BELIEFGRAPH-JEPA: STRUCTURED LATENT WORLD MODELS FOR ACTION-CONDITIONED TIME SERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Action-conditioned time-series forecasting requires accounting for how future actions and exogenous forcings influence multiple targets through partially observed effects with different delays and persistence. Direct conditioning leaves the evolution and target-specific influence of these effects implicit in the predictor, while static relational graphs specify connections without tracking evolving effects. This motivates representing future-driver influence through structured latent states that evolve over the forecast horizon and route information to individual targets. We introduce BELIEFGRAPH-JEPA, a structured latent world model that factorizes driver influence into typed latent-effect states. These states are rolled forward under future drivers and routed through a graph to target-specific nodes, forming the predictive base of a joint-embedding predictive architecture. A capacity-controlled residual supplements this base with direct driver information. On four multi-target clinical, agricultural, environmental, and industrial systems, the framework outperforms a range of pretrained and supervised known-future-covariate baselines. Matched controls isolate latent dynamics, future rollout, graph routing, and residual capacity; future rollout and graph-first residual routing improve forecasting across all four systems.

1 INTRODUCTION

Action-conditioned time-series forecasting predicts how future actions and exogenous forcings influence multiple targets. In clinical, agricultural, environmental, and industrial systems, a propofol infusion, a ventilation setpoint, a chemical dose, or a heating-valve command is a *driver*: an input that acts through delayed, persistent, mechanism-specific hidden effects. Actions are controllable inputs; exogenous forcings are inputs the system cannot set.

Those effects are only partially observed and target-specific: blood pressure and anesthetic depth respond to the same infusion with different lags, while indoor temperature, humidity, CO_2, and light evolve together under actuators and outdoor weather.

The question is *how to represent known future drivers whose effects are delayed, mechanism-specific, and shared across multiple targets*. Covariate-aware and pretrained forecasters (Salinas et al., 2020; Lim et al., 2021; Wang et al., 2024; Ansari et al., 2025) and relational variants (Qiu et al., 2025) incorporate available drivers as covariates, tokens, or relational inputs. Other approaches combine action-conditioned prediction with rollout or residual corrections (Meng et al., 2026; Liu et al., 2026). Across these approaches, typed driver-mediated latent effects remain implicit, as schematized in Figure 1(a). Graph-based forecasters model dependencies among observed series (Wu et al., 2019; 2020), but adjacency alone does not specify how driver-mediated latent effects evolve. Latent world models already roll generic states forward under future actions (Watter et al., 2015; Hafner et al., 2019; 2025). Our approach factorizes driver influence into typed latent-effect states, rolls them forward under future drivers, and routes them through a graph to target-specific nodes, using the structured trajectory as the predictive base (Figure 1(b)).

BELIEFGRAPH-JEPA is a structured latent world model organized around a pathway: driver $\rightarrow$ mechanism-specific latent effect $\rightarrow$ observed target. A domain-informed typed graph supplies a relational prior. Past drivers initialize mechanism-specific latent-effect states, and a relation-aware GNN reads out their horizon-specific evolution as target-specific graph representations. Actions and exogenous forcings share one driver interface. The graph represents the structured driver-mediated

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

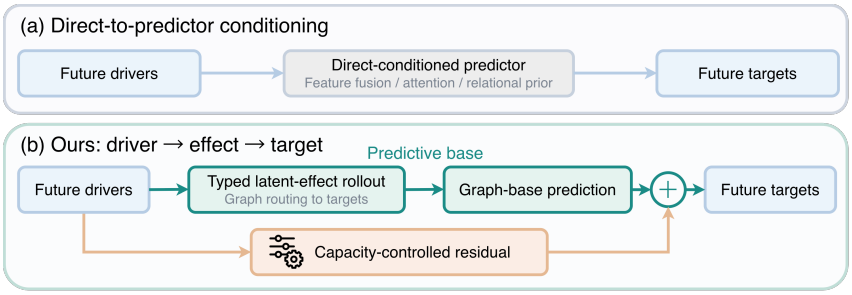

Figure 1: **Direct-to-predictor conditioning versus BeliefGraph-JEPA.** (a) Drivers condition predictions without explicit typed latent-effect states. (b) Drivers update typed latent-effect states, which are rolled to each horizon and routed to targets. Graph-base prediction of that trajectory is the predictive base; a capacity-controlled residual recovers omitted details.

component of the partially observed system, while the neural history representation captures remaining unstructured factors. *Graph-first residual routing* (GR) uses these graph representations together with observation history to form the predictive base in JEPA space (Assran et al., 2023; Ennadir et al., 2025). Direct driver tokens enter only through a gated residual adapter; a residual-magnitude penalty and graph-base supervision constrain the correction.

We evaluate this framework on four multi-target systems with shared driver-mediated mechanisms: intraoperative vital signs (VitalDB) (Lee et al., 2022), an autonomous greenhouse (Greenhouse) (Petropoulou et al., 2023), wastewater nutrient control (Wastewater) (Mohammadi et al., 2024), and a district-heating substation (PreDist) (Roelofs et al., 2026) (Table 1).

Our contributions are:

- **Structured latent driver dynamics.** We factor future-driver influence into typed latent-effect states and roll those states forward before target-specific graph routing; that structured trajectory is the predictive base. On a matched-dynamics system, this rollout recovers the correct persistence ordering and exhibits pathway-specific impulse responses.
- **Graph-first residual routing.** The graph pathway is the predictive base. Direct drivers enter only through a gated residual, whose contribution is constrained by a residual-magnitude penalty and graph-base supervision.
- **Mechanism-oriented evaluation.** We use four multi-target systems with shared driver-mediated mechanisms and matched controls that separately test latent dynamics, future rollout, graph structure, and residual capacity.

2 Related Work

Action-conditioned forecasting and known-future covariates. Covariate-aware forecasters send known future actions or exogenous inputs into the predictor. DeepAR, Temporal Fusion Transformer, and TiDE condition on known future covariates (Salinas et al., 2020; Lim et al., 2021; Das et al., 2023a); TimeXer attends over historically observed exogenous variables (Wang et al., 2024). Foundation models such as Chronos and TimesFM learn strong history extrapolation (Ansari et al., 2024; Das et al., 2023b); Chronos-2, ChronosX, and UniCA extend this line to future-known and heterogeneous covariates (Ansari et al., 2025; Arango et al., 2025; Han et al., 2026). ExoTimer integrates exogenous information with LLM-based text representations (Wu et al., 2026), and DAG injects discovered endogenous–exogenous correlations into the same forecast (Qiu et al., 2025). CEDAR augments action-conditioned rollout with event residuals (Meng et al., 2026); DRIFT adds recursive action-conditioned corrections to direct forecasts (Liu et al., 2026). BeliefGraph-JEPA represents driver influence through typed latent effects, rolled forward and routed to target-specific graph nodes, with the structured trajectory as the predictive base.

Latent world models, state-space models, and JEPA. Model-based control learns action-conditioned latent transitions $z_{t+1} = F(z_t, a_t)$ and rolls them forward under future actions (Watter

et al., 2015; Hafner et al., 2019; 2025; Zhou et al., 2024). Structured and selective state-space models provide efficient sequence backbones (Gu et al., 2021; Gu & Dao, 2023). Joint-embedding predictive architectures predict in representation space (Assran et al., 2023; Bardes et al., 2024; Assran et al., 2025), including time series, command-conditioned control, and glucose representation learning (Ennadir et al., 2025; Girgis et al., 2026; Muhammad et al., 2026). Phys-JEPA decomposes physical versus residual latent content (Nie et al., 2026); Delta-JEPA reconstructs actions from latent differences to encourage action-sensitive transitions (Zhang et al., 2026). Our rollout is factorized by typed driver-mediated effect families and routed to target-specific graph nodes, so future drivers enter as inputs to structured latent mechanisms rather than only as predictor covariates.

Graph-structured and neuro-symbolic world models. Graph neural forecasters learn or assume adjacency among observed variables (Wu et al., 2019; 2020; Cao et al., 2020); GCGNet uses graph consistency over endogenous–exogenous correlations (Li et al., 2026). LLM-constructed variable-level knowledge graphs add a static semantic prior over the same series (Sun et al., 2025). Graph world models use message passing over structured states with action nodes (Feng et al., 2025); causal representation models impose directed factor structure on static generative latents (Yang et al., 2021). Neuro-symbolic world models learn predicates and abstract dynamics for robot planning (Liang et al., 2025; 2026), and medical trajectory world models simulate clinical courses in language (Mu et al., 2026). Our BELIEFGRAPH-JEPA uses a typed graph to organize the driver–effect–target pathway for action-conditioned forecasting.

3 PROBLEM FORMULATION

At forecast anchor t, the model is given an observed context $X_{t-L_X+1:t} = (x_{t-L_X+1}, \dots, x_t)$ of length L_X, with $x_\tau \in \mathbb{R}^{d_x}$ at each time τ. The task is the continuation $Y_{t+1:t+H} = X_{t+1:t+H}$. The system is driven by $D_\tau = \{A_\tau, E_\tau\}$, where A_τ contains controllable actions and E_τ contains uncontrollable exogenous forcings. The model is conditioned on future actions over the prediction horizon; in our retrospective experiments, these are the logged actions. Future exogenous inputs, when used, hold the last context value (Table 1). Drivers act through delayed, mechanism-specific hidden effects. Those effects are represented as typed latent states and evolved under future drivers. We treat the setting as a partially observed controlled dynamical system. The model represents the structured driver-mediated component by typed latent-effect states Z_t^F, a deterministic summary of past driver effects that is evolved under future drivers and routed to observed targets. Remaining unstructured factors are captured by the neural history representation.

Joint-embedding prediction is performed in representation space (Assran et al., 2023; Ennadir et al., 2025). An online context encoder E_θ maps the observed history to H_t^X. An exponential-moving-average (EMA) target encoder $E_{\bar{\theta}}$ encodes the future sequence, $z^*_{t+1:t+H} = \mathrm{sg}\big(E_{\bar{\theta}}(X_{t+1:t+H})\big)$, with sg the stop-gradient. A predictor produces $\hat{z}_{t+1:t+H}$ from H_t^X, past drivers, and available future drivers $D_{t:t+H-1}$. An observation decoder maps $\hat{z}_{t+h}$ to $\hat{y}_{t+h}$. The latent and observation objectives are

$$\mathcal{L}_{\mathrm{JEPA}} = \frac{1}{H}\sum_{h=1}^{H} \|\hat{z}_{t+h} - z^*_{t+h}\|_2^2, \qquad \mathcal{L}_{\mathrm{obs}} = \frac{1}{H}\sum_{h=1}^{H} \|\hat{y}_{t+h} - y_{t+h}\|_2^2. \tag{1}$$

The base objective is $\mathcal{L} = \mathcal{L}_{\mathrm{JEPA}} + \lambda_{\mathrm{obs}}\,\mathcal{L}_{\mathrm{obs}}$. Section 4 instantiates the predictor by rolling Z_t^F under subsequent drivers and using the resulting graph trajectory as the predictive base.

4 METHOD

BELIEFGRAPH-JEPA has two parts (Figure 2). Section 4.1 turns a typed driver–effect–target prior into typed latent-effect states and rolls them under future drivers. Section 4.2 makes the resulting horizon-specific trajectory the predictive base, with direct driver information entering as a capacity-controlled residual. Implementation details are in Appendix F.

4.1 TYPED LATENT-EFFECT DYNAMICS AND GRAPH ROLLOUT

Typed relational prior. We assume a sparse typed relational prior $G_{\mathrm{static}} = (V_D, V_F, V_Y, \mathcal{E})$ with $V_D = V_A \cup V_E$, where V_A, V_E, V_F, and V_Y are action, exogenous, latent-effect, and target nodes.

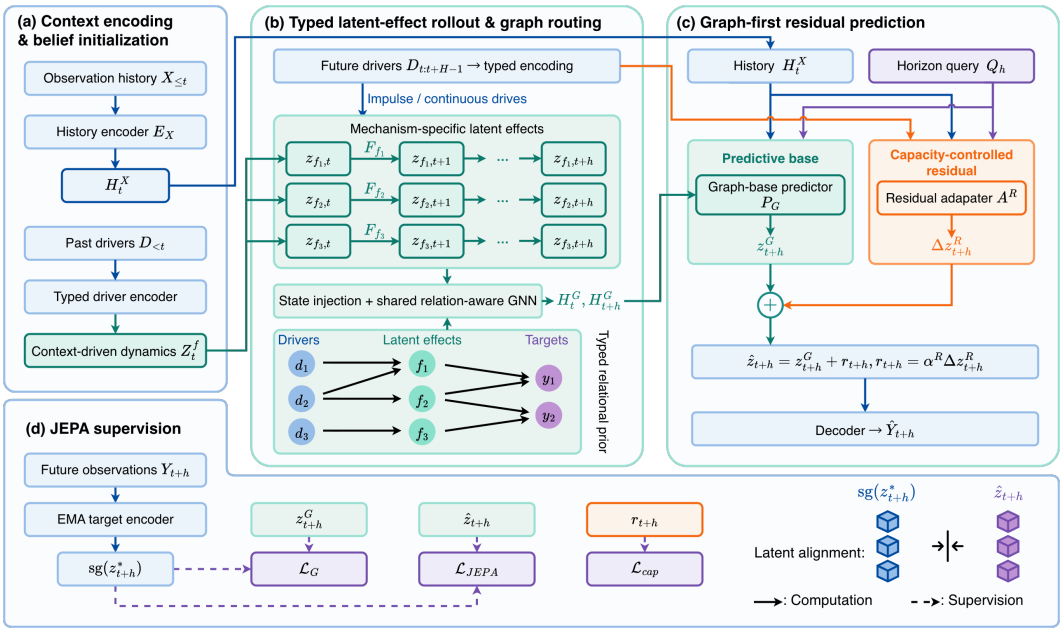

Figure 2: **BeliefGraph-JEPA: structured latent-effect rollout as the predictive base.** (a) Past observations and past drivers initialize a history state H_t^X and a family-effect state Z_t^f. (b) Future drivers roll typed latent-effect states; a shared relation-aware GNN, guided by the typed driver–effect–target prior, produces the horizon-specific graph trajectory. (c) z_{t+h}^G is the predictive base; a capacity-controlled residual $\alpha^R \Delta z_{t+h}^R$ recovers omitted details. (d) JEPA, graph-base, and capacity losses supervise the aligned latent.

$\mathcal{E}$ contains the directed driver→effect and effect→target edges, each with a relation label, sign, and confidence. An instantaneous-effect head extends the readout to contemporaneous algebraic targets on PreDist (Appendix E). The prior names which driver families may update which latent effects, and which effects may reach which targets, while sample-specific magnitude and temporal state are learned. V_E is a domain-selected exogenous subset; exogenous channels not placed in V_E remain available as direct predictor tokens. Appendix L records the frozen LLM-assisted construction used in experiments.

Unified driver representation. Actions and exogenous forcings share the interface $D_\tau = \{A_\tau, E_\tau\}$ and the same output dimension. Encoders keep three attributes separate: numeric type (categorical versus continuous), temporal modality (impulse versus continuous forcing), and role (controllable versus exogenous). Each encoded driver yields a token for the predictor and a latent drive u_τ that updates mechanism-specific effect states. Impulse events are one-shot jumps; continuous forcings remain active over physical time. Concurrent drives of the same family and modality are summed. Exogenous graph nodes are refreshed from observed history and, over the horizon, from the last context value; persistence lives in the effect state.

Current graph belief. For latent-effect family f and a generic time τ, impulse and continuous drives update a recurrent state

$$z_{f,\tau+1} = \lambda_f \odot z_{f,\tau} + \alpha_f^{\mathrm{imp}} \odot \phi_f^{\mathrm{imp}}(u_{f,\tau}^{\mathrm{imp}}) + (1-\lambda_f) \odot \alpha_f^{\mathrm{cont}} \odot \phi_f^{\mathrm{cont}}(u_{f,\tau}^{\mathrm{cont}}), \tag{2}$$

with σ the sigmoid and $\lambda_f = \sigma(\rho_f) \in (0,1)$ a learned per-family persistence, α_f^{imp} and α_f^{cont} learned injection scales in $(0,1)$, and ϕ_f^{imp}, ϕ_f^{cont} pointwise maps of the encoded drives. The impulse term is a state jump followed by decay. The continuous term is stable first-order forcing; $(1-\lambda_f)$ keeps a constant input from being amplified as persistence approaches one. Past drivers $D_{t-L_A:t-1}$ are rolled through Eq. (2) from a zero initial state to the forecast anchor t. The packed family-effect state $Z_t^F = (z_{f,t})_{f \in V_F}$ factorizes driver influence by mechanism family. The latent-effect window L_A equals the context length L_X.

At the forecast anchor, a gate injects each family state into its static node embedding e_f^{static},

$$h_{f,\tau}^{(0)} = e_f^{\text{static}} + \gamma_f W_f^G z_{f,\tau}, \qquad \gamma_f = \sigma(\delta_f). \tag{3}$$

The injection $\gamma_f W_f^G z_{f,\tau}$ keeps the magnitude of $z_{f,\tau}$, which encodes accumulation under the recurrence. History updates $z_{f,\tau}$ at every step and applies Eq. (3) to the anchor state $z_{f,t}$. $H_t^{(0)}$ initializes every node: effect nodes carry this injection, and action and exogenous nodes carry their encoded driver features. A two-layer relation-aware GNN (Scarselli et al., 2009; Schlichtkrull et al., 2018) then propagates along G_{static}, $H_t^G = \text{GNN}(G_{\text{static}}, H_t^{(0)})$. Action and exogenous neighbors of an effect node send role-conditioned messages; the two aggregates are concatenated before the effect update. Target nodes receive messages from these effect representations. Thus H_t^G is the GNN readout of Z_t^F, and the current graph belief is the pair (Z_t^F, H_t^G).

Future latent-effect rollout. Z_t^F summarizes drivers through time $t-1$. The same per-family transition (Eq. 2) is applied causally under subsequent drivers. Writing $\mathcal{F}$ for that map on the packed family state, $Z_{t+h}^F = \mathcal{F}(Z_{t+h-1}^F, D_{t+h-1})$ for $h = 1, \ldots, H$, the horizon-h state is obtained from Z_t^F and $D_{t:t+h-1}$. Each Z_{t+h}^F is injected by Eq. (3), and action and exogenous nodes at that step carry encodings of D_{t+h-1}. The same GNN returns the horizon-specific readout H_{t+h}^G. Latent recurrence is sequential and lightweight. The H GNN evaluations are independent given $\{Z_{t+h}^F\}$ and these driver encodings, and are batched over the horizon. Future exogenous inputs hold the last context value (Section 3).

History H_t^X, the anchor graph H_t^G, and the horizon readout H_{t+h}^G enter a predictor P with encoded tokens of the future drivers $D_{t:t+h-1}$, $\hat{z}_{t+h} = P(H_t^X,\, H_t^G,\, H_{t+h}^G,\, D_{t:t+h-1},\, Q_h)$, where Q_h is a learned query embedding for horizon h. The horizon queries are decoded jointly, with query h using H_{t+h}^G and the driver prefix $D_{t:t+h-1}$. We call this architecture *future rollout* (FR).

4.2 Structured Rollout as the Predictive Base

FR lets the horizon graph readout and encoded future driver tokens enter one predictor. GR designates the graph-mediated trajectory as the predictive base $z_{t+h}^G = P_G(H_t^X,\, H_t^G,\, H_{t+h}^G,\, Q_h)$. The base z_{t+h}^G is a JEPA-space prediction. History H_t^X and the anchor graph H_t^G enter P_G with the horizon readout H_{t+h}^G. Future drivers reach this base through the rollout $Z_{t+h}^F = \mathcal{F}(Z_{t+h-1}^F, D_{t+h-1})$ and the subsequent GNN, so P_G conditions on them through H_{t+h}^G. Encoded tokens of $D_{t:t+h-1}$ enter a residual adapter A_R with a stop-gradient copy of this base, $\Delta z_{t+h}^R = A_R(H_t^X,\, D_{t:t+h-1},\, \text{sg}(z_{t+h}^G),\, Q_h)$. The prediction is then

$$\hat{z}_{t+h} = z_{t+h}^G + \alpha^R \Delta z_{t+h}^R, \qquad \alpha^R = \alpha_{\max}^R\, \sigma(a), \tag{4}$$

where $\alpha_{\max}^R < 1$ is a fixed cap and a is a learned scalar. The residual adapter is zero-initialized, so training begins from $\hat{z}_{t+h} \approx z_{t+h}^G$.

Graph-base supervision trains the base to match the JEPA target on its own, and a residual-magnitude penalty keeps the residual a correction,

$$\mathcal{L}_G = \frac{1}{H}\sum_{h=1}^{H} \|z_{t+h}^G - z_{t+h}^*\|_2^2, \qquad \mathcal{L}_{\text{cap}} = \frac{1}{H}\sum_{h=1}^{H} \|\alpha^R \Delta z_{t+h}^R\|_2^2. \tag{5}$$

The training objective extends the base objective to $\mathcal{L} = \mathcal{L}_{\text{JEPA}} + \lambda_G\, \mathcal{L}_G + \lambda_{\text{cap}}\, \mathcal{L}_{\text{cap}} + \lambda_{\text{obs}}\, \mathcal{L}_{\text{obs}}$. Residual reliance is the magnitude of the gated correction relative to the combined magnitude of the base and correction,

$$\rho_R = \frac{1}{H}\sum_{h=1}^{H} \frac{\|\alpha^R \Delta z_{t+h}^R\|_2}{\|z_{t+h}^G\|_2 + \|\alpha^R \Delta z_{t+h}^R\|_2 + \varepsilon}. \tag{6}$$

where $\varepsilon > 0$ is a small constant. Because A_R learns the scale of Δz_{t+h}^R, reliance is read from ρ_R rather than from the gate α^R. The main method is FR-GR, which uses future rollout as the graph base and a capacity-controlled residual on that base.

Table 1: Multi-target forecasting protocols. †PreDist uses an instantaneous-effect head for algebraic heat power and flow (Appendix E).

Dataset	Targets	Actions	Exogenous
VitalDB	MBP, HR, SpO_2, BIS (4)	propofol, remifentanil rates	—
Greenhouse	indoor T, RH, CO_2, PAR (4)	heating, vents, screens, light, CO_2 dosing	outdoor weather (8)
Wastewater	DO, NH_4, PO_4 (3)	metal-salt dosing; process phases	inflow, water temperature
PreDist†	supply/return T, heat power, flow (5)	valve, setpoint; pump, mode, 3-way valve	outdoor T, grid supply T

5 EXPERIMENTS

Table 1 lists the four multi-target protocols, in each of which several targets share typed driver-mediated mechanisms. For retrospective evaluation, all action-conditioned methods receive the same logged future actions. Future exogenous inputs, when present, hold the last context value over the horizon. Sampling grids and forecast windows are in Table 6; preprocessing and training settings are in Appendix B and Appendix G.

Metrics and controls. Primary metrics are pooled normalized MAE and RMSE on the $[-1, 1]$ target scale, reported as mean $\pm$ sample standard deviation over three seeds $\{42, 43, 44\}$. Baselines are last-value persistence, history-only Chronos-2 and TimesFM 2.5, and covariate-aware Chronos-2-Cov, DAG, and TFT, the last three with oracle future actions and future exogenous inputs held at the last context value (Ansari et al., 2025; Das et al., 2023b; Qiu et al., 2025; Lim et al., 2021). The sections below isolate latent dynamics, future rollout, graph structure, and residual capacity in turn.

5.1 OVERALL FORECASTING PERFORMANCE

Table 2 reports the four-system comparison; per-target physical-unit errors are in Appendix D. TFT is one of the most competitive external baselines in our experiments and uses the same constructed-exogenous and oracle-action future protocol as DAG and BELIEFGRAPH-JEPA. FR-GR attains lower nMAE and nRMSE on Greenhouse, VitalDB, and PreDist; TFT attains slightly lower nMAE on Wastewater, with nRMSE essentially tied. Per-target gains vary across variables (Appendix D): FR-GR attains the lowest mean MAE and RMSE on Greenhouse indoor temperature, humidity, and CO_2, and on VitalDB BIS; on PreDist, it also improves over TFT on heat power and flow.

To isolate where those gains arise, a cumulative ladder adds one module at a time up to FR: history only (HO), future-driver tokens (DC), a static typed GNN (SG), instant effect nodes without temporal persistence (IE), and a current recurrent belief without future rollout (RB). Direct future-driver conditioning already provides a gain over history-only prediction: DC improves over HO on all four systems, most clearly on Greenhouse and Wastewater. A static typed GNN adds little beyond those tokens. Persistent latent-effect dynamics then improve consistently: RB improves over IE on every dataset, and rolling those states forward under future drivers improves again: FR improves over RB.

FLAT keeps the FR predictor but replaces GNN message passing with an unstructured aggregation of effect states. FLAT-GR applies the same base–residual decomposition, auxiliary base supervision, and residual-capacity penalty to that unstructured readout. UR-GR keeps the FLAT-GR readout and residual routing, and replaces the family-specific latent states with one shared state. RB-GR applies the same design to the current-belief model. FLAT stays close to FR, indicating that much of the forecasting benefit is retained without GNN message passing. This design improves both readouts: FLAT-GR over FLAT, and FR-GR over FR. FR-GR also improves over RB-GR, supporting future rollout under matched residual routing. UR-GR performs similarly to FLAT-GR overall. FR-GR achieves lower mean errors than UR-GR on all four datasets; its gains over FLAT-GR are clearest on Greenhouse, smaller on PreDist, and limited on Wastewater and VitalDB. These results indicate dataset-dependent gains from explicit graph routing. The remaining diagnostics use FR-GR as the main method.

5.2 MAKING STRUCTURED ROLLOUT THE PREDICTIVE BASE

Table 3 compares fusion controls. FR fuses the graph trajectory and direct future drivers in one predictor. A hard graph bottleneck (HB) uses the same predictor but withholds graph-covered future-

Table 2: Test nMAE / nRMSE (↓) on the four multi-target systems, mean ± sample std over three seeds. Best in each column in bold, second-best underlined. External baselines, a cumulative ladder up to FR, then unstructured-readout and GR controls. †Uses the instantaneous-effect head (Appendix E).

	Greenhouse		Wastewater		VitalDB		PreDist†	
Method	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE
Persistence	0.1552	0.2603	0.2911	0.4548	0.0655	0.1285	0.0677	0.1115
Chronos-2	0.1572	0.2620	0.2349	0.3836	0.0640	0.1244	0.0598	0.1004
Chronos-2-Cov	0.1464	0.2521	0.1765	0.3118	0.0639	0.1253	0.0554	0.0986
TimesFM 2.5	0.1562	0.2579	0.2231	0.3584	0.0611	0.1177	0.0563	0.0951
DAG	$0.1552_{\pm 0.0011}$	$0.2441_{\pm 0.0038}$	$0.1761_{\pm 0.0022}$	$0.2957_{\pm 0.0006}$	$0.0629_{\pm 0.0001}$	$0.1173_{\pm 0.0001}$	$0.0512_{\pm 0.0004}$	$0.0835_{\pm 0.0011}$
TFT	$\underline{0.1327}_{\pm 0.0014}$	$0.2007_{\pm 0.0026}$	$\mathbf{0.1201}_{\pm 0.0008}$	$\mathbf{0.2139}_{\pm 0.0031}$	$0.0588_{\pm 0.0017}$	$0.1083_{\pm 0.0018}$	$0.0576_{\pm 0.0048}$	$0.0890_{\pm 0.0073}$
HO	$0.1599_{\pm 0.0015}$	$0.2324_{\pm 0.0014}$	$0.2212_{\pm 0.0028}$	$0.3293_{\pm 0.0017}$	$0.0607_{\pm 0.0002}$	$0.1097_{\pm 0.0000}$	$0.0622_{\pm 0.0015}$	$0.0958_{\pm 0.0018}$
DC	$0.1486_{\pm 0.0008}$	$0.2137_{\pm 0.0004}$	$0.1520_{\pm 0.0033}$	$0.2493_{\pm 0.0043}$	$0.0596_{\pm 0.0005}$	$0.1056_{\pm 0.0001}$	$0.0578_{\pm 0.0015}$	$0.0889_{\pm 0.0011}$
SG	$0.1461_{\pm 0.0005}$	$0.2121_{\pm 0.0012}$	$0.1461_{\pm 0.0034}$	$0.2420_{\pm 0.0036}$	$0.0604_{\pm 0.0000}$	$0.1059_{\pm 0.0002}$	$0.0570_{\pm 0.0023}$	$0.0883_{\pm 0.0028}$
IE	$0.1479_{\pm 0.0036}$	$0.2139_{\pm 0.0040}$	$0.1479_{\pm 0.0030}$	$0.2444_{\pm 0.0035}$	$0.0600_{\pm 0.0003}$	$0.1057_{\pm 0.0002}$	$0.0578_{\pm 0.0017}$	$0.0891_{\pm 0.0019}$
RB	$0.1461_{\pm 0.0013}$	$0.2122_{\pm 0.0017}$	$0.1443_{\pm 0.0009}$	$0.2394_{\pm 0.0011}$	$0.0597_{\pm 0.0003}$	$0.1049_{\pm 0.0004}$	$0.0551_{\pm 0.0017}$	$0.0862_{\pm 0.0018}$
FR	$0.1392_{\pm 0.0028}$	$0.2030_{\pm 0.0020}$	$0.1304_{\pm 0.0018}$	$0.2230_{\pm 0.0025}$	$0.0587_{\pm 0.0002}$	$0.1040_{\pm 0.0002}$	$0.0536_{\pm 0.0021}$	$0.0816_{\pm 0.0029}$
FLAT	$0.1401_{\pm 0.0016}$	$0.2052_{\pm 0.0033}$	$0.1304_{\pm 0.0038}$	$0.2208_{\pm 0.0039}$	$0.0589_{\pm 0.0001}$	$0.1035_{\pm 0.0000}$	$0.0561_{\pm 0.0025}$	$0.0841_{\pm 0.0023}$
FLAT-GR	$0.1344_{\pm 0.0008}$	$0.2004_{\pm 0.0022}$	$0.1250_{\pm 0.0023}$	$\underline{0.2157}_{\pm 0.0042}$	$\mathbf{0.0580}_{\pm 0.0002}$	$\underline{0.1034}_{\pm 0.0002}$	$0.0511_{\pm 0.0020}$	$\underline{0.0776}_{\pm 0.0013}$
UR-GR	$0.1331_{\pm 0.0014}$	$\underline{0.1989}_{\pm 0.0022}$	$0.1293_{\pm 0.0018}$	$0.2208_{\pm 0.0026}$	$\underline{0.0583}_{\pm 0.0003}$	$0.1035_{\pm 0.0001}$	$\underline{0.0506}_{\pm 0.0002}$	$\underline{0.0776}_{\pm 0.0004}$
RB-GR	$0.1403_{\pm 0.0019}$	$0.2086_{\pm 0.0017}$	$0.1326_{\pm 0.0009}$	$0.2249_{\pm 0.0016}$	$\underline{0.0583}_{\pm 0.0002}$	$0.1037_{\pm 0.0002}$	$0.0535_{\pm 0.0018}$	$0.0832_{\pm 0.0023}$
FR-GR	$\mathbf{0.1285}_{\pm 0.0016}$	$\mathbf{0.1928}_{\pm 0.0033}$	$\underline{0.1245}_{\pm 0.0018}$	$\underline{0.2157}_{\pm 0.0028}$	$\mathbf{0.0580}_{\pm 0.0003}$	$\mathbf{0.1032}_{\pm 0.0001}$	$\mathbf{0.0495}_{\pm 0.0003}$	$\mathbf{0.0768}_{\pm 0.0004}$

Table 3: Routing and bottleneck diagnostics.

	Greenhouse		Wastewater		VitalDB		PreDist†	
Method	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE
FR	$0.1392_{\pm 0.0028}$	$0.2030_{\pm 0.0020}$	$0.1304_{\pm 0.0018}$	$0.2230_{\pm 0.0025}$	$0.0587_{\pm 0.0002}$	$0.1040_{\pm 0.0002}$	$0.0536_{\pm 0.0021}$	$0.0816_{\pm 0.0029}$
HB	$0.1370_{\pm 0.0014}$	$0.2009_{\pm 0.0011}$	$\underline{0.1277}_{\pm 0.0011}$	$0.2205_{\pm 0.0028}$	$0.0588_{\pm 0.0002}$	$0.1038_{\pm 0.0002}$	$0.0518_{\pm 0.0016}$	$0.0798_{\pm 0.0024}$
FR-GR-GO	$0.1381_{\pm 0.0009}$	$0.2014_{\pm 0.0035}$	$0.1354_{\pm 0.0014}$	$0.2262_{\pm 0.0018}$	$0.0593_{\pm 0.0003}$	$0.1044_{\pm 0.0002}$	$0.0534_{\pm 0.0007}$	$0.0813_{\pm 0.0011}$
FR-GR-U	$\underline{0.1319}_{\pm 0.0009}$	$\underline{0.1944}_{\pm 0.0004}$	$0.1280_{\pm 0.0004}$	$\underline{0.2172}_{\pm 0.0010}$	$\mathbf{0.0578}_{\pm 0.0001}$	$\underline{0.1033}_{\pm 0.0002}$	$\underline{0.0507}_{\pm 0.0025}$	$\underline{0.0777}_{\pm 0.0037}$
RB-GR	$0.1403_{\pm 0.0019}$	$0.2086_{\pm 0.0017}$	$0.1326_{\pm 0.0009}$	$0.2249_{\pm 0.0016}$	$0.0583_{\pm 0.0002}$	$0.1037_{\pm 0.0002}$	$0.0535_{\pm 0.0018}$	$0.0832_{\pm 0.0023}$
FR-GR	$\mathbf{0.1285}_{\pm 0.0016}$	$\mathbf{0.1928}_{\pm 0.0033}$	$\mathbf{0.1245}_{\pm 0.0018}$	$\mathbf{0.2157}_{\pm 0.0028}$	$\underline{0.0580}_{\pm 0.0003}$	$\mathbf{0.1032}_{\pm 0.0001}$	$\mathbf{0.0495}_{\pm 0.0003}$	$\mathbf{0.0768}_{\pm 0.0004}$

driver tokens from the direct path. FR-GR-GO evaluates the FR-GR graph base at $\alpha^R = 0$; FR-GR-U uses $\alpha^R = 1$ and no residual-capacity penalty.

On all four datasets, HB matches or slightly improves upon FR: graph-covered future driver information can be carried through the rolled graph pathway without those tokens on the direct path. The graph-only readout FR-GR-GO stays close to FR on Greenhouse and PreDist, so the rolled graph trajectory is already a structured predictive base. Removing the residual from FR-GR raises error on every dataset, so the direct correction still captures information beyond that base. An unrestricted residual (FR-GR-U) improves over FR on every dataset: graph-first prediction helps even without a residual-capacity penalty. Capacity control further improves Greenhouse, Wastewater, and PreDist, while FR-GR and FR-GR-U remain comparable on VitalDB. For FR-GR, the relative residual magnitude ρ_R is 1.4–3.9% (Appendix I). FR-GR also outperforms HB on all four datasets: a capacity-controlled residual complements the graph base better than a hard bottleneck.

5.3 What does graph structure contribute?

Where Table 2 varies whether readout is a typed GNN, Table 4 holds the FR-GR backbone fixed and varies which typed graph is used, by changing it during retraining or only at inference. These controls use Greenhouse and Wastewater, where the domain prior is sparse and a typed fully-connected alternative is well-defined. Panel A retrains FR-GR with a learned fully-connected graph (LFC) and a pre-specified slot-shuffled prior that remaps driver→effect slot inputs while keeping effect→target edges fixed. Panel B swaps the prior only at evaluation, testing the domain-graph checkpoint under the slot shuffle and the slot-shuffled checkpoint under the domain graph.

The domain graph attains the lowest error on both datasets. A learned fully-connected graph is worse, especially on Greenhouse, supporting the sparse typed prior over learning connectivity from a complete graph. Slot-shuffled retraining stays close, indicating that an alternative typed routing remains competitive after retraining. In contrast, swapping the same assignment only at inference raises error sharply, especially on Wastewater (nMAE 0.1245 to 0.1610); the reverse swap also

Table 4: Graph-structure controls.

Method	Greenhouse nMAE	Greenhouse nRMSE	Wastewater nMAE	Wastewater nRMSE
A. Retraining				
Domain graph (FR-GR)	$\mathbf{0.1285}_{\pm 0.0016}$	$\mathbf{0.1928}_{\pm 0.0033}$	$\mathbf{0.1245}_{\pm 0.0018}$	$\mathbf{0.2157}_{\pm 0.0028}$
Learned-FC (LFC)	$0.1357_{\pm 0.0013}$	$0.2029_{\pm 0.0010}$	$0.1256_{\pm 0.0007}$	$\underline{0.2170}_{\pm 0.0011}$
Shuffled-slot	$\underline{0.1309}_{\pm 0.0043}$	$\underline{0.1935}_{\pm 0.0055}$	$\underline{0.1254}_{\pm 0.0007}$	$0.2189_{\pm 0.0014}$
B. Inference-time prior swap				
Domain → shuffled-slot	$0.1356_{\pm 0.0014}$	$0.2013_{\pm 0.0005}$	$0.1610_{\pm 0.0019}$	$0.2639_{\pm 0.0014}$
Shuffled-slot → domain	$0.1365_{\pm 0.0046}$	$0.2008_{\pm 0.0053}$	$0.1472_{\pm 0.0046}$	$0.2503_{\pm 0.0081}$

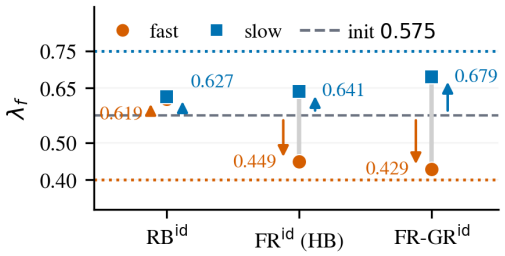

Figure 3: **Learned persistences under matched dynamics.** Learned λ_f from init 0.575 toward $(0.40, 0.75)$ (seed 42).

Table 5: **Peak future-pulse responses under matched dynamics.** Peak $|I_{d\to y}|$, the maximum absolute response of target y to a unit future pulse on driver d, on the two wired pathways and the off-path/null channels (seed 42).

	fast	slow	off-path
Truth	0.60	0.25	0
RB^{id}	0	0	$< 0.1\times10^{-3}$
FR^{id} (HB)	0.37	2×10^{-4}	9.6×10^{-3}
FR-GR^{id}	0.43	0.14	8.2×10^{-3}

increases error. The typed graph is therefore a relational inductive bias: predictions use the routing learned during training, while an alternative typed assignment remains competitive once retrained. The remaining experiments examine how the routed belief evolves under future drivers.

5.4 Can future rollout recover matched latent-effect dynamics?

Real forecasting data do not reveal ground-truth latent-effect states or persistences, and Appendix J shows that multiple persistence configurations can achieve nearly identical forecast error. We therefore probe capacity on a matched leaky-integrator with known fast and slow latent mechanisms: a class-matched check of whether future rollout recovers distinct persistences and pathway-specific driver-to-target responses (Figure 3, Table 5; Appendix K).

The two wired effect families have $\lambda_{\text{fast}} = 0.40$ and $\lambda_{\text{slow}} = 0.75$, both initialized at 0.575. Correct recovery therefore requires the learned persistences to move in opposite directions. A third driver has no target path and serves as a null control. Superscript id denotes the matched protocol of Appendix K, which isolates future driver-to-latent dynamics by restricting direct prediction shortcuts.

Figure 3 reports learned persistences against the shared initialization and the two ground-truth scales. Table 5 reports the corresponding peak driver-to-target pulse responses. Without future rollout, RB^{id} leaves the two persistences nearly equal and produces no pulse response. FR^{id} (HB) recovers the correct ordering ($0.449 < 0.641$) and halves MAE from 0.063 to 0.031, but only the fast pathway responds (0.371); the slow pathway stays near zero. FR-GR^{id} moves the estimates further toward the ground truth ($0.429 < 0.679$), reduces MAE to 0.013, and recovers both wired responses (0.430 and 0.141), with off-path and null channels near zero. Among the compared variants, future rollout recovers the correct persistence ordering, while the full FR-GR^{id} model exhibits responses on both wired pathways, including the slow pathway. Appendix K.4 examines a saturating $\tanh$ mismatch: the slow-path response remains near zero at 40 epochs but increases at 80 epochs, while both response magnitudes remain underestimated.

5.5 Do forecasts use future actions?

Figure 4 reports an input-dependence diagnostic: error sensitivity to later actions (Appendix H). At evaluation we add a fixed signed offset to every continuous future-action channel on the second half

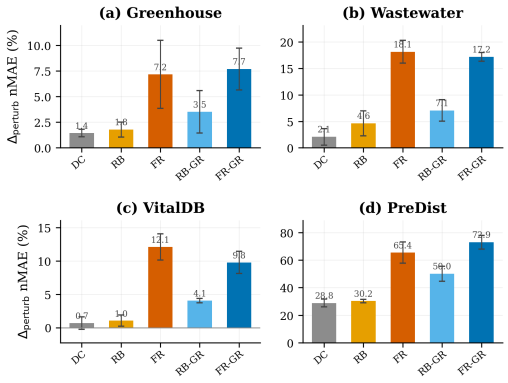

Figure 4: **Future-action perturbation.** Δ_{perturb} after a second-half offset (last 90 min / 3 min / 3 h; Appendix H). Bars are mean $\pm$ sample std over three seeds.

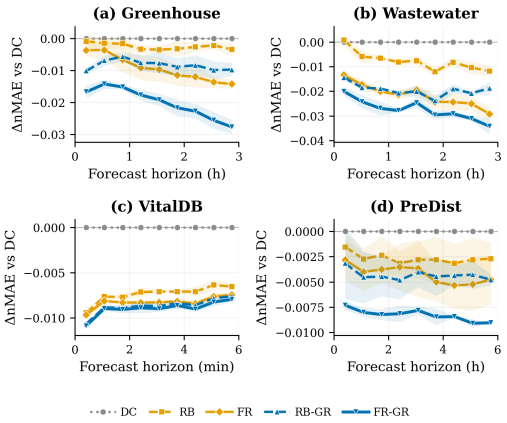

Figure 5: **Horizon-wise nMAE relative to DC.** Mean $\pm$ sample std over three seeds; negative is lower error. The x-axis is physical lead time.

of the horizon, hold exogenous inputs and categorical actions fixed, and report the percent nMAE increase $\Delta_{\text{perturb}} = (\text{nMAE}_{\text{perturb}} - \text{nMAE}_{\text{clean}})/\text{nMAE}_{\text{clean}}$ on the full window.

On Greenhouse, Wastewater, and VitalDB, DC and RB show only weak sensitivity to the second-half offset: Δ_{perturb} is 0.7–4.6%, despite DC receiving future-driver tokens directly. Explicitly rolling the latent-effect belief under those later drivers produces a much larger forecast-error increase: 7.2%, 18.1%, and 12.1%. RB-GR is intermediate (3.5%, 7.1%, 4.1%): the residual adapter sees the perturbed tokens, while the graph base remains the current belief. FR-GR matches or exceeds FR on Greenhouse (7.7% versus 7.2%) and remains comparable on Wastewater (17.2% versus 18.1%), and remains an order of magnitude above RB on VitalDB (9.8% versus 1.0%). On PreDist, Δ_{perturb} is already 28.8% for DC. Future rollout still produces the largest error increase (65.4% for FR, 72.9% for FR-GR). Future rollout therefore makes predictions depend on later actions, including offsets that token-only conditioning largely ignores.

5.6 Horizon-wise benefit of future rollout

Figure 5 reports per-horizon nMAE relative to DC. Negative values are lower error. FR-GR is lowest at every lead time. Across all four datasets, FR and FR-GR generally outperform their current-belief counterparts, RB and RB-GR, across lead times. The gap widens with lead time on Greenhouse and Wastewater, most steadily on Greenhouse. On PreDist, FR-GR remains lowest throughout. On VitalDB's six-minute horizon, FR-GR slightly improves over RB-GR. Together, these comparisons support evolving latent effects under future drivers, with the magnitude of the benefit varying across datasets and lead times.

6 Conclusion

We presented BeliefGraph-JEPA, a structured latent world model for action-conditioned forecasting. Future drivers enter latent dynamics through typed latent-effect states and graph-mediated rollout, with a capacity-controlled residual. On four multi-target systems across distinct domains, the framework outperforms a range of pretrained and supervised known-future-covariate baselines. Explicit graph routing provides additional, dataset-dependent gains. Matched ablations isolate latent dynamics, future rollout, graph routing, and residual capacity. On a matched-dynamics system, the model recovers the correct persistence ordering and exhibits pathway-specific impulse responses. On the real systems, later-action offsets change the forecast once the latent-effect states are rolled forward. Identifying latent mechanisms and extending the framework to action planning remain directions for future work.

AI USE STATEMENT

We used generative AI tools to polish the manuscript for grammar, clarity, and readability, to propose dataset-specific typed relational graph priors, to edit research code, and to suggest visual refinements of figures. The authors reviewed all AI-assisted work: translated and polished text was checked; proposed graph priors were schema-validated, reviewed, and frozen before training (Appendix L); modified code was inspected and tested; and figure suggestions were applied by the authors. Generative AI was not used to generate or clean datasets, develop the method or conceptual framework, formulate claims or hypotheses, design experiments, implement the method, or interpret results. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.

Abdul Fatir Ansari, Oleksandr Shchur, Jaris Küken, Andreas Auer, Boran Han, Pedro Mercado, Syama Sundar Rangapuram, Huibin Shen, Lorenzo Stella, Xiyuan Zhang, et al. Chronos-2: From univariate to universal forecasting. *arXiv preprint arXiv:2510.15821*, 2025.

Sebastian Pineda Arango, Pedro Mercado, Shubham Kapoor, Abdul Fatir Ansari, Lorenzo Stella, Huibin Shen, Hugo Senetaire, Caner Turkmen, Oleksandr Shchur, Danielle C Maddix, et al. Chronosx: Adapting pretrained time series models with exogenous variables. *arXiv preprint arXiv:2503.12107*, 2025.

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15619–15629. IEEE, 2023.

Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025.

Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. *arXiv preprint arXiv:2404.08471*, 2024.

Defu Cao, Yujing Wang, Juanyong Duan, Ce Zhang, Xia Zhu, Congrui Huang, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, and Qi Zhang. Spectral temporal graph neural network for multivariate time-series forecasting. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17766–17778, 2020.

Cursor Team. Introducing Grok 4.6. `https://cursor.com/blog/grok-4-6`, August 2026.

Abhimanyu Das, Weihao Kong, Andrew Leach, Shaan Mathur, Rajat Sen, and Rose Yu. Long-term forecasting with tide: Time-series dense encoder. *arXiv preprint arXiv:2304.08424*, 2023a.

Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. *arXiv preprint arXiv:2310.10688*, 2023b.

Anurag Das, David Kerr, Namino Glantz, Wendy Bevier, Rony Santiago, Ricardo Gutierrez-Osuna, and Bobak J. Mortazavi. Cgmacros: a pilot scientific dataset for personalized nutrition and diet monitoring. *Scientific Data*, 12(1):1557, 2025.

Sofiane Ennadir, Siavash Golkar, and Leopoldo Sarra. Joint embeddings go temporal. *arXiv preprint arXiv:2509.25449*, 2025.

Tao Feng, Yexin Wu, Guanyu Lin, and Jiaxuan You. Graph world model. In *International Conference on Machine Learning*, pp. 16931–16955. PMLR, 2025.

Abanoub M. Girgis, Alvaro Valcarce, and Mehdi Bennis. Time-series jepa for predictive remote control under capacity-limited networks. *IEEE Internet of Things Journal*, 13(7):14617–14632, 2026. doi: 10.1109/JIOT.2025.3650435.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021.

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

Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 640(8059):647–653, 2025.

Lu Han, Yu Liu, Lan Li, Qiwen Deng, Jian Jiang, Yinbo sun, Zhe Yu, Binfeng Wang, Xingyu Lu, Lintao Ma, Han-Jia Ye, and De-Chuan Zhan. Unica: Unified covariate adaptation for time series foundation model. In *International Conference on Learning Representations*, volume 2026, pp. 6403–6443, 2026.

Hyung-Chul Lee, Yoonsang Park, Soo Bin Yoon, Seong Mi Yang, Dongnyeok Park, and Chul-Woo Jung. Vitaldb, a high-fidelity multi-parameter vital signs database in surgical patients. *Scientific Data*, 9(1):279, 2022.

Zhengyu Li, Xiangfei Qiu, Yuhan Zhu, Xingjian Wu, Jilin Hu, Guo Guo, and Bin Yang. Gcgnet: Graph-consistent generative network for time series forecasting with exogenous variables. In *International Conference on Learning Representations*, volume 2026, pp. 27904–27926, 2026.

Yichao Liang, Nishanth Kumar, Hao Tang, Adrian Weller, Joshua B Tenenbaum, Tom Silver, Joao F. Henriques, and Kevin Ellis. Visualpredicator: Learning abstract world models with neuro-symbolic predicates for robot planning. In *International Conference on Learning Representations*, volume 2025, pp. 60416–60444, 2025.

Yichao Liang, Dat Nguyen, Cambridge Yang, Tianyang Li, Joshua B Tenenbaum, Carl Edward Rasmussen, Adrian Weller, Zenna Tavares, Tom Silver, and Kevin Ellis. Exopredicator: Learning abstract models of dynamic worlds for robot planning. In *International Conference on Learning Representations*, volume 2026, pp. 49174–49214, 2026.

Bryan Lim, Sercan Ö. Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International journal of forecasting*, 37(4): 1748–1764, 2021.

Weixin Liu, Juming Xiong, Congning Ni, Yanfan Zhu, Xingtao Lin, Bradley A. Malin, and Zhijun Yin. Drift: Direct-recursive intervention-conditioned forecasting of icu physiological trajectories. *arXiv preprint arXiv:2607.25864*, 2026.

Junjie Meng, Ranxu Zhang, Zi-an Zhang, Shujun Liu, Xiaoning Qi, Xiaozhou Xu, Yanyong Zhang, Hui Xiong, and Chao Wang. Cedar: Controlled and event-driven demand forecasting via residual decomposition. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 7797–7808, 2026.

Esmaeel Mohammadi, Anju Rani, Mikkel Stokholm-Bjerregaard, Daniel Ortiz-Arroyo, and Petar Durdevic. Wastewater treatment plant data for nutrient removal system. *arXiv preprint arXiv:2407.05346*, 2024.

Linjie Mu, Zhongzhen Huang, Yannian Gu, Shengqian Qin, Shaoting Zhang, and Xiaofan Zhang. Ehrworld: A patient-centric medical world model for long-horizon clinical trajectories. *arXiv preprint arXiv:2602.03569*, 2026.

Hada Melino Muhammad, Zechen Li, Flora Salim, and Ahmed A. Metwally. Cgm-jepa: Learning consistent continuous glucose monitor representations via predictive self-supervised pretraining. *arXiv preprint arXiv:2605.00933*, 2026.

Weizhi Nie, Weichao Liu, Honglin Guo, and Yuting Su. Phys-jepa: Physics-informed latent world models for multivariate time-series forecasting. *arXiv preprint arXiv:2606.16076*, 2026.

Anna Selini Petropoulou, Bart van Marrewijk, Feije de Zwart, Anne Elings, Monique Bijlaard, Tim van Daalen, Guido Jansen, and Silke Hemming. Lettuce production in intelligent greenhouses—3d imaging and computer vision for plant spacing decisions. *Sensors*, 23(6):2929, 2023.

Xiangfei Qiu, Yuhan Zhu, Zhengyu Li, Xingjian Wu, Bin Yang, and Jilin Hu. Dag: A dual correlation network for time series forecasting with exogenous variables. *arXiv preprint arXiv:2509.14933*, 2025.

Cyriana M.A. Roelofs, Edison Guevara Bastidas, Thomas Hugo, Stefan Faulstich, and Anna Cadenbach. Enabling predictive maintenance in district heating substations: A labelled dataset and fault detection evaluation framework based on service data. *Energy*, 355:141178, 2026. ISSN 0360-5442. doi: https://doi.org/10.1016/j.energy.2026.141178.

David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International journal of forecasting*, 36(3):1181–1191, 2020.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2009.

Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pp. 593–607. Springer, 2018.

Yifei Sun, Junming Liu, Yirong Chen, Xuefeng Yan, and Ding Wang. Timemkg: Knowledge-infused causal reasoning for multivariate time series modeling. *arXiv preprint arXiv:2508.09630*, 2025.

Yuxuan Wang, Haixu Wu, Jiaxiang Dong, Guo Qin, Haoran Zhang, Yong Liu, Yunzhong Qiu, Jianmin Wang, and Mingsheng Long. TimeXer: Empowering transformers for time series forecasting with exogenous variables. In *Advances in Neural Information Processing Systems*, volume 37, pp. 469–498, 2024.

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

Lan Wu, Xuebin Wang, Chenglong Ge, Ruijuan Chu, and LinYu Wang. Exotimer: Leveraging large language models for time series forecasting with exogenous variables. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 26956–26964, 2026.

Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. *arXiv preprint arXiv:1906.00121*, 2019.

Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 753–763, 2020.

Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. Causalvae: Disentangled representation learning via neural structural causal models. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9588–9597. IEEE, 2021.

Zhenghao Zhang, Yuanxiang Wang, Zhenyu Guan, Yujia Yang, Bingkang Shi, Tianyu Zong, Hongzhu Yi, Guoqing Chao, Xingchen Chen, Tiankun Yang, et al. Delta-jepa: Learning action-sensitive world models via latent difference decoding. *arXiv preprint arXiv:2606.31232*, 2026.

Qinpei Zhao, Jinhao Zhu, Xuan Shen, Chuwen Lin, Yinjia Zhang, Yuxiang Liang, Baige Cao, Jiangfeng Li, Xiang Liu, Weixiong Rao, and Congrong Wang. Chinese diabetes datasets for data-driven machine learning. *Scientific Data*, 10(1):35, 2023.

Gaoyue Zhou, Hengkai Pan, Yann LeCun, and Lerrel Pinto. Dino-wm: World models on pre-trained visual features enable zero-shot planning. *arXiv preprint arXiv:2411.04983*, 2024.

A DATASETS

The main evaluation uses four multi-target controlled systems (Table 1, Appendix A.1), in which several targets share typed driver-mediated mechanisms. We apply BELIEFGRAPH-JEPA to two single-target glucose protocols (Table 7, Appendix A.2), where the typed graph has one observed sink and fewer effect-to-target routes. Forecasting results for those protocols are in Appendix C. Per-target breakdowns of the four multi-target systems are in Appendix D. Train/val/test splits and remaining preprocessing are in Appendix B.

A.1 MULTI-TARGET DATASETS

VitalDB. Intraoperative vital signs from surgical cases (Lee et al., 2022). Targets are mean arterial pressure, heart rate, pulse oxygen saturation, and bispectral index. Continuous actions are propofol and remifentanil infusion rates. There is no exogenous channel. The same infusion can affect pressure and anesthetic depth with different lags, so the typed graph routes one driver family to several targets.

Greenhouse. The 3rd Autonomous Greenhouse Challenge (Petropoulou et al., 2023): lettuce compartments under closed-loop climate control. Targets are indoor air temperature, relative humidity, CO_2 concentration, and incoming PAR. Continuous actions are heating-pipe temperature, leeward and windward vents, thermal and shade screens, grow lights, and CO_2 dosing. Exogenous inputs are outdoor weather. Future outdoor weather holds the last context value over the horizon.

Wastewater nutrient control. A biological nutrient-removal tank with online sensors and dosing (Mohammadi et al., 2024). Targets are dissolved oxygen, ammonium, and phosphate. Continuous actions are metal-salt dosing channels; categorical actions are inlet and outlet process phases. Exogenous inputs are inflow rate and water temperature. Dosing and phase changes persist beyond a single step and affect the three nutrient targets through shared process mechanisms.

PreDist. District-heating substations (Roelofs et al., 2026). Targets include secondary supply temperature, return temperatures, heat power, and flow. Continuous actions are valve position and supply-temperature setpoint; categorical actions include pump status, control mode, and a three-way valve. Exogenous inputs are outdoor temperature and primary-grid supply temperature. Heat power and flow are algebraic: their instantaneous values are determined by contemporaneous temperatures and hydraulic drivers. This protocol mixes delayed inertial targets with same-horizon algebraic targets and uses an instantaneous-effect head (Appendix E).

Table 6 lists the sampling grid and forecast windows.

Table 6: Multi-target grids and forecast windows.

Dataset	Grid	Context, horizon
VitalDB	10 s	15 min, 6 min
Greenhouse	5 min	12 h, 3 h
Wastewater	2 min	6 h, 3 h
PreDist	10 min	12 h, 6 h

A.2 SINGLE-TARGET DATASETS

Table 7 summarizes the two single-target glucose protocols. Both use the same driver–graph–rollout–residual framework as the multi-target systems; there is no exogenous channel.

CGMacros. Continuous glucose monitoring with meal logs (Das et al., 2025). The sole target is interstitial glucose. Actions are sparse meal events with macronutrient annotations; there is no exogenous channel and no dense control input. The typed graph therefore has a single observed target and fewer effect-to-target routes than the multi-target systems above.

Shanghai Diabetes. Shanghai T1DM/T2DM CGM with meal, injection, and oral-drug logs (Zhao et al., 2023). The sole target is glucose. Actions are sparse clinical events; there is no exogenous channel. As in CGMacros, the graph topology is a single-target driver–effect–glucose routing structure.

Table 7: Single-target glucose protocols.

Dataset	Grid	Context, horizon	Actions
CGMacros	5 min	12 h, 2 h	sparse meal events
Shanghai Diabetes	15 min	12 h, 2 h	sparse meal, injection, and oral-drug events

B DATASET PREPROCESSING

Windows never cross session boundaries. A session break is a contiguous invalid-target gap (VitalDB, Greenhouse, Wastewater, PreDist) or a timestamp gap of more than 30 minutes (CGMacros, Shanghai Diabetes). Short sessions and windows whose target NaN ratio exceeds a dataset-specific threshold are dropped. Target and exogenous channels are min-max normalized to $[-1, 1]$ with fixed physical ranges. Continuous actions use channel-wise min-max. Categorical actions remain integer indices. Nutrition and dose features on the glucose datasets are z-scored from the training split only.

VitalDB. Native recordings are on a 2 s grid and are resampled to 10 s for training. Splits are by surgical case at 70/15/15 with seed 42: the same case never appears in more than one split.

Greenhouse. Splits are by team/compartment with seed 42: two teams are held out for test, one for validation, and the remainder for training. Compartments from the same team stay in the same split.

Wastewater nutrient control. The source is a single plant, so splits are chronological by calendar month at 70/15/15. Earlier months train, later months validate and test; there is no random subject shuffle.

PreDist. Splits are by substation at 70/15/15 with seed 42. Maintenance events are not treated as actions.

CGMacros. Native CGM is 1 min and is resampled to 5 min for training. Splits are by subject at 70/15/15 with seed 42, stratified by A1c phenotype (healthy, prediabetes, type 2 diabetes) so that each phenotype appears in every split when the cohort is large enough.

Shanghai Diabetes. Splits are by subject at 70/10/20 with seed 42, stratified by phenotype (T1DM, T2DM) and then merged. Adjacent records more than 30 minutes apart form distinct sessions.

C SINGLE-TARGET EVALUATION

Table 8 applies the HO–FR ladder and its GR variants from Table 2 to the two glucose protocols of Table 7. The typed graph has one observed target, so there are fewer effect-to-target routes than on the four multi-target systems. Metrics are physical-unit MAE and RMSE in mg/dL. Foundation-model rows are single-run; remaining rows are mean $\pm$ sample std over three seeds.

On CGMacros, future meal tokens produce the large step: DC lowers MAE from 14.72 (HO) to 13.24. Persistent effects and future rollout add smaller increments. FR-GR attains the lowest MAE and RMSE (12.76 and 19.95), essentially tied with TFT (12.79 and 20.81) and ahead of DAG (13.37) and TimesFM 2.5 (14.89). Relative to HO, that is a 13% MAE reduction.

On Shanghai Diabetes, actions are sparse typed clinical events (meal, injection, oral drug) and the graph has a single glucose sink. TFT attains the lowest MAE and RMSE (15.30 and 23.88), ahead of DAG (15.88). Within the shared ladder, RB-GR is lowest (16.24 MAE, 25.01 RMSE) and FR-GR is 16.52.

Table 9 evaluates the graph base z^G (α^R=0) and the full prediction $\hat{z} = z^G + \alpha^R \Delta z^R$ on the same checkpoints. Residual gain is $\mathrm{err}(z^G) - \mathrm{err}(\hat{z})$. The residual of RB-GR lowers MAE by 0.86. Future rollout improves the graph base by 0.32 (z^G MAE 17.10 to 16.78); the residual of FR-GR then lowers MAE by only 0.25. The two mechanisms are substitutes for future-driver information, and the current-belief residual is the stronger route. Future events occupy 40% of Shanghai test windows, so RB-GR already reads those tokens, while rolled Z^F is empty on the complementary windows.

Table 8: Test MAE / RMSE (↓) in mg/dL on the two single-target glucose datasets. Foundation-model rows are single-run; remaining rows are mean ± sample std over three seeds. The HO–FR ladder and its GR variants.

	CGMacros		Shanghai Diabetes	
Method	MAE	RMSE	MAE	RMSE
Persistence	17.25	28.11	18.71	29.55
Chronos-2	15.42	25.38	18.96	30.56
Chronos-2-Cov	15.91	26.31	18.75	30.37
TimesFM 2.5	14.89	24.59	17.25	27.72
DAG	$13.37_{\pm 0.12}$	$21.52_{\pm 0.21}$	$\underline{15.88}_{\pm 0.26}$	$25.05_{\pm 0.45}$
TFT	$\underline{12.79}_{\pm 0.21}$	$20.81_{\pm 0.39}$	$\mathbf{15.30}_{\pm 0.02}$	$\mathbf{23.88}_{\pm 0.13}$
HO	$14.72_{\pm 0.07}$	$23.42_{\pm 0.38}$	$17.08_{\pm 0.30}$	$25.80_{\pm 0.13}$
DC	$13.24_{\pm 0.15}$	$20.62_{\pm 0.15}$	$16.61_{\pm 0.14}$	$25.60_{\pm 0.25}$
SG	$13.20_{\pm 0.09}$	$20.69_{\pm 0.22}$	$16.69_{\pm 0.06}$	$25.49_{\pm 0.40}$
IE	$13.14_{\pm 0.04}$	$20.26_{\pm 0.13}$	$16.49_{\pm 0.05}$	$25.44_{\pm 0.12}$
RB	$12.92_{\pm 0.20}$	$20.09_{\pm 0.27}$	$16.51_{\pm 0.25}$	$25.57_{\pm 0.34}$
FR	$12.88_{\pm 0.17}$	$20.18_{\pm 0.25}$	$16.33_{\pm 0.31}$	$25.27_{\pm 0.15}$
RB-GR	$12.88_{\pm 0.03}$	$\underline{19.99}_{\pm 0.07}$	$16.24_{\pm 0.31}$	$\underline{25.01}_{\pm 0.39}$
FR-GR	$\mathbf{12.76}_{\pm 0.25}$	$\mathbf{19.95}_{\pm 0.16}$	$16.52_{\pm 0.19}$	$25.41_{\pm 0.13}$

Table 9: Shanghai residual gain. Graph-base error at α^R=0 versus the full prediction $\hat{z} = z^G + \alpha^R \Delta z^R$ on the same RB-GR and FR-GR checkpoints as Table 8. Residual gain $\Delta = \mathrm{err}(z^G) - \mathrm{err}(\hat{z})$ is computed per seed. Mean ± sample std over three seeds.

	MAE			RMSE		
Method	z^G	$\hat{z}$	Δ	z^G	$\hat{z}$	Δ
RB-GR	$17.10_{\pm 0.13}$	$16.24_{\pm 0.31}$	$0.86_{\pm 0.19}$	$25.81_{\pm 0.22}$	$25.01_{\pm 0.39}$	$0.81_{\pm 0.22}$
FR-GR	$16.78_{\pm 0.18}$	$16.52_{\pm 0.19}$	$0.25_{\pm 0.04}$	$25.62_{\pm 0.18}$	$25.41_{\pm 0.13}$	$0.21_{\pm 0.05}$

The same framework also applies to single-target action-conditioned forecasting. With one observed sink the gain over HO is smaller than in Table 2, and the best GR variant is not the same on both protocols: FR-GR on CGMacros, RB-GR on Shanghai. TFT is essentially tied with FR-GR on CGMacros and is strongest on Shanghai. The strongest benefits of the full framework are observed in the multi-target setting, where several targets share typed driver-mediated mechanisms. A plausible explanation is that multiple effect-to-target routes provide the graph pathway with target-specific routing structure that a single-sink graph does not provide.

D PER-TARGET EVALUATION

Table 2 pools every target on the $[-1, 1]$ scale. Tables 10–13 report physical-unit MAE and RMSE for each target, with the same methods and seeds. Foundation-model rows are single-run; remaining rows are mean ± sample std over seeds $\{42, 43, 44\}$.

Per-target errors reveal where the aggregate gains arise. On Greenhouse, FR-GR is lowest on indoor temperature, humidity, and CO_2. TFT is lowest on PAR MAE; UR-GR is lowest on PAR RMSE. On Wastewater, TFT is lowest on DO and PO_4 and on NH_4 MAE; FLAT-GR is lowest on NH_4 RMSE. On VitalDB, FR-GR is lowest on BIS, a graph-covered drug-sensitive target, and is comparable to FLAT-GR on MBP. TFT is lowest on HR MAE, while FR-GR is lowest on HR RMSE. SpO_2 has no drug-effect edge in the prior (Appendix L), and TimesFM 2.5 attains the lowest MAE. On PreDist, FR-GR is lower than FR, FLAT-GR, and UR-GR on net-return temperature and heat power, and lower than FR and FLAT-GR on flow. UR-GR is lowest on supply temperature. TFT is lowest on net-return temperature, and DAG is lower on the algebraic heat-power and flow channels. This distinction is reflected in the role-specific construction of Appendix E: temperatures evolve as inertial states, while heat power and flow are contemporaneous algebraic readouts (Table 14).

Table 10: Per-target test MAE / RMSE (↓) on Greenhouse, physical units. Same methods as Table 2.

	indoor T		RH		CO_2		PAR	
Method	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Persistence	1.489	2.273	5.51	8.09	81.5	151.5	73.8	143.2
Chronos-2	1.618	2.432	5.68	8.32	78.0	153.6	71.3	128.4
Chronos-2-Cov	1.450	2.290	5.18	7.89	76.7	151.7	67.7	125.0
TimesFM 2.5	1.537	2.317	5.57	8.08	77.0	148.6	77.1	138.1
DAG	$1.397_{\pm 0.026}$	$2.040_{\pm 0.035}$	$5.61_{\pm 0.08}$	$7.83_{\pm 0.08}$	$87.7_{\pm 0.5}$	$152.4_{\pm 3.1}$	$70.5_{\pm 1.8}$	$118.2_{\pm 2.4}$
TFT	$1.199_{\pm 0.046}$	$1.720_{\pm 0.066}$	$4.75_{\pm 0.17}$	$6.43_{\pm 0.22}$	$75.7_{\pm 4.0}$	$\underline{119.9}_{\pm 4.1}$	$\mathbf{60.3}_{\pm 1.5}$	$102.5_{\pm 3.8}$
HO	$1.350_{\pm 0.014}$	$1.881_{\pm 0.028}$	$5.52_{\pm 0.15}$	$7.41_{\pm 0.16}$	$100.2_{\pm 1.4}$	$151.1_{\pm 1.9}$	$72.9_{\pm 0.6}$	$108.0_{\pm 1.1}$
DC	$1.283_{\pm 0.017}$	$1.747_{\pm 0.021}$	$5.10_{\pm 0.09}$	$6.93_{\pm 0.07}$	$88.2_{\pm 1.9}$	$131.2_{\pm 1.9}$	$72.1_{\pm 2.5}$	$106.9_{\pm 2.5}$
SG	$1.238_{\pm 0.024}$	$1.713_{\pm 0.031}$	$5.14_{\pm 0.05}$	$6.94_{\pm 0.02}$	$84.6_{\pm 2.3}$	$128.9_{\pm 2.1}$	$71.8_{\pm 2.4}$	$107.4_{\pm 1.0}$
IE	$1.256_{\pm 0.031}$	$1.729_{\pm 0.037}$	$5.20_{\pm 0.08}$	$7.02_{\pm 0.07}$	$85.9_{\pm 1.8}$	$129.5_{\pm 2.7}$	$72.2_{\pm 3.3}$	$108.1_{\pm 3.7}$
RB	$1.265_{\pm 0.026}$	$1.744_{\pm 0.033}$	$5.17_{\pm 0.03}$	$6.95_{\pm 0.04}$	$84.2_{\pm 1.8}$	$127.9_{\pm 1.2}$	$70.2_{\pm 1.3}$	$106.9_{\pm 2.0}$
FR	$1.242_{\pm 0.021}$	$1.729_{\pm 0.013}$	$\underline{4.72}_{\pm 0.13}$	$\underline{6.38}_{\pm 0.16}$	$76.6_{\pm 1.1}$	$121.6_{\pm 1.0}$	$73.2_{\pm 2.6}$	$107.8_{\pm 3.0}$
FLAT	$1.194_{\pm 0.017}$	$1.669_{\pm 0.027}$	$4.89_{\pm 0.16}$	$6.59_{\pm 0.18}$	$82.4_{\pm 2.9}$	$127.2_{\pm 4.2}$	$68.0_{\pm 1.0}$	$103.2_{\pm 1.1}$
FLAT-GR	$1.158_{\pm 0.027}$	$1.620_{\pm 0.028}$	$4.85_{\pm 0.05}$	$6.58_{\pm 0.07}$	$74.7_{\pm 1.4}$	$\underline{119.9}_{\pm 4.3}$	$65.7_{\pm 1.4}$	$103.1_{\pm 1.8}$
UR-GR	$\underline{1.142}_{\pm 0.013}$	$\underline{1.605}_{\pm 0.027}$	$4.77_{\pm 0.04}$	$6.50_{\pm 0.06}$	$\underline{74.4}_{\pm 1.8}$	$120.6_{\pm 2.9}$	$65.6_{\pm 1.8}$	$\mathbf{101.3}_{\pm 2.3}$
RB-GR	$1.198_{\pm 0.014}$	$1.661_{\pm 0.016}$	$5.18_{\pm 0.13}$	$7.04_{\pm 0.14}$	$77.6_{\pm 1.4}$	$123.9_{\pm 0.2}$	$67.2_{\pm 1.9}$	$104.0_{\pm 1.9}$
FR-GR	$\mathbf{1.134}_{\pm 0.025}$	$\mathbf{1.588}_{\pm 0.047}$	$\mathbf{4.59}_{\pm 0.05}$	$\mathbf{6.29}_{\pm 0.08}$	$\mathbf{68.8}_{\pm 1.2}$	$\mathbf{112.8}_{\pm 2.6}$	$\underline{65.1}_{\pm 0.6}$	$\underline{102.2}_{\pm 1.7}$

Table 11: Per-target test MAE / RMSE (↓) on Wastewater, physical units. Same methods as Table 2.

	DO		NH_4		PO_4	
Method	MAE	RMSE	MAE	RMSE	MAE	RMSE
Persistence	0.813	1.223	0.768	1.105	0.549	0.825
Chronos-2	0.604	0.988	0.693	1.066	0.459	0.712
Chronos-2-Cov	0.484	0.828	0.474	0.785	0.338	0.572
TimesFM 2.5	0.567	0.920	0.664	0.995	0.439	0.668
DAG	$0.533_{\pm 0.011}$	$0.830_{\pm 0.009}$	$0.431_{\pm 0.011}$	$0.688_{\pm 0.012}$	$0.307_{\pm 0.006}$	$0.492_{\pm 0.008}$
TFT	$\mathbf{0.338}_{\pm 0.002}$	$\mathbf{0.574}_{\pm 0.007}$	$\mathbf{0.320}_{\pm 0.015}$	$0.574_{\pm 0.037}$	$\mathbf{0.222}_{\pm 0.002}$	$\mathbf{0.371}_{\pm 0.003}$
HO	$0.589_{\pm 0.011}$	$0.855_{\pm 0.012}$	$0.630_{\pm 0.008}$	$0.898_{\pm 0.014}$	$0.423_{\pm 0.004}$	$0.607_{\pm 0.001}$
DC	$0.403_{\pm 0.001}$	$0.648_{\pm 0.007}$	$0.425_{\pm 0.012}$	$0.651_{\pm 0.017}$	$0.296_{\pm 0.014}$	$0.468_{\pm 0.016}$
SG	$0.389_{\pm 0.008}$	$0.628_{\pm 0.007}$	$0.413_{\pm 0.017}$	$0.650_{\pm 0.028}$	$0.281_{\pm 0.007}$	$0.450_{\pm 0.005}$
IE	$0.395_{\pm 0.001}$	$0.636_{\pm 0.004}$	$0.415_{\pm 0.021}$	$0.659_{\pm 0.027}$	$0.285_{\pm 0.009}$	$0.451_{\pm 0.014}$
RB	$0.394_{\pm 0.005}$	$0.633_{\pm 0.006}$	$0.397_{\pm 0.005}$	$0.626_{\pm 0.005}$	$0.273_{\pm 0.004}$	$0.434_{\pm 0.003}$
FR	$0.363_{\pm 0.001}$	$0.598_{\pm 0.003}$	$0.351_{\pm 0.013}$	$0.573_{\pm 0.020}$	$0.244_{\pm 0.004}$	$0.396_{\pm 0.008}$
FLAT	$0.366_{\pm 0.007}$	$0.594_{\pm 0.007}$	$0.353_{\pm 0.024}$	$0.570_{\pm 0.028}$	$0.240_{\pm 0.006}$	$0.388_{\pm 0.009}$
FLAT-GR	$\underline{0.350}_{\pm 0.014}$	$0.584_{\pm 0.017}$	$0.330_{\pm 0.002}$	$\mathbf{0.529}_{\pm 0.011}$	$0.235_{\pm 0.002}$	$\underline{0.384}_{\pm 0.003}$
UR-GR	$0.358_{\pm 0.010}$	$0.594_{\pm 0.011}$	$0.345_{\pm 0.002}$	$0.559_{\pm 0.008}$	$0.245_{\pm 0.004}$	$0.393_{\pm 0.004}$
RB-GR	$0.363_{\pm 0.008}$	$0.600_{\pm 0.009}$	$0.356_{\pm 0.004}$	$0.567_{\pm 0.007}$	$0.254_{\pm 0.000}$	$0.409_{\pm 0.006}$
FR-GR	$0.353_{\pm 0.009}$	$\underline{0.582}_{\pm 0.012}$	$\underline{0.326}_{\pm 0.006}$	$\underline{0.537}_{\pm 0.017}$	$\underline{0.231}_{\pm 0.003}$	$\underline{0.384}_{\pm 0.005}$

E INSTANTANEOUS-EFFECT HEAD

The instantaneous-effect head is an extension for contemporaneous algebraic targets, used on PreDist.

An observation vector can mix two constitutive roles. Delayed inertial targets evolve as states with memory. Instantaneous algebraic targets are contemporaneous functions of those states and of the current drivers. A single JEPA latent z_{t+h} then packs both roles into one vector. The teacher, the predictor, and the decoder share one interface for a slow state and a same-horizon readout.

The instantaneous-effect head splits that interface by role $c \in \{\mathrm{dyn}, \mathrm{alg}\}$. Inertial channels keep family-effect memory through Eq. (2). Algebraic channels are evaluated on the same horizon and never enter $\mathcal{F}$. The construction follows the target property: delayed inertial channels together with instantaneous algebraic channels.

PreDist is the evaluation system with this mix. Supply and return temperatures store heat and lag valve, setpoint, and pump changes. Heat power and volume flow are algebraic functions of those

Table 12: Per-target test MAE / RMSE (↓) on VitalDB, physical units. Same methods as Table 2.

	MBP		HR		SpO_2		BIS	
Method	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Persistence	5.948	10.745	3.714	7.126	0.235	1.050	4.892	7.087
Chronos-2	6.063	10.455	3.470	6.416	0.229	1.025	4.693	6.995
Chronos-2-Cov	5.992	10.463	3.533	6.648	0.230	1.036	4.676	7.035
TimesFM 2.5	5.671	9.679	3.334	6.205	**0.214**	0.952	4.564	6.800
DAG	$5.608_{\pm 0.020}$	$9.335_{\pm 0.032}$	$3.526_{\pm 0.017}$	$6.318_{\pm 0.026}$	$0.280_{\pm 0.003}$	$1.050_{\pm 0.004}$	$4.638_{\pm 0.011}$	$6.864_{\pm 0.025}$
TFT	$5.117_{\pm 0.032}$	$8.685_{\pm 0.050}$	$\mathbf{3.294}_{\pm 0.040}$	$5.959_{\pm 0.038}$	$0.350_{\pm 0.109}$	$1.139_{\pm 0.186}$	$4.140_{\pm 0.017}$	$5.867_{\pm 0.034}$
HO	$5.401_{\pm 0.050}$	$8.867_{\pm 0.006}$	$3.401_{\pm 0.018}$	$5.885_{\pm 0.003}$	$0.282_{\pm 0.004}$	$0.896_{\pm 0.002}$	$4.441_{\pm 0.007}$	$6.442_{\pm 0.008}$
DC	$5.308_{\pm 0.062}$	$8.675_{\pm 0.014}$	$3.457_{\pm 0.056}$	$5.873_{\pm 0.027}$	$0.298_{\pm 0.006}$	$0.889_{\pm 0.004}$	$4.217_{\pm 0.007}$	$5.926_{\pm 0.009}$
SG	$5.335_{\pm 0.041}$	$8.696_{\pm 0.036}$	$3.416_{\pm 0.052}$	$5.873_{\pm 0.007}$	$0.341_{\pm 0.016}$	$0.899_{\pm 0.001}$	$4.232_{\pm 0.022}$	$5.944_{\pm 0.012}$
IE	$5.290_{\pm 0.018}$	$8.670_{\pm 0.022}$	$3.417_{\pm 0.052}$	$5.880_{\pm 0.013}$	$0.333_{\pm 0.017}$	$0.897_{\pm 0.003}$	$4.211_{\pm 0.009}$	$5.938_{\pm 0.010}$
RB	$5.276_{\pm 0.016}$	$8.657_{\pm 0.034}$	$3.393_{\pm 0.003}$	$5.875_{\pm 0.007}$	$0.348_{\pm 0.011}$	$0.901_{\pm 0.004}$	$4.122_{\pm 0.020}$	$5.809_{\pm 0.032}$
FR	$5.168_{\pm 0.027}$	$8.549_{\pm 0.017}$	$3.391_{\pm 0.023}$	$5.871_{\pm 0.016}$	$0.331_{\pm 0.017}$	$0.898_{\pm 0.001}$	$4.087_{\pm 0.013}$	$5.749_{\pm 0.019}$
FLAT	$5.176_{\pm 0.027}$	$8.501_{\pm 0.011}$	$3.391_{\pm 0.012}$	$5.881_{\pm 0.015}$	$0.337_{\pm 0.004}$	$0.895_{\pm 0.004}$	$4.086_{\pm 0.009}$	$5.723_{\pm 0.007}$
FLAT-GR	$\mathbf{5.097}_{\pm 0.027}$	$8.497_{\pm 0.011}$	$3.345_{\pm 0.003}$	$5.844_{\pm 0.015}$	$0.316_{\pm 0.010}$	$\mathbf{0.886}_{\pm 0.005}$	$4.089_{\pm 0.018}$	$5.735_{\pm 0.022}$
UR-GR	$5.128_{\pm 0.041}$	$8.495_{\pm 0.017}$	$3.350_{\pm 0.011}$	$5.850_{\pm 0.011}$	$0.317_{\pm 0.013}$	$0.890_{\pm 0.001}$	$4.103_{\pm 0.008}$	$5.744_{\pm 0.012}$
RB-GR	$5.148_{\pm 0.035}$	$8.542_{\pm 0.009}$	$3.350_{\pm 0.005}$	$5.843_{\pm 0.005}$	$0.320_{\pm 0.015}$	$0.888_{\pm 0.001}$	$4.084_{\pm 0.028}$	$5.732_{\pm 0.028}$
FR-GR	$5.103_{\pm 0.031}$	$\mathbf{8.487}_{\pm 0.006}$	$3.347_{\pm 0.019}$	$\mathbf{5.831}_{\pm 0.010}$	$0.319_{\pm 0.023}$	$0.888_{\pm 0.003}$	$\mathbf{4.071}_{\pm 0.014}$	$\mathbf{5.715}_{\pm 0.017}$

Table 13: Per-target test MAE / RMSE (↓) on PreDist†, physical units. Same methods as Table 2.

	Supply T		HC ret. T		Net ret. T		Heat power		Flow	
Method	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Persistence	2.641	4.164	2.470	3.732	2.450	3.958	21.85	37.40	0.617	1.064
Chronos-2	2.295	3.740	2.163	3.427	2.144	3.504	19.95	34.15	0.559	0.954
Chronos-2-Cov	1.772	3.061	2.168	3.521	2.117	3.630	18.83	35.93	0.524	0.984
TimesFM 2.5	2.182	3.591	2.024	3.205	2.014	3.313	18.84	32.18	0.527	0.904
DAG	$1.671_{\pm 0.022}$	$2.588_{\pm 0.058}$	$1.938_{\pm 0.025}$	$2.940_{\pm 0.045}$	$1.969_{\pm 0.030}$	$3.107_{\pm 0.045}$	$\mathbf{17.60}_{\pm 0.14}$	$\mathbf{30.30}_{\pm 0.28}$	$\mathbf{0.492}_{\pm 0.003}$	$\mathbf{0.840}_{\pm 0.006}$
TFT	$1.455_{\pm 0.089}$	$2.075_{\pm 0.103}$	$1.950_{\pm 0.113}$	$2.816_{\pm 0.117}$	$\mathbf{1.873}_{\pm 0.036}$	$\mathbf{2.852}_{\pm 0.044}$	$22.98_{\pm 2.83}$	$38.29_{\pm 4.60}$	$0.872_{\pm 0.243}$	$1.344_{\pm 0.305}$
HO	$2.191_{\pm 0.073}$	$3.271_{\pm 0.013}$	$2.088_{\pm 0.036}$	$3.118_{\pm 0.012}$	$2.277_{\pm 0.016}$	$3.360_{\pm 0.016}$	$22.08_{\pm 0.76}$	$34.79_{\pm 0.82}$	$0.676_{\pm 0.068}$	$1.083_{\pm 0.094}$
DC	$1.598_{\pm 0.041}$	$2.391_{\pm 0.042}$	$1.977_{\pm 0.032}$	$2.925_{\pm 0.027}$	$2.139_{\pm 0.083}$	$3.195_{\pm 0.088}$	$22.88_{\pm 1.71}$	$35.49_{\pm 2.43}$	$0.712_{\pm 0.028}$	$1.131_{\pm 0.055}$
SG	$1.595_{\pm 0.015}$	$2.371_{\pm 0.043}$	$1.965_{\pm 0.021}$	$2.937_{\pm 0.012}$	$2.054_{\pm 0.076}$	$3.124_{\pm 0.111}$	$22.55_{\pm 2.48}$	$35.66_{\pm 3.61}$	$0.713_{\pm 0.054}$	$1.128_{\pm 0.074}$
IE	$1.665_{\pm 0.045}$	$2.475_{\pm 0.050}$	$2.046_{\pm 0.008}$	$3.018_{\pm 0.021}$	$2.159_{\pm 0.054}$	$3.218_{\pm 0.027}$	$22.60_{\pm 1.51}$	$35.92_{\pm 1.96}$	$0.658_{\pm 0.062}$	$1.052_{\pm 0.083}$
RB	$1.530_{\pm 0.016}$	$2.338_{\pm 0.065}$	$1.967_{\pm 0.012}$	$2.932_{\pm 0.022}$	$2.022_{\pm 0.028}$	$3.085_{\pm 0.025}$	$22.15_{\pm 2.10}$	$35.30_{\pm 2.97}$	$0.643_{\pm 0.047}$	$1.033_{\pm 0.076}$
FR	$1.290_{\pm 0.040}$	$1.902_{\pm 0.029}$	$\mathbf{1.817}_{\pm 0.036}$	$2.602_{\pm 0.028}$	$1.980_{\pm 0.058}$	$2.921_{\pm 0.051}$	$23.28_{\pm 3.03}$	$36.68_{\pm 4.26}$	$0.689_{\pm 0.082}$	$1.074_{\pm 0.105}$
FLAT	$1.310_{\pm 0.047}$	$1.909_{\pm 0.010}$	$1.917_{\pm 0.102}$	$2.653_{\pm 0.071}$	$2.209_{\pm 0.209}$	$3.107_{\pm 0.162}$	$23.50_{\pm 1.05}$	$37.33_{\pm 1.62}$	$0.699_{\pm 0.112}$	$1.094_{\pm 0.163}$
FLAT-GR	$1.284_{\pm 0.022}$	$1.879_{\pm 0.027}$	$1.820_{\pm 0.066}$	$\mathbf{2.561}_{\pm 0.035}$	$2.050_{\pm 0.239}$	$2.936_{\pm 0.186}$	$20.04_{\pm 1.59}$	$32.28_{\pm 1.74}$	$0.583_{\pm 0.040}$	$0.944_{\pm 0.057}$
UR-GR	$\mathbf{1.252}_{\pm 0.008}$	$\mathbf{1.843}_{\pm 0.014}$	$1.871_{\pm 0.046}$	$2.622_{\pm 0.033}$	$2.036_{\pm 0.078}$	$2.949_{\pm 0.063}$	$19.60_{\pm 0.79}$	$31.99_{\pm 1.11}$	$0.561_{\pm 0.035}$	$0.923_{\pm 0.050}$
RB-GR	$1.496_{\pm 0.008}$	$2.168_{\pm 0.058}$	$2.040_{\pm 0.154}$	$2.960_{\pm 0.177}$	$1.967_{\pm 0.038}$	$3.031_{\pm 0.023}$	$20.26_{\pm 0.90}$	$32.53_{\pm 1.06}$	$0.592_{\pm 0.015}$	$0.965_{\pm 0.024}$
FR-GR	$1.290_{\pm 0.030}$	$1.905_{\pm 0.024}$	$1.834_{\pm 0.006}$	$2.616_{\pm 0.008}$	$1.928_{\pm 0.069}$	$2.871_{\pm 0.042}$	$18.85_{\pm 0.58}$	$30.88_{\pm 0.52}$	$0.562_{\pm 0.021}$	$0.926_{\pm 0.028}$

temperatures and of the hydraulic drivers at the same 10 min step. The typed prior follows the same split (Appendix L): G_{dyn} routes driver→effect→temperature; G_{alg} routes valve/pump/three-way→flow and flow+ΔT →heat power on the same horizon.

Independent observation head. The context encoder remains joint over the history window. The head indexes the EMA teacher, the predictor queries, and the decoder by role. Each role keeps its own JEPA latent, and each decoder head reconstructs only that role's channels. A shared-latent baseline that packs both roles into one z_{t+h}, treats algebraic channels as additional inertial states, and decodes every channel from that vector is the comparison in Table 14.

Family persistence λ_f. In Eq. (2), $\lambda_f = \sigma(\rho_f)$ is the learned retention of a *latent-effect family*. It says how long a driver injection remains in $z_{f,\tau}$. That persistence is the right object for delayed influence: a valve or setpoint command continues to move inertial temperatures after it is issued. Instantaneous algebraic targets have no such reservoir. Their value at $t + h$ is fixed once contemporaneous inertial states and drivers are known.

The IE ablation removes the $\lambda_f z_{f,\tau}$ term, so an effect node does not carry driver history. An algebraic observation is not an effect node with $\lambda_f = 0$. It is a constraint on the current inertial state and drivers, and it never enters Z^F. Family persistence governs how long those drivers remain available to the inertial targets. The instantaneous-effect head governs how the two kinds of *observation* are represented in JEPA space.

Construction. Let I_{dyn} index the inertial channels and I_{alg} the algebraic channels. On PreDist these are the three temperatures and the pair (heat power, flow). The context encoder E_θ remains joint over the full history. The EMA teacher is applied once per role to the same future sequence with a

causal mask and complementary channels zeroed:

$$z^{*,c}_{t+1:t+H} = \mathrm{sg}\big(E_{\bar{\theta}}(X^{c}_{t+1:t+H})\big), \qquad c \in \{\mathrm{dyn}, \mathrm{alg}\}, \tag{7}$$

where $X^{c}_{t+1:t+H}$ equals $X_{t+1:t+H}$ on I_c and is zero elsewhere, and $z^{*,c}_{t+h}$ depends only on $X^{c}_{t+1:t+h}$. A shared predictor, and under GR a shared residual adapter, emit $\hat{z}^{c}_{t+h}$ from a role query r_c. Role identity is that query embedding; driver tokens are shared across roles. When a dual-timescale prior is present, dyn queries attend the G_{dyn} readout with algebraic nodes zeroed, and alg queries attend the merged G_{alg} readout. At horizon h, G_{dyn} is evaluated from the rolled family states Z^{F}_{t+h}. G_{alg} then overwrites the action and exogenous nodes with the encodings of D_{t+h-1} and message-passes on algebraic edges. The heating-circuit return parent of heat power uses that G_{dyn} embedding as its source: static identity and type, updated by effect-to-temperature messages from those rolled states (Appendix F). Decoder heads map $\hat{z}^{\mathrm{dyn}}_{t+h}$ onto I_{dyn} and $\hat{z}^{\mathrm{alg}}_{t+h}$ onto I_{alg}. The JEPA term averages the two role residuals. The same driver–graph–rollout–residual stack of Section 4 is used throughout. HO and DC keep the role split without a GNN, reading roles from the channel inventory.

Table 14 compares that head to a shared-latent FR-GR that packs all five PreDist targets into one z_{t+h}. Removing the head raises pooled nMAE from 0.0495 to 0.0582, above DAG (0.0512) and TimesFM 2.5 (0.0563) in Table 2. The increase is concentrated on the algebraic channels: heat-power MAE rises from 18.85 to 28.52 and flow from 0.562 to 0.730. The three inertial temperatures remain close; heating-circuit return is unchanged. The current shared-latent packing is therefore a worse interface for this mix. The head keeps delayed temperatures and contemporaneous readouts in separate JEPA spaces, and that split recovers the algebraic maps.

Table 14: PreDist test error of FR-GR with a shared JEPA latent versus with the instantaneous-effect head. Pooled nMAE on $[-1, 1]$; per-target MAE in physical units. Mean $\pm$ sample std over seeds $\{42, 43, 44\}$. The headed row is the PreDist FR-GR entry of Tables 2 and 13.

	nMAE	Supply T	HC ret. T	Net ret. T	Heat	Flow
FR-GR, shared z	$0.0582_{\pm 0.0020}$	$1.447_{\pm 0.078}$	$1.823_{\pm 0.073}$	$2.104_{\pm 0.138}$	$28.52_{\pm 3.13}$	$0.730_{\pm 0.049}$
FR-GR	$\mathbf{0.0495}_{\pm 0.0003}$	$\mathbf{1.290}_{\pm 0.030}$	$1.834_{\pm 0.006}$	$\mathbf{1.928}_{\pm 0.069}$	$\mathbf{18.85}_{\pm 0.58}$	$\mathbf{0.562}_{\pm 0.021}$

F Architecture details

This appendix records the computational interfaces of Sections 3–4.

EMA teacher. $E_{\bar{\theta}}$ is a momentum-updated copy of E_θ with the same architecture. After each optimizer step, every parameter is updated by EMA with decay 0.996. E_θ and $E_{\bar{\theta}}$ use the same learned positional embeddings, evaluated at indices $0, \ldots, L-1$ of their respective input sequences. During training, $E_{\bar{\theta}}$ encodes the ground-truth future sequence $X_{t+1:t+H}$ with a causal mask. Its h-th output token defines the stop-gradient target z^{*}_{t+h} and depends only on $X_{t+1:t+h}$. On PreDist, the teacher is applied once per role to the same future sequence with complementary channels zeroed (Appendix E).

Non-effect node features. Every node is initialized as $e^{\mathrm{static}}_i = e^{\mathrm{id}}_i + e^{\mathrm{type}}_i$. Effect nodes then receive the gated injection of $z_{f,\tau}$ (Eq. 3). Exogenous nodes are source-only in the prior, so reverse edges into them are omitted; once written, their input encodings are restored after each GNN layer. Target nodes remain at their static embeddings until message passing. The GNN is re-evaluated separately at each horizon using $\{Z^{F}_{t+h}\}_{h=1}^{H}$ and $\{D_{t+h-1}\}_{h=1}^{H}$.

On a single-timescale prior, action and exogenous nodes are overwritten with their corresponding encodings from D_{t+h-1}, the last driver of the rollout prefix, before message passing. Horizon-h node features therefore comprise static identity and type, rolled family states Z^{F}_{t+h}, and those driver encodings. On a dual-timescale prior, G_{dyn} is first evaluated from Z^{F}_{t+h} without this horizon-driver overwrite. G_{alg} then overwrites only action and exogenous nodes with their corresponding encodings from D_{t+h-1} and message-passes on algebraic edges. On PreDist, the heating-circuit return parent of heat power is read from this G_{dyn} embedding.

GNN edge attributes. Relation labels become discrete edge types with a learned matrix W_r. Confidence initializes a learnable scalar gate $g = \sigma(\ell)$, with $\ell = \mathrm{logit}(\mathrm{clip}(c, \varepsilon, 1-\varepsilon))$. Messages are $g\, W_r h_{\mathrm{src}}$. Sign is stored as metadata in the prior.

Predictor attention. Query h conditions on H^G_{t+h} by cross-attention to that horizon's graph nodes. Cross-attention to encoded driver tokens is prefix-masked: query h attends $D_{t:t+h-1}$. Under GR, those tokens enter only through A_R, with the same prefix mask. Horizon queries then share decoder self-attention, with no causal mask among queries. The recurrence is sequential; the decoder is joint over the known-action horizon.

Residual adapter. A_R is a Transformer decoder over encoded tokens of $D_{t:t+H-1}$ and history H^X_t. Cross-attention uses the same prefix mask; query self-attention is unmasked. Its last linear layer is zero-initialized.

G TRAINING CONFIGURATION

Tables 15–16 record the FR-GR recipe used in Table 2. Wastewater uses 50 epochs rather than 30 because validation nMAE continues to decrease through the longer schedule; the remaining datasets use 30 epochs. The latent-effect window equals the context length L_X (Table 6). VitalDB uses a matched-window control: windows are indexed as if $L_A = 180$, then past drivers are truncated to $L_X = 90$. Per-family timescales τ_f initialize λ_f in steps: VitalDB propofol 24, remifentanil 12; Greenhouse thermal 24, ventilation 6, light 2, screen 4, CO_2 6; Wastewater dosing 8, aeration 6, phase 15; PreDist valve 6, setpoint 18, pump 1. A shared $\tau = 8$ is compared in Appendix J. The glucose protocols in Appendix C use the same FR-GR backbone with batch size 16 and learning rate 3×10^{-4}; Shanghai uses a typed event encoder and CGMacros a meal encoder.

Table 15: Shared FR-GR architecture and optimization. PreDist uses the large-batch row of Table 16; Wastewater uses the 50-epoch row.

Setting	Value
Width d / heads / FFN	128 / 4 / $4d$
History encoder	2-layer causal Transformer, Pre-LN, GELU
GNN / predictor / residual adapter	2 relation-aware layers; 2-layer Transformer decoder
Dropout / positional encoding	0.1; learned embeddings on the state encoder; sinusoidal Δt on driver events
EMA target decay	0.996
Optimizer	AdamW, weight decay 10^{-4}, gradient clip 1.0
Schedule	cosine, 30 epochs (50 on Wastewater)
Default learning rate	3×10^{-4}, min 10^{-5}, no warmup
Loss weights	$\lambda_G = 0.1$, $\lambda_{\mathrm{cap}} = 0.01$, $\lambda_{\mathrm{obs}} = 1$
Residual scale	$\alpha^R_{\max} = 0.5$, init 0.1; graph-gate logit -4
Seeds / hardware	$\{42, 43, 44\}$; one NVIDIA A100

Table 16: Dataset-specific FR-GR training settings. Train windows are the indexed training set. Wall-clock is one seed, train and validation. Wastewater uses 50 epochs, batch 32, learning rate 1.2×10^{-3}, min LR 4×10^{-5}, and a 3-epoch warmup; remaining datasets use 30 epochs. PreDist uses a larger batch, learning rate 10^{-3}, and min LR 3×10^{-5}.

Dataset	Batch	LR	LR warm.	Stride	Train N	Params	Time
VitalDB	8	3×10^{-4}	0	30	45,425	2.0M	4.0 h
Greenhouse	8	3×10^{-4}	0	36	2,144	2.9M	0.4 h
Wastewater	32	1.2×10^{-3}	3	90	4,093	2.9M	0.4 h
PreDist	64	1×10^{-3}	1	36	87,475	3.1M	2.1 h

Parameter counts stay near 2–3M. One run finishes in 0.4–4 hours on a single A100, with VitalDB taking the longest. Latent recurrence is sequential in L_X+H and linear in $|V_F|$. The H GNN evaluations are independent given the rolled states and horizon-specific driver encodings, and are batched over the horizon. The history encoder and predictor dominate wall-clock cost.

H FUTURE-ACTION PERTURBATION PROTOCOL

Section 5.5 offsets continuous future-action channels on the second half of the horizon, $I = [H/2, H)$, and reports the relative nMAE increase Δ_{perturb} on the full window. The offset is applied in the normalized action space used at training. Categorical actions and exogenous inputs are held fixed. Continuous channels admit a shared additive offset of controlled magnitude. Categorical substitutions have no common scale: a label change can be nearly inert or can switch the operating mode. The diagnostic therefore tests whether forecasts use later continuous-action inputs.

Table 17: Second-half continuous-action offset. δ is the per-channel magnitude; values are clipped to the listed range.

Dataset	H	Interval I	Range	δ
VitalDB	36 (6 min)	$[18, 36)$ (3 min)	$[0, 1]$	0.15
Greenhouse	36 (3 h)	$[18, 36)$ (90 min)	$[0, 1]$	0.15
Wastewater	90 (3 h)	$[45, 90)$ (90 min)	$[-1, 1]$	0.30
PreDist	36 (6 h)	$[18, 36)$ (3 h)	$[-1, 1]$	0.30

Per-channel signs $s_j \in \{+1, -1\}$ are drawn once and shared across methods. If more than 10% of steps in I clip for a sample-channel, that sample-channel sign is flipped. Clipped values remain in the training range.

I RESIDUAL SCALE AND RELIANCE

Table 18 reports the converged residual scale α^R and realized residual reliance ρ_R of FR-GR. Across datasets, ρ_R is 1.4–3.9%: the residual has small relative latent magnitude compared with the graph base. Removing it at evaluation (FR-GR-GO in Table 3) nevertheless raises nMAE and nRMSE on every dataset; the larger degradation on Greenhouse and Wastewater coincides with their larger ρ_R. A small residual is sufficient to recover information omitted by the graph abstraction. Sensitivity of the same quantities to residual-scale initialization and graph-base supervision is in Appendix J.

Table 18: Converged residual scale α^R and reliance ρ_R (Eq. 6) of FR-GR.

Dataset	α^R	ρ_R
Greenhouse	$0.107_{\pm 0.001}$	$0.0302_{\pm 0.0019}$
Wastewater	$0.104_{\pm 0.001}$	$0.0391_{\pm 0.0032}$
VitalDB	$0.050_{\pm 0.002}$	$0.0142_{\pm 0.0009}$
PreDist	$0.069_{\pm 0.001}$	$0.0143_{\pm 0.0016}$

J HYPERPARAMETER SENSITIVITY

We evaluate the sensitivity of FR-GR to four design choices across the four multi-target protocols: residual-scale initialization, graph-base supervision weight, latent-effect history length, and family-timescale initialization. All remaining settings follow Appendix G.

Table 19 tests sensitivity to the residual-scale initialization and graph-base supervision weight. Increasing the initial residual scale from $\alpha^R = 0.1$ to 0.3, or reducing λ_G tenfold from 0.1 to 0.01, leaves test nMAE and nRMSE essentially unchanged across all four protocols.

Table 20 reports the same residual-scale and reliance quantities as Table 18 under these perturbations. A larger initialization produces a substantially larger converged α^R, but the realized reliance ρ_R remains nearly unchanged. Reducing λ_G likewise has little effect on either forecasting error or residual reliance. Across all settings, ρ_R remains in the narrow range 0.014–0.040.

These results distinguish the learned scale parameter from realized residual magnitude: different values of α^R can produce similar ρ_R. The small and stable ρ_R shows that the residual remains small in latent magnitude relative to the graph predictive base.

Table 19: Sensitivity to residual-scale initialization and graph-base supervision. FR-GR uses α^R init 0.1 and $\lambda_G = 0.1$.

	Greenhouse		Wastewater		VitalDB		PreDist†	
Setting	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE
FR-GR	$0.1285_{\pm 0.0016}$	$0.1928_{\pm 0.0033}$	$0.1245_{\pm 0.0018}$	$0.2157_{\pm 0.0028}$	$0.0580_{\pm 0.0003}$	$0.1032_{\pm 0.0001}$	$0.0495_{\pm 0.0003}$	$0.0768_{\pm 0.0004}$
α^R init 0.3	$0.1296_{\pm 0.0014}$	$0.1930_{\pm 0.0017}$	$0.1246_{\pm 0.0019}$	$0.2158_{\pm 0.0027}$	$0.0579_{\pm 0.0003}$	$0.1032_{\pm 0.0000}$	$0.0502_{\pm 0.0009}$	$0.0773_{\pm 0.0009}$
$\lambda_G = 0.01$	$0.1272_{\pm 0.0004}$	$0.1915_{\pm 0.0009}$	$0.1242_{\pm 0.0004}$	$0.2151_{\pm 0.0008}$	$0.0580_{\pm 0.0003}$	$0.1032_{\pm 0.0001}$	$0.0497_{\pm 0.0009}$	$0.0768_{\pm 0.0010}$

Table 20: Converged residual scale α^R and reliance ρ_R (Eq. 6).

	Greenhouse		Wastewater		VitalDB		PreDist†	
Setting	α^R	ρ_R	α^R	ρ_R	α^R	ρ_R	α^R	ρ_R
FR-GR	$0.107_{\pm 0.001}$	$0.0302_{\pm 0.0019}$	$0.104_{\pm 0.001}$	$0.0391_{\pm 0.0032}$	$0.050_{\pm 0.002}$	$0.0142_{\pm 0.0009}$	$0.069_{\pm 0.001}$	$0.0143_{\pm 0.0016}$
α^R init 0.3	$0.305_{\pm 0.002}$	$0.0274_{\pm 0.0030}$	$0.287_{\pm 0.002}$	$0.0383_{\pm 0.0012}$	$0.084_{\pm 0.005}$	$0.0141_{\pm 0.0010}$	$0.138_{\pm 0.002}$	$0.0156_{\pm 0.0007}$
$\lambda_G = 0.01$	$0.108_{\pm 0.001}$	$0.0306_{\pm 0.0029}$	$0.105_{\pm 0.001}$	$0.0401_{\pm 0.0029}$	$0.050_{\pm 0.002}$	$0.0148_{\pm 0.0006}$	$0.072_{\pm 0.003}$	$0.0156_{\pm 0.0011}$

Table 21 tests whether initializing the current latent-effect belief requires driver history preceding the observation context. We set $L_A = 2L_X$, adding one full context length of past drivers before the observed history.

Test nMAE and nRMSE remain essentially unchanged on all four protocols. Additional pre-context driver history therefore provides no measurable forecasting benefit. The main model consequently uses context-aligned belief initialization, $L_A = L_X$, rather than a separate latent-dynamics warm-up window.

Table 21: Sensitivity to the latent-effect history window. FR-GR uses context-aligned belief initialization with $L_A = L_X$; the ablation adds one additional context length of pre-context driver history.

	Greenhouse		Wastewater		VitalDB		PreDist†	
Setting	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE
FR-GR ($L_A = L_X$)	$0.1285_{\pm 0.0016}$	$0.1928_{\pm 0.0033}$	$0.1245_{\pm 0.0018}$	$0.2157_{\pm 0.0028}$	$0.0580_{\pm 0.0003}$	$0.1032_{\pm 0.0001}$	$0.0495_{\pm 0.0003}$	$0.0768_{\pm 0.0004}$
$L_A = 2L_X$	$0.1282_{\pm 0.0023}$	$0.1918_{\pm 0.0014}$	$0.1254_{\pm 0.0013}$	$0.2173_{\pm 0.0017}$	$0.0580_{\pm 0.0003}$	$0.1032_{\pm 0.0001}$	$0.0500_{\pm 0.0010}$	$0.0771_{\pm 0.0007}$

Table 22 tests whether the reported forecasting gains depend on the hand-specified family timescales of Appendix G. We replace all family-specific initializations with a shared $\tau = 8$, corresponding to $\lambda_f = e^{-1/8}$. Test nMAE and nRMSE remain essentially unchanged from Table 2 across all four protocols.

To examine whether training recovers the same persistence parameters from the two initializations, we convert the learned λ_f back to an implied timescale,

$$\hat{\tau}_f = -1/\log \lambda_f,$$

the inverse of the initialization $\lambda_f = e^{-1/\tau_f}$. Table 23 reports these values for every latent-effect family.

The two initializations do not in general converge to a common set of timescales. Several families retain a strong dependence on their initialization, although some move substantially and others converge to similar values. Despite these different persistence solutions, forecast accuracy is nearly unchanged. Thus, on the real forecasting protocols, the supervised forecasting objective does not uniquely identify the family persistence parameters: multiple latent dynamics can support essentially the same predictive solution. The family-specific τ_f should therefore be interpreted as an inductive bias rather than as a recovered physical timescale.

Section 5.4 and Appendix K test the same future-rollout mechanism on a matched leaky-integrator system. Starting from a shared initialization, the model recovers the correct persistence ordering $\lambda_{\text{slow}} > \lambda_{\text{fast}}$ and exhibits pathway-specific impulse responses. This provides a controlled capacity check under matched dynamics. On the four real forecasting protocols, different persistence configurations yield nearly identical forecast errors, so predictive accuracy alone does not establish identification of physical timescales.

Accordingly, BELIEFGRAPH-JEPA does not rely on accurately specifying the physical family timescales to obtain its reported forecasting performance: replacing the family-specific initialization by the shared $\tau = 8$ initialization yields essentially the same error.

Table 22: Robustness to family-timescale initialization. FR-GR uses the family-specific τ_f values of Appendix G; the ablation initializes every family with the same $\tau = 8$. Forecast accuracy is essentially unchanged.

	Greenhouse		Wastewater		VitalDB		PreDist†	
Setting	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE	nMAE	nRMSE
FR-GR	$0.1285_{\pm 0.0016}$	$0.1928_{\pm 0.0033}$	$0.1245_{\pm 0.0018}$	$0.2157_{\pm 0.0028}$	$0.0580_{\pm 0.0003}$	$0.1032_{\pm 0.0001}$	$0.0495_{\pm 0.0003}$	$0.0768_{\pm 0.0004}$
shared $\tau = 8$	$0.1279_{\pm 0.0019}$	$0.1918_{\pm 0.0038}$	$0.1244_{\pm 0.0011}$	$0.2150_{\pm 0.0021}$	$0.0580_{\pm 0.0003}$	$0.1032_{\pm 0.0001}$	$0.0497_{\pm 0.0016}$	$0.0773_{\pm 0.0021}$

Table 23: Learned implied family timescales $\hat{\tau}_f = -1/\log \lambda_f$ under family-specific and shared initialization. The two initializations do not generally converge to the same persistence solution despite nearly identical forecasting accuracy in Table 22.

Dataset	Family	Prior τ_f	$\hat{\tau}$ (FR-GR)	$\hat{\tau}$ ($\tau = 8$)
Greenhouse	thermal	24	$21.4_{\pm 0.6}$	$7.0_{\pm 0.1}$
	ventilation	6	$5.7_{\pm 0.2}$	$7.4_{\pm 0.5}$
	light	2	$2.0_{\pm 0.0}$	$8.5_{\pm 0.1}$
	screen	4	$4.1_{\pm 0.1}$	$8.1_{\pm 0.1}$
	CO_2	6	$6.3_{\pm 0.2}$	$8.2_{\pm 0.3}$
Wastewater	dosing	8	$9.2_{\pm 0.5}$	$9.0_{\pm 1.0}$
	aeration	6	$4.3_{\pm 0.2}$	$7.3_{\pm 2.8}$
	phase	15	$11.3_{\pm 0.3}$	$6.5_{\pm 0.3}$
VitalDB	propofol	24	$19.7_{\pm 1.3}$	$11.8_{\pm 0.3}$
	remifentanil	12	$10.6_{\pm 1.1}$	$9.1_{\pm 1.1}$
PreDist†	valve	6	$3.8_{\pm 0.8}$	$4.4_{\pm 0.5}$
	setpoint	18	$6.1_{\pm 0.9}$	$3.8_{\pm 0.8}$
	pump	1	$1.2_{\pm 0.0}$	$9.1_{\pm 0.5}$

Taken together, these ablations show that the reported forecasting gains are not tied to a narrow optimization recipe. Residual reliance remains small under different residual initialization and graph-base supervision, additional pre-context driver history is unnecessary, and forecast accuracy is robust to replacing family-specific persistence initialization with a shared timescale.

K SYNTHETIC LATENT-DYNAMICS DIAGNOSTICS

Section 5.4, Figure 3, and Table 5 report a class-matched sanity check on a known instance of Eq. (2). This appendix specifies the leaky-integrator system, the probe protocol, and the full numerical tables, then examines a saturating $\tanh$ transition mismatch (Appendix K.4). The probes are the learned family persistences λ_f and the future-impulse map from each driver to each target. Results use seed 42.

K.1 MATCHED-DYNAMICS SYSTEM

Three continuous drivers u_{fast}, u_{slow}, and u_{null} update three latent-effect families. Two observed targets read a sparse gain:

$$z_{f,t+1} = \lambda_f \, z_{f,t} + (1 - \lambda_f) \, \alpha_f \, u_{\text{parent}(f),t}, \qquad y_t = W z_t + \varepsilon_t. \tag{8}$$

Equation (8) is the continuous term of Eq. (2) with scalar (λ_f, α_f) and delay d=0. The ground-truth parameters are $\lambda_{\text{fast}} = 0.40$, $\lambda_{\text{slow}} = 0.75$, $\alpha_{\text{fast}} = \alpha_{\text{slow}} = 1$, and $\lambda_{\text{null}} = \alpha_{\text{null}} = 0$. W maps the fast effect onto y_{fast} and the slow effect onto y_{slow}. The null family has no target edge. Observation noise is $\varepsilon_t \sim \mathcal{N}(0, 0.015^2)$. Drivers are a weak AR background with sparse pulses.

The instantaneous step gains are therefore $(1-\lambda)\alpha = 0.60$ on the fast channel and 0.25 on the slow channel. Both step gains are large relative to the observation-noise scale, so neither pathway is hidden by noise; the two mechanisms are instead distinguished primarily by their persistence. A unit pulse on u_{fast} peaks at $y_{\text{fast}} = 0.60$ and leaves $y_{\text{slow}} = 0$; a unit pulse on u_{slow} peaks at $y_{\text{slow}} = 0.25$ and leaves $y_{\text{fast}} = 0$; a pulse on u_{null} moves neither target. The typed prior supplied to the model is this ground-truth graph.

K.2 Identification protocol

Windows use context $L_X = 8$, horizon $H = 32$, and a 64-step latent warmup. All but the last context target are zeroed. Superscript id marks this identification setting: graph variants use the hard-bottleneck fusion of Section 5.2 (graph-covered driver tokens are withheld from the direct path), family states enter at $\gamma_f \equiv 1$, and each λ_f is a scalar initialized at 0.575, the midpoint of $(0.40, 0.75)$. Correct ranking therefore requires the two persistences to move in opposite directions. FR^{id} (HB) is future latent-effect rollout under that fusion.

The ladder is DC^{id}, future-driver tokens without a graph; IE^{id}, an instant graph with $\lambda_f \equiv 0$; RB^{id} and RB-GR^{id}, a current recurrent belief without and with graph-first residual routing; FR^{id} (HB) and FR-GR^{id}, future latent-effect rollout without and with that routing. Architecture and optimization follow Appendix G (width 128, batch 16, learning rate 3×10^{-4}, 40 epochs). Results are seed 42. The training split has 96 sessions of length 384.

After training, two probes are read from the checkpoint. The first is the learned $(\lambda_{\text{fast}}, \lambda_{\text{slow}})$ and whether $\lambda_{\text{slow}} > \lambda_{\text{fast}}$. The second is a unit future pulse on one driver with a zero context: the reported impulse $I_{d\to y}$ is the maximum absolute response of target y over the horizon. The ground-truth on-path step gains are 0.60 and 0.25; off-path entries are the crosstalk and the null driver.

K.3 Results

Figure 3 and Table 5 summarize learned persistences and future-pulse responses for the current-belief and rollout variants. Tables 24–25 report the full identification ladder.

Table 24: Learned persistence and test MAE on the matched-dynamics system (seed 42). Rank is $\lambda_{\text{slow}} > \lambda_{\text{fast}}$. IE^{id} fixes $\lambda_f = 0$. DC^{id} has no family persistence. FR^{id} (HB) is the hard bottleneck of Section 5.2. MAE is on the physical target scale.

Method	λ_{fast}	λ_{slow}	Rank	MAE	MAE$_{\text{fast}}$	MAE$_{\text{slow}}$
Truth	0.40	0.75	—	—	—	—
Init	0.575	0.575	—	—	—	—
DC^{id}	—	—	—	0.063	0.077	0.049
IE^{id}	0	0	No	0.063	0.077	0.049
RB^{id}	0.619	0.627	No	0.063	0.077	0.049
RB-GR^{id}	0.629	0.637	No	0.063	0.077	0.049
FR^{id} (HB)	0.449	0.641	Yes	0.031	0.013	0.049
FR-GR^{id}	0.429	0.679	Yes	0.013	0.013	0.014

Table 25: Peak absolute response to a unit future pulse (horizon 32, zero context). Truth peaks are 0.60 ($u_{\text{fast}}\to y_{\text{fast}}$) and 0.25 ($u_{\text{slow}}\to y_{\text{slow}}$).

Method	$I_{\text{fast}\to\text{fast}}$	$I_{\text{fast}\to\text{slow}}$	$I_{\text{slow}\to\text{slow}}$	I_{null}
Truth	0.600	0	0.250	0
DC^{id}	0.000	0.000	0.000	0.000
IE^{id}	0.000	0.000	0.000	0.000
RB^{id}	0.000	0.000	0.000	0.000
RB-GR^{id}	0.000	0.000	0.000	0.000
FR^{id} (HB)	0.371	0.010	0.000	0.001
FR-GR^{id}	0.430	0.008	0.141	0.002

The no-rollout variants fail to recover horizon-dependent responses to future driver pulses. DC^{id}, IE^{id}, RB^{id}, and RB-GR^{id} all record MAE 0.063, with y_{fast} MAE 0.077 and y_{slow} MAE 0.049, and all four produce zero future-pulse responses. IE^{id} fixes $\lambda_f \equiv 0$, and the two current-belief models leave the persistences nearly equal ($\lambda \approx 0.62$–0.64) rather than separating the fast and slow families. DC^{id} receives future-driver tokens directly, but likewise produces zero pulse responses. Thus, under this identification protocol, direct future conditioning alone does not recover the latent future driver-response dynamics.

The current graph belief H_t^G is computed at the forecast anchor and remains fixed over the prediction horizon. A driver pulse introduced at lead time h therefore never updates Z^F in the current-belief variants, so λ_f receives no gradient that identifies how later driver effects should persist across future horizons. This is why RB^{id} and RB-GR^{id} match on every probe.

Future rollout changes this identification path. Applying $Z^F_{t+h} = \mathcal{F}(Z^F_{t+h-1}, D_{t+h-1})$ at every lead time causes a later driver pulse to enter Z^F_{t+h} and propagate across subsequent horizons under λ_f. FR^{id} (HB) moves λ_{fast} from 0.575 to 0.449 and λ_{slow} from 0.575 to 0.641, recovering the correct persistence ordering and reducing MAE from 0.063 to 0.031. Its fast-path pulse response rises to 0.371 versus the ground-truth 0.60, with only 0.010 crosstalk onto y_{slow}. That gain is confined to the fast target: MAE falls from 0.077 to 0.013, while the slow-path pulse stays at 0 and y_{slow} MAE remains 0.049.

Graph-first residual routing recovers the remaining pathway. FR-GR^{id} moves the learned persistences to 0.429 and 0.679, closer to the ground-truth $(0.40, 0.75)$, reduces overall MAE to 0.013, and lowers the two per-target MAEs to 0.013 and 0.014. The corresponding future-pulse responses are 0.430 on the fast wired path and 0.141 on the slow wired path, while off-path crosstalk remains 0.008 and the null-driver response 0.002. Thus, under matched dynamics, the rollout variants recover the correct persistence ordering, while the full FR-GR^{id} model exhibits pathway-specific responses on both wired paths.

K.4 Saturating $\tanh$ mismatch

The matched system in Eq. (8) uses the same linear leaky integrator as the model. To test a mild class mismatch we keep (λ_f, α_f), $d \equiv 0$, observation noise, drivers, and the typed prior fixed, and replace the latent update by $z \leftarrow \tanh\big(\lambda z + (1-\lambda)\alpha u\big)$. A unit pulse then peaks at $\tanh(0.60) \approx 0.537$ on y_{fast} and $\tanh(0.25) \approx 0.245$ on y_{slow}. The trained model remains the linear integrator of Eq. (2). Optimization and initialization follow Appendix K.2 (width 128, batch 16, learning rate 3×10^{-4}, seed 42, α initialized at 0.5). The identification ladder is trained for 40 epochs, matching the linear system; the last row trains the same FR-GR^{id} recipe for 80 epochs.

Tables 26–27 report the same probes as the matched system.

Table 26: Learned persistence and test MAE on the saturating $\tanh$ system (seed 42). Rank is $\lambda_{\mathrm{slow}} > \lambda_{\mathrm{fast}}$. IE^{id} fixes $\lambda_f = 0$. DC^{id} has no family persistence. The ladder uses 40 epochs; the last row is FR-GR^{id} at 80 epochs. MAE is on the physical target scale.

Method	λ_{fast}	λ_{slow}	Rank	MAE	MAE$_{\mathrm{fast}}$	MAE$_{\mathrm{slow}}$
Truth	0.40	0.75	—	—	—	—
Init	0.575	0.575	—	—	—	—
DC^{id}	—	—	—	0.044	0.056	0.034
IE^{id}	0	0	No	0.044	0.055	0.034
RB^{id}	0.622	0.628	No	0.044	0.055	0.034
RB-GR^{id}	0.632	0.639	No	0.044	0.055	0.034
FR^{id} (HB)	0.580	0.591	No	0.044	0.055	0.034
FR-GR^{id} (40 ep)	0.401	0.609	Yes	0.023	0.013	0.034
FR-GR^{id} (80 ep)	0.343	0.615	Yes	0.012	0.012	0.012

Without future rollout the 40-epoch probes match the linear system: DC^{id}, IE^{id}, RB^{id}, and RB-GR^{id} leave MAE at 0.044 and produce no future-pulse response, and the current-belief persistences stay near 0.62–0.64. Unlike the matched system, FR^{id} (HB) also stays near the initialization (0.580,

Table 27: Peak absolute response to a unit future pulse on the saturating $\tanh$ system (horizon 32, zero context). Truth peaks are 0.537 ($u_{\text{fast}} \to y_{\text{fast}}$) and 0.245 ($u_{\text{slow}} \to y_{\text{slow}}$). The ladder uses 40 epochs; the last row is FR-GR^{id} at 80 epochs.

Method	$I_{\text{fast}\to\text{fast}}$	$I_{\text{fast}\to\text{slow}}$	$I_{\text{slow}\to\text{slow}}$	I_{null}
Truth	0.537	0	0.245	0
DC^{id}	0.000	0.000	0.000	0.001
IE^{id}	0.000	0.000	0.000	0.000
RB^{id}	0.000	0.000	0.000	0.000
RB-GR^{id}	0.000	0.000	0.000	0.000
FR^{id} (HB)	0.000	0.000	0.000	0.000
FR-GR^{id} (40 ep)	0.278	0.008	0.001	0.000
FR-GR^{id} (80 ep)	0.330	0.015	0.126	0.000

0.591), does not recover the rank, and records zero pulse responses. At 40 epochs, FR-GR^{id} separates the persistences ($0.401 < 0.609$) and halves MAE to 0.023 by fitting the fast target (0.013), but y_{slow} MAE remains 0.034 and the slow-path pulse is 0.001 against the saturated peak 0.245. The fast-path pulse is 0.278 against 0.537, with 0.008 crosstalk.

At 80 epochs the same FR-GR^{id} recipe recovers both wired pathways: the slow-path pulse rises to 0.126, the fast-path pulse to 0.330, and the two per-target MAEs both fall to 0.012. The learned persistences are then 0.343 and 0.615, still farther from $(0.40, 0.75)$ than the 40-epoch matched estimates (0.429, 0.679).

L TYPED-GRAPH GENERATION PROMPT

Cursor Grok 4.6 (Cursor Team, 2026) is queried once per dataset, offline. PreDist receives a second query that tags each edge dyn or alg. The input is the schema, variable descriptions, and the constraints below. The output is a JSON graph that is checked against a fixed schema, reviewed, and frozen; training loads that file and never calls the model. Self-loops and reverse message-passing edges are added programmatically after the proposal. Listing 1 is the shared instruction. Dataset-specific node inventories and routing constraints follow. CGMacros and Shanghai Diabetes use the same shared instruction with a single glucose target; their frozen prompts are stored with the graphs.

Listing 1: Shared prompt for typed-graph generation.

```
You are proposing a minimal symbolic relational graph prior for an
action-conditioned world model.

Given the dataset schema and the hard constraints below, return a
machine-readable JSON graph with:
1) nodes: id, display_name, type in {state, action, exo, latent-effect}
2) directed edges: source, target, relation, sign in {positive,
   negative, unknown}, confidence in [0,1], short rationale

Shared constraints:
- Driver nodes (action or exo) update latent-effect nodes; latent-effect
  nodes reach observed targets.
- Do not add direct action->state edges that skip latent-effect nodes.
- Do not add intra-state edges unless a dataset constraint explicitly
  allows an algebraic one.
- All exogenous channels remain available as direct future-driver
  tokens; only the listed subset become graph exo nodes.
- Exo nodes are source-only: do not add reverse message-passing edges
  into exo.
- Use the shared relation 'activates' for action-driver and exo-driver
  edges into effect nodes unless a dataset constraint assigns a unique
  relation id.
- Give each sign=unknown action-modulates edge a unique relation id so
  the GNN does not share W_r across them.
- Self-loops and reverse message-passing edges are added
  programmatically later; do not invent extra causal claims.
- Output JSON only, matching the repository schema.
```

VitalDB. Intraoperative vitals on a regularly sampled grid: MBP, HR, SpO_2, BIS. Dense continuous actions: propofol and remifentanil pump rates. No exogenous variables.

Listing 2: VitalDB graph constraints.

```
Hard constraints:
- Nodes: MBP, HR, SpO2, BIS (state); Propofol, Remifentanil (action);
  Propofol Effect, Remifentanil Effect (latent-effect).
- Edges: each drug activates its own effect; Propofol Effect decreases
  BIS, MBP, HR; Remifentanil Effect decreases BIS, MBP, HR.
- Do not add a drug-effect edge into SpO2.
- Do not merge the two drugs into one shared effect node.
- Do not create nodes for rate, Delta-rate, or on/off.
```

Greenhouse. Indoor climate on a 5 min grid: Tair, Rhair, CO_2air, PARin. Dense continuous actions: Tpipe, VentLee, VentWind, AssimLight, EnScr, BlckScr, CO_2dosing. All eight outdoor-weather channels stay on the direct path; only Tout, Windsp, PARout, and AbsHumOut are graph exo nodes.

Listing 3: Greenhouse graph constraints.

```
Hard constraints:
- State: tair, rhair, co2air, parin.
- Action drivers: tpipe, vent_lee, vent_wind, assim_light, en_scr,
  blck_scr, co2_dosing.
- Exo drivers: tout, windsp, parout, abshumout (refreshed each step).
- Effects: thermal, ventilation, light, screen, co2.
- Routes: tpipe+Tout -> thermal; vent_lee+vent_wind+Windsp+Tout+AbsHumOut
  -> ventilation; assim_light+PARout -> light;
  en_scr+blck_scr+PARout+Tout -> screen; co2_dosing -> co2.
- Do not add Rhout, Iglob, Winddir, or Rain as graph nodes.
- Do not add dense exo-to-target edges.
```

Wastewater nutrient control. Tank-1 nutrient state on a 2 min grid: DO, NH_4, PO_4. Continuous actions: metal-salt dosing and blower factor. Categorical actions: inlet and outlet process phases. Exogenous: inflow and water temperature; both are graph exo nodes.

Listing 4: Wastewater graph constraints.

```
Hard constraints:
- State: do_tank1, nh4, po4.
- Action drivers: inlet_metal_dosing, metal_dosing, max_chem_factor,
  processphase_inlet, processphase_outlet.
- Exo drivers: inflow, temperature (refreshed each step).
- Effects: dosing, aeration, phase.
- Routes: inlet_metal_dosing+metal_dosing+temperature -> dosing;
  max_chem_factor+inflow+temperature -> aeration;
  processphase_inlet+processphase_outlet+inflow -> phase.
- Do not add a temperature driver into phase_effect.
- Do not treat inflow or temperature as actions.
- Do not add dense exo-to-target edges.
```

PreDist. Substation thermodynamics on a 10 min grid. Continuous actions: valve position and supply-temperature setpoint. Categorical actions: pump status, control mode, three-way valve. Exogenous: outdoor temperature and primary-grid supply temperature. Temperatures are inertial (dyn); heat power and flow are algebraic (alg). A second query, given the same nodes, tags every edge with timescale dyn or alg and uses a disjoint relation vocabulary.

Listing 5: PreDist dual-timescale graph constraints.

```
Hard constraints:
- State: s_hc1_supply_temp, p_hc1_return_temp, p_net_return_temp
  (dyn); heat_power, flow (alg).
- Action drivers: valve_pos, supply_setpoint, pump_status,
  control_mode, dhw_3way.
- Exo drivers: outdoor_temp, p_net_supply (refreshed each step).
- Effects: valve, setpoint, pump. Effects drive temperatures only.
- Tag every edge timescale=dyn or timescale=alg. Use disjoint
  relation ids so G_dyn and G_alg do not share W_r.
- G_dyn routes: valve_pos+dhw_3way+p_net_supply+outdoor_temp -> valve;
  supply_setpoint+control_mode+outdoor_temp -> setpoint;
  pump_status -> pump;
  valve -> supply and both returns;
  setpoint -> supply and both returns;
  pump -> supply.
- G_alg routes: valve_pos+dhw_3way+pump_status -> flow (direct
```

```
   action->state, skipping latent-effect nodes);
   flow -> heat_power (determines);
   p_net_supply and p_hc1_return_temp -> heat_power (primary dT fan-in).
- Do not add effect -> flow or effect -> heat_power.
- Do not add action -> heat_power.
- Do not add supply_setpoint or control_mode into flow.
- Do not add reverse message-passing on G_alg.
- Do not add a weather driver into pump_effect.
- Do not encode maintenance or faults as actions or graph nodes.
- Do not add dense exo-to-target edges.
- Self-loops and reverse MP (dyn only, never into exo) are added
  programmatically later.
```