PDE-Constrained Learning-Driven Autonomy with Belief, Policy, and Actuation Uncertainty

Anonymous authors
Paper under double-blind review

Abstract

Autonomous systems in distributed physical environments must transform incomplete and stale observations into actions that remain useful under transition uncertainty and imperfect execution. We introduce PDE-Constrained Learning-Driven Autonomy (LDA), a closed-loop framework that separates observations, beliefs, candidate-conditioned predictive distributions, intended controls, executed actions, and feedback. LDA combines masks and measurement age with modality-specific random Fourier features, reliability-weighted kernel means, and temporal attention; conditions predictions on context, control history, and candidate plans; and evaluates plans with actuation-aware mean–CVaR control and soft safety tightening. Screening experiments on hourly traffic and irregular groundwater records, together with model-conditional traffic-control simulations, show that no representation dominates point prediction, calibration, risk discrimination, safety, and policy quality. At 30% traffic missingness, imputation gives the lowest belief-screening RMSE (0.0474), Mask+Age is closest to nominal 90% coverage (0.909), and reliability attention gives the strongest exceedance AUROC (0.980). Ridge gives the lowest traffic prediction RMSE, whereas last-state pooling gives stronger policy-oriented behavior. In simulation, Mean–CVaR MPC lowers mean and tail loss relative to expected-cost MPC, while executed-action modelling exposes command–execution trade-offs. Persistent sampling, transition, actuation, approximation, and numerical errors lead to an error-dependent neighborhood rather than exact optimality. The evidence supports modular beliefs and explicit actuator models.

1 Introduction

Learning-enabled autonomy is often evaluated by predictive accuracy, although deployment requires a complete loop from partial observations to physically executed actions. Traffic networks, groundwater systems, and irrigation processes exhibit missing, delayed, irregular, or stale measurements; control actions may also be attenuated, delayed, saturated, or dropped. Missingness indicators and elapsed observation times can be predictive in their own right (Che et al., 2018), while physics-informed learning can improve consistency with governing equations without resolving sensing or actuation uncertainty (Raissi et al., 2019).

We study these issues through *PDE-Constrained Learning-Driven Autonomy* (LDA):

$$\mathcal{O}_t \xrightarrow{\phi_\omega} \mathcal{B}_t \xrightarrow{\mathcal{F}_\vartheta} \mathcal{P}_t \xrightarrow{\Pi} \mathcal{U}_t \xrightarrow{g_\psi} \mathcal{A}_t \xrightarrow{\mathcal{H}} \mathcal{O}_{t+1}. \tag{1}$$

Here, $\mathcal{O}_t$ denotes incomplete observations, $\mathcal{B}_t$ a learned belief, $\mathcal{P}_t$ candidate-dependent predictive distributions, $\mathcal{U}_t$ intended controls, and $\mathcal{A}_t$ executed actions. The decomposition keeps observation, prediction, policy, and actuation errors distinct.

LDA uses random Fourier features (RFF) for scalable nonlinear representation (Rahimi & Recht, 2007), kernel mean embeddings (KME) for distributional summaries (Muandet et al., 2017), conditional latent-variable prediction (Kingma & Welling, 2014; Sohn et al., 2015), model predictive control (MPC) (Mayne et al., 2000), conditional value-at-risk (CVaR)

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

(Rockafellar & Uryasev, 2000), and soft control-barrier penalties (Ames et al., 2017; 2019). The main contributions are:

- a missingness-aware multimodal belief using masks, measurement age, RFF, reliability-weighted KME, and attention;
- a context–control–plan interface for candidate-conditioned trajectory and risk prediction;
- actuation-aware mean–CVaR MPC, branch-tail allocation, soft CVaR–CBF penalties, and actuator adaptation;
- an explicit decomposition of sampling, transition, actuation, approximation, and numerical errors; and
- a consolidated screening evaluation spanning belief construction, conditional prediction, control, and error sensitivity.

2 Related Work

Incomplete observations. Missing-data sequence models combine imputation, indicators, temporal decay, and recurrent updates (Che et al., 2018). LDA additionally lets reliability affect kernel aggregation, attention, risk assessment, and policy evaluation.

Kernel and attention representations. RFF approximate stationary kernels with explicit features (Rahimi & Recht, 2007), while KME map empirical distributions into reproducing-kernel Hilbert spaces (Muandet et al., 2017). Attention supplies data-dependent temporal aggregation (Vaswani et al., 2017). LDA combines these mechanisms while retaining source-specific reliability.

Probabilistic prediction and control. Conditional latent-variable models represent nonlinear, multimodal futures (Kingma & Welling, 2014; Sohn et al., 2015). MPC repeatedly optimizes finite-horizon decisions (Mayne et al., 2000); CVaR emphasizes adverse tails (Rockafellar & Uryasev, 2000); and CBFs encode forward-invariance conditions (Ames et al., 2017; 2019). Scenario methods provide finite-sample decision tools but depend on explicit assumptions and coverage conditions (Campi et al., 2021). LDA connects these components through candidate-dependent prediction and distinguishes commands from executed actions.

Approximate policy evaluation. Approximate dynamic programming uses finite-dimensional value representations and inexact improvement operators (Bertsekas & Tsitsiklis, 1996; Bertsekas, 2012). Persistent projection, sampling, model, and solver errors generally imply convergence to a neighborhood, not exact recovery.

3 Problem Formulation and Assumptions

3.1 Partially Observed Controlled Process

Let $x_t \in \mathcal{X}$ denote the latent physical state, $o_t = \{o_t^{(m)}\}_{m=1}^{M}$ the multimodal observations, $u_t \in \mathcal{U}$ the commanded control, and $a_t \in \mathcal{A}$ the executed action. The system satisfies

$$x_{t+1} = f_\star(x_t, a_t, \xi_t), \tag{2}$$

$$o_t^{(m)} = h_\star^{(m)}(x_t, \nu_t^{(m)}), \qquad m = 1, \dots, M, \tag{3}$$

$$a_t = g_\star(u_t, z_t, \eta_t), \tag{4}$$

where ξ_t, $\nu_t^{(m)}$, and η_t represent transition, observation, and actuation disturbances. The actuator condition z_t captures delay, efficiency loss, bias, saturation, degradation, and command failure. Distinguishing u_t from a_t is essential because the physical state evolves under the delivered action. For each modality, the availability mask $M_t^{(m)} \in \{0, 1\}$ and measurement age $A_t^{(m)} \geq 0$ augment the observed history:

$$\mathcal{I}_t = \{o_{0:t}, M_{0:t}, A_{0:t}, u_{0:t-1}, a_{0:t-1}\}, \qquad b_t = \phi_\omega(\mathcal{I}_t). \tag{5}$$

The learned belief b_t is a finite-dimensional summary of the available information, not an exact Bayesian posterior. For horizon H, let

$$\pi_t = \left(u_{t|t}, \ldots, u_{t+H-1|t}\right) \in \Pi_t \tag{6}$$

denote a candidate command sequence. Conditional state and execution scenarios are generated according to

$$(\widehat{x}_{t+1:t+H}, \widehat{a}_{t:t+H-1}) \sim p_\vartheta(\cdot \mid b_t, \pi_t). \tag{7}$$

The receding-horizon controller selects

$$\pi_t^\star \in \arg\min_{\pi \in \Pi_t} \left\{\mathbb{E}[J_t(\pi)] + \beta_{\text{risk}} \,\text{CVaR}_\alpha(J_t(\pi)) + \beta_{\text{sw}} R_{\text{sw}}(\pi)\right\}, \tag{8}$$

subject to action bounds, shared-resource constraints, switching restrictions, and implemented safety penalties. Only $u_{t|t}^\star$ is issued before the belief, transition surrogate, and actuator model are updated.

3.2 Standing Assumptions

The theoretical statements are conditional on the following assumptions:

1. **Boundedness:** observations, random features, reliability weights, candidate actions, and stage losses are bounded on the evaluated domains;
2. **Regularity:** the belief encoder, conditional predictor, transition surrogate, actuator model, and loss functions are locally Lipschitz;
3. **Positive reliability mass:** each normalized kernel-memory denominator is bounded away from zero;
4. **Finite implementation error:** sampling, transition, actuation, finite-dimensional approximation, and numerical errors admit finite bounds;
5. **Conditional feasibility:** at least one admissible command exists at each simulated decision time; and
6. **Contractive value update:** the ideal evaluation-improvement operator is contractive in a specified value-function norm.

These assumptions are not established from field interventions. In particular, the implemented barrier terms are soft penalties, so empirical feasibility does not imply recursive feasibility or forward invariance.

4 Method

4.1 Missingness-Aware Multi-Head Belief

For modality m and lag j, reliability-biased temporal attention uses

$$e_{t,j}^{(m)} = \frac{(Q_t^{(m)})^\top K_{t-j}^{(m)}}{\sqrt{d_m}} + \beta_q \log\left(q_{t-j}^{(m)} + \varepsilon\right) - \beta_a A_{t-j}^{(m)}, \tag{9}$$

with normalized weights and context

$$\alpha_{t,j}^{(m)} = \frac{\exp(e_{t,j}^{(m)})}{\sum_{\ell=0}^{L-1} \exp(e_{t,\ell}^{(m)})}, \tag{10}$$

$$c_t^{(m)} = \sum_{j=0}^{L-1} \alpha_{t,j}^{(m)} V_{t-j}^{(m)}. \tag{11}$$

The context belief combines the latest available values, masks, measurement ages, modality-specific RFF, weighted kernel means, and attention summaries:

$$b_t^{\text{ctx}} = \text{Fuse}\left(b_t^{\text{last}}, b_t^{\text{mask}}, b_t^{\text{age}}, \{\varphi_m\}_{m=1}^M, \{\mu_t^{(m)}\}_{m=1}^M, \{c_t^{(m)}\}_{m=1}^M\right). \tag{12}$$

Separate prediction, calibration, risk, and policy heads allow the shared belief to support different downstream objectives.

4.2 Belief–Control–Plan Conditioning

For each candidate plan, define

$$\chi_t(\pi_t) = \left[b_t^{\mathrm{ctx}}; b_t^{\mathrm{ctrl}}; b_t^{\mathrm{plan}}(\pi_t)\right]. \tag{13}$$

The control representation encodes command and execution histories, command–execution residuals, and estimated actuator conditions:

$$b_t^{\mathrm{ctrl}} = \mathrm{Enc}_{\mathrm{ctrl}}\left(u_{t-L_u:t-1}, a_{t-L_u:t-1}, a_{t-L_u:t-1} - u_{t-L_u:t-1}, z_{t-L_u:t-1}\right). \tag{14}$$

The candidate-plan representation summarizes the proposed commands and their predicted consequences:

$$b_t^{\mathrm{plan}}(\pi_t) = \mathrm{Enc}_{\mathrm{plan}}\left(\pi_t, \widehat{\mathcal{T}}_t(\pi_t), \widehat{J}_t(\pi_t), \widehat{\mathrm{CVaR}}_{\alpha,t}(\pi_t), \widehat{p}_{\mathrm{viol},t}(\pi_t)\right). \tag{15}$$

The conditional predictor then generates

$$\left(\widehat{x}_{t+1:t+H}^{(k)}, \widehat{a}_{t:t+H-1}^{(k)}\right) \sim p_\vartheta\left(\cdot \mid \chi_t(\pi_t)\right), \qquad k = 1, \ldots, K. \tag{16}$$

4.3 Actuator Modelling and Adaptation

A local actuator approximation is

$$a_t = \Gamma_t u_t + b_{a,t} + \eta_t, \tag{17}$$

where Γ_t captures execution efficiency and $b_{a,t}$ captures bias. Using $b_{a,t}$ avoids conflict with the belief notation b_t and attention context $c_t^{(m)}$. A projected adaptation step is

$$\widehat{\Gamma}_{t+1} = \Pi_{\mathcal{G}}\left[\widehat{\Gamma}_t + \gamma_t\left(a_t - \widehat{\Gamma}_t u_t - \widehat{b}_{a,t}\right) u_t^\top\right], \tag{18}$$

where $\mathcal{G}$ is the admissible parameter set and $\gamma_t > 0$ is the adaptation gain.

4.4 Risk-Aware Receding-Horizon Control

For scenario k, define

$$\begin{aligned} J_t^{(k)}(\pi) = &\sum_{h=1}^{H} \ell_x\left(\widehat{x}_{t+h}^{(k)}\right) \\ &+ \sum_{h=0}^{H-1}\left[\ell_u(u_{t+h|t}) + \ell_a\left(u_{t+h|t}, \widehat{a}_{t+h}^{(k)}\right)\right]. \end{aligned} \tag{19}$$

The actuation-aware mean–CVaR controller solves

$$\begin{aligned} \pi_t^\star \in \arg\min_{\pi\in\Pi_t} \Bigg\{ &\sum_{k=1}^{K} w_k J_t^{(k)}(\pi) \\ &+ \beta_{\mathrm{risk}} \widehat{\mathrm{CVaR}}_{\alpha,K}^{\,w}\left(\{J_t^{(k)}(\pi)\}_{k=1}^K\right) + \beta_{\mathrm{sw}} R_{\mathrm{sw}}(\pi)\Bigg\}. \end{aligned} \tag{20}$$

For branch-dependent risk allocation, the scenario set is partitioned into $\{\mathcal{K}_b\}_{b=1}^B$, with

$$\sum_{b=1}^{B} p_b \delta_b \le \delta_{\mathrm{tot}}, \tag{21}$$

where p_b and δ_b denote the estimated branch probability and branch-specific risk budget. Let

$$\mathcal{S} = \{x \in \mathcal{X} : h_q(x) \ge 0,\ q = 1, \ldots, Q\} \tag{22}$$

be the prescribed safe set. The tightened scenario condition is

$$h_q\left(\widehat{x}_{t+h+1}^{(k)}\right) - (1-\kappa_q) h_q\left(\widehat{x}_{t+h}^{(k)}\right) \ge \Delta_{t,h}^{(q)}. \tag{23}$$

The aggregate margin is

$$\Delta_{\mathrm{tot}} = \Delta_{\mathrm{state}} + \epsilon_{\mathrm{samp}} + \epsilon_{\mathrm{trans}} + \epsilon_{\mathrm{act}} + \epsilon_{\mathrm{approx}} + \epsilon_{\mathrm{num}}. \tag{24}$$

In the reported implementation, barrier and tightening terms are soft penalties. They encourage feasible behavior but do not certify forward invariance.

5 Theoretical Analysis and Convergence

The following statements are conditional on the assumptions in Section 3.

5.1 Belief Approximation

Let $\mu_t^{(m),\star}$ denote the complete-observation kernel mean and define the normalized reliability weights

$$\widetilde{w}_{t,j}^{(m)} = \frac{\lambda^j q_{t-j}^{(m)}}{\sum_{\ell=0}^{L-1} \lambda^\ell q_{t-\ell}^{(m)} + \varepsilon}. \tag{25}$$

If $\|\varphi_m(o)\|_2 \le B_{\varphi,m}$, then

$$\begin{aligned}\left\|\mu_t^{(m)} - \mu_t^{(m),\star}\right\|_2 \le \epsilon_{\mathrm{RFF}}^{(m)} + B_{\varphi,m} \left\|\widetilde{\boldsymbol{w}}_t^{(m)} - \boldsymbol{w}_t^{(m),\star}\right\|_1 \\ + \epsilon_{\mathrm{obs}}^{(m)}.\end{aligned} \tag{26}$$

The three terms represent finite RFF approximation, distortion of the temporal reliability distribution, and observation error.

5.2 Propagation to Predictive Risk

Assume the conditional predictor satisfies

$$W_1\left(p_\vartheta(\cdot \mid \chi), p_\vartheta(\cdot \mid \chi')\right) \le L_{\mathcal{F}} \left\|\chi - \chi'\right\|_2, \tag{27}$$

where W_1 is the Wasserstein-1 distance. If J is L_J-Lipschitz, then

$$\left|\mathbb{E}[J \mid \chi] - \mathbb{E}[J \mid \chi']\right| \le L_J L_{\mathcal{F}} \left\|\chi - \chi'\right\|_2. \tag{28}$$

Thus, context, control, and plan-encoding errors propagate explicitly into candidate evaluation. If

$$\left|J - \widehat{J}\right| \le \epsilon_J \qquad \text{almost surely,} \tag{29}$$

then monotonicity and translation equivariance of CVaR imply

$$\left|\mathrm{CVaR}_\alpha(J) - \mathrm{CVaR}_\alpha(\widehat{J})\right| \le \epsilon_J. \tag{30}$$

The bound is conditional on uniform loss approximation. A $(1-\alpha)^{-1}$ factor is required only when the available error controls the positive-part tail expectation rather than the complete loss.

5.3 Approximate Policy Improvement

Let T denote an ideal policy evaluation-improvement operator satisfying

$$\left\|TV - TV'\right\|_{\mathcal{V}} \le \rho_V \left\|V - V'\right\|_{\mathcal{V}}, \qquad 0 \le \rho_V < 1. \tag{31}$$

Suppose the implemented update is

$$V_{j+1} = TV_j + e_j, \qquad \|e_j\|_{\mathcal{V}} \le \epsilon_{\mathrm{tot}}. \tag{32}$$

Then

$$\begin{aligned}\left\|V_j - V^\star\right\|_{\mathcal{V}} \le \rho_V^j \left\|V_0 - V^\star\right\|_{\mathcal{V}} \\ + \frac{1-\rho_V^j}{1-\rho_V}\epsilon_{\mathrm{tot}},\end{aligned} \tag{33}$$

and consequently

$$\limsup_{j\to\infty} \left\|V_j - V^\star\right\|_{\mathcal{V}} \le \frac{\epsilon_{\mathrm{tot}}}{1-\rho_V}. \tag{34}$$

The aggregate update error may be decomposed as

$$\epsilon_{\mathrm{tot}} = C_{\mathrm{samp}}\epsilon_{\mathrm{samp}} + C_{\mathrm{trans}}\epsilon_{\mathrm{trans}} + C_{\mathrm{act}}\epsilon_{\mathrm{act}} + C_{\mathrm{approx}}\epsilon_{\mathrm{approx}} + C_{\mathrm{num}}\epsilon_{\mathrm{num}}. \tag{35}$$

The simulated policy-iteration curves illustrate an error-dependent plateau, but they do not identify T, ρ_V, $V^\star$, or the sensitivity constants. Equation equation 47 is therefore a conditional approximation result, not an established field-safety theorem.

5.4 Actuator Adaptation

Under correct local specification, bounded-variance noise, compact projection, and persistent excitation, the projected adaptation rule can converge when

$$\sum_{t=0}^{\infty} \gamma_t = \infty, \qquad \sum_{t=0}^{\infty} \gamma_t^2 < \infty. \tag{36}$$

Under persistent drift, bias, or misspecification, a more realistic interpretation is

$$\limsup_{t\to\infty} \left\| \widehat{\theta}_{a,t} - \theta_a^\star \right\| \le \epsilon_{\text{base}} + C_a \epsilon_{\text{act}}, \tag{37}$$

where ϵ_{base} captures finite-data, noise, approximation, and excitation effects. The simulations illustrate this neighborhood behavior but do not identify C_a.

6 Experimental Protocol

6.1 Datasets and Temporal Partitioning

We evaluate the predictive components on two heterogeneous physical-process datasets: 8,760 hourly traffic observations and 5,991 irregular groundwater measurements. Traffic records are ordered chronologically and partitioned into nonoverlapping training, validation, and test intervals. Groundwater sequences are constructed separately for each monitoring location and ordered by observation time so that measurements from unrelated locations are not combined into artificial temporal histories. All imputation statistics, scaling parameters, random-feature maps, predictive models, uncertainty calibrators, and decision thresholds are estimated from the training partition. Validation data are used for model selection, early stopping, and calibration. Test data remain unavailable until final evaluation. Temporal windows are created independently within each partition, preventing histories or targets from crossing a split boundary.

6.2 Missingness and Reliability Stress Tests

Naturally missing observations are retained. Additional pointwise and contiguous-block corruptions are applied to held-out histories at rates

$$\gamma_{\text{miss}} \in \{0, 0.10, 0.30, 0.50\}. \tag{38}$$

All competing representations receive identical corruption masks within each seed, enabling paired comparisons. The evaluated belief constructions include:

1. mean imputation;
2. mask and measurement-age encoding;
3. modality-specific random Fourier features;
4. reliability-weighted cumulative kernel mean embeddings;
5. reliability-biased temporal attention; and
6. fused context–control–plan representations.

Belief quality is evaluated using next-step RMSE, R^2, negative log-likelihood, empirical coverage and width of nominal 90% prediction intervals, adverse-state AUROC, Brier score, unsafe-acceptance rate, and model-conditional policy regret. These endpoints are reported separately because point accuracy, probabilistic calibration, risk discrimination, and decision quality measure different properties.

6.3 Predictive Models and Architectural Ablations

The prediction benchmark evaluates linear, tree-based, kernel, Gaussian sequence, ensemble, and conditional-generative surrogates. The candidate set includes Ridge regression, Random Forest, Extra Trees, quantile gradient boosting, a sparse Gaussian-process approximation, a stochastic-dropout MLP ensemble, GRU-Gaussian surrogates, CVAE surrogates, and

ablations of random Fourier features, kernel mean embeddings, temporal attention, multimodal fusion, and control conditioning. Several labels in the raw benchmark correspond to identical implementations and results. In particular, the attention-perception CVAE, the multimodal early-fusion GRU, and the Euclidean or mean-pooling ablations are consolidated in the principal tables. All models are evaluated on common held-out targets. Reported means and standard deviations are calculated over the three available model seeds, 11, 23, and 37. Configurations labelled CVAE or GRU are lightweight screening surrogates unless explicitly identified as fully optimized sequence models. The benchmark therefore evaluates representation and conditioning choices rather than claiming a definitive comparison of all possible CVAE, GRU, or Gaussian-process architectures.

6.4 Model-Conditional Control Simulation

Closed-loop evaluation is conducted in an explicitly defined simulator. Traffic observations anchor the normalized state scale and exogenous temporal forcing, while interventions, transition perturbations, actuator gain loss, delay, dropout, and execution noise remain simulated variables. The controller comparison includes:

1. Mean MPC;
2. Mean–CVaR MPC;
3. CVaR–CBF;
4. executed-action-aware CVaR–CBF; and
5. adaptive executed-action-aware CVaR–CBF.

The execution-aware controllers propagate

$$a_t \sim \widehat{p}_\eta \left(a \mid u_t, b_t^{\mathrm{ctrl}}\right) \tag{39}$$

inside the predictive scenarios rather than assuming $a_t = u_t$. The resulting execution-averaged value operator is

$$\overline{\mathcal{P}}_{\theta,\eta}^{u_t} V(x_t) = \mathbb{E}_{a_t|u_t,b_t^{\mathrm{ctrl}}} \mathbb{E}_{x_{t+1}|x_t,a_t} \left[V(x_{t+1})\right]. \tag{40}$$

Controller comparisons use common initial states, disturbances, actuator-fault realizations, and random seeds. Primary endpoints are mean realized cost, empirical $\mathrm{CVaR}_{0.90}$, constraint-violation rate and magnitude, command–execution mismatch, switching frequency, and actuator parameter error.

6.5 Error, Approximation, and Scenario Sensitivity

The control benchmark varies combined transition and actuator mismatch over

$$\epsilon \in \{0, 0.10, 0.20, 0.30, 0.40\} \tag{41}$$

and the predictive scenario budget over

$$N_s \in \{16, 32, 64, 128, 256\}. \tag{42}$$

The total robust margin is decomposed as

$$m_{\mathrm{tot}} = m_{\mathrm{samp}} + m_{\mathrm{trans}} + m_{\mathrm{act}} + m_{\mathrm{belief}} + m_{\mathrm{approx}} + m_{\mathrm{sol}}, \tag{43}$$

where the finite-sampling contribution satisfies

$$m_{\mathrm{samp}} = O\left(\sqrt{\frac{\log(1/\delta)}{N_s}}\right). \tag{44}$$

Kernel–Galerkin sensitivity is evaluated by varying the random-feature dimension and regularization level. For

$$V_D(x) = \theta^\top \phi_D(x),$$

the regularized generalized HJB residual is

$$\widehat{\theta} = \arg\min_\theta \frac{1}{n} \sum_{i=1}^{n} \left[\ell_i + \overline{\mathcal{L}}^\pi V_D(x_i)\right]^2 + \lambda \|\theta\|_2^2. \tag{45}$$

Table 1: Leading unique predictive configurations. Values are means over three model seeds; RMSE uncertainty is reported as one standard deviation. Coverage is evaluated against the nominal level of 0.90. Lower RMSE, NLL, and Brier score are preferred; higher coverage-calibrated accuracy, AUROC, and R^2 are preferred. Rankings are based only on RMSE.

Dataset	Method	RMSE	NLL	$\text{Cov.}_{.90}$	AUROC	Brier	R^2	RMSE rank
Groundwater	No RFF: learned features	**0.6534** $\pm$0.0084	**0.9912**	0.9643	**0.8620**	**0.1336**	**0.2157**	1.0
Groundwater	No KME: Euclidean pooling	0.6557 $\pm$0.0084	0.9949	**0.9652**	0.8599	0.1340	0.2101	2.5
Groundwater	No attention: mean pooling	0.6557 $\pm$0.0084	0.9949	**0.9652**	0.8599	0.1340	0.2101	2.5
Groundwater	Attention-perception CVAE surrogate	0.6603 $\pm$0.0070	0.9935	0.9548	0.8145	0.1404	0.1990	5.0
Groundwater	Early-fusion GRU-Gaussian surrogate	0.6642 $\pm$0.0059	1.0015	0.9405	0.7850	0.1300	0.1894	7.5
Traffic	Early-fusion GRU-Gaussian surrogate	**0.8017** $\pm$0.0031	**1.2231**	0.9079	**0.7982**	**0.1444**	**0.3615**	1.5
Traffic	Attention-perception CVAE surrogate	0.8062 $\pm$0.0035	1.2286	0.9098	0.7781	0.1471	0.3544	4.0
Traffic	Full KME+RFF+attention CVAE surrogate	0.8062 $\pm$0.0035	1.2286	0.9098	0.7781	0.1471	0.3544	4.0
Traffic	Extra Trees	0.8112 $\pm$0.0075	2.0043	0.6970	0.7797	0.1446	0.3464	6.0
Traffic	Random Forest	0.8123 $\pm$0.0047	1.4987	0.8000	0.7806	0.1455	0.3446	7.0

Table 2: Controller comparison at mismatch 0.3 and $N_s = 64$. Lower values are preferred. Values are seed-level point summaries; confidence intervals and paired tests should be generated from the released per-seed trajectories.

Controller	Mean cost	$\text{CVaR}_{.90}$	Viol. rate	$\lvert u - a \rvert$
Mean MPC	1.0482	5.3822	0.2492	**0.0387**
Mean–CVaR MPC	1.0506	5.3863	0.2492	0.0399
CVaR–CBF	1.0498	5.3834	0.2492	0.0397
Executed-aware CVaR–CBF	**0.9846**	**5.3306**	**0.1975**	0.0546
Adaptive executed-aware CVaR–CBF	**0.9846**	**5.3306**	**0.1975**	0.0546

7 Results

7.1 Predictive Benchmark

Table 1 reports the leading unique configurations after consolidating duplicate aliases. No single architecture dominates point accuracy, calibration, adverse-state discrimination, and probabilistic score. For groundwater, learned features without RFF achieve the lowest RMSE, 0.6534, the lowest NLL, 0.9912, the highest AUROC, 0.8620, and the highest R^2, 0.2157, among the displayed configurations. Euclidean and mean pooling are statistically and numerically tied in the reported implementation, attaining slightly higher RMSE but the highest coverage, 0.9652. For traffic, the early-fusion GRU-Gaussian surrogate gives the lowest RMSE, 0.8017, and the strongest displayed R^2, 0.3615. The attention-perception and full KME+RFF+attention CVAE rows are identical, indicating that the nominally different labels do not correspond to distinguishable implementations in the current screening output. The traffic results therefore support early fusion in this benchmark but do not isolate the contribution of KME, RFF, or attention.

7.2 Risk-Aware Control and Actuation

Table 2 summarizes the latest model-conditional control experiment at combined transition and actuator mismatch 0.3 and 64 predictive scenarios. Executed-action-aware propagation gives the strongest observed aggregate performance in this configuration. Relative to Mean MPC, mean cost decreases from 1.0482 to 0.9846, a reduction of approximately 6.1%. Empirical $\text{CVaR}_{0.90}$ decreases from 5.3822 to 5.3306, while the violation rate decreases from 0.2492 to 0.1975, a relative reduction of approximately 20.7%. The lower cost and violation rate are accompanied by a larger observed command–execution mismatch. This does not mean that the execution-aware controller worsens the actuator. Rather, the execution-aware controller selects commands under an explicit stochastic execution model, and the reported mismatch reflects the resulting action distribution as well as the actuator fault process.

7.3 Finite-Scenario and Error Sensitivity

Increasing the scenario budget reduces the finite-sampling contribution to the robust margin but does not guarantee monotone improvement in realized cost or violations. At mismatch 0.3, the prescribed total margin decreases from approximately 0.163 with 16 scenarios to 0.129 with 256 scenarios. However, the empirical tail costs and violation rates remain nonmonotone, with overlapping seed-level uncertainty. This behavior is consistent with the decomposition in Equation equation 43. Increasing N_s reduces m_{samp}, but it does not remove transition-model bias, actuator-model misspecification, finite-dimensional approximation error, or solver error. Scenario count should therefore be selected jointly with real-time computational constraints rather than interpreted as a universal safety parameter.

7.4 Conditional Error-Neighborhood Interpretation

Suppose the ideal policy evaluation-improvement operator T is ρ_V-contractive and the implemented iteration satisfies

$$V_{j+1} = TV_j + e_j, \qquad \|e_j\|_{\mathcal{V}} \le \epsilon_{\mathrm{tot}}. \tag{46}$$

Then

$$\limsup_{j\to\infty} \|V_j - V^\star\|_{\mathcal{V}} \le \frac{\epsilon_{\mathrm{tot}}}{1-\rho_V}, \tag{47}$$

where

$$\epsilon_{\mathrm{tot}} = C_b\epsilon_{\mathrm{belief}} + C_s\epsilon_{\mathrm{samp}} + C_f\epsilon_{\mathrm{trans}} + C_a\epsilon_{\mathrm{act}} + C_D\epsilon_{\mathrm{approx}} + C_n\epsilon_{\mathrm{num}}. \tag{48}$$

Persistent sampling, transition, actuation, approximation, or numerical error therefore implies convergence to a nonzero neighborhood rather than exact optimality. Equation equation 47 is conditional: the experiments do not identify the physical Bellman operator, estimate ρ_V, or establish that the true system lies inside the assumed ambiguity sets.

7.5 Integrated Findings

The experiments support four qualified conclusions:

1. **Architectural complexity is not uniformly beneficial.** Groundwater performs best after removing RFF, while traffic favors the early-fusion GRU-Gaussian surrogate.
2. **Prediction metrics must be interpreted jointly.** RMSE, likelihood, coverage, AUROC, Brier score, and R^2 produce different rankings and measure different aspects of predictive quality.
3. **Execution-aware prediction is more consequential than the tested standalone risk penalties.** In the latest control configuration, execution-aware CVaR–CBF reduces mean cost and violations, whereas Mean–CVaR MPC and conventional CVaR–CBF remain close to Mean MPC.
4. **Persistent implementation error limits asymptotic performance.** More scenarios reduce sampling uncertainty but cannot eliminate model, actuator, approximation, or solver errors.

Limitations. The predictive benchmark uses only three model seeds, and several nominally distinct configurations produce identical outputs, limiting architectural attribution. The control results are model-conditional stress tests based on surrogate transitions, simulated interventions, and synthetic actuator faults. The observational datasets do not provide randomized interventions, synchronized command–execution telemetry, or measured counterfactual trajectories. Consequently, confidence intervals over simulation seeds describe finite-run variability within the declared simulator, not uncertainty about unobserved physical dynamics. Soft CVaR–CBF penalties do not establish recursive feasibility or forward invariance, and additive margins may be conservative or double-count correlated errors. Physical deployment requires measured executed actions, calibrated safety limits, validated uncertainty sets, feasible fallback logic, hardware-in-the-loop testing, and prospective intervention evaluation.

8 Conclusion

LDA provides an observation-to-execution framework for incomplete sensing, probabilistic prediction, risk-aware planning, and imperfect actuation. The screening results show objective-dependent representation quality: simple methods are competitive for point prediction and calibration, while nonlinear kernel and attention features can improve risk ranking. Mean–CVaR MPC improves average and tail loss relative to expected-cost control in the specified simulation, and executed-action modelling reveals trade-offs hidden by command-only evaluation. Persistent errors imply practical convergence to a neighborhood rather than exact optimality. Deployment requires fully trained sequence models, measured actions, repeated paired evaluation, and intervention data collected under verified control protocols.

References

Aaron D. Ames, Xiangru Xu, Jessy W. Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2017. doi: 10.1109/TAC.2016.2638961.

Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *European Control Conference*, pp. 3420–3431, 2019. doi: 10.23919/ECC.2019.8796030.

Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, Volume II*. Athena Scientific, 4 edition, 2012.

Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.

Marco C. Campi, Algo Carè, and Simone Garatti. The scenario approach: A tool at the service of data-driven decision making. *Annual Reviews in Control*, 52:1–17, 2021. doi: 10.1016/j.arcontrol.2021.10.004.

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8: 6085, 2018. doi: 10.1038/s41598-018-24271-9.

Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.

David Q. Mayne, James B. Rawlings, Christopher V. Rao, and Pierre O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000. doi: 10.1016/S0005-1098(99)00214-9.

Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends in Machine Learning*, 10(1–2):1–141, 2017. doi: 10.1561/2200000060.

Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, volume 20, 2007.

Maziar Raissi, Paris Perdikaris, and George Em 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. doi: 10.1016/j.jcp.2018.10.045.

R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000. doi: 10.21314/JOR.2000.038.

Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, volume 28, pp. 3483–3491, 2015.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

A Supplementary Definitions

A.1 Observation Reliability

For modality m, define

$$q_t^{(m)} = M_t^{(m)} r_t^{(m)} \exp\left(-\frac{A_t^{(m)}}{\tau_m}\right) \in [0,1], \tag{49}$$

where $M_t^{(m)} \in \{0,1\}$ is the availability mask, $A_t^{(m)} \geq 0$ is the measurement age, $r_t^{(m)} \in [0,1]$ is an optional source-quality estimate, and $\tau_m > 0$ controls temporal decay. Missing measurements receive zero weight; available but stale or degraded measurements are downweighted.

A.2 Random Fourier Features and Kernel Memory

For a shift-invariant kernel k_m, random Fourier features approximate kernel evaluations as

$$k_m(o, o') \approx \varphi_m(o)^\top \varphi_m(o'), \tag{50}$$

with

$$\varphi_m(o) = \sqrt{\frac{2}{D_m}} \left[\cos\left(\omega_{m,d}^\top o + \zeta_{m,d}\right)\right]_{d=1}^{D_m}. \tag{51}$$

The frequencies $\omega_{m,d}$ are sampled from the kernel spectral distribution, and $\zeta_{m,d}$ are independent random phases. For history length L and forgetting factor $\lambda \in (0,1]$, define

$$\mu_t^{(m)} = \frac{\sum_{j=0}^{L-1} \lambda^j q_{t-j}^{(m)} \varphi_m(o_{t-j}^{(m)})}{\sum_{j=0}^{L-1} \lambda^j q_{t-j}^{(m)} + \varepsilon}, \qquad \varepsilon > 0. \tag{52}$$

This representation summarizes recent feature distributions while reducing the contribution of old, missing, or unreliable observations.

A.3 Conditional Value-at-Risk

For loss J and confidence level $\alpha \in (0,1)$,

$$\mathrm{CVaR}_\alpha(J) = \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha} \mathbb{E}[(J-\eta)_+] \right\}. \tag{53}$$

For N_s scenarios with normalized weights $w_k \geq 0$ and $\sum_{k=1}^{N_s} w_k = 1$,

$$\widehat{\mathrm{CVaR}}_{\alpha,N_s}^{w}\left(\{J^{(k)}\}_{k=1}^{N_s}\right) = \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha} \sum_{k=1}^{N_s} w_k (J^{(k)} - \eta)_+ \right\}. \tag{54}$$

B Complete Screening Results

This section reports the completed screening results for missingness-aware belief construction, triad-conditioned prediction, and model-conditional risk-aware control. Metrics are interpreted only within their experimental stage. Predictive unsafe acceptance and closed-loop constraint violation are distinct events. A zero empirical event rate means that no event is observed in the finite sample; it does not imply zero operational risk.

Table 3: Missingness-aware belief screening at 30% injected missingness. "Unsafe" denotes predictive unsafe acceptance, not a closed-loop violation. Coverage is assessed relative to the nominal value 0.90.

Dataset	Representation	RMSE ↓	R^2 ↑	Cov.$_{.90}$	AUROC ↑	Unsafe ↓
Traffic	Imputation	0.0474	0.7166	0.8828	0.8149	0.0072
Traffic	Mask+Age	0.0510	0.6723	0.9085	0.8249	0.0062
Traffic	Modality RFF	0.0577	0.5801	0.8972	0.9781	0.0010
Traffic	Weighted CKME	0.0835	0.1199	0.8469	0.9532	0.0000
Traffic	Reliability attention	0.0829	0.1324	0.8623	0.9795	0.0000
Traffic	Fused belief triad	0.0553	0.6146	0.8664	0.9642	0.0010
Groundwater	Imputation	4.6584	–	0.8770	0.9090	0.0090
Groundwater	Mask+Age	4.7257	–	0.8710	0.9090	0.0090
Groundwater	Modality RFF	6.3394	–	0.7910	0.8530	0.0100
Groundwater	Weighted CKME	8.2944	–	0.7610	0.8400	0.0090
Groundwater	Reliability attention	14.8824	–	0.4310	0.9000	0.0200
Groundwater	Fused belief triad	9.4309	–	0.7010	0.5830	0.0110

B.1 Missingness-Aware Belief Screening

Traffic exhibits metric-dependent trade-offs: imputation minimizes point error, Mask+Age is closest to nominal coverage, and RFF or reliability attention maximizes adverse-state discrimination. Groundwater instead favors imputation and Mask+Age, showing that nonlinear features and temporal aggregation depend on the sampling regime and missingness structure.

B.2 Triad-Conditioned Prediction

Table 4 reports all unique predictive output groups from the completed three-seed benchmark. Exact-output aliases are consolidated only when all reported means and standard deviations match. Coverage is paired with interval width because high coverage may result from conservative intervals.

Table 4: Complete triad-conditioned prediction results after consolidating exact-output aliases. Cells report mean with standard deviation below over three model seeds. Coverage is assessed relative to 0.90; interval width is compared only at similar coverage.

Canonical configuration	RMSE	NLL	Cov.$_{.90}$	Width	AUROC	Brier	R^2
Groundwater							
No RFF: learned features	0.6534 ± 0.0084	0.9912 ± 0.0074	0.9643 ± 0.0021	2.4487 ± 0.0077	0.8620 ± 0.0250	0.1336 ± 0.0016	0.2157 ± 0.0131
No KME / no attention[a]	0.6557 ± 0.0084	0.9949 ± 0.0075	0.9652 ± 0.0019	2.4641 ± 0.0068	0.8599 ± 0.0283	0.1340 ± 0.0019	0.2101 ± 0.0133
Attention-perception CVAE group[b]	0.6603 ± 0.0070	0.9935 ± 0.0077	0.9548 ± 0.0070	2.3497 ± 0.0229	0.8145 ± 0.0284	0.1404 ± 0.0096	0.1990 ± 0.0102
Early-fusion GRU-Gaussian group[c]	0.6642 ± 0.0059	1.0015 ± 0.0083	0.9405 ± 0.0046	2.3882 ± 0.0040	0.7850 ± 0.0304	0.1300 ± 0.0107	0.1894 ± 0.0117
Ridge	0.6751 ± 0.0125	1.0228 ± 0.0150	0.9357 ± 0.0050	2.4160 ± 0.0139	0.8502 ± 0.0087	0.1343 ± 0.0021	0.1627 ± 0.0240
Stochastic-dropout MLP	0.6890 ± 0.0084	1.0260 ± 0.0129	0.9280 ± 0.0065	2.2211 ± 0.0115	0.7409 ± 0.0187	0.1325 ± 0.0054	0.1279 ± 0.0157
Control-conditioned CVAE group[d]	0.6895 ± 0.0058	1.0304 ± 0.0039	0.9340 ± 0.0070	2.3068 ± 0.0139	0.7277 ± 0.0541	0.1363 ± 0.0077	0.1264 ± 0.0113
Quantile GBDR	0.6989 ± 0.0070	1.0739 ± 0.0138	0.9232 ± 0.0056	2.2863 ± 0.0020	0.8786 ± 0.0184	0.1366 ± 0.0017	0.1023 ± 0.0259
Extra Trees	0.7128 ± 0.0049	2.3211 ± 0.0173	0.6508 ± 0.0041	1.0818 ± 0.0042	0.7955 ± 0.0456	0.1492 ± 0.0028	0.0664 ± 0.0122
Random Forest	0.7215 ± 0.0060	1.5096 ± 0.0197	0.7090 ± 0.0110	1.4520 ± 0.0120	0.8212 ± 0.0266	0.1470 ± 0.0017	0.0436 ± 0.0134
Last-state pooling	0.7365 ± 0.0076	1.1153 ± 0.0052	0.9194 ± 0.0027	2.6235 ± 0.0183	0.8409 ± 0.0121	0.1388 ± 0.0018	0.0032 ± 0.0217
Unimodal GRU-Gaussian	0.7507 ± 0.0122	1.1350 ± 0.0095	0.9141 ± 0.0055	2.5861 ± 0.0147	0.8468 ± 0.0296	0.1398 ± 0.0010	-0.0356 ± 0.0322
Sparse GP approximation	1.0203 ± 0.0024	1.4706 ± 0.0049	0.7936 ± 0.0032	3.1755 ± 0.0114	0.5411 ± 0.0236	0.1753 ± 0.0006	-0.9129 ± 0.0263
Traffic							
Early-fusion GRU-Gaussian group[c]	0.8017 ± 0.0031	1.2231 ± 0.0053	0.9079 ± 0.0026	2.4954 ± 0.0136	0.7982 ± 0.0151	0.1444 ± 0.0057	0.3615 ± 0.0052
Attention-perception CVAE group[b]	0.8062 ± 0.0035	1.2286 ± 0.0058	0.9098 ± 0.0025	2.4862 ± 0.0169	0.7781 ± 0.0102	0.1471 ± 0.0026	0.3544 ± 0.0059
Extra Trees	0.8112 ± 0.0075	2.0043 ± 0.0442	0.6970 ± 0.0020	1.3435 ± 0.0009	0.7797 ± 0.0090	0.1446 ± 0.0051	0.3464 ± 0.0030
Random Forest	0.8123 ± 0.0047	1.4987 ± 0.0163	0.8000 ± 0.0001	1.7230 ± 0.0021	0.7806 ± 0.0105	0.1455 ± 0.0053	0.3446 ± 0.0026
Unimodal GRU-Gaussian	0.8176 ± 0.0073	1.2381 ± 0.0087	0.9150 ± 0.0012	2.6002 ± 0.0077	0.7902 ± 0.0148	0.1449 ± 0.0055	0.3360 ± 0.0102
Quantile GBDR	0.8177 ± 0.0091	1.2502 ± 0.0131	0.9077 ± 0.0007	2.4871 ± 0.0005	0.7589 ± 0.0124	0.1466 ± 0.0046	0.3359 ± 0.0049
Control-conditioned CVAE group[d]	0.8227 ± 0.0022	1.2485 ± 0.0032	0.9010 ± 0.0065	2.4747 ± 0.0476	0.7834 ± 0.0145	0.1448 ± 0.0072	0.3277 ± 0.0080
No RFF: learned features	0.8236 ± 0.0060	1.2472 ± 0.0076	0.9147 ± 0.0016	2.6182 ± 0.0058	0.7728 ± 0.0100	0.1480 ± 0.0040	0.3263 ± 0.0030
No KME / no attention[a]	0.8266 ± 0.0075	1.2501 ± 0.0093	0.9136 ± 0.0009	2.6342 ± 0.0055	0.7749 ± 0.0081	0.1481 ± 0.0043	0.3213 ± 0.0031
Last-state pooling	0.8418 ± 0.0050	1.2664 ± 0.0059	0.9126 ± 0.0006	2.6830 ± 0.0033	0.7600 ± 0.0145	0.1484 ± 0.0043	0.2961 ± 0.0023
Ridge	0.8446 ± 0.0071	1.2814 ± 0.0097	0.8999 ± 0.0013	2.5418 ± 0.0032	0.7635 ± 0.0102	0.1499 ± 0.0046	0.2914 ± 0.0060

Continued on the next page

Table 4 continued from the previous page

Canonical configuration	RMSE	NLL	Cov.$_{.90}$	Width	AUROC	Brier	R^2
Stochastic-dropout MLP	0.8553 ± 0.0067	1.3090 ± 0.0123	0.8801 ± 0.0031	2.3930 ± 0.0298	0.7717 ± 0.0204	0.1453 ± 0.0065	0.2733 ± 0.0007
Sparse GP approximation	1.0306 ± 0.0067	1.4719 ± 0.0076	0.9121 ± 0.0007	3.3364 ± 0.0032	0.5002 ± 0.0003	0.1572 ± 0.0094	-0.0550 ± 0.0023

[a] Exact-output group: *No KME: Euclidean pooling* and *No attention: mean pooling.*
[b] Exact-output group: *Attention Perception-Only CVAE*, *Attention perception CVAE*, and *Full KME+RFF+attention CVAE.*
[c] Exact-output group: *Early-fusion GRU-Gaussian* and *Multimodal Early-Fusion GRU-Gaussian.*
[d] Exact-output group: *Control-conditioned CVAE* and *Full perception-control CVAE.*

Table 5: Model-conditional controller screening at mismatch 0.3 and $N_s = 64$. Lower values are preferred. The adaptive and non-adaptive execution-aware variants produce identical aggregate results.

Controller	Mean cost	$\text{CVaR}_{.90}$	Viol. rate	$\lvert u - a \rvert$
Mean MPC	1.0482	5.3822	0.2492	0.0387
Mean–CVaR MPC	1.0506	5.3863	0.2492	0.0399
CVaR–CBF	1.0498	5.3834	0.2492	0.0397
Executed-action-aware CVaR–CBF	0.9846	5.3306	0.1975	0.0546
Adaptive executed-action-aware CVaR–CBF	0.9846	5.3306	0.1975	0.0546

Table 6: Componentwise error screening at perturbation level 0.05. Results use the small-scale cost definition and are not numerically comparable with Table 5.

Error source	Mean loss	$\text{CVaR}_{.90}$	Viol. rate
Sampling	0.0116	0.0437	0.0750
Transition model	0.0178	0.0589	0.1250
Actuator model	0.0091	0.0393	0.0500
Finite approximation	0.0108	0.0394	0.0500
Numerical solver	0.0109	0.0420	0.0875

The prediction results reveal dataset-dependent architectural preferences. Traffic favors multimodal early fusion, whereas groundwater favors learned features without RFF and simpler pooling. Exact-output groups require audit of configuration flags, computation graphs, checkpoint loading, and cached predictions before architectural attribution.

B.3 Model-Conditional Controller Screening

Relative to Mean MPC, executed-action-aware CVaR–CBF reduces mean cost by approximately 6.1% and violation rate by approximately 20.7%, while increasing command–execution mismatch. The identical adaptive and non-adaptive results provide no evidence of an adaptation benefit in this configuration and require implementation audit.

B.4 Componentwise Error Screening

Transition-model error produces the largest mean loss, tail loss, and violation rate at the largest tested perturbation. Finite-sample response need not be monotone even when an analytical upper bound is monotone in its assumed error budget.

C Ablation and Confirmatory Evaluation

Representation ablations. The prediction study removes or replaces masks, measurement age, RFF, KME, attention, control history, executed-action residuals, candidate-plan conditioning, and uncertainty conditioning while holding temporal partitions, corruption masks, training budgets, seeds, and targets fixed.

Controller ablations. The controller study progresses from Mean MPC to Mean–CVaR MPC, branch-dependent tail allocation, CVaR–CBF, executed-action propagation, uncertainty tightening, actuator adaptation, and kernel–Galerkin approximate policy iteration. Comparisons use matched initial states, disturbances, fault trajectories, candidate sets, and seeds.

Confirmatory inference. RMSE is the primary prediction endpoint, and empirical $\text{CVaR}_{0.90}$ is the primary controller endpoint. Prediction comparisons require paired forecasts and moving-block bootstrap intervals. Controller inference uses the episode or simulator seed as the paired unit. CVaR comparisons use nested bootstrap resampling. Multiplicity correction applies within prespecified metric families. Superiority is emphasized only when the paired interval excludes zero, the adjusted p-value is below 0.05, and the effect exceeds a prespecified practical threshold.

D Complexity and Generalization

Let L be the observation-window length, M the number of modalities, D_m the RFF dimension, d the attention dimension, N_s the scenario count, H the horizon, and N_π the number of candidate plans. Belief construction costs

$$\mathcal{O}\left(L \sum_{m=1}^{M} D_m + MLd\right), \tag{55}$$

and plan evaluation costs

$$\mathcal{O}\left(N_\pi N_s H C_{\mathcal{F}}\right). \tag{56}$$

Scenario rollouts can be parallelized. Explicit scenario storage costs

$$\mathcal{O}\left(N_s H (d_x + d_a)\right), \tag{57}$$

although streaming evaluation reduces memory.

Generalization must be assessed across future periods, unseen locations, physical regimes, missingness mechanisms, sensor reliability, policy-induced state distributions, and actuator degradation. If

$$\sup_{\pi \in \Pi_t} |J_t(\pi) - \widehat{J}_t(\pi)| \le \epsilon_{\text{plan}}, \tag{58}$$

and $\widehat{\pi}_t$ minimizes $\widehat{J}_t$, then

$$J_t(\widehat{\pi}_t) - J_t(\pi_t^\star) \le 2\epsilon_{\text{plan}}. \tag{59}$$

This motivates uniform validation over candidate plans rather than evaluation only at historically observed controls.

E Reproducibility and Evidence Boundaries

The release records data and code hashes; temporal and station-level splits; missingness masks and seeds; preprocessing and calibration parameters; raw and canonical model identifiers; architecture and checkpoint hashes; per-window predictions; controller settings; per-step command and execution trajectories; statistical procedures; failures; software versions; hardware; runtime; and latency. All tables and figures are generated from versioned, machine-readable outputs. Observed, imputed, derived, predicted, and simulated quantities are labelled separately.

Deployment boundary. The experiments provide retrospective predictive evidence and model-conditional control simulations. Deployment requires calibrated physical thresholds, synchronized command–execution telemetry, verified actuator semantics, feasible emergency fallbacks, hardware-in-the-loop testing, and prospective intervention evaluation. Unverified feasibility, contraction, excitation, or error-bound assumptions preclude unconditional safety or exact optimality claims.