Support-Safe Variational Hybrid Filtering for Contact-Mode and Sparse-Law Recovery

Marios Papamichalis[*]
marios.papamichalis@yale.edu

Regina Ruane[†]
jeanneruane@gmail.com

Abstract

Contact-rich robot dynamics are hybrid: a single observation can match several latent states and contact regimes (free, impact, stick–slip). A standard amortized filter that places no probability on a feasible contact transition will permanently lose the branch the robot actually follows. We introduce VHYDRO, a variational hybrid dynamics learner that prevents this branch loss. At each step, VHYDRO mixes the learned proposal with a feasible transition law before sampling and importance weighting, ensuring that every transition retained by the model-feasible carrier remains covered. VHYDRO jointly infers a continuous latent state and a discrete contact mode, and fits a sparse port-Hamiltonian law to each recovered regime. On top of this, three guarantees connect: support coverage stabilizes filtering, the stabilized filter concentrates the discrete contact posterior on coherent regimes, and mode-pure segments admit sparse port-Hamiltonian recovery. The recovery error separates cleanly into filtering, derivative, mode-impurity, and physics-residual parts. Three empirical findings track the same mechanism. Under heavy occlusion the support-safe filter stays usable while a non-defensive proposal collapses. On ManiSkill demonstrations and on four Sawyer/BridgeData task families the discrete state forms temporally coherent contact-regime segments and yields a stronger joint profile across ARI, change-point F1, and segment purity than post-hoc and mode-free baselines. On hybrid systems with known equations the mode-conditioned sparse fit recovers the active physical terms; purely predictive baselines do not.

1 Introduction

Contact-rich manipulation exposes a failure mode that average one-step error can hide: a support-mismatched amortized filter may delete a model-feasible contact branch before later observations reveal it. Latent dynamics models and variational state-space models are useful because they maintain hidden state, produce action-conditioned predictive rollouts, and reason under uncertainty. Modern latent dynamics models and variational state-space models provide scalable predictive states from images, proprioception, and recorded robot-object state streams[Watter et al., 2015, Ha and Schmidhuber, 2018, Hafner et al., 2019b,a, Karl et al., 2016, Maddison et al., 2017, Naesseth et al., 2018]. The challenge addressed here is to combine scalable variational filtering with hybrid contact structure, interpretable mode-conditioned dynamics, and explicit protection against proposal-level support loss.

Contact-rich dynamics, however, are not globally smooth. A trajectory may switch between free motion, sticking, sliding, and impact, and each regime follows different local physics. Treating these regimes as one smooth transition can make a model confident in the wrong future, which is dangerous for planning because rollout errors compound over imagined trajectories. Hybrid and contact-aware

[*]Human Nature Lab, Yale University, New Haven, CT 06511, marios.papamichalis@yale.edu

[†]This work was finished while Regina Ruane was working in the Department of Statistics and Data Science, The Wharton School, University of Pennsylvania, 3733 Spruce Street, Philadelphia, PA 19104-6340, jeanneruane@gmail.com.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

models address parts of this issue through switching state-space models, Hybrid-SINDy, ContactNets, and physically structured contact dynamics [Linderman et al., 2017, Mangan et al., 2019, Pfrommer et al., 2021, Hochlehnert et al., 2021]. Particle and sequential Monte Carlo methods provide stronger posterior inference for nonlinear state-space models [Doucet et al., 2001, Andrieu et al., 2010], and defensive mixtures reduce support mismatch by ensuring that proposals cover the target support [Hesterberg, 1995]; yet these methods are often expensive to amortize at the scale of vision-based robot learning.

A complementary line of work seeks interpretable physical structure. Sparse and symbolic discovery methods recover governing equations from trajectories or learned coordinates [Brunton et al., 2016, Rudy et al., 2017, Schmidt and Lipson, 2009, Champion et al., 2019], while Bayesian extensions add uncertainty quantification and robustness under noisy or partial observations [Mars Gao and Nathan Kutz, 2024, Fung et al., 2025, Rosafalco et al., 2024, Conti et al., 2026]. Physics-informed, continuous-time, Hamiltonian, Lagrangian, and port-Hamiltonian models inject useful inductive biases for stability, passivity, and extrapolation [Raissi et al., 2019, Chen et al., 2018, Greydanus et al., 2019, Cranmer et al., 2020, Lutter et al., 2019, Desai et al., 2021, Rettberg et al., 2025]. Relational dynamics and object-centric models capture multi-object and articulated interactions [Battaglia et al., 2016, Kipf et al., 2018, Sanchez-Gonzalez et al., 2018, 2020, Bishnoi et al., 2024, Jatavallabhula et al., 2023], and contact-rich manipulation benchmarks provide natural evaluation settings [Mandlekar et al., 2018, 2021, Walke et al., 2023, Gu et al., 2023, Xiang et al., 2020, Mo et al., 2019].

Despite this progress, the ingredients needed for contact-rich hybrid filtering remain largely separated across method families. Scalable variational latent models are typically not explicitly hybrid or physically interpretable. Physics-discovery methods are usually not designed for raw observations, actions, and partial observability. Posterior-correct particle methods provide accurate inference, but are difficult to deploy as end-to-end learned perception models. This leaves a need for hybrid variational state-space models that combine scalable perception, interpretable dynamics, uncertainty diagnostics, and model-relative support-safe inference.

We propose VHYDRO (*variational hybrid dynamics with robust proposal support*), a support-safe hybrid variational state-space model for partially observed contact dynamics. VHYDRO represents the manipulation scene by a continuous latent state z_t, and represents the active predictive regime by a discrete mode s_t. The transition associated with each mode is parameterized by a port-Hamiltonian core with sparse library corrections, yielding dynamics that are physically structured and inspectable.

Inference in VHYDRO uses a support-safe variational family. A standard amortized proposal may assign zero or very small probability to model-feasible contact states, producing unstable importance weights and degrading learning. VHYDRO mitigates this failure mode by mixing the learned proposal with a feasible transition carrier. The resulting mixture satisfies $P_t \ll Q_{t,\lambda}$ and gives positive $Q_{t,\lambda}$-mass to every measurable transition set with positive P_t-mass.

The contribution is not a new defensive-IS inequality; it is a deployable support-safety mechanism for amortized hybrid filtering: choose the mixture mass before sampling, sample from the same mixture used in the importance weight, and audit the resulting representation through filtering stability, mode coherence, and sparse-law recovery.

Contributions. **(i) Support-safe IW/FIVO gate.** We sample from

$$Q_{t,\lambda} = (1-\lambda)Q_t + \lambda P_t$$

and weight using the same realized mixture. Specializing the classical defensive-mixture inequality [Hesterberg, 1995] to the one-step IW/FIVO increment gives $P_t \ll Q_{t,\lambda}$, $dP_t/dQ_{t,\lambda} \le 1/\lambda$, and the pre-sampling rule $\lambda_t \ge \bar{\rho}_t/(1+N\tau^2)$ whenever the requested relative-variance budget is certifiable. The paper-specific point is operational rather than a new defensive-IS inequality: the theoretical rule uses the certificate before sampling inside an amortized hybrid variational filter, with the same λ_t in the proposal and the importance weight.

(ii) Predictive contact modes under support repair. The stabilized filter maintains a continuous latent state and a discrete predictive mode. Theorem 2 states the standard predictive-separation condition under which such modes concentrate. The primary segmentation audit measures alignment with frozen kinematic proxy regimes. An additional PushCube audit evaluates agreement with simulator contact-force labels.

(iii) Controlled sparse-law audit. Once mode-pure segments are available, Theorem 3 applies a classical Lasso oracle bound [Bickel et al., 2009] to mode-conditioned port-Hamiltonian regression with aggregate plug-in perturbations. Experiment 3 measures aggregate support and coefficient recovery on known-equation systems; it does not separately certify the individual filtering, derivative, mode-impurity, and port-Hamiltonian residual terms.

Headline result and falsification target. The clearest predicted-then-verified signature is the high-occlusion regime of Experiment 1. Theorem 1 directly controls the relative normalizer-variance and ESS frontier; NLL and coverage are reported as downstream collapse diagnostics, not as quantities directly bounded by the theorem. The $\lambda = 0$ baseline uses the same architecture, particle count, objective, observation likelihood, resampling rule, training budget, and numerical stabilizers as the support-safe variants; it removes only the defensive carrier mass. If the gains came merely from generic clipping or implementation stabilization rather than support repair, this matched $\lambda = 0$ baseline should retain the same ESS/N and relative-variance frontier. Instead, at 90% occlusion, the conservative support-safe variant improves ESS/N from 0.042 to 0.141 on Contact-style tasks and from 0.046 to 0.151 on Sawyer/BridgeData-style diagnostics.

2 Theory

Classical tools appear here as certifiable components, but the way they are wired together is specific to VHYDRO. Theorem 1 turns defensive mixing into a history-conditioned, pre-sampling IW/FIVO support gate; the same realized λ_t is used in both proposal and weight. Theorem 2 states the predictive-separation condition under which the discrete posterior concentrates on a coherent mode using only likelihood-ratio evidence. Theorem 3 then applies sparse recovery to mode-pure port-Hamiltonian segments with aggregate filtering, derivative, mode-impurity, and plug-in perturbations. The empirical section tests these three links separately.

Model and training. VHYDRO parameterizes a hybrid latent state $x_t = (s_t, z_t)$, where z_t is a continuous belief state and s_t is a discrete predictive contact regime. The amortized proposal Q_t is observation-adaptive and sharp. The carrier P_t is the model's one-step transition law before conditioning on o_{t+1}:

$$P_t(dx_{t+1}) = \pi_\theta(ds_{t+1} \mid x_t, a_t)\, p_\theta(dz_{t+1} \mid s_{t+1}, z_t, a_t).$$

Its discrete support is the model-feasible mode support, and its continuous support is the corresponding mode-conditioned transition support with the validation-fixed noise floor shared across variants. Thus P_t is not assumed to be the true physical law; it declares the model-feasible support that the learned proposal is not allowed to delete. Training samples from

$$Q_{t,\lambda} = (1 - \lambda)Q_t + \lambda P_t,$$

and weights particles using the same realized mixture density. Support safety is therefore model-relative: VHYDRO preserves every measurable transition set assigned positive mass by P_t, but it cannot recover true physical branches omitted by P_t. After filtering, mode-pure segments are fit with a sparse port-Hamiltonian library by regularized regression. The mode count and candidate library are fixed in the present experiments.

The learned parameters are optimized with the sequential IW/FIVO objective: at step t, the one-step normalizer estimate is

$$\widehat{Z}_t = \frac{1}{N}\sum_{i=1}^{N} p_\theta(o_{t+1} \mid z_{t+1}^i)\frac{dP_t}{dQ_{t,\lambda_t}}(x_{t+1}^i), \qquad x_{t+1}^i \sim Q_{t,\lambda_t},$$

and training maximizes the corresponding sum of log normalizer increments. Thus the same realized mixture is used for sampling and weighting in both training and evaluation. The theorem requires an h_t-measurable upper certificate $\bar{\rho}_t \geq \rho_t$ before particles are drawn. In the experiments, λ_t is selected by a validation-frozen plug-in rule based on empirical likelihood moments. Its variance-frontier plots are diagnostics and are not used for post-hoc retuning. When the plug-in carrier-mass demand exceeds 1, upper saturation activates the validation-fixed fallback $\lambda_{\mathrm{fb}} = 1$. The empirical rule does not certify the requested τ, including when the fallback is inactive.

Support-safe hybrid filtering step.
Input: particles $\{(s_t^i, z_t^i, w_t^i)\}_{i=1}^N$, action a_t, observation o_{t+1}, amortized proposal Q_t, feasible transition law P_t, h_t-measurable upper certificate $\bar{\rho}_t \geq \rho_t$, tolerance τ, and validation-fixed fallback $\lambda_{\text{fb}} \in (0, 1]$.

1. Choose λ_t before sampling. Let $\lambda_{\min,t} := \bar{\rho}_t/(1+N\tau^2)$. If $\lambda_{\min,t} \leq 1$, set $\lambda_t = \lambda_{\min,t}$, the smallest mixture mass that certifies the requested budget. Otherwise mark the requested budget as uncertified and use the validation-fixed fallback λ_{fb}; this fallback is reported as an uncertified operating point, not as a proof of the requested τ.
2. Sample each next particle from
$$Q_{t,\lambda_t} = (1 - \lambda_t)Q_t + \lambda_t P_t.$$
3. Weight the sample using the realized mixture density:
$$W_t(x) = p_\theta(o_{t+1} \mid z_{t+1}) \frac{dP_t}{dQ_{t,\lambda_t}}(x).$$
4. Normalize weights, compute ESS/N and relative weight variance, and resample when the ESS rule triggers.
5. Return the filtered particles and the mode posterior $q_\phi(s_{t+1} \mid o_{\leq t+1}, a_{\leq t})$.

Figure 1: Conditional rule for the VHYDRO filtering step. The same λ_t is used for sampling and weighting. No post-hoc proposal adaptation is used. The empirical plug-in implementation is described in Appendix A.1.

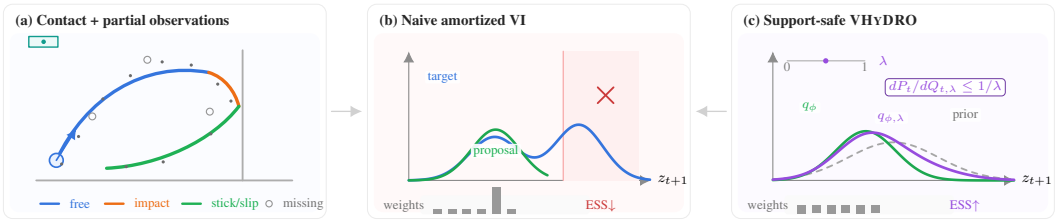

Figure 2: Why support-safe hybrid inference matters. (a) Under partial observation, multiple contact-consistent futures remain feasible; filled dots denote observed states and hollow dots denote missing observations. (b) A narrow amortized proposal can miss a feasible posterior branch, concentrating importance weights and lowering ESS. (c) VHYDRO mixes the learned proposal with a feasible transition law, retaining support and bounding the transition-to-proposal density ratio $dP_t/dQ_{t,\lambda} \leq 1/\lambda$.

The empirical rule and its variance-frontier plots do not provide the conditional certificate in Theorem 1. VHYDRO combines support-safe hybrid filtering, mode concentration, and sparse-law recovery; Theorems 1–3 formalize these measured links.

At filtering step t, write $x_{t+1} := (s_{t+1}, z_{t+1}) \in \mathsf{X} := \{1, \ldots, M\} \times \mathbb{R}^d$. Conditioned on a particle history $h_t = (s_t, z_t, a_t, o_{\leq t+1})$, let P_t be the model one-step transition law used as the feasible support carrier, and let Q_t be the observation-adaptive amortized proposal. For $\lambda \in (0, 1]$, define the defensive proposal

$$Q_{t,\lambda} := (1 - \lambda)Q_t + \lambda P_t. \tag{1}$$

Let $g_t(x) := p_\theta(o_{t+1} \mid z_{t+1})$, $Z_t := \mathbb{E}_{P_t}[g_t(X)]$, $M_{2,t} := \mathbb{E}_{P_t}[g_t(X)^2]$, and $\Pi_t(dx) := g_t(x)P_t(dx)/Z_t$.

Theorem 1 (Support-safe budgeted IW/FIVO increment)**.** *Assume $0 < Z_t < \infty$ and $M_{2,t} < \infty$. Then $P_t \ll Q_{t,\lambda}$, $\Pi_t \ll Q_{t,\lambda}$, and*

$$0 \leq \frac{dP_t}{dQ_{t,\lambda}}(x) \leq \frac{1}{\lambda} \qquad Q_{t,\lambda}\text{-a.s.} \tag{2}$$

Consequently, for the one-step weight $W_t(x) := g_t(x)dP_t/dQ_{t,\lambda}(x)$,

$$\chi^2(\Pi_t \| Q_{t,\lambda}) \leq \frac{1}{\lambda}\rho_t - 1, \qquad \rho_t := \frac{M_{2,t}}{Z_t^2}. \tag{3}$$

Since $M_{2,t} \geq Z_t^2$ by Jensen's inequality, $\rho_t \geq 1$, so the right-hand side is nonnegative for $\lambda \in (0, 1]$.

If $X^{1:N} \overset{\text{i.i.d.}}{\sim} Q_{t,\lambda}$ and $\widehat{Z}_t = N^{-1} \sum_{i=1}^N W_t(X^i)$, then

$$\begin{aligned} \mathbb{E}[\widehat{Z}_t \mid h_t] &= Z_t, \\ \frac{\mathrm{Var}(\widehat{Z}_t \mid h_t)}{Z_t^2} &\leq \frac{1}{N}\left(\frac{\rho_t}{\lambda} - 1\right), \\ \frac{\mathrm{ESS}_N}{N} &\to \frac{1}{1 + \chi^2(\Pi_t \| Q_{t,\lambda})} \geq \frac{\lambda}{\rho_t}. \end{aligned} \tag{4}$$

Moreover, if an h_t-measurable certificate $\bar{\rho}_t \geq \rho_t$ is available, the relative-variance budget $\mathrm{Var}(\widehat{Z}_t \mid h_t)/Z_t^2 \leq \tau^2$ is certified whenever

$$\lambda \geq \frac{\bar{\rho}_t}{1 + N\tau^2}, \qquad \frac{\bar{\rho}_t}{1 + N\tau^2} \leq 1. \tag{5}$$

Theorem 1 makes the support role explicit: if $g_t \in L^2(P_t)$, the one-step IW/FIVO normalizer estimate has relative variance bounded by $(\rho_t/\lambda - 1)/N$. We test this through ESS/N, relative weight variance, empirical relative variance, NLL, ECE, and Cov90.

Relation to prior bounds. The density-ratio inequality is the classical defensive-mixture argument specialized to an IW/FIVO filtering increment. The VHYDRO-specific contribution is operational and sequential: under the theorem's certificate assumption, λ_t is chosen at each history h_t before particles are drawn, the same realized mixture is used for weighting, and Eq. (5) becomes a support-coverage gate inside the hybrid filter rather than a post-hoc static importance-sampling diagnostic.

Theorem 2 (Exponential concentration of predictive modes)**.** *Consider a length-L segment generated by a single latent segment mode m. Actions are treated as fixed interventions in the logged evaluation segment, or are drawn by a nonanticipative behavior policy whose conditional action law is common to all candidate modes; under this intervention/off-policy convention, action probabilities cancel from mode odds. If the policy likelihood depends on the candidate mode, that likelihood must be included in Eq.* (6)*, and the theorem below does not apply in this form. Let S be the segment mode and let*

$$\Pi_L(S = s) \propto \pi_s \prod_{\ell=0}^{L-1} p_s(O_{\ell+1} \mid \mathcal{F}_\ell) \tag{6}$$

be the clamped-mode VHYDRO *posterior after integrating or particle-marginalizing the continuous latent state z. The theorem uses only predictive likelihood increments and does not assume simulator contact labels. For each wrong mode $s \neq m$, define*

$$R_{\ell,s} := \log \frac{p_m(O_{\ell+1} \mid \mathcal{F}_\ell)}{p_s(O_{\ell+1} \mid \mathcal{F}_\ell)}. \tag{7}$$

Assume $\sum_{\ell=0}^{L-1} \mathbb{E}_\star[R_{\ell,s} \mid \mathcal{F}_\ell] \geq \Gamma_{L,s}$ and that the centered increments are conditionally sub-Gaussian with variance proxy $V_s := \sum_\ell \sigma_{\ell,s}^2$. Let $B_s := \log(\pi_s/\pi_m)$ and $u_{s,\delta} := \sqrt{2V_s \log((M-1)/\delta)}$. For every $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\begin{aligned} \Pi_L(S \neq m) &\leq \frac{A_{L,\delta}^\sharp}{1 + A_{L,\delta}^\sharp} \leq A_{L,\delta}^\sharp, \\ A_{L,\delta}^\sharp &:= \sum_{s \neq m} \exp\left(B_s - \Gamma_{L,s} + u_{s,\delta}\right). \end{aligned} \tag{8}$$

In particular, if $\Gamma_{L,s} \geq L\Delta$, $B_s \leq B$, and $V_s \leq L\sigma^2$, then, with $u_{L,\delta} := \sqrt{2L \log((M-1)/\delta)}$,

$$\Pi_L(S \neq m) \leq (M-1) \exp\left(B - L\Delta + \sigma u_{L,\delta}\right). \tag{9}$$

If the amortized mode posterior satisfies $\mathrm{KL}(q_\phi(S \mid O_{1:L}, A_{0:L-1}) \| \Pi_L) \leq \varepsilon_q$, then on the same event

$$q_\phi(S \neq m \mid \cdot) \leq \frac{A_{L,\delta}^\sharp}{1 + A_{L,\delta}^\sharp} + \sqrt{\varepsilon_q/2}. \tag{10}$$

Theorem 2 states the predictive-separation condition needed for posterior concentration of a predictive mode up to label permutation. We test this with Mode F1, ARI, change-point F1, and segment purity computed against frozen kinematic proxy labels.

Connection to Experiment 2. The theorem identifies the sufficient quantities for concentration: log-likelihood separation $\Gamma_{L,s}$, variance proxy V_s, and prior log-odds B_s. We do not claim a fully numerical per-task certificate of these constants. Instead, Table 10 reports held-out metrics computed with the frozen kinematic proxy labels, auditing whether the recovered discrete state behaves like a coherent predictive contact regime.

Theorem 3 (Filtering- and mode-robust sparse port-Hamiltonian recovery)**.** *Fix a recovered mode m with index set I_m and $n_m := |I_m|$. Suppose the true mode law has sparse port-Hamiltonian form*

$$\begin{aligned} \dot{z}_i &= (J_m - R_m)\nabla H_m(z_i) + G_m a_i + \varepsilon_i, \\ H_m(z) &= \Theta(z)^\top \xi_m^\star. \end{aligned} \tag{11}$$

where $J_m^\top = -J_m$, $R_m \succeq 0$, and $\xi_m^\star$ is k-sparse on support S. Let VHYDRO *produce the stacked plug-in regression*

$$\begin{aligned} \widehat{b} &= \widehat{A}\xi_m^\star + \zeta_m, \\ \widehat{A}_i &= (\widehat{J}_m - \widehat{R}_m)\nabla\Theta(\widehat{z}_i)^\top, \\ \widehat{b}_i &= \widehat{\dot{z}}_i - \widehat{G}_m a_i. \end{aligned} \tag{12}$$

Assume the score and restricted-strong-convexity conditions

$$\begin{aligned} &\left\| \frac{1}{n_m} \widehat{A}^\top \zeta_m \right\|_\infty \le \frac{\lambda_m}{2}, \\ &\frac{1}{n_m}\|\widehat{A}\Delta\|_2^2 \ge \kappa_m \|\Delta\|_2^2 \quad \textit{for } \|\Delta_{S^c}\|_1 \le 3\|\Delta_S\|_1. \end{aligned} \tag{13}$$

For the Lasso estimator

$$\widehat{\xi}_m \in \arg\min_{\xi} \left\{ \frac{1}{2n_m}\|\widehat{b} - \widehat{A}\xi\|_2^2 + \lambda_m \|\xi\|_1 \right\}, \tag{14}$$

we have

$$\begin{aligned} \|\widehat{\xi}_m - \xi_m^\star\|_2 &\le \frac{4\sqrt{k}\lambda_m}{\kappa_m}, \\ \frac{1}{n_m}\|\widehat{A}(\widehat{\xi}_m - \xi_m^\star)\|_2^2 &\le \frac{16 k \lambda_m^2}{\kappa_m}. \end{aligned} \tag{15}$$

If $\beta_{\min,m} := \min_{j\in S} |\xi_{m,j}^\star| > 8\sqrt{k}\lambda_m/\kappa_m$, thresholding $\widehat{\xi}_m$ at any level in $(4\sqrt{k}\lambda_m/\kappa_m,\ \beta_{\min,m} - 4\sqrt{k}\lambda_m/\kappa_m)$ recovers S. If $\zeta_m = e_m + u_m$, e_m is conditionally sub-Gaussian with scale σ_m, and $\|(1/n_m)\widehat{A}^\top u_m\|_\infty \le \nu_m$, then the score part of condition (13) *holds with probability at least $1-\delta$ for*

$$\begin{aligned} \lambda_m &\ge 2\left[L_{A,m}\sigma_m \sqrt{\frac{2\log(2p/\delta)}{n_m}} + \nu_m \right], \\ L_{A,m} &:= \max_j \frac{\|\widehat{A}_{\cdot j}\|_2}{\sqrt{n_m}}. \end{aligned} \tag{16}$$

The restricted-strong-convexity part of condition (13) is assumed for the realized design; if it holds only on a separate event, the probability statement is combined with that event by a union bound.

Under the score and restricted-strong-convexity assumptions, Theorem 3 gives finite-sample sparse-recovery bounds for mode-conditioned port-Hamiltonian regression with aggregate filtering, derivative, mode-impurity, and plug-in perturbations. Once mode-pure segments are available, VHYDRO can recover interpretable equations evaluated by sparse support F1, coefficient error, vector-field NRMSE, and physical-constant error.

Relation to prior bounds. The result is a Lasso oracle bound applied to a mode-conditioned port-Hamiltonian library with plug-in errors. The four-term decomposition is used as claim bookkeeping: it identifies the aggregate sources of perturbation that can affect sparse recovery, but the experiments measure aggregate sparse support and coefficient recovery rather than separately certifying each perturbation component.

The four components of

$$\nu_m = \nu_m^{\text{filt}} + \nu_m^{\text{der}} + \nu_m^{\text{mode}} + \nu_m^{\text{pH}}$$

are defined in Appendix B.3. They should be read as an aggregate error accounting device, not as four independently measured empirical certificates.

Assumptions and claim boundary. All three guarantees are model-relative. The carrier P_t must assign positive mass to the physical branch one wants to preserve; a branch omitted by P_t cannot be recovered by support mixing. The mode set and sparse library are finite and fixed in the present experiments. The action-cancellation condition in Theorem 2 matches the logged-intervention evaluation protocol used here; mode-dependent closed-loop policies require including the policy likelihood in the mode odds. Finally, the results below audit filtering, segmentation, and sparse-law recovery, not certified closed-loop control.

3 Experiments

The experiments audit the three mechanism links: support-safe IW/FIVO under occlusion, proxy-label contact-regime segmentation, and sparse-law recovery on known-equation systems. We report ESS/N, relative weight variance, empirical relative variance, NLL, ECE, Cov90, Mode F1, ARI, change-point F1, segment purity, support F1, coefficient error, physical-constant error, and vector-field NRMSE. ESS/N is a normalized effective sample size. Cov90 measures nominal 90% interval coverage. ECE measures coverage gaps. ARI is the adjusted Rand index. NRMSE normalizes RMSE by target RMS. Appendix A gives the full definitions.

Shared protocol and matched ablations. Each applicable method/task/condition cell uses 20 independent seeds, with no seed pruning. Within each matched-ablation comparison, variants share architecture family, latent dimension, particle count, training budget, observation likelihood, resampling rule, occlusion schedule when applicable, and proxy-label construction. The $\lambda = 0$ / no-support ablation removes only defensive support mass; the smooth/no-mode/no-sparsity/no-pH ablations remove the indicated mechanism. Proxy labels are deterministic functions of recorded state/action fields and are used only for evaluation. These are mechanism-isolation audits, not closed-loop task-success, or real-robot claims.

Experiment 1: high-occlusion support-mismatch audit. Contact-style denotes PickCube-v1, PokeCube-v1, PullCube-v1, and PushCube-v1 under controlled occlusion. We report two support-safe operating points on the same support-mass/sharpness frontier. VHYDRO-adaptive uses the default validation-frozen rule and spends only the support mass selected before sampling. VHYDRO-conservative uses a deliberately larger fixed support mass, buying variance margin at the cost of a less sharp proposal. The conservative variant is therefore not a separate architecture or a post-hoc diagnostic; it is a second operating point predicted by Theorem 1. The non-defensive $\lambda = 0$ / no-support baseline is matched: it uses the same architecture, particle count, objective, observation likelihood, resampling rule, training budget, and numerical stabilizers, and removes only the defensive carrier mass.

Experiment 1 uses sampled Contact-style transition families and combines state-based latent-coordinate masking with controlled proposal stress while keeping the carrier unchanged. Additional 20-seed comparisons use matched evaluation settings whose scale differs from the main tables. At 90% occlusion, adaptive clean-carrier NLL/Cov90 is $0.66/0.72$, versus $10.09/0.43$ without defense. Generic proposal controls have NLL at least 5.3. Appendix A.1 gives the observation encoding, protocols, carrier-stress tests, and sensitivity results.

At 90% occlusion, defensive support mixing prevents the predictive collapse seen in the non-defensive filter. On Contact-style tasks, NLL drops from 11.286 for $\lambda = 0$ to 0.161 for the conservative support-safe operating point; on Sawyer/BridgeData-style diagnostics, NLL drops from 9.401 to

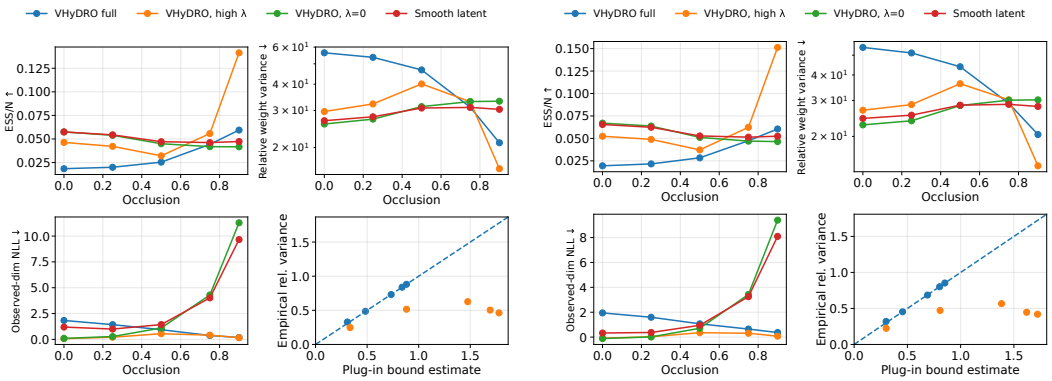

(a) Contact-style family. **(b)** Sawyer/BridgeData family.

Figure 3: Support-safe filtering under partial observation. Occlusion makes the one-step posterior multimodal and exposes proposal support mismatch. Defensive support mixing increases high-occlusion ESS and lowers weight instability. NLL is shown as a downstream collapse diagnostic; the finite-variance certificate in Theorem 1 directly concerns the normalizer-estimation variance and ESS frontier.

Table 1: **Experiment 1: high-occlusion filtering metrics.** At 90% occlusion, support-safe proposals retain feasible contact branches. NLL is reported as a downstream collapse diagnostic; the theorem directly controls relative normalizer variance and the ESS frontier. Each family row aggregates four tasks or diagnostic families with 20 seeds each (80 runs per method-family cell). Appendix A.1 reports the full sweep and mean±SEM values.

Family	Method	ESS/N ↑	Rel. w-var ↓	Emp. rel-var ↓	NLL ↓	ECE ↓	Cov90 ↑
Contact-style	VHYDRO-conservative	**0.141**	**15.851**	**0.248**	**0.161**	**0.125**	**0.756**
Contact-style	VHYDRO-adaptive	0.059	21.057	0.329	0.186	0.152	0.727
Contact-style	Smooth latent	0.047	30.366	0.474	9.660	0.469	0.356
Contact-style	VHYDRO-no-support	0.042	33.195	0.519	11.286	0.467	0.358
Sawyer/BridgeData-style	VHYDRO-conservative	**0.151**	**14.355**	**0.224**	**0.074**	**0.111**	**0.772**
Sawyer/BridgeData-style	VHYDRO-adaptive	0.060	20.432	0.319	0.369	0.147	0.732
Sawyer/BridgeData-style	Smooth latent	0.052	28.050	0.438	8.095	0.440	0.391
Sawyer/BridgeData-style	VHYDRO-no-support	0.046	30.205	0.472	9.401	0.435	0.397

0.074. ESS/N, which is directly tied to the theorem's normalizer-variance frontier, rises from 0.042 to 0.141 and from 0.046 to 0.151, respectively. The adaptive operating point also avoids the large NLL collapse but remains closer to the variance frontier. Thus the claim is high-occlusion support repair, not uniform calibration dominance: Appendix A.1 shows that low-occlusion coverage can be sharper for adaptive support mixing.

Experiment 2: contact-regime segmentation with proxy labels on manipulation demonstrations.

Datasets. Experiment 2 uses released ManiSkill 2 HDF5 demonstrations [Gu et al., 2023] on PushCube-v1, PokeCube-v1, PullCubeTool-v1, and PegInsertionSide-v1. The main table reports this released-demonstration audit. Appendix A.3 reports the same segmentation protocol on four author-defined Sawyer/BridgeData-style diagnostic families: BridgeDataOcclusion, SawyerTowerCreation, SawyerLaundryLayout, and SawyerObjectSearch. These names denote fixed diagnostic groupings used in this paper, not canonical released BridgeData benchmark task labels.

This experiment uses frozen, deterministic kinematic proxy labels only as an evaluation target; they are matched to predicted modes after training and should not be read as physical contact-state ground truth. The HMM baseline is a post-hoc segmenter fit to the same kinematic features, while VHYDRO uses the discrete state inside the predictive filter; the no-support and no-mode ablations remove defensive mixing and discrete contact structure, respectively.

On held-out episodes, VHYDRO has the strongest joint profile across ARI, change-point F1, Mode F1, and segment purity. The most reliable gains are on ARI and change-point F1. The Mode F1 gain over HMM is positive but smaller relative to its SEM, so we interpret the result as improved predictive segmentation rather than uniform dominance on every proxy-label metric.

Table 2: **Experiment 2: contact-regime segmentation with proxy labels on ManiSkill HDF5 demonstrations.** Metrics are computed on held-out test episodes across four tasks. Mode F1 uses denoised kinematic proxy labels, not simulator contact-force ground truth. The aggregate mean±SEM is computed over the held-out task/seed cells used in the main-table aggregation; Table 10 reports per-task means with seed-level SEMs, so the aggregate SEM and per-task SEMs summarize different sources of variation.

Method	Mode F1 ↑	ARI ↑	Change-point F1 ↑	Segment purity ↑
VHYDRO	**0.724±0.026**	**0.647±0.019**	**0.652±0.019**	**0.919±0.007**
no-support	0.612±0.014	0.465±0.018	0.433±0.018	0.915±0.005
HMM	0.675±0.033	0.490±0.032	0.527±0.028	0.890±0.005
no-mode	0.221±0.002	0.000±0.000	0.000±0.000	0.496±0.006

Additional 20-seed audits in Appendix A.2 compare rSLDS and vary proxy construction. The autoregressive-HMM fallback selected post hoc has lower Mode F1 than VHYDRO (0.620 versus 0.779). The default rSLDS has higher segment purity in a separate implementation diagnostic. VHYDRO ranks first on ARI and change-point F1 across $9/9$ proxy settings. Direct force-label validation is limited to PushCube, where agreement is moderate and VHYDRO retains the highest Mode F1 (0.58, versus 0.55 for HMM).

Experiment 3: sparse physical-law recovery. We next test a controlled identifiability question: given sufficiently mode-pure filtered states, can the sparse port-Hamiltonian library recover the active physical terms? Because the manipulation demonstration datasets do not provide ground-truth port-Hamiltonian coefficients, Experiment 3 uses a known-equation audit over four controlled hybrid systems. It is not an additional manipulation-demonstration dataset and is not evidence of broad real-robot system identification. The candidate library has at most $p = 12$ terms per system and the true active support has size at most $k = 3$. We compare VHYDRO with no-mode, no-sparsity, and no-pH ablations to separate accurate prediction from accurate mechanism recovery.

All sparse-regression variants use the same candidate library, normalization, state and derivative estimates, validation split, and regularization grid. No method is given the true active support. Thus Table 3 tests whether mode-conditioned segments make the shared sparse-recovery problem easier, not whether VHYDRO receives more favorable sparse-regression tuning.

Table 3: **Experiment 3: sparse physical-law recovery.** Values are mean±standard deviation over the four known-equation systems, twenty seeds per system, and the fixed perturbation conditions described in Appendix A.4. Dashes denote metrics that are undefined because the no-pH baseline has no port-Hamiltonian coefficient vector.

Method	Support F1 ↑	Rel. coeff. error ↓	Vector-field NRMSE ↓	Phys.-constant error ↓
VHYDRO-full	0.978 ± 0.042	0.006 ± 0.010	0.001 ± 0.001	0.001 ± 0.001
no-mode	0.713 ± 0.039	1.087 ± 0.330	0.287 ± 0.196	0.188 ± 0.234
no-sparsity	0.771 ± 0.073	0.851 ± 0.559	0.043 ± 0.043	0.014 ± 0.008
no-pH	–	–	0.031 ± 0.015	–

Table 3 reports the controlled known-equation audit; predictive baselines do not provide the same sparse mechanism recovery.

An exploratory Hybrid-SINDy adaptation achieves support F1 0.814, versus 0.976 for VHYDRO. The baseline discovers regimes unsupervised, whereas our sparse fit uses filtered assignments. In a separate prediction audit, 50-step rollout NRMSE is 0.0009, versus 0.010–0.146 for the ablations. A planar-pushing bridge obtains support F1 0.78, versus 1.00 with oracle segmentation and a 5% median error for the dominant kinetic coefficient. This controlled bridge uses an HMM E-step in place of the amortized mode posterior. Appendix A.4 gives the full results, library-sensitivity analysis, and sparse-recovery diagnostics.

The planning evidence consists of conditional theory in Theorem 4, MPC diagnostics, and predictive-belief calibration in Appendix A.5. These results do not establish broad robotic performance gains.

4 Conclusion and limitations

Amortized variational filters can catastrophically delete feasible contact branches under partial observation. VHYDRO prevents this by using a model-relative defensive transition carrier with a pre-sampling variance budget. This support repair stabilizes filtering, and the stabilized filter yields coherent predictive modes that can support downstream sparse-law recovery. Technically, the defensive mixture ensures $P_t \ll Q_{t,\lambda}$ and bounds the transition-to-proposal density ratio for transition sets retained by the feasible carrier P_t. Empirically, the paper tests the three links of this mechanism chain: high-occlusion IW/FIVO diagnostics, proxy-label contact-regime segmentation, and known-equation sparse-law recovery.

Limitations. The guarantees are model-relative: if the carrier P_t omits a real physical branch, VHYDRO cannot recover it. Branch deletion raises wrong-mode errors from 3.2% to 6.1%, including high-confidence errors. The pre-sampling demand and evidence alarm do not reliably detect this failure at individual steps. Accumulated predictive NLL reveals the mismatch after observations arrive. The mode set and sparse library are fixed in the present experiments. The action-cancellation assumption in Theorem 2 matches logged interventions or mode-independent behavior policies; mode-dependent closed-loop policies require a different likelihood-ratio statement. The experiments are matched mechanism audits, not claims of closed-loop task-success dominance, external SOTA ranking, simulator contact-force ground-truth recovery, real-robot transfer, or real-robot sparse-law identification. A natural follow-up is closed-loop MPC with a held-out predictive KL or calibration certificate, which would let recovered mode-wise laws support chance-constrained planning rather than only offline mechanism audits.

Acknowledgments and Disclosure of Funding

No funding was received for this work. The authors declare no competing interests.

References

Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3): 269–342, 2010.

Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. *Advances in neural information processing systems*, 29, 2016.

Peter J Bickel, Ya'acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of lasso and dantzig selector. 2009.

Suresh Bishnoi, Ravinder Bhattoo, Jayadeva, Sayan Ranu, and NM Anoop Krishnan. Discovering symbolic laws directly from trajectories with hamiltonian graph neural networks. *Machine Learning: Science and Technology*, 5(3):035049, 2024.

Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.

Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45): 22445–22451, 2019.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Paolo Conti, Jonas Kneifl, Andrea Manzoni, Attilio Frangi, Jörg Fehr, Steven L Brunton, and J Nathan Kutz. Veni, vindy, vici: A generative reduced-order modeling framework with uncertainty quantification. *Neural Networks*, page 108543, 2026.

Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. *arXiv preprint arXiv:2003.04630*, 2020.

Shaan A Desai, Marios Mattheakis, David Sondak, Pavlos Protopapas, and Stephen J Roberts. Port-hamiltonian neural networks for learning explicit time-dependent dynamical systems. *Physical Review E*, 104(3):034312, 2021.

Arnaud Doucet, Nando De Freitas, and Neil Gordon. An introduction to sequential monte carlo methods. In *Sequential Monte Carlo methods in practice*, pages 3–14. Springer, 2001.

Lloyd Fung, Urban Fasel, and Matthew Juniper. Rapid bayesian identification of sparse nonlinear dynamics from scarce and noisy data. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 481(2307), 2025.

Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *Advances in neural information processing systems*, 32, 2019.

Jiayuan Gu, Fanbo Xiang, Xuanlin Li, Zhan Ling, Xiqiang Liu, Tongzhou Mu, Yihe Tang, Stone Tao, Xinyue Wei, Yunchao Yao, et al. Maniskill2: A unified benchmark for generalizable manipulation skills. *arXiv preprint arXiv:2302.04659*, 2023.

David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2(3):440, 2018.

Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019a.

Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, pages 2555–2565. PMLR, 2019b.

Tim Hesterberg. Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37(2):185–194, 1995.

Andreas Hochlehnert, Alexander Terenin, Steindór Sæmundsson, and Marc Deisenroth. Learning contact dynamics using physically structured neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 2152–2160. PMLR, 2021.

Krishna Murthy Jatavallabhula, Miles Macklin, Dieter Fox, Animesh Garg, and Fabio Ramos. Bayesian object models for robotic interaction with differentiable probabilistic programming. In *Conference on Robot Learning*, pages 1563–1574. PMLR, 2023.

Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick Van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.

Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *International conference on machine learning*, pages 2688–2697. Pmlr, 2018.

Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial intelligence and statistics*, pages 914–922. PMLR, 2017.

Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. *arXiv preprint arXiv:1907.04490*, 2019.

Chris J Maddison, John Lawson, George Tucker, Nicolas Heess, Mohammad Norouzi, Andriy Mnih, Arnaud Doucet, and Yee Teh. Filtering variational objectives. *Advances in neural information processing systems*, 30, 2017.

Ajay Mandlekar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao, John Emmons, Anchit Gupta, Emre Orbay, et al. Roboturk: A crowdsourcing platform for robotic skill learning through imitation. In *Conference on Robot Learning*, pages 879–893. PMLR, 2018.

Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. *arXiv preprint arXiv:2108.03298*, 2021.

Niall M Mangan, Travis Askham, Steven L Brunton, J Nathan Kutz, and Joshua L Proctor. Model selection for hybrid dynamical systems via sparse regression. *Proceedings. Mathematical, physical, and engineering sciences*, 475(2223):20180534, 2019.

L Mars Gao and J Nathan Kutz. Bayesian autoencoders for data-driven discovery of coordinates, governing equations and fundamental constants. *Proceedings of the Royal Society A*, 480(2286): 20230506, 2024.

Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 909–918, 2019.

Christian Naesseth, Scott Linderman, Rajesh Ranganath, and David Blei. Variational sequential monte carlo. In *International conference on artificial intelligence and statistics*, pages 968–977. PMLR, 2018.

Samuel Pfrommer, Mathew Halm, and Michael Posa. Contactnets: Learning discontinuous contact dynamics with smooth, implicit representations. In *Conference on Robot Learning*, pages 2279–2291. PMLR, 2021.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

Johannes Rettberg, Jonas Kneifl, Julius Herb, Patrick Buchfink, Jörg Fehr, and Bernard Haasdonk. Data-driven identification of latent port-hamiltonian systems. *Computational Science and Engineering*, 2(1):4, 2025.

Luca Rosafalco, Paolo Conti, Andrea Manzoni, Stefano Mariani, and Attilio Frangi. Ekf–sindy: Empowering the extended kalman filter with sparse identification of nonlinear dynamics. *Computer Methods in Applied Mechanics and Engineering*, 431:117264, 2024.

Mark Rudelson and Shuheng Zhou. Reconstruction from anisotropic random measurements. In *Conference on Learning Theory*, pages 10–1. JMLR Workshop and Conference Proceedings, 2012.

Samuel H Rudy, Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Data-driven discovery of partial differential equations. *Science advances*, 3(4):e1602614, 2017.

Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell, and Peter Battaglia. Graph networks as learnable physics engines for inference and control. In *International conference on machine learning*, pages 4470–4479. PMLR, 2018.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR, 2020.

Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *science*, 324(5923):81–85, 2009.

Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pages 1723–1736. PMLR, 2023.

Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *Advances in neural information processing systems*, 28, 2015.

Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11097–11107, 2020.

A Supplementary material for experiments

Experiment 1 reports the full occlusion sweep, plug-in variance diagnostics, calibration diagnostics, and hard-occlusion tables. Experiment 2 reports proxy-label construction, validation-only hyperparameter selection, per-task mode metrics, and Sawyer/BridgeData diagnostics. Experiment 3 reports sparse-law recovery diagnostics.

Seed and data convention. The experiments use 20 seeds. Additional comparisons retain their stated task coverage and protocol scale. Methods are matched within each comparison. Their absolute values can therefore differ from those in the main tables. In particular, the additional high-occlusion filtering comparison has an ESS/NLL/Cov90 anchor of $0.045/10.1/0.43$ for $\lambda = 0$, compared with $0.042/11.3/0.36$ in Table 1. The word "sample" in algorithms and proofs refers to internal Monte Carlo particle draws from the IW/FIVO proposal, not to dataset subsampling. Proxy labels in Experiment 2 and the Sawyer/BridgeData diagnostic are deterministic functions of recorded state/action fields and are used only for validation diagnostics and held-out evaluation.

The supplementary proofs give the full derivations for Theorems 1–3, which formalize the measured filtering, mode-concentration, and sparse-recovery links.

Additional comparison protocols. External-baseline and sensitivity studies use the released pipeline with 20 seeds per condition. Experiment 3 comparisons use the known-equation protocol. Filtering and carrier-stress studies use 90% occlusion on the Contact-style family unless Sawyer is specified. Evaluation settings differ from Table 1. The $\lambda = 0$ ESS/N, NLL, and Cov90 values are $0.045/10.1/0.43$, compared with $0.042/11.3/0.36$ in that table. Methods within each comparison share data and evaluation settings.

Metric definitions. For normalized weights $\widetilde{w}_i = W_i / \sum_j W_j$, $\mathrm{ESS}/N = (N \sum_i \widetilde{w}_i^2)^{-1}$, with the zero-weight convention in the proof of Theorem 1. The empirical relative-variance estimator is $(N \sum_i \widetilde{w}_i^2 - 1)/N$. Its plug-in bound, $(\widehat{\rho}_t/\lambda - 1)/N$, is undefined at $\lambda = 0$. NLL is the negative log mixture marginal per transition on observed coordinates. The occlusion audit normalizes it per observed dimension. Cov90 measures the fraction of observed coordinates within ± 1.645 posterior standard deviations. ECE averages absolute gaps between empirical and nominal coverage at $0.50, 0.68, 0.90, 0.95$. This is a coverage-gap ECE, not a binned-confidence ECE. ARI is the permutation-invariant adjusted Rand index on per-frame mode labels. NRMSE divides RMSE by the target RMS.

A.1 Experiment 1: support-safe filtering audit

Protocol. All methods use the same architecture family, latent dimension, training/evaluation budget, number of particles, occlusion levels $\{0, 0.25, 0.50, 0.75, 0.90\}$, and the same 20 random seeds. The variants differ only in the proposal-support mechanism: adaptive support mixing, $\lambda = 0$, high fixed λ, and a smooth single-mode latent baseline. We report ESS/N, relative weight variance, empirical estimator relative variance, NLL, ECE, and 90% coverage at every occlusion level. Bold table entries mark the best value within each occlusion or task block. The occlusion audit combines masking with controlled proposal-support stress. Observations are state-based next-latent vectors $y = \mathrm{mask} \odot (z_1 + \sigma\epsilon)$, with $\sigma = 0.08$, in a 16-dimensional PCA space. Images and a separate proprioceptive channel are not used. Coordinates are retained independently with probability $1 - r$, subject to retention of at least one coordinate. The mask is supplied to the proposal. NLL is normalized per observed dimension. Coverage uses observed coordinates. Each variant receives proposal stress of strength $s = 1.35 \max\{0, (r - 0.20)/0.80\}(0.35 + 0.65c)$, where c is the contact proxy divided by its batchwise 90th percentile, with denominator at least 10^{-6}, and clipped to $[0, 1.5]$. With $d = (\tanh z_{0,1}, -\tanh a_1, 0.5, 0, \ldots, 0)$, each component mean is shifted by $0.70sd$, each log standard deviation is reduced by $1.25s$ and clipped to $[-5, 1]$, and every component logit except the first is reduced by $3s$. The carrier remains unchanged, isolating λ. Experiment 1 uses sampled contact-style transition families identified by task names. Experiment 2 uses ManiSkill HDF5 demonstrations with denoised kinematic proxy labels. Experiment 3 uses synthetic systems with known supports.

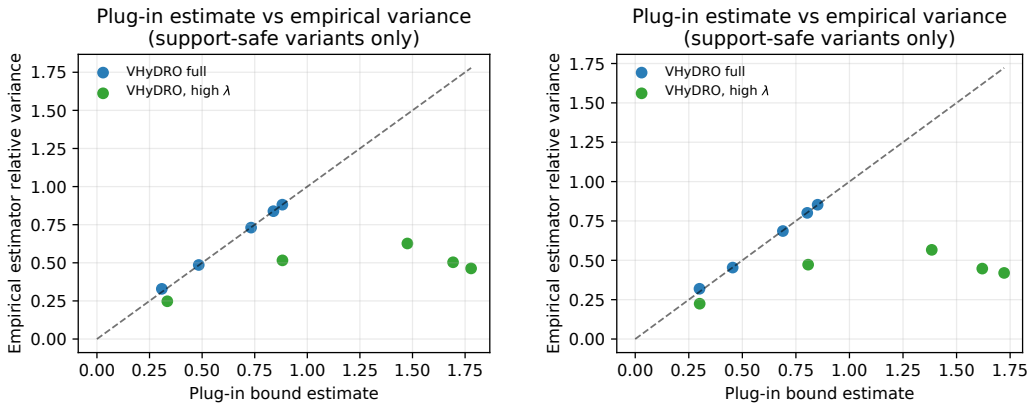

(a) Contact-style family.

(b) Sawyer/BridgeData task family.

Figure 4: Support-safe variance diagnostic. Points below the diagonal have empirical variance below the plug-in bound estimate. These plots diagnose the realized variance frontier and do not certify the requested budget.

Why the comparisons are mechanistic. Table 4 summarizes what each comparison removes. The goal is to isolate the role of support preservation, hybrid modes, and port-Hamiltonian structure, not to claim broad task-success dominance over all latent-dynamics model families.

Mechanism-isolation controls. External latent-dynamics and contact-dynamics model families differ in architecture, feature extraction, filtering objective, and training protocol. For that reason, the main comparisons use matched ablations that change one mechanism at a time while holding the remaining implementation fixed. This design isolates proposal-support repair; it should not be read as a broad architecture-ranking benchmark.

Scope of external comparisons. ContactNets-style methods identify contact dynamics from fully observed state transitions. These methods differ from latent filtering under occlusion. We compare overlapping objectives through rSLDS segmentation and Hybrid-SINDy-style sparse recovery.

Table 4: **Mechanistic comparison scope.** Each comparison removes one component of the proposed chain.

Comparison	Removed mechanism	Tested claim
$\lambda = 0$	Support preservation	Support mismatch causes IW/FIVO instability
High-λ	Proposal sharpness	Conservative support mass gives variance margin
Smooth latent	Hybrid modes and support mixing	A single smooth state cannot preserve contact branches
No support	Defensive proposal	Modes alone are insufficient under support mismatch
No mode	Discrete contact state	Predictive regime segmentation uses a hybrid state
No sparsity	Sparse support selection	Dense pH fits obscure active laws
No-pH	Interpretable pH coefficients	Prediction is not mechanism recovery
HMM	Predictive latent dynamics	Post-hoc segmentation is not a predictive filtering model.

Certificate diagnostic. Figure 4 compares empirical estimator relative variance with the plug-in bound estimate. The finite-sample plug-in estimate is not a proven upper bound on ρ_t, so the plotted frontier does not certify the requested variance budget under Theorem 1. The adaptive proposal tracks the bound more tightly, using less defensive mass while preserving low likelihood error in the main paper. This is the expected tradeoff in Theorem 1: larger support mass buys variance margin; adaptive support mass buys sharper filtering.

Table 5: **Claim boundary.** The paper evaluates intermediate mechanisms, not broad task-success dominance.

Claim tested	Evidence reported	Not claimed
Support-safe filtering controls one-step IW/FIVO weight diagnostics under occlusion	ESS/N, relative variance, NLL, coverage audits	General calibration dominance
Hybrid state forms coherent contact-regime segments	ARI, change-point F1, Mode F1 vs. kinematic proxy labels	Simulator contact-force ground truth
Mode-pure segments enable sparse-law recovery	Controlled systems with known sparse laws	Real-robot system identification

Empirical support-mass rule. The plug-in estimate $\widehat{\rho}_t = m_2/m_1^2$ uses $N = 64$ carrier draws. Here m_1 and m_2 are the empirical first and second likelihood moments. This finite-sample estimate is not a proven upper bound on ρ_t, so the empirical rule and its variance-frontier plots do not provide the conditional certificate in Theorem 1. Under the theorem's upper-certificate assumption, the rule based on Eq. (5) selects the smallest certifiable carrier mass or the uncertified fallback λ_{fb}. The empirical implementation uses $\lambda_t = \min(\max(\widehat{\rho}_t/(1 + N\tau^2), 0.02), 1)$, with a floor of 0.02 and upper saturation for the fallback. The tolerance and fallback/floor constants are fixed once on validation and shared across tasks and data families. The conservative operating point fixes $\lambda = 0.50$.

Jensen's inequality gives $\widehat{\rho}_t \geq 1$. With $\tau^2 = 0.25$ and $N\tau^2 = 16$, the budgeted branch gives $\lambda_t \geq 1/17 \approx 0.059$. The tested floors $0.005, 0.02, 0.05$ therefore cannot activate. Identical results to three decimal places check implementation consistency. The sweep does not test an active floor. At 0.90 occlusion, the plug-in carrier-mass demand exceeds 1 on 22% of steps, rising to 37% under severe carrier bias. These steps saturate at $\lambda_t = 1$ and are reported as uncertified.

Branch activation at 90% *occlusion,* 20 *seeds.*

Carrier condition	Uncertified upper fallback	Lower floor
Clean carrier	22% of steps	Inactive
Severe mean bias	37% of steps	Inactive

Sensitivity to carrier mass. NLL decreases smoothly and monotonically in the fixed-λ sweep. NLL decreases from 2.32 to 0.08 on the Contact-style family and from 1.93 to 0.02 on Sawyer as λ increases from 0.05 to 0.70. The plug-in bound is 3.17, versus empirical relative variance 0.65, at $\lambda = 0.05$ and tightens to 0.211 versus 0.181 at $\lambda = 0.70$.

At 90% occlusion, larger λ further improves performance. The frozen default is conservative but is not optimal at this occlusion level.

Full occlusion diagnostics. Figure 5 exposes where support mismatch enters. At easy occlusion, non-defensive proposals can be competitive. At high occlusion, the IW/FIVO log-normalizer variance and calibration curves separate sharply: support-safe variants retain usable coverage, while $\lambda = 0$ and the smooth latent baseline lose calibration as the hidden contact branch becomes ambiguous.

Low-occlusion calibration audit. The adaptive support-safe variant is not uniformly calibrated across the sweep. At 0–25% occlusion, it under-covers: Cov90 is 0.162–0.222 on the Contact-style family and 0.198–0.266 on the Sawyer/BridgeData task family. This failure occurs in the easy-observation regime where the adaptive proposal remains sharp. The high-λ variant preserves broader uncertainty and gives higher coverage, but at the cost of a less targeted proposal. Accordingly, the paper claims robustness to support ambiguity at high occlusion, not uniform calibration dominance.

λ-selection audit. The adaptive rule uses a finite-sample plug-in estimate rather than a proven upper bound on ρ_t. Its realized variance frontier therefore does not provide the conditional certificate in Theorem 1. The same realized mixture is used in the proposal and the importance weight. When the plug-in carrier-mass demand exceeds 1, the empirical rule saturates at the fixed fallback $\lambda_{\mathrm{fb}} = 1$. These steps are reported as uncertified. The empirical rule provides no conditional certificate when the fallback is inactive either. The high-λ variant is included only to show the variance margin obtained when the support budget is made deliberately conservative.

Clean-carrier comparison. Under heavy masking and proposal stress, the amortized proposal may assign vanishing probability to feasible contact transitions. The carrier mixture retains these transitions, allowing mode recovery when observations return. At $r = 0.9$, NLL/Cov90 is $0.66/0.72$ for the adaptive point and $0.47/0.75$ for the conservative point, versus $10.09/0.43$ for $\lambda = 0$. The

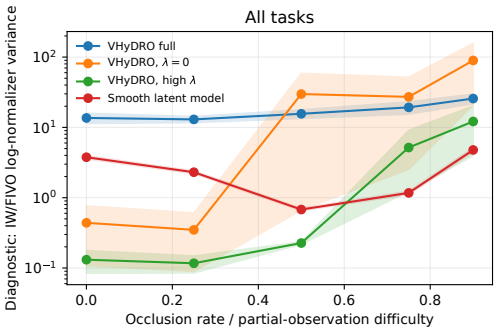

(a) Contact-style: IW log-normalizer variance.

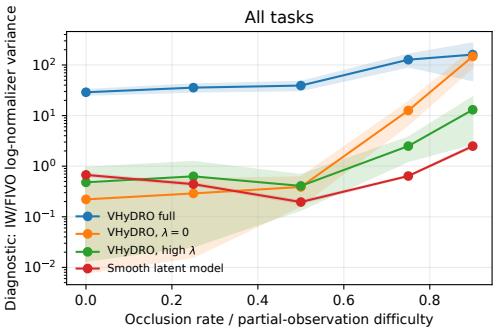

(b) Sawyer/BridgeData task families: IW log-normalizer variance.

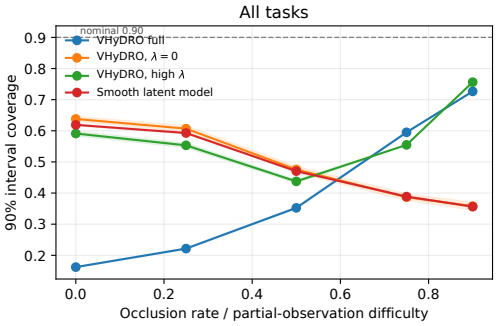

(c) Contact-style: 90% coverage.

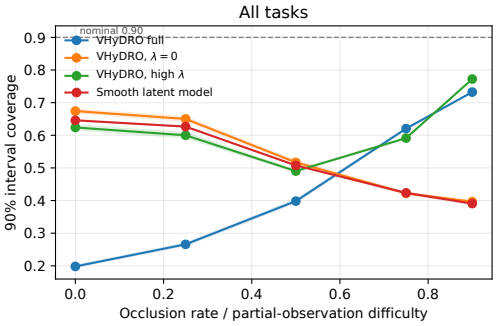

(d) Sawyer/BridgeData task families: 90% coverage.

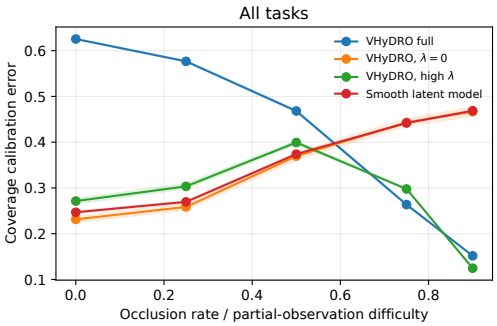

(e) Contact-style: ECE.

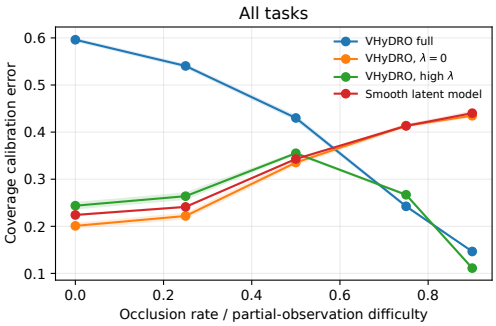

(f) Sawyer/BridgeData task families: ECE.

Figure 5: Full occlusion diagnostics. The informative regime is the high-occlusion end of the sweep. Support-safe proposals keep uncertainty usable where non-defensive and smooth proposals under-cover and incur unstable normalizer estimates.

architecture and carrier are identical. Varying λ solely yields an approximately 15-fold adaptive NLL improvement.

Carrier misspecification. We corrupt the trained carrier at evaluation time consistently in sampling, density evaluation, importance weighting, and the ρ_t plug-in estimate. Variants share the corrupted carrier and differ only in λ. The bias and branch-deletion comparisons use 320 generated episodes per task of 100 transitions, 25 training epochs, and 48 fresh evaluation episodes per task. Perturbations delete a contact-related mixture component, bias carrier means at $0.35/0.7/1.4$, or reduce carrier log standard deviations by $0.5/1.0/1.5$. From a clean-carrier NLL of 0.59, increasing bias gives $0.69/1.24/4.50$ and narrowing gives $0.98/1.66/2.14$. Coverage decreases through 0.73, 0.59, and 0.39. Branch deletion raises NLL to 13.7, approximately 23-fold on the Contact-style family and 28-fold on Sawyer.

Both monitors are fixed before evaluation. The plug-in demand in Eq. (5) is most sensitive to bias. Required carrier mass increases from 0.71 to 1.20 and the uncertified-step rate increases from 22% to 37%. The evidence monitor thresholds prior-predictive log evidence at the fifth percentile of clean validation and is most sensitive to narrowing. Step-level alarms increase from 5.0% to 7.0–22.7%. Episode-level alarms increase from 0% to 24% under severe bias and 11.5% under severe narrowing.

Under branch deletion, wrong-mode rates increase from 3.2% to 6.1%, including high-confidence errors, while step alarms increase only from 5.0% to 7.8%. Neither the pre-sampling plug-in demand nor the evidence alarm reliably identifies this failure at individual steps. Accumulated predictive NLL and evidence collapse reveal the mismatch after observations arrive, but support mixing cannot recover a branch absent from P_t.

Across corruptions, realized weight variance remains consistent with the plug-in bound evaluated against the corrupted P_t. This consistency is relative to the carrier and does not establish agreement with the data-generating process. Predictive monitors assess that discrepancy subject to the omitted-branch limitation.

Additional proposal controls. The controls use the same trained models and data, with inference-time differences only. Unless specified otherwise, evaluations use 20 seeds at 90% occlusion on the Contact-style family. Sawyer/BridgeData-style results follow the same ordering.

We compare the model-feasible carrier with weight clipping, proposal variance inflation, and a generic mixture with a mode-agnostic broad-Gaussian carrier. The generic mixture uses either the same budgeted λ-selection rule or the model carrier's exact realized λ schedule. Each control retains target transition P_t and weights samples using its actual proposal density. Weight clipping truncates each weight at $\sqrt{N}$ times the sample mean before normalization, so its normalizer and NLL estimates are biased. Variance inflation multiplies proposal variances by 2 or 4. The generic carrier is a diagonal Gaussian with the training next-state mean and coordinatewise standard deviations $2\widehat{\sigma}_j + 10^{-4}$, where $\widehat{\sigma}_j$ is the training next-state standard deviation.

Contact-style family, 90% *occlusion,* 20 *seeds.*

Control	ESS/N	NLL	Cov90
Model carrier, conservative	0.132	0.47	0.75
Model carrier, adaptive	0.059	0.66	0.72
Weight clipping	0.076	10.89	0.53
Variance inflation 2× / 4×	0.048 / 0.051	7.31 / 5.34	0.50 / 0.56
Generic mixture, same rule	0.021	125.8	0.12
Generic mixture, matched schedule	0.036	37.6	0.33
No defense ($\lambda = 0$)	0.045	10.09	0.43

Generic controls have NLL at least 5.3 and Cov90 at most 0.56. Both model-carrier points achieve NLL at most 0.66 and Cov90 at least 0.72. On Sawyer, NLL/Cov90 is $0.29/0.76$ for the carrier versus $32.7/0.35$ for the best generic mixture. Matching defensive mass and its schedule does not recover the model-carrier performance. Mass placed away from feasible branches increases the proposal density without assigning sufficient probability to the relevant target branches.

Full numeric audit. Tables 6–9 give the complete sweep and per-task hard-occlusion results. The pattern is deliberately not claimed as uniform dominance at low occlusion. The evidence appears

when observations hide the contact branch: at 90% occlusion, support-safe filtering increases ESS/N, reduces empirical relative variance, and prevents the NLL and coverage collapse seen for $\lambda = 0$ and the smooth latent baseline.

Table 6: **Experiment 1A: contact-style occlusion sweep.** Full occlusion sweep for the contact-style family. Bold entries mark the best value within each occlusion block and metric. Means±SEM over twenty seeds.

Occ.	Method	ESS/N ↑	Rel. w-var ↓	Emp. rel-var ↓	NLL ↓	ECE ↓	Cov90 ↑
0.00	VHYDRO full	0.018±0.000	56.407±0.060	0.881±0.001	1.833±0.119	0.625±0.001	0.162±0.001
0.00	VHYDRO, high λ	0.046±0.001	29.642±0.544	0.463±0.009	**0.071±0.014**	0.271±0.006	0.591±0.007
0.00	VHYDRO, $\lambda = 0$	**0.057±0.001**	**25.858±0.594**	**0.404±0.009**	0.098±0.066	**0.231±0.007**	**0.638±0.008**
0.00	Smooth latent	0.057±0.001	26.808±0.240	0.419±0.004	1.193±0.064	0.247±0.002	0.619±0.003
0.25	VHYDRO full	0.020±0.000	53.675±0.088	0.839±0.001	1.427±0.098	0.577±0.002	0.222±0.002
0.25	VHYDRO, high λ	0.042±0.001	32.223±0.557	0.503±0.009	**0.213±0.013**	0.303±0.006	0.553±0.007
0.25	VHYDRO, $\lambda = 0$	0.054±0.001	**27.275±0.658**	**0.426±0.010**	0.275±0.070	**0.258±0.007**	**0.607±0.009**
0.25	Smooth latent	**0.054±0.001**	27.986±0.264	0.437±0.004	0.999±0.039	0.270±0.002	0.593±0.002
0.50	VHYDRO full	0.025±0.000	46.771±0.167	0.731±0.003	0.936±0.068	0.468±0.002	0.352±0.003
0.50	VHYDRO, high λ	0.032±0.001	40.113±0.426	0.627±0.007	**0.551±0.014**	0.399±0.004	0.437±0.005
0.50	VHYDRO, $\lambda = 0$	0.045±0.001	31.287±0.752	0.489±0.012	1.103±0.062	**0.369±0.008**	**0.476±0.010**
0.50	Smooth latent	**0.047±0.000**	**30.804±0.259**	**0.481±0.004**	1.434±0.023	0.374±0.002	0.470±0.002
0.75	VHYDRO full	0.044±0.000	31.086±0.142	0.486±0.002	**0.377±0.027**	**0.264±0.002**	**0.595±0.002**
0.75	VHYDRO, high λ	**0.056±0.000**	33.022±0.058	0.516±0.001	0.410±0.020	0.298±0.001	0.555±0.001
0.75	VHYDRO, $\lambda = 0$	0.042±0.002	33.068±0.957	0.517±0.015	4.299±0.069	0.443±0.009	0.387±0.011
0.75	Smooth latent	0.046±0.000	**31.015±0.211**	**0.485±0.003**	4.016±0.030	0.442±0.002	0.389±0.002
0.90	VHYDRO full	0.059±0.000	21.057±0.020	0.329±0.000	0.186±0.019	0.152±0.001	0.727±0.001
0.90	VHYDRO, high λ	**0.141±0.001**	**15.851±0.107**	**0.248±0.002**	**0.161±0.023**	**0.125±0.001**	**0.756±0.001**
0.90	VHYDRO, $\lambda = 0$	0.042±0.002	33.195±1.128	0.519±0.018	11.286±0.151	0.467±0.010	0.358±0.012
0.90	Smooth latent	0.047±0.001	30.366±0.290	0.474±0.005	9.660±0.098	0.469±0.001	0.356±0.001

Table 7: **Experiment 1B: Sawyer/Bridge occlusion sweep.** Full occlusion sweep for the Sawyer/Bridge family. Bold entries mark the best value within each occlusion block and metric. Means±SEM over twenty seeds.

Occ.	Method	ESS/N ↑	Rel. w-var ↓	Emp. rel-var ↓	NLL ↓	ECE ↓	Cov90 ↑
0.00	VHYDRO full	0.019±0.000	54.593±0.200	0.853±0.003	1.945±0.065	0.596±0.003	0.198±0.003
0.00	VHYDRO, high λ	0.052±0.001	26.870±0.594	0.420±0.009	-0.086±0.028	0.244±0.006	0.624±0.008
0.00	VHYDRO, $\lambda = 0$	**0.067±0.002**	**22.779±0.494**	**0.356±0.008**	**-0.115±0.017**	**0.201±0.005**	**0.674±0.006**
0.00	Smooth latent	0.065±0.000	24.502±0.143	0.383±0.002	0.328±0.028	0.224±0.001	0.646±0.002
0.25	VHYDRO full	0.022±0.000	51.286±0.282	0.801±0.004	1.595±0.064	0.540±0.003	0.266±0.004
0.25	VHYDRO, high λ	0.049±0.001	28.631±0.626	0.447±0.010	0.026±0.028	0.264±0.007	0.600±0.008
0.25	VHYDRO, $\lambda = 0$	**0.064±0.002**	**23.852±0.547**	**0.373±0.009**	**0.013±0.024**	**0.222±0.005**	**0.650±0.006**
0.25	Smooth latent	0.062±0.000	25.418±0.136	0.397±0.002	0.370±0.025	0.241±0.001	0.627±0.001
0.50	VHYDRO full	0.028±0.000	43.903±0.319	0.686±0.005	1.067±0.052	0.430±0.004	0.398±0.005
0.50	VHYDRO, high λ	0.037±0.001	36.249±0.389	0.566±0.006	**0.354±0.029**	0.355±0.002	0.490±0.003
0.50	VHYDRO, $\lambda = 0$	0.051±0.001	**28.336±0.535**	**0.443±0.008**	0.707±0.029	**0.335±0.004**	**0.517±0.005**
0.50	Smooth latent	**0.053±0.000**	28.455±0.161	0.445±0.003	0.949±0.019	0.343±0.001	0.508±0.001
0.75	VHYDRO full	0.047±0.000	29.046±0.233	0.454±0.004	0.646±0.123	**0.242±0.002**	**0.621±0.002**
0.75	VHYDRO, high λ	**0.062±0.001**	30.229±0.163	0.472±0.003	**0.309±0.020**	0.267±0.002	0.591±0.002
0.75	VHYDRO, $\lambda = 0$	0.047±0.001	30.110±0.511	0.470±0.008	3.426±0.054	0.414±0.003	0.422±0.003
0.75	Smooth latent	0.051±0.000	**28.715±0.037**	**0.449±0.001**	3.246±0.027	0.413±0.002	0.423±0.002
0.90	VHYDRO full	0.060±0.000	20.432±0.043	0.319±0.001	0.369±0.097	0.147±0.000	0.732±0.001
0.90	VHYDRO, high λ	**0.151±0.002**	**14.355±0.071**	**0.224±0.001**	**0.074±0.019**	**0.111±0.001**	**0.772±0.001**
0.90	VHYDRO, $\lambda = 0$	0.046±0.001	30.205±0.325	0.472±0.005	9.401±0.067	0.435±0.002	0.397±0.003
0.90	Smooth latent	0.052±0.000	28.050±0.120	0.438±0.002	8.095±0.016	0.440±0.004	0.391±0.004

Table 8: **Experiment 1A: contact-style hard-occlusion audit.** Per-task hard-occlusion results for the contact-style family at 90% occlusion. Bold entries mark the best value within each task block and metric. Means over twenty seeds.

Task	Method	ESS/N ↑	Rel. w-var ↓	Emp. rel-var ↓	NLL ↓	ECE ↓	Cov90 ↑
PickCube-v1	VHYDRO full	0.059	21.140	0.330	**0.162**	0.154	0.724
PickCube-v1	VHYDRO, high λ	**0.141**	**15.883**	**0.248**	0.167	**0.125**	**0.756**
PickCube-v1	VHYDRO, $\lambda = 0$	0.041	33.337	0.521	11.444	0.468	0.357
PickCube-v1	Smooth latent	0.047	30.444	0.476	9.712	0.469	0.356
PokeCube-v1	VHYDRO full	0.060	21.027	0.329	0.168	0.150	0.728
PokeCube-v1	VHYDRO, high λ	**0.141**	**15.881**	**0.248**	**0.148**	**0.125**	**0.756**
PokeCube-v1	VHYDRO, $\lambda = 0$	0.042	33.164	0.518	11.189	0.467	0.359
PokeCube-v1	Smooth latent	0.047	30.374	0.475	9.644	0.469	0.356
PullCube-v1	VHYDRO full	0.059	21.087	0.329	**0.158**	0.152	0.727
PullCube-v1	VHYDRO, high λ	**0.142**	**15.763**	**0.246**	0.161	**0.125**	**0.756**
PullCube-v1	VHYDRO, $\lambda = 0$	0.042	33.179	0.518	11.262	0.466	0.358
PullCube-v1	Smooth latent	0.047	30.345	0.474	9.620	0.469	0.356
PushCube-v1	VHYDRO full	0.060	20.971	0.328	0.255	0.151	0.727
PushCube-v1	VHYDRO, high λ	**0.141**	**15.874**	**0.248**	**0.167**	**0.124**	**0.758**
PushCube-v1	VHYDRO, $\lambda = 0$	0.042	33.106	0.517	11.246	0.467	0.358
PushCube-v1	Smooth latent	0.048	30.292	0.473	9.662	0.468	0.356

Table 9: **Experiment 1B: Sawyer/Bridge hard-occlusion audit.** Per-task hard-occlusion results for the Sawyer/Bridge family at 90% occlusion. Bold entries mark the best value within each task block and metric. Means over twenty seeds.

Task	Method	ESS/N ↑	Rel. w-var ↓	Emp. rel-var ↓	NLL ↓	ECE ↓	Cov90 ↑
BridgeDataOcclusion	VHYDRO full	0.058	21.814	0.341	0.351	0.156	0.719
BridgeDataOcclusion	VHYDRO, high λ	**0.130**	**17.182**	**0.268**	**0.312**	**0.138**	**0.743**
BridgeDataOcclusion	VHYDRO, $\lambda = 0$	0.046	30.752	0.481	10.950	0.448	0.381
BridgeDataOcclusion	Smooth latent	0.051	28.568	0.446	9.765	0.452	0.376
SawyerLaundryLayout	VHYDRO full	0.061	20.154	0.315	0.297	0.144	0.736
SawyerLaundryLayout	VHYDRO, high λ	**0.152**	**14.085**	**0.220**	**0.050**	**0.107**	**0.777**
SawyerLaundryLayout	VHYDRO, $\lambda = 0$	0.048	29.461	0.460	8.683	0.425	0.408
SawyerLaundryLayout	Smooth latent	0.054	27.525	0.430	7.489	0.436	0.396
SawyerObjectSearch	VHYDRO full	0.060	20.577	0.322	0.308	0.145	0.734
SawyerObjectSearch	VHYDRO, high λ	**0.143**	**15.027**	**0.235**	**0.109**	**0.116**	**0.766**
SawyerObjectSearch	VHYDRO, $\lambda = 0$	0.046	30.422	0.475	9.918	0.440	0.390
SawyerObjectSearch	Smooth latent	0.052	28.425	0.444	8.752	0.447	0.383
SawyerTowerCreation	VHYDRO full	0.062	19.372	0.303	0.527	0.143	0.739
SawyerTowerCreation	VHYDRO, high λ	**0.177**	**11.585**	**0.181**	**-0.141**	**0.089**	**0.798**
SawyerTowerCreation	VHYDRO, $\lambda = 0$	0.046	30.244	0.473	8.253	0.429	0.404
SawyerTowerCreation	Smooth latent	0.053	27.689	0.433	6.598	0.427	0.406

A.2 Experiment 2: ManiSkill contact-regime segmentation audit with proxy labels

Protocol. Experiment 2 uses released ManiSkill HDF5 demonstrations from PushCube-v1, PokeCube-v1, PullCubeTool-v1, and PegInsertionSide-v1. The evaluation labels are denoised kinematic proxies derived from raw HDF5 states and actions. They distinguish free motion, impact-like transitions, and stick/slip-like contact using smoothed object/end-effector kinematic changes. They are used for validation diagnostics and held-out evaluation only; the model is not trained with simulator contact-force labels. Hyperparameters are selected on validation episodes, and all reported metrics are computed on held-out test episodes.

Kinematic proxy-label construction. For each trajectory, the deterministic proxy uses only recorded state/action fields and is shared by all methods. Let p_t^{obj} and p_t^{ee} denote object and end-effector positions, and let a_t denote the action. We compute normalized finite-difference terms

$$r_t^{\mathrm{obj}} = \frac{\|\Delta p_t^{\mathrm{obj}}\|_2}{\mathrm{MAD}(\|\Delta p^{\mathrm{obj}}\|_2) + \epsilon}, \qquad r_t^{\mathrm{ee}} = \frac{\|\Delta p_t^{\mathrm{ee}}\|_2}{\mathrm{MAD}(\|\Delta p^{\mathrm{ee}}\|_2) + \epsilon}, \qquad r_t^{a} = \frac{\|\Delta a_t\|_2}{\mathrm{MAD}(\|\Delta a\|_2) + \epsilon}.$$

The scalar score is

$$c_t = \alpha_{\mathrm{obj}} r_t^{\mathrm{obj}} + \alpha_{\mathrm{ee}} r_t^{\mathrm{ee}} + \alpha_a r_t^{a}.$$

The weights α, smoothing window, minimum-run denoising length, and thresholds $\theta_1 < \theta_2$ are selected on validation episodes and then frozen. After smoothing, labels are assigned as free if $c_t < \theta_1$, impact-like if $c_t \geq \theta_2$, and stick/slip-like otherwise, followed by the fixed minimum-run denoising. Clipping at the 99.5th percentile is used only for plotting histograms, not for metric computation. The test split is never used to choose thresholds. Proxy labels are used only for metric computation after label-permutation matching; they are not provided to the VHYDRO mode posterior during training.

External switching-state baseline. We evaluate the recurrent switching linear dynamical system (rSLDS) of Linderman et al., implemented in `ssm`, on PushCube-v1 and PegInsertionSide-v1 using 20 seeds at the default $(K, D) = (3, 6)$. Episode splits and standardization follow Experiment 2. Averaging the two tasks within each seed, the mean$\pm$sample SD across 20 seeds is 0.506 ± 0.079 for Mode F1, 0.250 ± 0.110 for ARI, and 0.313 ± 0.055 for change-point F1. All 40 task and seed cells complete under fixed budgets of 25 initialization, 100 training, and 50 inference iterations. Optimizer convergence is not established. The configuration and implementation diagnostics below are separate from this evaluation. The configuration study requests $K \in \{3, 4, 6\}$ modes and latent-dimension settings $D \in \{6, 10\}$, with $K = 3$ matching the labeled contact modes. The configurations labeled $D = 10$, $K = 4$, and $K = 6$ use an autoregressive-HMM fallback and therefore do not establish rSLDS sensitivity to these settings. The fallback configuration labeled $(K, D) = (3, 10)$, selected post hoc on the final evaluation metric, has macro Mode F1/ARI/change-point F1 of $0.620/0.423/0.452$, compared with $0.779/0.650/0.674$ for VHYDRO-full and $0.763/0.577/0.605$ for HMM. Separate completion diagnostics for the fallback configurations cover seed identifiers 0–4 and report $10/10$ completed task and seed cells per configuration. The genuine default rSLDS initially completes 7 of 10 cells for these seed identifiers, with 3 failed cells. Its completed cells have mean Mode F1 0.460, below the completed fallback configurations, and segment purity 0.93, compared with 0.92 for VHYDRO-full. The no-support ablation has Mode F1 0.661. These external comparisons do not isolate the contribution of defensive support mass.

Robustness to proxy construction. We vary four label-constructor parameters one at a time on the four tasks in Table 2, using 20 seeds, i.e. contact quantile $\{0.50, 0.58, 0.66\}$, smoothing window $\{3, 5, 9\}$, minimum-run length $\{2, 4, 8\}$, and hysteresis fill-gap $\{1, 3, 6\}$. Each of the nine settings re-derives the labels and re-fits every method, with the constructor frozen before held-out evaluation. VHYDRO ARI ranges from 0.28 to 0.60 and ranks first on ARI and change-point F1 in all $9/9$ settings. At smoothing window 3, the ARI margin narrows to 0.278 versus 0.270 for HMM. The only ordering change is a swap between HMM and no-support in second and third place at the largest smoothing window.

Mode confidence and transition errors. A separate 20-seed audit on real PushCube and PegInsertionSide demonstrations evaluates mode confidence against kinematic proxy labels. Its smaller-data protocol preserves the main segmentation ranking, with different absolute values. Mode F1 is 0.779

for VHYDRO-full, 0.763 for HMM, 0.661 for no-support, and 0.220 for no-mode. Mean maximum posterior probability is 0.861 on incorrect frames and 0.959 on correct frames. The confidently-wrong rate is 0.69 for VHYDRO-full and 0.98 for HMM. Of the full model's errors, 54% occur within ± 2 frames of a true transition. The confidence gate at 0.60 down-weights these boundary errors before sparse fitting.

Validation against simulator contact forces. We replay 40 held-out PushCube episodes and derive regime labels from simulator contact reports. Contact requires an impulse exceeding 2% of the object's resting support impulse. The threshold is frozen on validation data, with proxy labels and predictions unchanged. Proxy-to-force agreement is moderate: Mode F1 0.55, ARI 0.31, and change-point F1 0.41. Proxy contact segments cover 52% of steps, compared with 31% for force labels, extending mainly across approach and release. Force-label evaluation retains the ranking in Table 2 on all three metrics. Mode F1/ARI/change-point F1 are $0.58/0.47/0.40$ for VHYDRO, $0.55/0.45/0.35$ for HMM, $0.45/0.24/0.19$ for no-support, and $0.27/0.00/0.00$ for no-mode. SEM is at most 0.03.

Mode-timeline diagnostic. Figure 6 shows representative trajectories. VHYDRO produces contiguous mode segments that align with the denoised kinematic proxy labels. The no-support ablation keeps a discrete state but fragments transitions more often. The no-mode baseline cannot represent regime switches. The HMM is a useful segmentation baseline, but it does not share a predictive latent state with the dynamics model.

Proxy-label diagnostic. Figure 7 makes the evaluation labels auditable. The score distributions are not used as supervised contact labels for VHYDRO. They provide a reproducible proxy against which mode F1, ARI, change-point F1, and segment purity can be computed.

Per-task contact-regime segmentation. Table 10 reports held-out per-task metrics. VHYDRO is best on ARI and change-point F1 on every task, and best on purity in three of four tasks. The higher HMM Mode F1 on PushCube-v1, and the higher no-support purity on PokeCube-v1, do not carry over to the full predictive segmentation profile. This distinction matters for Theorem 2: the recovered state must be a coherent predictive mode, not only a post-hoc label sequence.

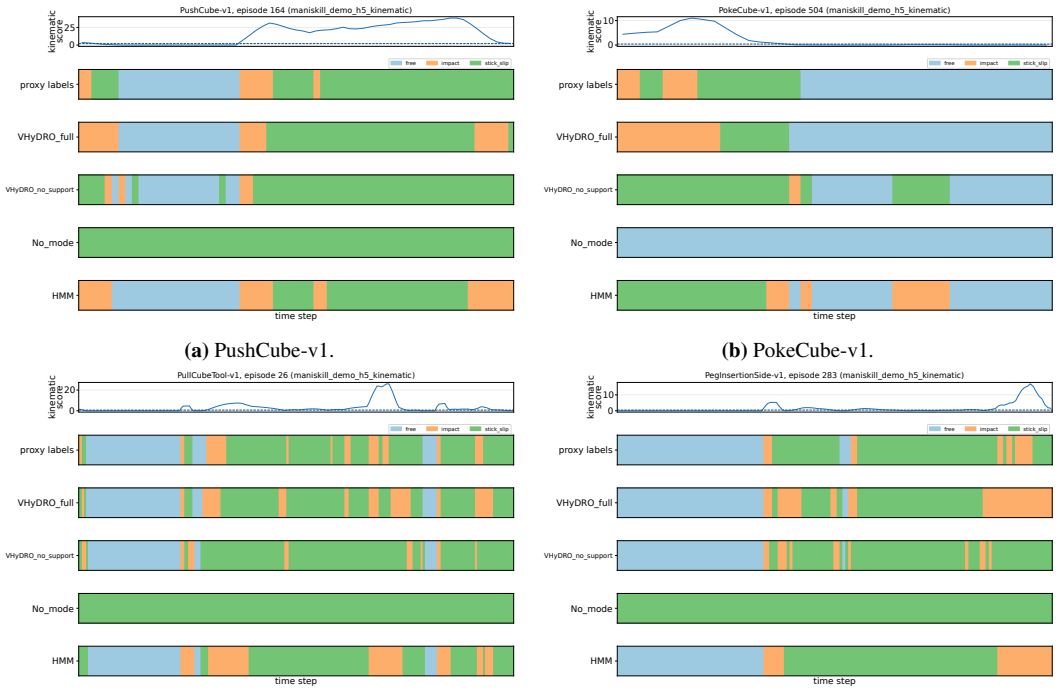

(a) PushCube-v1. **(b)** PokeCube-v1.

(c) PullCubeTool-v1. **(d)** PegInsertionSide-v1.

Figure 6: Representative held-out mode timelines. VHYDRO infers temporally coherent mode segments that align with denoised kinematic proxy labels. Alignment is measured against denoised kinematic proxy labels; no simulator contact-force labels are used.

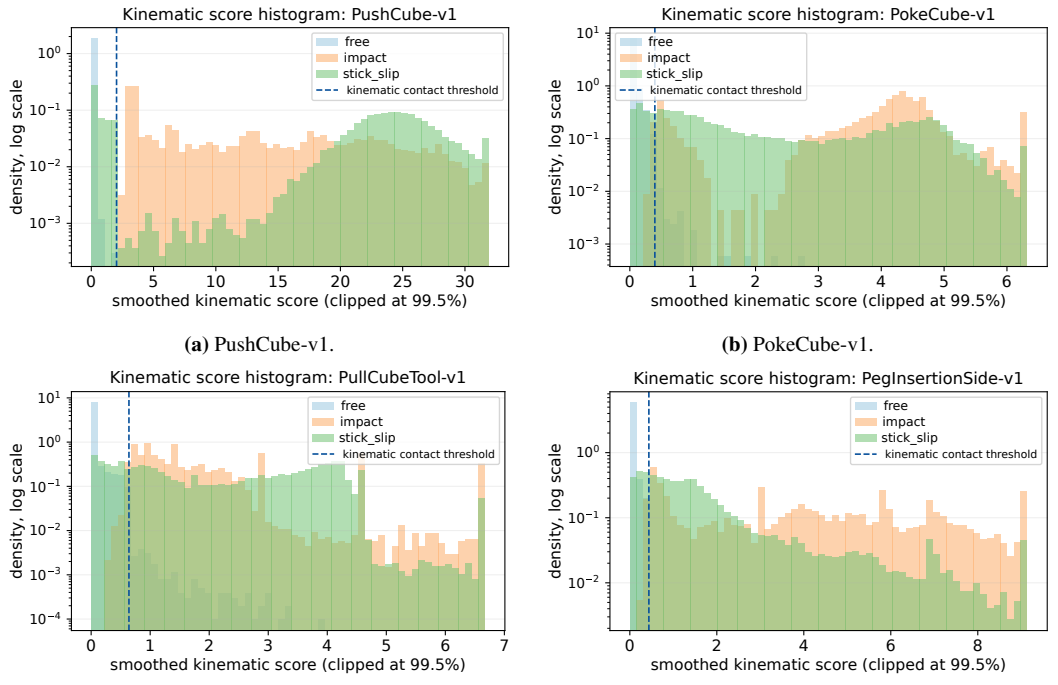

(a) PushCube-v1. **(b)** PokeCube-v1.

(c) PullCubeTool-v1. **(d)** PegInsertionSide-v1.

Figure 7: Kinematic proxy-label construction. Histograms of the smoothed kinematic score used to form denoised free/impact/stick-slip proxy labels from HDF5 states and actions. The labels audit mode segmentation; they are not simulator contact-force labels.

Table 10: **Experiment 2: per-task proxy-label contact-regime segmentation.** Held-out test split across ManiSkill HDF5 demonstrations. Labels are denoised kinematic proxies, not simulator contact-force labels. Means±SEM over twenty seeds.

Task	Method	Mode F1 ↑	ARI ↑	Change-point F1 ↑	Purity ↑
PushCube-v1	VHYDRO	0.810±0.004	**0.714±0.015**	**0.716±0.007**	**0.915±0.002**
PushCube-v1	no support	0.635±0.040	0.536±0.023	0.371±0.028	0.901±0.008
PushCube-v1	HMM	**0.828±0.008**	0.647±0.054	0.664±0.016	0.907±0.013
PushCube-v1	no mode	0.218±0.002	0.000±0.000	0.000±0.000	0.485±0.006
PokeCube-v1	VHYDRO	**0.583±0.006**	**0.599±0.014**	**0.557±0.013**	0.886±0.001
PokeCube-v1	no support	0.558±0.007	0.398±0.016	0.420±0.029	**0.914±0.007**
PokeCube-v1	HMM	0.523±0.008	0.403±0.018	0.409±0.006	0.897±0.004
PokeCube-v1	no mode	0.229±0.001	0.000±0.000	0.000±0.000	0.524±0.005
PullCubeTool-v1	VHYDRO	**0.770±0.002**	**0.697±0.006**	**0.698±0.002**	**0.942±0.001**
PullCubeTool-v1	no support	0.648±0.006	0.495±0.021	0.509±0.001	0.915±0.008
PullCubeTool-v1	HMM	0.672±0.009	0.488±0.003	0.530±0.022	0.876±0.006
PullCubeTool-v1	no mode	0.217±0.002	0.000±0.000	0.000±0.000	0.482±0.005
PegInsertionSide-v1	VHYDRO	**0.732±0.003**	**0.578±0.009**	**0.637±0.002**	**0.934±0.000**
PegInsertionSide-v1	no support	0.607±0.022	0.430±0.012	0.431±0.026	0.932±0.009
PegInsertionSide-v1	HMM	0.678±0.006	0.422±0.013	0.504±0.005	0.879±0.002
PegInsertionSide-v1	no mode	0.220±0.003	0.000±0.000	0.000±0.000	0.494±0.011

A.3 Sawyer/BridgeData-style contact-regime diagnostics with proxy labels

We run the Experiment 2 contact-regime segmentation protocol on four Sawyer/BridgeData task families: BridgeDataOcclusion, SawyerTowerCreation, SawyerLaundryLayout, and SawyerObject-Search. The evaluation labels are denoised kinematic proxy labels derived from trajectory state and action changes. They separate free motion, impact-like transitions, and sustained interaction, and are used only for validation diagnostics and held-out evaluation.

Table 11: **Additional Sawyer/BridgeData-style predictive mode-segmentation diagnostic.** Mean±SEM over four task families and twenty seeds. Higher is better. Bold entries mark the quantities emphasized by the mode-concentration claim.

Method	Mode F1 ↑	ARI ↑	Change-point F1 ↑	Purity ↑
VHYDRO	**0.708±0.003**	**0.560±0.008**	**0.706±0.004**	0.918±0.002
no support	0.535±0.009	0.171±0.010	0.389±0.004	**0.920±0.003**
HMM	0.470±0.005	0.174±0.004	0.348±0.003	0.882±0.001
no mode	0.222±0.001	0.000±0.000	0.000±0.000	0.498±0.003

Representative timelines. Figure 8 shows representative Sawyer/BridgeData task timelines. VHYDRO produces longer coherent segments than the no-support ablation, which switches rapidly on several trajectories. The no-mode baseline collapses to one regime and therefore cannot represent contact changes.

Proxy-label audit. Figure 9 shows the smoothed kinematic score distributions used to construct the kinematic proxy labels. These histograms document the evaluation target; they are not contact-force histograms and are not used as supervised training labels.

Per-task audit. Table 12 reports task-level metrics for the Sawyer/BridgeData task families using frozen kinematic proxy labels. The aggregate pattern in Table 11 is consistent across tasks: VHYDRO improves the metrics tied to mode concentration, while purity alone is not sufficient to identify coherent predictive modes.

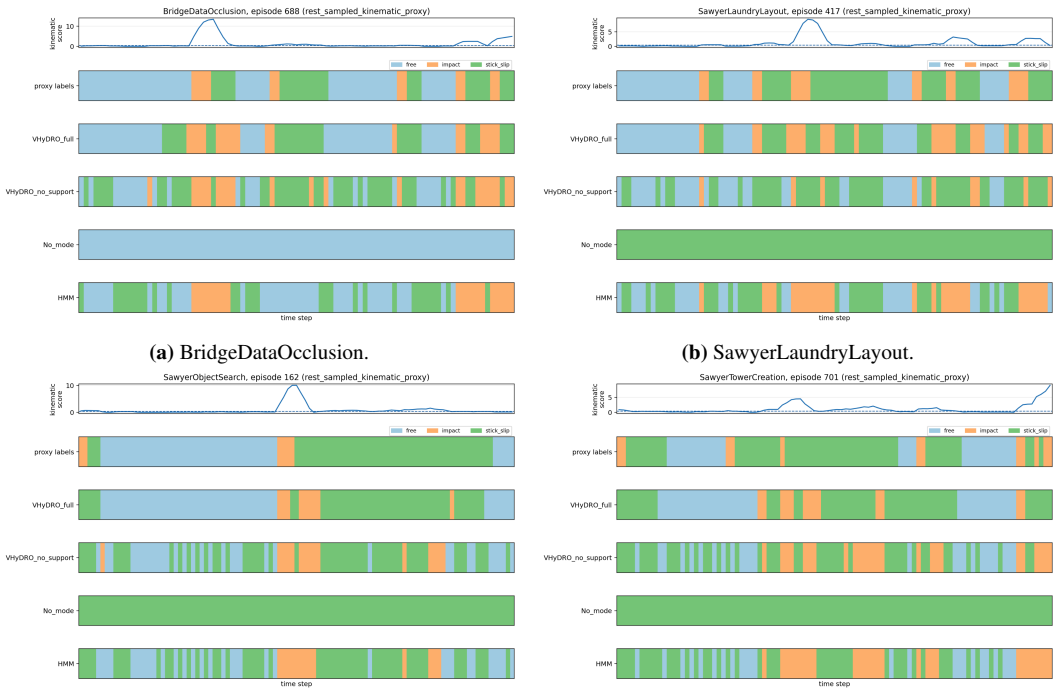

(a) BridgeDataOcclusion. **(b)** SawyerLaundryLayout.

(c) SawyerObjectSearch. **(d)** SawyerTowerCreation.

Figure 8: Sawyer/BridgeData mode timelines with proxy labels. Each panel shows the smoothed kinematic score, denoised kinematic proxy labels, and inferred mode sequences. The labels audit mode segmentation; they are not simulator contact-force labels.

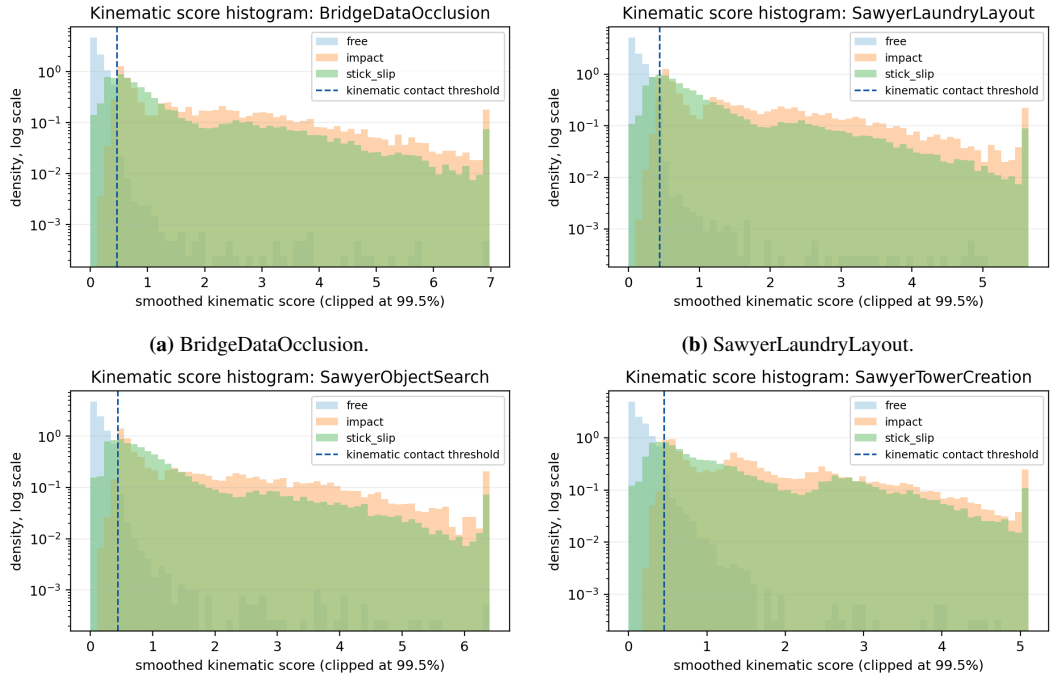

(a) BridgeDataOcclusion. **(b)** SawyerLaundryLayout.

(c) SawyerObjectSearch. **(d)** SawyerTowerCreation.

Figure 9: Sawyer/BridgeData kinematic-score histograms. Smoothed kinematic-score distributions used to form free, impact-like, and stick/slip-like kinematic proxy labels. The labels audit mode segmentation; they are not simulator contact-force labels.

Table 12: **Per-task Sawyer/BridgeData-style predictive mode-segmentation diagnostic.** Held-out test metrics for Sawyer/BridgeData-style proxy trajectories. Labels are denoised kinematic proxies, not simulator contact-force labels. Means±SEM over twenty seeds. Bold entries mark the best predictive segmentation value within each task block.

Task	Method	Mode F1 ↑	ARI ↑	Change-point F1 ↑	Purity ↑
SawyerTowerCreation	VHYDRO	**0.697±0.001**	**0.528±0.004**	**0.697±0.009**	0.908±0.001
SawyerTowerCreation	no support	0.516±0.009	0.165±0.024	0.398±0.002	**0.917±0.009**
SawyerTowerCreation	HMM	0.448±0.014	0.193±0.006	0.363±0.001	0.880±0.002
SawyerTowerCreation	no mode	0.221±0.002	0.000±0.000	0.000±0.000	0.495±0.005
SawyerLaundryLayout	VHYDRO	**0.724±0.005**	**0.592±0.014**	**0.722±0.009**	**0.928±0.002**
SawyerLaundryLayout	no support	0.579±0.009	0.187±0.024	0.403±0.011	0.922±0.005
SawyerLaundryLayout	HMM	0.489±0.001	0.161±0.002	0.346±0.003	0.881±0.003
SawyerLaundryLayout	no mode	0.222±0.000	0.000±0.000	0.000±0.000	0.499±0.001
BridgeDataOcclusion	VHYDRO	**0.707±0.001**	**0.558±0.007**	**0.705±0.004**	0.919±0.001
BridgeDataOcclusion	no support	0.513±0.001	0.192±0.003	0.381±0.000	**0.930±0.002**
BridgeDataOcclusion	HMM	0.472±0.001	0.177±0.003	0.350±0.001	0.886±0.001
BridgeDataOcclusion	no mode	0.220±0.001	0.000±0.000	0.000±0.000	0.492±0.005
SawyerObjectSearch	VHYDRO	**0.704±0.000**	**0.563±0.008**	**0.700±0.004**	**0.915±0.002**
SawyerObjectSearch	no support	0.530±0.010	0.139±0.018	0.375±0.005	0.912±0.007
SawyerObjectSearch	HMM	0.470±0.003	0.165±0.003	0.335±0.003	0.881±0.002
SawyerObjectSearch	no mode	0.224±0.003	0.000±0.000	0.000±0.000	0.505±0.009

Audit reproducibility. The aggregate and per-task tables report seed-level means, task-level means, and standard errors computed from the fixed held-out evaluation protocol.

A.4 Experiment 3: sparse physical-law recovery audit

This appendix audits the sparse-law link in the theory chain: *mode concentration* $\Rightarrow$ *sparse physical-law recovery*. The benchmark uses known sparse port-Hamiltonian laws, so support recovery, coefficient error, vector-field error, and physical-constant error can be measured directly.

Protocol. We evaluate four controlled known-equation hybrid systems with known sparse laws: a bouncing puck with wall contacts, a sliding block with Coulomb friction, a pendulum with contact stops, and a planar pushing system. The same twenty-seed evaluation protocol is applied to every method and includes observation noise, missing data, hidden velocities, numerical derivative noise, and small mode-assignment perturbations. We compare VHYDRO-full, *no-mode*, *no-sparsity*, and *no-pH*. The first three produce port-Hamiltonian coefficients and can therefore be evaluated by sparse support recovery, relative coefficient error, and physical-constant error; the no-pH baseline is evaluated only as a vector-field predictor. The candidate libraries contain at most $p = 12$ terms per system, and the true active support has size at most $k = 3$. These values match the low-dimensional sparse setting used when interpreting Theorem 3.

Summary metrics. Table 13 is the main quantitative audit. VHYDRO is best on every metric that reflects sparse physical-law recovery: support F1, relative coefficient error, vector-field NRMSE, and physical-constant error. The no-mode ablation merges incompatible regimes and degrades both support recovery and coefficient estimation. The no-sparsity ablation predicts more accurately than no-mode, but the dense fit obscures the active support. The no-pH baseline reaches a moderate vector-field error but has no sparse port-Hamiltonian coefficients to inspect.

External sparse hybrid baseline. We adapt Hybrid-SINDy (Mangan et al.) to Experiment 3, sharing the data generator, random-number stream, candidate library, threshold grid, and metrics. The adaptation uses BIC-selected Gaussian-mixture regime discovery and per-cluster STLSQ, with joint AICc selection of structure and threshold and no labels at fit time. The adaptation does not reproduce the original measurement-coordinate clustering pipeline. This is an exploratory comparison because Hybrid-SINDy discovers regimes unsupervised, while our sparse fit uses filtered mode assignments. Over 20 seeds, Hybrid-SINDy obtains support F1 0.814, precision 0.702, relative coefficient error 0.391, and vector-field NRMSE 0.116. The corresponding VHYDRO-full values are 0.976, 0.956, 0.006, and 0.0012. The VHYDRO support-F1 reference differs across matched runs: 0.976 here, 0.981 in the library-sensitivity run, and 0.978 in Table 3. Unsupervised regime discovery achieves ARI 0.54. Assignment errors propagate into the per-regime sparse fits, consistent with the mode-purity requirement of Theorem 3.

Held-out prediction and ensemble uncertainty. We extend the known-equation protocol in Table 3 using the same library, normalization, and regularization grid, 20 seeds, and fixed held-out rollouts. The sparse mode-conditioned law reduces 50-step rollout NRMSE by approximately 11–160 times. The 90% seed-ensemble bands describe variation across seeds, not posterior predictive uncertainty. Their coverage is evaluated at step 50. Cov90 and ECE in Table 1 remain the filtering-calibration evidence.

Additional held-out prediction comparison, 20 *seeds.*

Method	One-step NRMSE	50-step NRMSE	50-step band coverage
VHYDRO	0.001	0.0009	0.40
no-mode	0.301	0.146	0.01
no-sparsity	0.033	0.013	0.21
no-pH	0.032	0.010	0.18

Experiment 5: filtered-state-to-law recovery in planar pushing. We evaluate simulated planar pushing with known mass normalization and mode-wise friction/contact coefficients. The support-safe hybrid filter processes occluded position-only observations using the Experiment 1 support-mass rule. An HMM E-step over fitted per-mode dynamics supplies assignments for the sparse

Table 13: **Experiment 3 summary.** Values are mean±standard deviation over 20 seeds for each known-equation system and fixed evaluation-condition batch, aggregated over the four systems. For the dense no-sparsity baseline, Support F1 is computed after applying the same post-hoc coefficient threshold used for the sparse-support audit. The no-pH model is evaluated only by vector-field prediction because it has no sparse port-Hamiltonian coefficients.

Method	Support F1 ↑	Rel. coeff. error ↓	Vector-field NRMSE ↓	Phys.-constant error ↓
VHYDRO-full	0.978 ± 0.042	0.006 ± 0.010	0.001 ± 0.001	0.001 ± 0.001
no-mode	0.713 ± 0.039	1.087 ± 0.330	0.287 ± 0.196	0.188 ± 0.234
no-sparsity	0.771 ± 0.073	0.851 ± 0.559	0.043 ± 0.043	0.014 ± 0.008
no-pH	N/A	N/A	0.031 ± 0.015	N/A

port-Hamiltonian fit. The HMM replaces the amortized mode posterior. This controlled filtered-state-to-law experiment therefore does not evaluate the full end-to-end VHYDRO pipeline. Fitting uses no oracle segmentation. Labels are permutation-matched to ground truth only for evaluation. The confidence gate at 0.60 precedes the sparse fit, as in Experiment 3.

Over 20 seeds, the dominant inverse-mass coefficient is recovered at 0.984 against a true value of 1.00, with a reported median error of 5%. The remaining kinetic terms have mean errors of 8–32$\%$. Support F1 is 0.78, compared with 1.00 under oracle segmentation. A velocity cross-term with true coefficient 0.35 falls below the sparsity threshold at this noise level. In a stick-slip sliding system, rapid switching degrades derivative estimation and support F1 is 0.60. Appendix A.4 reports the accompanying diagnostics for this experiment and Experiment 3.

Sensitivity to the candidate library. We vary the candidate library under the Experiment 3 protocol using 20 seeds. Here, c denotes the number of added correlated distractor columns. This comparison uses $2{,}500$ training and 600 test trajectories per system, 260 steps per trajectory, and at most $60{,}000$ selected rows per mode. The held-out fraction is 0.20. Correlated distractors reduce VHYDRO-full support F1 from 0.981 to 0.979 at $c = 2$, 0.786 at $c = 4$, and 0.714 at $c = 8$. Precision falls from 0.965 to 0.598. The ridge ablation degrades faster, from 0.912 to 0.586, indicating greater robustness from sparsity without eliminating the failure mode. Conditioning deteriorates while support F1 remains high, i.e. at $c = 2$, F1 remains 0.979, but the smallest eigenvalue on the estimated support, $\widehat{\kappa}$, falls from 4.4×10^{-3} to 10^{-4}. Mutual coherence rises to 0.998, reaching 0.9997 at $c = 4$. These diagnostics use no ground-truth support and do not verify the restricted-eigenvalue assumption in Theorem 3 and Appendix B.3. Removing a true active term improves $\widehat{\kappa}$ to 0.067, while held-out one-step residual NRMSE increases by 3.1 times on average and 1.1–5.5 across systems. The in-library fit is unchanged. Conditioning and coherence flag redundant libraries, while held-out residuals flag missing physical terms.

Sparse-recovery diagnostics and perturbation sensitivity. We measure mode impurity, derivative error, residual port-Hamiltonian error, and design conditioning using the sparse fit's mode and confidence selection rule, with 20 seeds for each of four systems. The diagnostics independently subsample modes exceeding the $80{,}000$-row cap, so their capped subsets need not coincide with the fitted subsets. Each run uses $1{,}200$ training and 300 test trajectories with 220 steps. These controlled diagnostics perturb the known states and derivatives to construct proxy regression inputs. Unperturbed assignments receive confidence in $[0.92, 1]$, while injected flips receive confidence in $[0.05, 0.40]$, so the 0.60 gate removes the injected low-confidence assignments by construction. Post-gate impurity is exactly 0, compared with raw impurity 0.013 under 2% injected label flips. Derivative NRMSE is 0.019–0.027, matching the injected perturbation 0.02. Residual port-Hamiltonian error is 0.02–0.05. For each mode, $G = A^\top A / n_A$, where n_A is the number of scalar regression equations. The reported minimum eigenvalue is $\max\{\lambda_{\min}(G), 0\}$, and the floored condition diagnostic is

$$\widehat{C} = \sqrt{\frac{\max\{\lambda_{\max}(G), 10^{-12}\}}{\max\{\lambda_{\min}(G), 10^{-12}\}}}.$$

Table 14 reports these diagnostics by mode.

A 20-seed perturbation sweep examines the sensitivity described by Theorem 3. At 20% injected label flips, raw impurity rises to 0.133, while post-gate impurity remains 0 and recovery is unchanged.

Table 14: Per-mode design conditioning in the controlled sparse-recovery diagnostic. Entries are means $\pm$ standard deviations over 20 seeds per system, comprising 80 system–seed runs. The condition diagnostic uses the 10^{-12} eigenvalue floor defined in the text. It is not the true condition number of a rank-deficient design.

System	Mode	Minimum eigenvalue	Floored condition diagnostic
Bouncing puck	Free	0	$7.071 \times 10^5 \pm 4.750 \times 10^1$
Bouncing puck	Right wall	0	$1.019 \times 10^6 \pm 1.270 \times 10^3$
Bouncing puck	Left wall	0	$1.019 \times 10^6 \pm 1.152 \times 10^3$
Sliding block	Free	$5.070 \times 10^{-5} \pm 2.394 \times 10^{-7}$	105.198 ± 0.588
Sliding block	Stick	$3.673 \times 10^{-3} \pm 2.293 \times 10^{-4}$	20.645 ± 0.628
Sliding block	Slip	$2.543 \times 10^{-2} \pm 2.673 \times 10^{-4}$	6.445 ± 0.054
Pendulum stop	Free swing	0	$9.886 \times 10^5 \pm 7.698 \times 10^3$
Pendulum stop	Positive stop	0	$7.867 \times 10^5 \pm 3.262 \times 10^3$
Pendulum stop	Negative stop	$1.591 \times 10^{-17} \pm 2.979 \times 10^{-17}$	$7.866 \times 10^5 \pm 2.570 \times 10^3$
Planar pushing	No contact	$2.638 \times 10^{-3} \pm 5.621 \times 10^{-5}$	13.890 ± 0.222
Planar pushing	Sticking push	$3.133 \times 10^{-3} \pm 9.058 \times 10^{-5}$	12.449 ± 0.167
Planar pushing	Sliding push	$3.276 \times 10^{-3} \pm 7.553 \times 10^{-5}$	12.971 ± 0.187

Under this constructed confidence model, the gate reduces the mode-error contribution to approximately 0. Derivative perturbation produces gradual, monotone degradation. On the hardest system, coefficient error increases from 0.031 to 0.114 and support F1 falls from 0.90 to 0.87 at 10 times the base perturbation.

Two systems have a full-design minimum eigenvalue of approximately 0 and floored condition diagnostic approximately 10^6 because some contact features are structurally inactive within particular mode regions. A finite floored diagnostic does not imply a finite true condition number for a rank-deficient design. These full-space diagnostics do not verify the restricted-eigenvalue condition of Theorem 3, which concerns the sparse cone. Accurate recovery can therefore coexist with poor full-space conditioning. The measurements do not separately certify all four terms of the theorem's perturbation budget.

Coefficient-recovery diagnostics. Figure 10 shows estimated versus true active coefficients. The VHYDRO points stay close to the identity line across the active terms, while the dense no-sparsity fit spreads mass across nonessential terms and inflates coefficient magnitude. Figure 11 reports the corresponding physical-constant error: VHYDRO stays near zero, whereas the no-mode ablation varies widely because it tries to explain multiple regimes with a single law.

Support-recovery diagnostics. Figure 12 decomposes support recovery into selectivity. Precision is the informative component: VHYDRO substantially exceeds the dense and no-mode ablations, indicating that the recovered support is not only inclusive but selective. The support heatmap is included in Figure 13.

Interpretation. The sparse-law benchmark is designed so that a purely predictive model can perform well without identifying the governing mechanism. The no-pH baseline demonstrates this possibility: it achieves a moderate vector-field error, yet it offers no recoverable port-Hamiltonian coefficients, no sparse support, and no physical constants. On this controlled known-equation audit, VHYDRO simultaneously predicts accurately and recovers the sparse, mode-wise physical mechanism.

A.5 Planning theorem

Theorem 4 (Passivity-calibrated rollouts and KL-robust MPC)**.** *For each action sequence $a \in \mathcal{A}$, let Q_a be the* VHYDRO *predictive law over hybrid trajectories generated, between switching times, by*

$$dz_t = \big[(J_{s_t} - R_{s_t})\nabla H_{s_t}(z_t) + G_{s_t} a_t\big]dt + \Sigma_{s_t}^{1/2} dW_t. \tag{17}$$

Assume the rollout remains in certified regions $\mathcal{D}_m$ where $h_m^\star := \inf_{z \in \mathcal{D}_m} H_m(z) > -\infty$ and $U_m(z) := H_m(z) - h_m^\star$. Uniformly over modes, suppose $J_m^\top = -J_m$, $R_m \succeq rI$, $\|G_m\|_{\mathrm{op}} \le g$,

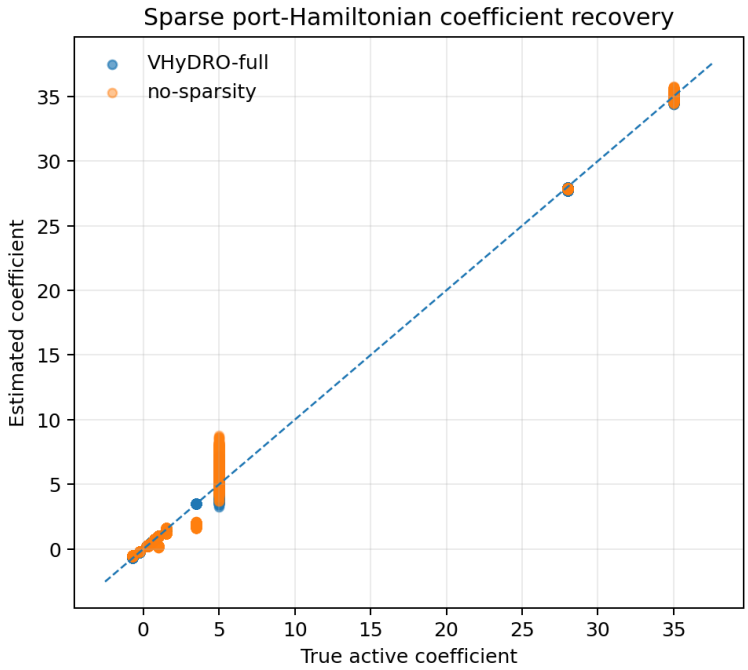

Figure 10: Coefficient recovery scatter. Estimated active coefficients versus true active coefficients. VHyDRO concentrates tightly around the identity line, while the no-sparsity ablation exhibits larger spread and systematic inflation.

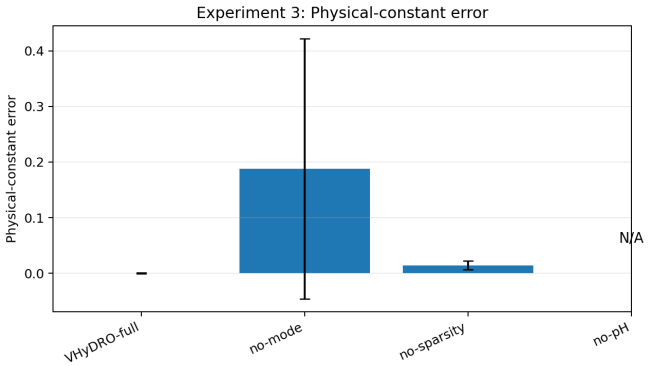

Figure 11: Physical-constant error. VHyDRO preserves the physical constants implied by the recovered law. The no-mode ablation is unstable because it conflates distinct regimes, and the no-pH baseline is not defined in this metric.

$\mathrm{tr}(\Sigma_m) \le \nu^2$, $\|a_t\|_2 \le A$, $\|\nabla^2 H_m\|_{\mathrm{op}} \le L_H$, *and* $\|\nabla H_m(z)\|_2^2 \ge 2\mu U_m(z)$ *on* $\mathcal{D}_m$. *Assume* $\Sigma_m \succeq 0$, *finitely many switches almost surely on* $[0,T]$, *and sufficient integrability so that the stochastic integrals in the hybrid Itô decomposition are true martingales. Assume also that contact switches are energy-nonexpansive:* $U_{s_\tau}(z_\tau) \le U_{s_{\tau^-}}(z_{\tau^-})$. *With*

$$C_E := \frac{g^2 A^2}{2r} + \frac{L_H \nu^2}{2}, \qquad \alpha := r\mu, \tag{18}$$

we have, for all a *and horizons* T,

$$\mathbb{E}_{Q_a}[U_{s_T}(z_T)] \le e^{-\alpha T}\mathbb{E}_{Q_a}[U_{s_0}(z_0)] + \frac{1-e^{-\alpha T}}{\alpha} C_E. \tag{19}$$

Let $P_a^\star$ *be the true trajectory law and assume uniform predictive calibration* $D_{\mathrm{KL}}(P_a^\star \| Q_a) \le \varepsilon_H$. *With* $\epsilon_{\mathrm{TV}} := \min\{1, \sqrt{\varepsilon_H/2}\}$, *every bounded cost* $C \in [0, C_{\max}]$ *and every* η_{opt}*-optimal* VHYDRO *plan* $\widehat{a}$ *satisfy*

$$\mathbb{E}_{P_{\widehat{a}}^\star}[C] - \inf_{a \in \mathcal{A}} \mathbb{E}_{P_a^\star}[C] \le 2C_{\max}\epsilon_{\mathrm{TV}} + \eta_{\mathrm{opt}}. \tag{20}$$

For any failure event $\mathcal{F}$ *and any* $\delta \in [\epsilon_{\mathrm{TV}}, 1]$,

$$Q_a(\mathcal{F}) \le \delta - \epsilon_{\mathrm{TV}} \quad \Longrightarrow \quad P_a^\star(\mathcal{F}) \le \delta. \tag{21}$$

If $Q_a(\mathcal{F})$ *is estimated from* n *rollouts over a finite evaluated set of* L *action sequences, the implication remains valid with probability at least* $1-\beta$ *after replacing* ϵ_{TV} *in* (21) *by*

$$\epsilon_{\mathrm{TV}} + \sqrt{\frac{\log(L/\beta)}{2n}}. \tag{22}$$

Theorem 4 is a conditional downstream statement: if the learned rollout law is KL-calibrated, the recovered port-Hamiltonian structure transfers to bounded-energy and chance-constraint guarantees. The empirical section focuses only on the three measured links above.

Relation to prior bounds. The result combines standard dissipative energy estimates, Gronwall-style control, and Pinsker transfer from predictive KL to cost and chance-constraint error. The VHYDRO-specific role is to connect recovered sparse hybrid laws to conditional planning guarantees, not to provide a numerical planning certificate in the present experiments.

This section reports supplementary MPC diagnostics motivated by Theorem 4: one-step rollout error, standardized energy-proxy drift, safety-adjusted return, and failure-risk calibration. These diagnostics do not instantiate the theorem as a numerical certificate because the required held-out KL bound $D_{\mathrm{KL}}(P_a^\star \| Q_a)$ is not estimated.

Protocol. We evaluate PushCube-v1, PickCube-v1, PegInsertionSide-v1, and StackCube-v1, the four tasks that instantiate reliably in our local ManiSkill build. All training data come from rollout HDF5 files collected by this experiment. External ManiSkill demonstration files are not mixed into training because observation layouts differ across tasks and control modes. Rollout collection and MPC evaluation are executed serially in the main process to avoid native SAPIEN/ManiSkill subprocess crashes; PyTorch training uses two GPUs.

Diagnostics. The main diagnostic reports quantities motivated by Theorem 4: rollout RMSE, common energy-proxy drift, safety-adjusted return, and episode-level risk-calibration gap. The standardized energy proxy is a diagnostic for bounded standardized energy growth and is not an exact Hamiltonian energy. Because Theorem 4 is stated for bounded cost minimization, return diagnostics should be read after the standard transformation $C = C_{\max} - R$; the reported returns are descriptive and are not used to instantiate Eq. (20).

Table 15: **Supplementary Experiment 4 MPC diagnostics.** Values are from the online MPC audit. The run uses twenty seeds and should be interpreted as descriptive diagnostics motivated by Theorem 4, not as a numerical theorem certificate or task-success evidence.

Task	RMSE↓	Energy drift↓	Return↑	Safe return↑	Pred. risk	Risk gap↓	ECE↓
PegInsertionSide	0.246	0.857	29.643	18.786	0.295	0.705	0.705
PickCube	0.493	2.452	11.638	-0.813	0.516	0.484	0.484
PushCube	0.502	1.658	11.171	1.179	0.637	0.196	0.308
StackCube	0.460	3.210	20.038	6.828	0.603	0.397	0.397

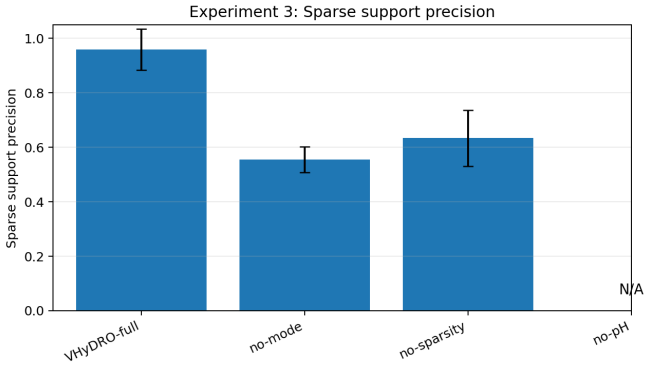

Figure 12: Sparse support precision. The key difference between VHYDRO and the ablations is selectivity: VHYDRO recovers the active support while suppressing inactive library terms.

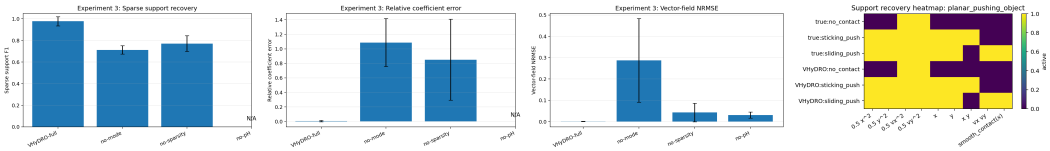

Figure 13: Sparse physical-law recovery. On hybrid systems with known sparse port-Hamiltonian laws, VHYDRO recovers the active library terms with high support F1 and low coefficient error. The no-mode ablation confounds regimes, the no-sparsity ablation predicts with dense coefficients but loses sparse support recovery, and the no-pH ablation fits a vector field without recoverable port-Hamiltonian coefficients.

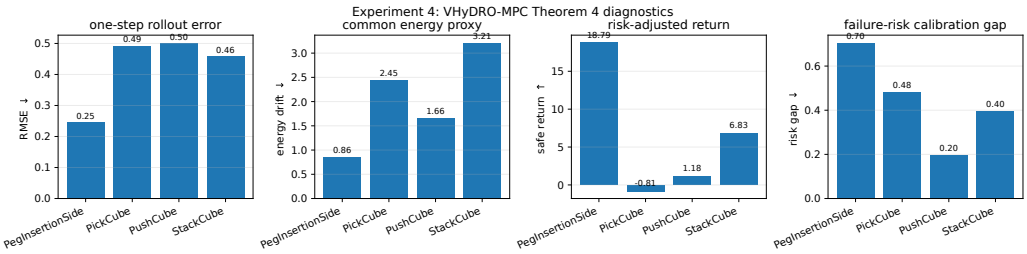

Figure 14: Supplementary Experiment 4 MPC diagnostics. Values are from the online MPC audit. The run uses twenty seeds and should be interpreted as descriptive diagnostics motivated by Theorem 4, not as a numerical theorem certificate or task-success evidence.

Interpretation. Across twenty seeds, the reported supplementary diagnostics are finite on all four tasks: mean rollout RMSE 0.425, mean common energy-proxy drift 2.044, mean safety-adjusted return 6.495, and mean absolute risk-calibration gap 0.445. These values are a descriptive MPC diagnostic, not a non-vacuity check on the constants of Theorem 4. We do not estimate a held-out upper bound on $D_{\mathrm{KL}}(P_a^\star \| Q_a)$, so the resulting Pinsker ϵ_{TV} is not numerically reported and Eq. (20) is not a certified suboptimality interval; this, together with a per-component empirical instantiation of $(\nu^{\text{filt}}, \nu^{\text{der}}, \nu^{\text{mode}}, \nu^{\text{pH}})$ from Theorem 3, are the principal follow-ups.

Where the diagnostics remain loose. Table 15 also exposes where the present MPC diagnostics remain loose. PickCube exhibits a negative safe-return mean (-0.813), and PegInsertionSide shows a 0.705 risk-calibration gap together with a 0.705 ECE; both are consistent with our *not* yet having a held-out upper bound on $D_{\mathrm{KL}}(P_a^\star \,\|\, Q_a)$, which leaves the Pinsker ϵ_{TV} unestimated and Eq. (20) uncertified numerically. We therefore present Experiment 4 strictly as supplementary descriptive diagnostics, and identify a calibrated KL upper bound and the matching per-component $(\nu^{\text{filt}}, \nu^{\text{der}}, \nu^{\text{mode}}, \nu^{\text{pH}})$ audit for Theorem 3 as the principal next experiment, not as an unstated limitation.

Connection to planning. A planner uses the filter's predictive distribution, whose deterioration under proposal-support loss is measured in Table 1. Theorem 4 connects KL-calibrated rollouts to conditional bounds on MPC suboptimality and chance-constraint transfer. Supplementary Experiment 4 evaluates online VHYDRO-MPC on four ManiSkill tasks using 20 seeds (Table 15 and Figure 14).

The planning and predictive-belief audits test the proposed mechanism. They do not establish broad robotic performance gains or a numerical certificate for the conditional planning theorem.

Predictive-belief calibration. We filter identical observations for 30 steps and roll the same model open-loop for 25 steps. The conservative support-safe proposal improves calibration relative to no defense at every evaluated horizon. Rollout coverage is 0.60 versus 0.50 at step 5 and 0.75 versus 0.71 at step 25. Rollout NLL at step 25 is 2.32 versus 2.92.

Compute and hardware disclosure. All experiments were run on a single workstation with an AMD Ryzen Threadripper PRO 5975WX CPU (30 physical cores / 60 logical threads) and two NVIDIA GeForce RTX 4090 GPUs with approximately 24 GB VRAM each. We used approximately 30 CPU worker threads where parallel data loading was safe. Hybrid-latent training was parallelized across the two GPUs by running independent seeds concurrently, while rollout collection and MPC evaluation in Experiment 4 were executed serially in the main process to avoid native SAPIEN/ManiSkill subprocess crashes.

The full Experiment 4 audit, including rollout collection, MPC evaluation, HDF5 serialization, model fitting, and repeated environment-instantiation attempts across the four ManiSkill tasks and twenty seed configurations, took approximately 80 hours on this machine.

B Supplementary proofs

The main text states the three measured results in compact form. The proofs below keep the full technical derivations and explicitly map each conclusion back to the corresponding theorem and equation numbers. Displays labeled with prefix `eq:app` are auxiliary equations used only inside the Supplementary Information; equations without this prefix refer to the numbered equations in the main paper.

B.1 Proof of Theorem 1: support-safe budgeted IW/FIVO increment

Proof. The theorem is stated around the defensive proposal in Eq. (1). We prove, in order, the domination and Radon–Nikodym bound in Eq. (2), the chi-square bound in Eq. (3), the unbiasedness/variance/ESS statement in Eq. (4), and the budgeted choice of λ in Eq. (5).

All quantities are conditional on the realized particle history

$$h_t = (s_t, z_t, a_t, o_{\leq t+1}). \tag{23}$$

For readability, write

$$P := P_t, \qquad Q := Q_t, \qquad Q_\lambda := Q_{t,\lambda}, \qquad g := g_t, \qquad Z := Z_t, \qquad M_2 := M_{2,t}. \tag{24}$$

With this notation, Eq. (1) becomes

$$Q_\lambda = (1-\lambda)Q + \lambda P. \tag{25}$$

Defensive-mixture domination. For any measurable set A, Eq. (25) gives

$$Q_\lambda(A) = (1-\lambda)Q(A) + \lambda P(A) \geq \lambda P(A). \tag{26}$$

Hence, if $Q_\lambda(A) = 0$, then $P(A) = 0$. Therefore

$$P \ll Q_\lambda. \tag{27}$$

By the Radon–Nikodym theorem, the derivative

$$r(x) := \frac{dP}{dQ_\lambda}(x) \tag{28}$$

exists Q_λ-almost surely.

Bounded prior-to-proposal ratio. The domination inequality (26) implies

$$P(A) \leq \frac{1}{\lambda} Q_\lambda(A) \qquad \forall A. \tag{29}$$

We claim that

$$r(x) \leq \frac{1}{\lambda} \qquad Q_\lambda\text{-a.s.} \tag{30}$$

Suppose, for contradiction, that

$$B := \left\{ x : r(x) > \frac{1}{\lambda} \right\} \tag{31}$$

has positive Q_λ-measure. Then

$$P(B) = \int_B r(x)\, Q_\lambda(dx) > \frac{1}{\lambda} Q_\lambda(B), \tag{32}$$

which contradicts Eq. (29). Thus Eq. (30) holds. Together with nonnegativity of Radon–Nikodym derivatives, this is precisely the main-text bound Eq. (2).

Domination of the filtering target. Define the unnormalized target measure

$$\widetilde{\Pi}(A) := \int_A g(x)\, P(dx). \tag{33}$$

Because $g \geq 0$, if $P(A) = 0$, then $\widetilde{\Pi}(A) = 0$. Hence $\widetilde{\Pi} \ll P$. Since $0 < Z < \infty$, the normalized measure

$$\Pi(A) := \frac{\widetilde{\Pi}(A)}{Z} = \frac{1}{Z} \int_A g(x)\, P(dx) \tag{34}$$

is a probability measure. Moreover,

$$\Pi \ll P \ll Q_\lambda, \tag{35}$$

so $\Pi \ll Q_\lambda$. This proves the target-support part of Theorem 1.

Derivative of the filtering target. Since

$$\Pi(dx) = \frac{g(x)}{Z} P(dx) \tag{36}$$

and

$$P(dx) = r(x) Q_\lambda(dx), \tag{37}$$

the Radon–Nikodym chain rule gives

$$\frac{d\Pi}{dQ_\lambda}(x) = \frac{g(x)}{Z} r(x) \qquad Q_\lambda\text{-a.s.} \tag{38}$$

The assumption that $g(x) < \infty$ $(P+Q)$-a.s. implies $g(x) < \infty$ Q_λ-a.s., because $Q_\lambda = (1-\lambda)Q + \lambda P$. Together with Eq. (30), this gives

$$W(x) := g(x) r(x) < \infty, \qquad 0 \leq W(x) \leq \frac{g(x)}{\lambda} \qquad Q_\lambda\text{-a.s.} \tag{39}$$

This is the full SI version of the finite-weight statement summarized in Theorem 1.

Chi-square identity. By definition,

$$1 + \chi^2(\Pi \| Q_\lambda) = \int_{\mathsf{X}} \left(\frac{d\Pi}{dQ_\lambda}(x) \right)^2 Q_\lambda(dx). \tag{40}$$

Using Eq. (38),

$$1 + \chi^2(\Pi \| Q_\lambda) = \frac{1}{Z^2} \int_{\mathsf{X}} g(x)^2 r(x)^2 \, Q_\lambda(dx). \tag{41}$$

Guide-independent chi-square bound. Since $0 \le r(x) \le \lambda^{-1}$ Q_λ-a.s. by Eq. (2),

$$r(x)^2 \le \frac{1}{\lambda} r(x) \qquad Q_\lambda\text{-a.s.} \tag{42}$$

Therefore,

$$\begin{aligned} 1 + \chi^2(\Pi \| Q_\lambda) &\le \frac{1}{\lambda Z^2} \int_{\mathsf{X}} g(x)^2 r(x) \, Q_\lambda(dx) \\ &= \frac{1}{\lambda Z^2} \int_{\mathsf{X}} g(x)^2 \, P(dx) = \frac{M_2}{\lambda Z^2}. \end{aligned} \tag{43}$$

Equivalently,

$$\chi^2(\Pi \| Q_\lambda) \le \frac{1}{\lambda} \frac{M_2}{Z^2} - 1. \tag{44}$$

Since $Z = \mathbb{E}_P[g(X)]$, $M_2 = \mathbb{E}_P[g(X)^2]$, and $\rho_t = M_{2,t}/Z_t^2$, Eq. (44) is exactly the main-text chi-square bound Eq. (3).

Second-moment identity. From Eq. (38),

$$W(x) = g(x) r(x) = Z \frac{d\Pi}{dQ_\lambda}(x). \tag{45}$$

Therefore,

$$\begin{aligned} \mathbb{E}_{Q_\lambda}[W(X)^2] &= Z^2 \int_{\mathsf{X}} \left(\frac{d\Pi}{dQ_\lambda}(x) \right)^2 Q_\lambda(dx) \\ &= Z^2 \left(1 + \chi^2(\Pi \| Q_\lambda) \right). \end{aligned} \tag{46}$$

Combining Eq. (46) with Eq. (43) gives

$$\mathbb{E}_{Q_\lambda}[W(X)^2] \le \frac{1}{\lambda} M_2 < \infty. \tag{47}$$

Unbiasedness of the one-step normalizer estimate. Let $X \sim Q_\lambda$. Then

$$\mathbb{E}_{Q_\lambda}[W(X)] = \int_{\mathsf{X}} g(x) r(x) \, Q_\lambda(dx) = \int_{\mathsf{X}} g(x) \, P(dx) = Z. \tag{48}$$

Thus, for conditionally i.i.d. samples $X^1, \dots, X^N \sim Q_\lambda$ and

$$\widehat{Z}^{(N)} := \frac{1}{N} \sum_{i=1}^{N} W(X^i), \tag{49}$$

we have

$$\mathbb{E}[\widehat{Z}^{(N)} \mid h_t] = Z. \tag{50}$$

This is the first line of Eq. (4).

Variance identity. Because $X^1, \ldots, X^N$ are conditionally i.i.d. given h_t,

$$\mathrm{Var}(\widehat{Z}^{(N)} \mid h_t) = \frac{1}{N} \mathrm{Var}_{Q_\lambda}(W(X)). \tag{51}$$

Using Eqs. (46) and (48),

$$\mathrm{Var}_{Q_\lambda}(W(X)) = Z^2 \left(1 + \chi^2(\Pi \| Q_\lambda)\right) - Z^2 = Z^2 \chi^2(\Pi \| Q_\lambda). \tag{52}$$

Therefore,

$$\frac{\mathrm{Var}(\widehat{Z}^{(N)} \mid h_t)}{Z^2} = \frac{1}{N} \chi^2(\Pi \| Q_\lambda). \tag{53}$$

Applying Eq. (44) yields

$$\frac{\mathrm{Var}(\widehat{Z}^{(N)} \mid h_t)}{Z^2} \leq \frac{1}{N} \left[\frac{1}{\lambda} \frac{\mathbb{E}_P[g(X)^2]}{\mathbb{E}_P[g(X)]^2} - 1 \right], \tag{54}$$

which is the variance part of Eq. (4).

Asymptotic normalized ESS. Let

$$W_i := W(X^i), \qquad \overline{W}_N := \frac{1}{N} \sum_{i=1}^N W_i, \qquad \overline{W^2}_N := \frac{1}{N} \sum_{i=1}^N W_i^2. \tag{55}$$

By Eqs. (46) and (48),

$$\mathbb{E}[W_i] = Z \in (0, \infty), \qquad \mathbb{E}[W_i^2] < \infty. \tag{56}$$

The strong law of large numbers gives

$$\overline{W}_N \to Z, \qquad \overline{W^2}_N \to Z^2 \left(1 + \chi^2(\Pi \| Q_\lambda)\right) \qquad \text{a.s.} \tag{57}$$

The limiting denominator is strictly positive because $Z > 0$. Hence, on the almost-sure event where both limits hold,

$$\frac{\mathrm{ESS}_N}{N} = \frac{\left(N^{-1} \sum_{i=1}^N W_i\right)^2}{N^{-1} \sum_{i=1}^N W_i^2} = \frac{\overline{W}_N^2}{\overline{W^2}_N} \to \frac{1}{1 + \chi^2(\Pi \| Q_\lambda)}. \tag{58}$$

The finite-N convention $\mathrm{ESS}_N = 0$ when $\sum_i W_i^2 = 0$ does not affect the limit because $\overline{W}_N \to Z > 0$. From Eq. (43),

$$\frac{1}{1 + \chi^2(\Pi \| Q_\lambda)} \geq \lambda \frac{\mathbb{E}_P[g(X)]^2}{\mathbb{E}_P[g(X)^2]}. \tag{59}$$

Equations (58)–(59) are exactly the ESS portion of Eq. (4).

Variance-budget certificate. Define

$$\rho := \frac{M_2}{Z^2} = \frac{\mathbb{E}_P[g(X)^2]}{\mathbb{E}_P[g(X)]^2}. \tag{60}$$

Because P is a probability measure and $g \in L^2(P)$, Jensen's inequality gives $M_2 \geq Z^2$. Thus $\rho \in [1, \infty)$. Let $\bar{\rho} \geq \rho$ be a certified upper bound. From Eq. (54),

$$\frac{\mathrm{Var}(\widehat{Z}^{(N)} \mid h_t)}{Z^2} \leq \frac{1}{N} \left(\frac{\rho}{\lambda} - 1 \right) \leq \frac{1}{N} \left(\frac{\bar{\rho}}{\lambda} - 1 \right). \tag{61}$$

Therefore the guide-independent bound certifies

$$\frac{\mathrm{Var}(\widehat{Z}^{(N)} \mid h_t)}{Z^2} \leq \tau^2 \tag{62}$$

whenever

$$\frac{1}{N} \left(\frac{\bar{\rho}}{\lambda} - 1 \right) \leq \tau^2. \tag{63}$$

This is equivalent to

$$\lambda \geq \frac{\bar{\rho}}{1+N\tau^2}. \tag{64}$$

If $\bar{\rho} \leq 1 + N\tau^2$, then the right-hand side of Eq. (64) is at most one, so every

$$\lambda \in \left[\frac{\bar{\rho}}{1+N\tau^2}, 1\right] \tag{65}$$

certifies the target budget. If $\bar{\rho} > 1 + N\tau^2$, then no $\lambda \in (0, 1]$ can certify the target using this guide-independent upper bound. The function

$$\lambda \mapsto \frac{1}{N}\left(\frac{\bar{\rho}}{\lambda} - 1\right) \tag{66}$$

is decreasing on $(0, 1]$, so the best certified bound from this inequality is attained at $\lambda = 1$ and equals

$$\frac{1}{N}(\bar{\rho} - 1). \tag{67}$$

The sufficient condition Eq. (64), together with feasibility $\bar{\rho}/(1 + N\tau^2) \leq 1$, is the main-text budget rule Eq. (5).

Density-level version. This paragraph is an implementation-level SI extension of Theorem 1; it is not needed to state the compact main-text theorem, but it justifies the density formula used in code. Assume that P and Q admit tractable densities p and q with respect to a common σ-finite reference measure μ. Then Q_λ admits the density

$$q_\lambda(x) = (1 - \lambda)q(x) + \lambda p(x). \tag{68}$$

Since

$$q_\lambda(x) \geq \lambda p(x) \qquad \mu\text{-a.e.}, \tag{69}$$

we have, with the convention $0/0 := 0$,

$$\frac{p(x)}{q_\lambda(x)} \leq \frac{1}{\lambda} \qquad q_\lambda\text{-a.e.} \tag{70}$$

Moreover, for any measurable A,

$$\int_A \frac{p(x)}{q_\lambda(x)} q_\lambda(x)\, \mu(dx) = \int_A p(x)\, \mu(dx) = P(A), \tag{71}$$

so p/q_λ is a valid density-level version of dP/dQ_λ. Therefore the density-level importance weight

$$W(x) = g(x)\frac{p(x)}{q_\lambda(x)} \tag{72}$$

agrees with the measure-theoretic weight above up to Q_λ-null sets.

State-dependent or adaptive λ. This is another SI extension of the compact theorem. The preceding argument was conditional on h_t and on a fixed $\lambda \in (0, 1]$. If λ is chosen as an h_t-measurable, possibly random quantity before drawing $X^1, \ldots, X^N$, then conditional on (h_t, λ) the same proof applies verbatim. Taking expectations over the randomness used to choose λ preserves conditional unbiasedness, provided the same realized value of λ is used in both the proposal $Q_{t,\lambda}$ and the weight W_t.

The result is therefore a one-step conditional statement for the proposal-and-weighting increment. In an IW/FIVO/SMC implementation, it applies to each particle ancestor after conditioning on that ancestor and on the information available before sampling the next particle. □

B.2 Proof of Theorem 2: exponential concentration of predictive modes

Proof. The main text defines the clamped-mode posterior in Eq. (6) and the predictive log-ratio in Eq. (7). We prove the sharp posterior bound Eq. (8), its simplified exponential form Eq. (9), and the variational transfer bound Eq. (10).

All probabilities and expectations are taken under the true data-generating law $\mathbb{P}^\star$ for the considered constant-mode segment, conditional on the realized action process when actions are treated as interventions. Fix a wrong mode $s \neq m$ and define

$$S_{L,s} := \sum_{\ell=0}^{L-1} R_{\ell,s}, \tag{73}$$

where $R_{\ell,s}$ is the main-text log-ratio in Eq. (7). By definition,

$$R_{\ell,s} = \mu_{\ell,s} + \xi_{\ell,s}, \qquad \mathbb{E}_\star[\xi_{\ell,s} \mid \mathcal{F}_\ell] = 0. \tag{74}$$

Moreover, by the conditional sub-Gaussian assumption,

$$\mathbb{E}_\star[\exp\{t\xi_{\ell,s}\}|\mathcal{F}_\ell] \leq \exp\left(\frac{t^2\sigma_{\ell,s}^2}{2}\right) \qquad \forall t \in \mathbb{R}. \tag{75}$$

The variable $\xi_{\ell,s}$ is measurable with respect to the sigma-field generated by $\mathcal{F}_\ell$ and the next observation $O_{\ell+1}$, hence with respect to $\mathcal{F}_{\ell+1}$ under the filtration above. Therefore the partial sums

$$M_{j,s} := \sum_{\ell=0}^{j-1} \xi_{\ell,s}, \qquad j = 0, \ldots, L, \tag{76}$$

form a martingale with conditionally sub-Gaussian increments.

We first prove a lower-tail inequality for $S_{L,s}$. By cumulative predictive separation,

$$\sum_{\ell=0}^{L-1} \mu_{\ell,s} \geq \Gamma_{L,s} \qquad \text{almost surely.} \tag{77}$$

Hence, for any $u > 0$,

$$\begin{aligned} \mathbb{P}^\star[S_{L,s} \leq \Gamma_{L,s} - u] &= \mathbb{P}^\star\left[\sum_{\ell=0}^{L-1}(\mu_{\ell,s} + \xi_{\ell,s}) \leq \Gamma_{L,s} - u\right] \\ &\leq \mathbb{P}^\star\left[\sum_{\ell=0}^{L-1} \xi_{\ell,s} \leq -u\right]. \end{aligned} \tag{78}$$

Assume first that $V_s > 0$. For any $\eta > 0$, Markov's inequality gives

$$\begin{aligned} \mathbb{P}^\star\left[\sum_{\ell=0}^{L-1} \xi_{\ell,s} \leq -u\right] &= \mathbb{P}^\star\left[\exp\left\{-\eta \sum_{\ell=0}^{L-1} \xi_{\ell,s}\right\} \geq e^{\eta u}\right] \\ &\leq e^{-\eta u}\mathbb{E}_\star\left[\exp\left\{-\eta \sum_{\ell=0}^{L-1} \xi_{\ell,s}\right\}\right]. \end{aligned} \tag{79}$$

Using the tower property and applying Eq. (75) backwards from $\ell = L-1$ to $\ell = 0$,

$$\mathbb{E}_\star\left[\exp\left\{-\eta \sum_{\ell=0}^{L-1} \xi_{\ell,s}\right\}\right] \leq \exp\left(\frac{\eta^2}{2}\sum_{\ell=0}^{L-1} \sigma_{\ell,s}^2\right) = \exp\left(\frac{\eta^2 V_s}{2}\right). \tag{80}$$

Therefore,

$$\mathbb{P}^\star[S_{L,s} \leq \Gamma_{L,s} - u] \leq \exp\left(-\eta u + \frac{\eta^2 V_s}{2}\right). \tag{81}$$

Optimizing over $\eta > 0$ with $\eta = u/V_s$ yields

$$\mathbb{P}^\star[S_{L,s} \le \Gamma_{L,s} - u] \le \exp\left(-\frac{u^2}{2V_s}\right). \tag{82}$$

If $V_s = 0$, then $\sigma_{\ell,s} = 0$ for every ℓ. In this case, for any fixed nonzero t, Jensen's inequality and Eq. (75) imply

$$1 = \exp(t\,\mathbb{E}_\star[\xi_{\ell,s} \mid \mathcal{F}_\ell]) \le \mathbb{E}_\star[\exp\{t\xi_{\ell,s}\} \mid \mathcal{F}_\ell] \le 1. \tag{83}$$

Thus equality holds in Jensen's inequality. Since the exponential function is strictly convex, $\xi_{\ell,s}$ must be conditionally almost surely constant given $\mathcal{F}_\ell$. Its conditional mean is zero, so $\xi_{\ell,s} = 0$ almost surely. Consequently,

$$S_{L,s} = \sum_{\ell=0}^{L-1} \mu_{\ell,s} \ge \Gamma_{L,s} \qquad \text{almost surely.} \tag{84}$$

Thus the desired lower bound also holds in the degenerate case $V_s = 0$.

Now define, for every $s \neq m$,

$$u_s := \sqrt{2V_s \log \frac{M-1}{\delta}}. \tag{85}$$

For modes with $V_s > 0$, Eq. (82) gives

$$\mathbb{P}^\star[S_{L,s} \le \Gamma_{L,s} - u_s] \le \frac{\delta}{M-1}. \tag{86}$$

For modes with $V_s = 0$, the event $S_{L,s} \ge \Gamma_{L,s} - u_s$ holds almost surely because $u_s = 0$. Hence, by a union bound over the $M-1$ wrong modes, with probability at least $1-\delta$, the event

$$\mathcal{E} := \left\{ S_{L,s} \ge \Gamma_{L,s} - \sqrt{2V_s \log \frac{M-1}{\delta}} \quad \forall s \neq m \right\} \tag{87}$$

holds.

On $\mathcal{E}$, Bayes' rule and the constant-segment posterior from Eq. (6) give, for each $s \neq m$,

$$\begin{aligned} \frac{\Pi_L(S=s)}{\Pi_L(S=m)} &= \frac{\pi_s \prod_{\ell=0}^{L-1} p_s(O_{\ell+1} \mid \mathcal{F}_\ell)}{\pi_m \prod_{\ell=0}^{L-1} p_m(O_{\ell+1} \mid \mathcal{F}_\ell)} \\ &= \frac{\pi_s}{\pi_m} \exp\left(-\sum_{\ell=0}^{L-1} \log \frac{p_m(O_{\ell+1} \mid \mathcal{F}_\ell)}{p_s(O_{\ell+1} \mid \mathcal{F}_\ell)}\right) \\ &= \frac{\pi_s}{\pi_m} \exp(-S_{L,s}) \\ &\le \exp\left(B_s - \Gamma_{L,s} + \sqrt{2V_s \log \frac{M-1}{\delta}}\right). \end{aligned} \tag{88}$$

Therefore,

$$\sum_{s \neq m} \frac{\Pi_L(S=s)}{\Pi_L(S=m)} \le \sum_{s \neq m} \exp\left(B_s - \Gamma_{L,s} + \sqrt{2V_s \log \frac{M-1}{\delta}}\right) = A^\sharp_{L,\delta}. \tag{89}$$

Let

$$T_L := \sum_{s \neq m} \frac{\Pi_L(S=s)}{\Pi_L(S=m)}. \tag{90}$$

Since

$$1 = \Pi_L(S=m) + \sum_{s \neq m} \Pi_L(S=s) = \Pi_L(S=m)(1+T_L), \tag{91}$$

we have

$$\Pi_L(S \neq m) = \sum_{s \neq m} \Pi_L(S=s) = \frac{T_L}{1+T_L}. \tag{92}$$

On $\mathcal{E}$, $T_L \le A^\sharp_{L,\delta}$. Since $x \mapsto x/(1+x)$ is increasing on $[0,\infty)$,

$$\Pi_L(S \neq m) \le \frac{A^\sharp_{L,\delta}}{1 + A^\sharp_{L,\delta}} \le A^\sharp_{L,\delta}. \tag{93}$$

This is exactly the sharp exact-posterior recovery bound in Eq. (8).

We now prove the simplified bound. If

$$\Gamma_{L,s} \ge L\Delta, \qquad B_s \le B, \qquad V_s \le V_{\max}, \tag{94}$$

for all $s \neq m$, then

$$\begin{aligned} A^\sharp_{L,\delta} &= \sum_{s \neq m} \exp\left(B_s - \Gamma_{L,s} + \sqrt{2V_s \log \frac{M-1}{\delta}} \right) \\ &\le (M-1) \exp\left(B - L\Delta + \sqrt{2V_{\max} \log \frac{M-1}{\delta}} \right) = A_{L,\delta}. \end{aligned} \tag{95}$$

Since $x \mapsto x/(1+x)$ is increasing,

$$\Pi_L(S \neq m) \le \frac{A_{L,\delta}}{1 + A_{L,\delta}} \le A_{L,\delta}. \tag{96}$$

If additionally $\sigma_{\ell,s} \le \sigma$ for all ℓ and $s \neq m$, then

$$V_s = \sum_{\ell=0}^{L-1} \sigma^2_{\ell,s} \le L\sigma^2, \tag{97}$$

so $V_{\max} \le L\sigma^2$, which gives

$$\Pi_L(S \neq m) \le (M-1) \exp\left(B - L\Delta + \sigma\sqrt{2L \log \frac{M-1}{\delta}} \right). \tag{98}$$

This is the main-text simplified bound Eq. (9).

It remains to transfer the bound from the constant-segment model posterior Π_L to the amortized variational posterior q_ϕ. Suppose that, on the realized observations,

$$\mathrm{KL}(q_\phi(S \mid O_{1:L}, A_{0:L-1}) \,\|\, \Pi_L(S)) \le \varepsilon_q. \tag{99}$$

Using the convention

$$\mathrm{TV}(P,Q) := \sup_{\mathcal{A}} |P(\mathcal{A}) - Q(\mathcal{A})|, \tag{100}$$

Pinsker's inequality gives

$$\mathrm{TV}(q_\phi, \Pi_L) \le \sqrt{\frac{1}{2}\mathrm{KL}(q_\phi \| \Pi_L)} \le \sqrt{\frac{\varepsilon_q}{2}}. \tag{101}$$

Apply this to the event $\mathcal{A} := \{S \neq m\}$. Then

$$q_\phi(\mathcal{A}) \le \Pi_L(\mathcal{A}) + \mathrm{TV}(q_\phi, \Pi_L). \tag{102}$$

Combining Eq. (102) with Eq. (93) yields

$$q_\phi(S \neq m \mid O_{1:L}, A_{0:L-1}) \le \frac{A^\sharp_{L,\delta}}{1 + A^\sharp_{L,\delta}} + \sqrt{\frac{\varepsilon_q}{2}}. \tag{103}$$

This is exactly the variational transfer statement in Eq. (10). Using $A^\sharp_{L,\delta}/(1 + A^\sharp_{L,\delta}) \le A^\sharp_{L,\delta} \le A_{L,\delta}$ gives the simplified variational bound

$$q_\phi(S \neq m \mid O_{1:L}, A_{0:L-1}) \le A_{L,\delta} + \sqrt{\frac{\varepsilon_q}{2}}. \tag{104}$$

Finally, if the KL condition in Eq. (99) holds only with probability at least $1 - \delta_q$, then a union bound over the high-probability event $\mathcal{E}$ in Eq. (87) and the KL event gives probability at least $1 - \delta - \delta_q$. This last union-bound variant is a useful SI extension of the compact theorem statement. □

B.3 Proof of Theorem 3: filtering- and mode-robust sparse port-Hamiltonian recovery

Proof. The main text specifies the sparse port-Hamiltonian model in Eq. (11), the plug-in regression in Eq. (12), the score/RSC assumptions in Eq. (13), the estimator in Eq. (14), the recovery bounds in Eq. (15), and the high-probability penalty choice in Eq. (16). The proof below keeps the full Lasso argument and explicitly maps the final deterministic and probabilistic conclusions to those equations.

All deterministic statements below are conditional on the realized VHYDRO design $(\widehat{A}, \widehat{b})$. Let

$$\widehat{\Delta} := \widehat{\xi}_m - \xi_m^\star. \tag{105}$$

If $\widehat{\Delta} = 0$, then the coefficient and prediction bounds in Eq. (15) are immediate. Thus, whenever we divide by $\|\widehat{\Delta}\|_2$, we assume $\widehat{\Delta} \neq 0$.

By optimality of $\widehat{\xi}_m$ in Eq. (14),

$$\frac{1}{2n_m}\|\widehat{b} - \widehat{A}\widehat{\xi}_m\|_2^2 + \lambda_m\|\widehat{\xi}_m\|_1 \leq \frac{1}{2n_m}\|\widehat{b} - \widehat{A}\xi_m^\star\|_2^2 + \lambda_m\|\xi_m^\star\|_1. \tag{106}$$

Using Eq. (12), namely $\zeta_m = \widehat{b} - \widehat{A}\xi_m^\star$, and using $\widehat{\xi}_m = \xi_m^\star + \widehat{\Delta}$, we have

$$\widehat{b} - \widehat{A}\widehat{\xi}_m = \zeta_m - \widehat{A}\widehat{\Delta}. \tag{107}$$

Substituting Eq. (107) into Eq. (106), expanding the square, and cancelling $(2n_m)^{-1}\|\zeta_m\|_2^2$ from both sides yields the basic inequality

$$\frac{1}{2n_m}\|\widehat{A}\widehat{\Delta}\|_2^2 \leq \frac{1}{n_m}\zeta_m^\top\widehat{A}\widehat{\Delta} + \lambda_m\left(\|\xi_m^\star\|_1 - \|\xi_m^\star + \widehat{\Delta}\|_1\right). \tag{108}$$

We first control the perturbation term. By Holder's inequality and the score part of Eq. (13),

$$\begin{aligned}\frac{1}{n_m}\zeta_m^\top\widehat{A}\widehat{\Delta} &\leq \left|\left\langle \frac{1}{n_m}\widehat{A}^\top\zeta_m, \widehat{\Delta}\right\rangle\right| \\ &\leq \left\|\frac{1}{n_m}\widehat{A}^\top\zeta_m\right\|_\infty \|\widehat{\Delta}\|_1 \leq \frac{\lambda_m}{2}\|\widehat{\Delta}\|_1.\end{aligned} \tag{109}$$

Next, since $S = \operatorname{supp}(\xi_m^\star)$, we have $\xi_{m,S^c}^\star = 0$. Therefore,

$$\begin{aligned}\|\xi_m^\star\|_1 - \|\xi_m^\star + \widehat{\Delta}\|_1 &= \|\xi_{m,S}^\star\|_1 - \|\xi_{m,S}^\star + \widehat{\Delta}_S\|_1 - \|\widehat{\Delta}_{S^c}\|_1 \\ &\leq \|\widehat{\Delta}_S\|_1 - \|\widehat{\Delta}_{S^c}\|_1,\end{aligned} \tag{110}$$

where the final inequality follows from the triangle inequality.

Combining Eqs. (108), (109), and (110), and using $\|\widehat{\Delta}\|_1 = \|\widehat{\Delta}_S\|_1 + \|\widehat{\Delta}_{S^c}\|_1$, we obtain

$$\begin{aligned}\frac{1}{2n_m}\|\widehat{A}\widehat{\Delta}\|_2^2 &\leq \frac{\lambda_m}{2}\left(\|\widehat{\Delta}_S\|_1 + \|\widehat{\Delta}_{S^c}\|_1\right) + \lambda_m\left(\|\widehat{\Delta}_S\|_1 - \|\widehat{\Delta}_{S^c}\|_1\right) \\ &= \frac{3\lambda_m}{2}\|\widehat{\Delta}_S\|_1 - \frac{\lambda_m}{2}\|\widehat{\Delta}_{S^c}\|_1.\end{aligned} \tag{111}$$

The left-hand side is nonnegative and $\lambda_m > 0$. Hence the right-hand side must be nonnegative, which implies the cone condition

$$\|\widehat{\Delta}_{S^c}\|_1 \leq 3\|\widehat{\Delta}_S\|_1. \tag{112}$$

Thus $\widehat{\Delta}$ belongs to the cone on which the RSC part of Eq. (13) holds.

Multiplying Eq. (111) by 2 and dropping the negative term gives

$$\frac{1}{n_m}\|\widehat{A}\widehat{\Delta}\|_2^2 \leq 3\lambda_m\|\widehat{\Delta}_S\|_1. \tag{113}$$

By the restricted strong-convexity assumption in Eq. (13), applied using Eq. (112),

$$\kappa_m\|\widehat{\Delta}\|_2^2 \leq \frac{1}{n_m}\|\widehat{A}\widehat{\Delta}\|_2^2. \tag{114}$$

Combining Eqs. (113) and (114), and using

$$\|\widehat{\Delta}_S\|_1 \le \sqrt{k}\|\widehat{\Delta}_S\|_2 \le \sqrt{k}\|\widehat{\Delta}\|_2, \tag{115}$$

we get

$$\kappa_m\|\widehat{\Delta}\|_2^2 \le 3\lambda_m\sqrt{k}\|\widehat{\Delta}\|_2. \tag{116}$$

If $\widehat{\Delta} \neq 0$, division by $\|\widehat{\Delta}\|_2$ gives

$$\|\widehat{\Delta}\|_2 \le \frac{3\sqrt{k}\lambda_m}{\kappa_m} \le \frac{4\sqrt{k}\lambda_m}{\kappa_m}. \tag{117}$$

Together with the trivial case $\widehat{\Delta} = 0$, Eq. (117) proves the coefficient-error bound in Eq. (15).

The empirical vector-field bound follows from Eqs. (113) and (117):

$$\frac{1}{n_m}\|\widehat{A}\widehat{\Delta}\|_2^2 \le 3\lambda_m\sqrt{k}\|\widehat{\Delta}\|_2 \le \frac{9k\lambda_m^2}{\kappa_m} \le \frac{16k\lambda_m^2}{\kappa_m}. \tag{118}$$

This proves the second bound in Eq. (15). The equivalent displayed form in terms of the per-sample port-Hamiltonian vector field follows directly from the stacked definition of $\widehat{A}$ in Eq. (12).

We now prove exact support recovery after thresholding. Since

$$\|\widehat{\Delta}\|_\infty \le \|\widehat{\Delta}\|_2 \le \frac{4\sqrt{k}\lambda_m}{\kappa_m}, \tag{119}$$

every inactive coordinate $j \notin S$ satisfies

$$|\widehat{\xi}_{m,j}| = |\widehat{\xi}_{m,j} - \xi^\star_{m,j}| = |\widehat{\Delta}_j| \le \frac{4\sqrt{k}\lambda_m}{\kappa_m}. \tag{120}$$

Therefore, if $\tau > 4\sqrt{k}\lambda_m/\kappa_m$, no inactive coordinate is selected. For any active coordinate $j \in S$,

$$|\widehat{\xi}_{m,j}| \ge |\xi^\star_{m,j}| - |\widehat{\Delta}_j| \ge \beta_{\min,m} - \frac{4\sqrt{k}\lambda_m}{\kappa_m}. \tag{121}$$

If $\beta_{\min,m} > 8\sqrt{k}\lambda_m/\kappa_m$, then the interval

$$\left(\frac{4\sqrt{k}\lambda_m}{\kappa_m}, \beta_{\min,m} - \frac{4\sqrt{k}\lambda_m}{\kappa_m}\right) \tag{122}$$

is nonempty. For any τ in this interval, Eqs. (120) and (121) imply

$$\{j : |\widehat{\xi}_{m,j}| > \tau\} = S. \tag{123}$$

This proves the support-recovery part of Theorem 3.

It remains to prove the high-probability score bound and connect it to Eq. (16). Suppose

$$\zeta_m = e_m + u_m, \tag{124}$$

where e_m is conditionally sub-Gaussian with scale σ_m, and suppose

$$\left\|\frac{1}{n_m}\widehat{A}^\top u_m\right\|_\infty \le \nu_m. \tag{125}$$

For each column $\widehat{A}_{\cdot j}$, define

$$X_j := \frac{1}{n_m}\widehat{A}_{\cdot j}^\top e_m. \tag{126}$$

If $L_{A,m} = 0$, then $\widehat{A}^\top e_m = 0$, and the stochastic score bound is immediate. Otherwise, conditional on $\widehat{A}$, the sub-Gaussian assumption gives, for every $s \in \mathbb{R}$,

$$\mathbb{E}\Big[\exp(sX_j) \mid \widehat{A}\Big] \le \exp\left(\frac{\sigma_m^2 s^2\|\widehat{A}_{\cdot j}\|_2^2}{2n_m^2}\right). \tag{127}$$

By the definition of $L_{A,m}$ in Eq. (16),

$$\frac{\|\widehat{A}_{\cdot j}\|_2}{n_m} \le \frac{L_{A,m}}{\sqrt{n_m}}, \tag{128}$$

so X_j, conditional on $\widehat{A}$, is sub-Gaussian with variance proxy $\sigma_m^2 L_{A,m}^2/n_m$. Hence, for every $t > 0$,

$$\mathbb{P}\left(|X_j| > t \mid \widehat{A}\right) \le 2\exp\left(-\frac{n_m t^2}{2\sigma_m^2 L_{A,m}^2}\right). \tag{129}$$

Set

$$t_\delta := L_{A,m}\sigma_m\sqrt{\frac{2\log(2p/\delta)}{n_m}}. \tag{130}$$

Applying the union bound over $j = 1, \dots, p$, we obtain

$$\mathbb{P}\left(\left\|\frac{1}{n_m}\widehat{A}^\top e_m\right\|_\infty > t_\delta \;\middle|\; \widehat{A}\right) \le \delta. \tag{131}$$

Therefore, with conditional probability at least $1-\delta$,

$$\left\|\frac{1}{n_m}\widehat{A}^\top e_m\right\|_\infty \le L_{A,m}\sigma_m\sqrt{\frac{2\log(2p/\delta)}{n_m}}. \tag{132}$$

On this same event, using Eq. (124) and Eq. (125),

$$\begin{aligned}\left\|\frac{1}{n_m}\widehat{A}^\top \zeta_m\right\|_\infty &\le \left\|\frac{1}{n_m}\widehat{A}^\top e_m\right\|_\infty + \left\|\frac{1}{n_m}\widehat{A}^\top u_m\right\|_\infty \\ &\le L_{A,m}\sigma_m\sqrt{\frac{2\log(2p/\delta)}{n_m}} + \nu_m.\end{aligned} \tag{133}$$

Thus, any penalty satisfying the main-text choice Eq. (16) ensures

$$\left\|\frac{1}{n_m}\widehat{A}^\top \zeta_m\right\|_\infty \le \frac{\lambda_m}{2}, \tag{134}$$

which is exactly the score condition in Eq. (13). Since the conditional probability bound holds for every realized $\widehat{A}$, the same probability bound holds unconditionally by the tower property.

If the RSC part of Eq. (13) holds deterministically for the realized design, then the deterministic part of the proof applies on the score event above, so Eq. (15) and the support-recovery statement hold with probability at least $1-\delta$. If instead RSC holds on an event $\mathcal{E}_{\mathrm{rsc}}$ with probability at least $1-\delta_{\mathrm{rsc}}$, then the union bound gives

$$\mathbb{P}\big(\text{score condition} \cap \mathcal{E}_{\mathrm{rsc}}\big) \ge 1 - \delta - \delta_{\mathrm{rsc}}, \tag{135}$$

and the deterministic conclusions apply on this intersection.

Finally, the main text summarizes all filtering, derivative, mode, and plug-in effects inside the perturbation ν_m. The full SI decomposition is as follows. If

$$u_m = u_m^{\mathrm{filt}} + u_m^{\mathrm{der}} + u_m^{\mathrm{mode}} + u_m^{\mathrm{pH}}, \tag{136}$$

then by the triangle inequality,

$$\left\|\frac{1}{n_m}\widehat{A}^\top u_m\right\|_\infty \le \sum_{q\in\{\mathrm{filt},\mathrm{der},\mathrm{mode},\mathrm{pH}\}} \left\|\frac{1}{n_m}\widehat{A}^\top u_m^q\right\|_\infty. \tag{137}$$

Thus, if each component is bounded by ν_m^q, then the deterministic perturbation part of Eq. (16) holds with

$$\nu_m = \nu_m^{\mathrm{filt}} + \nu_m^{\mathrm{der}} + \nu_m^{\mathrm{mode}} + \nu_m^{\mathrm{pH}}. \tag{138}$$

This final decomposition is more detailed than the compact main-text theorem but is useful for ablations that separately measure filtering error, derivative-estimation error, mode impurity, and port-Hamiltonian plug-in error. □

What the bound assumes, and what would instantiate it quantitatively. Theorem 3 is conditional on two ingredients that are not directly certified in Experiment 3. *(i) Restricted-eigenvalue (RE) constant* κ_m. We assume Eq. (13) holds with some $\kappa_m > 0$ on the cone $\|\Delta_{S^c}\|_1 \le 3\|\Delta_S\|_1$; we do not verify this constant directly. On the four controlled hybrid systems used in Experiment 3, the candidate library has at most $p \le 12$ terms and the true active support has size at most $k \le 3$, which is the low-sparsity regime in which Lasso oracle bounds are typically informative [Bickel et al., 2009]. An explicit RE certificate via a sparse-eigenvalue lower bound is left as future work [Rudelson and Zhou, 2012]. *(ii) Four-component perturbation budget* $\nu_m = \nu_m^{\mathrm{filt}} + \nu_m^{\mathrm{der}} + \nu_m^{\mathrm{mode}} + \nu_m^{\mathrm{pH}}$. The decomposition in Appendix B.3 separates filtering error, derivative-estimation error, mode impurity, and port-Hamiltonian plug-in error. The empirical diagnostics in Appendix A.4 measure mode impurity, derivative error, and residual port-Hamiltonian error, but do not separately certify all four perturbation-budget terms. Support F1 and relative coefficient error assess aggregate recovery.

B.4 Proof of Theorem 4: passivity-calibrated rollouts and KL-robust MPC

Proof. Appendix A.5 defines the learned hybrid SDE in Eq. (17), the constants C_E and α in Eq. (18), the energy inequality in Eq. (19), the MPC suboptimality bound in Eq. (20), the deterministic chance certificate in Eq. (21), and the Monte Carlo margin in Eq. (22). We prove these claims and include two SI-level details not displayed in the compact theorem: the Markov high-energy tail bound and the full hybrid Itô decomposition across switches.

Fix an admissible action sequence $a \in \mathcal{A}$ and work under the corresponding learned predictive law Q_a. Let

$$0 < \tau_1 < \cdots < \tau_N \le T \tag{139}$$

be the genuine mode-switching times up to horizon T, and set

$$\tau_0 := 0, \qquad \tau_{N+1} := T. \tag{140}$$

By assumption, $N < \infty$ almost surely. On each interval $[\tau_j, \tau_{j+1})$, the mode is constant; denote it by m_j.

We first derive a drift bound on a single constant-mode interval. Fix one such interval and abbreviate

$$H = H_{m_j}, \quad U = U_{m_j}, \quad J = J_{m_j}, \quad R = R_{m_j}, \quad G = G_{m_j}, \quad \Sigma = \Sigma_{m_j}. \tag{141}$$

For $z \in \mathcal{D}_{m_j}$ and a predictable action value a_t, the infinitesimal generator corresponding to Eq. (17) applied to H is

$$\mathcal{L}_{m_j}^{a_t} H(z) = \nabla H(z)^\top \Big[(J - R)\nabla H(z) + G a_t\Big] + \frac{1}{2}\operatorname{tr}\big(\Sigma \nabla^2 H(z)\big). \tag{142}$$

The skew-symmetry of J gives

$$\nabla H(z)^\top J \nabla H(z) = 0. \tag{143}$$

Since $R \succeq rI$,

$$-\nabla H(z)^\top R \nabla H(z) \le -r\|\nabla H(z)\|_2^2. \tag{144}$$

The input term is bounded by Cauchy's inequality and Young's inequality:

$$\nabla H(z)^\top G a_t \le \|\nabla H(z)\|_2 \, \|G\|_{\mathrm{op}} \|a_t\|_2 \le gA\|\nabla H(z)\|_2 \le \frac{r}{2}\|\nabla H(z)\|_2^2 + \frac{g^2 A^2}{2r}. \tag{145}$$

For the diffusion term, because $\Sigma \succeq 0$, $\operatorname{tr}(\Sigma) \le \nu^2$, and $\|\nabla^2 H(z)\|_{\mathrm{op}} \le L_H$, we have

$$\frac{1}{2}\operatorname{tr}\big(\Sigma \nabla^2 H(z)\big) \le \frac{1}{2}\big|\operatorname{tr}\big(\Sigma \nabla^2 H(z)\big)\big| \le \frac{1}{2}\|\nabla^2 H(z)\|_{\mathrm{op}} \operatorname{tr}(\Sigma) \le \frac{L_H \nu^2}{2}. \tag{146}$$

Combining Eqs. (142)–(146) yields

$$\mathcal{L}_{m_j}^{a_t} H(z) \le -\frac{r}{2}\|\nabla H(z)\|_2^2 + \frac{g^2 A^2}{2r} + \frac{L_H \nu^2}{2}. \tag{147}$$

Since $U(z) = H(z) - h_{m_j}^\star$ differs from $H(z)$ only by a constant,

$$\mathcal{L}_{m_j}^{a_t} U(z) = \mathcal{L}_{m_j}^{a_t} H(z). \tag{148}$$

By the Polyak–Lojasiewicz condition,

$$\|\nabla H(z)\|_2^2 \geq 2\mu U(z). \tag{149}$$

Therefore, for all $z \in \mathcal{D}_{m_j}$,

$$\mathcal{L}_{m_j}^{a_t} U(z) \leq -r\mu U(z) + \frac{g^2 A^2}{2r} + \frac{L_H \nu^2}{2} = -\alpha U(z) + C_E, \tag{150}$$

where α and C_E match their main-text definitions in Eq. (18).

We now sum Itô's formula over the constant-mode intervals. Since the rollout remains in the certified regions almost surely, the drift bound Eq. (150) applies along the path. On each interval $[\tau_j, \tau_{j+1})$, Itô's formula gives

$$U_{m_j}(z_{\tau_{j+1}^-}) = U_{m_j}(z_{\tau_j}) + \int_{\tau_j}^{\tau_{j+1}} \mathcal{L}_{m_j}^{a_t} U_{m_j}(z_t)\, dt + M_j, \tag{151}$$

where M_j is the stochastic integral over that interval. The assumed square-integrability condition implies that the sum of these stochastic integrals is a true martingale with mean zero; equivalently, one may localize on compact sublevel sets and then pass to the limit.

At a genuine switching time τ_j, define the active-energy jump

$$\Delta_j := U_{s_{\tau_j}}(z_{\tau_j}) - U_{s_{\tau_j^-}}(z_{\tau_j^-}). \tag{152}$$

By the passive-switch assumption,

$$\Delta_j \leq 0 \qquad \text{a.s.} \tag{153}$$

Summing over all intervals and genuine switches up to T yields the hybrid Itô decomposition

$$U_{s_T}(z_T) = U_{s_0}(z_0) + \int_0^T \mathcal{L}_{s_t}^{a_t} U_{s_t}(z_t)\, dt + M_T + \sum_{0<\tau_j \leq T} \Delta_j, \tag{154}$$

where $\mathbb{E}_{Q_a}[M_T] = 0$ and $\Delta_j \leq 0$ almost surely. Taking expectations and using Eq. (150) gives

$$\mathbb{E}_{Q_a}[U_{s_T}(z_T)] \leq \mathbb{E}_{Q_a}[U_{s_0}(z_0)] + \int_0^T \left(-\alpha\, \mathbb{E}_{Q_a}[U_{s_t}(z_t)] + C_E\right) dt. \tag{155}$$

Define

$$\psi(t) := \mathbb{E}_{Q_a}[U_{s_t}(z_t)]. \tag{156}$$

Then Eq. (155) reads

$$\psi(T) \leq \psi(0) + \int_0^T (-\alpha\psi(t) + C_E) dt. \tag{157}$$

By the integral form of Gronwall's inequality,

$$\psi(T) \leq e^{-\alpha T}\psi(0) + \frac{1 - e^{-\alpha T}}{\alpha} C_E. \tag{158}$$

Substituting back the definition of ψ proves

$$\mathbb{E}_{Q_a}[U_{s_T}(z_T)] \leq e^{-\alpha T}\mathbb{E}_{Q_a}[U_{s_0}(z_0)] + \frac{1 - e^{-\alpha T}}{\alpha} C_E, \tag{159}$$

which is the main-text energy certificate Eq. (19).

The following tail bound is more detailed than the compact main-text statement but is useful for experiments that measure high-energy failures. Define

$$\overline{U}_T(a) := e^{-\alpha T}\mathbb{E}_{Q_a}[U_{s_0}(z_0)] + \frac{1 - e^{-\alpha T}}{\alpha} C_E. \tag{160}$$

Since $U_{s_T}(z_T) \geq 0$, Markov's inequality gives, for every $B > 0$,

$$Q_a(U_{s_T}(z_T) \geq B) \leq \frac{\mathbb{E}_{Q_a}[U_{s_T}(z_T)]}{B} \leq \frac{\overline{U}_T(a)}{B}. \tag{161}$$

We now prove the KL-robust risk and MPC claims. Let

$$d_{\mathrm{TV}}(P,Q) := \sup_{\mathcal{B}} |P(\mathcal{B}) - Q(\mathcal{B})| \tag{162}$$

denote total variation distance. By Pinsker's inequality and the uniform predictive calibration assumption,

$$d_{\mathrm{TV}}(P_a^\star, Q_a) \le \sqrt{\frac{1}{2} D_{\mathrm{KL}}(P_a^\star \| Q_a)} \le \sqrt{\frac{\varepsilon_H}{2}}. \tag{163}$$

Since total variation is always at most one,

$$d_{\mathrm{TV}}(P_a^\star, Q_a) \le \epsilon_{\mathrm{TV}} := \min\left\{1, \sqrt{\frac{\varepsilon_H}{2}}\right\} \qquad \forall a \in \mathcal{A}. \tag{164}$$

Therefore, for every measurable event $\mathcal{F}$,

$$P_a^\star(\mathcal{F}) \le Q_a(\mathcal{F}) + d_{\mathrm{TV}}(P_a^\star, Q_a) \le Q_a(\mathcal{F}) + \epsilon_{\mathrm{TV}}. \tag{165}$$

Applying this to the high-energy event $\mathcal{E}_B = \{U_{s_T}(z_T) \ge B\}$ and using Eq. (161) gives the SI-level true high-energy certificate

$$P_a^\star(\mathcal{E}_B) \le \min\left\{1, \frac{\overline{U}_T(a)}{B} + \epsilon_{\mathrm{TV}}\right\}. \tag{166}$$

This bound is additional theory beyond the compact theorem statement, but it follows immediately from Eq. (19) and the KL calibration assumption.

Next, let $C(\tau) \in [0, C_{\max}]$ be any bounded measurable planning cost. Since $C/C_{\max} \in [0,1]$, the variational characterization of total variation gives

$$\left|\mathbb{E}_{P_a^\star}[C(\tau)] - \mathbb{E}_{Q_a}[C(\tau)]\right| \le C_{\max}\, d_{\mathrm{TV}}(P_a^\star, Q_a) \le C_{\max}\epsilon_{\mathrm{TV}} \qquad \forall a \in \mathcal{A}. \tag{167}$$

Using Eq. (167) for $\widehat{a}$ gives

$$\mathbb{E}_{P_{\widehat{a}}^\star}[C] \le \mathbb{E}_{Q_{\widehat{a}}}[C] + C_{\max}\epsilon_{\mathrm{TV}}. \tag{168}$$

By η_{opt}-optimality of $\widehat{a}$ under the VHYDRO predictive law,

$$\mathbb{E}_{Q_{\widehat{a}}}[C] \le \inf_{a\in\mathcal{A}} \mathbb{E}_{Q_a}[C] + \eta_{\mathrm{opt}}. \tag{169}$$

Moreover, for every $a \in \mathcal{A}$, Eq. (167) also gives

$$\mathbb{E}_{Q_a}[C] \le \mathbb{E}_{P_a^\star}[C] + C_{\max}\epsilon_{\mathrm{TV}}. \tag{170}$$

Taking the infimum over $a \in \mathcal{A}$ yields

$$\inf_{a\in\mathcal{A}} \mathbb{E}_{Q_a}[C] \le \inf_{a\in\mathcal{A}} \mathbb{E}_{P_a^\star}[C] + C_{\max}\epsilon_{\mathrm{TV}}. \tag{171}$$

Combining Eqs. (168), (169), and (171) gives

$$\mathbb{E}_{P_{\widehat{a}}^\star}[C] - \inf_{a\in\mathcal{A}} \mathbb{E}_{P_a^\star}[C] \le 2C_{\max}\epsilon_{\mathrm{TV}} + \eta_{\mathrm{opt}}, \tag{172}$$

which is the main-text MPC bound Eq. (20). Since $\epsilon_{\mathrm{TV}} \le \sqrt{\varepsilon_H/2}$, Eq. (172) also implies the slightly looser form

$$\mathbb{E}_{P_{\widehat{a}}^\star}[C] - \inf_{a\in\mathcal{A}} \mathbb{E}_{P_a^\star}[C] \le 2C_{\max}\sqrt{\frac{\varepsilon_H}{2}} + \eta_{\mathrm{opt}}. \tag{173}$$

If the true infimum is attained by $a^\star$, then

$$\inf_{a\in\mathcal{A}} \mathbb{E}_{P_a^\star}[C] = \mathbb{E}_{P_{a^\star}^\star}[C], \tag{174}$$

which gives the oracle-comparison form.

For the deterministic chance constraint, let $\mathcal{F}$ be any measurable failure event. If

$$Q_a(\mathcal{F}) \le \delta - \epsilon_{\mathrm{TV}}, \qquad \delta \in [\epsilon_{\mathrm{TV}}, 1], \tag{175}$$

then Eq. (165) gives

$$P_a^\star(\mathcal{F}) \le Q_a(\mathcal{F}) + \epsilon_{\mathrm{TV}} \le \delta. \tag{176}$$

This is the deterministic implication in Eq. (21).

It remains to prove the rollout certificate and connect it to Eq. (22). Let

$$\mathcal{A}_{\mathrm{eval}} = \{a^1, \dots, a^L\} \tag{177}$$

be the finite evaluated action set. For each $a \in \mathcal{A}_{\mathrm{eval}}$, define

$$Y_{a,i} := \mathbf{1}\{\tau_a^{(i)} \in \mathcal{F}\}, \qquad \tau_a^{(i)} \overset{\text{i.i.d.}}{\sim} Q_a. \tag{178}$$

Then $Y_{a,i} \in [0,1]$,

$$\mathbb{E}[Y_{a,i}] = Q_a(\mathcal{F}), \qquad \widehat{Q}_{a,n}(\mathcal{F}) = \frac{1}{n}\sum_{i=1}^{n} Y_{a,i}. \tag{179}$$

For a fixed a, Hoeffding's inequality gives

$$\Pr\Big(Q_a(\mathcal{F}) > \widehat{Q}_{a,n}(\mathcal{F}) + \epsilon_n(L,\beta)\Big) \le \exp\big(-2n\epsilon_n(L,\beta)^2\big) = \frac{\beta}{L}, \tag{180}$$

where

$$\epsilon_n(L,\beta) := \sqrt{\frac{\log(L/\beta)}{2n}}. \tag{181}$$

Taking a union bound over the L evaluated actions yields, with probability at least $1-\beta$, the simultaneous event

$$Q_a(\mathcal{F}) \le \widehat{Q}_{a,n}(\mathcal{F}) + \epsilon_n(L,\beta) \qquad \forall a \in \mathcal{A}_{\mathrm{eval}}. \tag{182}$$

On this simultaneous event, if the possibly data-dependent selected action $\widetilde{a} \in \mathcal{A}_{\mathrm{eval}}$ satisfies

$$\widehat{Q}_{\widetilde{a},n}(\mathcal{F}) \le \delta - \epsilon_{\mathrm{TV}} - \epsilon_n(L,\beta), \tag{183}$$

then

$$Q_{\widetilde{a}}(\mathcal{F}) \le \delta - \epsilon_{\mathrm{TV}}. \tag{184}$$

Applying the deterministic chance certificate in Eq. (21) to $\widetilde{a}$ gives

$$P^{\star}_{\widetilde{a}}(\mathcal{F}) \le \delta. \tag{185}$$

Equations (181)–(185) prove the Monte Carlo margin statement Eq. (22), completing the proof of Theorem 4. □

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction state the proposed support-safe hybrid variational filtering mechanism, the mode-segmentation evaluation, and the sparse-law recovery audit. The Positioning, Experiments, Conclusion, and Limitations sections explicitly bound the claims to matched internal ablations, proxy-label evaluation, and known-equation sparse-law recovery.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The paper includes a dedicated Conclusion and Limitations section. It states that closed-loop task-success dominance, real-robot transfer, external SOTA comparison, simulator contact-force ground truth, and real-robot sparse-law identification are not claimed.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The main text states Theorems 1–3 with their assumptions, and Appendix B provides full proofs. Appendix A.5 and Appendix B.4 state and prove the conditional planning theorem separately, while the main text clarifies that this theorem is not part of the empirical claim.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper, regardless of whether the code and data are provided or not?

Answer: [Yes]

Justification: The paper reports the tasks, metrics, seed count, ablations, matching protocol, occlusion settings, proxy-label construction, known-equation sparse-law audit, and compute disclosure in Section 3 and Appendix A. Code is included in this submission, which is addressed separately in Item 5.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The paper uses released benchmark datasets and describes the experimental protocol and the submission includes an anonymized code release and complete run scripts. The paper is not presented as a new open-source benchmark.

6. Experimental setting/details

Question: Does the paper specify all the training and test details necessary to understand the results?

Answer: [Yes]

Justification: The paper specifies the evaluation tasks, occlusion levels, ablations, metrics, seed count, held-out evaluation protocol, proxy-label construction, and known-equation sparse-law audit in Section 3 and Appendix A. Hyperparameters for proxy labels are selected on validation episodes and then frozen.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The paper reports means with SEM or standard deviation where appropriate and states that the main experiment cells use twenty independent random seeds. Figures and tables identify whether uncertainty is SEM or standard deviation.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources needed to reproduce the experiments?

Answer: [Yes]

Justification: The appendix reports the workstation hardware and gives detailed runtime estimates for the supplementary MPC audit. We are happy to provide further information regarding the running times of each experiment separately.

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics?

Answer: [Yes]

Justification: The work uses simulation/benchmark datasets and controlled known-equation systems, and does not involve deception, harmful deployment, human-subject experiments conducted by the authors, or release of high-risk models.

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [No]

Justification: The paper focuses on intermediate filtering, segmentation, and sparse-law recovery mechanisms and does not include a separate broader-impact discussion. It explicitly states that closed-loop control and real-robot deployment claims are out of scope.

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse?

Answer: [N/A]

Justification: The paper does not release a high-risk pretrained model, image generator, scraped dataset, or other artifact requiring controlled release safeguards.

12. Licenses for existing assets

Question: Are the creators or original owners of assets used in the paper properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [No]

Justification: The paper cites the creators of the benchmark datasets and simulators used, including ManiSkill, SAPIEN, BridgeData, RoboTurk, and related resources. However, the current text does not explicitly list the license names or terms of use for each asset.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The supplementary ZIP includes anonymized code as a new software asset, with a README describing dependencies, expected data inputs, run commands, and scripts for reproducing the reported experiments. No new dataset or model checkpoint is released.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation?

Answer: [N/A]

Justification: The authors do not conduct new crowdsourcing experiments or new human-subject studies in this work. The paper evaluates released manipulation datasets and controlled known-equation systems.

15. Institutional review board approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board approvals or an equivalent approval/review were obtained?

Answer: [N/A]

Justification: The paper does not report new human-subject experiments conducted by the authors. Therefore IRB or equivalent approval is not applicable to the work as submitted.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research?

Answer: [N/A]

Justification: LLMs are not an important, original, or non-standard component of the core method, experiments, or theoretical results. Any use limited to writing, editing, or formatting does not require declaration under this checklist item.