TRIAD-CONDITIONED FORECASTING FOR RISK-AWARE IRRIGATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Autonomous irrigation operates under incomplete sensing, measurement staleness, uncertain forecasts, and imperfect actuation. We study a triad-conditioned framework that separates environmental belief, executed-control history, and candidate-plan information. The environmental state is represented by adaptively routing among a direct temporal expert, a reliability-weighted random-feature kernel-mean expert, and a conservative fallback. The routed belief conditions conditional variational autoencoders (CVAEs) and recurrent Gaussian predictors. Because tail-risk control is meaningful only for calibrated scenarios, forecasts pass through chronological conformal calibration and state- and loss-space readiness checks before entering executed-action-aware mean–CVaR model predictive control (MPC). Experiments show that complexity is not uniformly beneficial: a unimodal GRU–Gaussian attains the best neural RMSE (0.1939) and energy score (0.3214), Random Forest gives the best overall point RMSE (0.1854), and the No-KME CVAE outperforms the full KME–RFF–attention CVAE. Principal CVAEs cover only 32.21–57.54% of nominal 95% intervals, and simulated controllers violate moisture constraints in 42.38–50.71% of decision steps. We therefore present the system as an auditable diagnostic architecture, not a deployed safety certificate. The results motivate adaptive representation routing, calibration-gated control, loss-tail backtesting, measured actuator models, and validated fallback operation.

1 INTRODUCTION

Irrigation control couples soil-water dynamics, weather, crop development, hydraulic constraints, and incomplete sensing. Reference-evapotranspiration scheduling remains a strong operational baseline (Allen et al., 1998); machine learning broadens predictive capability (Liakos et al., 2018; Kamilaris & Prenafeta-Boldú, 2018), but forecast accuracy alone does not establish closed-loop reliability. Sensor missingness may be informative (Rubin, 1976; Little & Rubin, 2019), and elapsed-time-aware models show the value of retaining masks and measurement age (Che et al., 2018). A controller must additionally distinguish recent executed actions from candidate future plans and represent uncertainty well enough to evaluate tail losses.

We organize the predictor around

$$\chi_t^\pi = [b_t; c_t^u; q_t^\pi], \tag{1}$$

where b_t is an environmental belief, c_t^u summarizes recent command and execution history, and q_t^π represents a candidate plan. The initial fixed KME–RFF–attention pathway was not empirically dominant: direct and ablated models often performed better. We therefore introduce adaptive routing among direct, distributional-memory, and fallback beliefs. We also make calibration a control gate, since under-dispersed scenarios can omit precisely the trajectories required by CVaR optimization.

Contributions. (1) A triad context separating environment, executed control, and future plan; (2) an adaptive router using missingness, age, outage length, expert disagreement, out-of-distribution (OOD) score, actuator residual, and rolling coverage; (3) a broad forecasting benchmark with transparent negative results; (4) chronological conformal calibration and direct VaR/CVaR backtesting before control; and (5) an executed-action-aware MPC analysis that explicitly reports unresolved violations. Detailed architectures, proofs, hyperparameters, and additional plots are moved to the appendix.

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.

2 Related Work

GRU-D combines masking, elapsed time, and decay for irregular sequences (Che et al., 2018). Kernel mean embeddings represent distributions in an RKHS (Smola et al., 2007; Muandet et al., 2017); random Fourier features (RFFs) approximate shift-invariant kernels with explicit finite-dimensional maps (Rahimi & Recht, 2007). CVAEs provide conditional latent-variable generation (Kingma & Welling, 2014; Sohn et al., 2015), while proper scores and calibration distinguish useful uncertainty from narrow but misleading intervals (Gneiting & Raftery, 2007; Guo et al., 2017). MPC, tube tightening, CVaR, and barrier functions provide complementary control tools under explicit assumptions (Mayne et al., 2000; 2005; Rockafellar & Uryasev, 2000; Ames et al., 2017). Our focus is their empirical alignment: routing chooses a representation, calibration determines scenario eligibility, and control acts only within validated regimes.

3 Problem Formulation and Assumptions

Partially observed dynamics. Let $x_t \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ denote the latent agro-hydrological state, $u_t \in \mathcal{U} \subseteq \mathbb{R}^{d_u}$ the commanded irrigation input, $a_t \in \mathcal{A} \subseteq \mathbb{R}^{d_a}$ the executed action, and w_t an exogenous disturbance. The system evolves as

$$x_{t+1} = f_\star(x_t, a_t, w_t), \qquad a_t \sim G_\star(\cdot \mid u_t, x_t), \tag{2}$$

where $G_\star$ captures discrepancies among the command, valve and pump responses, and delivered flow.

At time t, modality $m \in \{1, \dots, M\}$ provides an observation history

$$\mathcal{O}_t^{(m)} = \left\{ o^{(m)}_{t-L+1:t}, M^{(m)}_{t-L+1:t}, A^{(m)}_{t-L+1:t} \right\},$$

where $o_s^{(m)} \in \mathbb{R}^{d_m}$ is the recorded value, $M_s^{(m)} \in \{0,1\}^{d_m}$ is the observation mask, and $A_s^{(m)} \in \mathbb{R}_+^{d_m}$ is the measurement age. The complete partially observed history is denoted by $\mathcal{I}_t$.

For a horizon H, the controller evaluates the candidate plan

$$\pi_t = \left(u_{t|t}, u_{t+1|t}, \dots, u_{t+H-1|t}\right) \in \Pi_t. \tag{3}$$

Prediction is conditioned on the triad

$$\chi_t^\pi = [b_t; c_t^u; q_t^\pi], \tag{4}$$

where b_t is the environmental belief, c_t^u summarizes recent commanded and executed actions, and q_t^π encodes the candidate plan. Explicit plan conditioning permits counterfactual plans with the same observed history to induce different predictive distributions.

Risk-aware objective. Let $Y_{t+1:t+H}$ be a future state or target trajectory and let $J(\pi_t, Y_{t+1:t+H})$ include moisture-band deviation, water or energy use, actuator effort, switching, and constraint penalties. The controller solves

$$\min_{\pi_t \in \Pi_t} \ \mathbb{E}[J(\pi_t)] + \eta \, \mathrm{CVaR}_\alpha \left(J(\pi_t)\right), \tag{5}$$

subject to

$$x_{t+h} \in \mathcal{X}, \qquad u_{t+h|t} \in \mathcal{U}, \qquad \Delta u_{t+h|t} \in \Delta\mathcal{U}, \quad h = 0, \dots, H-1, \tag{6}$$

together with pump, pressure, flow, terminal, and switching constraints. The coefficient $\eta \geq 0$ controls risk sensitivity and $\alpha \in (0,1)$ specifies the upper-tail level.

Scenario-based optimization is enabled only when the predictive distribution, calibration regime, actuator model, and solver satisfy the validity checks in Section 5. Failure of any check activates restricted or fallback control.

Standing assumptions. The main results rely on the following conditions.

Assumption 1 (Bounded validated domain)**.** *The state, action, observation, and plan spaces are compact. Encoded observations, expert beliefs, routing variables, and finite-horizon losses are bounded within the validated operating domain.*

Assumption 2 (Regular representation and prediction maps)**.** *The modality encoders, recurrent updates, attention maps, projections, routing network, and control and plan encoders are Lipschitz continuous. The reliability-weighted kernel-mean denominator is bounded below by $\epsilon > 0$. Moreover,*

$$W_1\left(P_\theta(\cdot \mid \chi), P_\theta(\cdot \mid \chi')\right) \le L_P \|\chi - \chi'\|, \tag{7}$$

and the finite-horizon loss is L_J-Lipschitz in the predicted trajectory.

Assumption 3 (Bounded plant and actuator errors)**.** *On a certified event, the nominal dynamics $\widehat{f}$ and learned actuator model $\widehat{G}$ satisfy*

$$\left\| f_\star(x,a,w) - \widehat{f}(x,a,w) \right\| \le \bar{\varepsilon}_f, \tag{8}$$

$$\|a - \widehat{\mu}_G(u,b)\| \le \bar{\varepsilon}_a, \tag{9}$$

and $\widehat{f}$ is Lipschitz in state and executed action with constants L_x and L_a.

Assumption 4 (Calibration and terminal validity)**.** *The calibration data follow the training period chronologically and satisfy the exchangeability or dependence conditions required by the selected calibration procedure. The tightened MPC admits a robust terminal set and a backup controller for the same state, transition, and actuator-error bounds.*

These assumptions provide sufficient conditions for the subsequent analysis. The empirical study examines prediction and simulated control performance but does not assert that every condition has been certified for field deployment. Detailed validity regimes, dependence conditions, and terminal-set requirements are given in Appendix A.

4 PRELIMINARIES

Observation provenance and reliability. For modality m, sequence position s, and decision time t, define

$$\xi_{t,s}^{(m)} = \left[M_{t,s}^{(m)} \odot o_{t,s}^{(m)}; M_{t,s}^{(m)}; \widetilde{A}_{t,s}^{(m)} \right], \qquad r_{t,s}^{(m)} = M_{t,s}^{(m)} \odot \exp\left(-\lambda_m A_{t,s}^{(m)}\right), \tag{10}$$

where $\widetilde{A}_{t,s}^{(m)}$ is normalized measurement age. The mask distinguishes an unavailable value from an observed zero, while the age term reduces the influence of stale measurements, following the broader use of masks and elapsed time in irregular time-series modelling (Che et al., 2018).

Random features and reliability-weighted kernel means. A modality-specific encoder maps the provenance-aware input to

$$z_{t,s}^{(m)} = e_m\left(\xi_{t,s}^{(m)}\right) \in \mathbb{R}^{d_m^z}. \tag{11}$$

For a shift-invariant kernel, random Fourier features provide the explicit map

$$\phi_m(z) = \sqrt{\frac{2}{D_m}} \cos\left(W_m^\top z + b_m\right), \tag{12}$$

where W_m is sampled from the kernel spectral distribution and $b_m \sim \mathrm{Unif}[0, 2\pi]$ (Rahimi & Recht, 2007). Over a local memory window $\mathcal{W}_{t,s}$, the reliability-weighted empirical kernel mean is

$$\widehat{\mu}_{t,s}^{(m)} = \frac{\sum_{j \in \mathcal{W}_{t,s}} \bar{r}_{t,j}^{(m)} \phi_m\left(z_{t,j}^{(m)}\right)}{\epsilon + \sum_{j \in \mathcal{W}_{t,s}} \bar{r}_{t,j}^{(m)}}, \tag{13}$$

where $\bar{r}_{t,j}^{(m)}$ aggregates channel-wise reliability. This is a finite-dimensional approximation to a reliability-weighted kernel mean embedding (Smola et al., 2007; Muandet et al., 2017).

Reliability-biased temporal attention. Given recurrent states $h_{t,s}$, temporal relevance is computed by

$$e_{t,s} = v^\top \tanh(W h_{t,s}) + \gamma \log(\epsilon + \bar{r}_{t,s}) + \rho \nu_s, \tag{14}$$

$$\omega_{t,s} = \frac{\exp(e_{t,s})}{\sum_{\ell=1}^{L} \exp(e_{t,\ell})}, \tag{15}$$

where ν_s is a normalized recency index. Reliability biases attention toward recent and trustworthy states without replacing the learned relevance score.

Conditional variational forecasting. For future trajectory $Y = Y_{t+1:t+H}$ and context χ, the conditional prior and recognition distributions are

$$p_\theta(z \mid \chi) = \mathcal{N}\left(\mu_\theta^p(\chi), \operatorname{diag}\exp \ell_\theta^p(\chi)\right), \tag{16}$$

$$q_\varphi(z \mid Y, \chi) = \mathcal{N}\left(\mu_\varphi^q(Y, \chi), \operatorname{diag}\exp \ell_\varphi^q(Y, \chi)\right). \tag{17}$$

Training minimizes

$$\mathcal{L}_{\mathrm{CVAE}} = \mathbb{E}_{q_\varphi}\left[-\log p_\theta(Y \mid z, \chi)\right] + \beta_e D_{\mathrm{KL}}\left(q_\varphi \| p_\theta\right), \tag{18}$$

where β_e increases during KL warm-up. At inference, trajectories are generated from $p_\theta(z \mid \chi)$; the future-conditioned recognition model is used only during training and posterior diagnostics (Kingma & Welling, 2014; Sohn et al., 2015).

Chronological conformal calibration. For raw lower and upper predictive quantiles $\widehat{q}_{\ell,h,d}$ and $\widehat{q}_{u,h,d}$, define the calibration score

$$s_{i,h,d} = \max\left\{\widehat{q}_{\ell,h,d}(x_i) - y_{i,h,d},\ y_{i,h,d} - \widehat{q}_{u,h,d}(x_i),\ 0\right\}. \tag{19}$$

If $c_{h,d}^{(1-\alpha)}$ is the corresponding finite-sample calibration quantile, the corrected interval is

$$\mathcal{C}_{h,d}^{(1-\alpha)}(x) = \left[\widehat{q}_{\ell,h,d}(x) - c_{h,d}^{(1-\alpha)},\ \widehat{q}_{u,h,d}(x) + c_{h,d}^{(1-\alpha)}\right]. \tag{20}$$

Marginal finite-sample coverage follows under the calibration assumptions; horizon- and regime-specific coverage remain empirical validation criteria. Dependence-aware calibration and block-resampling details are provided in Appendix E.

Conditional value-at-risk. For loss J, CVaR has the variational form

$$\mathrm{CVaR}_\alpha(J) = \min_{\tau \in \mathbb{R}} \left\{\tau + \frac{1}{1-\alpha}\mathbb{E}[(J-\tau)_+]\right\}, \tag{21}$$

and its K-scenario estimator is

$$\widehat{\mathrm{CVaR}}_\alpha(J) = \min_{\tau \in \mathbb{R}} \left\{\tau + \frac{1}{K(1-\alpha)} \sum_{k=1}^{K} [J_k - \tau]_+\right\}. \tag{22}$$

The factor $(1-\alpha)^{-1}$ amplifies both genuine tail losses and distributional error. Consequently, CVaR is evaluated only after predictive calibration and loss-tail backtesting. The practical tail-resolution criterion $K(1-\alpha) \geq 20$, bootstrap uncertainty, and exceedance diagnostics are specified in Appendix E.

5 METHOD

5.1 DIRECT AND DISTRIBUTIONAL-MEMORY EXPERTS

The direct expert preserves temporally resolved modality features:

$$d_{t,s} = \mathrm{concat}_{m=1}^{M} z_{t,s}^{(m)}, \tag{23}$$

$$h_{t,s}^{\mathrm{dir}} = \mathrm{GRU}_{\mathrm{dir}}\left(d_{t,s}, h_{t,s-1}^{\mathrm{dir}}\right), \tag{24}$$

$$b_t^{\mathrm{dir}} = P_{\mathrm{dir}} \sum_{s=1}^{L} \omega_{t,s}^{\mathrm{dir}} h_{t,s}^{\mathrm{dir}}. \tag{25}$$

This pathway is intended to retain short-term variation when sensing is dense and recent. The distributional-memory expert concatenates the modality-specific embeddings from equation 13:

$$g_{t,s} = \mathrm{concat}_{m=1}^{M} \widehat{\mu}_{t,s}^{(m)}. \tag{26}$$

A second recurrent encoder produces

$$h_{t,s}^{\mathrm{kme}} = \mathrm{GRU}_{\mathrm{kme}} \left(g_{t,s}, h_{t,s-1}^{\mathrm{kme}} \right), \tag{27}$$

$$b_t^{\mathrm{kme}} = P_{\mathrm{kme}} \left[\sum_s \omega_{t,s}^{\mathrm{kme}} h_{t,s}^{\mathrm{kme}}; \sum_s \omega_{t,s}^{\mathrm{kme}} g_{t,s} \right]. \tag{28}$$

This expert provides distributional memory when recent observations are sparse but informative history remains within the validated memory window.

5.2 Adaptive belief routing

The router receives

$$s_t^{\mathrm{route}} = \left[\bar{M}_t; \bar{A}_t; R_t; D_t^b; U_t^{\mathrm{dir}}; U_t^{\mathrm{kme}}; S_t^{\mathrm{OOD}}; e_t^a; C_t^{\mathrm{roll}} \right], \tag{29}$$

where $\bar{M}_t$ is recent completeness, $\bar{A}_t$ is normalized age, R_t is the longest outage, D_t^b is expert disagreement, U_t^e is expert uncertainty, S_t^{OOD} is an out-of-distribution score, e_t^a is the actuator residual, and C_t^{roll} is recent calibrated coverage. Routing weights are

$$\boldsymbol{\pi}_t = \left[\pi_t^{\mathrm{dir}}, \pi_t^{\mathrm{kme}}, \pi_t^{\mathrm{fb}} \right] = \mathrm{softmax} \left(f_g(s_t^{\mathrm{route}}) \right), \tag{30}$$

and the routed belief is

$$b_t = \pi_t^{\mathrm{dir}} b_t^{\mathrm{dir}} + \pi_t^{\mathrm{kme}} b_t^{\mathrm{kme}} + \pi_t^{\mathrm{fb}} b_t^{\mathrm{fb}}. \tag{31}$$

The fallback belief b_t^{fb} combines verified measurements, persistence, conservative uncertainty inflation, and bounded crop-stage rules. The router is trained with

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{CVAE}} + \lambda_{\mathrm{bal}} \mathcal{L}_{\mathrm{bal}} + \lambda_{\mathrm{sw}} \mathcal{L}_{\mathrm{sw}}, \tag{32}$$

where

$$\mathcal{L}_{\mathrm{sw}} = \frac{1}{T-1} \sum_{t=2}^{T} \| \boldsymbol{\pi}_t - \boldsymbol{\pi}_{t-1} \|_1 \tag{33}$$

discourages route chattering. Hysteresis and minimum dwell time are applied during control. The explicit fallback mode prevents the gate from blending two unreliable neural experts when the current regime lies outside calibration support.

5.3 Control-history and candidate-plan representations

The control encoder summarizes commands, measured execution, action increments, switching, and residual behavior:

$$c_t^u = E_u \left(u_{t-L+1:t}, a_{t-L+1:t}, \Delta a_{t-L+1:t}, e_{t-L+1:t}^a \right). \tag{34}$$

The plan encoder represents the candidate future input sequence:

$$q_t^\pi = E_\pi \left(u_{t|t}, \ldots, u_{t+H-1|t} \right). \tag{35}$$

The final predictive context is the triad in equation 4. Distinct plans therefore induce distinct conditional trajectory distributions.

5.4 Scenario generation and calibration gate

The CVAE generates

$$Y_{t+1:t+H}^{(k)} \sim P_\theta \left(\cdot \mid \chi_t^\pi \right), \qquad k = 1, \ldots, K. \tag{36}$$

The raw predictive intervals are calibrated using equation 19. Control eligibility requires that lower confidence bounds for nominal 90% and 95% coverage exceed 87% and 92%, respectively, within every safety-relevant horizon and validated operating regime. A coverage error between three and five percentage points activates restricted control and uncertainty inflation; a larger error activates fallback. Loss-space eligibility additionally requires:

1. empirical VaR exceedance frequency consistent with $1-\alpha$;
2. acceptable agreement between predicted and realized mean excess loss;
3. current routing and actuator regimes represented in calibration data;
4. no unresolved OOD, actuator, or solver alarm.

The practical tail-resolution requirement

$$K(1-\alpha) \geq 20 \tag{37}$$

is used for comparative CVaR evaluation. It is a numerical stability criterion, not a formal safety guarantee.

5.5 Executed-action-aware MPC and fallback

For a candidate command sequence, executed actions are sampled as

$$a_{t+h}^{(k)} \sim \widehat{G}\left(\cdot \mid u_{t+h|t}, b_t\right). \tag{38}$$

The resulting predictive law is

$$\widehat{P}(Y \mid b, u) = \int \widehat{P}_\theta(Y \mid b, a)\widehat{G}(da \mid u, b). \tag{39}$$

The calibrated MPC objective is

$$\widehat{J}_{\mathrm{MPC}}(\pi) = \frac{1}{K}\sum_{k=1}^{K} J_k(\pi) + \eta\widehat{\mathrm{CVaR}}_\alpha\left(J(\pi)\right) + \lambda_{\Delta u}\sum_{h=0}^{H-1}\|u_{t+h|t} - u_{t+h-1|t}\|_1. \tag{40}$$

For safe set $\mathcal{S} = \{x : h(x) \geq 0\}$, the risk-tightened condition is

$$\widehat{\mathrm{CVaR}}_\alpha\left[-h(\widehat{X}_{t+1}) \mid b_t, u_t\right] + m_t^{\mathrm{tot}} \leq -\kappa h(\widehat{x}_t), \tag{41}$$

where m_t^{tot} aggregates certified state-estimation, transition, actuator, calibration, scenario, approximation, and solver margins. When the calibration or feasibility gate fails, the system moves through a graded hierarchy: restricted MPC, validated crop-stage scheduling, bounded hold or stop action, and human intervention. Hysteresis, minimum dwell time, and bounded action increments support bumpless mode transitions.

5.6 Kernel–Galerkin policy evaluation

Let

$$V_D(x) = \theta^\top \varphi_D(x)$$

be an RFF value approximation. Given transitions (x_i, a_i, x_i^+) and costs ℓ_i, policy evaluation solves

$$\widehat{\theta} = \arg\min_\theta \sum_{i=1}^{n}\left[\ell_i + \gamma\theta^\top\varphi_D(x_i^+) - \theta^\top\varphi_D(x_i)\right]^2 + \lambda\|\theta\|_2^2. \tag{42}$$

A candidate policy is accepted only if its estimated improvement exceeds the combined uncertainty margin. Extended branch-dependent allocation, fallback pseudocode, and solver details are provided in Appendix **??**.

6 Theoretical Analysis, Generalization, and Complexity

This section summarizes the stability, generalization, feasibility, and computational properties needed to interpret the proposed framework. Complete proofs, constants, dependence conditions, and complexity derivations are given in Appendices C–J.

Representation and routing stability. Let

$$b = \sum_{e \in \mathcal{E}} \pi_e b^e, \qquad \mathcal{E} = \{\mathrm{dir}, \mathrm{kme}, \mathrm{fb}\},$$

where b^e is an expert belief and $\boldsymbol{\pi} = (\pi_e)_{e \in \mathcal{E}}$ is the routing distribution. Under Assumptions 1 and 2, each expert satisfies

$$\|b^e - b^{e\prime}\| \le C_o^e \|o - o'\| + C_M^e \|M - M'\| + C_A^e \|A - A'\| + C_D^e \varepsilon_{\mathrm{RFF}}, \tag{43}$$

where the final term is omitted for experts that do not use random features. For bounded shift-invariant kernels,

$$\varepsilon_{\mathrm{RFF}} = O_{\mathbb{P}}\left(\sqrt{\frac{\log(1/\delta)}{D_{\min}}}\right)$$

on compact domains (Rahimi & Recht, 2007). If $\|b^{e\prime}\| \le B_b$, then the routed belief satisfies

$$\|b - b'\| \le \sum_{e \in \mathcal{E}} \pi_e \|b^e - b^{e\prime}\| + B_b \|\boldsymbol{\pi} - \boldsymbol{\pi}'\|_1. \tag{44}$$

When the router is L_g-Lipschitz,

$$\|\boldsymbol{\pi} - \boldsymbol{\pi}'\|_1 \le L_g \|s^{\mathrm{route}} - s^{\mathrm{route}\prime}\|.$$

Thus, routed-belief robustness depends jointly on expert stability and routing smoothness. This motivates route-switching regularization, hysteresis, and minimum dwell time.

Forecast and tail-risk sensitivity. Under Assumption 2, suppose that the conditional trajectory law is L_P-Lipschitz in context under the 1-Wasserstein metric and that the finite-horizon loss is L_J-Lipschitz. Then

$$|\mathbb{E}J_\chi - \mathbb{E}J_{\chi'}| \le L_J L_P \|\chi - \chi'\|, \tag{45}$$

$$|\mathrm{CVaR}_\alpha(J_\chi) - \mathrm{CVaR}_\alpha(J_{\chi'})| \le \frac{L_J L_P}{1-\alpha} \|\chi - \chi'\|. \tag{46}$$

The first bound follows from Wasserstein duality, while the second follows from the variational representation of CVaR (Rockafellar & Uryasev, 2000). The factor $(1-\alpha)^{-1}$ shows that extreme tail levels amplify representation and distributional errors. Increasing α therefore cannot compensate for trajectories omitted by an under-dispersed predictor.

Statistical generalization. Let $\mathcal{F}$ denote the forecasting class, and suppose that the prediction loss is bounded by B_ℓ and is L_ℓ-Lipschitz in the model output. For independent observations, standard Rademacher-complexity arguments give, with probability at least $1 - \delta$,

$$\sup_{f \in \mathcal{F}} \left| \mathcal{R}(f) - \widehat{\mathcal{R}}_n(f) \right| \le 2L_\ell \widehat{\mathfrak{R}}_n(\mathcal{F}) + 3B_\ell \sqrt{\frac{\log(2/\delta)}{2n}}. \tag{47}$$

For temporally dependent observations, n is replaced by an effective sample size obtained using a blocking or mixing argument. The bound expresses the main representation trade-off: additional experts, random features, and routing capacity may reduce approximation error while increasing estimation error. For a fixed policy with bounded loss $0 \le J \le B_J$, the finite-scenario CVaR estimation error scales as

$$\left| \mathrm{CVaR}_\alpha(J) - \widehat{\mathrm{CVaR}}_\alpha(J) \right| = O_{\mathbb{P}}\left(\frac{B_J}{1-\alpha} \sqrt{\frac{\log(1/\delta)}{K}} \right). \tag{48}$$

A simultaneous finite-policy bound follows by applying a union bound over the candidate set. This result characterizes scenario-sampling error but does not correct predictive misspecification, which must be addressed through calibration and held-out loss-tail backtesting.

Calibration-aware improvement. If the true and estimated losses satisfy

$$|J - \widehat{J}| \le \varepsilon_J \quad \text{almost surely,}$$

then the monotonicity and translation equivariance of CVaR imply

$$\left|\mathrm{CVaR}_\alpha(J) - \mathrm{CVaR}_\alpha(\widehat{J})\right| \le \varepsilon_J. \tag{49}$$

We collect the principal decision-level discrepancies as

$$\varepsilon_{\mathrm{tot}} = \varepsilon_{\mathrm{cal}} + \varepsilon_{\mathrm{samp}} + \varepsilon_{\mathrm{tr}} + \varepsilon_{\mathrm{act}} + \varepsilon_{\mathrm{RFF}} + \varepsilon_{\mathrm{Gal}} + \varepsilon_{\mathrm{sol}}. \tag{50}$$

A candidate policy is accepted only when its estimated improvement over the incumbent exceeds $2\varepsilon_{\mathrm{tot}}$. Smaller gains remain unresolved at the available calibration, modelling, sampling, approximation, and numerical accuracy. Split-conformal calibration provides marginal finite-sample coverage under its exchangeability conditions. Horizon-wise and regime-wise coverage checks are used as conservative control-readiness diagnostics rather than universal conditional-coverage guarantees.

Conditional practical feasibility. Under Assumptions 3 and 4, define the state-error tube recursively by

$$\rho_{h+1} = L_x \rho_h + L_a \bar{\varepsilon}_a + \bar{\varepsilon}_f, \qquad \rho_0 = \varepsilon_t^x, \quad h = 0, \dots, H-1. \tag{51}$$

If the nominal trajectory satisfies

$$\widehat{x}_{t+h|t} \in \mathcal{X} \ominus \mathbb{B}(\rho_h), \qquad h = 1, \dots, H, \tag{52}$$

and the terminal set is robustly positively invariant under the backup controller for the same error bounds, then

$$x_{t+h} \in \mathcal{X}, \qquad h = 1, \dots, H,$$

on the joint certified event. The tightened problem remains recursively feasible while the state-estimation error, routing regime, calibration status, transition-model error, and actuator residual remain within the validity sets used to construct the tube. Leaving these sets activates fallback rather than extrapolating the feasibility result (Mayne et al., 2000; 2005).

Approximate value convergence. If the exact policy-evaluation operator is a γ-contraction and each approximate evaluation or improvement step incurs error at most $C\varepsilon_{\mathrm{tot}}$, then

$$\limsup_{k \to \infty} \|V_k - V^\star\| \le \frac{C\varepsilon_{\mathrm{tot}}}{1-\gamma}. \tag{53}$$

Persistent calibration, sampling, transition, actuator, random-feature, Galerkin, or solver errors therefore imply convergence to a quantified neighborhood rather than exact optimality (Bertsekas, 2012; Sutton & Barto, 2018).

Computational complexity. Let

$$D_\Sigma = \sum_{m=1}^{M} D_m, \qquad d_\chi = d_b + d_c + d_p,$$

and let B, L, H, K, N_π denote batch size, context length, prediction horizon, scenario count, and number of candidate plans. Using cumulative kernel statistics, the dominant costs are

$$\text{RFF–KME encoding:} \quad O\left(BL \sum_{m=1}^{M} d_m^z D_m\right), \tag{54}$$

$$\text{recurrent experts:} \quad O\left(BLd_h(D_\Sigma + d_h)\right), \tag{55}$$

$$\text{CVAE scenario generation:} \quad O\left(BKHD_{\mathrm{dec}}\right), \tag{56}$$

$$\text{candidate-plan evaluation:} \quad O(N_\pi K H C_f), \tag{57}$$

$$\text{empirical CVaR:} \quad O(N_\pi K \log K), \tag{58}$$

where

$$D_{\mathrm{dec}} = d_{\mathrm{dec}}\left(d_{\mathrm{dec}} + d_\chi + d_z\right)$$

Table 1: Selected probabilistic forecasting results at 25% synthetic missingness. Errors use the standardized target scale. Coverage refers to nominal 95% intervals. Lower RMSE and energy score are preferred; coverage should be assessed relative to 0.95.

Model	RMSE	Coverage	Energy
Random Forest	**0.1854**	0.9277	0.2983
Extra Trees	0.1977	0.9241	**0.2967**
Full KME–RFF–attention CVAE	0.2350	0.3467	0.4377
Control-conditioned CVAE	0.2162	0.5754	0.3790
No-KME CVAE	0.2068	0.4374	0.3731
Unimodal soil CVAE	0.2120	0.9785	0.3717
Unimodal GRU–Gaussian	0.1939	0.9826	0.3214
Early-fusion GRU–Gaussian	0.2292	0.9538	0.3923

and C_f is the cost of one transition and constraint evaluation. Scenario generation and candidate-plan rollout dominate when K and N_π are large. The leading memory requirement is

$$O\left(BLD_\Sigma + BKHD_y + D_V^2\right), \tag{59}$$

corresponding to expert activations, generated trajectories, and the Kernel–Galerkin system. Streaming scenario evaluation can remove the $BKHD_y$ storage term when only running moments, quantiles, and tail statistics are required. Hard routing reduces expected inference cost by evaluating only the selected expert, whereas soft routing provides smoother transitions at a higher computational cost.

7 Experiments and Results

Experimental protocol. We evaluate chronological irrigation forecasting, robustness to synthetic missingness, finite-scenario estimation, and closed-loop control. Forecasting uses a context length of $L = 12$ and a prediction horizon of $H = 6$. Data are partitioned chronologically into training, validation, calibration, and test periods to avoid temporal leakage; random cross-validation is not used because it can yield optimistic estimates for dependent time series (Bergmeir et al., 2018). Standardization parameters are estimated from the training period only. Missingness rates are

$$0, \quad 0.10, \quad 0.25, \quad \text{and} \quad 0.40.$$

Neural models are trained with AdamW, gradient clipping, validation-based early stopping, and KL warm-up to $\beta_{\max} = 0.1$. CVAE scenarios are sampled from the conditional prior at inference. Forecasts are evaluated using RMSE, MAE, empirical coverage and width of nominal 95% intervals, and the energy score (Gneiting & Raftery, 2007). Forecasting metrics are reported on the standardized target scale, whereas control quantities use physical or simulation-specific units and are interpreted separately. Complete preprocessing, model configurations, tuning ranges, random seeds, and structured-outage protocols are reported in Appendix F.

Forecasting performance. Table 1 reports representative results at 25% missingness. Complete classical, neural, and ablation results are provided in Appendix G. Random Forest achieves the lowest overall point RMSE, 0.1854, while Extra Trees obtains the best energy score, 0.2967. The unimodal GRU–Gaussian is the strongest neural predictor, with RMSE 0.1939 and energy score 0.3214, indicating that recent soil-moisture dynamics provide a strong predictive signal. Early multimodal fusion does not improve upon the unimodal model under the present data regime. The full KME–RFF–attention CVAE is not empirically dominant. Removing KME reduces RMSE from 0.2350 to 0.2068 and improves the energy score from 0.4377 to 0.3731. No-RFF and no-attention variants also improve point accuracy, as reported in Appendix G. Accordingly, the revised architecture treats the KME pathway as a routed distributional-memory expert rather than a mandatory representation. The principal CVAE variants are substantially under-dispersed, attaining only 32.21% to 57.54% coverage for nominal 95% intervals. For example, coverage of 44.36% implies a 55.64% miss rate, approximately 11.1 times the intended 5% miss frequency. Coverage of 34.67% corresponds to approximately 13.1 times the intended frequency. These ratios are diagnostics of predictive miscalibration, not direct estimates of CVaR bias. Nevertheless, they exclude the raw CVAE scenarios from tail-risk control until chronological calibration and held-out loss-space backtesting have been completed.

Table 2: Ten-seed controller diagnostics. Violation is the fraction of decision steps outside the prescribed moisture band. Costs are simulation-specific and are not comparable with standardized forecasting errors.

Policy	Mean cost	$\text{CVaR}_{0.9}$	Violation	Mean action
Nominal MPC	583.80	3777.55	**0.4238**	0.8213
Mean–CVaR MPC	586.72	**3767.83**	0.4494	0.8840
Robust MPC	613.18	3819.03	0.5071	0.7940
Tube MPC	**583.80**	3777.60	**0.4238**	0.8220

Robustness to missing observations. For the full belief representation, RMSE increases from 0.2027 under complete observations to 0.2456 at 40% missingness. Mean reliability decreases from 1.0000 to 0.5930, while latent displacement from the complete-observation representation increases from 0 to 0.8657. Last-state pooling deteriorates more sharply, with RMSE increasing from 0.2179 to 0.3633. This supports temporally distributed aggregation when the most recent observations are unavailable. No-KME and No-RFF representations remain competitive at several missingness levels. The results therefore support adaptive routing: the direct expert is favored under dense and recent sensing, while distributional memory is reserved for validated outage regimes. Complete missingness curves, expert ablations, and latent diagnostics are provided in Appendix G.

Finite-scenario sensitivity. Increasing the scenario count reduces Monte Carlo variability. The standard deviation of the estimated objective decreases from 468.9 at $K = 5$ to 215.3 at $K = 30$, while modal-action selection frequency increases from 0.583 to 0.917 by $K = 20$. This indicates numerical stabilization but does not constitute a scenario-optimization guarantee, which would require additional convexity, support-rank, and sampling conditions (Calafiore & Campi, 2006; Campi & Garatti, 2008). For empirical CVaR, the expected number of upper-tail samples is

$$K_{\text{tail}} \approx K(1-\alpha).$$

We therefore use

$$K(1-\alpha) \geq 20$$

as a practical minimum for comparative tail-risk analysis and accompany CVaR estimates with bootstrap uncertainty. Increasing K reduces sampling error but cannot correct a misspecified or under-dispersed predictive distribution. Full scenario-count and risk-level sensitivity results are reported in Appendix H.

Closed-loop controller diagnostics. Table 2 summarizes nominal, mean–CVaR, robust, and tube MPC over ten random seeds. Mean–CVaR MPC lowers the reported tail cost by approximately 9.7 relative to nominal MPC but increases mean cost, mean action, and violation frequency. Robust MPC has the highest mean cost, tail cost, and violation rate, while tube MPC is nearly indistinguishable from nominal MPC. Thus, adding risk or robustness terms does not by itself improve empirical constraint satisfaction. Violation rates between 42.38% and 50.71% preclude deployment-level safety claims. Likely causes include predictive undercoverage, transition and actuator mismatch, poorly scaled constraints, and unverified terminal ingredients. The controller results are therefore interpreted as closed-loop failure diagnostics. In the revised pipeline, scenario MPC is enabled only after calibration, OOD, actuator, and feasibility checks pass; otherwise, the system activates restricted MPC or a validated fallback policy. Detailed violation severity, duration, time to first violation, water use, switching, solver failures, actuator residuals, and fallback behavior are reported in Appendix H.

Model selection and latent diagnostics. Hyperparameters are selected using randomized search on the chronological validation period; the calibration and test periods remain untouched during selection. The selected model is refitted using the designated development data and evaluated once on the held-out test period. Multiple seeds, moving-block bootstrap intervals, and rolling-origin evaluation are used for final comparisons. PCA and t-SNE projections of the belief, conditional prior, and recognition posterior are used only as descriptive diagnostics (van der Maaten & Hinton, 2008). Optional K-means and silhouette analyses examine local structure without treating clusters as causal or physically discrete regimes (Lloyd, 1982; Rousseeuw, 1987). Complete search spaces, seed-level results, latent plots, model sizes, stopping rules, and compute requirements are provided in Appendices B and F.

8 Discussion and Limitations

The triad context remains structurally useful because candidate plans must enter the predictor to support counterfactual evaluation. However, the experiments reject mandatory use of every representation component. Adaptive routing offers a testable alternative: use the direct expert under clean sensing, distributional memory during validated outages, and fallback outside calibration support.

Calibration changes real-time actions by widening or reshaping scenario tails, altering candidate-plan rankings, and exposing infeasibility previously hidden by under-dispersion. Fallback improves auditability and availability but trades off water efficiency, adaptation, and possible crop stress. We therefore use a graded ladder: calibrated MPC, restricted MPC, validated crop-stage scheduling, bounded hold/stop, and human intervention, with hysteresis and bumpless transfer.

The study remains limited by retrospective data, synthetic missingness, incomplete intervention logging, uncalibrated soil-water simulation, and unverified actuator bounds. Theoretical statements are conditional foundations, not novel deployment guarantees. Future work should evaluate structured outages, use APSIM/DSSAT/SWAP-class calibrated dynamics, identify command-to-flow distributions on a pump–valve testbed, and require near-zero violations before field trials.

9 Conclusion

We presented an auditable framework for triad-conditioned irrigation forecasting and diagnostic risk control. The revised architecture adaptively routes among direct, kernel-memory, and fallback beliefs and admits scenarios to MPC only after calibration and loss-tail checks. Current results favor simpler predictors, reveal severe CVAE undercoverage, and show unacceptable simulated violation rates. The framework therefore defines a validation pathway rather than a safety claim: reliable deployment requires calibrated uncertainty, measured actuator responses, verified invariant backup control, and physically grounded closed-loop testing.

Reproducibility Statement

Appendix F specifies preprocessing, chronological splits, model variants, hyperparameter ranges, metrics, scenario counts, and control settings. The supplementary code reproduces all tables and plots from machine-readable outputs; identifying repository metadata must remain anonymous during review.

Ethics Statement

The work uses retrospective sensor data and simulation and does not involve human subjects. The principal risk is overclaiming safety from uncalibrated forecasts or synthetic plant models. We mitigate this risk by reporting failures, gating control on calibration, requiring fallback operation, and explicitly withholding field-safety claims.

AI Use Statement

Generative AI tools were used for language editing,

References

Richard G. Allen, Luis S. Pereira, Dirk Raes, and Martin Smith. *Crop Evapotranspiration: Guidelines for Computing Crop Water Requirements*. Number 56 in FAO Irrigation and Drainage Paper. Food and Agriculture Organization of the United Nations, Rome, Italy, 1998. ISBN 9789251042199.

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.

Christoph Bergmeir, Rob J. Hyndman, and Bonsoo Koo. A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis*, 120:70–83, 2018. doi: 10.1016/j.csda.2017.11.003.

Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, Belmont, Massachusetts, 4 edition, 2012.

Giuseppe C. Calafiore and Marco C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006. doi: 10.1109/TAC.2006.875041.

Marco C. Campi and Simone Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008. doi: 10.1137/07069821X.

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.

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330, 2017.

Andreas Kamilaris and Francesc X. Prenafeta-Boldú. Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147:70–90, 2018. doi: 10.1016/j.compag.2018.02.016.

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

Konstantinos G. Liakos, Patrizia Busato, Dimitrios Moshou, Simon Pearson, and Dionysis Bochtis. Machine learning in agriculture: A review. *Sensors*, 18(8):2674, 2018. doi: 10.3390/s18082674.

Roderick J. A. Little and Donald B. Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, Hoboken, New Jersey, 3 edition, 2019. ISBN 9781119482260. doi: 10.1002/9781119482260.

Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28 (2):129–137, 1982. doi: 10.1109/TIT.1982.1056489.

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.

David Q. Mayne, María M. Seron, and Saša V. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005. doi: 10.1016/j.automatica.2004.08.019.

Krikamol Muandet, Kenji Fukumizu, Bharath K. 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, pp. 1177–1184. Curran Associates, Inc., 2007.

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

R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002. doi: 10.1016/S0378-4266(02)00271-6.

Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987. doi: 10.1016/0377-0427(87)90125-7.

Donald B. Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976. doi: 10.1093/biomet/63.3.581.

Alex J. Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A hilbert space embedding for distributions. In *Algorithmic Learning Theory*, volume 4754 of *Lecture Notes in Computer Science*, pp. 13–31. Springer, 2007. doi: 10.1007/978-3-540-75225-7_5.

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.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, Massachusetts, 2 edition, 2018.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

A Notation, Assumptions, and Validity Regimes

This appendix formalizes the notation, regularity conditions, and operational regimes used in the main analysis.

A.1 Notation

Table 3 summarizes the principal symbols.

Table 3: Principal notation.

Symbol	Description
$x_t \in \mathbb{R}^{d_x}$	latent agro-hydrological state
$u_t \in \mathbb{R}^{d_u}$	commanded irrigation input
$a_t \in \mathbb{R}^{d_a}$	executed irrigation action
$o_t^{(m)} \in \mathbb{R}^{d_m}$	observation from modality m
$M_t^{(m)}$	observation-availability mask
$A_t^{(m)}$	measurement-age vector
L, H	observation-window and forecast-horizon lengths
D_m	RFF dimension for modality m
$b_t^{\mathrm{dir}}, b_t^{\mathrm{kme}}$	direct and KME expert beliefs
b_t^{fb}	conservative fallback belief
$\boldsymbol{\pi}_t$	routing distribution over experts
c_t^u, q_t^π	control-history and candidate-plan representations
χ_t^π	triad-conditioned predictive context
K, α, η	scenario count, CVaR level, and risk weight
ρ_h	propagated state-error tube radius
$\varepsilon_{\mathrm{tot}}$	aggregate decision-level discrepancy

A.2 Operational domain and fallback validity

The validated domain specifies physical ranges for soil moisture, evapotranspiration, temperature, humidity, solar radiation, pressure, flow, valve position, and irrigation command. A state, observation pattern, actuator residual, or routing vector outside its validation range activates fallback rather than extrapolation of the predictive or feasibility guarantees. The fallback-validity set is

$$\mathcal{V}_{\mathrm{fb}} = \mathcal{V}_x \cap \mathcal{V}_{\mathrm{obs}} \cap \mathcal{V}_{\mathrm{act}} \cap \mathcal{V}_{\mathrm{cal}}, \tag{60}$$

where the component sets encode state-estimation validity, observation availability, actuator consistency, and predictive calibration.

A.3 Regularity and dependence conditions

The modality encoders, recurrent maps, attention operators, projections, routing network, and conditional generator are assumed Lipschitz continuous on the validated compact domain. Their constants may be bounded through spectral norms of the corresponding affine maps and Lipschitz constants of the nonlinear activations. For chronological observations, generalization results use a blocking argument under an explicit β-mixing or α-mixing condition. Calibration claims are marginal unless stronger regime-wise exchangeability or dependence-aware assumptions are established.

A.4 Plant, actuator, and terminal validity

Transition and actuator discrepancies are bounded on the certified event by

$$\left\| f_\star(x,a,w) - \widehat{f}(x,a,w) \right\| \le \bar{\varepsilon}_f, \tag{61}$$

$$\|a - \widehat{\mu}_G(u,b)\| \le \bar{\varepsilon}_a. \tag{62}$$

These bounds must be estimated from held-out transitions and measured command-to-flow responses. The bounds are invalidated when state, weather, pressure, command, or actuator residual leaves the corresponding validation range. The terminal set $\mathcal{X}_f$ must be robustly positively invariant under the backup controller κ_{fb}:

$$x \in \mathcal{X}_f \quad \Longrightarrow \quad f_\star\left(x, \kappa_{\mathrm{fb}}(x), w\right) \in \mathcal{X}_f \tag{63}$$

for all disturbances and actuator discrepancies inside the certified sets.

B Architecture and Routing Details

B.1 Expert architectures

The direct expert processes temporally resolved modality features:

$$z_{t,s}^{(m)} = e_m\left(\xi_{t,s}^{(m)}\right), \tag{64}$$

$$d_{t,s} = \mathrm{concat}_{m=1}^{M} z_{t,s}^{(m)}, \tag{65}$$

$$h_{t,s}^{\mathrm{dir}} = \mathrm{GRU}_{\mathrm{dir}}\left(d_{t,s}, h_{t,s-1}^{\mathrm{dir}}\right), \tag{66}$$

$$b_t^{\mathrm{dir}} = P_{\mathrm{dir}} \sum_{s=1}^{L} \omega_{t,s}^{\mathrm{dir}} h_{t,s}^{\mathrm{dir}}. \tag{67}$$

The distributional-memory expert first computes reliability-weighted random feature means:

$$\phi_m(z) = \sqrt{\frac{2}{D_m}} \cos\left(W_m^\top z + b_m\right), \tag{68}$$

$$\widehat{\mu}_{t,s}^{(m)} = \frac{\sum_{j \in \mathcal{W}_{t,s}} \bar{r}_{t,j}^{(m)} \phi_m(z_{t,j}^{(m)})}{\epsilon + \sum_{j \in \mathcal{W}_{t,s}} \bar{r}_{t,j}^{(m)}}, \tag{69}$$

$$g_{t,s} = \mathrm{concat}_{m=1}^{M} \widehat{\mu}_{t,s}^{(m)}. \tag{70}$$

A separate GRU and reliability-biased attention mechanism produce b_t^{kme}. The fallback expert combines the most recent verified state, persistence forecasts, conservative interval inflation, crop-stage rules, and bounded control recommendations.

B.2 Routing variables

The routing network receives

$$s_t^{\mathrm{route}} = \left[\bar{M}_t; \bar{A}_t; R_t; D_t^b; U_t^{\mathrm{dir}}; U_t^{\mathrm{kme}}; S_t^{\mathrm{OOD}}; e_t^a; C_t^{\mathrm{roll}}\right], \tag{71}$$

where:

- $\bar{M}_t$ is recent observation completeness;
- $\bar{A}_t$ is normalized mean measurement age;
- R_t is the longest current observation outage;
- $D_t^b = \|b_t^{\mathrm{dir}} - b_t^{\mathrm{kme}}\|_2$;
- U_t^e summarizes predictive uncertainty from expert e;
- S_t^{OOD} is an out-of-distribution score;
- e_t^a is the command-to-execution residual;
- C_t^{roll} is rolling empirical coverage.

The soft routing distribution is

$$\boldsymbol{\pi}_t = \mathrm{softmax}\left(f_g(s_t^{\mathrm{route}})\right), \tag{72}$$

and the final belief is

$$b_t = \pi_t^{\mathrm{dir}} b_t^{\mathrm{dir}} + \pi_t^{\mathrm{kme}} b_t^{\mathrm{kme}} + \pi_t^{\mathrm{fb}} b_t^{\mathrm{fb}}. \tag{73}$$

B.3 Routing regularization and mode transitions

The route-switching penalty is

$$\mathcal{L}_{\mathrm{sw}} = \frac{1}{T-1} \sum_{t=2}^{T} \|\boldsymbol{\pi}_t - \boldsymbol{\pi}_{t-1}\|_1. \tag{74}$$

A weak branch-balance term prevents premature gate collapse:

$$\mathcal{L}_{\mathrm{bal}} = \sum_e \left(\frac{1}{T} \sum_{t=1}^{T} \pi_t^e - \rho_e \right)^2. \tag{75}$$

Entry and exit thresholds are separated by hysteresis. Once fallback is activated, a minimum dwell time is imposed before normal control can resume. Action changes during route transitions satisfy

$$u_t = \mathrm{clip}\left[u_{t-1} + \mathrm{clip}\left(u_t^{\mathrm{new}} - u_{t-1}, -\Delta u_{\max}, \Delta u_{\max}\right), u_{\min}, u_{\max}\right]. \tag{76}$$

B.4 Training and inference procedures

The recommended training sequence is:

1. train the direct expert independently;
2. train the distributional-memory expert independently;
3. evaluate and calibrate both experts by sensing regime;
4. freeze the experts and train the router;
5. fine-tune the complete architecture with a reduced learning rate;
6. recalibrate the final predictive distribution;
7. evaluate the held-out test period once.

At inference, the controller computes route diagnostics, selects or blends experts, generates trajectories, applies calibration, evaluates control readiness, and either solves scenario MPC or invokes the fallback hierarchy.

B.5 Architecture ablations

The complete evaluation includes:

- direct expert only;
- KME expert only;

- soft-routed direct and KME experts;
- routed direct, KME, and fallback experts;
- no-RFF, no-attention, and last-state variants;
- hard versus soft routing;
- routing without calibration and routing without OOD diagnostics.

C Complete Theoretical Analysis

C.1 Expert-belief perturbation

For each expert, the perturbation analysis propagates errors through the missingness-aware encoder, reliability map, random-feature approximation, KME quotient, recurrent operator, attention layer, and final projection. For the KME expert, the denominator is bounded below by $\epsilon > 0$. Combining the quotient perturbation bound with the uniform random-feature approximation of Rahimi & Recht (2007) gives

$$\|b^e - b^{e\prime}\| \le C_o^e \|o - o'\| + C_M^e \|M - M'\| + C_A^e \|A - A'\| + C_D^e \varepsilon_{\mathrm{RFF}}. \tag{77}$$

For bounded shift-invariant kernels,

$$\varepsilon_{\mathrm{RFF}} = O_{\mathbb{P}}\left(\sqrt{\frac{\log(1/\delta)}{D_{\min}}}\right) \tag{78}$$

uniformly on compact sets.

C.2 Routed-belief perturbation

Using

$$b - b' = \sum_e \pi_e (b^e - b^{e\prime}) + \sum_e (\pi_e - \pi_e') b^{e\prime}, \tag{79}$$

the triangle inequality yields

$$\|b - b'\| \le \sum_e \pi_e \|b^e - b^{e\prime}\| + B_b \|\boldsymbol{\pi} - \boldsymbol{\pi}'\|_1, \tag{80}$$

provided that $\|b^{e\prime}\| \le B_b$. If the router is L_g-Lipschitz,

$$\|\boldsymbol{\pi} - \boldsymbol{\pi}'\|_1 \le L_g \|s^{\mathrm{route}} - s^{\mathrm{route}\prime}\|. \tag{81}$$

This establishes that stable experts do not guarantee a stable routed belief unless the router is also regularized.

C.3 Forecast and CVaR propagation

If

$$W_1\left(P_\theta(\cdot \mid \chi), P_\theta(\cdot \mid \chi')\right) \le L_P \|\chi - \chi'\| \tag{82}$$

and the loss is L_J-Lipschitz, Kantorovich-Rubinstein duality implies

$$|\mathbb{E}J_\chi - \mathbb{E}J_{\chi'}| \le L_J L_P \|\chi - \chi'\|. \tag{83}$$

Using the variational representation of CVaR (Rockafellar & Uryasev, 2000; 2002) and the Lipschitz continuity of the positive-part operator gives

$$|\mathrm{CVaR}_\alpha(J_\chi) - \mathrm{CVaR}_\alpha(J_{\chi'})| \le \frac{L_J L_P}{1-\alpha} \|\chi - \chi'\|. \tag{84}$$

C.4 Bounded-loss CVaR tightening

If

$$\widehat{J} - \varepsilon_J \le J \le \widehat{J} + \varepsilon_J, \tag{85}$$

then monotonicity and translation equivariance imply

$$\left|\mathrm{CVaR}_\alpha(J) - \mathrm{CVaR}_\alpha(\widehat{J})\right| \le \varepsilon_J. \tag{86}$$

Consequently,

$$\mathrm{CVaR}_\alpha(\widehat{J}) \le c - \varepsilon_J \tag{87}$$

is sufficient for $\mathrm{CVaR}_\alpha(J) \le c$ on the certified event.

C.5 Conditional recursive feasibility

Define the state-prediction error

$$e_h = x_{t+h} - \widehat{x}_{t+h|t}. \tag{88}$$

Under the plant and actuator bounds,

$$\|e_{h+1}\| \le L_x\|e_h\| + L_a\bar{\varepsilon}_a + \bar{\varepsilon}_f. \tag{89}$$

With

$$\rho_{h+1} = L_x\rho_h + L_a\bar{\varepsilon}_a + \bar{\varepsilon}_f, \qquad \rho_0 = \varepsilon_t^x, \tag{90}$$

induction gives $\|e_h\| \le \rho_h$. Hence,

$$\widehat{x}_{t+h|t} \in \mathcal{X} \ominus \mathbb{B}(\rho_h) \tag{91}$$

implies

$$x_{t+h} \in \mathcal{X}. \tag{92}$$

The shifted candidate sequence and robust terminal invariance establish recursive feasibility while all modelling, routing, calibration, and actuator conditions remain in their certified sets (Mayne et al., 2000; 2005).

C.6 Approximate value convergence

If the exact policy-evaluation operator is a γ-contraction and approximation error is bounded by $C\varepsilon_{\mathrm{tot}}$, then

$$\|V_{k+1} - V^\star\| \le \gamma\|V_k - V^\star\| + C\varepsilon_{\mathrm{tot}}. \tag{93}$$

Unrolling the recursion gives

$$\|V_k - V^\star\| \le \gamma^k\|V_0 - V^\star\| + \frac{1-\gamma^k}{1-\gamma}C\varepsilon_{\mathrm{tot}}, \tag{94}$$

and therefore

$$\limsup_{k\to\infty} \|V_k - V^\star\| \le \frac{C\varepsilon_{\mathrm{tot}}}{1-\gamma}. \tag{95}$$

D Generalization under Temporal Dependence

D.1 Forecast-risk generalization

Let $\mathcal{F}$ denote the forecasting class and suppose the prediction loss is bounded by B_ℓ and L_ℓ-Lipschitz. For independent samples, standard symmetrization and contraction arguments yield

$$\sup_{f\in\mathcal{F}} \left|\mathcal{R}(f) - \widehat{\mathcal{R}}_n(f)\right| \le 2L_\ell\widehat{\mathfrak{R}}_n(\mathcal{F}) + 3B_\ell\sqrt{\frac{\log(2/\delta)}{2n}} \tag{96}$$

with probability at least $1-\delta$.

D.2 Blocking construction

For dependent sequences, divide the chronological observations into blocks separated by gaps. Under a suitable β-mixing or α-mixing condition, the number of approximately independent blocks defines an effective sample size n_{eff}. The independent-sample bound is then applied with n_{eff} in place of n. The block length and gap should be selected from the empirical dependence range and reported with a sensitivity analysis.

D.3 Finite-policy CVaR concentration

For bounded loss $0 \le J \le B_J$, the fixed-policy CVaR estimation error satisfies

$$\left|\mathrm{CVaR}_\alpha(J) - \widehat{\mathrm{CVaR}}_\alpha(J)\right| = O_{\mathbb{P}}\left(\frac{B_J}{1-\alpha}\sqrt{\frac{\log(1/\delta)}{K}}\right). \tag{97}$$

For a finite candidate set Π, a simultaneous bound follows by replacing δ with $\delta/|\Pi|$. This result controls Monte Carlo estimation under the learned forecast law. It does not control mismatch between that law and the physical system.

D.4 Calibration scope

The analysis distinguishes:

1. finite-sample marginal conformal coverage;
2. horizon-specific empirical coverage;
3. coverage by missingness, crop stage, and action regime;
4. loss-space VaR exceedance frequency;
5. realized mean excess loss and empirical CVaR.

Exact conditional coverage is not claimed without additional structural assumptions.

E Calibration and Tail-Risk Backtesting

E.1 Chronological calibration partition

The data are divided chronologically into training, validation, calibration, and test periods. The training period estimates model parameters, the validation period selects architecture and hyperparameters, the calibration period estimates conformal corrections, and the test period is reserved for final evaluation. The calibration period is not reused for hyperparameter selection because such reuse can invalidate nominal coverage.

E.2 Finite-sample conformal correction

For n_{cal} nonconformity scores, the correction uses rank

$$k_\alpha = \lceil (n_{\text{cal}} + 1)(1-\alpha) \rceil, \tag{98}$$

with clipping to the available order-statistic range. Corrections are estimated separately by target and forecast horizon.

E.3 Coverage uncertainty and operating zones

Moving-block bootstrap intervals account for temporal dependence. The lower coverage confidence bound is used for control readiness. Define

$$E_{\text{cov}} = \left|\widehat{C}^{(1-\alpha)} - (1-\alpha)\right|. \tag{99}$$

The operating zones are:

$$E_{\text{cov}} \leq 0.03 \quad \text{continued control evaluation,} \tag{100}$$

$$0.03 < E_{\text{cov}} \leq 0.05 \quad \text{restricted control,} \tag{101}$$

$$E_{\text{cov}} > 0.05 \quad \text{fallback.} \tag{102}$$

These thresholds are engineering validation criteria, not universal distribution-free guarantees.

E.4 Loss-tail backtesting

For predicted $\widehat{\text{VaR}}_{\alpha,i}$ and realized loss J_i, define

$$E_i^{\alpha} = \mathbb{I}\left[J_i > \widehat{\text{VaR}}_{\alpha,i}\right]. \tag{103}$$

A calibrated VaR forecast should satisfy

$$\frac{1}{N}\sum_{i=1}^{N} E_i^{\alpha} \approx 1 - \alpha. \tag{104}$$

CVaR validation additionally compares predicted and realized mean excess loss:

$$m_{\alpha}^{\text{real}} = \mathbb{E}\left[J - \widehat{\text{VaR}}_{\alpha} \mid J > \widehat{\text{VaR}}_{\alpha}\right]. \tag{105}$$

E.5 Scenario adequacy

The expected empirical tail size is

$$K_{\text{tail}} = K(1-\alpha). \tag{106}$$

The study uses

$$K(1-\alpha) \geq 20 \tag{107}$$

as a practical minimum for comparative tail-risk evaluation. Results should also report bootstrap uncertainty and sensitivity to both K and α.

F Complete Experimental Protocol

F.1 Datasets and preprocessing

The study uses the semicolon-delimited drip-irrigation dataset and the multivariate tomato dataset. For each dataset, report:

- source, sampling interval, and chronological coverage;
- sensor modalities, targets, command, and executed-action variables;
- physical validity ranges and invalid-value rules;
- interpolation and imputation procedures;
- categorical encoding;
- training-only means and standard deviations;
- chronological train, validation, calibration, and test boundaries.

Datasets without measured actuation support forecasting comparison but not causal identification of irrigation effects.

F.2 Missingness and fault protocols

The evaluation includes:

- independent missingness at rates 0, 0.10, 0.25, and 0.40;

- block outages lasting 2, 4, 8, 12, and 24 intervals;
- complete loss of individual sensing modalities;
- stale, frozen, and biased observations;
- delayed valve response and pump saturation;
- simultaneous sensing and actuation faults.

F.3 Training configuration

Report:

- AdamW learning rate and weight decay;
- batch size and maximum epochs;
- KL warm-up schedule;
- gradient-clipping threshold;
- early-stopping patience and minimum improvement;
- model-selection criterion;
- scenario count;
- random seeds;
- CPU, GPU, operating system, and software versions;
- training time and peak memory.

F.4 Hyperparameter search

The search space includes

$$d_z \in \{8, 16, 24, 32\}, \qquad D \in \{32, 48, 64\}, \tag{108}$$

$$d_b \in \{64, 96, 128\}, \qquad \eta_{\mathrm{lr}} \in \{3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}\}, \tag{109}$$

$$\beta_{\max} \in \{0.03, 0.1, 0.3\}, \qquad \lambda_{\mathrm{wd}} \in \{0, 10^{-5}, 10^{-4}\}. \tag{110}$$

Report the number of randomized trials, tuning epochs, random seeds, selection metric, and final refitting protocol.

G Complete Forecasting and Representation Results

The complete results should include:

- every classical and neural forecasting model;
- RMSE and MAE by target and horizon;
- nominal 80%, 90%, and 95% empirical coverage;
- interval width and interval score;
- energy score and negative log-likelihood;
- means and standard deviations across seeds;
- paired moving-block bootstrap comparisons;
- performance at each missingness rate;
- direct, KME, routed, and fallback ablations;
- no-RFF, no-attention, no-KME, and last-state ablations;
- route probabilities by observation rate and outage duration.

Results from standardized forecasting and physical-unit control experiments must remain in separate tables.

H Scenario, Control, and Fallback Diagnostics

H.1 Scenario-count sensitivity

For each combination of K and α, report:

- objective mean and variance;
- action-selection stability;
- effective tail count $K(1-\alpha)$;
- bootstrap CVaR interval;
- run time and memory usage.

H.2 Controller evaluation

For nominal, mean-CVaR, robust, tube, restricted, and fallback controllers, report:

- mean and tail loss;
- violation frequency, magnitude, and duration;
- time to first violation;
- water and energy consumption;
- mean action and total action variation;
- valve-switching count;
- solver success, infeasibility, and latency;
- actuator residual;
- terminal-set entry;
- fallback activation and recovery.

H.3 Fallback hierarchy

The graded fallback sequence is:

1. calibrated scenario MPC;
2. restricted MPC;
3. validated crop-stage and weather-adjusted scheduling;
4. bounded hold, pulse, or stop action;
5. human intervention.

Report entry and exit triggers, hysteresis, minimum dwell time, bumpless transfer, activation frequency, mean duration, recovery time, and operational cost.

I Latent-Space Diagnostics

Belief, conditional-prior, and recognition-posterior representations are examined using PCA and t-SNE (van der Maaten & Hinton, 2008). Optional K-means clustering and silhouette analysis are used only as descriptive diagnostics (Lloyd, 1982; Rousseeuw, 1987). The visual analysis includes overlays by:

- crop stage;
- missingness rate and outage duration;
- measurement age;
- forecast target;

- commanded and executed action;
- actuator residual;
- routing weight.

Report PCA explained variance, t-SNE settings and random seeds, K-means selection metrics, and cluster sizes. Prior-posterior separation is interpreted as potential conditional-prior mismatch, while a nearly constant posterior is treated as a possible collapse diagnostic. Latent clusters are descriptive and do not establish causal, agronomic, or physically discrete regimes.

J Complexity and Runtime Analysis

J.1 Representation complexity

For batch size B, context length L, and $D_\Sigma = \sum_m D_m$, report the costs of:

- modality encoding;
- RFF construction;
- cumulative KME updates;
- direct and KME recurrent experts;
- attention and routing;
- fallback diagnostics.

Cumulative KME updates require

$$O(BLD_\Sigma) \tag{111}$$

after random-feature construction, avoiding the naive $O(BL^2D_\Sigma)$ implementation.

J.2 Forecasting and control complexity

Generating K trajectories of horizon H has cost

$$O\left(BKHD_{\mathrm{dec}}\right), \tag{112}$$

where D_{dec} summarizes recurrent-decoder dimensions. For N_π candidate plans, rollout and constraint evaluation require

$$O(N_\pi KHC_f), \tag{113}$$

while empirical CVaR computation by sorting requires

$$O(N_\pi K \log K). \tag{114}$$

J.3 Kernel-Galerkin complexity

For n transitions and D_V value features, feature construction costs

$$O(nd_xD_V). \tag{115}$$

Direct solution of the regularized normal equations costs

$$O(nD_V^2 + D_V^3) \tag{116}$$

time and $O(D_V^2)$ memory. Iterative solvers are preferable when D_V is large.

J.4 Hard and soft routing

If expert e has cost C_e, gate cost is C_g, and hard-routing usage probability is p_e, the expected hard-routing cost is

$$C_g + \sum_e p_e C_e. \tag{117}$$

Soft routing costs

$$C_g + \sum_e C_e. \tag{118}$$

Hard routing reduces latency but can introduce discontinuous expert switching. Soft routing provides smoother transitions at higher computational cost.

J.5 STREAMING TAIL EVALUATION

Chunked scenario generation, running moments, online quantile estimation, and tail buffers avoid retaining all K trajectories. Streaming removes the full

$$O(BKHD_y)$$

scenario-storage term when only predictive moments and tail statistics are required.

J.6 MEASURED RUNTIME

Report:

- CPU and GPU specifications;
- operating system and software versions;
- trainable parameter count;
- model-training time;
- mean and upper-quantile decision latency;
- peak memory;
- scaling with L, H, K, D_Σ, and N_π.