PROCEEDINGS, 51st Workshop on Geothermal Reservoir Engineering
Stanford University, Stanford, California, February 9-11, 2026
SGP-TR-230

Thermal–Hydraulic–Mechanical Embed-to-Control-and-Observe (THM-E2CO): Toward Real Time Surrogate Modeling of Geothermal Systems

Taha Yehia[1], Moamen Gasser[1], Hossam Ebaid[1], Elias Ghaly[1], Youssef Falola[2], D. Nathan Meehan[1]

[1]Harold Vance Department of Petroleum Engineering, Texas A&M University, College Station, TX, USA

[2]Halliburton, Houston, TX, USA

Taha.Yehia@tamu.edu

Keywords: Geothermal; Modeling; Surrogate; Machine Learning; E2CO

ABSTRACT

This study presents a deep learning framework, Thermal-Hydraulic-Mechanical Embed-to-Control-and-Observe (THM-E2CO), designed for high-fidelity reduced-order modeling of geothermal reservoirs governed by tightly coupled poro-thermo-elastic processes. The architecture extends the E2CO paradigm to address the complex nonlinear interactions between heat transfer, fluid flow, and rock deformation. Unlike static surrogate models, this framework is specifically engineered to predict and capture the temporal evolution of porosity and permeability as they respond to thermal drawdown and mechanical loading within the reservoir.

The THM-E2CO architecture integrates a 3D convolutional encoder-decoder to compress and reconstruct full field pressure and temperature distributions while preserving critical thermal-hydraulic-mechanical couplings. A nonlinear latent-state transition network models the temporal evolution of the reservoir under varying operational controls. Crucially, the model is trained to observe the feedback loops where temperature gradients and pore pressure fluctuations drive mechanical deformation, subsequently altering the flow properties of the rock matrix. The learning process is regularized through a composite loss function that enforces data reconstruction fidelity, latent consistency, and conservation of mass and energy, ensuring physical plausibility across the coupled domains.

Using CMG STARS as the high-fidelity simulator, a diverse ensemble of training data was generated via CMOST-driven variations in injection rates and producer bottom hole pressures. The dataset comprises over 1,000 simulations spanning a 30-year production period, featuring a six-well configuration. This wide operational space allows the model to learn robust coupled behavior and the nonlinear degradation or enhancement of reservoir transmissibility over time.

The novelty of this work lies in the framework's ability to explicitly capture the poro-thermo-elastic response of the reservoir, providing a rapid alternative to computationally expensive geomechanical solvers. By embedding the physics of stress-dependent permeability into the latent space, the model achieves high-fidelity reproduction of full physics simulations with computational accelerations exceeding 30,000 times. This development provides a scalable foundation for the prediction of rapid geothermal performance, uncertainty quantification, and real time optimization under complex geomechanical conditions.

1. INTRODUCTION

Geothermal energy is increasingly recognized as a firm, low-carbon resource capable of delivering continuous electricity and direct use heat without reliance on intermittent meteorological conditions (Elsayed and Okoroafor, 2024b; Gasser et al., 2025). By extracting heat from the subsurface, geothermal systems can achieve high-capacity utilization when reservoirs and surface facilities are appropriately engineered and managed (Elsayed and Okoroafor, 2024a; Oliver et al., 2024; Gasser et al., 2025). This firm generation characteristic is particularly valuable in deeply decarbonized energy systems that must maintain reliability while reducing emissions (Olasolo et al., 2016; Okoroafor et al., 2023; Tayyib et al., 2025). Despite its potential, geothermal deployment remains constrained by subsurface uncertainty, high upfront capital requirements, and the need to demonstrate long term reservoir performance, motivating the use of predictive modeling and decision support tools to reduce development risk (Carter et al., 2025; Yehia et al., 2024; Carter et al.).

High-fidelity geothermal reservoir simulation becomes especially computationally demanding when thermo–hydro–mechanical (THM) coupling is included (Okoroafor and Horne, 2024; Ebaid et al., 2024). In such systems, pressure, temperature, and stress evolve in a tightly coupled manner, and thermo-poroelastic effects can induce time-dependent changes in porosity and permeability that feed back into flow capacity and heat extraction. From a numerical standpoint, THM coupling increases state dimensionality, stiffness, and nonlinearity, often requiring fine spatial resolution and small time steps (Falola et al., 2025a; Yehia et al., 2025). These costs are amplified in practical workflows that rely on ensembles of simulations for uncertainty quantification, calibration, and optimization of operational controls, rendering full-physics simulation prohibitive for iterative decision-making (Ghaly and Shor, 2025; Falola et al., 2025b, 2022; Yehia et al., 2023).

These challenges have motivated extensive research into reduced-order models (ROMs) and data-driven surrogate approaches for subsurface systems. Classical projection-based ROMs can provide speedups near training trajectories but often struggle under strong nonlinearity and control perturbations, limitations that are exacerbated in geothermal settings (Falola et al., 2024, 2026, 2023). Learning-based surrogates have therefore gained prominence, particularly latent state-space models that compress high-dimensional reservoir states and evolve them forward under control actions. The Embed-to-Control (E2C) framework introduced a principled approach by learning locally linear latent dynamics suitable for control while retaining nonlinear encoders and decoders for complex spatial fields (Watter et al., 2015; Le and Tao, 2024).

Early E2C-style subsurface surrogates focused primarily on state prediction, with well responses computed indirectly from reconstructed fields using explicit well models. This indirect coupling can introduce additional error sources and limit robustness in optimization workflows. To address this limitation, the E2CO framework augments latent dynamics with a learned observation operator that directly predicts well level quantities from the latent state, enabling end-to-end evaluation of control strategies without explicit well equations (Coutinho et al., 2021). Subsequent studies demonstrated the practicality of E2CO-style surrogates in 3D settings, life-cycle production optimization, and control under geological uncertainty (Atadeger et al., 2022, 2024; Nguyen and Onur, 2024a, b).

E2CO-based surrogates have since been applied to a range of subsurface decision-making problems, including well control optimization (Kompantsev et al., 2024), CO2 storage and pressure management (Chen et al., 2025; Noh and Swidinsky, 2025), and long-horizon control scenarios requiring stable multi-step rollouts. Recent work has further highlighted the importance of mitigating rollout drift through multi-step latent training when surrogates are deployed for long-term forecasting and control (Chen et al., 2024a; Adeyemi and Onur, 2025; Chen et al., 2024b).

Despite these advances, most existing E2CO-style applications treat porosity and permeability as static spatial fields that parameterize the dynamics but are not themselves predicted states. This assumption becomes limiting for geothermal reservoirs, where THM coupling induces sustained, time-dependent property evolution that directly influences transmissibility, pressure support, and heat extraction efficiency. The central contribution of this work is therefore to extend the E2CO paradigm to geothermal THM systems by explicitly modeling porosity and permeability as dynamic state variables alongside pressure and temperature. By learning the coupled evolution of these quantities under time-varying operational controls, the proposed surrogate is designed to capture key THM feedback mechanisms while remaining directly usable for long-horizon forecasting, monitoring, and optimization.

Building on this framework, the paper develops and evaluates a THM-aware geothermal E2CO reduced-order model trained on a long-horizon, control-driven dataset comprising approximately 1,000 simulation cases spanning 30 annual time steps. The model provides unified prediction of full field reservoir states and control-relevant well outputs and is systematically assessed in terms of predictive accuracy and computational efficiency for geothermal decision-support applications.

2. METHODOLOGY

2.1 Base Model and Dataset Generation

This study constructs a supervised learning dataset from high-fidelity thermal reservoir simulations to train a THM-E2CO for geothermal systems. The workflow is designed to preserve the governing physics of coupled THM processes while spanning a commercially relevant range of long-term injection and production strategies. All simulations are performed using CMG STARS with CMOST used to automate the sampling of operational controls via Latin Hyper-cube Sampling (LHS). The resulting dataset integrates static geological properties with time-dependent controls and records both full field state evolution and well level performance metrics over a 30-year production horizon. We made the model to be open source (available at https://github.com/hossam2179/geothermal-stars-model/tree/ main).

2.1.1 Brady-Derived Geothermal Reservoir Model

The base reservoir model is derived from the Brady Hot Springs geothermal field in Nevada, USA, a liquid-dominated geothermal system characterized by elevated subsurface temperatures and long-term commercial production. The model is configured to represent a single-phase hot water geothermal reservoir and is intended to capture coupled pressure, temperature, and rock property evolution under sustained injection and production. This configuration provides a realistic field scale setting for generating training data representative of geothermal operations. The base model is displayed in Fig. 1, which shows the location of six wells (three injectors and three producers). The model discretization is summarized in Table 1.

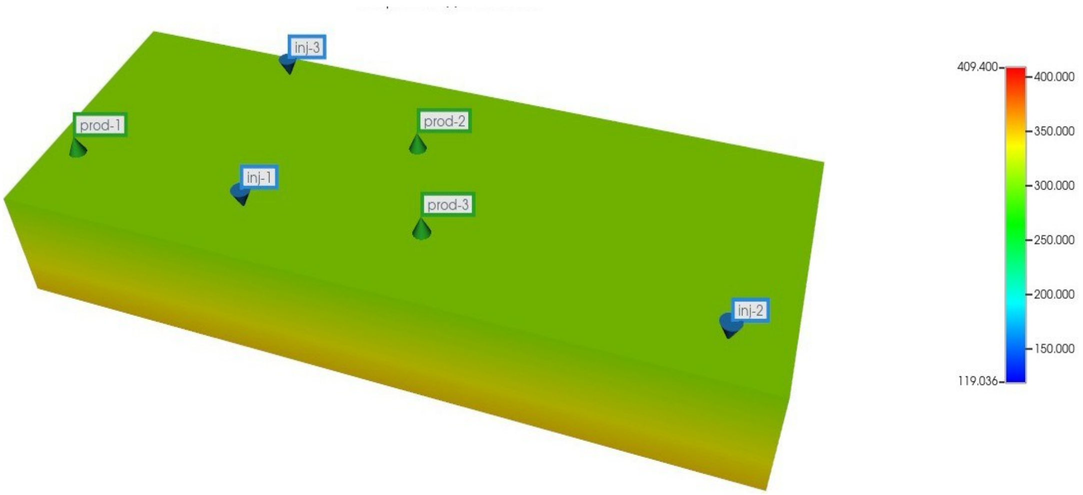

Figure 1: 3D view of the base geothermal reservoir model showing the spatial distribution of injection and production wells.

Table 1: Grid discretization and stratigraphic sizing for the Brady derived model.

Category	Parameter	Value	Unit/Description
Grid Resolution	Total Nodes ($I \times J \times K$)	$34 \times 16 \times 45$	24,480 Cells
Lateral Sizing	Block Dimension ($DI \times DJ$)	500×400	ft × ft
Vertical Structure	Overburden (Layers 1-10)	3,000	ft (300 ft/layer)
	Primary Reservoir (Layers 11-35)	500	ft (20 ft/layer)
	Underburden (Layers 36-45)	5,500	ft (550 ft/layer)
Global Geometry	Top of Grid (DTOP)	2,000	ft

Petrophysical characteristics are assigned via a tiered rock-type approach to reflect the heterogeneous conductivity of deep geothermal systems. Horizontal permeability within the reservoir interval ranges from 11.35 to 162.5 mD, while vertical permeability is constrained by a global anisotropy ratio (k_v/k_h) of 0.1 to simulate stratigraphic baffles. A 3D geomechanical module is active throughout the simulation to capture the thermoelastic response of the rock matrix to cold-water injection. The stress state is initialized in a normal faulting regime, where the vertical stress represents the maximum principal stress. Rock deformation is primarily driven by thermal contraction (α), which decreases effective rock volume as the cooling front advances. Detailed input properties, including thermal conductivity, heat capacity, and Young's modulus, are provided in Table 2.

Table 2: Petrophysical, thermal, and geomechanical input properties for different stratigraphic units.

Property	Unit	Overburden	Reservoir	Underburden
Matrix Porosity (φ)	fraction	0.08	0.25	0.08
Horizontal Perm (k_h)	mD	0.01	11.35 – 162.5	0.01
Pore Compressibility	psi^{-1}	5.0×10^{-6}	2.27×10^{-6}	1.5×10^{-7}
Thermal Conductivity	Btu/ft-day-°F	25.0	38.0	45.0
Volumetric Heat Capacity	$\text{Btu/ft}^3\text{-°F}$	35.0	40.0	42.0
Young's Modulus (E)	psi	0.4×10^{6}	3.0×10^{6}	8.7×10^{6}
Poisson's Ratio (v)	fraction	0.25	0.25	0.10
Vertical Stress (σz)	psi	—	5,000	—
Min. Horiz. Stress (σx)	psi	—	2,500	—
Max. Horiz. Stress (σy)	psi	—	3,900	—

Geothermal fluid is pure water, modeled as a slightly compressible liquid. Water density ρw(p,T) and enthalpy hw(p,T) are computed internally by STARS approximating steam-table data for subcooled water. Initial conditions are representative of an undisturbed, liquid-dominated reservoir with a hydrostatic pressure gradient (~0.45 psi/ft) and a conductive geothermal gradient of 8-9◦F/100 ft. Fluid-flow boundary conditions represent an open system with lateral boundaries coupled to infinite-acting analytical aquifers, while the bottom boundary supplies a constant basal heat flux of approximately 0.1 W/m2 (nearly 193◦C).

2.1.2 CMOST Dataset Generation Strategy

A design-of-experiments approach is used to generate a diverse ensemble of injection and production scenarios for surrogate model training. CMOST is employed to sample well control parameters using LHS, generating a total of 1,000 simulation cases. The primary control variables are injector water rates (sampled between 20,000 and 80,000 BWPD) and producer bottom-hole pressures (sampled between 1,500 and 2,000 psi). An example of the time-dependent control schedules and resulting responses is illustrated in Fig. 2. For each case, well level outputs (BHP, water rates, and energy production rates) and field level outputs (full field 3D pressure, temperature, porosity, and permeability snapshots) are recorded. Fig. 3 shows the representative spatial evolution of these state variables over the 30-year field life. Simulations were performed using CMG STARS (version 2025.2) on a workstation featuring a multi-core Intel® Xeon® processor and 768 GB of RAM. The total ensemble of 1,000 cases required approximately 10 days of wall-clock time, with individual runs averaging 15 minutes.

2.2 E2CO Model Architecture

We define the geothermal reservoir dynamics as a discrete-time control system. The dataset is structured into three distinct components: the high-dimensional spatial state ($\mathbf{x}_t$), the operational control vector ($\mathbf{u}_t$), and the observation vector ($\mathbf{y}_t$).

The reservoir state at any time step t is represented as a high-dimensional tensor $\mathbf{x}_t \in \mathrm{R}^{N_c\times N_x \times N_y\times N_z}$ defined on a structured grid with dimensions $34\times 16\times 25$. To capture the coupled THM behavior, we define the state tensor as a concatenation of four distinct physical ($N_c = 4$):

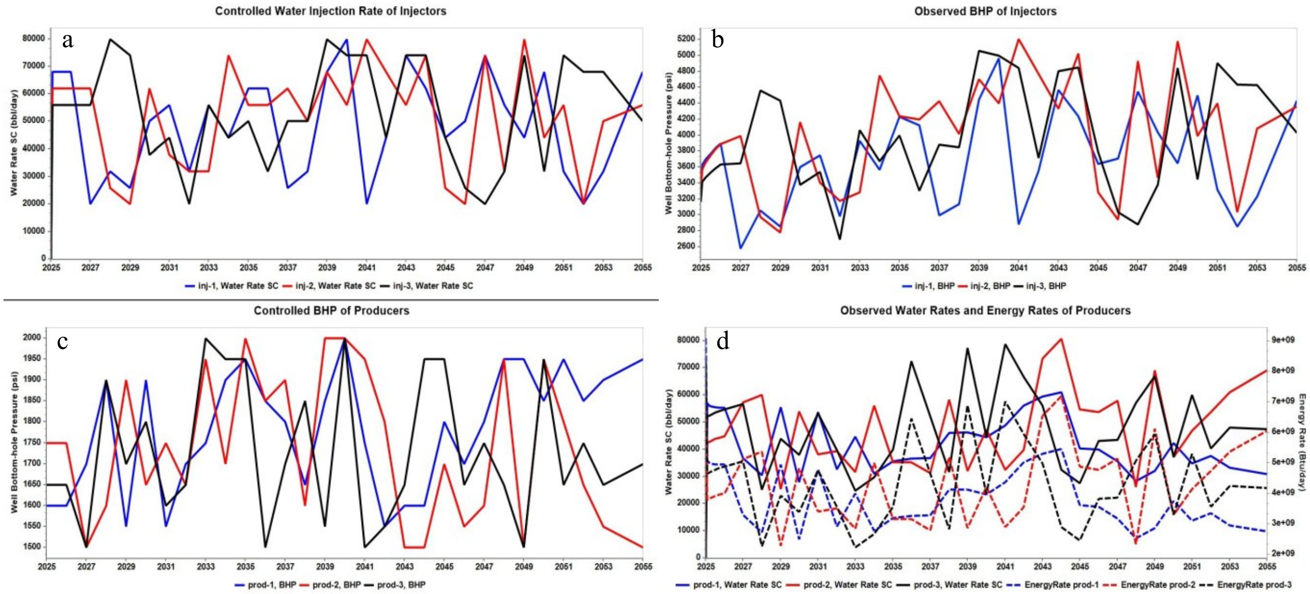

Figure 2: Example CMOST-generated control schedules: (a) Injection rates, (b) observed injector BHP, (c) producer BHP setpoints, and (d) resulting water production rates.

$$\mathbf{x}_t = [\mathrm{P}_t,\ \mathrm{T}_t,\ \phi_t,\ \mathrm{k}_t] \tag{1}$$

where each component represents a 3D field:

1. Pressure (Pt): The fluid pressure distribution driving Darcy flow.
2. Temperature (Tt): The thermal energy distribution governing heat transport.
3. Porosity (φt): The dynamic porosity field, evolving via thermo-poroelastic stress.
4. Permeability (kt): The dynamic permeability field, updated at each step to reflect stress-dependent changes.

The system evolution is driven by an operator-defined control vector ut ∈ Rnu where nu = 6. This vector con-catenates the injection and production targets for the six active wells:

$$\mathbf{u}_t = [q_{w,inj}^{(1)}, \dots, q_{w,inj}^{(3)},\ p_{wf,prod}^{(1)}, \dots, p_{wf,prod}^{(3)}]^\top \tag{2}$$

The components correspond to:

- Injectors ($I_1 – I_3$): Water injection rates ($q_{w,in\ j}$).
- Producers ($P_1 – P_3$): Bottom-hole pressure (BHP) targets ($p_{wf,prod}$).

For decision-making and closed-loop monitoring, the model explicitly predicts an observation vector $y_T \in R^{ny}$ where $n_y = 9$. This vector aggregates the critical performance metrics derived from the well responses:

$$\mathbf{y}_t = [p_{wf,inj}^{(1\dots3)},\ q_{w,prod}^{(1\dots3)},\ \dot{E}_{prod}^{(1\dots3)}]^\top \tag{3}$$

- Injectors (I1 – I3): Resulting BHP (pwf,in j), critical for monitoring formation integrity.
- Producers (P1 – P3): Produced water rates (qw,prod).
- Producers (P1 – P3): Produced energy rates (E˙prod), representing the instantaneous thermal power extraction.

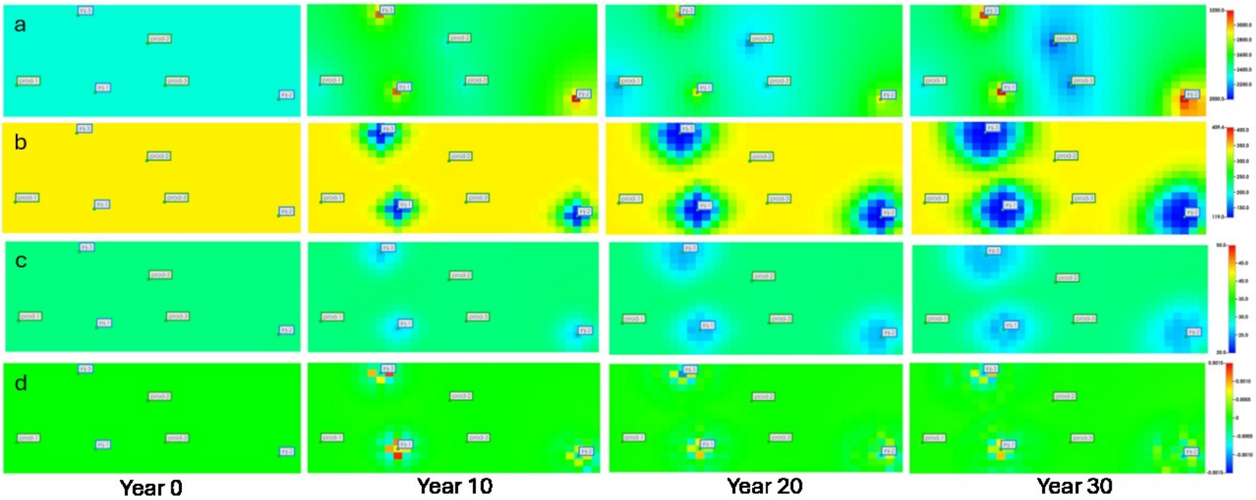

Figure 3: Spatial evolution of reservoir state variables (pressure, temperature, permeability, and porosity) at the top of the reservoir for Years 0, 10, 20, and 30.

The E2CO framework employs a deep 3D convolutional autoencoder coupled with a recurrent latent-dynamics block. The complete network architecture is illustrated in Fig. 4, detailing the tensor transformations from the high-dimensional physical space to the compressed latent manifold.

The architecture comprises four coupled components trained end-to-end:

2.2.1 Encoder (E) and Decoder (D)

The Encoder maps the instantaneous reservoir state $x_t \in R^{Nc\times Nx\times Ny\times Nz}$ to a low-dimensional latent vector $z_t \in R_{dz}$. As shown in Fig. 4, it utilizes four 3D convolutional layers followed by a ResNet block to extract non-linear spatial correlations between pressure, temperature, and rock properties. Conversely, the Decoder mirrors this structure using

transposed convolutions (deconvolutions) to map the latent vector back to the physical domain:

$$\mathbf{z}_t = \mathcal{E}(\mathbf{x}_t), \quad \hat{\mathbf{x}}_t = \mathcal{D}(\mathbf{z}_t) \tag{4}$$

2.2.2 Linear Transition Model (T)

To enable rapid control optimization, we enforce a locally linear structure on the latent dynamics. The transition block (Center panel of Fig. 4) uses a dense network to generate state-dependent system matrices At and Bt, which evolve the state linearly:

$$\mathbf{z}_{t+1} = \mathbf{A}_t(\mathbf{z}_t)\mathbf{z}_t + \mathbf{B}_t(\mathbf{z}_t)\mathbf{u}_t \tag{5}$$

This structure ensures that the control inputs u_t (Injection Rates and Producer BHPs) have a direct and interpretable influence on the system evolution via the matrix B_t.

2.2.3 Observation Network (Well Performance)

To bypass the computational cost of calculating well indices from grid blocks, the Observer network predicts well outputs directly from the latent state and controls:

$$\hat{\mathbf{y}}_t = \mathbf{C}_t \mathbf{z}_t + \mathbf{D}_t \mathbf{u}_t \tag{6}$$

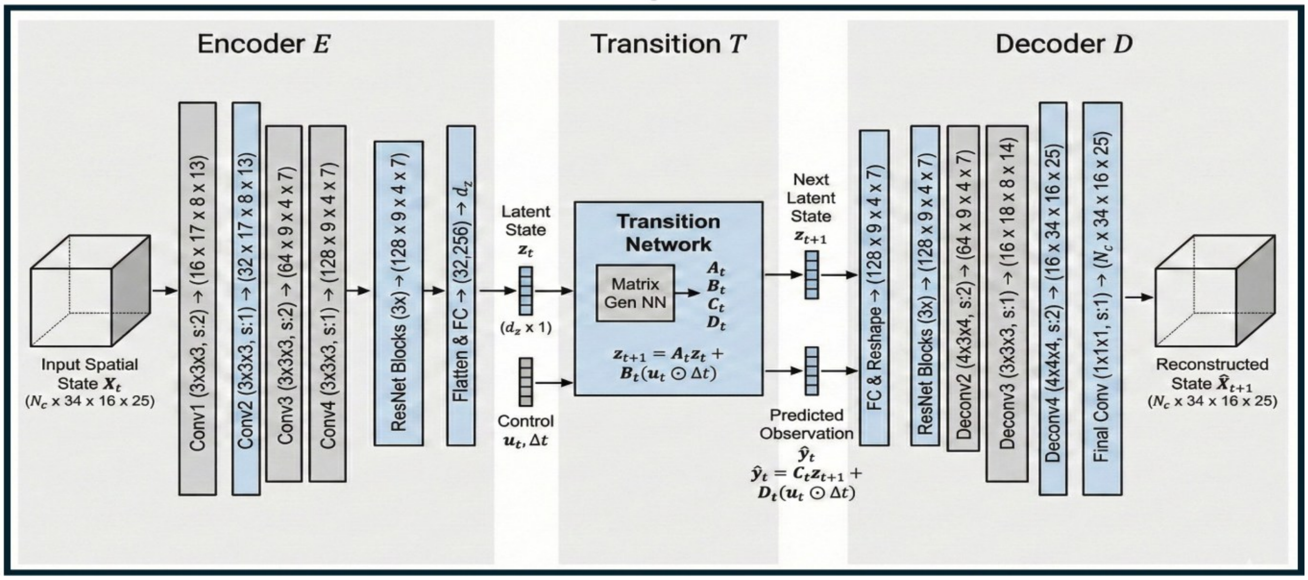

Figure 4: E2CO Network Architecture. The framework consists of three coupled modules: (Left) The Encoder compresses the Nc = 4 channel geothermal state into a latent vector zt; (Center) The Transition network evolves the latent state linearly based on controls ut; (Right) The Decoder reconstructs the spatial pressure, temperature, porosity, and permeability fields. The specific layer dimensions (filters, kernels, strides) ensure the conservation of spatial topology during compression (modified after (Nguyen et al., 2024).

Here, Ct maps the latent reservoir state to the observation space, while Dt captures the immediate feed-through effect of control actions on well responses. This formulation allows the model to instantly forecast energy production and pressure constraints without fully reconstructing the 3D grid at every time step.

2.3 Training and Optimization

The THM-E2CO model is trained using a supervised learning approach, utilizing a multi-objective loss function designed to balance convergence across the coupled poro-thermo-elastic domains.

The composite loss function L_{total} is formulated to ensure physical consistency across spatial, temporal, and operational dimensions:

$$\mathcal{L}_{total} = \sum_{c=1}^{N_c} \lambda_{rec,c} \mathcal{L}_{rec,c} + \lambda_{dyn} \mathcal{L}_{dyn} + \lambda_{obs} \mathcal{L}_{obs} \tag{7}$$

Where $\mathcal{L}_{rec,c}$ represents the reconstruction loss for the c -th physical channel (Pressure, Temperature, Porosity, and Permeability).

To mitigate "loss competition", a phenomenon where gradients from easily learnable features dominate the optimization. we implement a per-channel weighting strategy ($\lambda_{rec,c}$). Static or slowly evolving properties, such as porosity and permeability, are assigned lower relative weights. Conversely, pressure and temperature fields, which exhibit high-frequency transients and tight THM coupling, are assigned higher weights to ensure the model captures complex fluid dynamics and avoids vanishing gradients during the learning process.

In this study, a dedicated flux loss term is deliberately omitted. While flux-based regularization can improve local mass balance, previous literature suggests it significantly increases computational overhead during training without a proportional gain in global reconstruction accuracy. By prioritizing the reconstruction of state variables and well observations, the framework achieves computational accelerations exceeding 10,000 times. Future work will investigate the integration of physics-informed flux and heat-loss components to enhance stability in high-velocity fracture flow regimes.

Once trained, the E2CO model offers two distinct prediction modalities, as illustrated in Fig. 5.

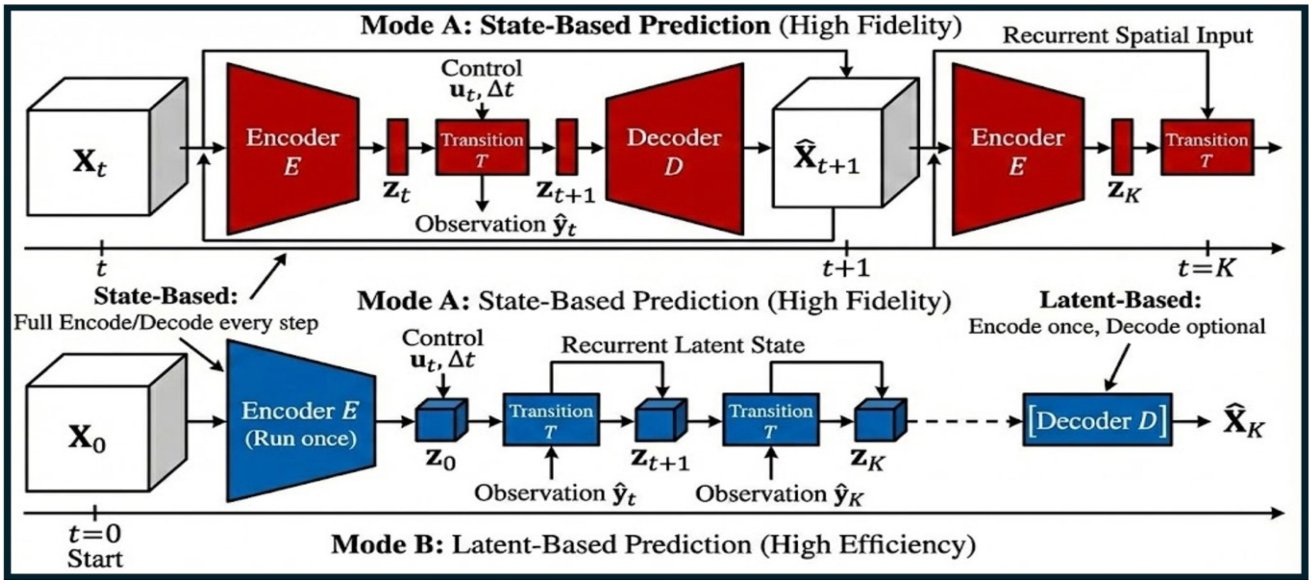

Figure 5: Prediction Modes Workflow. Mode A: State-Based Prediction (Red path) reconstructs the full 3D spatial state at every time step, re-encoding it to drive the next iteration. Mode B: Latent-Based Prediction (Blue path) executes the encoder only once at initialization (t = 0). The system then evolves exclusively within the compressed latent space, decoding only when a spatial snapshot is required. Mode B offers a speedup over the high-fidelity simulator (modified after (Coutinho et al., 2021).

1. Mode A: State-Based Prediction (High Fidelity) - Shown in the top panel (Red blocks) of Fig. 5, this mode performs a full closed-loop update.

- **Workflow:** At each step t, the predicted latent state $\mathbf{z}_{t+1}$ is decoded to form the full spatial field $\hat{\mathbf{X}}_{t+1}$. This reconstructed field is then fed back into the Encoder to initialize the next step.

- **Application:** This mode is computationally intensive but necessary for detailed reservoir surveillance, allowing engineers to inspect thermal fronts, pressure anomalies, and porosity and permeability changes at every time step.

2. Mode B: Latent-Based Prediction (High Efficiency) - Shown in the bottom panel (Blue blocks) of Fig. 5, this mode decouples the dynamics from the high-dimensional grid.

- **Workflow:** The Encoder is run only once at $t = 0$. For all subsequent steps, the model iterates purely in the latent space ($\mathbf{z}_t \rightarrow \mathbf{z}_{t+1}$) using the linear transition matrices. The Decoder is bypassed entirely during the time-stepping loop and is invoked only if a final visualization is needed.

- **Application:** By avoiding expensive 3D convolutions at every step, this mode achieves extreme computational speedups, making it the standard choice for real time control and iterative optimization of energy rates.

3. RESULTS AND DISCUSSIONS

Before presenting time series results for well observations, we first provide an illustrative example of the model ability to predict full field 3D reservoir properties as they evolve over time. Figs. 6 through 13 show representative predicted spatial distributions of permeability, porosity, temperature, and pressure for a case (Case 780) at the same simulation time, extracted on two different layers (Layer 12 and Layer 22). These figures are intended to demonstrate the structure, coherence, and physical consistency of the predicted volumetric fields, as well as to highlight layer dependent variability, rather than to serve as a pointwise error assessment.

In contrast to purely static geological descriptors, permeability and porosity in this setting evolve in response to pressure and temperature changes induced by long-term injection and production. The predicted permeability and porosity fields (Figs. 6, 7, 10, and 11) exhibit pronounced spatial heterogeneity that differs across layers, while also reflecting temporal modifications associated with poroelastic deformation and evolving effective stress. This indicates that the model is able to capture not only the underlying geological structure, but also its dynamic alteration under changing operational conditions. The predicted temperature and pressure fields (Figs. 8, 9, 12, and 13) reflect the dynamic system response to injection and production over the 30-year horizon. Pressure fields display laterally continuous gradients and localized anomalies consistent with well influence and depth dependent transmissibility, while temperature fields exhibit smooth thermal patterns whose magnitude and spatial extent vary between layers. Together, these trends indicate that the model resolves vertically heterogeneous thermo-hydraulic behavior and its coupling with evolving petrophysical properties.

Importantly, these 3D property predictions are produced directly by the learned spatiotemporal dynamics without any physics-based post processing. Although a high-fidelity full simulation run using CMG STARS requires approximately 15 minutes per case, the THM-

E2CO model generates the average prediction per case in only 0.03 seconds. This represents a computational acceleration of approximately 30,000 times, enabling rapid uncertainty quantification and real time optimization that would otherwise be computationally prohibitive. While near well regions remain challenging due to sharp gradients and strong nonlinear coupling, the layer-scale predictions remain smooth, physically interpretable, and mutually consistent across properties. In the following subsection, we quantify the predictive fidelity of the same model in the observation space via direct, post processing, and free prediction of injector BHPs and producer level energy and water rate time series over the full 30-year rollout.

In addition to predicting the 3D state evolution (pressure, temperature, porosity, and permeability), the proposed model is trained to directly predict well level observations through an observation head, without any post processing or physics-based well model evaluation (e.g., no Peaceman-based rate reconstruction). In this subsection, we evaluate the fidelity of these direct observation predictions on three cases (Case 780, Case 880, and Case 980) out of 250 test cases during the full 30-year rollout. For each case, we report two prediction modes shown in Fig. 14 through 22: (i) state-based prediction, where the observation is inferred from the predicted state trajectory, and (ii) latent-based prediction, where the observation head predicts the observation directly from the latent representation. These two modes help diagnose whether residual mismatch is dominated by (a) state reconstruction/rollout error or (b) the learned observation operator. Each observation is assessed using the coefficient of determination (R^2), root-mean-error (RMSE), mean absolute error (MAE), and average percent error (APE). For a time series $\{y_t\}^T_{t=1}$ and predictions $\{\hat{y}\}^T_{t=1}$, we define

$$\mathrm{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(\hat{y}_t - y_t)^2} \tag{8}$$

$$\mathrm{MAE} = \frac{1}{T}\sum_{t=1}^{T}|\hat{y}_t - y_t| \tag{9}$$

$$\mathrm{APE} = \frac{100}{T}\sum_{t=1}^{T}\left|\frac{\hat{y}_t - y_t}{y_t}\right| \tag{10}$$

$$R^2 = 1 - \frac{\sum_{t=1}^{T}(\hat{y}_t - y_t)^2}{\sum_{t=1}^{T}(y_t - \bar{y})^2} \tag{11}$$

Where $\bar{y}$ is the ground ruth series mean. Parity plots shown alongside the time series further illustrate the dispersion about the perfect rediction line and the ±10% envelopes.

3.1 Overall Performance Across the Three Test Cases

Across the three representative test cases, the model exhibits strong agreement for injector BHP trajectories and consistently good performance for producer energy and water rate observations, despite the absence of any explicit post processing. Aggregating all nine observations across the three cases yields:

- **Injector BHP:** $R^2 \in [0.92, 1.00]$ and APE $\in$ [0.72%, 3.28%].
- **Producer energy:** $R^2 \in [0.85, 1.00]$ and APE $\in$ [0.57%, 9.42%].
- **Producer water rate:** $R^2 \in [0.89, 0.98]$ and APE $\in$ [4.52%, 8.04%].

These ranges are consistent with the qualitative behavior in the time series overlays: BHP is comparatively smooth and primarily controlled by operational settings and connectivity, while energy and water rates are more sensitive to localized THM feedbacks, near well gradients, and nonlinear coupling between temperature-driven viscosity/density changes and evolving porosity/permeability.

3.1.1 Key patterns and implications

Three practical observations emerge from Table 3 and the corresponding figures (i.e Fig. 14 through Fig. 22):

1. **BHP predictions are near exact across cases**. The observation head captures the dominant control response relationship for injectors, with parity points tightly clustered around the 1:1 line and APE typically below ~ 2%. This indicates that the latent dynamics retains sufficient information about pressure propagation and effective transmissibility under varying operational schedules.
2. **Energy rate predictions are accurate but more variable across wells**. The largest errors occur for the most volatile producer energy trajectories (notably in Case 780), where sharper peaks and drops are present. In geothermal THM systems, energy rate depends on both flow rate and enthalpy (hence temperature) and therefore compounds any small mismatch in predicted temperature near producing intervals.
3. **Water rate predictions show the expected difficulty of near well nonlinearities.** While R^3 remains high, APE is systematically larger than that for BHP. This is consistent with rate sensitivity to (i) evolving near well permeability/porosity, (ii) strong coupling between pressure drawdown and mobility, and (iii) operational switching that can create sharp transients at yearly resolution.

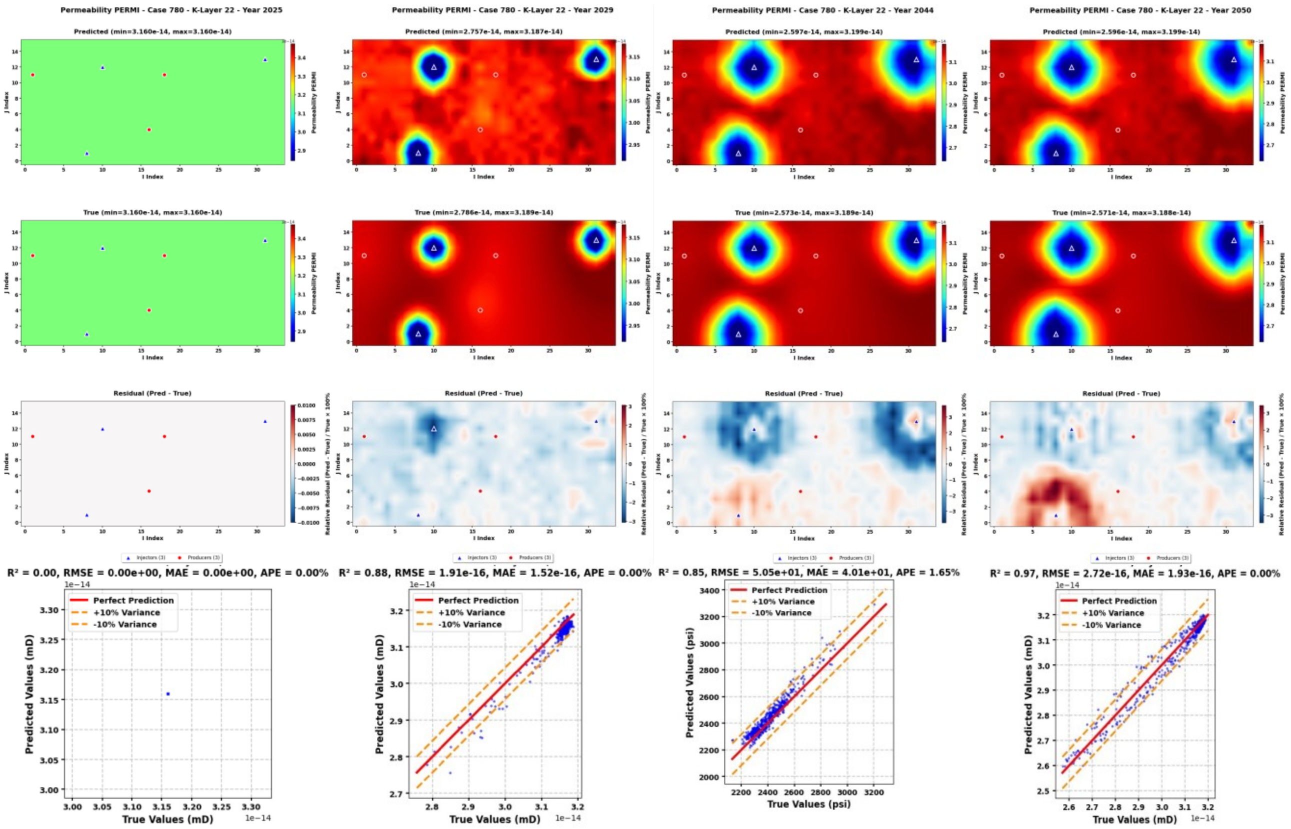

Figure 6: Predicted permeability field for Case 780 on Layer 22.

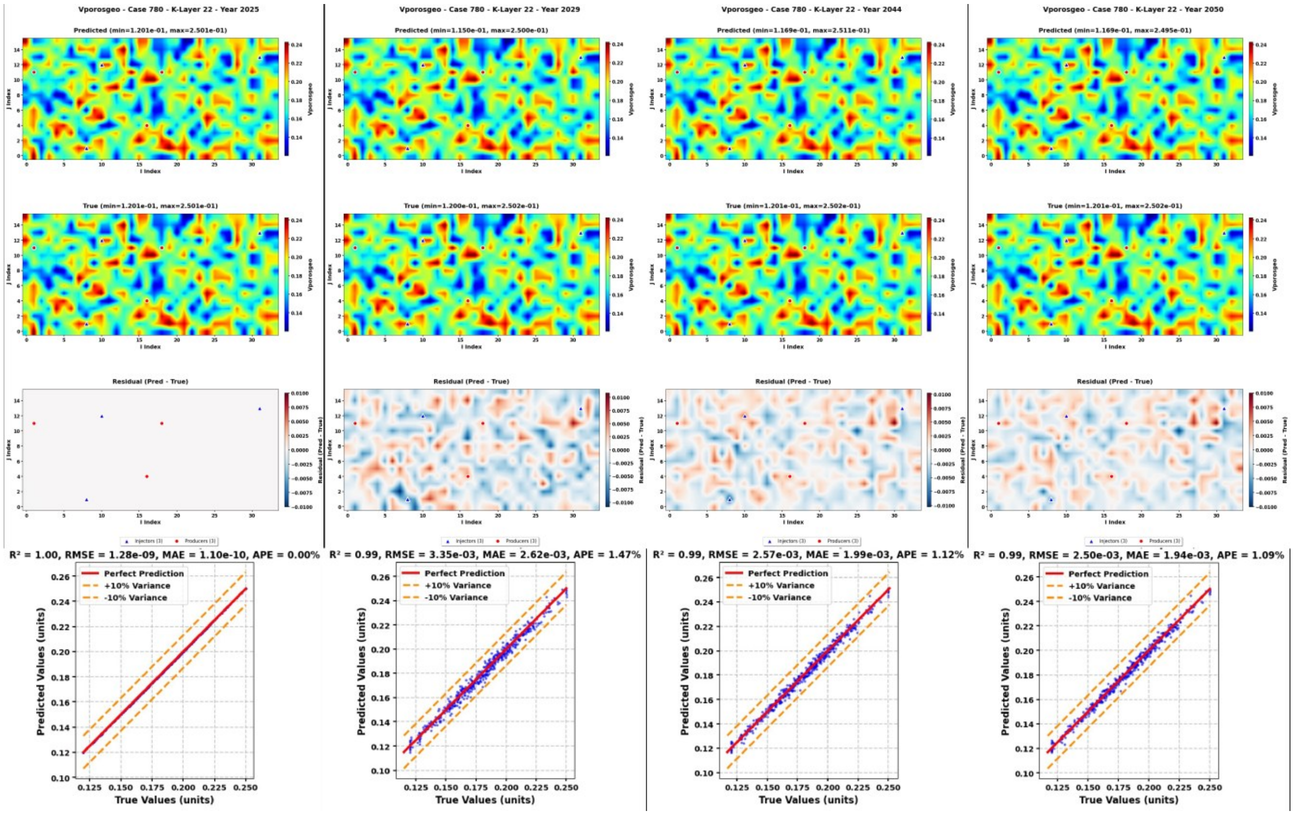

Figure 7: Predicted porosity field for Case 780 on Layer 22.

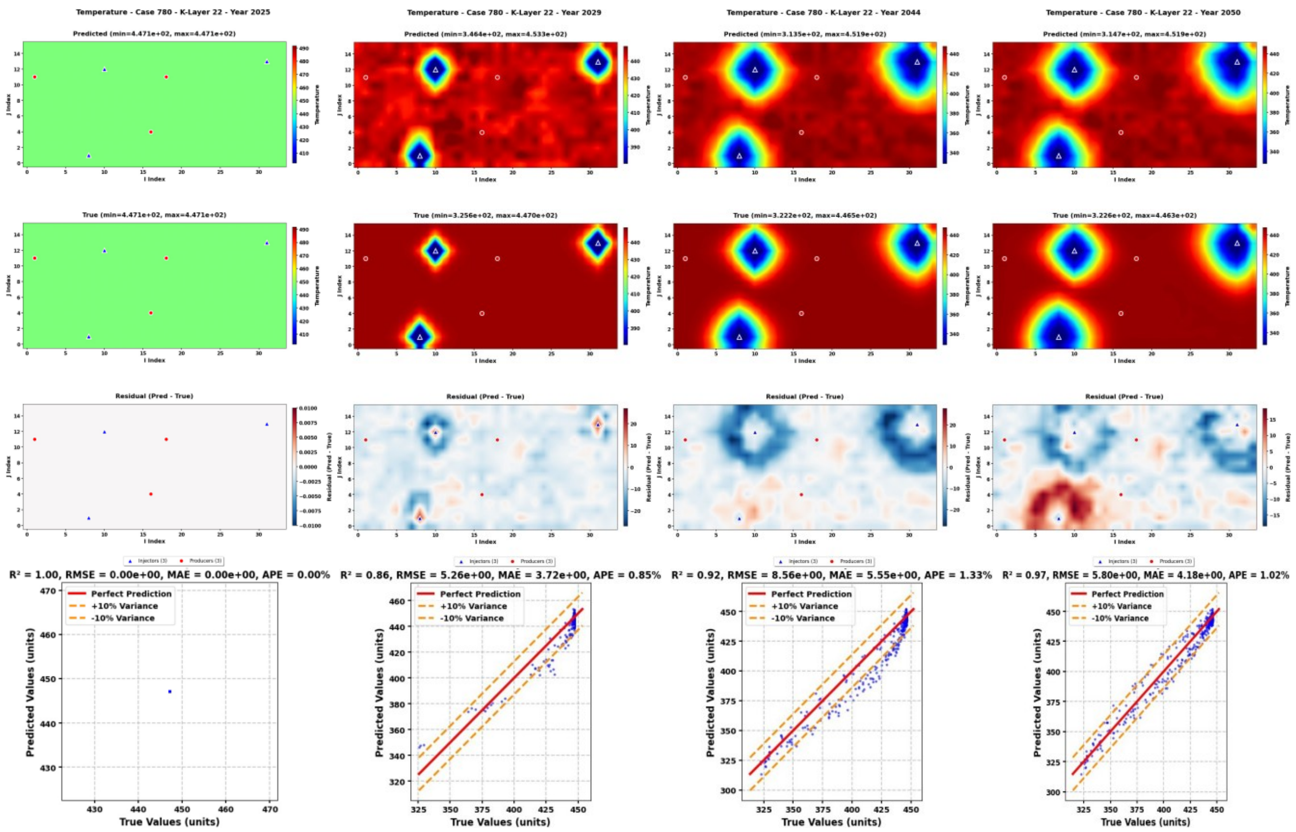

Figure 8: Predicted temperature field for Case 780 on Layer 22.

Table 3: Summary of direct observation prediction performance for three heldout cases. Values are averaged across the three wells within each observation group.

	R^2 (mean)	APE (mean, %)	R^2 (mean)	APE (mean, %)	R^2 (mean)	APE (mean, %)
780	0.963	1.94	0.920	6.01	0.933	6.64
880	0.950	2.05	0.973	3.95	0.973	4.69
980	0.977	1.88	0.967	4.87	0.967	6.02

3.2 Latent-Based vs. State-Based Observation Prediction

A consistent feature across the three cases is that the latent-based observation head and the state-based observation inference are generally close, but not identical, which is informative diagnostically. When the two predictions diverge, it typically occurs near high-gradient events (e.g., sharp peaks in producer energy or water rate). This behavior suggests that:

- State-based inference can inherit amplified errors from local state reconstruction near wells (where rates are most sensitive), whereas

- The latent-based observation head can act as a learned correction operator that maps the latent representation directly to well responses, effectively regularizing local errors without explicit post processing.

From an engineering perspective, this is valuable because field workflows often prioritize accurate well level observables for optimization and control, even when local grid block errors remain non-negligible. The results here show that the proposed model can deliver optimization-ready observables directly.

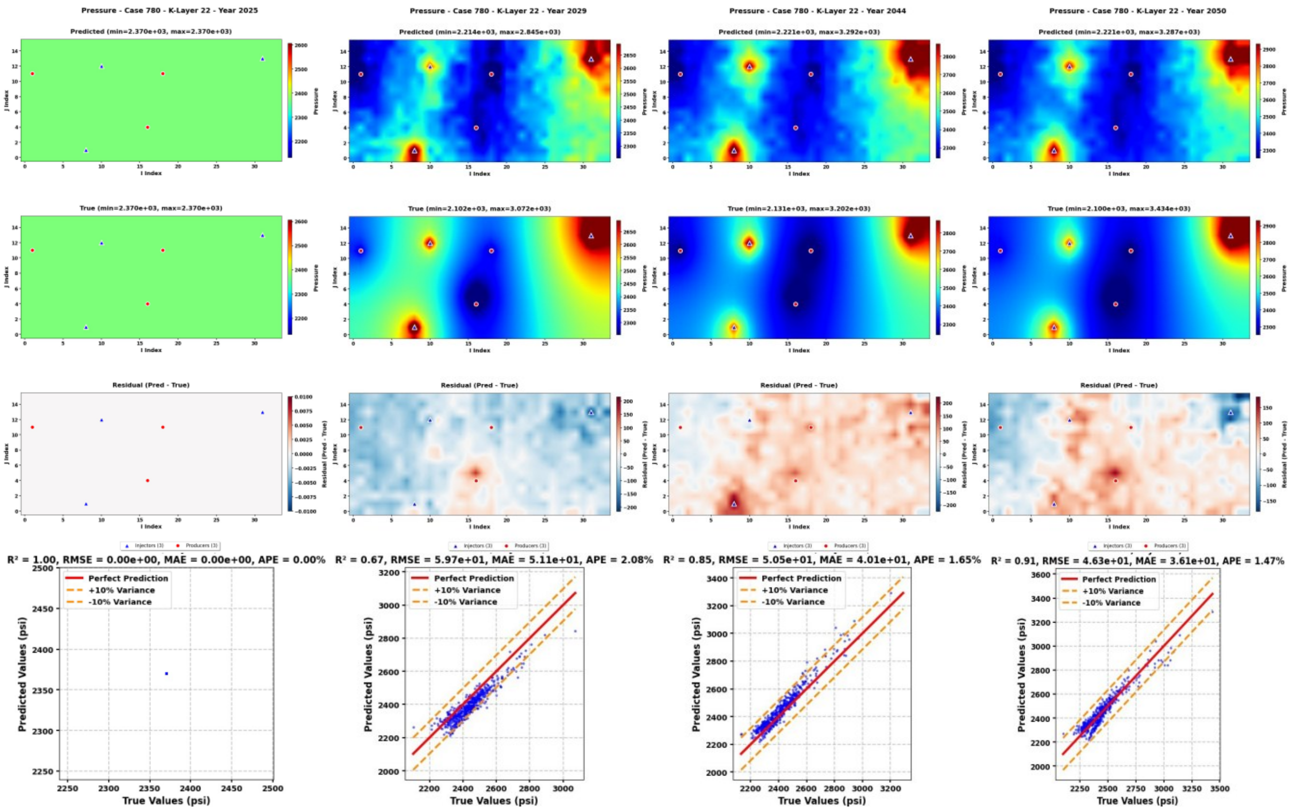

Figure 9: Predicted pressure field for Case 780 on Layer 22.

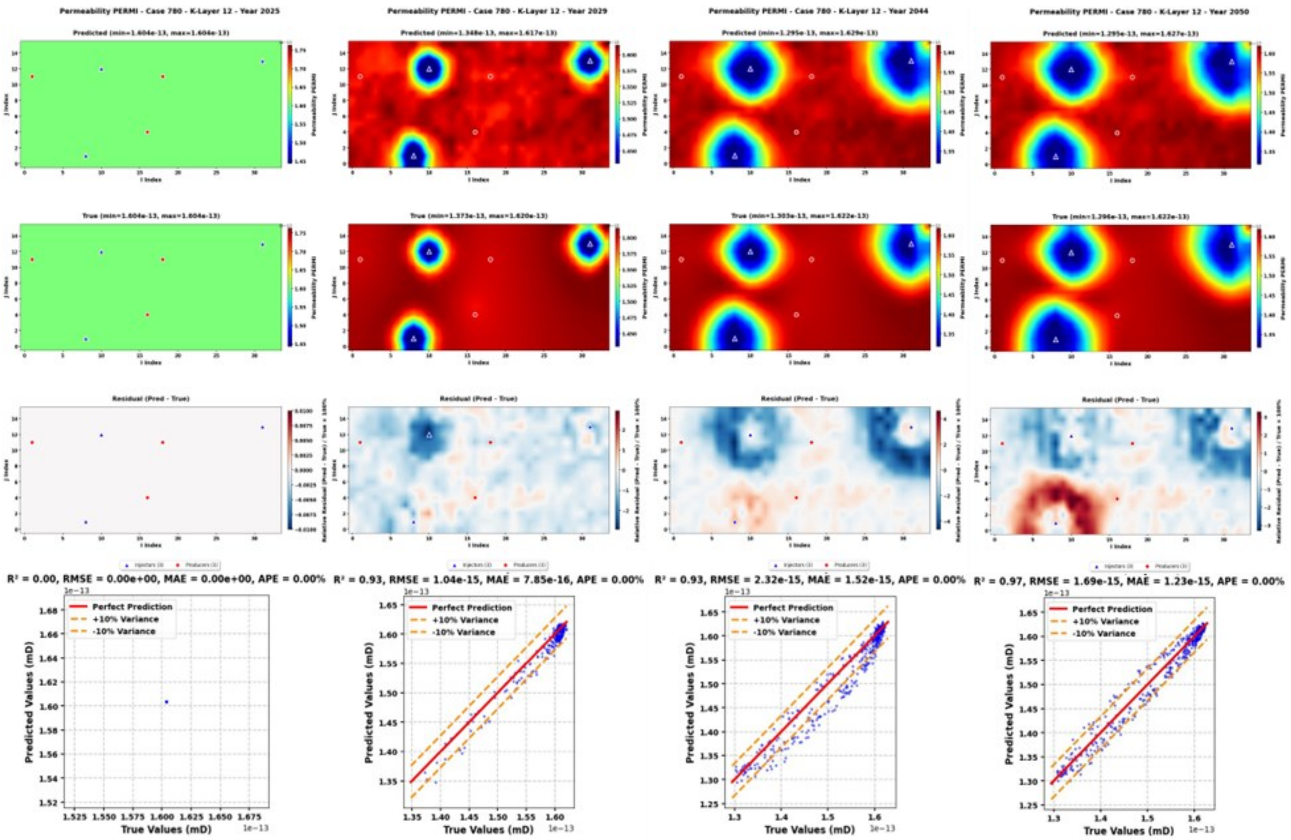

Figure 10: Predicted permeability field for Case 780 on Layer 12.

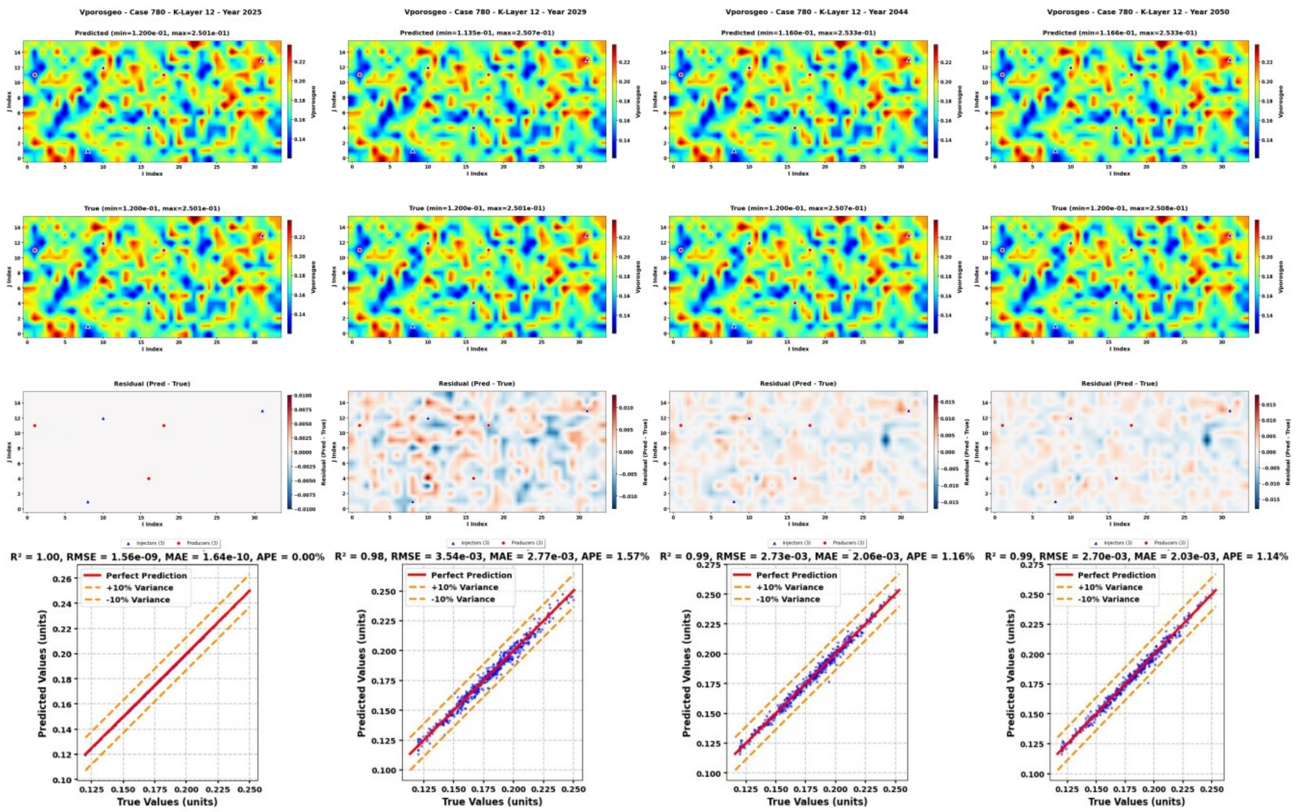

Figure 11: Predicted porosity field for Case 780 on Layer 12.

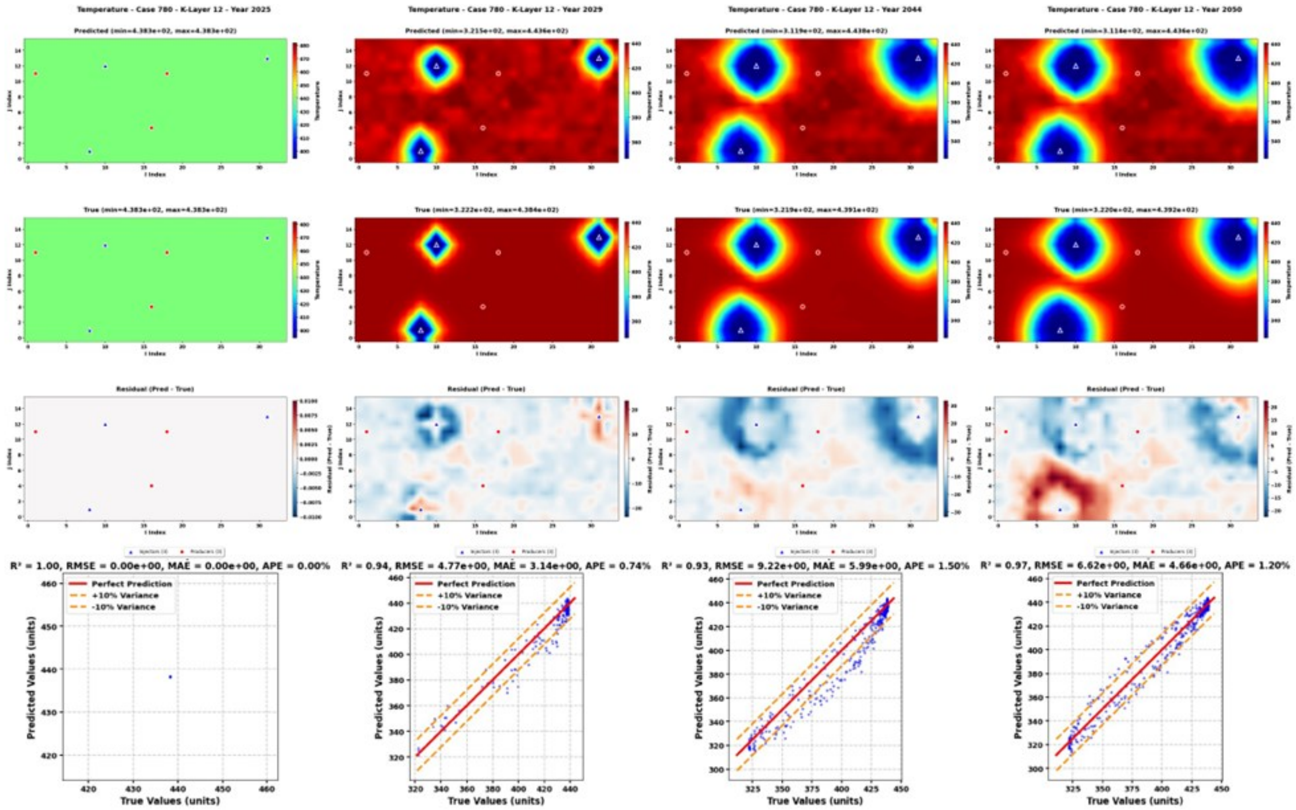

Figure 12: Predicted temperature field for Case 780 on Layer 12.

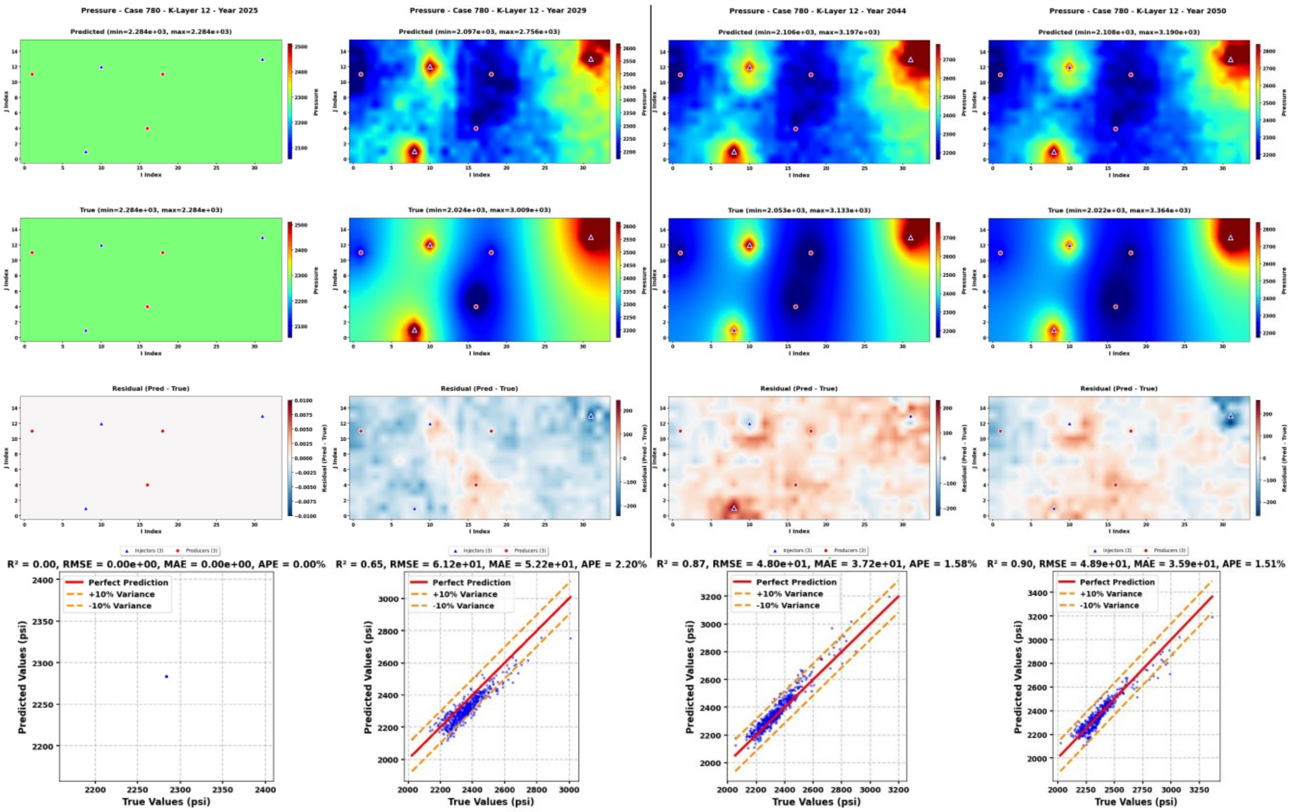

Figure 13: Predicted pressure field for Case 780 on Layer 12.

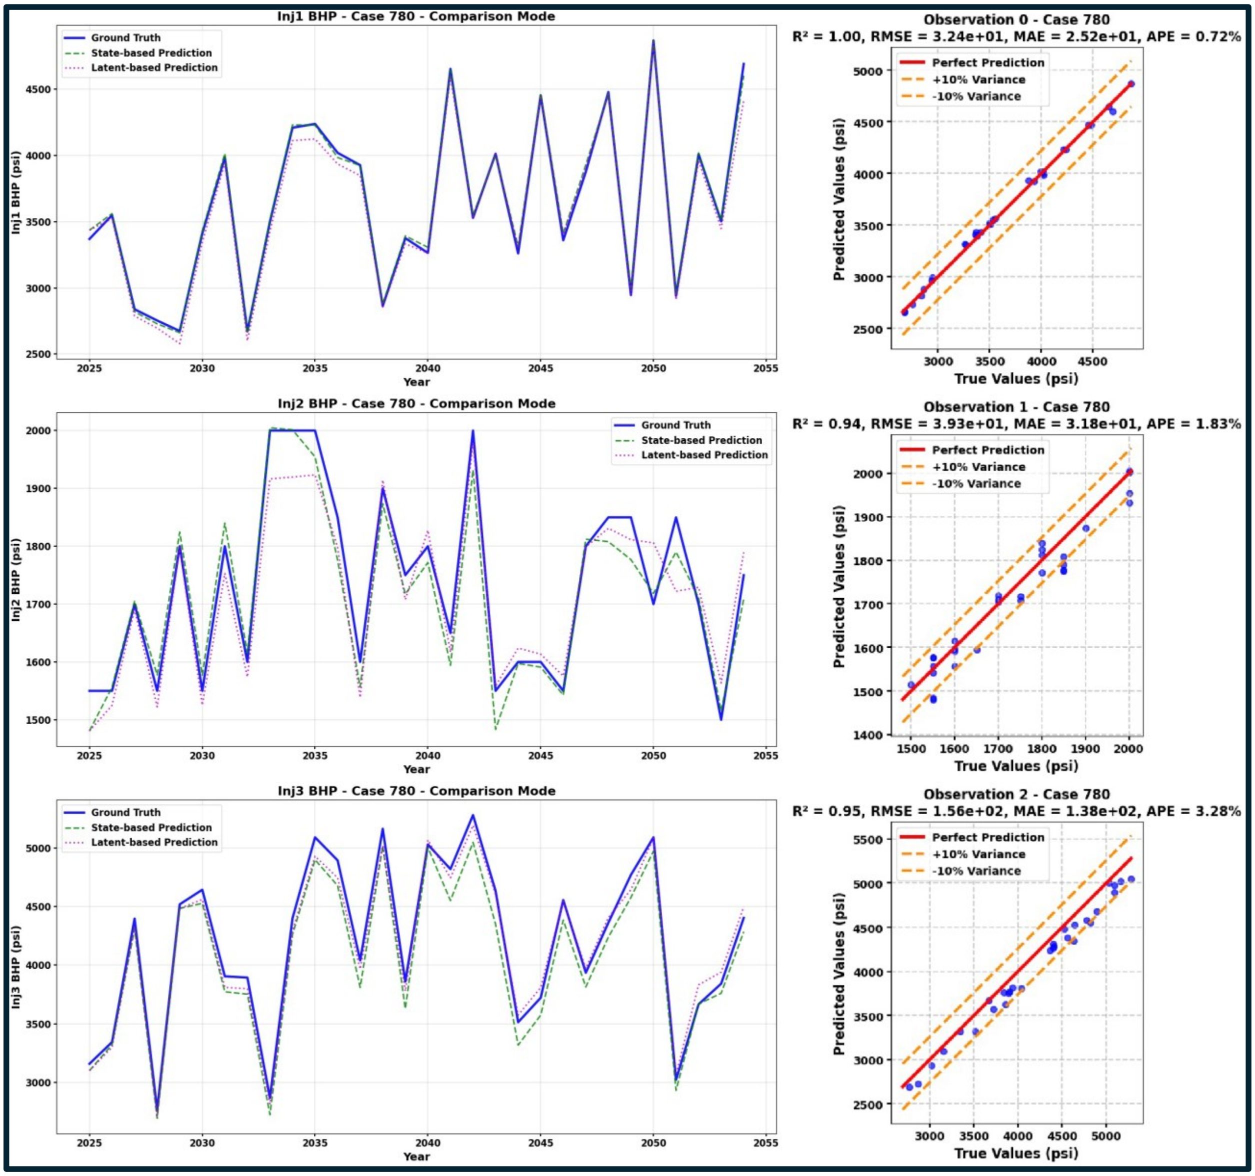

Figure 14: Direct observation prediction for Case 780: injector BHPs (time series and parity plots).

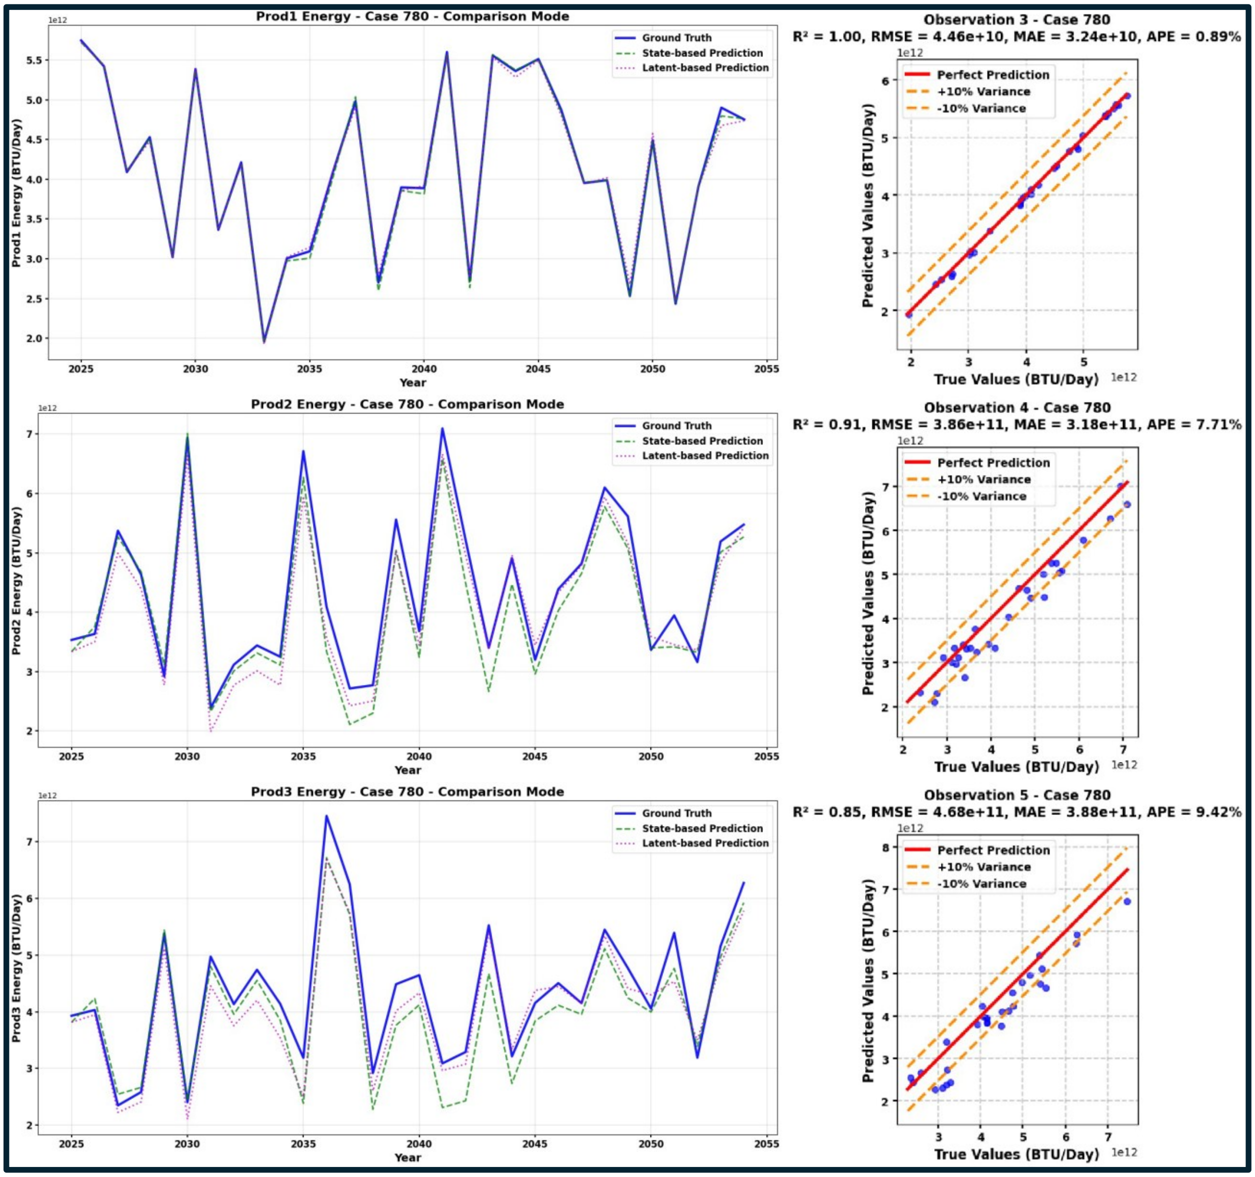

Figure 15: Direct observation prediction for Case 780: producer energy rates (time series and parity plots).

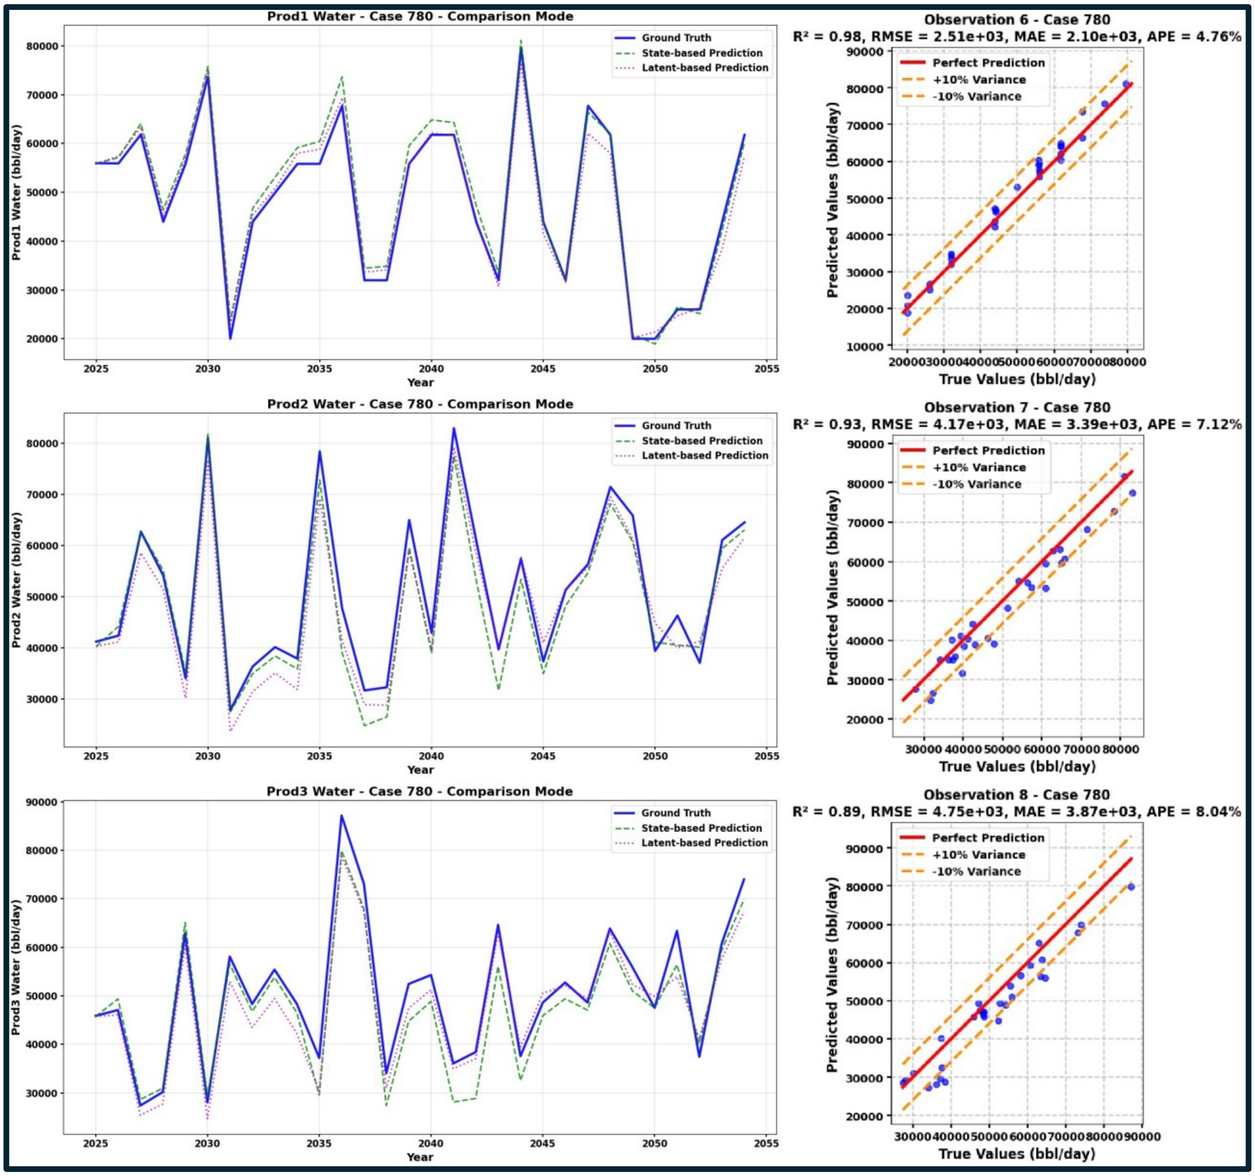

Figure 16: Direct observation prediction for Case 780: producer water rates (time series and parity plots).

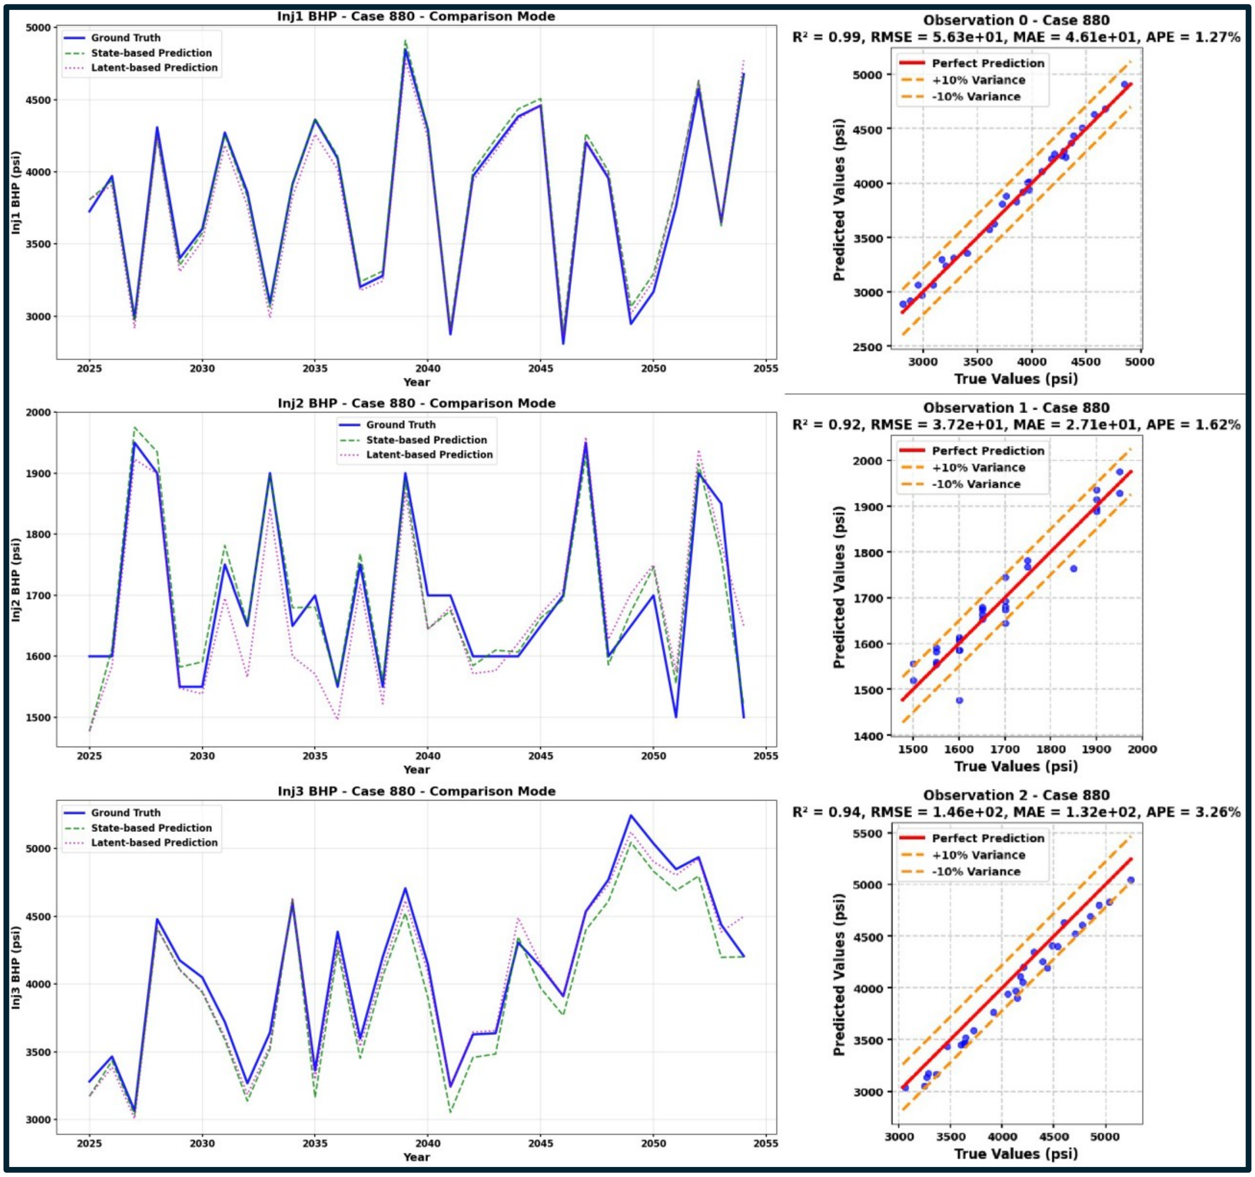

Figure 17: Direct observation prediction for Case 880: injector BHPs (time series and parity plots).

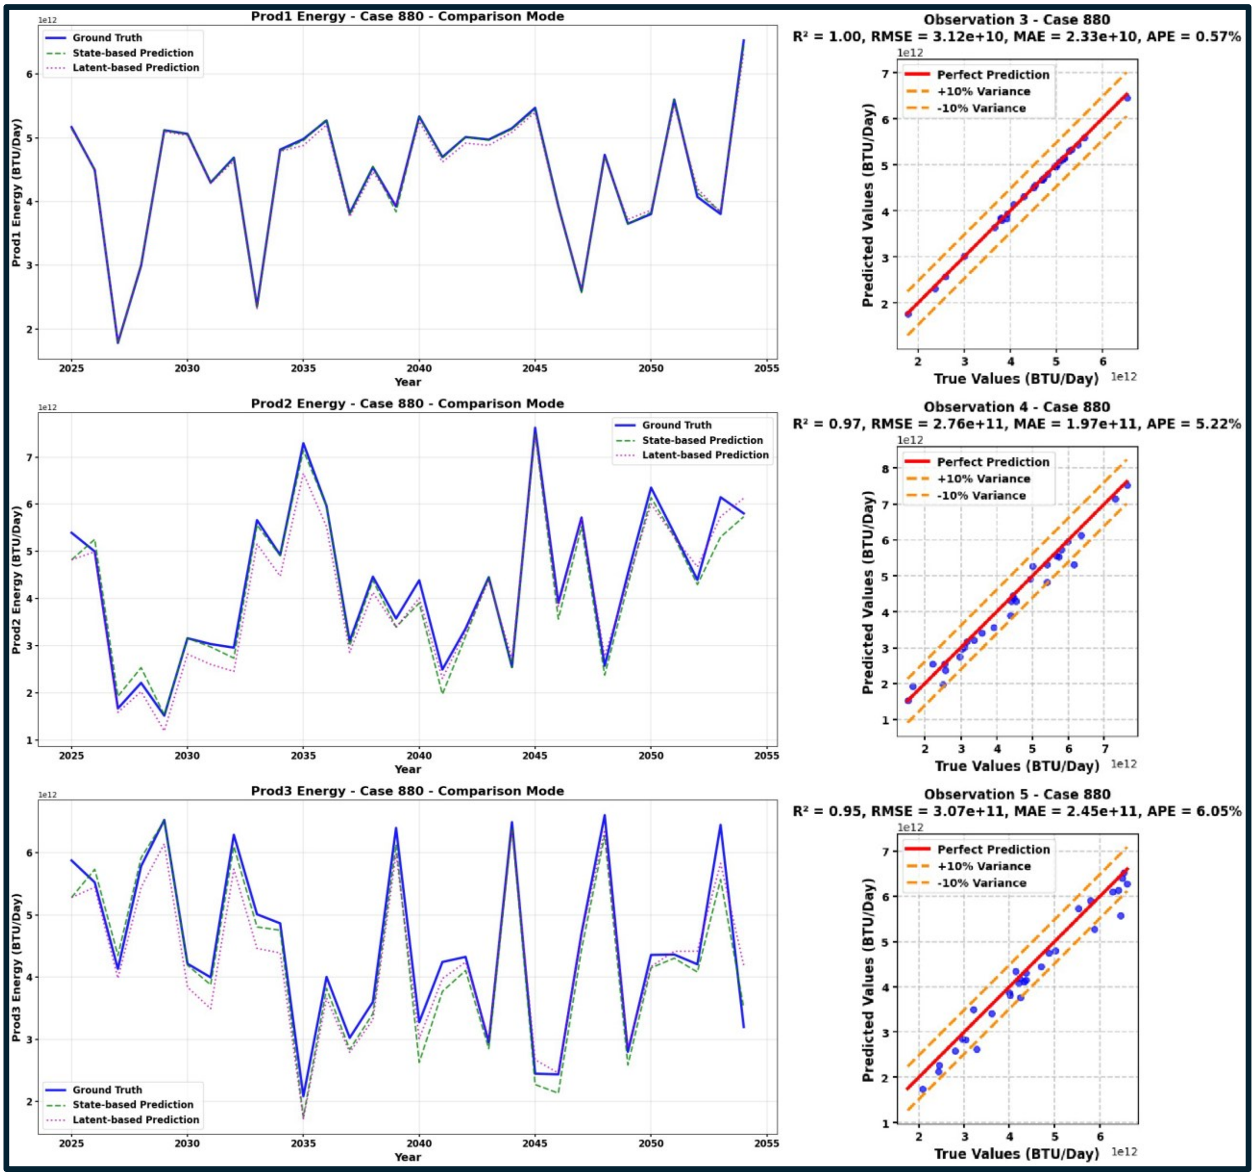

Prod1 Energy - Case 880 - Comparison Mode

Observation 3 - Case 880

R² = 1.00, RMSE = 3.12e+10, MAE = 2.33e+10, APE = 0.57%

Prod2 Energy - Case 880 - Comparison Mode

Observation 4 - Case 880

R² = 0.97, RMSE = 2.76e+11, MAE = 1.97e+11, APE = 5.22%

Prod3 Energy - Case 880 - Comparison Mode

Observation 5 - Case 880

R² = 0.95, RMSE = 3.07e+11, MAE = 2.45e+11, APE = 6.05%

Figure 18: Direct observation prediction for Case 880: producer energy rates (time series and parity plots).

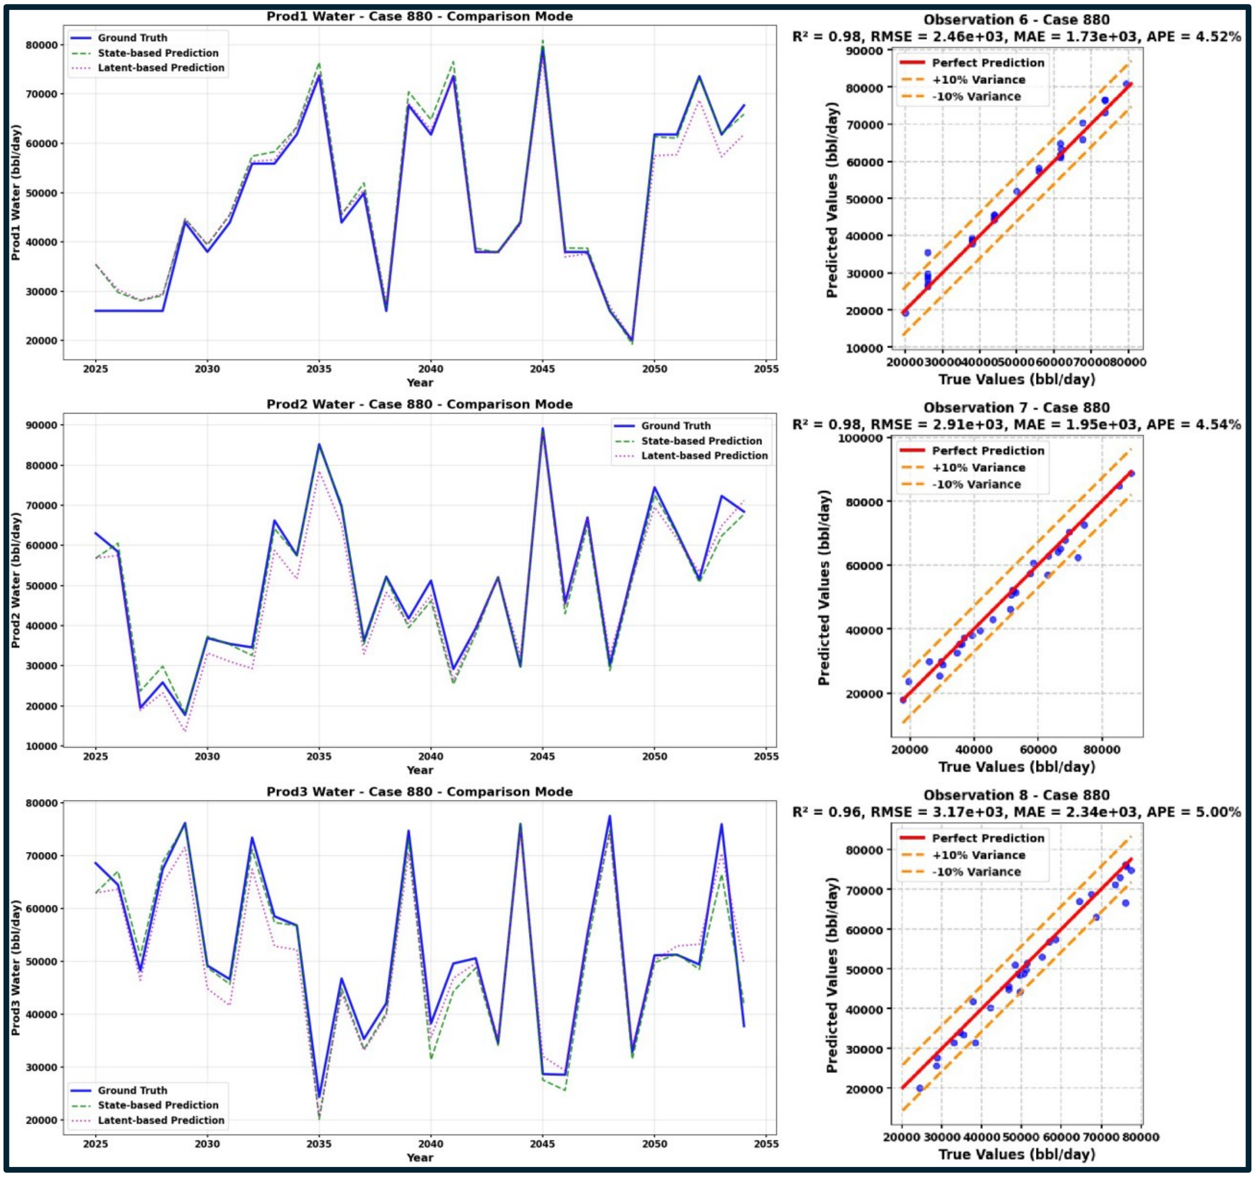

Figure 19: Direct observation prediction for Case 880: producer water rates (time series and parity plots).

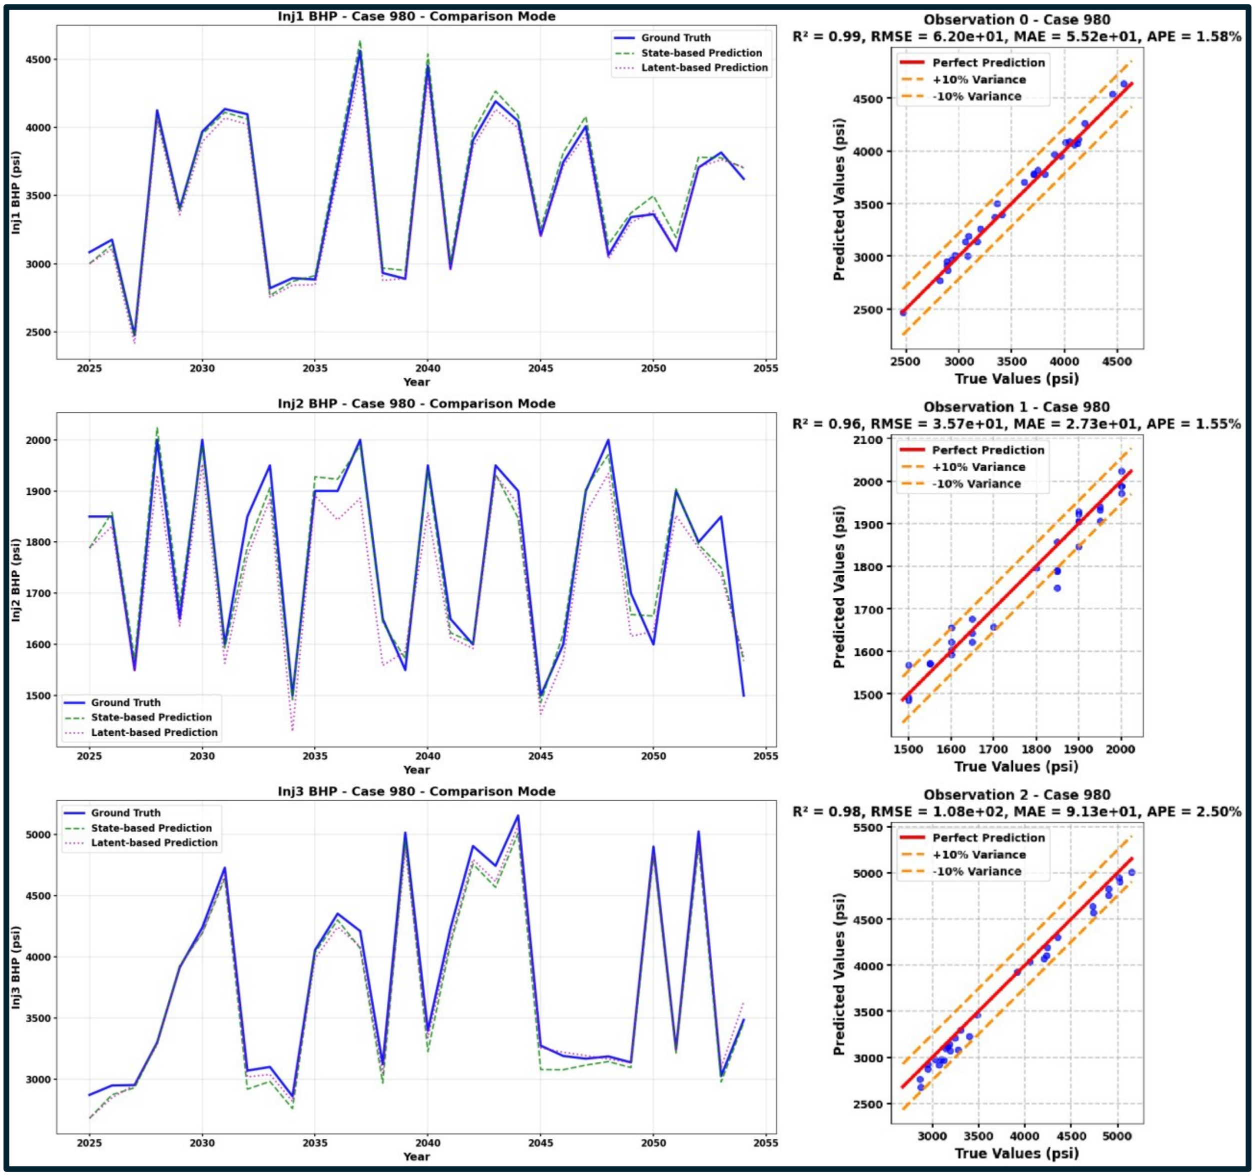

Figure 20: Direct observation prediction for Case 980: injector BHPs (time series and parity plots).

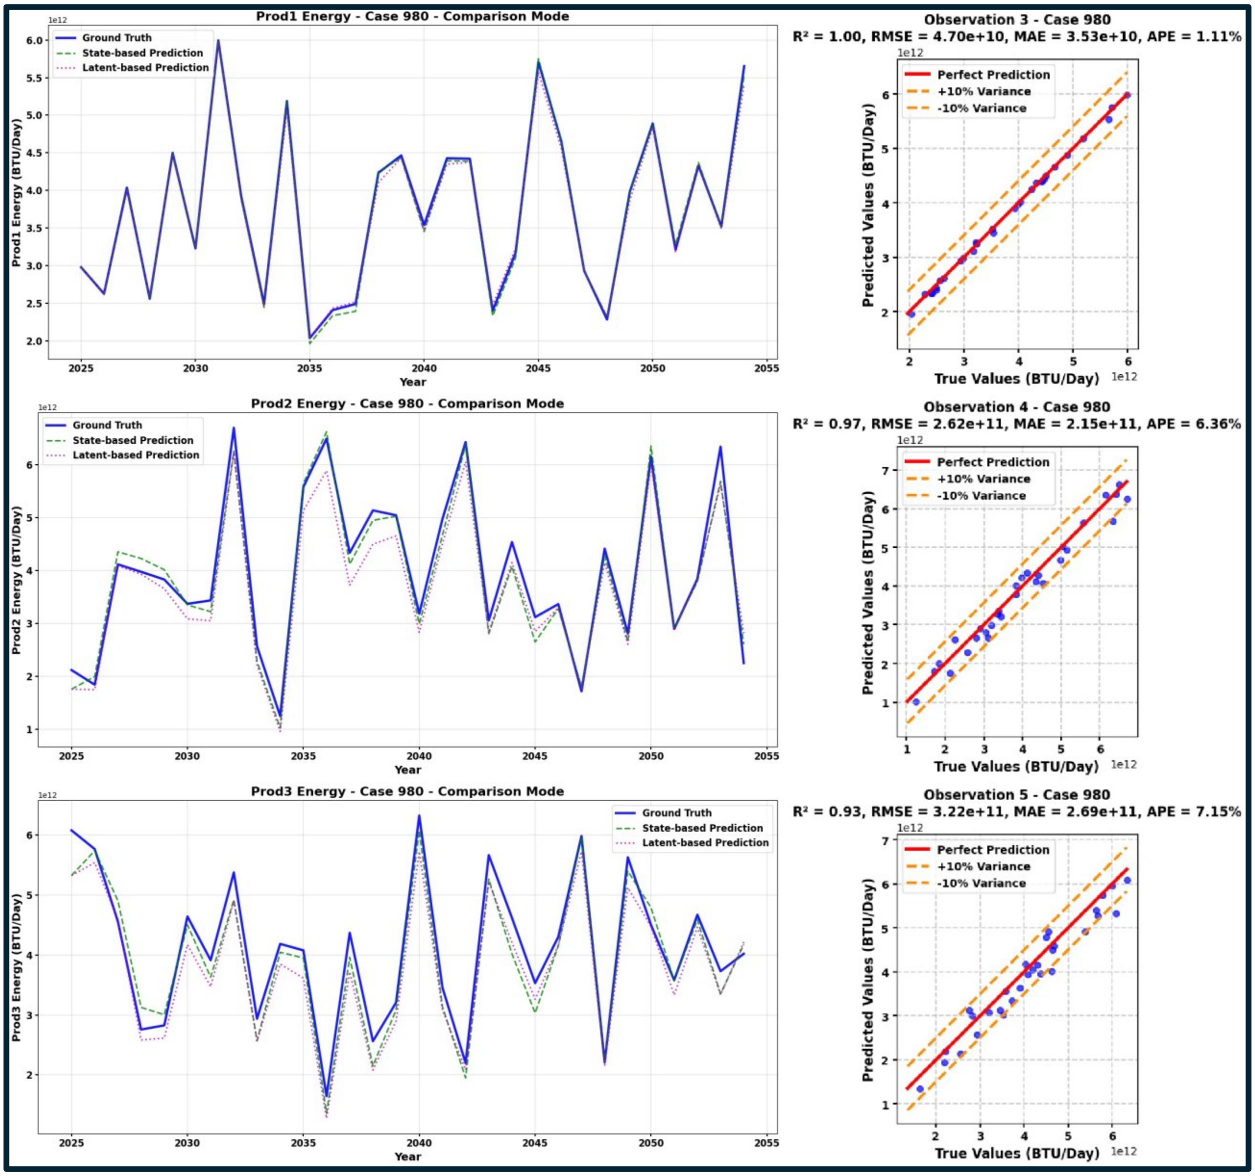

Figure 21: Direct observation prediction for Case 980: producer energy rates (time series and parity plots).

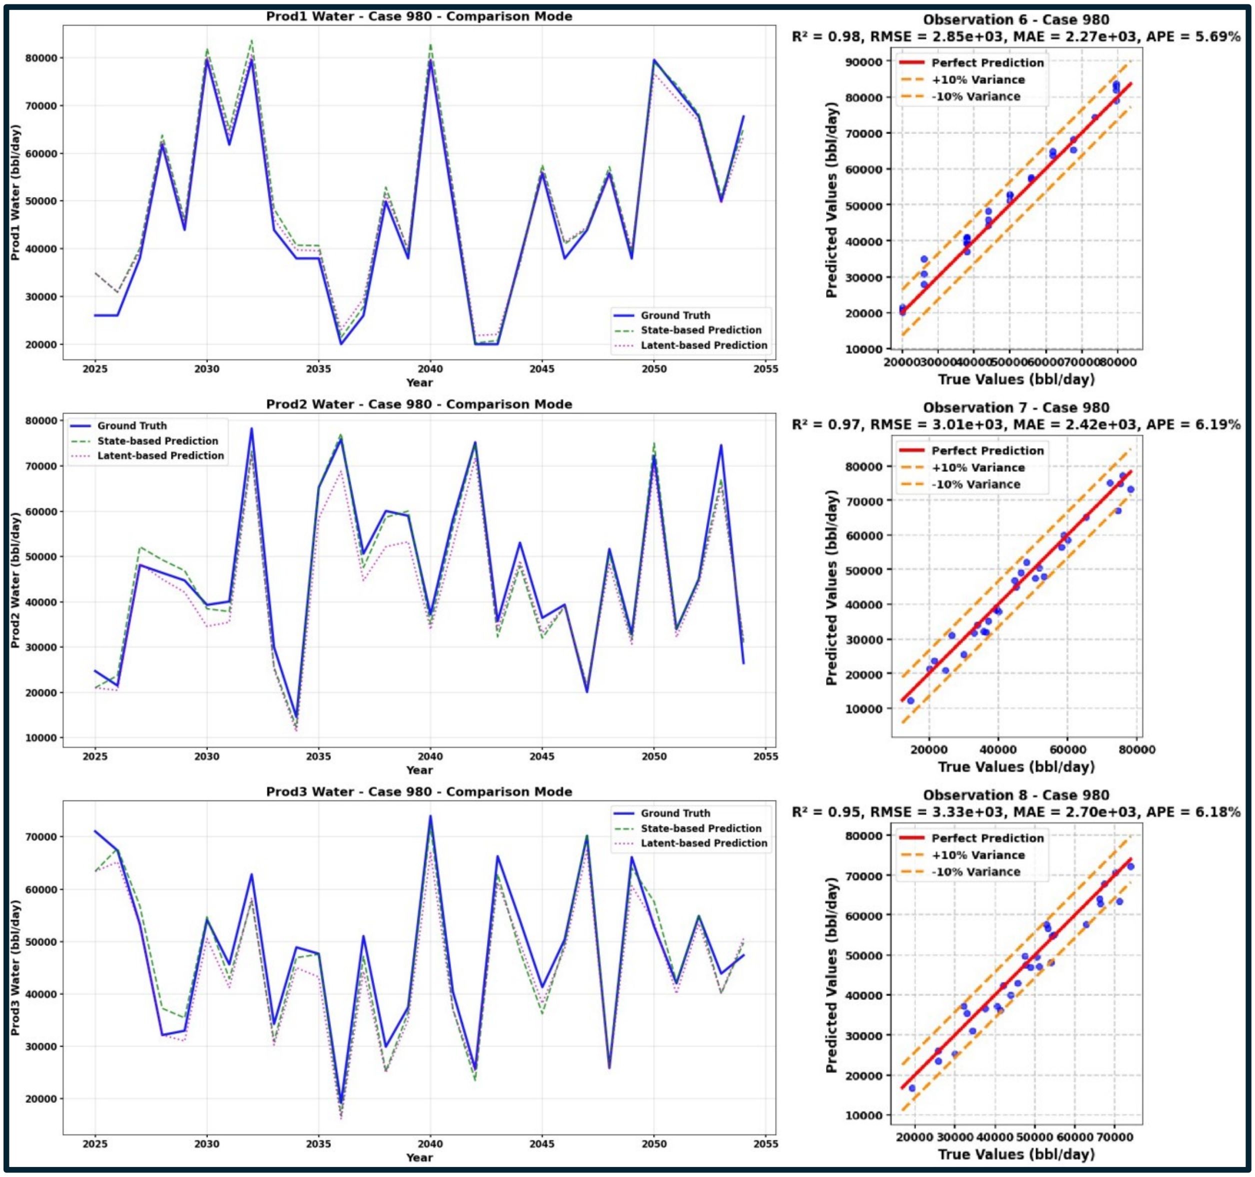

Figure 22: Direct observation prediction for Case 980: producer water rates (time series and parity plots).

4. LIMITATIONS AND MODEL CONSTRAINTS OF THE E2CO FRAMEWORK

Despite achieving substantial computational acceleration (on the order of 10^3-10^4 times) while maintaining high predictive accuracy, the E2CO framework remains subject to several intrinsic limitations (Chen et al., 2024a). These limitations arise from fundamental tradeoffs between representation capacity, control embedding, physical consistency, and computational scalability. In what follows, we systematically categorize these constraints into architectural, operational, and computational limitations, drawing connections to observations reported across prior E2CO and related reduced-order modeling studies (He et al., 2016; Nanga et al., 2021).

4.1 Architectural Limitations

4.1.1 Latent Dimensionality Trade-off Between Spatial Fidelity and Control Embedding

A central architectural limitation of the E2CO framework lies in the selection of the latent-space dimensionality dz. This parameter governs a fundamental trade-off between accurate high-dimensional state reconstruction and effective control embedding within the latent dynamics. On one hand, larger latent dimensions ($d_z \approx 64$–128) improve the expressive capacity of the encoder–decoder pair, enabling the model to retain high-frequency spatial features such as sharp thermal fronts, pressure gradients, and permeability-driven heterogeneity (Abdulkareem et al., 2025a,b). This

capability becomes increasingly critical for 3D geothermal and subsurface flow models. However, high-dimensional latent manifolds are inherently sparse, which degrades the ability of the linear transition operator ($z_{t+1} = A_t z_t + B_t u_t$) to learn globally stable, control-sensitive dynamics across the operational envelope.

Conversely, smaller latent dimensions (dz ≈ 8–16) promote compact representations that facilitate learning of the latent transition matrices At and Bt, thereby improving the accuracy and stability of observation predictions (e.g., well level production rates and bottom-hole pressures). However, excessive compression acts as a low-pass spatial filter, suppressing localized features essential for detecting breakthrough events and spatially resolved thermal behavior. This irreducible trade-off reflects a structural limitation of using a single shared latent space to simultaneously satisfy reconstruction, control, and observation objectives.

4.1.2 Autoregressive Error Accumulation in Long-Horizon Forecasting

In latent-only prediction mode, E2CO performs recursive autoregressive rollouts without intermittent correction from ground-truth states. As a result, even small approximation errors in the latent transition model accumulate over time. When the spectral radius of the learned transition matrix satisfies $\rho(A_t) \geq 1$, or when the linearization assumption breaks down in strongly nonlinear flow regimes, the latent trajectory may progressively diverge from the true physical manifold. Although multi-step training and temporal regularization reduce short-term drift, empirical evidence indicates that prediction accuracy degrades for long horizons ($T \gtrsim$ 20-30 years), particularly in highly heterogeneous systems. These accumulated latent errors ultimately propagate to the observation head, leading to biased long-term forecasts of energy production and reservoir response.

4.2 Operational and Topological Constraints

4.2.1 Assumption of Fixed Well Locations

The E2CO framework implicitly assumes fixed well locations throughout training and inference. The encoder learns spatial features conditioned on the specific pressure, temperature, and flow patterns induced by these fixed injector and producer configurations. Consequently, the learned latent representation is not location-invariant and cannot generalize to unseen well placements without retraining. This limitation restricts the applicability of E2CO to problems involving joint optimization of well placement and control. Extending the framework to variable well topologies would require abandoning grid-aligned convolutional encoders in favor of representations capable of handling variable geometry, such as graph neural networks, coordinate-based encodings, or operator-learning architectures.

4.2.2 Dependence on Large Training Datasets

As a purely data-driven reduced-order model, E2CO relies heavily on extensive high-fidelity simulation data to achieve robust generalization. Capturing coupled thermo-hydro-mechanical dynamics across multi-dimensional control spaces often requires on the order of 103 full-physics simulations. This requirement reflects the curse of dimensionality: expanding the operational envelope (e.g., wider injection rate ranges or additional controllable wells) increases the required training dataset size exponentially to maintain adequate state-space coverage. In such cases, the cost of dataset generation can offset the computational benefits of the surrogate model itself.

4.3 Computational and Training Constraints

4.3.1 Cost of Physics-Aware Loss Augmentation

Augmenting the E2CO training objective with physics-aware loss terms improves physical plausibility and long-horizon stability but introduces substantial computational overhead. Flux-consistency losses require computing spatial pressure gradients and transmissibilities across all grid interfaces, significantly increasing memory traffic and backpropagation cost. Similarly, well locationbased penalties require repeated indexing and extraction operations at specific spatial coordinates. Moreover, incorporating multiple auxiliary losses complicates the optimization land-scape and necessitates careful tuning of weighting coefficients to prevent dominance over the primary reconstruction objective. This tuning process further increases training time and reduces scalability, particularly for 3D models.

4.3.2 Sequence Length Versus Memory Scalability

Capturing long-term temporal dependencies requires training on extended sequences of states and controls. Increasing the sequence length n steps improves transition-model stability and reduces latent drift but introduces severe memory and computational constraints. Backpropagation through time requires storing the full computational graph across all timesteps. Given that each state is represented as a high-dimensional tensor (e.g., 4 × 34 × 16 × 25), GPU memory consumption scales linearly with n steps and batch size. This scaling often precludes training on full operational horizons (e.g., 30 years), forcing the use of truncated sequences or small batch sizes, which in turn can lead to noisy gradient estimates and slower convergence.

5. FUTURE RESEARCH DIRECTIONS BEYOND THE E2CO FRAMEWORK

The limitations of the E2CO framework serve as indicators of the inherent complexity of subsurface control problems. To advance this paradigm toward robust, field-deployable decision-making systems, several research directions are prioritized:

- Hierarchical and Corrective Latent Dynamics: To mitigate autoregressive drift, future models should employ hierarchical temporal representations that decouple short-term physical evolution from long-term trends.
- Corrective schemes that periodically anchor surrogate predictions to sparse high-fidelity simulations can further stabilize long-horizon forecasts.

- Multi-Objective and Adaptive Representations: Future architectures should transition toward structured la-tent spaces that explicitly separate state reconstruction, control sensitivity, and observation prediction. Task-conditioned subspaces offer a principled way to resolve the spatial–control trade-offs inherent in current formulations.

- Topology-Aware and Geometry-Invariant Modeling: Moving beyond fixed well location assumptions re-quires representations that encode well geometry and reservoir topology. Graph-based models and neural operator formulations (e.g., FNO or DeepONet) provide promising avenues for achieving geometry-invariant, control-aware surrogates.

- Hybrid Physics–Data Learning: Reducing data requirements calls for hybrid strategies such as selective physical regularization and conservation-aware latent transitions. These approaches offer a practical compromise between the data efficiency of E2CO and the physical consistency of PINNs.

- Scalable Physics-Aware Training: Future work should emphasize localized constraint enforcement and curriculum-based training. Introducing physical losses progressively can improve model robustness without the prohibitive computational costs typically associated with physics-informed training.

- Long-Horizon Control via Temporal Abstraction: Integrating E2CO surrogates with hierarchical reinforcement learning (HRL) enables policies to operate on compressed belief states. This temporal abstraction reduces memory requirements and facilitates adaptive, closed-loop control for complex management horizons.

Addressing these challenges requires a shift toward modular and physics-aware learning systems. Such hybrid architectures represent the critical pathway toward the reliable deployment of reduced-order models in real world geothermal and subsurface energy applications.

6. CONCLUSION

This work successfully extended the E2CO framework to geothermal reservoirs by incorporating fully THM processes. By explicitly modeling porosity and permeability as dynamic state variables alongside pressure and temperature, the THM-E2CO model accurately captured critical petrophysical feedbacks induced by thermal drawdown and mechanical loading. Preliminary results demonstrate that the proposed architecture achieves high-fidelity spatial and observation predictions with computational accelerations exceeding 30,000 times compared to full physics simulators. While challenges remain regarding long-horizon error accumulation and fixed well topologies, this framework provides a scalable foundation for real time geothermal performance prediction, uncertainty quantification, and closed loop reservoir management.

7. ACKNOWLEDGMENT

Portions of this research were conducted using the advanced computing resources provided by Texas A&M High-Performance Research Computing (HPRC). The authors would like to express their sincere gratitude to Dr. Eduardo Gildin and Dr. Jungang Chen for their invaluable guidance and technical support throughout this work.

Special recognition is given to the foundational work of Chen et al. regarding the E2CO framework. This study specifically acknowledges the use of the open-source 2D E2CO implementation (available at https://github. com/jungangc/CCS_E2CO-RL), which served as the structural basis for the extension into the 3D THM-E2CO architecture presented in this research.

REFERENCES

Abdulkareem, U., Adeyemi, A., and Onur, M.: History Matching Geological Models Using an Embed-To-Control Observe Deep-Learning Reservoir Surrogate, SPE Annual Technical Conference and Exhibition, SPE, (2025a). https://doi.org/10.2118/228029-MS.

Abdulkareem, U., Nguyen, Q. M., Adeyemi, A., and Onur, M.: Deep-Learning-Based History Matching Framework Using an Embed-To-Control Observe Reservoir Surrogate, SPE Annual Technical Conference and Exhibition, SPE, (2025b). https://doi.org/10.2118/223904-MS.

Adeyemi, A., and Onur, M.: Accelerated Deep Reinforcement Learning in Subsurface Reservoir Production Optimization, *SPE Journal*, 30(11), 6603–6621 (2025). https://doi.org/10.2118/230310-PA.

Atadeger, A., Sheth, S., Vera, G., Banerjee, R., and Onur, M.: Deep Learning-Based Proxy Models to Simulate Subsurface Flow of Three-Dimensional Reservoir Systems, Proceedings, ECMOR 2022, EAGE, The Hague, Netherlands / Online (2022). https://doi.org/10.3997/2214-4609.202244049.

Atadeger, A., Onur, M., and Banerjee, R.: Comparative Study of Kernel- and Deep Learning-Based Proxy Models for Nonlinearly Constrained Life-Cycle Production Optimization, *Geoenergy Science and Engineering*, 242, 213273 (2024). https://doi.org/10.1016/j.geoen.2024.213273.

Caruso Carter, G., Gasser, M., Albabawat, Z., Vijapurapu, R., Alharbi, E., Yehia, T., Gardella, B., Emera, R., Badejo, A., Wilson, P., Okoroafor, E. R., and Shor, R.: A Techno-Economic, Policy, and Regulation Assessment of an Integrated Campus Geothermal Energy System, (2025).

Caruso Carter, G., Eckstein, G., Yehia, T., Gasser, M., Okoroafor, E. R., and Shor, R. J.: An Economic, Policy and Regulatory Assessment of Repurposing Oil and Gas Wells for Geothermal Energy Applications, (2025).

Chen, J., Gildin, E., and Killough, J.: Multi-Step Embed-to-Control: A Novel Deep Learning-Based Approach for Surrogate Modelling in Reservoir Simulation, arXiv preprint arXiv:2409.09920 (2024a).

Chen, J., Gildin, E., and Kompantsev, G.: Optimization of Pressure Management Strategies for Geological CO_2 Storage Using Surrogate Model-Based Reinforcement Learning, *International Journal of Greenhouse Gas Control*, 138, 104262 (2024b). https://doi.org/10.1016/j.ijggc.2024.104262.

Chen, Y., Nguyen, Q. M., Abdulkareem, U., Onur, M., and Alpak, F. O.: Application of Embed-to-Control Observe Deep-Learning-Based Reservoir Surrogate in CO_2 Storage Monitoring and Multi-Objective Optimization, *SPE Journal*, 1–22 (2025). https://doi.org/10.2118/223914-PA.

Coutinho, E. J. R., Dall'Aqua, M., and Gildin, E.: Physics-Aware Deep-Learning-Based Proxy Reservoir Simulation Model Equipped With State and Well Output Prediction, *Frontiers in Applied Mathematics and Statistics*, 7, 651178 (2021). https://doi.org/10.3389/fams.2021.651178.

Ebaid, H., Gasser, M., Yehia, T., and Okoroafor, E. R.: Enhancing Oil Recovery in Heterogeneous Reservoirs with CO_2 Flooding: A Simulation Study, Proceedings, SPE/AAPG/SEG Carbon Capture, Utilization, and Storage Conference and Exhibition, SPE (2024).

Elsayed, T., and Okoroafor, E. R.: From Capturing to Sequestration: A Comprehensive Techno-Economic Analysis of CCS Projects, SPE Technical Paper (2024a). https://doi.org/10.2118/221395-MS.

Elsayed, T., and Okoroafor, E. R.: Integrated Thermo-Hydro-Chemo-Mechanical (THCM) Analysis for Evaluating CO_2 Plume Geothermal System Potential, SPE Technical Paper (2024b). https://doi.org/10.2118/220927-MS.

Falola, Y., Nair, G. G., Toms, J., and Osypov, K.: Optimizing Oil Field Development with Multi-Action Decision-Making in Deep Reinforcement Learning, SPE Technical Paper (2025a). https://doi.org/10.2118/227966-MS.

Falola, Y., Misra, S., Foster, J., and Bhatia, M.: Data-Driven Workflow for the Pre-Emptive Detection of Excess Water Producing Wells Drilled in Unconventional Shales, *Journal of Natural Gas Science and Engineering*, 103, 104609 (2022). https://doi.org/10.1016/j.jngse.2022.104609.

Falola, Y., Misra, S., and Calvo Nunez, A.: Rapid High-Fidelity Forecasting for Geological Carbon Storage Using Neural Operator and Transfer Learning, SPE Technical Paper (2023). https://doi.org/10.2118/216135-MS.

Falola, Y., Rathore, P. S., Nair, G. G., and Toms, J.: Application of Neural Operator Technique for Rapid Forecast of CO_2 Pressure and Saturation Distribution, SPE Technical Paper (2024). https://doi.org/10.4043/35230-MS.

Falola, Y., Churilova, P., Liu, R., Huang, C.-K., Delgado, J. F., and Misra, S.: Generating Extremely Low-Dimensional Representation of Subsurface Earth Models Using Vector Quantization and Deep Autoencoder, *Petroleum Research*, 10(1), 28–44 (2025b). https://doi.org/10.1016/j.ptlrs.2024.07.001.

Falola, Y., Misra, S., and Nunez, A.: Leveraging Transfer Learning for Reliable CO_2 Storage Forecasting Across Diverse Operational Conditions, *Fuel*, 403, 135990 (2026). https://doi.org/10.1016/j.fuel.2025.135990.

Gasser, M., Rehg, D., Oesterberg, M., and Meehan, N.: Harnessing AI and Computer Vision for Efficient Geothermal Field Exploration and Prospect Evaluation, URTeC, Houston, TX (2025). https://doi.org/10.15530/urtec-2025-4265623.

Ghaly, E., and Shor, R. J.: Bit-Rock Interaction for PDC Drill Bits: A Review and an Open-Source Library, SPE Technical Paper (2025). https://doi.org/10.2118/223715-MS.

He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, Proceedings, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV (2016).

Kompantsev, G., Gildin, E., and Rabbani, H.: Novel Transfer Learning Workflow Allowing Flexible Well Configurations for Physics-Aware Deep-Learning-Based Proxy Reservoir Simulation Models, Proceedings, ECMOR 2024, EAGE, Oslo, Norway (2024). https://doi.org/10.3997/2214-4609.202437046.

Le, H., and Tao, J.: Hierarchical Autoencoder-Based Lossy Compression for Large-Scale High-Resolution Scientific Data, *Computing & AI Connect*, 1(1) (2024). https://doi.org/10.69709/CAIC.2024.193132.

Nanga, S., Bawah, A. T., Acquaye, B. A., Billa, M.-I., Baeta, F. D., Odai, N. A., Obeng, S. K., and Nsiah, A. D.: Review of Dimension Reduction Methods, *Journal of Data Analysis and Information Processing*, 9(3), 189–231 (2021). https://doi.org/10.4236/jdaip.2021.93013.

Nguyen, Q. M., and Onur, M.: Reservoir Forecast and Production Optimization Under Geological Uncertainty Using an Embed-to-Control Deep-Learning Surrogate Model, Proceedings, ECMOR 2024, EAGE (2024a). https://doi.org/10.3997/2214-4609.202437090.

Nguyen, Q. M., and Onur, M.: A Deep-Learning-Based Reservoir Surrogate for Performance Forecast and Nonlinearly Constrained Life-Cycle Production Optimization Under Geological Uncertainty, Proceedings, SPE Europe Energy Conference and Exhibition, Turin, Italy (2024b). https://doi.org/10.2118/220002-MS.

Nguyen, Q. M., Onur, M., and Alpak, F. O.: Embed-to-Control-Based Deep-Learning Surrogate for Robust Nonlinearly Constrained Life-Cycle Production Optimization: A Realistic Deepwater Application, SPE Technical Paper (2024). https://doi.org/10.2118/220783-MS.

Noh, K., and Swidinsky, A.: Geophysical Control of Geologic Carbon Storage Using Deep Reinforcement Learning: Sensitivity to Multigeophysical Noise and to the Uncertainty of Digital Twins, *Geophysics*, 90(4), H1–H14 (2025). https://doi.org/10.1190/geo2024-0402.1.

Okoroafor, E. R., Kang, Y., Elsayed, T., and Indro, A.: Enabling CO₂ Plume Geothermal: Key Geological, Reservoir, and Operational Considerations, ARMA Conference Proceedings (2023). https://doi.org/10.56952/ARMA-2023-0816.

Okoroafor, E. R., and Horne, R. N.: Fracture Aperture Anisotropic Effects on Field-Scale Enhanced Geothermal System Thermal Performance, *Geothermics*, 118, 102897 (2024). https://doi.org/10.1016/j.geothermics.2023.102897.

Olasolo, P., Juárez, M. C., Morales, M. P., D’Amico, S., and Liarte, I. A.: Enhanced Geothermal Systems (EGS): A Review, *Renewable and Sustainable Energy Reviews*, 56, 133–144 (2016). https://doi.org/10.1016/j.rser.2015.11.031.

Oliver, G., Gasser, M., Oesterberg, M., Pennino, A., and Craig, P.: Mount Augustine: Exploring for a new geothermal resource in south-central Alaska, *Geothermal Resources Council Transactions*, 48, 1435–1462 (2024).

Tayyib, D. M., Elsayed, T., and Okoroafor, E. R.: Techno-Economic and Life-Cycle Assessment of Subsurface Hydrogen and Synthetic Geothermal Storage Technologies, *Journal of Energy Storage*, 118, 116020 (2025). https://doi.org/10.1016/j.est.2025.116020.

Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M.: Embed to Control: A Locally Linear Latent Dynamics Model for Control from Raw Images, Proceedings, Advances in Neural Information Processing Systems (NeurIPS 28), Montreal, Canada (2015).

Yehia, T. A., Freestone, B., Nair, G. G., Falola, Y., Ailneni, R. R., Rathore, P. S., Toms, J., and Meehan, D. N.: An Ensemble-Based Automated Workflow for Decline Curve Analysis and Probabilistic Estimated Ultimate Recovery Estimation in Shale Gas Wells, *SPE Journal*, 30(11), 6638–6655 (2025). https://doi.org/10.2118/230302-PA.

Yehia, T., Abdelhafiz, M. M., Hegazy, G. M., Elnekhaily, S. A., and Mahmoud, O.: A Comprehensive Review of Deterministic Decline Curve Analysis for Oil and Gas Reservoirs, *Geoenergy Science and Engineering*, 211775 (2023). https://doi.org/10.1016/j.geoen.2023.211775.

Yehia, T., Gasser, M., Ebaid, H., Meehan, N., and Okoroafor, E. R.: Comparative Analysis of Machine Learning Techniques for Predicting Drilling Rate of Penetration (ROP) in Geothermal Wells: A Case Study of the FORGE Site, *Geothermics*, 121, 103028 (2024). https://doi.org/10.1016/j.geothermics.2024.103028.