,j}$	Context-specific chart, static-context bijection, coordinate permutation, and componentwise transforms.
Equivalence and Target Refinement	
$[\theta]_X,\ [\theta]_{XY}$	Observational-equivalence classes under the laws of X and (X, Y), respectively.
$J_C,\ C,\ V$	Target-relevant index set, target-relevant factors, and complementary factors.
HiDICE Estimator	
$\widetilde{b}_e,\ \widetilde{w}_r$	Learned environment and recurrent-state embeddings.
$Z^{\mathrm{cand}},\ Z^{\mathrm{resid}},\ Z^{\mathrm{sel}}$	Candidate, residual, and target-selected representations.

B RELATED WORK

Temporal representation learning under unknown nonstationarity. Nonlinear temporal representation learning obtains identifiability from temporal dependence, observed auxiliary variables, interventions, or changing transition mechanisms (Schölkopf et al., 2021; Hyvärinen et al., 2019; Khemakhem et al., 2020; Yao et al., 2022b; Lippe et al., 2022; Chen et al., 2024). Unknown-intervention CRL identifies latent causal structure from partitioned datasets (von Kügelgen et al., 2023); ours leaves the environment partition, recurrent states, and both cardinalities unlabeled. Label-free extensions infer nonstationarity from observations: hidden Markov nonlinear ICA uses a latent Markov state (Hälvä & Hyvärinen, 2020), NCTRL uses a Markov domain prior (Song et al., 2023), and CtrlNS uses transition sparsity to infer time-indexed domains (Song et al., 2024). Related work studies support structure, endogenous contexts, or conditional stationarity (Oublal et al., 2024; Günther et al., 2024; Balsells-Rodas et al., 2025). SPACETIME discovers temporal/dataset partitions for observed-variable causal discovery (Mameche et al., 2025); these results do not identify the latent persistent–recurrent hierarchy, both cardinalities, and cross-environment state correspondence jointly.

Switching and hierarchical latent dynamics. Switching dynamical systems model regime-dependent dynamics through one switching variable (Linderman et al., 2017; Balsells-Rodas et al., 2024). The analysis of Balsells-Rodas et al. (2024) is closest to our within-sequence stage, but does not further decompose a persistent environment and recurrent state with shared cross-environment correspondence. Related models disentangle transition and observation changes or learn nonlinear

switching dynamics (Yao et al., 2022a; 2024; González Laiz et al., 2025; Hızlı et al., 2025). Others learn time-scale/domain hierarchies, sequence-specific dynamics, or intermittent factors (Gumbsch et al., 2024; Brenner et al., 2025; Zhou & Xie, 2025; Li et al., 2025a). CHiLD identifies a hierarchy among time-indexed variables (Li et al., 2025b); our upper state is instead sequence fixed and both hierarchy cardinalities are unknown.

Invariant and target-relevant representations. Invariant prediction, risk extrapolation, and environment inference seek stable predictors or expose predictive instability (Peters et al., 2016; Krueger et al., 2021; Creager et al., 2021; Li et al., 2024). COGS jointly discovers domains for label-free temporal OOD generalization (Song et al., 2026); related methods model drift, reward sufficiency, evolving domains, or distribution extrapolation (He et al., 2025; Yan et al., 2025; Xie et al., 2024; Zeng et al., 2024). Complementary identifiability results use multiple distributions, grouped observations, or sparsity (Zhang et al., 2024; Morioka & Hyvärinen, 2024; Xu et al., 2024), and content–style models separate shared and domain-specific factors (Shrestha & Fu, 2025). These objectives do not generally identify the persistent–recurrent organization studied here.

C Formal Assumptions and Observational Equivalence

This appendix formalizes the population regime and conditions underlying the identification results in Sections 2–3.

C.1 Population regime and observational equivalence

Population observation regime. Equation (3) defines the population experiment used throughout: N is fixed, every $T_n \to \infty$, and the conditional stationary window laws are observed for every fixed $W < \infty$. Let $\mathcal{K}_E, \mathcal{K}_R \subset \mathbb{N}$ be finite admissible candidate order sets containing the data-generating orders, and let Θ_{adm} be the union of the corresponding admissible model classes over $(k_E, k_R) \in \mathcal{K}_E \times \mathcal{K}_R$. The candidate sets are not fixed to the unknown true pair. For $\vartheta \in \Theta_{\mathrm{adm}}$, write $K_E(\vartheta)$ and $K_R(\vartheta)$ for its environment and recurrent state cardinalities, shared across all covered sequences within that parameterization, so $E_n^\vartheta \in [K_E(\vartheta)]$ and $R_{n,t}^\vartheta \in [K_R(\vartheta)]$. At the data-generating parameterization, $K_E := K_E(\theta)$ and $K_R := K_R(\theta)$. For any admissible parameterization ϑ, define the covered context sets

$$\begin{aligned} \mathcal{C}_{\mathrm{cov}}(\vartheta) &:= \{(E_n^\vartheta, S_n^\vartheta) : n \in [N]\}, \\ \mathcal{S}_{\mathrm{cov}}^{+}(\vartheta) &:= \{S_n^\vartheta : n \in [N]\}. \end{aligned} \tag{19}$$

At the data-generating parameterization θ, abbreviate these sets by $\mathcal{C}_{\mathrm{cov}}$ and $\mathcal{S}_{\mathrm{cov}}^{+}$. All conclusions are restricted to these sets; unrepresented static contexts are neither identified nor aligned.

Observed-sequence environment coverage. Every $\vartheta \in \Theta_{\mathrm{adm}}$ satisfies observed-sequence coverage: each $e \in [K_E(\vartheta)]$ occurs as E_n^ϑ for some observed sequence n. This is population coverage, not merely positive probability under an environment law.

Target law convention. The target subset corollary assumes the population law of (X, Y) is observed. Partial target observation requires additional missingness and support conditions and is outside that conclusion.

Observational equivalence. For the target-agnostic hierarchy and coordinate results, define

$$\theta \sim_X \widetilde{\theta} \quad \Longleftrightarrow \quad p_\theta(\boldsymbol{X}) = p_{\widetilde{\theta}}(\boldsymbol{X}), \qquad [\theta]_X := \{\vartheta \in \Theta_{\mathrm{adm}} : \vartheta \sim_X \theta\}. \tag{20}$$

When Assumption 2 is invoked, the full population joint law is observed and we additionally define

$$\theta \sim_{XY} \widetilde{\theta} \quad \Longleftrightarrow \quad p_\theta(\boldsymbol{X}, \boldsymbol{Y}) = p_{\widetilde{\theta}}(\boldsymbol{X}, \boldsymbol{Y}), \qquad [\theta]_{XY} := \{\vartheta \in \Theta_{\mathrm{adm}} : \vartheta \sim_{XY} \theta\}. \tag{21}$$

Thus $[\theta]_{XY} \subseteq [\theta]_X$. The hierarchy and coordinate results, including dynamic indecomposability, quantify over $[\theta]_X$; the target-dependent refinement restricts comparison to $[\theta]_{XY}$. Unless explicitly qualified otherwise, "observationally equivalent" below means X-equivalent. The identification lemmas below establish $K_E(\vartheta) = K_E(\theta) = K_E$ and $K_R(\vartheta) = K_R(\theta) = K_R$ for every $\vartheta \in [\theta]_X$;

these equalities are conclusions, not restrictions in the definition of Θ_{adm}. After this order alignment, the recovered discrete states may differ only by

$$\left.\begin{aligned}\widetilde{E}_n &= \pi_E(E_n), \quad \pi_E \in \mathfrak{S}_{K_E},\\ \widetilde{R}_{n,t} &= \pi_R(R_{n,t}), \quad \pi_R \in \mathfrak{S}_{K_R},\end{aligned}\right\} \qquad \theta \sim_X \widetilde{\theta}, \tag{22}$$

The static contexts satisfy

$$\widetilde{S}_n = \kappa(S_n), \qquad n \in [N], \tag{23}$$

for a bimeasurable bijection κ on the covered context set. Let I_j and $\widetilde{I}_j$ denote the coordinate projections of the covered dynamic supports in the two models. After global alignment,

$$\widetilde{Z}_{n,t,j} = a_j\big(Z_{n,t,\pi_Z(j)}\big), \qquad \pi_Z \in \mathfrak{S}_{d_Z}, \qquad a_j \in \mathrm{Diff}^2(I_{\pi_Z(j)}, \widetilde{I}_j). \tag{24}$$

Here $\mathrm{Diff}^2(I, J)$ denotes C^2 diffeomorphisms from I onto J. Without cross-context anchoring, replace (π_Z, a_j) by $(\pi_{Z,s}, a_{j,s})$ and use the corresponding context-specific support projections. All relations in Eqs. (23)– (24) are restricted to $\mathcal{S}^+_{\mathrm{cov}}$. The state relations in Eq. (22) compare latent-state models under a label-normalized coupling; they do not make a realized finite state path observable.

Observation regularity and convention. On the covered joint support, $F : (z, s) \mapsto F(z, s)$ is injective and bimeasurable onto its image. For every covered static context s, the slice $f_s : z \mapsto F(z, s)$ is a C^2 embedding on the covered dynamic support. The first clause makes the common-observation coupling used below well defined up to null sets; the second supplies the slice-wise inverse and derivatives used in the coordinate proof. For $\mathcal{M}_s := f_s(\operatorname{supp} Z)$, densities are taken with respect to its induced volume measure, so

$$\begin{aligned} J^{\mathrm{vol}}_{f_s}(z) &:= \det\big(Jf_s(z)^\top Jf_s(z)\big)^{1/2},\\ p_X(f_s(z) \mid \xi) &= \frac{p_Z(z \mid \xi)}{J^{\mathrm{vol}}_{f_s}(z)}, \qquad \xi = (h^Z, e, r, s). \end{aligned} \tag{25}$$

If $X^{\mathrm{obs}} = X + \varepsilon$ with known independent noise whose characteristic function is everywhere nonzero, the noise-free law is identified by characteristic function deconvolution.

C.2 Switching-state identification

Autoregressive switching family. Section 2 defines the observed transition kernel in Eq. (2). For every admissible parameterization ϑ and covered pair $(e, s) \in \mathcal{C}_{\mathrm{cov}}(\vartheta)$, the conditional joint process $(R_t, X_t)_{t\ge1}$ is stationary and its observed marginal is ergodic. Under P_ϑ, $R_t \in [K_R(\vartheta)]$. Conditional on $E = e$, (R_t) is a time-homogeneous Markov chain independent of S, with transition matrix $A^{(e)}_\vartheta = (A^{(e)}_{\vartheta,rr'})_{r,r'\in[K_R(\vartheta)]}$ and stationary law $\pi^{(e)}_\vartheta = (\pi^{(e)}_{\vartheta,r})_{r\in[K_R(\vartheta)]}$, so $\pi^{(e)}_\vartheta A^{(e)}_\vartheta = \pi^{(e)}_\vartheta$.

For $W \ge 1$ and $\boldsymbol{r} = (r_1, \dots, r_W) \in [K_R(\vartheta)]^W$, define

$$\begin{aligned} \omega^{(W)}_{\vartheta,e,\boldsymbol{r}} &:= P_\vartheta(R_{1:W} = \boldsymbol{r} \mid E = e) = P_\vartheta(R_{1:W} = \boldsymbol{r} \mid E = e, S = s)\\ &= \pi^{(e)}_{\vartheta,r_1} \prod_{t=2}^{W} A^{(e)}_{\vartheta,r_{t-1},r_t}, \qquad \mathcal{R}^+_{\vartheta,e,W} := \{\boldsymbol{r} \in [K_R(\vartheta)]^W : \omega^{(W)}_{\vartheta,e,\boldsymbol{r}} > 0\}. \end{aligned}$$

For each $\boldsymbol{r} \in \mathcal{R}^+_{\vartheta,e,W}$, let

$$\begin{aligned} P^{(W)}_{\vartheta,e,s} &:= \mathcal{L}_\vartheta(X_{1:W} \mid E = e, S = s)\,,\\ P^{(W)}_{\vartheta,e,s,\boldsymbol{r}} &:= \mathcal{L}_\vartheta(X_{1:W} \mid E = e, S = s, R_{1:W} = \boldsymbol{r})\,,\\ P_{\vartheta,e,s,r} &:= \mathcal{L}_\vartheta(X_1 \mid E = e, S = s, R_1 = r)\,. \end{aligned}$$

The law of total probability gives, for every admissible ϑ, the positive-path mixture

$$\begin{aligned} P^{(W)}_{\vartheta,e,s} &= \sum_{\boldsymbol{r}\in\mathcal{R}^+_{\vartheta,e,W}} \omega^{(W)}_{\vartheta,e,\boldsymbol{r}} P^{(W)}_{\vartheta,e,s,\boldsymbol{r}},\\ \sum_{\boldsymbol{r}\in\mathcal{R}^+_{\vartheta,e,W}} \omega^{(W)}_{\vartheta,e,\boldsymbol{r}} &= 1. \end{aligned} \tag{26}$$

For the following factorization, fix any admissible ϑ and suppress its parameterization index; at $\vartheta = \theta$ this agrees with the notation in Section 2. Every admissible parameterization satisfies the switching-autoregressive property: for each positive path, $t = 2, \dots, W$, Borel set B, and almost every history under the corresponding component history law,

$$P\big(X_t \in B \mid X_{1:t-1} = x_{1:t-1}, E = e, S = s, R_{1:W} = \mathbf{r}\big) = \mu_{e,r_t,s}\big(B \mid h_t^X(x_{1:t-1})\big)\,. \tag{27}$$

and $P(X_1 \in B \mid E = e, S = s, R_{1:W} = \boldsymbol{r}) = P_{e,s,r_1}(B)$. Consequently, the positive-path component law factorizes as

$$P^{(W)}_{e,s,\boldsymbol{r}}(dx_{1:W}) = P_{e,s,r_1}(dx_1) \prod_{t=2}^{W} \mu_{e,r_t,s}\big(dx_t \mid h_t^X(x_{1:t-1})\big)\,, \tag{28}$$

Thus disintegration of an identified component recovers the state-conditional kernel at each covered path position. Let $\mathbb{W}_{\mathrm{id}}$ be a singleton containing a window that is uniformly identifying over Θ_{adm} in a finite-dimensional family, or an unbounded sequence of identifying windows in the nonparametric formulation; every $W \in \mathbb{W}_{\mathrm{id}}$ satisfies $W \geq 3$. For every $W \in \mathbb{W}_{\mathrm{id}}$, define the common model-class family

$$\mathfrak{P}_W := \Big\{ P^{(W)}_{\vartheta,e,s,\boldsymbol{r}} : \vartheta \in \Theta_{\mathrm{adm}}, \quad (e,s) \in \mathcal{C}_{\mathrm{cov}}(\vartheta), \quad \boldsymbol{r} \in \mathcal{R}^+_{\vartheta,e,W} \Big\}\,.$$

This is a union of probability measures; the path-index alphabets may differ across parameterizations of different recurrent orders. The path-injectivity clause below holds within every admissible representation, whereas the linear-independence clause applies to the full model-class family. Unless stated otherwise, fix $\vartheta \in \Theta_{\mathrm{adm}}$, $(e,s) \in \mathcal{C}_{\mathrm{cov}}(\vartheta)$, and $W \in \mathbb{W}_{\mathrm{id}}$ for the remaining clauses.

(i) Path-component identifiability. Distinct positive paths induce distinct component measures. In addition, for every $W \in \mathbb{W}_{\mathrm{id}}$, every $L \geq 1$, and every collection of distinct measures $Q_1, \dots, Q_L \in \mathfrak{P}_W$, linear independence gives

$$\boldsymbol{r}, \boldsymbol{r}' \in \mathcal{R}^+_{\vartheta,e,W}, \quad \boldsymbol{r} \neq \boldsymbol{r}' \quad \Longrightarrow \quad P^{(W)}_{\vartheta,e,s,\boldsymbol{r}} \neq P^{(W)}_{\vartheta,e,s,\boldsymbol{r}'}, \tag{29a}$$

$$\sum_{\ell=1}^{L} \alpha_\ell Q_\ell = 0, \quad (\alpha_1, \dots, \alpha_L) \in \mathbb{R}^L \quad \Longrightarrow \quad \alpha_1 = \cdots = \alpha_L = 0. \tag{29b}$$

(ii) Path coverage. Define the positive-transition support $\mathrm{supp}_+(A^{(e)}_\vartheta) := \{(r,r') : A^{(e)}_{\vartheta,rr'} > 0\}$. Positive paths cover every recurrent state and every positive transition:

$$\Big\{ r_t : \boldsymbol{r} \in \mathcal{R}^+_{\vartheta,e,W},\ 2 \leq t \leq W \Big\} = [K_R(\vartheta)], \tag{30a}$$

$$\Big\{ (r_t, r_{t+1}) : \boldsymbol{r} \in \mathcal{R}^+_{\vartheta,e,W},\ 2 \leq t \leq W-1 \Big\} = \mathrm{supp}_+(A^{(e)}_\vartheta). \tag{30b}$$

(iii) State-kernel separation. For $\boldsymbol{r}, \boldsymbol{r}' \in \mathcal{R}^+_{\vartheta,e,W}$, and $t, t' \in \{2, \dots, W\}$, write $\mu_{\vartheta,e,r,s}$ for the corresponding candidate kernel and define the two component history laws and their reference measure by

$$\eta_{\boldsymbol{r},t} := \mathcal{L}_{P^{(W)}_{\vartheta,e,s,\boldsymbol{r}}}(H^X_t), \qquad \eta_{\boldsymbol{r}',t'} := \mathcal{L}_{P^{(W)}_{\vartheta,e,s,\boldsymbol{r}'}}(H^X_{t'}),$$

$$\lambda_{\boldsymbol{r},t;\boldsymbol{r}',t'} := \tfrac{1}{2}(\eta_{\boldsymbol{r},t} + \eta_{\boldsymbol{r}',t'}), \qquad q_{\boldsymbol{r},t} := \frac{d\eta_{\boldsymbol{r},t}}{d\lambda_{\boldsymbol{r},t;\boldsymbol{r}',t'}}, \qquad q_{\boldsymbol{r}',t'} := \frac{d\eta_{\boldsymbol{r}',t'}}{d\lambda_{\boldsymbol{r},t;\boldsymbol{r}',t'}}.$$

Their essential history overlap is

$$\mathcal{H}^+_{\boldsymbol{r},t;\boldsymbol{r}',t'} := \{h : q_{\boldsymbol{r},t}(h) > 0,\ q_{\boldsymbol{r}',t'}(h) > 0\}\,.$$

Distinct recurrent states have distinguishable conditional kernels on a positive-measure subset of this overlap:

$$r_t \neq r'_{t'} \ \Longrightarrow \ \lambda_{\boldsymbol{r},t;\boldsymbol{r}',t'}\Big(\Big\{h \in \mathcal{H}^+_{\boldsymbol{r},t;\boldsymbol{r}',t'} : \mu_{\vartheta,e,r_t,s}(\cdot \mid h) \neq \mu_{\vartheta,e,r'_{t'},s}(\cdot \mid h)\Big\}\Big) > 0. \tag{31}$$

(iv) Nondegenerate minimal switching order. The switching representation has positive occupancy and full-rank transition dynamics, and its order is minimal:

$$\pi^{(e)}_{\vartheta,r} > 0 \quad \text{for every } r \in [K_R(\vartheta)], \qquad \operatorname{rank}(A^{(e)}_{\vartheta}) = K_R(\vartheta), \tag{32a}$$

$$\min\Big\{K \in \mathbb{N} : P^{(W)}_{\vartheta,e,s} \text{ admits an admissible } K\text{-state switching representation}\Big\} = K_R(\vartheta). \tag{32b}$$

The injectivity and linear-independence clauses in Eqs. (29a) and (29b) ensure that component matching is pathwise rather than only after aggregating duplicate component measures. The linear-independence clause is model-class, not true-parameter-only, and thus permits comparison of observationally equivalent mixtures. The reference measure in Eq. (31) exposes the essential overlap of each pair of component history laws. Equations (29)–(32) are used in Appendix D. These conditions are imposed at the process level. For a finite-dimensional identifiable switching family, one sufficiently long identifying window determines the initial law, transition matrix, and conditional-kernel parameters. In a nonparametric formulation, component matching and path coverage are required along an unbounded sequence of window lengths. This convention is used for posterior-process recovery below.

C.3 Coordinate-recovery conditions

Contrast derivatives. Fix an arbitrary covered context (e, s); all clauses in this paragraph are required for every $(e, s) \in \mathcal{C}_{\mathrm{cov}}$, and the context indices are suppressed. Let Ω_q be the common support of contrast q and $\Omega_\star := \bigcap_{q=1}^{m} \Omega_q$. Conditional innovation factorization means that, on the covered support,

$$p_Z(z \mid h^Z, e, r) = \prod_{j=1}^{d_Z} p_j(z_j \mid h^Z, e, r). \tag{33}$$

Consequently, for $q = 1, \ldots, m$, each log ratio between recurrent states in Eq. (7) has the component-additive form below. With $m \geq 2d_Z$ such contrasts, the required transition variation is the separate full-rank clause:

$$\rho_q(z \mid h^Z_q) = \sum_{j=1}^{d_Z} \rho_{qj}(z_j \mid h^Z_q), \qquad q = 1, \ldots, m, \tag{34a}$$

$$\operatorname{rank}\big(M_\rho(z, h^Z_{1:m})\big) = 2d_Z \quad \text{for a.e. } z \in \Omega_\star. \tag{34b}$$

Equation (34a) supplies the additive functional equation used in coordinate recovery, whereas Eq. (34b) eliminates its mixed-derivative null space. The support clause of Condition B is split into global connectedness and connected coordinate slices:

$$\varnothing \neq \Omega_\star \subset \mathbb{R}^{d_Z} \quad \text{is open and path-connected,} \tag{35a}$$

$$\Omega_{\star,-j}(u) := \{z_{-j} : (u, z_{-j}) \in \Omega_\star\} \quad \text{is connected whenever nonempty,} \qquad j = 1, \ldots, d_Z. \tag{35b}$$

Transition-configuration and null-set conventions. Let $P^{\vartheta}_{e,s}$ denote the stationary conditional process law at a covered pair (e, s) under parameterization ϑ, and set

$$\Xi_t := (H^Z_t, R_t, R_{t+1}), \qquad \nu^{\vartheta}_{e,s} := \mathcal{L}_{P^{\vartheta}_{e,s}}(\Xi_t), \qquad \mu^{\vartheta}_{e,s} := \mathcal{L}_{P^{\vartheta}_{e,s}}(Z_t),$$

$$\mathcal{Z}^{\vartheta}_{e,s} := \operatorname{supp} \mu^{\vartheta}_{e,s}, \qquad \mathcal{T}^{+}_{e,s}(\vartheta) := \operatorname{supp} \nu^{\vartheta}_{e,s}.$$

Regular conditional laws given $\Xi_t = \tau$ are understood $\nu^{\vartheta}_{e,s}$-almost everywhere on $\mathcal{T}^{+}_{e,s}(\vartheta)$. All support conditions above are imposed in each admissible parameterization ϑ. For every covered (e, s), the corresponding common contrast domain additionally satisfies $\mu^{\vartheta}_{e,s}(\Omega_\star) = 1$. The "a.e." qualifier in Eq. (34b) is with respect to d_Z-dimensional Lebesgue measure on $\Omega_\star$.

Dynamic indecomposability. For each covered s, define

$$\mathcal{E}_s(\vartheta) := \{e : (e,s) \in \mathcal{C}_{\mathrm{cov}}(\vartheta)\}, \qquad \mathcal{Z}_s^{\vartheta} := \bigcup_{e \in \mathcal{E}_s(\vartheta)} \mathcal{Z}_{e,s}^{\vartheta}.$$

For measurable maps $a : \mathcal{Z}_s^{\vartheta} \to \mathsf{A}$, where A is a standard Borel space, let

$$\mathcal{I}_s(\vartheta) := \left\{ a : \begin{array}{l} P_{e,s}^{\vartheta}(a(Z_t) = a(Z_{t+1}) \mid \Xi_t = \tau) = 1 \\ \text{for } \nu_{e,s}^{\vartheta}\text{-a.e. } \tau \in \mathcal{T}_{e,s}^{+}(\vartheta), \\ \text{for every } e \in \mathcal{E}_s(\vartheta) \end{array} \right\}. \tag{36}$$

Condition B requires

$$\mathcal{I}_s(\vartheta) = \left\{ a : \begin{array}{l} \exists c \in \mathsf{A}, \quad \mu_{e,s}^{\vartheta}(\{z : a(z) \neq c\}) = 0 \\ \text{for every } e \in \mathcal{E}_s(\vartheta) \end{array} \right\} \tag{37}$$

for every covered s and every $\vartheta \in [\theta]_X$.

C.4 Cross-context alignment and target conditions

Cross-context anchoring. For source and target coordinate-support projections $\mathcal{I} = (I_1, \ldots, I_{d_Z})$ and $\mathcal{J} = (J_1, \ldots, J_{d_Z})$, define the admissible support-restricted coordinate class

$$\begin{aligned} \mathcal{G}_{d_Z}(\mathcal{I}, \mathcal{J}) := \big\{ T : (T(z))_j = a_j(z_{\pi(j)}), \quad & j \in [d_Z], \\ & \pi \in \mathfrak{S}_{d_Z}, \\ & a_j \in \mathrm{Diff}^2(I_{\pi(j)}, J_j), \quad j \in [d_Z] \big\}. \end{aligned} \tag{38}$$

When $\mathcal{I} = \mathcal{J}$, this is a group; we suppress its support arguments when unambiguous. For a candidate chart family $T = (T_s)_{s \in \mathcal{S}_{\mathrm{cov}}^+}$, let $T_s^{(H)}$ apply T_s to every latent lag in H_t^Z. When applied to a history, T_s^{-1} abbreviates $(T_s^{(H)})^{-1}$. For $s \in \mathcal{S}_{\mathrm{cov}}^+$, let $N_s := |\{n \in [N] : S_n = s\}|$ and define the covered configuration law

$$Q_{s,T_s} := \frac{1}{N_s} \sum_{n : S_n = s} \left[\delta_{E_n} \otimes \mathcal{L}_{P_{E_n,s}^{\theta}}\big(R_t, T_s^{(H)} H_t^Z\big) \right].$$

For each pair (s, s'), define

$$\begin{aligned} \lambda_{ss'}^T &:= \tfrac{1}{2}\big(Q_{s,T_s} + Q_{s',T_{s'}}\big), \qquad q_{s|ss'}^T := \frac{dQ_{s,T_s}}{d\lambda_{ss'}^T}, \qquad q_{s'|ss'}^T := \frac{dQ_{s',T_{s'}}}{d\lambda_{ss'}^T}, \\ \mathcal{A}_{ss'}(T) &:= \Big\{ (e, r, \widetilde{h}) : q_{s|ss'}^T(e, r, \widetilde{h}) > 0,\ q_{s'|ss'}^T(e, r, \widetilde{h}) > 0 \Big\}, \\ \mathcal{E}_{\mathrm{ctx}}(T) &:= \big\{ \{s, s'\} : \lambda_{ss'}^T(\mathcal{A}_{ss'}(T)) > 0 \big\}, \\ \mathcal{G}_{\mathrm{ctx}}(T) &:= \big(\mathcal{S}_{\mathrm{cov}}^+, \mathcal{E}_{\mathrm{ctx}}(T)\big). \end{aligned} \tag{39}$$

We suppress the superscript T and the argument (T) when unambiguous. Equalities of transition kernels over $\mathcal{A}_{ss'}$ are understood $\lambda_{ss'}$-almost everywhere. Let $\overline{\mu}_Z := N^{-1} \sum_{n=1}^{N} \mathcal{L}_{\theta}(Z_{n,t} \mid E_n, S_n)$. The conclusion $T_s = T_{s'}$ a.e. means equality with respect to $\overline{\mu}_Z$ on the common covered dynamic support. Condition C requires these properties for every admissible chart family T arising from comparison with a parameterization in $[\theta]_X$: $\mathcal{G}_{\mathrm{ctx}}(T)$ is connected, and the implication in Eq. (10) holds on every edge.

Target minimality. Assumption 2 requires one unique inclusion-minimal coordinate set that renders the complementary dynamic coordinates conditionally independent of the target.

D Proofs of Hierarchical State Identification

Proof roadmap. Finite-window mixture matching and disintegration first recover a sequence-specific recurrent state quotient. Translated signature rows then separate environments and align these quotients globally; the lemmas below follow this order.

D.1 QUOTIENT OF RECURRENT STATES WITHIN A SEQUENCE

Lemma D.1 (Finite mixture component matching)**.** *Fix an observed sequence n and $W \in \mathbb{W}_{\mathrm{id}}$. Under Eqs.* (29a) *and* (29b)*, its window law identifies the positive path component laws and their weights up to one bijection of the two component index sets.*

Proof. Path injectivity makes every positive path a distinct member of $\mathfrak{P}_W$, while linear independence makes the coefficients in any two finite representations of the same mixture unique. The population regime identifies the positive-path mixture in Eq. (26). Fix $\vartheta \in [\theta]_X$ and write $(e, s) = (E_n, S_n)$ and $(e^\vartheta, s^\vartheta) = (E_n^\vartheta, S_n^\vartheta)$. The true and candidate mixtures are two finite representations of the same observed window law. Collect the distinct component measures from their union as $Q_1, \dots, Q_L \in \mathfrak{P}_W$. Subtracting the representations gives

$$\sum_{\ell=1}^{L} \alpha_\ell Q_\ell = 0,$$

where α_ℓ is the difference between the two weights assigned to Q_ℓ. Equation (29b) forces every α_ℓ to vanish. Hence the two representations contain the same component measures with the same weights, and there is a bijection $\tau : \mathcal{R}^+_{\theta,e,W} \to \mathcal{R}^+_{\vartheta,e^\vartheta,W}$ such that

$$P^{(W)}_{\theta,e,s,\boldsymbol{r}} = P^{(W)}_{\vartheta,e^\vartheta,s^\vartheta,\tau(\boldsymbol{r})}, \qquad \omega^{(W)}_{\theta,e,\boldsymbol{r}} = \omega^{(W)}_{\vartheta,e^\vartheta,\tau(\boldsymbol{r})}. \tag{40}$$

This is the finite mixture criterion of Proposition A.3 in Balsells-Rodas et al. (2024), applied to joint autoregressive path laws rather than independent observations. It identifies components and weights only; labels of recurrent states are recovered next. □

Lemma D.2 (Quotient of recurrent states from path components)**.** *Under path coverage, state-kernel separation, and switching minimality, the matched path components identify K_R and the state-conditional kernels up to one sequence-specific permutation.*

Proof. Fix an observed sequence n, $\vartheta \in [\theta]_X$, and the paired contexts used in Lemma D.1. Because the two admissible representations have the same observed window law, switching minimality in Eq. (32b) gives $K_R(\vartheta) = K_R(\theta) = K_R$; the candidate order was not fixed in advance. Equation (28) implies that, for every positive path $\boldsymbol{r}$, covered time $t \in \{2, \dots, W\}$, measurable B, and almost every common-support history,

$$P^{(W)}_{e,s,\boldsymbol{r}}\big(X_t \in B \mid X_{1:t-1} = x_{1:t-1}\big) = \mu_{e,r_t,s}\big(B \mid h^X_t(x_{1:t-1})\big)\,.$$

Thus disintegration recovers the one-step kernel at each covered path position. The same disintegration of the matched measure in Eq. (40) gives the same kernel. Using the reference measure and essential-overlap set in Eq. (31), declare two covered positions $(\boldsymbol{r}, t)$ and $(\boldsymbol{r}', t')$ equivalent when these kernels agree almost everywhere on that overlap. State-kernel separation gives

$$(\boldsymbol{r}, t) \sim (\boldsymbol{r}', t') \quad \Longleftrightarrow \quad r_t = r'_{t'}.$$

State coverage in Eq. (30a) ensures that every recurrent state occurs in at least one such position, so the true and candidate constructions have $K_R(\theta)$ and $K_R(\vartheta)$ quotient classes, respectively. Component matching and its inverse preserve equality of disintegrated kernels, so τ induces a bijection between these quotient sets. Since their orders have now been identified, this bijection can be written as one permutation $\pi_{e,s} \in \mathfrak{S}_{K_R}$, independent of the path and time position, satisfying

$$[\tau(\boldsymbol{r})]_t = \pi_{e,s}(r_t) \quad \text{for every positive } \boldsymbol{r} \text{ and } t \in \{2, \dots, W\}.$$ □

Lemma D.3 (Transition and posterior laws)**.** *Given the quotient in Lemma D.2, the matched path weights identify $A^{(e)}$ and every finite-window posterior law of the recurrent state process, up to the same permutation.*

Proof. Fix $W \in \mathbb{W}_{\mathrm{id}}$ and, for each covered transition under consideration, an interior position $t \in \{2, \dots, W-1\}$ supplied by the positive-transition coverage in Eq. (30b). Both endpoints then

have recovered quotient labels. The identified weights determine the state marginals and adjacent-state probabilities:

$$P(R_t = r \mid E = e) = \sum_{\substack{\boldsymbol{r}' \in \mathcal{R}^+_{e,W}: \\ r'_t = r}} \omega^{(W)}_{e,\boldsymbol{r}'},$$

$$P(R_t = r, R_{t+1} = r' \mid E = e) = \sum_{\substack{\boldsymbol{q} \in \mathcal{R}^+_{e,W}: \\ (q_t, q_{t+1}) = (r, r')}} \omega^{(W)}_{e,\boldsymbol{q}}.$$

Positive occupancy in Eq. (32a) makes the denominator below nonzero, hence

$$A^{(e)}_{r,r'} = \frac{P(R_t = r, R_{t+1} = r' \mid E = e)}{P(R_t = r \mid E = e)}.$$

An adjacent pair absent from every positive path has probability zero and therefore identifies a zero transition.

For the posterior claim, let $P^{(W)}_{e,s}$ dominate every identified component measure. Bayes' formula for finite mixtures gives, for $P^{(W)}_{e,s}$-almost every window $x_{1:W}$,

$$P(R_{1:W} = \boldsymbol{r} \mid X_{1:W} = x_{1:W}, e, s) = \omega^{(W)}_{e,\boldsymbol{r}} \frac{dP^{(W)}_{e,s,\boldsymbol{r}}}{dP^{(W)}_{e,s}}(x_{1:W}).$$

Both the weight and Radon–Nikodym derivative are identified, so the posterior over component paths is identified up to τ. Applying the quotient map from Lemma D.2 to positions $2{:}W$ and aggregating component paths with the same interior labels identifies the posterior of $R_{2:W}$ up to $\pi_{e,s}$. In a finite-dimensional identifiable switching family, the recovered initial law, transition matrix, and conditional-kernel parameters determine every finite-window posterior by Bayes recursion. In the nonparametric formulation, for any target block length L, choose an identifying window $W \geq L+2$. Marginalizing latent and observed coordinates outside an interior length-L block yields its joint law; stationarity transfers this law to $(R_{1:L}, X_{1:L})$, and Bayes' formula gives the desired posterior. The quotient classes align their interior labels across identifying windows, and the recovered laws are projectively consistent because each is a restriction of the identified window-law family. They therefore determine the population posterior process law. This does not make a realized finite state path observable unless its posterior is degenerate.

□

Lemma D.4 (Identification of recurrent states within a sequence)**.** *Under Condition A, each sequence-specific window law identifies the quotient of recurrent states, K_R, $A^{(E_n)}$, $\{\mu_{E_n,r,S_n}\}_{r=1}^{K_R}$, and the posterior process law up to one sequence-specific permutation. In particular, $K_R(\vartheta) = K_R(\theta) = K_R$ for every $\vartheta \in [\theta]_X$.*

Proof. Apply Lemmas D.1– D.3; their common quotient permutation gives all claims simultaneously. □

D.2 Global–local hierarchy separation

Lemma D.5 (Rigidity of finite translations)**.** *If a finite nonempty set $\mathcal{A} \subset \mathbb{R}^d$ satisfies $\mathcal{A} = \mathcal{A} + d$, then $d = 0$.*

Proof. Choose $a \in \mathcal{A}$. Repeated use of the set equality gives $a + md \in \mathcal{A}$ for every integer $m \geq 0$. If $d \neq 0$, these points are all distinct, contradicting finiteness. Hence $d = 0$. □

Lemma D.6 (Environment partition and cardinality)**.** *Under Assumption 1 and observed-sequence environment coverage, equality of recovered signature rows identifies the environment partition and K_E; in particular, $K_E(\vartheta) = K_E(\theta) = K_E$ for every $\vartheta \in [\theta]_X$.*

Proof. Let $\mathcal{W} = \{w_1, \dots, w_{K_R}\}$. By Lemma D.4, sequence n identifies its unordered kernel set and hence the signature row $\mathcal{Q}_n = b_{E_n} + \mathcal{W}$. If the rows for e and e' agree, then

$$\mathcal{W} = \mathcal{W} + (b_{e'} - b_e).$$

Lemma D.5 forces $b_e = b_{e'}$, and injectivity of $e \mapsto b_e$ then gives $e = e'$. The converse is immediate. Row equality therefore identifies the environment partition. Coverage ensures that every environment contributes a row, so the number of distinct rows is K_E. The same observable row partition is recovered under every $\vartheta \in [\theta]_X$; candidate-wise coverage makes its number of classes $K_E(\vartheta)$, proving $K_E(\vartheta) = K_E$.

□

Lemma D.7 (Global alignment of recurrent states)**.** *The recovered signature rows determine one correspondence between recurrent states across all covered environments, up to the arbitrary naming of one reference row.*

Proof. Lemmas D.4 and D.6 first establish common recurrent and environment cardinalities across observationally equivalent parameterizations, so the correspondences below are between equal finite label sets. Fix a reference environment e_0. The translation

$$t_e := b_{e_0} - b_e \quad \text{satisfies} \quad \mathcal{Q}_e + t_e = b_{e_0} + \mathcal{W}.$$

It is unique. Indeed, if t_e and t'_e both map $\mathcal{Q}_e$ to the reference row, that row is invariant under $t_e - t'_e$; by Lemma D.5, $t_e = t'_e$. After applying the unique translation, every row contains the same points $b_{e_0} + w_r$. Injectivity of $r \mapsto w_r$ gives a unique bijection from the recurrent state classes in row e to those in the reference row. Fixing arbitrary names for the reference classes therefore induces one global permutation π_R.

Because the correspondence is defined on the quotient classes themselves, it transfers every object attached to those classes. Thus, after an arbitrary environment permutation π_E,

$$\begin{aligned} \widetilde{A}^{(\pi_E(e))}_{\pi_R(r),\pi_R(r')} &= A^{(e)}_{r,r'}, \\ \widetilde{\mu}_{\pi_E(e),\pi_R(r),\widetilde{s}} &= \mu_{e,r,s}. \end{aligned} \tag{41}$$

on common histories, and the posterior finite-dimensional laws agree under the same relabeling. No additional, path-specific permutation is available.

□

Proof of Theorem 1. Lemma D.4 identifies K_R and the within-sequence objects; Lemma D.6 identifies K_E, and Lemma D.7 aligns the states globally. Only the names of environment classes and reference row states remain arbitrary, giving Eq. (22). □

E Proofs of Coordinate and Target Identification

Proof roadmap. Starting from globally aligned recurrent states, the proof cancels the observation Jacobian, excludes cross-coordinate mixing, and obtains within-context componentwise recovery. Cross-context anchoring then propagates one chart, after which target minimality transfers under the same equivalence.

E.1 Static context preservation

Lemma E.1 (Static context preservation)**.** *Suppose two admissible parameterizations are X-equivalent and satisfy Condition B. Then there is a bimeasurable bijection κ on the covered context set, up to null sets, such that $\widetilde{S} = \kappa(S)$ almost surely. For every covered s, the dynamic representations are related by a well-defined context-specific diffeomorphism.*

Proof. Let $\vartheta \in [\theta]_X$ denote the untilded parameterization. On the common noise-free observation support, draw X from the shared observed law and set $(Z, S) = F^{-1}(X)$ and $(\widetilde{Z}, \widetilde{S}) = \widetilde{F}^{-1}(X)$. This common-$X$ coupling defines the bimeasurable latent isomorphism

$$H := \widetilde{F}^{-1} \circ F, \qquad H(z, s) = \big(h_Z(z, s), h_S(z, s)\big).$$

Fix a covered $(e, s) \in \mathcal{C}_{\mathrm{cov}}(\vartheta)$. In the equivalent model, $\widetilde{S}$ is a single variable fixed throughout the sequence, not a time-indexed state. Therefore the coupled two-time law satisfies

$$P^{\vartheta}_{e,s}(h_S(Z_t, s) = h_S(Z_{t+1}, s) \mid \Xi_t = \tau) = 1 \quad \text{for } \nu^{\vartheta}_{e,s}\text{-a.e. } \tau \in \mathcal{T}^{+}_{e,s}(\vartheta).$$

Since this holds for every $e \in \mathcal{E}_s(\vartheta)$, $a_s(z) := h_S(z, s)$ belongs to $\mathcal{I}_s(\vartheta)$ in Eq. (36). Condition B makes a_s almost surely constant, with one common value across these environments. The conditional law of $\widetilde{S}$ given $S = s$ is therefore a point mass, defining a measurable map κ such that

$$h_S(z, s) = \kappa(s) \quad \text{for almost every } z, \qquad \widetilde{S} = \kappa(S) \quad \text{almost surely.}$$

Applying the same argument to the inverse isomorphism $H^{-1} = F^{-1} \circ \widetilde{F}$ gives a measurable λ with $S = \lambda(\widetilde{S})$ almost surely. Thus $\lambda \circ \kappa$ and $\kappa \circ \lambda$ are the identity outside null sets on the covered context sets, so κ is bimeasurable there. The first component of H now defines

$$u_s(z) := \widetilde{f}^{-1}_{\kappa(s)}(f_s(z)), \qquad v_s := u_s^{-1} = f_s^{-1} \circ \widetilde{f}_{\kappa(s)}.$$

The slice-wise C^2 embedding clause of Condition B makes these maps well-defined C^2 diffeomorphisms between the common latent supports. □

E.2 Observable density ratio contrasts

Lemma E.2 (Within-sequence observable density ratio contrasts)**.** *Fix a sequence n, two identified recurrent states $r_q, r_{0,q}$, and a history h^X_q on which their conditional kernels have common support. The Radon–Nikodym contrast*

$$\Lambda_{n,q}(x \mid h^X_q) = \log \frac{d\mu_{E_n, r_q, S_n}(\cdot \mid h^X_q)}{d\mu_{E_n, r_{0,q}, S_n}(\cdot \mid h^X_q)}(x) \tag{42}$$

is identified without observing S_n and satisfies Eq. (6)*.*

Proof. Because the two states occur in the same sequence, they share $s = S_n$ and f_s. Set $z = f_s^{-1}(x)$ and transform the lag history by the same inverse. For either $r \in \{r_q, r_{0,q}\}$, change of variables on $\mathcal{M}_s$ gives

$$p_X(x \mid h^X_q, E_n, r, s) = p_Z(z \mid h^Z_q, E_n, r)\{J^{\mathrm{vol}}_{f_s}(z)\}^{-1}.$$

Hence, on common support,

$$\begin{aligned} \frac{p_X(x \mid h^X_q, E_n, r_q, s)}{p_X(x \mid h^X_q, E_n, r_{0,q}, s)} &= \frac{p_Z(z \mid h^Z_q, E_n, r_q)\{J^{\mathrm{vol}}_{f_s}(z)\}^{-1}}{p_Z(z \mid h^Z_q, E_n, r_{0,q})\{J^{\mathrm{vol}}_{f_s}(z)\}^{-1}} \\ &= \frac{p_Z(z \mid h^Z_q, E_n, r_q)}{p_Z(z \mid h^Z_q, E_n, r_{0,q})}. \end{aligned} \tag{43}$$

Taking logarithms gives Eq. (6); Lemma D.4 identifies both observed kernels and hence their Radon–Nikodym contrast in Eq. (42). □

E.3 Within-context coordinate recovery

Lemma E.3 (Contrast functional equation)**.** *After the state labels and static contexts have been aligned, every identified contrast q satisfies, on its common support,*

$$\sum_{i=1}^{d_Z} \rho_{qi}\big(v_{s,i}(\widetilde{z}) \mid h^Z_q\big) = \sum_{i=1}^{d_Z} \widetilde{\rho}_{qi}\big(\widetilde{z}_i \mid \widetilde{h}^Z_q\big). \tag{44}$$

Proof. Use the common state indexing from Theorem 1, fix a covered s, and set $z = v_s(\widetilde{z})$ and $x = f_s(z) = \widetilde{f}_{\kappa(s)}(\widetilde{z})$, with corresponding lag histories. Lemma E.2 and the contrast additivity in

Eq. (34a) give

$$\begin{aligned}\sum_{i=1}^{d_Z} \rho_{qi}\big(v_{s,i}(\widetilde{z}) \mid h_q^Z\big) &= \rho_q\big(v_s(\widetilde{z}) \mid h_q^Z\big) \\ &= \Lambda_{n,q}(x \mid h_q^X) \\ &= \widetilde{\rho}_q\big(\widetilde{z} \mid \widetilde{h}_q^Z\big) \\ &= \sum_{i=1}^{d_Z} \widetilde{\rho}_{qi}\big(\widetilde{z}_i \mid \widetilde{h}_q^Z\big).\end{aligned}$$

The two endpoints give Eq. (44). □

Lemma E.4 (Mixed derivatives vanish)**.** *Under the full-rank clause in Eq.* (34b)*, for every i and every pair $k \neq \ell$,*

$$\partial_{k\ell} v_{s,i} = 0, \qquad \partial_k v_{s,i}\, \partial_\ell v_{s,i} = 0 \tag{45}$$

throughout the open common comparison support.

Proof. Fix $k \neq \ell$. Applying the (k, ℓ) mixed derivative to both sides of Eq. (44) gives the following chain for each contrast q. The first equality uses coordinate additivity on the right-hand side, and the last equality is the chain rule:

$$\begin{aligned}0 &= \partial_{k\ell} \sum_{i=1}^{d_Z} \widetilde{\rho}_{qi}\big(\widetilde{z}_i \mid \widetilde{h}_q^Z\big) \\ &= \partial_{k\ell} \sum_{i=1}^{d_Z} \rho_{qi}\big(v_{s,i}(\widetilde{z}) \mid h_q^Z\big) \\ &= \sum_{i=1}^{d_Z} \partial_{k\ell} \rho_{qi}\big(v_{s,i}(\widetilde{z}) \mid h_q^Z\big) \\ &= \sum_{i=1}^{d_Z} \left[\rho'_{qi}(v_{s,i}) \partial_{k\ell} v_{s,i} + \rho''_{qi}(v_{s,i}) \partial_k v_{s,i} \partial_\ell v_{s,i}\right].\end{aligned}$$

Stacking over $q = 1, \ldots, m$ yields

$$M_\rho(v_s(\widetilde{z}), h^Z_{1:m}) c_{k\ell}(\widetilde{z}) = 0, \qquad c_{k\ell} := \Big((\partial_{k\ell} v_{s,i})_{i=1}^{d_Z}, (\partial_k v_{s,i} \partial_\ell v_{s,i})_{i=1}^{d_Z}\Big)^\top.$$

Equation (34b) gives a trivial null space almost everywhere in z. Because v_s^{-1} is locally Lipschitz, the preimage under v_s of the rank-deficient null set is also null. Hence $c_{k\ell} = 0$ almost everywhere in $\widetilde{z}$. Its entries are continuous on the open support, so they vanish throughout that support, proving Eq. (45).

□

Thus, each component of v_s depends on at most one coordinate of $\widetilde{z}$. Invertibility and connected support next show that the surviving coordinate assignments form one fixed permutation.

Lemma E.5 (Componentwise rigidity)**.** *At every point of the common comparison support, Jv_s has exactly one nonzero entry in each row and column. Under Eqs.* (35a) *and* (35b)*, these locations define one permutation $\pi_s \in \mathfrak{S}_{d_Z}$ over the support, and there are scalar C^2 diffeomorphisms $a_{s,i}$ such that*

$$v_{s,i}(\widetilde{z}) = a_{s,i}\big(\widetilde{z}_{\pi_s(i)}\big), \qquad i = 1, \ldots, d_Z. \tag{46}$$

Proof. For fixed $\widetilde{z}$, Eq. (45) gives

$$k \neq \ell \quad \Longrightarrow \quad \partial_k v_{s,i}(\widetilde{z})\, \partial_\ell v_{s,i}(\widetilde{z}) = 0 \quad \Longrightarrow \quad \big|\{k : \partial_k v_{s,i}(\widetilde{z}) \neq 0\}\big| \leq 1.$$

Since v_s is a diffeomorphism,

$$\det Jv_s(\widetilde{z}) \neq 0 \quad \Longrightarrow \quad [Jv_s(\widetilde{z})]_{ik} = d_{s,i}(\widetilde{z}) \mathbf{1}\{k = \pi_s(i, \widetilde{z})\}, \qquad d_{s,i}(\widetilde{z}) \neq 0,$$

where $\pi_s(\cdot, \widetilde{z})$ is a permutation. Continuity and the global connectedness in Eq. (35a) yield

$$\pi_s(\cdot, \widetilde{z}) \text{ locally constant}, \qquad \Omega_\star \text{ connected} \quad \Longrightarrow \quad \pi_s(\cdot, \widetilde{z}) \equiv \pi_s.$$

For each i, the connected-slice clause of Eq. (35b) gives

$$\partial_k v_{s,i} = 0 \quad (k \neq \pi_s(i)) \quad \Longrightarrow \quad v_{s,i}(\widetilde{z}) = a_{s,i}\big(\widetilde{z}_{\pi_s(i)}\big).$$

Finally, the active-coordinate projection of $\Omega_\star$ is an interval and

$$a'_{s,i}\big(\widetilde{z}_{\pi_s(i)}\big) = d_{s,i}(\widetilde{z}) \neq 0.$$

Thus $a_{s,i}$ is strictly monotone, and the inverse function theorem makes it a C^2 scalar diffeomorphism. This proves Eq. (46). □

Proposition E.6 (Within-context coordinate identification)**.** *Under Condition B, every covered context identifies Z up to one context-specific coordinate permutation and componentwise scalar diffeomorphisms.*

Proof. Combine Lemmas E.1 and E.3 with Lemmas E.4–E.5. The common contrast domain has full probability, as does its image under u_s. Thus the coordinate equivalence holds almost surely in each covered context. □

E.4 Cross-context alignment

Lemma E.7 (Global alignment over the context graph)**.** *Under Condition C, the context-specific maps in Proposition E.6 agree almost everywhere over the connected graph of covered contexts.*

Proof. Let $T_s = v_s^{-1}$. The dynamic transition family in either admissible model conditions on (e, r, h) but not on the static context. Therefore, whenever (s, s') is an edge of $\mathcal{G}_{\mathrm{ctx}}$, the shared transition configurations give

$$(T_s)_{\#} p_Z\big(\cdot \mid T_s^{-1}(\widetilde{h}), e, r\big) = (T_{s'})_{\#} p_Z\big(\cdot \mid T_{s'}^{-1}(\widetilde{h}), e, r\big), \quad \lambda_{ss'}\text{-a.e. } (e, r, \widetilde{h}) \in \mathcal{A}_{ss'}. \tag{47}$$

Equation (47) is the premise of Eq. (10); its edgewise uniqueness clause gives $T_s = T_{s'}$ almost everywhere.

For arbitrary covered contexts s and s', connectivity supplies a path $s = s_0, s_1, \ldots, s_L = s'$ in $\mathcal{G}_{\mathrm{ctx}}$. Applying edgewise equality along the path gives $T_{s_0} = T_{s_1} = \cdots = T_{s_L}$. Hence one coordinate permutation and one collection of scalar transformations apply over the entire connected graph. If the graph is disconnected, the same conclusion holds only within each connected component. □

Proposition E.8 (A sufficient anchoring criterion)**.** *For an overlap edge (s, s'), let $\mathrm{Stab}_{ss'}$ be the set of $U \in \mathcal{G}_{d_Z}$ satisfying*

$$U_{\#} p_Z(\cdot \mid h, e, r) = p_Z(\cdot \mid Uh, e, r).$$

This equality is required for $\lambda_{ss'}$-almost every shared configuration for which both sides are defined. The same U is applied to every latent lag collected in h. If $\mathrm{Stab}_{ss'} = \{\mathrm{id}\}$ on every edge and $\mathcal{G}_{\mathrm{ctx}}$ is connected, then Condition C holds.

Proof. Suppose the two pushforward families in Eq. (10) agree and set $U = T_{s'}^{-1} \circ T_s$. For a history h, take $\widetilde{h} = T_s h$, so that $T_{s'}^{-1}(\widetilde{h}) = Uh$. Applying $T_{s'}^{-1}$ to the common pushforward law gives

$$\begin{aligned} U_{\#} p_Z(\cdot \mid h, e, r) &= (T_{s'}^{-1})_{\#} (T_s)_{\#} p_Z(\cdot \mid h, e, r) \\ &= (T_{s'}^{-1})_{\#} (T_{s'})_{\#} p_Z(\cdot \mid Uh, e, r) \\ &= p_Z(\cdot \mid Uh, e, r). \end{aligned}$$

This is precisely the defining relation for $U \in \mathrm{Stab}_{ss'}$. Triviality of the stabilizer gives $U = \mathrm{id}$, hence $T_s = T_{s'}$. Connectivity then propagates the equality as in Lemma E.7. □

A practical contrast diagnostic. Under factorized transitions, one can screen for obvious stabilizers by checking whether one-dimensional contrast vectors for distinct coordinates are related by scalar diffeomorphisms and whether each coordinate has an injective shared contrast. Failure of the first check suggests a coordinate-permutation symmetry; failure of the second leaves a scalar reparameterization unresolved. Passing these checks is diagnostic evidence, not a proof that the full shared transition family has a trivial stabilizer.

E.5 Target-minimality transfer

Lemma E.9 (Preservation of the minimal target set)**.** *Among parameterizations in $[\theta]_{XY}$, under a global coordinate permutation and componentwise bijections, the family of inclusion-minimal target-sufficient coordinate sets is relabeled by the same permutation. Under Assumption 2, the unique set is therefore J_C up to that permutation.*

Proof. Coordinatewise bijections preserve the sigma-field generated by each coordinate, while a permutation only relabels these sigma-fields. Thus, for every $J \subseteq [d_Z]$,

$$\sigma(Z_J) = \sigma(\widetilde{Z}_{\pi^{-1}(J)}), \qquad \sigma(Z_{-J}) = \sigma(\widetilde{Z}_{-\pi^{-1}(J)}).$$

Conditional independence depends only on the conditioning and conditioned sigma-fields, so

$$Y \perp\!\!\!\perp Z_{-J} \mid Z_J \quad \Longleftrightarrow \quad Y \perp\!\!\!\perp \widetilde{Z}_{-\pi^{-1}(J)} \mid \widetilde{Z}_{\pi^{-1}(J)}.$$

Thus $J \mapsto \pi^{-1}(J)$ is an inclusion-preserving bijection between the two families of sufficient coordinate sets and therefore between their inclusion-minimal elements. Assumption 2 then maps the unique element J_C to its relabeled counterpart. □

Proof of Theorem 2. Theorem 1 identifies the state hierarchy up to environment and recurrent state permutations, and Lemma E.1 gives $\widetilde{S} = \kappa(S)$ on the covered set. Proposition E.6 then gives Eq. (46). Under Condition C, Lemma E.7 globalizes these maps and proves Eq. (24); without it, the context-specific result is sharp. All conclusions are restricted to the covered context graph and the equivalence class in Eqs. (46)–(24). □

Proof of the target-relevant-subset corollary. Lemma E.9, Assumption 2, and the global equivalence class from Theorem 2 identify J_C up to the same coordinate permutation. The causal interpretation of C requires the additional target mechanism assumption stated in Section 3.2. □

F Identification Boundaries and Counterexamples

Table 3 provides a roadmap of the identification boundaries: for each violated condition, it records the guarantee that is lost and the weaker recovery object that remains. The following subsections give corollaries and counterexamples establishing these boundaries.

F.1 State hierarchy boundaries

Corollary 2 (Incomplete recurrent state support)**.** *If ergodicity or positive occupancy of every recurrent state is removed, the largest identified state cardinality within a sequence is only a lower bound on K_R. Recurrent state labels align within connected components of the row translation graph; disconnected components retain separate permutations, and K_E need not be identifiable.*

Proof. For sequence n, let

$$\mathcal{R}_n^+ := \{r : P(R_{n,t} = r) > 0 \text{ for some covered } t\}, \qquad K_n^{\text{obs}} := |\mathcal{R}_n^+|.$$

A state outside $\mathcal{R}_n^+$ contributes no positive path component to Eq. (26). It therefore leaves neither a conditional kernel nor a transition weight. Arbitrarily many such states can therefore be added without changing the window law, so

$$\max_n K_n^{\text{obs}} \le K_R, \tag{48}$$

Table 3: Identification boundaries. When each condition is violated, we show which recovery guarantee is lost and what remains identifiable or becomes ambiguous.

Condition violated	**Guarantee lost**	**Identifiable remainder / ambiguity**
	Panel A: State hierarchy.	
Switching minimality	Recurrent-state decomposition and K_R	Observational state classes only
State coverage	Exact recurrent-state cardinality K_R	Observed-state lower bound $\max_n K_n^{\mathrm{obs}} \leq K_R$
Signature separation	Environment partition and global R alignment	Unseparated global–local state index
	Panel B: Coordinate recovery.	
Structured innovations	Componentwise coordinate recovery	General within-context diffeomorphic chart
Dynamic indecomposability	Separation of static and dynamic factors	Joint static–dynamic invariant σ-field
Full contrast rank, $\mathrm{rank}(M_\rho) = 2d_Z$	Unique componentwise coordinate basis	Rotational ambiguity (Counterexample 4)
Connected dynamic support	One chart over the covered support	Separate chart for each support component
Cross-context anchoring	One globally aligned coordinate chart	Context-specific coordinate charts
	Panel C: Target and observation.	
Unique minimal target sufficiency	Unique target-relevant subset J_C	Multiple inclusion-minimal sufficient subsets
Joint observation injectivity	Pointwise latent state (Z, S)	Observation fibers only

with equality unprovable without coverage. In the row translation graph on $\mathcal{Q}_n^+ = \{\Delta_{n,r} : r \in \mathcal{R}_n^+\}$, unique edgewise matches compose only within connected components. Disconnected components admit independent permutations, while unobserved row completions can change the environment partition; hence K_E need not be identified. □

Counterexample 1 (Global–local signature aliasing)**.** Let $E, R \in \{0, 1\}$ and suppose the only observable transition contrast is $\Delta = E \,\mathsf{xor}\, R$. The two environment-indexed rows coincide:

$$\mathcal{Q}_0 = \{0 \,\mathsf{xor}\, 0, 0 \,\mathsf{xor}\, 1\} = \{0, 1\}, \qquad \mathcal{Q}_1 = \{1 \,\mathsf{xor}\, 0, 1 \,\mathsf{xor}\, 1\} = \{0, 1\}.$$

The correspondence is $R \mapsto R$ in environment 0 but $R \mapsto 1 - R$ in environment 1. When the local chain and its conditional kernels are invariant under this swap, replacing (E, R) by the single flat state $U = E \,\mathsf{xor}\, R$ preserves every state-conditional window law and mixture weight. The factorization of U into (E, R) is therefore not observable, so neither the environment partition nor a single global alignment of R is identifiable.

F.2 Coordinate-recovery boundaries

Counterexample 2 (Nonadditive transition innovations)**.** Remove the innovation-factorization clause of Condition B and let $B \in \mathbb{R}^{d_Z \times d_Z}$ be any invertible non-monomial matrix. Define

$$\widetilde{Z}_t = BZ_t, \qquad \widetilde{f}_s(\widetilde{z}) = f_s(B^{-1}\widetilde{z}),$$

and transform every transition kernel by

$$p_{\widetilde{Z}}(\widetilde{z}_t \mid \widetilde{h}, e, r) = |\det B^{-1}|\, p_Z(B^{-1}\widetilde{z}_t \mid B^{-1}\widetilde{h}, e, r).$$

Then $\widetilde{f}_s(\widetilde{Z}_t) = f_s(Z_t)$ and all observed transition laws are unchanged. Without the component-factorized innovation restriction, the transformed joint kernels remain admissible for arbitrary B. Hence the observed law generally identifies only a mixed within-context coordinate system, not componentwise coordinates.

Counterexample 3 (A nonconstant dynamic invariant)**.** Let S, U, and the dynamic process W_t be mutually independent, with $S, U \sim \mathcal{N}(0,1)$, and set $Z_t = (U, W_t)$ so that U is constant over time. Define the orthogonal reparameterization

$$\widetilde{S} = \frac{S+U}{\sqrt{2}}, \qquad \widetilde{U} = \frac{S-U}{\sqrt{2}}, \qquad \widetilde{Z}_t = (\widetilde{U}, W_t).$$

The variables $\widetilde{S}$ and $\widetilde{U}$ are again independent standard Gaussians, and $\widetilde{U}$ is still time-invariant. Composing the observation map with the inverse orthogonal transformation preserves the observed process and the required static–dynamic factorization. However, $\widetilde{S}$ is not a function of S alone. Thus, when Condition B fails, a dynamic invariant that persists throughout the sequence can be reallocated between S and Z; only their joint invariant sigma-field is identified.

Counterexample 4 (Rank-deficient variation)**.** Let $Z_t = (Z_{t,1}, Z_{t,2}, Z_{t,3})$, let $(Z_{t,2}, Z_{t,3})$ be independent standard-Gaussian white-noise coordinates, and let every environment change only the transition law of $Z_{t,1}$. For any non-permutation $B \in \mathrm{O}(2)$, set $\widetilde{Z}_{t,1} = Z_{t,1}$, $\widetilde{Z}_{t,2:3} = BZ_{t,2:3}$, and $\widetilde{f}(\widetilde{z}) = f(z_1, B^{-1}\widetilde{z}_{2:3})$. Then $\widetilde{f}(\widetilde{Z}_t) = f(Z_t)$, and all contrast ratios are preserved because the corresponding columns of M_ρ vanish. Thus the internal coordinate basis of this Gaussian block is not identifiable, although the displayed rotations preserve its span.

Counterexample 5 (Disconnected dynamic support)**.** Fix one context and let $d_Z = 2$ with $\Omega_\star = \Omega_1 \dot\cup \Omega_2$, where $\Omega_1 = B((-2,-2), 1/4)$ and $\Omega_2 = B((2,2), 1/4)$. Suppose $P(Z_{t+1} \in \Omega_{\ell'} \mid Z_t \in \Omega_\ell, e, r) = 0$ for $\ell' \neq \ell$. Let T be the identity on Ω_1 and swap the coordinates on Ω_2, with $\widetilde{Z} = T(Z)$ and $\widetilde{f} = f \circ T^{-1}$. Transforming the two within-component kernels preserves the observed law, but the recovered permutation differs across support components.

F.3 Cross-context and target boundaries

Counterexample 6 (Unanchored static contexts)**.** Let $S \in \{0,1\}$ and suppose every conditional transition law of a scalar coordinate Z_t is invariant under $z \mapsto -z$. Define $\widetilde{Z}_t = (-1)^S Z_t$ and $\widetilde{f}_s(\widetilde{z}) = f_s((-1)^s \widetilde{z})$. Then $\widetilde{f}_S(\widetilde{Z}_t) = f_S(Z_t)$, while sign invariance preserves every transition configuration under the same history transformation. The required coordinate map is the identity for $s = 0$ and sign reversal for $s = 1$, so no single map covers both contexts. This is the ambiguity excluded by Condition C.

Counterexample 7 (Illustrative nonunique target-minimal set)**.** At the target selection stage, let $Z = (Z_1, Z_2, V) = (U, b(U), V)$ for an invertible b, and let $Y = g(U, \varepsilon_Y)$ with $\varepsilon_Y \perp\!\!\!\perp (U, V)$. Then $Y \perp\!\!\!\perp (Z_2, V) \mid Z_1$ and $Y \perp\!\!\!\perp (Z_1, V) \mid Z_2$, so both $\{1\}$ and $\{2\}$ are inclusion-minimal sufficient coordinate sets; the target law cannot determine which redundant coordinate is designated as C. This construction has singular support because $Z_2 = b(Z_1)$. It is therefore not a single-condition sharpness example under the open-support clause of Condition B. It isolates the target law ambiguity that Assumption 2 excludes.

G HiDICE Estimation and Shared Evaluation Protocol

This appendix specifies HiDICE's finite-sample inference procedure, complete training objective, model-selection rule, and shared evaluation metrics.

G.1 Optimization and inference

Encoder parameterization. Let $U_{n,t} = g_\phi(\boldsymbol{X}_n)_t$ and $\overline{U}_n = T_n^{-1} \sum_t U_{n,t}$, excluding padded positions. The environment posterior and local-state emission logits are

$$q_\phi(E_n = e \mid \boldsymbol{X}_n) = \mathrm{softmax}\big(a_E(\overline{U}_n)\big)_e\,,$$
$$\ell^{(e)}_{n,t,r} = \Big[a_R\Big(U_{n,t} - \overline{U}_n, \widetilde{b}_e\Big)\Big]_r\,.$$

Here $\widetilde{b}_e$ is the learned environment embedding; through $\ell^{(e)}_{n,t,r}$, it parameterizes the chain posterior in Eq. (12). The transition matrix $A^{(e)}$ is row-stochastic; a nonnegative diagonal bias encourages persistence while all off-diagonal transitions remain learnable.

Fitted signature extraction. For each (e, r), let $\widehat{\boldsymbol{m}}_{e,r}$ and $\widehat{\boldsymbol{\ell}}_{e,r}$ denote the conditional-mean and log-marginal-variance feature vectors extracted from the fitted state-conditional transition distributions. The fitted signature concatenates these statistics:

$$\widehat{\Delta}_{e,r} := \operatorname{concat}\left(\widehat{\boldsymbol{m}}_{e,r}, \widehat{\boldsymbol{\ell}}_{e,r}\right). \tag{49}$$

This same extraction rule supplies the signatures in $\mathcal{L}_{\mathrm{sig}}$ and in the hierarchy residual $\widehat{\mathcal{R}}_\Gamma$ below.

Auxiliary hierarchy and factor regularizers. Stacking the branch-specific component maps in Eq. (13) into $\overline{B}^\beta$ and $\overline{L}_R^\beta$ gives

$$\mathcal{L}_{\mathrm{sep}} = \sum_{\beta \in \mathcal{B}} \left\| \overline{B}^{\beta\top} \overline{L}_R^\beta \right\|_F^2.$$

An auxiliary inverse a_ψ and mini-batch cross-covariances define

$$\mathcal{L}_{\mathrm{cyc}} = \mathbb{E}_{p_\theta(X \mid Z^{\mathrm{cand}}, Z^{\mathrm{resid}}, S)} \left\| a_\psi(X) - (Z^{\mathrm{cand}}, Z^{\mathrm{resid}}, S) \right\|_2^2,$$
$$\mathcal{L}_{\mathrm{fac}} = \| \operatorname{Cov}(Z^{\mathrm{cand}}, Z^{\mathrm{resid}}) \|_F^2 + \| \operatorname{Cov}(Z^{\mathrm{cand}}, S) \|_F^2 + \| \operatorname{Cov}(Z^{\mathrm{resid}}, S) \|_F^2.$$

Covariances are computed from standardized mini-batch samples. $\mathcal{L}_{\mathrm{cyc}}$ discourages many-to-one decoding, whereas $\mathcal{L}_{\mathrm{fac}}$ penalizes cross-branch posterior duplication on sampled support.

Let $\widehat{\pi}_e = N^{-1} \sum_n q_\phi(E_n = e \mid \boldsymbol{X}_n)$ and let $\widehat{\rho}_{e,r}$ be the posterior-weighted occupancy of mechanism state r under environment e. For a fixed floor $\delta > 0$, the occupancy and optional sparse-change penalties are

$$\mathcal{L}_{\mathrm{occ}} = \sum_e [\delta - \widehat{\pi}_e]_+^2 + \sum_{e,r} [\delta - \widehat{\rho}_{e,r}]_+^2,$$
$$\mathcal{L}_{\mathrm{sp}} = \sum_{\beta \in \mathcal{B}} \sum_{e,r,j} \left\| \eta_{e,r,j}^\beta - \eta_{0,j}^\beta \right\|_2.$$

For observed target indices $\mathcal{I}_Y := \{(n,t) : t \in \mathcal{O}_n\}$, define

$$\gamma_{n,t}^{e,r} := q_\phi(E_n = e, R_{n,t} = r \mid \boldsymbol{X}_n), \qquad \ell_{n,t}^Y := -\log p_\theta(Y_{n,t} \mid Z_{n,t}^{\mathrm{sel}}),$$
$$M_{e,r} := \sum_{(n,t) \in \mathcal{I}_Y} \gamma_{n,t}^{e,r}, \qquad \overline{\ell}_{e,r} := M_{e,r}^{-1} \sum_{(n,t) \in \mathcal{I}_Y} \gamma_{n,t}^{e,r} \ell_{n,t}^Y,$$
$$\omega_{e,r} := \frac{M_{e,r}}{\sum_{e',r'} M_{e',r'}}, \qquad \overline{\ell} := \sum_{e,r} \omega_{e,r} \overline{\ell}_{e,r},$$
$$\mathcal{L}_{\mathrm{inv}} := \sum_{e,r : M_{e,r} > 0} \omega_{e,r} (\overline{\ell}_{e,r} - \overline{\ell})^2.$$

For unsupervised tasks, we set $\mathcal{L}_{\mathrm{inv}} = \mathcal{L}_{\mathrm{min}} = 0$. The predictor is spectrally normalized, and validation data determine τ_C.

Rank monitoring and complete objective. Let $\widehat{M}_\rho$ be the derivative contrast matrix obtained by automatic differentiation of fitted state- and history-conditioned log-density ratios, with singular values $\widehat{\sigma}_1 \geq \widehat{\sigma}_2 \geq \cdots$. The optional rank bias is

$$\mathcal{L}_{\mathrm{rank}} = \sum_{k=1}^{2d_Z} [\epsilon - \widehat{\sigma}_k]_+^2, \qquad d_Z = d_C + d_V,$$

where ϵ is selected on validation data. The complete minimization objective is

$$\begin{aligned} \mathcal{J}(\theta, \phi, \psi, \alpha) = -\sum_{n=1}^N \mathcal{E}_n &+ \underbrace{\lambda_{\mathrm{occ}} \mathcal{L}_{\mathrm{occ}} + \lambda_{\mathrm{sp}} \mathcal{L}_{\mathrm{sp}} + \lambda_{\mathrm{sig}} \mathcal{L}_{\mathrm{sig}}}_{\text{hierarchy-recovery bias}} \\ &+ \underbrace{\lambda_{\mathrm{sep}} \mathcal{L}_{\mathrm{sep}} + \lambda_{\mathrm{fac}} \mathcal{L}_{\mathrm{fac}} + \lambda_{\mathrm{cyc}} \mathcal{L}_{\mathrm{cyc}} + \lambda_{\mathrm{rank}} \mathcal{L}_{\mathrm{rank}}}_{\text{factor-recovery bias}} \\ &+ \underbrace{\lambda_{\mathrm{inv}} \mathcal{L}_{\mathrm{inv}} + \lambda_{\mathrm{min}} \mathcal{L}_{\mathrm{min}}}_{\text{target-refinement bias}}. \end{aligned} \tag{50}$$

All λ values are selected using validation data. The activation schedule is

$$\begin{aligned}\mathcal{A}_1 &= \{-\textstyle\sum_n \mathcal{E}_n, \mathcal{L}_{\mathrm{occ}}\},\\ \mathcal{A}_2 &= \mathcal{A}_1 \cup \{\mathcal{L}_{\mathrm{sig}}, \mathcal{L}_{\mathrm{sep}}, \mathcal{L}_{\mathrm{sp}}\},\\ \mathcal{A}_3 &= \mathcal{A}_2 \cup \{\mathcal{L}_{\mathrm{fac}}, \mathcal{L}_{\mathrm{cyc}}, \mathcal{L}_{\mathrm{rank}}, \mathcal{L}_{\mathrm{inv}}, \mathcal{L}_{\mathrm{min}}\}.\end{aligned}$$

Stage lengths and stopping rules are fixed before test access. Forward–backward inference yields state posteriors; Viterbi decoding is optional. The outputs are $(q(E \mid \boldsymbol{X}), q(\boldsymbol{R} \mid \boldsymbol{X}), \widehat{J}_C)$ and posterior predictive samples.

G.2 Reproducibility, model selection, and computational reporting

Let $m_{a,s,u}$ be method a's metric for replicate unit u under matched seed s. Seed-level aggregation is

$$m_{a,s} := \frac{1}{|\mathcal{U}_s|} \sum_{u \in \mathcal{U}_s} m_{a,s,u},$$

Here $\mathcal{U}_s$ contains sequences for the controlled SCM and HPR-MNIST. For PPG-DaLiA, window errors are first averaged within each held-out participant and participant-level values are then averaged within seed. Frames and overlapping windows are not statistical replicates. Reported method differences are descriptive comparisons of the resulting aggregate seed-level estimates.

For each candidate capacity pair $c = (k_E, k_R)$, we compute the mean held-out ELBO $\bar{\mathcal{E}}_c$ and standard error s_c across matched seeds. The one-standard-error selection rule is

$$\begin{aligned}c^\star &:= \arg\max_c \bar{\mathcal{E}}_c,\\ \mathcal{C}_{1\mathrm{SE}} &:= \left\{c : \bar{\mathcal{E}}_c \geq \bar{\mathcal{E}}_{c^\star} - s_{c^\star}\right\},\\ \widehat{c} &:= \operatorname*{lex\,argmin}_{c \in \mathcal{C}_{1\mathrm{SE}}} \left(k_E k_R,\, k_E + k_R,\, k_E,\, k_R\right).\end{aligned}$$

The effective state counts follow Eq. (18) with the fixed occupancy floor $\delta = 0.02$. All exact-cardinality metrics use these effective orders after capacity selection. True cardinalities are used only for post-hoc scoring and never for model selection. Splits, seeds, state capacity grids, probes, primary metrics, hyperparameter grids, and stopping rules are frozen before test evaluation. Failed runs remain in the reliability summary. Complete run manifests record configurations, software and hardware metadata, resource use, and failure status.

All experiments were run on four NVIDIA A40 GPUs, each with 48 GB of memory, using Ubuntu 20.04, Python 3.9, PyTorch 2.8.0, and CUDA 12.6.

Per sequence, the discrete hierarchy costs $O(T_n K_E K_R^2)$ time and $O(T_n K_E K_R)$ memory; state-conditional factor evaluation adds $O(T_n K_E K_R d)$ time for latent width d.

G.3 Metrics and alignment

State and cardinality recovery. The state metrics are

$$\begin{aligned}\mathcal{M}_{\mathrm{state}} := \{ &\underbrace{\mathrm{ARI}_E, \mathrm{NMI}_E, \mathrm{macroF1}_R^{\pi_R}, \mathrm{IoU}_{\mathrm{seg}}}_{\text{recovery metrics}},\\ &\underbrace{\mathbf{1}\{(\widehat{K}_E, \widehat{K}_R) = (K_E, K_R)\}, |\widehat{K}_E - K_E|, |\widehat{K}_R - K_R|}_{\text{cardinality errors}}\}.\end{aligned}$$

The Hungarian permutation π_R is shared across environments; separate environment-wise alignments would not test the claimed hierarchy.

Kernel and hierarchy recovery. Let $\widehat{\mu}_{e,r}$ denote the posterior-weighted fitted state-conditional observed transition kernel. Let P_{π_R} be the permutation matrix for a single global alignment of

recurrent states, and let $\mathcal{H}_\star$ denote a common history grid. Define

$$D_A := \min_{\pi_E,\pi_R} \frac{1}{K_E}\sum_{e=1}^{K_E} \frac{\left\| P_{\pi_R}^\top \widehat{A}^{(\pi_E(e))} P_{\pi_R} - A^{(e)} \right\|_F}{\|A^{(e)}\|_F},$$

$$D_\mu := \min_{\pi_E,\pi_R} \frac{1}{K_E K_R |\mathcal{H}_\star|} \sum_{e,r}\sum_{h\in\mathcal{H}_\star} \mathrm{MMD}_k^2\big(\widehat{\mu}_{\pi_E(e),\pi_R(r)}(\cdot \mid h), \mu_{e,r}(\cdot \mid h)\big).$$

The kernel k, its bandwidth, and $\mathcal{H}_\star$ are fixed on validation data. The fitted signatures follow Eq. (49). The hierarchy residual is

$$\widehat{\mathcal{R}}_\Gamma = \min_{\{u_e, v_r, \pi_e\}} \frac{1}{K_E K_R}\sum_{e=1}^{K_E}\sum_{r=1}^{K_R} \|\widehat{\Delta}_{e,\pi_e(r)} - u_e - v_r\|_2^2, \tag{51}$$

where π_e is an environment-specific candidate permutation. We report this residual, separation between translated row sets, and whether the minimizing permutations induce one common alignment of local states.

Coordinate and target subset recovery. Automatic differentiation constructs $\widehat{M}_\rho$. Report

$$\mathcal{M}_{\mathrm{rank}} := \left(\widehat{\sigma}_{2d_Z}, \frac{\widehat{\sigma}_1}{\widehat{\sigma}_{2d_Z}}, \mathrm{rank}_{\mathrm{eff}}\, \widehat{M}_\rho\right),$$

using a relative threshold selected on validation data and fixed before test evaluation. For scalar-component latent vectors, coordinate recovery is

$$\mathrm{MCC}_{\mathrm{S}}(Z, \widehat{Z}) = \frac{1}{d_Z}\max_{\pi}\sum_{j=1}^{d_Z} \left|\rho_{\mathrm{S}}(Z_j, \widehat{Z}_{\pi(j)})\right|. \tag{52}$$

We report Eq. (52) for Z^{cand}, Z^{resid}, and their concatenation. To audit cross-context alignment, we compare MCC under context-specific permutations with MCC under one pooled permutation. At the rank-path endpoint, we compare the column spaces of the ground-truth and estimated dynamic trajectory matrices, with rows aligned at the same evaluated time points. Let Q and $\widehat{Q}$ be orthonormal bases for these spaces, and let $c_1, \ldots, c_q$ be the singular values of $Q^\top \widehat{Q}$. The principal angles are

$$\theta_i = \arccos\big(\mathrm{clip}(c_i, 0, 1)\big).$$

We report their RMS-sine subspace error:

$$\mathcal{E}_{\mathrm{sub}} := \left(\frac{1}{q}\sum_{i=1}^{q} \sin^2\theta_i\right)^{1/2}. \tag{53}$$

Target selection reports precision, recall, and F1 against J_C under unique minimality. At the non-unique-minimality boundary, designated J_C F1 is computed against one prespecified minimal sufficient subset; held-out sufficiency agreement evaluates the boundary-specific recovery object.

Task and leakage metrics. For HPR-MNIST, we report interpolation and shortcut-reversal accuracy. Frozen probes measure digit information in Z^{cand}, corruption information in Z^{resid}, rendered color information in their concatenation, and target leakage from (Z^{resid}, S). PPG-DaLiA reports heart rate MAE, worst-activity MAE, activity leakage above majority from Z^{cand}, and heart rate MAE gain from (Z^{resid}, S) relative to a training mean predictor. As a complementary representation-sensitivity diagnostic, HPR-MNIST also evaluates support-respecting nuisance swaps. We report the swap acceptance rate together with selected code, target prediction, and nuisance representation changes (D_C, D_Y, D_N); Appendix H.3.3 defines the protocol.

H Experimental Protocols and Complete Results

H.1 Baselines and implementation parity

We use three comparison groups. Architecture-matched controls isolate HiDICE's components while holding fixed the temporal backbone, decoder family, latent width, optimizer, data, and validation budget. When an ablation changes model size, we evaluate both the original-width and

parameter-matched variants. External representation methods are compared only on objects they define, and task baselines are compared under identical inputs and splits.

Superscripts ^{I}, ^{L}, and ^{P}follow the convention defined in Section 5. Temporal ERM and the flat AR-HMM are in-house baselines defined and implemented for this study.

The no-signature-decomposition variant retains both discrete branches while removing alignment based on translated signatures. It therefore serves as an architecture-matched hierarchical control for separating the contribution of the hierarchy itself from that of signature-based cross-environment alignment.

The no-persistent-branch variant retains $q_\phi(E_n \mid \boldsymbol{X}_n)$ solely for diagnostic evaluation of E ARI. The no-recurrent-branch variant retains its time-indexed recurrent-state head solely for diagnostic evaluation of R F1, using the full model's single global Hungarian alignment. Neither variant uses a post-hoc hierarchy map.

The flat AR-HMM uses the HiDICE state cardinality grid and held-out likelihood selection but receives no environment labels. We apply one frozen hierarchy map, based on Eq. (51), to both the flat AR-HMM and eligible HM-NLICA (Hälvä & Hyvärinen, 2020) outputs. This separates errors in flat-state estimation from errors in hierarchy decomposition. CHiLD (Li et al., 2025b) is scored through its native hierarchical layer assignment, whereas rSLDS (Linderman et al., 2017) and CtrlNS (Song et al., 2024) are scored only on their recurrent-state, kernel, or coordinate outputs.

On HPR-MNIST, all methods share the image encoder and temporal head. Temporal ERM predicts the target directly, COGS (Song et al., 2026) infers domains without labels, and REx (Krueger et al., 2021) receives the observed (E, R) groups. On PPG-DaLiA, every task baseline uses the same wrist channels and partitions. SAB-HRT (Crepaldi et al., 2025) is reimplemented with the frozen preprocessing, while Deep PPG (Reiss et al., 2019) and PULSE (Kasnesis et al., 2026) are retrained on our splits. The REx comparison uses subject and activity labels and is therefore marked label-assisted. Oracle-E, oracle-R, and oracle-(E, R) variants are used only to diagnose headroom.

A dash in a result table therefore means that the method does not define the corresponding object, rather than that its run failed. Only HiDICE returns the complete unlabeled tuple (E, R, Z, J_C); the other methods are evaluated on the largest native subset of that tuple.

H.2 Controlled SCM: protocol and results

H.2.1 Reference data-generating process

The reference SCM exposes every recovery object and permits the hierarchy, signature, rank, support, and anchoring conditions to be checked before fitting. Environments are balanced in expectation; each environment-specific transition matrix $A^{(e)}$ is ergodic, starts at stationarity, and uses the same recurrent state labels while assigning different transition probabilities.

For $Z \in \{C, V\}$, coordinate j follows

$$Z_{n,t,j} = g_j^Z(H_{n,t}^Z) + \alpha_{E_n,j}^Z + \beta_{R_{n,t},j}^Z + \gamma_{ER}\xi_{E_n,R_{n,t},j}^Z + \sigma_{E_n,R_{n,t},j}^Z(H_{n,t}^Z)\varepsilon_{n,t,j}^Z. \tag{54}$$

and the observation and target are

$$\begin{aligned} X_{n,t} &= f_\omega([C_{n,t}; V_{n,t}; S_n]) + \sigma_X \epsilon_{n,t}^X, \\ Y_{n,t} &= g_Y(C_{n,t}) + \sigma_Y \epsilon_{n,t}^Y. \end{aligned} \tag{55}$$

The functions g_j^Z are stable nonlinear autoregressions. We rescale each linear lag operator to spectral norm below 0.8 before applying smooth elementwise nonlinearities. Additive signatures for environments and recurrent states modulate both conditional means and log variances. To certify the reference condition, we retain generators only when translated signature sets are separated by at least 0.5 and $\sigma_{\min}(M_\rho) \geq 0.25$. These checks are never provided to the estimator.

An injective random network maps the six latent coordinates to twelve observed coordinates through invertible affine layers, leaky-ReLU activations, and a full-column-rank expansion. We add isotropic Gaussian observation noise with standard deviation 0.05. The target is a noisy nonlinear function of the first coordinate of each C component, so the minimal target-relevant subset is unique and

available for evaluation. Unless otherwise stated, the reference controlled SCM uses $(K_E, K_R) = (3, 4)$, $(d_C, d_V, d_S) = (2, 2, 2)$, lag order 2, and $N = 200$ training sequences of length $T = 400$, with 50 validation and 100 test sequences.

H.2.2 IDENTIFICATION SETTING

For the controlled SCM, we fit overcomplete models over the prespecified Cartesian product capacity grid

$$K_E^{\max} \in \{2, 3, 4\}, \qquad K_R^{\max} \in \{3, 4, 5\}.$$

Selection and effective order estimation follow Appendix G.2, including the one-standard error rule on held-out ELBO and the fixed occupancy threshold $\delta = 0.02$. The ground-truth cardinalities $(K_E, K_R) = (3, 4)$ are used only for post-hoc exact recovery scoring.

The reference experiment evaluates the identification chain under matched latent-process conditions. Every environment and recurrent state has positive probability, and the training trajectories contain every positive-probability (e, r) configuration. The finite-sample decoder is fit directly to noisy observations, so the added Gaussian noise forms part of the estimation setting rather than a direct implementation of the population deconvolution argument.

H.2.3 BOUNDARY CONSTRUCTIONS AND EVALUATION

The signature path decreases the minimum distance between translated row sets while preserving identifiable local switching kernels and the overall signature scale.

The anchoring test isolates global coordinate alignment. Its reference generator contains three static context strata with connected support overlap and no nontrivial context-dependent coordinate symmetry. In the unanchored condition, each stratum receives a different coordinate permutation and componentwise invertible map; the observation and transition mechanisms are composed accordingly, preserving the observed law within each context. The transformed transition kernels remain S-independent, with all other state, signature, rank, noise, and target parameters fixed. We compare $\mathrm{MCC}_{\mathrm{S}}$ under a separate matching per stratum with $\mathrm{MCC}_{\mathrm{S}}$ under one matching shared across strata. The former measures the boundary-specific collection of context-specific charts; the latter measures a single global chart.

The rank path removes one independent direction without making the contrasts sparser. We construct $M_\rho(\lambda) = U\,\mathrm{diag}(1, \ldots, 1, \lambda)W^\top$, where U and W are fixed dense orthogonal factors within seed and $\lambda \in \{1, .75, .5, .25, .10, 0\}$. The resulting contrasts are mapped to mean and log-variance signature coefficients, and instances with a singleton contrast are rejected. The Frobenius norm and average pairwise contrast magnitude are normalized across λ so that the path primarily controls the weakest independent direction rather than overall signal strength. Coordinate recovery is evaluated for $\lambda > 0$; at $\lambda = 0$, we report the trajectory subspace error in Eq. (53).

We evaluate four boundary constructions with matched seeds, targeting signature separation, contrast rank, cross-context anchoring, and unique target minimality. For each construction, Table 5 reports the evaluation object and endpoint diagnostics. Reference-setting results are reported in Table 4; the three continuous paths used in the main text are shown in Figure 3. State aliasing and many-to-one mixing are instead represented by the explicit failure constructions in Appendix F; no finite benchmark can resolve objects that are observationally equivalent by construction.

H.2.4 COMPLETE RESULTS

Reference recovery. Table 4 extends the main-text comparison with additional hierarchy, transition, signature, and target-subset diagnostics. HiDICE uniquely recovers the complete unlabeled tuple (E, R, Z, J_C). Across comparable metrics, it attains the best mean recovery, with lower transition, kernel, and translated-signature errors, together with higher continuous-factor and target-subset recovery. Lower translated-signature error and better cross-environment alignment localize the gain to persistent–recurrent separation before the final coordinate-recovery stage.

Identification boundaries. Table 5 gives the complete endpoint diagnostics for the three boundary paths summarized in Figure 3 and additionally reports the target-minimality boundary. At this fourth boundary, low agreement with one designated J_C coexists with high held-out sufficiency agreement,

Table 4: End-to-end recovery on the reference controlled SCM. Panel A reports state count and hierarchy recovery; Panel B reports kernel, signature, and factor recovery.

Panel A: state counts and hierarchy.

Method	**Exact ↑**	E **ARI ↑**	E **NMI ↑**	R **F1 ↑**	**Seg. IoU ↑**
Flat AR-HMM[I,P]	2/5	0.595 ± 0.051	0.641 ± 0.052	0.644 ± 0.042	0.525 ± 0.047
rSLDS (2017)	–	–	–	0.688 ± 0.051	0.554 ± 0.055
HM-NLICA (2020)[P]	3/5	0.682 ± 0.054	0.719 ± 0.049	0.792 ± 0.045	0.597 ± 0.043
CHiLD (2025b)	3/5	0.767 ± 0.047	0.779 ± 0.043	0.759 ± 0.041	0.621 ± 0.049
CtrlNS (2024)	–	–	–	0.786 ± 0.036	0.651 ± 0.043
HiDICE (ours)	**5/5**	$\mathbf{0.886 \pm 0.046}$	$\mathbf{0.898 \pm 0.034}$	$\mathbf{0.852 \pm 0.033}$	$\mathbf{0.744 \pm 0.037}$

Panel B: kernel, signature, and factors.

Method	$D_A \downarrow$	$D_\mu \downarrow$	$\widehat{\mathcal{R}}_\Gamma \downarrow$	Z **MCC**$_S$ ↑	J_C **F1 ↑**
Flat AR-HMM[I,P]	0.218 ± 0.029	0.186 ± 0.025	0.161 ± 0.028	0.664 ± 0.043	0.472 ± 0.084
rSLDS (2017)	0.193 ± 0.026	0.163 ± 0.023	–	0.706 ± 0.047	–
HM-NLICA (2020)[P]	0.165 ± 0.023	0.141 ± 0.021	0.124 ± 0.024	0.787 ± 0.035	–
CHiLD (2025b)	0.147 ± 0.021	0.126 ± 0.018	0.108 ± 0.022	0.779 ± 0.038	–
CtrlNS (2024)	0.134 ± 0.018	0.116 ± 0.016	–	0.812 ± 0.024	–
HiDICE (ours)	$\mathbf{0.079 \pm 0.015}$	$\mathbf{0.066 \pm 0.012}$	$\mathbf{0.057 \pm 0.013}$	$\mathbf{0.862 \pm 0.027}$	$\mathbf{0.806 \pm 0.061}$

Notes. Exact (K_E, K_R) recovery counts matched seeds (out of five) recovering both cardinalities. A dash denotes an undefined object.

as expected when multiple inclusion-minimal sufficient subsets are compatible with the target law. The dense-rank construction varies contrast rank without relying on singleton contrasts.

Table 5: Stagewise recovery at designed identification boundaries. Each panel reports the identification target, the boundary-specific recoverable object, and endpoint diagnostics.

Identification target	**Recoverable object at boundary**	**Endpoint diagnostics**
	Panel A: signature aliasing.	
Environment partition and globally shared R labels	Sequence-local recurrent state classes	E ARI ↑: 0.389 ± 0.083; shared R F1 ↑: 0.601 ± 0.052; local R F1 ↑: $\mathbf{0.776 \pm 0.036}$
	Panel B: rank-deficient contrasts.	
Coordinate recovery	Trajectory subspace	Fitted $\sigma_{\min}(M_\rho)$: 0.028 ± 0.009; Subspace error ↓: $\mathbf{0.176 \pm 0.041}$
	Panel C: no cross-context anchoring.	
One globally aligned coordinate chart	Context-specific coordinate charts	Global MCC_S ↑: 0.593 ± 0.071; alignment stability: $1/5$; within-context MCC_S ↑: $\mathbf{0.844 \pm 0.033}$
	Panel D: non-unique minimal sufficient subsets.	
Unique inclusion-minimal J_C	Alternative sufficient subsets	Designated J_C F1 ↑: 0.515 ± 0.121; Sufficiency agreement ↑: $\mathbf{0.941 \pm 0.026}$

Notes. Bold values highlight the designated recovery metrics. At exact rank deficiency, the reported $\sigma_{\min}(M_\rho)$ is fitted (the generated value is zero). Designated J_C F1 uses one prespecified minimal sufficient subset and is diagnostic only.

H.3 HPR-MNIST: PROTOCOL AND RESULTS

H.3.1 SOURCE PARTITION AND SEQUENCE ASSEMBLY

HPR-MNIST uses $K_E = 3$ persistent environments and $K_R = 4$ recurrent mechanism states. Before any rendering or sequence construction, we partition the underlying MNIST images to prevent

an image exemplar from leaking across splits. The 60,000 original training images are divided into 50,000 training and 10,000 validation images, and the official 10,000-image test set is retained. Each split maintains an independently shuffled pool for every digit and reshuffles a pool only after exhaustion. We then generate 500 training, 100 validation, and 100 test sequences of length $T = 100$ with separate random seeds.

In the full-support interpolation and reversal settings, each local chain is ergodic, has self-transition probability 0.8, and assigns positive probability to every state. Independently of $(E_n, R_{n,1:T})$, digit identity follows the stationary first-order Markov chain

$$\begin{aligned} D_{n,1} &\sim \mathrm{Uniform}\{0, \dots, 9\}, \\ \Pr(D_{n,t} = j \mid D_{n,t-1} = i) &= \begin{cases} 0.6, & j = i, \\ 0.4/9, & j \neq i. \end{cases} \end{aligned} \tag{56}$$

This symmetric transition matrix has a uniform stationary distribution and is shared across environments, recurrent states, and evaluation distributions. It therefore introduces temporal dependence without changing target prevalence or revealing (E, R). For digit $D_{n,t}$, the target, shortcut corruption, and rendered color are

$$\begin{aligned} Y_{n,t} &= \mathbf{1}\{D_{n,t} \geq 5\}, \\ V_{n,t} &\sim \mathrm{Bernoulli}(1 - \rho_{E_n, R_{n,t}}), \\ B_{n,t} &= Y_{n,t} \oplus V_{n,t}. \end{aligned} \tag{57}$$

where $B_{n,t}$ indexes the rendered color and

$$\rho = \begin{pmatrix} .95 & .85 & .75 & .65 \\ .90 & .80 & .70 & .60 \\ .85 & .75 & .65 & .55 \end{pmatrix}. \tag{58}$$

All training reliabilities exceed chance, but neither E nor R is identical to the target.

No single visual cue deterministically identifies an environment or recurrent state. Both color classes use the same base hues in every environment, and environment-dependent hue and saturation distributions overlap. The four recurrent states apply no rotation, positive rotation, negative rotation, or partial occlusion, with continuous jitter around each transformation. The static context fixes the background texture and stroke-width perturbation within a sequence, but its distribution also overlaps across environments.

H.3.2 Evaluation distributions and frozen probes

The interpolation split samples all 12 (e, r) pairs. The shortcut-reversal test replaces $\rho_{e,r}$ with $1 - \rho_{e,r}$ only at test time, leaving the digit and state processes, style support, and marginal color frequency unchanged.

To keep the factor diagnostics separate from representation learning, all probes are trained after freezing HiDICE. Two-layer probes predict digit identity from Z^{cand}, the corruption bit from Z^{resid}, rendered color from their concatenation, and the target from (Z^{resid}, S). Width, regularization, and early stopping are selected on validation sequences only. Target leakage is the corresponding target accuracy above chance, 0.5.

H.3.3 Support-respecting nuisance-swap diagnostic

Frozen leakage probes ask whether target information is decodable from a learned nuisance branch; the swap diagnostic instead measures sensitivity to an explicit change in the complete learned nuisance representation. Consider equal-length posterior segments a and b, including their prescribed lag histories. With the validation-fixed gate g, define

$$\begin{aligned} N_t &:= \big((1 - g) \odot Z_t^{\mathrm{cand}},\, Z_t^{\mathrm{resid}}, S\big), \\ Z_t^{\mathrm{cand}, a \leftarrow b} &:= g \odot Z_t^{\mathrm{cand}, a} + (1 - g) \odot Z_t^{\mathrm{cand}, b}. \end{aligned}$$

Thus the hybrid retains the selected target path from source a while taking the nonselected candidate coordinates, residual branch, and static branch from donor b. It is decoded and re-encoded using the frozen model.

To avoid evaluating arbitrary off-support hybrids, let $\widehat{J}_C := \{j : g_j \geq \tau_C\} \neq \varnothing$ and let $\widetilde{H}_t^{\mathrm{cand},a\leftarrow b}$ denote the candidate history formed from the hybrid path. In terms of the coordinatewise transition factors in Eq. (15), a prespecified donor pair is accepted only if

$$\Gamma_{a\leftarrow b} := \frac{1}{T|\widehat{J}_C|} \sum_{t=1}^{T} \sum_{j\in\widehat{J}_C} \log p_\theta\Big(Z_{t,j}^{\mathrm{cand},a} \mid \widetilde{H}_t^{\mathrm{cand},a\leftarrow b}; \eta_{E^b,R_t^b,j}^{\mathrm{cand}}\Big) \geq \tau_{\mathrm{swap}}. \tag{59}$$

Both τ_C and the validation-quantile threshold τ_{swap} are fixed before test evaluation. For the prespecified donor pair set $\mathcal{P}$, let $\mathcal{A} := \{(a,b) \in \mathcal{P} : \Gamma_{a\leftarrow b} \geq \tau_{\mathrm{swap}}\}$; the acceptance rate is $R_{\mathrm{acc}} = |\mathcal{A}|/|\mathcal{P}|$.

For a length-T, d-dimensional path, write $\|U\|_{T,d}^2 := (Td)^{-1}\sum_{t=1}^{T}\|U_t\|_2^2$. Let $\widehat{Z}_{a\leftarrow b}^{\mathrm{sel}}$ and $\widehat{N}_{a\leftarrow b}$ be the re-encoded selected and nuisance paths of the hybrid, and let p_t^a and $\widehat{p}_t^{a\leftarrow b}$ be the corresponding target predictions. We report

$$D_C := \mathbb{E}_{\mathcal{A}}\Big[\|\widehat{Z}_{a\leftarrow b}^{\mathrm{sel}} - Z^{\mathrm{sel},a}\|_{T,d_{\mathrm{sel}}}^2\Big],$$
$$D_Y := \mathbb{E}_{\mathcal{A}}\Big[T^{-1}\sum_{t=1}^{T} d_Y\big(\widehat{p}_t^{a\leftarrow b}, p_t^a\big)\Big],$$
$$D_N := \mathbb{E}_{\mathcal{A}}\Big[\|\widehat{N}_{a\leftarrow b} - N_a\|_{T,d_N}^2\Big],$$

where the expectations are uniform over accepted pairs, d_{sel} and d_N are the corresponding path dimensions, and d_Y is total-variation distance for classification. The four prespecified controls are unrestricted swaps, state-matched unfiltered swaps, support-respecting swaps, and swaps after constructed candidate/residual mixing. Donor rules, support models, gate and support thresholds, and the fixed mixing map are selected on validation data and frozen before test evaluation.

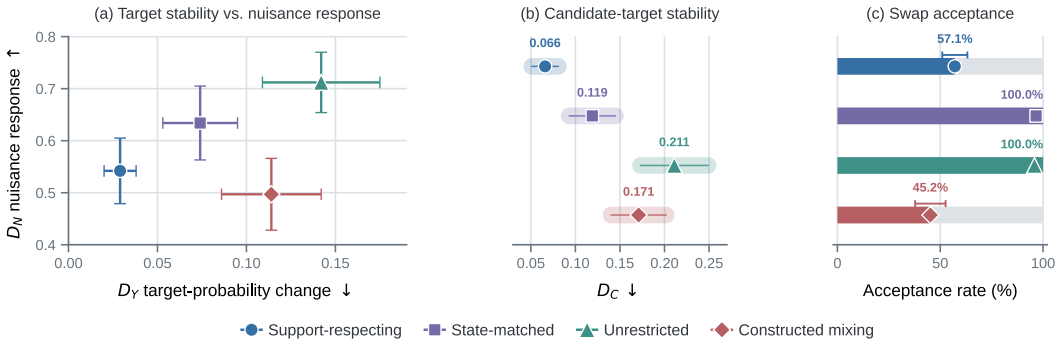

Figure 5: Support-respecting nuisance-swap diagnostic on HPR-MNIST. (a) Target prediction change D_Y versus nuisance response D_N; lower D_Y with nonzero D_N indicates stronger prediction stability under an effective nuisance perturbation. (b) Selected code change D_C. (c) Swap acceptance rate; gray segments denote rejected swaps. Constructed mixing applies the same swaps after a fixed invertible mixing of candidate and residual factors.

Support-respecting swaps in Figure 5 produced the smallest prediction change ($D_Y = 0.029 \pm 0.009$) and selected code change while retaining a nonzero nuisance response, with a 57.1% acceptance rate. The unrestricted and constructed-mixing controls argue against explaining this stability by ineffective perturbations alone. Together, the frozen probes and support-respecting swaps provide evidence for the stability of the selected target path under effective nuisance perturbations on observed support.

Complete results. Table 6 complements the shortcut-reversal endpoint with hierarchy and factor-allocation probes. Relative to the flat AR-HMM, HiDICE improved both environment and recurrent-state recovery, achieved higher frozen-probe accuracy for digit prediction from Z^{cand} and corruption prediction from Z^{resid}, and reduced target leakage from the nuisance representation. HiDICE also retained high interpolation accuracy, improved reversal accuracy by 0.417 over Temporal ERM, and

achieved the highest label-free reversal accuracy without access to the true (E, R) groups. Together with the swap diagnostic above, the branch-specific probe accuracies and lower target leakage provide representation-level evidence of selective factor allocation rather than uniform suppression of nuisance variation.

Table 6: HPR-MNIST hierarchy recovery, factor allocation, and shortcut-reversal robustness. Panel A reports hierarchy and factor probes; Panel B reports in-distribution and reversal accuracy.

Panel A: hierarchy and factor probes.

Method	$\boldsymbol{E}$ **ARI** ↑	$\boldsymbol{R}$ **F1** ↑	**Digit**| $\boldsymbol{Z}^{\mathbf{cand}}$ ↑	**Corrupt.**| $\boldsymbol{Z}^{\mathbf{resid}}$ ↑	**Target leakage (Δ acc.)** ↓
Temporal ERM[I]	–	–	–	–	–
REx (2021)[L]	–	–	–	–	–
COGS (2026)	0.703 ± 0.082	–	–	–	–
Flat AR-HMM[I,P]	0.602 ± 0.074	0.698 ± 0.051	0.833 ± 0.045	0.724 ± 0.052	0.203 ± 0.044
HiDICE (ours)	$\mathbf{0.801 \pm 0.049}$	$\mathbf{0.831 \pm 0.044}$	$\mathbf{0.904 \pm 0.025}$	$\mathbf{0.823 \pm 0.038}$	$\mathbf{0.091 \pm 0.032}$

Panel B: generalization.

Method	**ID acc.** ↑	**Reversal acc.** ↑
Temporal ERM[I]	$\mathbf{0.966 \pm 0.002}$	0.371 ± 0.067
REx (2021)[L]	0.951 ± 0.004	$\mathbf{0.812 \pm 0.036}$
COGS (2026)	0.958 ± 0.002	0.704 ± 0.051
Flat AR-HMM[I,P]	0.956 ± 0.005	0.619 ± 0.074
HiDICE (ours)	0.953 ± 0.006	0.788 ± 0.049

Notes. Target leakage is probe accuracy from (Z^{resid}, S) above chance; reversal accuracy uses the flipped target–color association. COGS infers domains without hierarchy labels.

H.4 PPG-DaLiA: protocol and results

Evaluation protocol. PPG-DaLiA provides wrist blood volume pulse at 64 Hz, wrist acceleration at 32 Hz, and activity annotations for 15 participants (Reiss et al., 2019). We use only wrist PPG and triaxial acceleration as inputs. Acceleration is resampled to 64 Hz, PPG is band-pass filtered between 0.5 and 4 Hz, and all channels are standardized with training partition statistics. Chest signals, EDA, temperature, and raw ECG are excluded; ECG is used only through the released heart rate target.

We assign partitions before extracting eight-second windows with a two-second stride. An eight-second guard interval at every boundary prevents overlapping samples from crossing partitions. Windows that intersect an activity boundary remain eligible for heart rate evaluation but are excluded from activity proxy metrics.

Seen-participant evaluation. This split measures interpolation without random-window leakage. Every sufficiently long subject–activity block is divided, in acquisition order, into contiguous 60%/20%/20% training, validation, and test segments, with the guard interval applied at both boundaries.

Unseen participant–activity pairs. This split tests composition while retaining every subject and activity during training. With subjects $s_1, \ldots, s_{15}$ and activities $a_1, \ldots, a_8$ in released order, the test pairs are $(s_i, a_{1+((i-1) \bmod 8)})$ and the validation pairs are $(s_i, a_{1+(i \bmod 8)})$; all remaining pairs form the training set. If a pair lacks the minimum number of windows, we select the next activity in cyclic order before fitting and freeze that list for all methods.

Leave-one-subject-out evaluation. LOSO is the primary test of subject generalization. Fold i reserves s_i for testing, $s_{1+(i \bmod 15)}$ for validation, and the remaining 13 subjects for training. The test participant contributes to neither normalization nor model selection, early stopping, or calibra-

tion. Subject and activity labels are hidden from HiDICE and all unsupervised controls; only the explicitly label-assisted REx baseline receives them.

The participant, not the window, is the statistical unit. We average window errors within each participant before computing the cross-participant mean. Worst-activity MAE is the largest participant-averaged error among activities observed for at least ten participants. Activity leakage from Z^{cand} and heart rate information in (Z^{resid}, S) are descriptive factor diagnostics. Because activity and acquisition order may be correlated, these diagnostics are not interpreted as identification of the underlying activity mechanism; the primary real-data endpoint is held-out heart rate prediction.

Complete results. Table 7 extends the main-text comparison with seen-participant interpolation and reports all four evaluation endpoints together. Without subject or activity labels, HiDICE attained the lowest mean MAE in every setting. On seen participants, its MAE was 2.863 ± 0.095 BPM, compared with 2.916 ± 0.083 for label-assisted REx. This seen-participant comparison complements the shift-focused main-text result by showing that the same leading position for HiDICE was already observed under interpolation.

Table 7: PPG-DaLiA heart rate MAE under seen-participant, unseen participant–activity pair, leave-one-subject-out (LOSO), and worst-activity evaluation (BPM).

Method	**Seen MAE ↓**	**Pair MAE ↓**	**LOSO MAE ↓**	**Worst-act. MAE ↓**
SAB-HRT (2025)	3.892 ± 0.117	5.106 ± 0.144	6.217 ± 0.182	11.624 ± 0.315
Deep PPG (2019)	4.424 ± 0.134	5.979 ± 0.172	7.412 ± 0.247	13.283 ± 0.372
PULSE (2026)	2.978 ± 0.096	3.852 ± 0.127	4.785 ± 0.145	8.918 ± 0.281
Temporal ERM^{I}	3.011 ± 0.109	4.031 ± 0.168	4.938 ± 0.163	9.181 ± 0.343
REx (2021)^{L}	2.916 ± 0.083	3.794 ± 0.133	4.661 ± 0.154	8.634 ± 0.296
HiDICE (ours)	$\mathbf{2.863 \pm 0.095}$	$\mathbf{3.687 \pm 0.143}$	$\mathbf{4.513 \pm 0.126}$	$\mathbf{8.346 \pm 0.339}$

H.5 Cross-dataset ablations and diagnostics

Table 8: HiDICE component ablations across the controlled SCM, HPR-MNIST, and PPG-DaLiA. Panel A evaluates hierarchy recovery components; Panel B evaluates factor allocation and target selection.

Panel A: hierarchy ablations.

Variant	**Controlled SCM**		**HPR-MNIST**
	$\boldsymbol{E}$ **ARI ↑**	$\boldsymbol{R}$ **F1 ↑**	**Reversal acc. ↑**
No persistent branch	0.401 ± 0.112	–	0.629 ± 0.083
No recurrent branch	–	0.548 ± 0.096	0.603 ± 0.089
No signature decomposition	0.647 ± 0.091	0.702 ± 0.068	–
HiDICE (ours)	$\mathbf{0.886 \pm 0.046}$	$\mathbf{0.852 \pm 0.033}$	$\mathbf{0.788 \pm 0.049}$

Panel B: factor allocation and target selection ablations.

Variant	**Controlled SCM**	**HPR-MNIST**		**PPG-DaLiA**	
	$\boldsymbol{J_C}$ **F1 ↑**	**Reversal acc. ↑**	**Target leakage (Δ acc.) ↓**	**LOSO MAE ↓**	**Activity leakage (Δ acc.) ↓**
No factor split	0.438 ± 0.118	0.651 ± 0.072	0.246 ± 0.064	4.760 ± 0.180	0.151 ± 0.032
No target selector	0.522 ± 0.137	0.713 ± 0.058	0.164 ± 0.047	4.640 ± 0.160	0.121 ± 0.026
HiDICE (ours)	$\mathbf{0.806 \pm 0.061}$	$\mathbf{0.788 \pm 0.049}$	$\mathbf{0.091 \pm 0.032}$	$\mathbf{4.513 \pm 0.126}$	$\mathbf{0.069 \pm 0.018}$

Notes. Target leakage is measured from (Z^{resid}, S) and activity leakage from Z^{cand}; both are probe accuracies above their respective baselines. Factor split denotes the separation into Z^{cand}, Z^{resid}, and S.

Table 8 separates hierarchy-recovery failures from factor-allocation and target-selection failures. Removing the persistent, recurrent, or signature component primarily degrades the corresponding E/R recovery measures. On PPG-DaLiA, removing the factor split increases activity leakage from

0.069 ± 0.018 to 0.151 ± 0.032 and LOSO MAE from 4.513 ± 0.126 to 4.760 ± 0.180 BPM. Removing the target selector produces smaller increases to 0.121 ± 0.026 and 4.640 ± 0.160, respectively. These ablations expose distinct failure modes: hierarchy components primarily affect structural recovery, whereas factor allocation and target selection affect nuisance separation and target sufficiency. The paired leakage and predictive endpoints provide complementary diagnostics of the latter two failures.