TRAJECTORY TRANSLATOR FOR AMORTIZED FIELD RECONSTRUCTORS UNDER OBSERVATION-INTERFACE SHIFT

Anonymous authors
Paper under double-blind review

ABSTRACT

Amortized field reconstruction models learn to infer complete physical fields from partial observations under the sensing configurations represented during training. When training is restricted to a fixed sensing trajectory, the learned reconstruction map may become specialized to its observation interface. At deployment, however, novel or adaptively selected trajectories can introduce observation-interface distribution shift, constituting an out-of-distribution (OOD) setting. We introduce **Trajectory Translator (TT)**, an add-on module that maps deployment observations to virtual observations on the original observation interface, guiding field reconstruction without modifying the frozen reconstructor. Under a finite-state Bayesian formulation, we characterize when bounded virtual observations preserve deployment evidence exactly. For deterministic reconstructors, we relate squared reconstruction risk to approximation of the conditional mean and evaluate TT through empirical reconstruction errors. Experiments across physical-field benchmarks show that TT improves reconstruction under the evaluated observation-interface shifts while keeping the underlying reconstructors frozen, reducing relative reconstruction error by up to 88.9% for the Transformer on the plate-strain benchmark. These results demonstrate that virtual observations on the original interface can effectively steer a frozen reconstructor toward more accurate field estimates. At deployment, TT translates each observation sequence or block in a single forward pass. In the measured offline settings, TT adds only 0.39–0.57 ms per complete sequence, introducing little additional computational overhead. Code is available at `https://anonymous.4open.science/r/Trajectory-Translator-3D38/`.

1 INTRODUCTION

In real-time applications, such as closed-loop flow control(Nonomura et al., 2025) and structural health monitoring(Roy et al., 2022), physical-field reconstruction tasks require reconstructors to efficiently process continuous sensor observations and promptly update their field estimates. A common approach is to use **amortized field reconstructors**, which learn to map partial observations to complete physical fields during training, enabling fast inference at deployment. For example, SHRED uses an LSTM to reconstruct high-dimensional states from sparse sensor histories (Williams et al., 2024); FLRONet uses operator learning to reconstruct flow fields from sparse measurements (Dang & Nguyen, 2026). Sequential flow matching further enables recursive state estimation by using incoming observations to update the previous belief through a learned probability flow (Huang et al., 2026). These approaches shift much of the computational burden to training, enabling fast inference at deployment.

However, amortized field reconstructors may become specialized to observation configurations encountered during training, including sensor counts, ordering, and sensing trajectories. At deployment, different observation interfaces may reduce reconstruction accuracy, exhibiting the effects of **observation-interface out-of-distribution (OOD) shifts**. Some existing works highlight this dependence: Mobile-SHRED incorporates the chosen observation trajectories into its training data (Ebers et al., 2024), while DiffSRDA examines the effects of deployment-time sensor-layout changes on reconstruction (Ayapilla et al., 2026).

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.

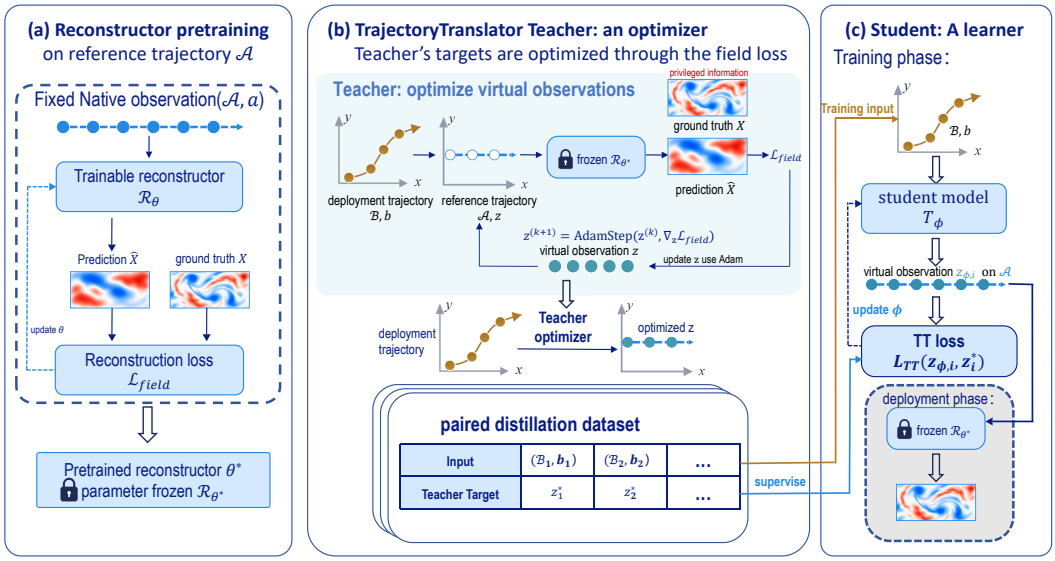

Figure 1: Trajectory Translator. A teacher optimizes virtual observations using full-field supervision; a student predicts them at deployment to drive the frozen reconstructor.

Moreover, in adaptive sensing, we need to assess the observation value of a sensing trajectory for the field reconstructor(Deng et al., 2026). However, observation-interface OOD errors confound this assessment. Figure 2 shows that each recursively updated conditional flow-matching (RU-CFM) reconstructor performs best on its matched trajectory. This motivates distinguishing the value of the observations from the reconstructor's familiarity with their interface.

Adapting to new sensing configurations can involve retraining the reconstructor, adding trainable internal branches as in ControlNet (Zhang et al., 2023), or performing iterative field refinement at deployment. However, these approaches require retraining, additional internal branches, or iterative computation at deployment.

We introduce *Trajectory Translator* (TT), an add-on module that maps deployment configurations and measurements to **virtual observations on the original interface**, keeping the reconstructor frozen. Rather than reproducing true measurements on the original trajectory, TT constructs inputs that convey deployment evidence to the existing model. These inputs steer reconstruction without changing model parameters (Appendix B.2).

Our contributions are:

- **Add-on module TT.** We introduce TT as an attempt to address observation-interface generalization across multiple reconstructor architectures, without retraining or fine-tuning the underlying reconstructors.
- **Theoretical analysis.** We characterize bounded evidence-preserving translation in finite-state Bayesian models and relate squared reconstruction risk to conditional-mean approximation for deterministic reconstructors under a fixed evidence–target distribution.
- **Evaluation on multiple physical benchmarks.** We validate exact posterior matching in a finite-state toy experiment and demonstrate that TT improves reconstruction in most evaluated settings across multiple physical-field benchmarks and reconstructor architectures. These results demonstrate effective input control of frozen reconstructors through their original observation interfaces.

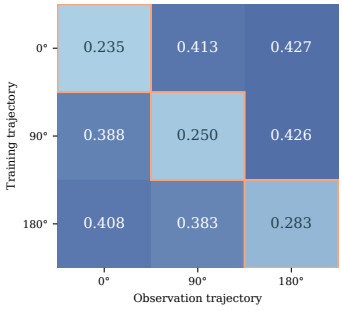

Figure 2: RU-CFM reconstruction error (lower is better). RU-CFM reconstruction error (↓). Rows/columns denote training/deployment trajectories; outlined cells mark each row's lowest error.

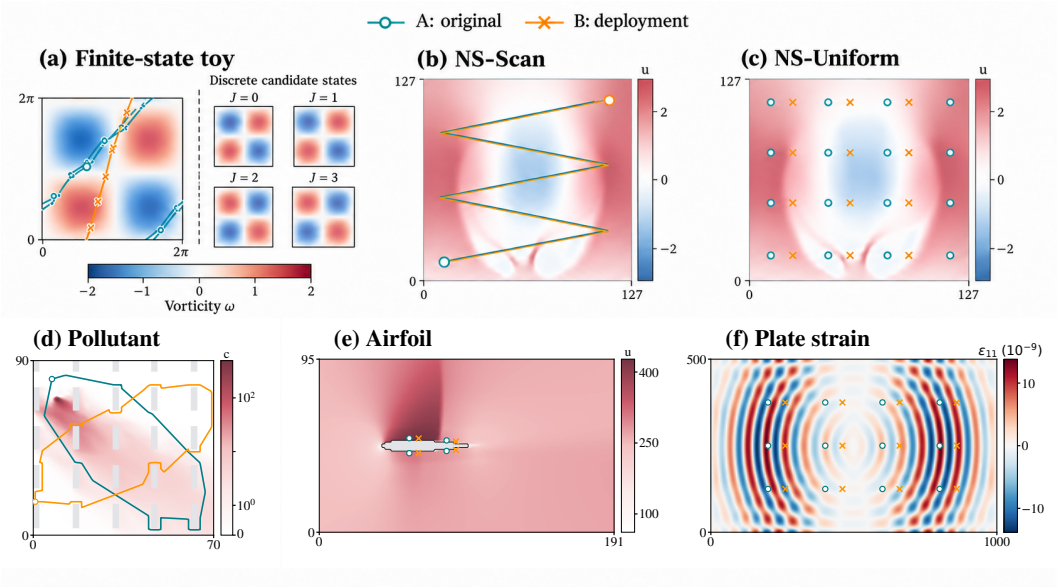

Figure 3: **Experimental visualization**: Panels (a,b,c,e,f) show representative snapshots of time-dependent fields; panel (d) shows a top-down horizontal slice at $z = 2\,\mathrm{m}$ of a three-dimensional time-averaged pollutant field. The original interface $\mathcal{A}$ and deployment interface $\mathcal{B}$ show the complete sensing trajectories over the observation sequence, rather than only the positions at the displayed frame; stationary sensors appear as fixed points. In (d), the three-dimensional trajectories are projected onto the horizontal plane. All backgrounds are ground-truth fields, not reconstructions.

2 Motivation and Method

2.1 Problem Setting

Let $X_t \in \mathcal{X}$ denote the true state of a time-varying physical field at time t. The reconstructor R_θ is trained with a fixed sensing trajectory $\mathcal{A} = (A_0, A_1, \ldots, A_T)$, with the measurements

$$a_t = y_t^{\mathcal{A}} = \mathcal{H}_{A_t}(X_t) + \epsilon_t^{\mathcal{A}}, \qquad a = (a_0, \ldots, a_T). \tag{1}$$

Here, "fixed" refers to the predefined sensing schedule, not to a stationary sensor: the locations A_t may vary with time. Writing $X = (X_0, \ldots, X_T)$ for the field sequence, we suppress time indices in the trajectory-level reconstruction notation:

$$\hat{X}^{\mathcal{A}} = F_{\mathcal{A}}(a) := R_\theta(\mathcal{A}, a). \tag{2}$$

The reconstructor parameters θ remain frozen at deployment.

At deployment, suppose the reconstructor instead uses a trajectory $\mathcal{B} = (B_0, B_1, \ldots, B_T)$, with measurements

$$b_t = y_t^{\mathcal{B}} = \mathcal{H}_{B_t}(X_t) + \epsilon_t^{\mathcal{B}}, \qquad b = (b_0, \ldots, b_T). \tag{3}$$

When direct evaluation on the new interface is defined, we denote its result by $\hat{X}^{\mathcal{B}} = F_{\mathcal{B}}(b) := R_\theta(\mathcal{B}, b)$. Even if the reconstructor takes positional encodings as inputs, directly feeding observations from an unseen trajectory $\mathcal{B}$ may introduce observation-interface OOD and degrade reconstruction accuracy. Meanwhile, its prediction results may fail to faithfully reflect the true conditional distribution of the physical field given the observations along trajectory $\mathcal{B}$.

$$E = (\mathcal{B}, b), m(E) = \mathbb{E}[X|E]. \tag{4}$$

Here, E denotes the evidence provided by the deployment trajectory and its measurements, and $m(E)$ denotes the corresponding conditional mean, which is the ideal estimator under squared error (Bishop & Nasrabadi, 2006, Section 1.5.5). The resulting reconstruction error may therefore reflect both limited observational information and the reconstructor's failure to exploit the available evidence.

2.2 Design of the Translation Target

To exploit deployment observations without modifying the reconstructor, we seek to translate them into virtual observations on the interface with which the reconstructor is most familiar, namely the original observation trajectory $\mathcal{A}$. We choose to translate the deployment trajectory $\mathcal{B}$ and observations b to virtual observations z on trajectory $\mathcal{A}$. Accordingly, our **Trajectory Translator (TT)** learns a mapping from $(\mathcal{B}, b, \mathcal{A})$ to $z^{\mathcal{A}}$, parameterized by ϕ:

$$z^{\mathcal{A}} = T_\phi(\mathcal{B}, b; \mathcal{A}), \qquad \widehat{X}^{TT} = R_\theta(\mathcal{A}, z^{\mathcal{A}}). \tag{5}$$

In our design, we do not aim to learn the mapping $(\mathcal{B}, b) \mapsto a$. Instead, we construct virtual observations z, without inherently requiring them to be close to a. Nevertheless, the frozen reconstructor processes z as observations on trajectory $\mathcal{A}$, even when they differ from the true measurements. Inspired by adaptive control (Mareels et al., 1987), when the reconstructor itself incurs reconstruction error, appropriate local input adjustments can reduce this error under the stated conditions. Under this formulation, the virtual observations z serve as control inputs to the frozen reconstructor, rather than estimates of the true observation values, as formalized in Proposition 1.

Proposition 1 (Local input compensation)**.** *Given a target field X and an initial input z_0, suppose that $F_{\mathcal{A}}$ is continuously differentiable in a neighborhood of z_0. Define*

$$e = F_{\mathcal{A}}(z_0) - X, \qquad J = DF_{\mathcal{A}}(z_0).$$

Assume that $J^\top e \neq 0$ and that z_0 lies in the interior of the admissible input set. Then, for all sufficiently small $\eta > 0$, the input

$$z_\eta = z_0 - \eta J^\top e$$

is admissible and satisfies

$$\|F_{\mathcal{A}}(z_\eta) - X\|_2^2 < \|F_{\mathcal{A}}(z_0) - X\|_2^2. \tag{6}$$

The proof is provided in Appendix B.2.

2.3 Teacher–Student Framework

Our objective is to learn a mapping from the deployment configuration $\mathcal{B}$ and observations b to virtual observations z that improve full-field reconstruction. During training, the ground-truth physical field X is available, although it will not be available at deployment. We exploit this *privileged information* (Vapnik & Vashist, 2009) through a **Teacher–Student framework** (Lopez-Paz et al., 2015) (Figure 1): the teacher uses X to optimize virtual observations, and the student learns to predict them using only deployment-accessible inputs.

The teacher evaluates virtual observations through the full-field reconstruction loss:

$$\mathcal{L}_F(z) = \ell_F\big(F_{\mathcal{A}}(z), X\big). \tag{7}$$

Here ℓ_F denotes the experiment-specific reconstruction loss, including its temporal aggregation and normalization. The following theorem motivates full-field supervision: under squared loss, it relates the reconstruction objective to approximation of the conditional mean given deployment evidence.

Theorem 1 (Squared-risk decomposition)**.** Let $X \in \mathbb{R}^d$ have finite second moment, and let E denote the deployment evidence. Define the conditional mean $m(E) = \mathbb{E}[X \mid E]$. Any E-measurable reconstruction $\widehat{X}(E)$ with finite squared risk satisfies, almost surely,

$$\underbrace{\mathbb{E}[\|\widehat{X}(E) - X\|_2^2 \mid E]}_{\text{reconstruction risk}} = \underbrace{\mathbb{E}[\|X - m(E)\|_2^2 \mid E]}_{\text{irreducible risk}} + \underbrace{\|\widehat{X}(E) - m(E)\|_2^2}_{\text{conditional-mean error}}. \tag{8}$$

Minimizing expected squared full-field loss is therefore equivalent to minimizing the expected squared distance to the conditional mean within the realizable reconstruction class, without explicitly computing that mean. This motivates using the privileged field X to supervise the teacher through $\mathcal{L}_F$. The identity applies to squared loss; it does not imply the same equivalence for relative ℓ_2 objectives. Appendix B.1 provides the proof, a distillation counterexample, and the corresponding identity for independently randomized decoders. Table 5 reports lower empirical MSE with TT in all 18 settings; Appendix F.3.2 explains the conditional-mean interpretation and its limits.

Starting from a deployment-accessible initialization, the teacher backpropagates $\mathcal{L}_F$ through the frozen reconstructor to optimize z using a fixed budget of constrained Adam updates (Kingma & Ba, 2014):

$$\nabla_z \mathcal{L}_F = \left(\frac{\partial R_\theta(\mathcal{A}, z)}{\partial z} \right)^\top \nabla_{\widehat{X}} \ell_F(\widehat{X}, X). \tag{9}$$

$$\begin{aligned} g^{(k)} &= \nabla_z \mathcal{L}_F(z^{(k)}), \\ z^{(k+1)} &= \Pi_{\mathcal{Z}} \Big(\mathrm{AdamUpdate}(z^{(k)}, g^{(k)}; s^{(k)}, \eta) \Big). \end{aligned} \tag{10}$$

Here $s^{(k)}$ is the Adam state and η the learning rate; $\Pi_{\mathcal{Z}}$ enforces the input constraint around the reference initialization. We denote the resulting feasible teacher candidate by z^{T}; it is not assumed to be a globally optimal control.

However, the teacher requires privileged full-field information and iterative optimization, making it unsuitable for direct deployment. We therefore train a GRU-based student (Cho et al., 2014) to predict the teacher's virtual observations from deployment-accessible inputs. The student encodes the observation sequence together with the deployment and reference coordinates, and predicts a constrained residual correction to a reference initialization:

$$z_\phi^{\mathcal{A}} = z_{\mathrm{init}}^{\mathcal{A}} + \Delta z_\phi(\mathcal{A}, \mathcal{B}, b). \tag{11}$$

The label-supervision objective is the mean squared discrepancy between student and teacher virtual inputs, written schematically as

$$\mathcal{L}_{\mathrm{distill}}(\phi) = \mathbb{E}_{\mathrm{train}} \left[\| T_\phi(E; \mathcal{A}) - z^{\mathrm{T}} \|_2^2 \right], \tag{12}$$

with averaging over the supervised coordinates and times. Any additional output-loss refinement is separate from this label objective.

2.4 Online and Offline TT

For simplicity, we have omitted explicit time indices in the preceding formulation. We now distinguish offline and online TT according to the observation availability and update requirements of the reconstructor.

2.4.1 Offline Trajectory Translation

The offline version applies when the complete deployment observation sequence is available. In the sequence formulas below, we index measurement updates by $1, \ldots, T$. The translator jointly uses the entire trajectory and its measurements to generate a virtual observation sequence on the reference interface:

$$z_{1:T}^{\mathcal{A}} = T_\phi^{\mathrm{off}}(\mathcal{A}, \mathcal{B}, b_{1:T}). \tag{13}$$

The frozen reconstructor then produces

$$\widehat{X}_{1:T}^{TT} = R_\theta(\mathcal{A}, z_{1:T}^{\mathcal{A}}). \tag{14}$$

Because the complete sequence is accessible, we use a bidirectional GRU to encode the observation sequence, allowing the virtual observation at time t to depend on deployment measurements from later times. This version is intended for offline sequence reconstruction and does not require an output as soon as each measurement arrives.

2.4.2 Online Trajectory Translation

Online TT addresses practical tasks that require continually incorporating new observations and updating field estimates. Using only evidence available at each update, we consider a block of length k starting at time τ:

$$I_{k,\tau} = \{\tau, \tau + 1, \ldots, \tau + k - 1\}.$$

Once the block has been collected, the translator generates

$$z_{I_{k,\tau}}^{\mathcal{A}} = T_\phi^{\mathrm{on}}\big(A_{1:\tau+k-1}, B_{1:\tau+k-1}, b_{1:\tau+k-1}\big). \tag{15}$$

During training, the teacher jointly optimizes the virtual observations using an aggregation of full-field losses within the block:

$$\mathcal{L}_{F,I_{k,\tau}}\big(z^{\mathcal{A}}_{I_{k,\tau}}\big) = \sum_{t\in I_{k,\tau}} w_t\, \ell_{F,t}\Big(\widehat{X}_t(z^{\mathcal{A}}_{I_{k,\tau}}), X_t\Big)\,, \tag{16}$$

where fixed weights $w_t \geq 0$ sum to one and specify the temporal aggregation used in each experiment. Here, reconstruction starts from the state inherited from the preceding block, which is held fixed during the current block's optimization. Gradients are backpropagated through the within-block reconstruction updates to adjust only $z^{\mathcal{A}}_{I_{k,\tau}}$:

$$\begin{aligned} g^{(\ell)}_{I_{k,\tau}} &= \nabla_{z^{\mathcal{A}}_{I_{k,\tau}}} \mathcal{L}_{F,I_{k,\tau}}\Big|_{z^{\mathcal{A}}_{I_{k,\tau}} = z^{\mathcal{A},(\ell)}_{I_{k,\tau}}}\,, \\ z^{\mathcal{A},(\ell+1)}_{I_{k,\tau}} &= \Pi_{\mathcal{Z}_{I_{k,\tau}}}\Big(\mathrm{AdamUpdate}\Big(z^{\mathcal{A},(\ell)}_{I_{k,\tau}}, g^{(\ell)}_{I_{k,\tau}}; s^{(\ell)}_{\mathrm{Adam}}, \eta\Big)\Big)\,. \end{aligned} \tag{17}$$

Here, ℓ indexes optimization iterations, $s^{(\ell)}_{\mathrm{Adam}}$ denotes the optimizer state, and $\Pi_{\mathcal{Z}_{I_{k,\tau}}}$ enforces the prescribed input constraint within the current block. Previously submitted virtual observations and the reconstructor parameters remain unchanged.

3 Theory: Evidence Preservation and Interface Reachability

We first ask what a virtual input must preserve in an exact Bayesian model, then determine whether the reference interface can express it. The finite-state case makes this question a linear system. Appendix C develops its solution, its Taylor–Green instantiation, and its extension to general state spaces before explaining the connection to neural reconstructors.

From posterior matching to an evidence equation. Let $J_t \in \{0, \dots, K-1\}$ have transition matrix P_t, and let $E_t = (B_{1:t}, b_{1:t})$ denote the deployment evidence. In the reference model, the target and controlled beliefs satisfy

$$\pi^B_t = U_{g^B_t}(\pi^B_{t-1}P_t), \qquad \widetilde{\pi}_t = U_{g^A_t(\cdot;z_t)}(\widetilde{\pi}_{t-1}P_t), \tag{18}$$

where $U_g(\mu)_j = \mu_j g(j)/\sum_i \mu_i g(i)$ and $g^B_t(j) = p(b_t \mid J_t = j, B_t, E_{t-1})$. For a common strictly positive prediction and positive potentials, the updates agree exactly when the potentials are proportional across states. Hence, with $\beta(g)_j = \log[g(j)/g(0)]$, $j = 1, \dots, K-1$, the requirement is

$$\beta^A_t(z_t) = \beta^B_t. \tag{19}$$

Thus, the object to preserve is relative evidence, not the true measurements that would have been collected on A.

Solving the equation under an input constraint. For shared-covariance Gaussian observations, expansion of the log likelihood gives an affine map $\beta^A_t(z) = M_t z - q_t$; its coefficients are derived in Appendix C.1. Consider

$$\mathcal{Z}_t = \{z^{\mathrm{ref}}_t + u : \|u\|_2 \leq r_t\}, \qquad v_t = \beta^B_t + q_t - M_t z^{\mathrm{ref}}_t. \tag{20}$$

The translation problem becomes $M_t u = v_t$ with $\|u\|_2 \leq r_t$.

Theorem 2 (Exact reachability of bounded evidence translation)**.** *An evidence-preserving input exists in $\mathcal{Z}_t$ if and only if*

$$v_t \in \mathrm{range}(M_t), \qquad \|M_t^\dagger v_t\|_2 \leq r_t. \tag{21}$$

Its minimum-displacement solution is $z^*_t = z^{\mathrm{ref}}_t + M_t^\dagger v_t$. If the two recursions share their initial law and transitions, have positive predictions, and satisfy these conditions at every update, this construction gives $\widetilde{\pi}_t = \pi^B_t$ at every prefix. It is causal when all quantities used at update t are measurable from E_t and the known reference interface.

The two conditions separate an unavailable evidence direction from an insufficient input radius. When exact translation is infeasible, constrained least squares finds the smallest evidence residual; a one-step KL bound relates that residual to posterior error (Appendix C.2). These are properties of the specified interface, not guarantees that a learned translator finds its best input.

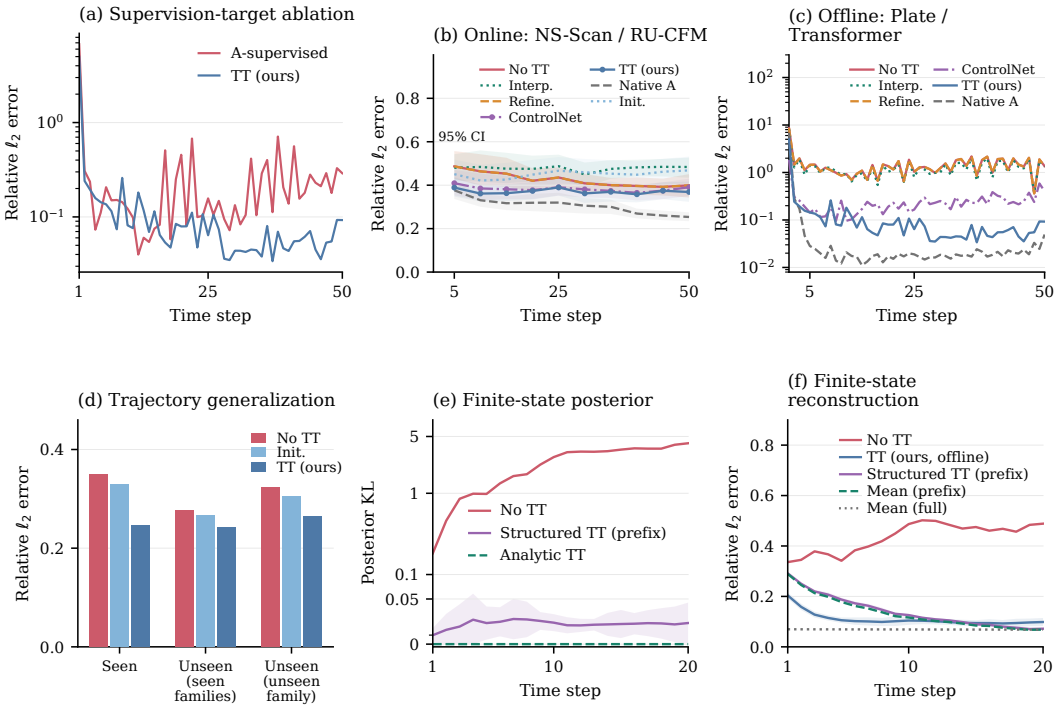

Figure 4: **Interface shift and trajectory translation.** (a) A-supervised versus TT on Plate–Transformer (example independent_000; 50 nonzero-target frames; Appendix F.2.1). (b,c) Online/offline reconstruction with parameter-matched ControlNet; shading: 95% field-bootstrap CIs; Init.: TT initialization; Interp.: spatial interpolation. (d) Generalization to unseen trajectories (Native A: 0.1715). (e,f) Finite-state posterior KL and reconstruction error. Structured/Analytic TT use learned/exact likelihood ratios with a fixed inverse. Mean denotes the conditional field mean; prefix/full use observations through the current/final time.

From the toy model to physical reconstruction. The Taylor–Green toy directly tests the exact Bayesian construction (Appendix C.3); the general-state extension is given in Appendix C.4. Our deterministic neural reconstructors instead produce point estimates of the physical field. Because their reconstruction maps need not implement the assumed Bayesian update, the reachability theorem does not directly provide a virtual-observation construction for these models. We therefore evaluate empirically whether TT improves reconstruction. Theorem 1 provides the squared-error interpretation: with the same deployment evidence and evaluation distribution, a reduction in population squared risk equals a reduction in mean squared distance to the conditional mean. The corresponding empirical MSE results and their interpretation are reported in Appendix F.3, Table 5.

4 Experiments

We evaluate TT from two complementary perspectives: reconstruction under observation-interface shifts in physical-field benchmarks, and direct verification of evidence preservation in a controlled finite-state model.

4.1 Controlled Validation in a Finite-State Setting

Experiment Setup We construct a toy dataset with four phase-shifted Taylor–Green flow trajectories. Each sequence samples a fixed state identity $J \in \{0, 1, 2, 3\}$, while its physical field evolves over time:

$$X_t = X_t^{(J)}, \qquad \widehat{X}_t = \sum_{j=0}^{3} \widetilde{\pi}_t(j) X_t^{(j)}. \tag{22}$$

Here, $\widetilde{\pi}_t$ is the state probability vector updated by a frozen Bayesian reconstructor with known Gaussian observation likelihoods on interface $\mathcal{A}$.

At deployment, noisy observations $b_{1:t}$ are collected along a rotated trajectory $\mathcal{B}$ and translated into virtual observations on $\mathcal{A}$. To assess evidence-preserving reachability, we compare the resulting probabilities with the exact deployment posterior:

$$\pi_t^{\mathcal{B}}(j) = \Pr(J = j \mid B_{1:t}, b_{1:t}), \qquad D_{\mathrm{KL}}\big(\pi_t^{\mathcal{B}} \,\|\, \widetilde{\pi}_t\big)\,. \tag{23}$$

We additionally report full-field relative ℓ_2 error. The evidence equation and its correspondence to the toy are detailed in Appendix C.3. The complete data and training protocol is given in Appendix E.

Results. This experiment examines the posterior-matching prediction of Theorem 2. The Gaussian observation model yields the affine evidence interface in equation 20, with M_t constructed to be invertible at every step. Without an input-radius constraint, the theorem's construction reduces to

$$z_t^* = M_t^{-1}(\beta_t^{\mathcal{B}} + q_t), \tag{24}$$

which we call *Analytic TT*. It satisfies the evidence-matching condition in equation 19 and therefore predicts $\widetilde{\pi}_t = \pi_t^{\mathcal{B}}$ throughout the sequence.

Figure 4(e) confirms this prediction: Analytic TT maintains posterior KL at numerical zero, whereas directly feeding deployment observations through the original interface does not preserve the posterior. The learned variants additionally illustrate approximate evidence translation and its reconstruction benefits (Figure 4(f)). The analytic reference is unconstrained; learned input bounds and information availability need not match it (Appendix C.3).

4.2 Evaluation on Physical-Field Benchmarks

Experiment Setup We consider five dataset–sensing configuration combinations (Figure 3): Navier–Stokes flows around three cylinders with a moving scan sensor (NS–Scan) or a fixed 16-sensor array (NS–Uniform), airfoil flows with four fixed sensors, plate strain fields with 12 fixed sensors, and three-dimensional urban pollutant fields with a moving UAV.

We use Transformer (Vaswani et al., 2017), SHRED (Mobile-SHRED for moving sensors), and FNO (Li et al., 2020) as reconstructors for offline TT, and RU-CFM as an amortized Bayesian reconstructor for online TT. We evaluate reconstruction with and without TT and compare against various add-on modules. A-supervised translation is studied in Appendix F.2.1. Dataset and sensing details are in Appendix D; model and baseline implementations are in Appendix G.

Results. Table 1 reports the available comparisons. TT outperforms parameter-matched ControlNet (Zhang et al., 2023) adaptation in 14 of 18 settings. With a larger ControlNet adapter, TT still achieves lower error in seven settings, despite the baseline using 21–250 times as many trainable parameters in those settings (Appendix F.4). Figure 4(b,c) shows temporal reconstruction errors for online RU-CFM and an offline Transformer example, respectively. TT also improves consistency with the deployment observations in 17 of the 18 settings (Appendix F.6, Table 9).

Generalization to unseen sensing trajectories. We evaluate a single GRU-based TT attached to a frozen Transformer on NS–Uniform, considering seen trajectories, unseen trajectories within seen families, and trajectories from an unseen family. As shown in Figure 4(d), TT reduces reconstruction error relative to both No TT and interpolation in all three groups. On the two unseen-trajectory groups, TT reduces error by 12.1% and 18.3% relative to No TT, respectively. These results demonstrate empirical cross-trajectory generalization.

4.3 Ablation and Inference Cost

Figure 4(a) compares A-supervised translation with TT on Plate–Transformer, the setting with the largest aggregate improvement (73.5%); the displayed sequence has the largest relative improvement among three archived paired examples. Appendix F.2 reports all 18 comparisons and the teacher-budget ablation, with spatial comparisons in Appendix F.1. Table 2 reports four offline timings:

Table 1: Relative ℓ_2 reconstruction error under observation-interface shifts (lower is better). Native A uses different observations and serves only as a reference. Reduction is $(\text{No TT} - \text{TT})/\text{No TT} \times 100\%$. Bold and underline mark the lowest and second-lowest errors, excluding Native A. Pollutant errors are averaged across the three channel-wise relative ℓ_2 errors.

Physical field	Reconstructor	Native A	No TT	Interp.	Refine.	ControlNet	TT (ours)	Reduction (↑)
NS–Scan	Transformer	0.2672	0.6185	0.6327	0.6185	**0.3676**	<u>0.3735</u>	39.6%
	Mobile-SHRED	0.2939	0.6607	0.6607	0.6607	**0.3649**	<u>0.3755</u>	43.2%
	FNO	0.3339	0.5042	0.5829	0.5042	<u>0.4678</u>	**0.3449**	31.6%
	RU-CFM	0.3114	0.4299	0.4769	0.4299	<u>0.3842</u>	**0.3774**	12.2%
NS–Uniform	Transformer	0.1715	0.4416	0.4080	0.4415	<u>0.3105</u>	**0.2312**	47.6%
	SHRED	0.2125	0.4213	0.3884	0.4212	<u>0.3007</u>	**0.2944**	30.1%
	FNO	0.2689	<u>0.3051</u>	0.3985	<u>0.3051</u>	0.3386	**0.2882**	5.5%
	RU-CFM	0.1586	0.4432	0.4318	0.4430	**0.2126**	<u>0.2966</u>	33.1%
Airfoil ($\times 10^{-2}$)	Transformer	0.1900	2.1350	1.1410	2.1310	**0.2795**	<u>0.4030</u>	81.1%
	SHRED	0.3550	2.1480	1.3670	2.1450	<u>0.6866</u>	**0.5070**	76.4%
	FNO	0.4620	2.8420	1.9140	2.8390	<u>0.7724</u>	**0.5510**	80.6%
Plate	Transformer	0.0287	1.3907	1.2382	1.3888	<u>0.3857</u>	**0.1538**	88.9%
	SHRED	0.0549	1.2912	1.2401	1.2891	<u>0.5490</u>	**0.3751**	70.9%
	FNO–128	0.1995	1.7894	1.6838	1.7872	<u>0.9991</u>	**0.7434**	58.5%
	FNO–512	0.1995	1.7894	1.6838	1.7872	<u>0.9991</u>	**0.6897**	61.5%
Pollutant	Transformer	0.1335	0.5910	0.5894	0.5903	<u>0.2004</u>	**0.1703**	71.2%
	SHRED	0.1818	1.0785	0.3733	1.0778	<u>0.2397</u>	**0.2303**	78.6%
	FNO	0.1615	0.9744	0.5874	0.9735	<u>0.5626</u>	**0.2427**	75.1%

ControlNet matches TT in trainable adapter parameters (within 0.3%); frozen backbone parameters are excluded. Full-size results appear in Appendix F.4. The A-supervised translator is evaluated separately in Appendix F.2.1. Field refinement optimizes the directly reconstructed field using deployment-observation consistency, without TT. RU-CFM uses online TT; the other reconstructors use offline TT. For Plate, FNO–128 and FNO–512 denote teacher optimization budgets of 128 and 512 Adam steps, respectively.

Table 2: Inference cost (ms). Additional time and overhead are measured relative to the original model. Field refinement includes the initial reconstruction and 32 optimization steps on X.

Setting	Original	Model + Refine	Model + TT(ours)	Added time	Overhead
Airfoil Transformer	2.70	44.16	3.09	0.39	14.3%
Airfoil SHRED	3.75	44.95	4.19	0.44	11.7%
Airfoil FNO	16.13	56.44	16.54	0.41	2.5%
Plate FNO	34.01	108.46	34.59	0.57	1.7%

Times are per complete sequence: 50 frames for Airfoil and 51 for Plate. Measurements use GPU-resident inputs and cached tokens, excluding data loading, token construction, and teacher optimization.

TT adds 0.39–0.57 ms (1.7–14.3%) per sequence. Model-plus-TT timing includes translation and reconstruction; field-refinement timing includes the initial reconstruction and 32 optimization steps on X. These results indicate a small absolute inference overhead.

5 Conclusion

TT adapts frozen reconstructors through virtual observations, using a full-field-supervised teacher and a deployment-observation student. Our theory characterizes finite-state Bayesian reachability and links squared risk to conditional-mean approximation. Experiments demonstrate control through the original interface across sensing trajectories: TT outperforms parameter-matched ControlNet in 14/18 settings, adding 0.39–0.57 ms per complete sequence in offline timing tests.

Limitations and future directions. Our configuration-to-configuration adaptation does not yet enable zero-shot use across arbitrary configurations, leaving a gap to closed-loop active sensing. Beyond the Bayesian toy, full recovery of the field distribution conditioned on a sensing trajectory and its measurements remains unestablished. Future work will target active exploration and trajectory observation-value assessment for amortized reconstructors, separating information value from interface-shift errors.

AI USE STATEMENT

Generative AI tools supported the formulation and refinement of research hypotheses, methodological development, feedback on experimental designs, and dataset cleaning and reformatting. We also used these tools for brainstorming, information search, code development and revision, and editing the manuscript for clarity. Mathematical assistance included refining formal statements and working through supporting derivations. The authors take full responsibility for the accuracy and integrity of the final manuscript, analyses, and accompanying artifacts.

REPRODUCIBILITY STATEMENT

The assumptions and proofs are provided in the theoretical appendices; physical datasets, sensing protocols, normalization, and metrics are described in Appendix D. Appendix G details the reconstructors, translators, baselines, and timing protocol. The experimental appendices provide toy protocols, ablations, and squared-error evaluations. Code and experiment configurations are available in our anonymous repository, whose README documents external prerequisites and the scope of the released implementation.

REFERENCES

Aditya Sai Pranith Ayapilla, Kazuya Miyashita, Yuki Yasuda, and Ryo Onishi. Uncertainty-aware spatiotemporal super-resolution data assimilation with diffusion models. *arXiv preprint arXiv:2604.21180*, 2026.

Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006.

Kyunghyun Cho, Bart Van Merriënboer, Çağlar Gulçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1724–1734, 2014.

Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022.

Hiep Vo Dang and Phong CH Nguyen. Deep operator learning for high-fidelity fluid flow field reconstruction from sparse sensor measurements. *Journal of Computing and Information Science in Engineering*, 26(1):011007, 2026.

Dassault Systèmes SIMULIA. Abaqus/Explicit. `https://www.3ds.com/products/simulia/abaqus/explicit`, n.d. Official software documentation and product description.

Huayu Deng, Jinghui Zhong, Xiangming Zhu, Yunbo Wang, and Xiaokang Yang. Laser: Learning active sensing for continuum field reconstruction. *arXiv preprint arXiv:2604.19355*, 2026.

Megan R Ebers, Jan P Williams, Katherine M Steele, and J Nathan Kutz. Leveraging arbitrary mobile sensor trajectories with shallow recurrent decoder networks for full-state reconstruction. *IEEE Access*, 12:97428–97439, 2024.

Geir Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367, 2003.

Christopher J Greenshields, Henry G Weller, Luca Gasparini, and Jason M Reese. Implementation of semi-discrete, non-staggered central schemes in a colocated, polyhedral, finite volume framework, for high-speed viscous flows. *International journal for numerical methods in fluids*, 63(1): 1–21, 2010.

Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-solving under partial observation. *Advances in Neural Information Processing Systems*, 37: 130291–130323, 2024.

Yinan Huang, Hans Hao-Hsun Hsu, Junran Wang, Bo Dai, and Pan Li. Accelerated sequential flow matching: A bayesian filtering perspective. *arXiv preprint arXiv:2602.05319*, 2026.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

François-Xavier Le Dimet and Olivier Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A: Dynamic Meteorology and Oceanography*, 38(2):97–110, 1986.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. *arXiv preprint arXiv:1511.03643*, 2015.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

E Lumet, T Jaravel, and MC Rochoux. Ppmles–perturbed-parameter ensemble of must large-eddy simulations. Dataset. Zenodo, 2024a.

Eliott Lumet, Thomas Jaravel, Mélanie C Rochoux, Olivier Vermorel, and Simon Lacroix. Assessing the internal variability of large-eddy simulations for microscale pollutant dispersion prediction in an idealized urban environment. *Boundary-Layer Meteorology*, 190(2):9, 2024b.

Eliott Lumet, Thomas Jaravel, and Mélanie C Rochoux. Dataset of microscale atmospheric flow and pollutant concentration large-eddy simulations for varying mesoscale meteorological forcing in an idealized urban environment. *Data in Brief*, 58:111285, 2025.

Yuezhou Ma, Haixu Wu, Hang Zhou, Huikun Weng, Jianmin Wang, and Mingsheng Long. Physense: Sensor placement optimization for accurate physics sensing. *Advances in Neural Information Processing Systems*, 38:35697–35725, 2025.

Agnese Marcato, Eric Guiltinan, Hari Viswanathan, Daniel O'Malley, Nicholas Lubbers, and Javier E Santos. Journey over destination: dynamic sensor placement enhances generalization. *Machine Learning: Science and Technology*, 5(2):025070, 2024.

Iven MY Mareels, Brian DO Anderson, Robert R Bitmead, Marc Bodson, and Shankar S Sastry. Revisiting the mit rule for adaptive control. In *Adaptive Systems in Control and Signal Processing 1986*, pp. 161–166. Elsevier, 1987.

Michael D McKay, Richard J Beckman, and William J Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.

Franck Nicoud and Frédéric Ducros. Subgrid-scale stress modelling based on the square of the velocity gradient tensor. *Flow, turbulence and Combustion*, 62(3):183–200, 1999.

Taku Nonomura, Chihaya Abe, Ryo Naramura, and Yasuo Sasaki. Real-time feedback control of flow velocity field using sparse processing particle image velocimetry and plasma actuators. *Experiments in Fluids*, 66(7):136, 2025.

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

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.

Rinto Roy, Alexander Tessler, Cecilia Surace, and Marco Gherlone. Efficient shape sensing of plate structures using the inverse finite element method aided by strain pre-extrapolation. *Thin-Walled Structures*, 180:109798, 2022.

Javier E Santos, Zachary R Fox, Arvind Mohan, Daniel O'Malley, Hari Viswanathan, and Nicholas Lubbers. Development of the senseiver for efficient field reconstruction from sparse observations. *Nature Machine Intelligence*, 5(11):1317–1325, 2023.

Mikhail L Shur, Philippe R Spalart, Mikhail Kh Strelets, and Andrey K Travin. A hybrid rans-les approach with delayed-des and wall-modelled les capabilities. *International journal of heat and fluid flow*, 29(6):1638–1649, 2008.

Lieven Vandenberghe and Stephen Boyd. *Convex optimization*, volume 1. Cambridge university press Cambridge, 2004.

Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5-6):544–557, 2009.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.

Gabriel D Weymouth. Lily pad: Towards real-time interactive computational fluid dynamics. *arXiv preprint arXiv:1510.06886*, 2015.

Gabriel D Weymouth and Dick KP Yue. Boundary data immersion method for cartesian-grid simulations of fluid-body interaction problems. *Journal of Computational Physics*, 230(16):6233–6247, 2011.

Jan P Williams, Olivia Zahn, and J Nathan Kutz. Sensing with shallow recurrent decoder networks. *Proceedings of the Royal Society A*, 480(2298):20240054, 2024.

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3813–3824. IEEE, 2023.

A Related Work

A.1 Physical-Field Reconstruction

Recovering a complete physical field from sparse observations can be approached through inference performed for the current measurements or through a reusable reconstruction map learned offline. Non-amortized approaches explicitly combine the observation operator, physical model, and prior when solving a given instance. For example, variational data assimilation optimizes states or initial conditions to reconcile model trajectories with observations Le Dimet & Talagrand (1986), while physics-informed neural networks (PINNs) trained for individual problems combine observation errors and equation residuals to infer fields and unknown parameters Raissi et al. (2019). Because measurement locations enter the inference objective explicitly, these formulations generally do not require deployment observations to follow a fixed input interface used to train a conditional reconstruction map. They therefore avoid that particular source of interface mismatch. Repeated optimization, model integration, or gradient evaluation can nevertheless increase the cost per instance and complicate frequent reconstruction. These methods remain susceptible to errors in the physical model, prior, or observation model.

Amortized reconstructors move part of this computation into offline training by learning a reusable mapping from partial measurements to complete fields. SHRED encodes sparse sensor histories with an LSTM and reconstructs high-dimensional states through a shallow decoder Williams et al. (2024). FLRONet uses a branch–trunk operator-learning architecture to map sparse measurements to flow fields represented continuously in space and time Dang & Nguyen (2026). Senseiver uses positional encoding and cross-attention to represent sparse observations and decode field values at query coordinates, supporting variable observation counts Santos et al. (2023). These methods illustrate how learned reconstruction maps can reduce repeated inference costs, while deployment performance still depends on training coverage, input representation, and generalization.

Within amortized reconstruction, online and offline use should be distinguished by the evidence available at inference time. Online reconstruction uses current and past observations and updates estimates as measurements arrive; offline reconstruction may jointly use an entire sequence, including measurements later than the state being reconstructed. SHRED's reconstruction from sensor histories ending at the current time exemplifies the former information pattern Williams et al. (2024). This distinction is independent of whether training occurs offline and cannot be inferred from the use of an LSTM, Transformer, or generative model alone.

Online field estimation also includes classical sequential data assimilation. The ensemble Kalman filter (EnKF) propagates an ensemble of states and then corrects its forecast using current observations and sample covariance estimates Evensen (2003). Its updates can accommodate observation operators that vary over time, without relying on an end-to-end supervised map tied to a fixed sensing trajectory. Performance nevertheless depends on factors such as dynamics errors, ensemble size, and covariance estimation, while computational cost depends on ensemble propagation and model complexity. Non-amortized and amortized approaches thus make different tradeoffs among model assumptions, deployment computation, and training investment.

A.2 Adaptation to Changing Observation Configurations

Amortized reconstruction need not be restricted to fixed sensors, but handling configuration changes typically requires suitable representations and training. Mobile SHRED extends history-based reconstruction to moving sensors and explicitly requires training measurements along the chosen trajectories; its original architecture does not directly extrapolate to unseen sensor trajectories Ebers et al. (2024). Other approaches encode observation coordinates, represent measurements as sets, or vary sensor counts and locations during training. Senseiver provides an attention-based representation for variable numbers of observations Santos et al. (2023). *Journey over Destination* studies moving and optimizing sensor locations during training to exploit spatial data more fully and improve reconstruction generalization Marcato et al. (2024). These approaches incorporate configuration diversity into model design or training, without guaranteeing performance on arbitrary unseen trajectories.

A complementary route imposes observation constraints during deployment-time inference. Diffusion posterior sampling (DPS) incorporates measurement-dependent gradients into diffusion sam-

pling to solve inverse problems using a pretrained generative prior Chung et al. (2022). Diffusion-PDE combines a generative prior with observation and physics constraints for PDE problems under partial observation Huang et al. (2024). DiffSRDA further studies deployment-time changes in sensor layouts: directly changing the layout in the original conditioning channels can reduce performance, while sampling-time guidance helps incorporate additional measurements or measurements from different layouts Ayapilla et al. (2026). Such methods can introduce new constraints without retraining the generative prior, but require the associated sampling and guidance computations at inference time; their costs depend on the implementation and iteration budget.

Our focus is post-hoc input adaptation for a reconstructor trained on a fixed interface. TT conveys deployment information through virtual observations on the original interface while preserving the reconstructor's parameters and internal reconstruction procedure. This design can complement training with diverse configurations.

Conditional add-on modules. ControlNet introduces trainable conditional branches with zero-initialized connections to a frozen text-to-image diffusion model (Zhang et al., 2023). We use this conditional-adaptation principle to construct field-reconstruction baselines under different parameter budgets. Their task-specific implementations are described in Appendix G.4.

A.3 Sensor Placement and Active Sensing

Sensor placement and active sensing additionally require decisions about where observations are most useful. Evaluating candidate configurations requires a reconstruction model or surrogate objective that can handle those configurations. If a reconstructor suffers from interface mismatch, its scores may conflate insufficient observational information with failure to exploit the information available. Existing approaches therefore develop observation selection mechanisms together with reconstruction models and training procedures suited to changing observations.

PhySense trains a conditional flow reconstructor with cross-attention and randomized sensor counts and positions. It then freezes the reconstructor and optimizes sensor placements with projected gradients under geometric constraints, using a flow-matching objective as a surrogate Ma et al. (2025). Placement optimization is thus supported by a reconstructor trained across configurations, rather than one trained only on a single fixed trajectory.

LASER addresses online active sensing of dynamic fields by formulating sensor movement as a partially observable decision process. Its world model uses a set representation of observation coordinates and values, is trained with randomly sampled sensor locations, and combines historical memory with latent dynamics prediction to support action selection. Policy learning uses feedback based on reconstruction error Deng et al. (2026). Its support for changing observations therefore comes from configuration diversity during training, coordinate-aware representations, and a predictive model. It does not establish that the error of an arbitrary frozen, fixed-interface reconstructor reliably measures the value of an unseen trajectory.

These studies primarily develop observation selection together with suitable reconstruction models. We address a complementary question: how to make effective use of measurements from a new interface when a reconstructor has already been trained on a particular interface and frozen. Our focus is on mitigating interface mismatch as a prerequisite for such use, without equating reconstruction error directly with a trajectory's information value.

B Supporting Results for the Reconstruction Objective

This appendix supports the reconstruction objective and the teacher's local input optimization. These results do not establish evidence reachability; that analysis is developed separately in Appendix C.

B.1 Theorem 1: squared reconstruction risk and its scope

Proof of Theorem 1. Write $f(E) = \widehat{X}(E)$ and $m(E) = \mathbb{E}[X \mid E]$. Expanding gives

$$\|f(E) - X\|_2^2 = \|f(E) - m(E)\|_2^2 + \|X - m(E)\|_2^2 - 2\langle f(E) - m(E), X - m(E)\rangle.$$

The cross term has conditional expectation zero because $f(E) - m(E)$ is E-measurable and $\mathbb{E}[X - m(E) \mid E] = 0$. Taking expectations again yields

$$\mathcal{R}(f) = \mathcal{R}_{\text{Bayes}} + \mathbb{E}\|f(E) - m(E)\|_2^2, \tag{25}$$

where $\mathcal{R}(f) = \mathbb{E}\|f(E) - X\|_2^2$ and $\mathcal{R}_{\text{Bayes}} = \mathbb{E}\|X - m(E)\|_2^2$. □

The learned parameters and the joint evaluation law of (E, X) are fixed. E may be a complete offline sequence or the prefix available when an output is issued. Changing an active sensing policy can change this law and its Bayes risk. An unsquared or relative reconstruction loss does not generally have the same conditional-mean minimizer.

Randomized reconstruction. For an output $f(E, U)$, suppose the decoder randomness satisfies $U \perp\!\!\!\perp X \mid E$ and the squared risk is finite. The same calculation gives

$$\mathbb{E}[\|f(E,U) - X\|_2^2 \mid E] = \mathbb{E}[\|X - m(E)\|_2^2 \mid E] + \mathbb{E}[\|f(E,U) - m(E)\|_2^2 \mid E]. \tag{26}$$

The last term includes both the error of the conditional output mean and the conditional variance of the output. A single sample and an ensemble mean are therefore different estimators. Fixing a decoder seed gives a deterministic predictor to which Theorem 1 applies.

What the identity does not imply. Output-level squared-risk minimization is not equivalent to regression of privileged teacher controls. For example, let E be constant, let X equal 1 or 9 with equal probability, and let $F_A(z) = z^2$ on $[0, 4]$. The unique nonnegative exact teacher controls are 1 and 3. Their squared-label-loss minimizer is $z = 2$, giving output 4 and field risk 17. The conditional mean is 5, realizable by $z = \sqrt{5}$ with risk 16. Thus even exact teachers and optimal label regression need not minimize field risk. Our distillation procedure is evaluated empirically; the risk identity does not establish its consistency.

B.2 Local input compensation

Proof of Proposition 1. For a fixed target X, let $L(z) = \frac{1}{2}\|F_A(z) - X\|_2^2$. Suppose F_A is continuously differentiable near an interior feasible point z_0. With $e = F_A(z_0) - X$ and $J = DF_A(z_0)$, the chain rule gives $\nabla L(z_0) = J^\top e$. If this gradient is nonzero, then for sufficiently small $\eta > 0$, the input $z_\eta = z_0 - \eta J^\top e$ is feasible and

$$L(z_\eta) = L(z_0) - \eta\|J^\top e\|_2^2 + o(\eta) < L(z_0).$$

□

More generally, the same local descent argument applies to a differentiable scalar loss with nonzero input gradient. It explains a possible mechanism for input compensation, not global reachability or convergence of projected Adam. The direction uses the training target X and is not a deployment rule for the student.

A classical adaptive-control example. Model-reference adaptive control adjusts controller parameters using a tracking error and its sensitivity. The MIT gradient rule takes the form $\dot{\vartheta} = -\gamma e\, \partial e/\partial \vartheta$ for the objective $e^2/2$ (Mareels et al., 1987). This motivates treating an adjustable input as a control variable when a fixed model has a systematic output error. For a scalar illustration of Proposition 1, let the frozen reconstruction map be $F_A(z) = \alpha z + \beta$, with $\alpha \neq 0$, and let a genuine input a have error $e_0 = \alpha a + \beta - X$. The virtual input $z_\eta = a - \eta\alpha e_0$ gives

$$F_A(z_\eta) - X = (1 - \eta\alpha^2)e_0.$$

If $e_0 \neq 0$ and $0 < \eta < 2/\alpha^2$, its squared reconstruction error is strictly smaller, provided the input remains feasible. At $\eta = 1/\alpha^2$ the scalar error vanishes. Thus correcting the input can compensate a fixed gain or bias error, even when the original observation is accurate. In TT, the teacher uses training fields to construct such corrections, and the student learns to predict them from deployment observations. The adaptive-control analogy concerns sensitivity-based correction; the proposition supplies the local guarantee used here.

C Proofs and Extensions of the Reachability Theorem

We organize the analysis around one question: can a virtual input make the reference interface express the evidence supplied at deployment? We derive the equation, solve it under a norm constraint, instantiate it in the toy, and then separate its general-state extension from its empirical use with neural reconstructors. Superscripts A, B denote $\mathcal{A}, \mathcal{B}$.

C.1 What must a virtual observation preserve?

Let $J_t \in \{0, \dots, K-1\}$ follow a hidden Markov model with row-stochastic transition P_t. Sensing actions depend only on past visible evidence and independent randomness and do not alter the state dynamics. Thus, conditional on the observed history, selecting B_t contributes no additional hidden-state information before observing b_t. With $g_t^B(j) = p(b_t \mid J_t = j, B_t, E_{t-1})$, the target and controlled recursions are those in equation 18. The reference potential $g_t^A(j; z_t)$ defines the controlled update; z_t need not be a naturally generated measurement on A.

Assume positive potentials and strictly positive prediction distributions. A positive initial law and at least one nonzero entry in every column of each P_t suffice for positivity of all predictions. For a common positive prediction μ, normalization gives

$$U_{g^A(\cdot;z)}(\mu) = U_{g^B}(\mu) \quad \Longleftrightarrow \quad g^A(j; z) = c\, g^B(j) \quad \text{for every } j, \tag{27}$$

where $c > 0$ is independent of j. Indeed, equality of the normalized probabilities and $\mu_j > 0$ imply $g^A(j; z)/g^B(j) = \sum_i \mu_i g^A(i; z) / \sum_i \mu_i g^B(i)$. Conversely, a common factor cancels. Relative log likelihoods remove this irrelevant factor, giving equation 19. If a prior has smaller support, equality is required only on that support.

The Gaussian interface produces the linear system. For $g_t^A(j; z) = \mathcal{N}(z; m_{t,j}, \Sigma_t)$ with a shared positive-definite covariance, the term quadratic in z cancels between states. Consequently,

$$\beta_{t,j}^A(z) = (M_t)_j z - (q_t)_j, \tag{28}$$

$$(M_t)_j = (m_{t,j} - m_{t,0})^\top \Sigma_t^{-1}, \qquad (q_t)_j = \tfrac{1}{2}\big(m_{t,j}^\top \Sigma_t^{-1} m_{t,j} - m_{t,0}^\top \Sigma_t^{-1} m_{t,0}\big). \tag{29}$$

State-dependent covariances generally leave quadratic terms in z and do not satisfy this affine parameterization. The next step addresses whether the target β_t^B lies in the evidence set expressed by the allowed inputs.

C.2 How is the translation equation solved?

Suppress the time index and set $z = z^{\mathrm{ref}} + u$ and $v = \beta^B + q - M z^{\mathrm{ref}}$. Exact translation is $Mu = v$ with $\|u\|_2 \le r$.

Proof of Theorem 2. The equation has a solution exactly when $v \in \mathrm{range}(M)$. Every solution is $u = M^\dagger v + w$, where $w \in \ker M$ (Vandenberghe & Boyd, 2004). Orthogonality gives $\|u\|_2^2 = \|M^\dagger v\|_2^2 + \|w\|_2^2$. Hence the minimum-displacement solution is unique and lies in the ball exactly when equation 21 holds. Equal previous beliefs predict to equal beliefs under a common P_t; equation 27 then gives equality after the next update. Induction proves prefix-wise equality. Conversely, equality at adjacent prefixes forces evidence equality at that update under the positive common prediction. The pseudoinverse construction is causal when its inputs are causally available; for fixed interface matrices it is a measurable linear function of the required evidence and reference input. □

When exact matching is infeasible. A natural approximate control solves the convex problem

$$d_*^2 = \min_{\|u\|_2 \le r} \|Mu - v\|_2^2. \tag{30}$$

If $u_0 = M^\dagger v$ lies in the ball, it is a minimum-norm minimizer. Otherwise, for $r > 0$, the solution has the form

$$u_\lambda = (M^\top M + \lambda I)^{-1} M^\top v, \qquad \lambda > 0, \quad \|u_\lambda\|_2 = r. \tag{31}$$

These conditions follow from the Lagrangian of $\frac{1}{2}\|Mu - v\|_2^2$ with constraint $\frac{1}{2}(\|u\|_2^2 - r^2) \le 0$. Convexity makes them sufficient (Vandenberghe & Boyd, 2004, Section 5.5.3). An SVD reduces the norm condition to a one-dimensional root search. If $r = 0$, the only choice is $u = 0$. Radially clipping $M^\dagger v$ is not generally the constrained optimum.

What does the remaining residual imply? For a fixed common positive prediction μ, write $p_\mu(\beta) = U_{(1,e^{\beta_1},\dots,e^{\beta_{K-1}})}(\mu)$ and define

$$L_* = \min_{z \in \mathcal{Z}} D_{\mathrm{KL}}\big(p_\mu(\beta^B) \,\|\, p_\mu(Mz - q)\big).$$

Both minima are attained on the compact ball, and

$$L_* = 0 \iff d_* = 0, \qquad L_* \le \tfrac{1}{2} d_*^2. \tag{32}$$

To see the bound, let $F_\mu(\beta) = \log(\mu_0 + \sum_{j=1}^{K-1} \mu_j e^{\beta_j})$. Direct substitution gives the exponential-family identity (Wainwright & Jordan, 2008)

$$D_{\mathrm{KL}}(p_\mu(\beta^B) \| p_\mu(\beta)) = F_\mu(\beta) - F_\mu(\beta^B) - \nabla F_\mu(\beta^B)^\top (\beta - \beta^B). \tag{33}$$

Its Hessian is $H = \mathrm{diag}(p_{1:K-1}) - p_{1:K-1} p_{1:K-1}^\top$, so $0 \preceq H \preceq I$. Taylor's integral formula bounds the KL by $\frac{1}{2}\|\beta - \beta^B\|_2^2$; evaluation at a least-squares solution proves equation 32. Zero KL is equivalent to identical evidence coordinates under the common positive prior, proving the equivalence. The KL objective is also convex in z, but its minimizer need not equal the least-squares control.

If every probability coordinate is at least $\kappa > 0$ along all segments from the target to feasible evidence, the same argument gives $L_* \ge \frac{1}{2}\kappa^2 d_*^2$: for $w_0 = 0$, $w^\top H w = \sum_{i<j} p_i p_j (w_i - w_j)^2 \ge \kappa^2 \|w\|_2^2$. Although such a κ exists for each fixed compact problem, it may be extremely small. These are one-step statements with a common prediction; they are not a uniform long-horizon guarantee for approximate filters.

C.3 How does the Taylor–Green toy instantiate the equation?

The toy in Section 4 has four candidate Taylor–Green trajectories. Their identity $J \in \{0, 1, 2, 3\}$ is fixed within a sequence, while each physical field $X_t^{(j)}$ evolves. Thus $P_t = I$, and the positive initial probabilities and Gaussian likelihoods preserve positive predictions. At a given update, $m_{t,j} = \mathcal{H}_{A_t}(X_t^{(j)})$ contains the candidate's three (u, v, p) values on the reference interface. Applying equation 29 produces three evidence coordinates and a 3×3 matrix M_t. The reference scan is chosen to make this matrix invertible at every step.

The noisy deployment measurement gives

$$(\beta_t^B)_j = \log \frac{p(b_t \mid J = j, B_t)}{p(b_t \mid J = 0, B_t)}, \qquad j = 1, 2, 3.$$

The unbounded analytic reference uses

$$z_t^* = M_t^{-1}(\beta_t^B + q_t). \tag{34}$$

Substitution gives exact evidence equality. With the same initial law, the reference and deployment posteriors therefore coincide at each prefix in exact arithmetic. Their KL, evaluated against the same deployment prefix, is the direct numerical test of this prediction. The induced conditional field mean is $\sum_j \pi_t^B(j) X_t^{(j)}$; it need not equal the realized field because observational uncertainty remains.

This experiment instantiates the full-rank, unconstrained positive case. It does not test all rank/radius obstructions or a mixing-transition bound. The evidence-structured learned reference predicts relative evidence before applying the fixed inverse, whereas the practical offline student predicts bounded virtual inputs from the full observation sequence. Their controls, supervision, and information timing differ. Only like-for-like prefix posteriors test the equality above; offline reconstruction gains are a separate evaluation of learned translation.

C.4 What survives beyond finite states?

The finite-state assumption makes evidence equality a finite linear system; it is not needed for the underlying normalization argument. Let μ be a common probability measure on a measurable state space. Assume that $g^A(x;z)$ and $g^B(x)$ are finite and strictly positive μ-almost everywhere, with finite, nonzero normalizing constants $Z_A(z)$ and Z_B. Define

$$\Pi_A^z(dx) = \frac{g^A(x;z)}{Z_A(z)}\mu(dx), \qquad \Pi_B(dx) = \frac{g^B(x)}{Z_B}\mu(dx). \tag{35}$$

Equality of their densities relative to μ proves

$$\Pi_A^z = \Pi_B \quad \Longleftrightarrow \quad \log g^A(x;z) - \log g^B(x) = c \quad \mu\text{-a.e.}, \tag{36}$$

where c is independent of x. Conversely, exponentiating this identity and normalizing proves sufficiency. Thus the general-state problem is to match an evidence *function* up to an additive constant.

For a Gaussian reference likelihood $g^A(x;z) = \mathcal{N}(z; m_A(x), \Sigma_A)$ with fixed positive-definite Σ_A, this condition becomes

$$m_A(x)^\top \Sigma_A^{-1} z - \tfrac{1}{2} m_A(x)^\top \Sigma_A^{-1} m_A(x) = \log g^B(x) + c \quad \mu\text{-a.e.} \tag{37}$$

The quadratic term in z is constant with respect to x and is absorbed into c. This is the direct counterpart of $Mz - q = \beta^B$: finitely many input coefficients must now match a function on the prior support.

The same range and radius test in function space. To make the analogy precise, set $h(x) = \Sigma_A^{-1} m_A(x)$ and

$$b(x) = \log g^B(x) + \tfrac{1}{2} m_A(x)^\top \Sigma_A^{-1} m_A(x) - h(x)^\top z^{\mathrm{ref}}.$$

Assume additionally that h has square-integrable components and $b \in L^2(\mu)$. Center functions by $C_\mu f = f - \mathbb{E}_\mu f$ and define

$$(\mathsf{M}u)(x) = (C_\mu h(x))^\top u, \qquad v = C_\mu b.$$

Equation equation 37 is then exactly $\mathsf{M}u = v$ for $z = z^{\mathrm{ref}} + u$. The bounded linear map $\mathsf{M} : \mathbb{R}^d \to L^2(\mu)$ has finite, closed range. The orthogonal minimum-norm argument used above therefore gives the same criterion: $v \in \mathrm{range}(\mathsf{M})$ and $\|\mathsf{M}^\dagger v\|_2 \le r$. For explicit computation, let

$$G = \mathbb{E}_\mu[(C_\mu h)(C_\mu h)^\top], \qquad a = \mathbb{E}_\mu[(C_\mu h)v]. \tag{38}$$

Expanding $\|\mathsf{M}u - v\|^2_{L^2(\mu)}$ gives $u^\top G u - 2a^\top u + \mathbb{E}_\mu v^2$. Moreover, $a \in \mathrm{range}(G)$: if $w \in \ker G$, then $w^\top C_\mu h = 0$ almost everywhere and hence $w^\top a = 0$. The minimum-norm least-squares solution is $u_0 = G^\dagger a$, with residual $\mathbb{E}_\mu v^2 - a^\top G^\dagger a$. Exact bounded translation requires this residual to vanish and $\|u_0\|_2 \le r$. This characterization presumes the required likelihoods and integrals are available; it is not itself a learned deployment algorithm. A small $L^2(\mu)$ residual does not, without additional assumptions, inherit the finite-coordinate KL bound in equation 32.

Extension does not mean universal reachability. For a scalar example, take $x \sim \mathcal{N}(0,1)$, reference likelihood $z \mid x \sim \mathcal{N}(x,1)$, and deployment likelihood $b \mid x \sim \mathcal{N}(x, 1/4)$. Completing the square gives

$$\Pi_A^z = \mathcal{N}(z/2, 1/2), \qquad \Pi_B = \mathcal{N}(4b/5, 1/5). \tag{39}$$

Changing z cannot change the reference posterior variance, so no virtual value matches the full target distribution. Nevertheless, $z = 8b/5$ matches its conditional mean when no input constraint rules out that value. This distinguishes distribution reachability from reconstruction-mean reachability and motivates the separate evaluation of physical reconstruction.

C.5 What do the neural reconstruction experiments test?

There are two different changes from the toy to the physical benchmarks. First, the field is not restricted to four candidates; the general-state evidence condition above addresses this change within

an exact Bayesian model. Second, a learned reconstructor need not implement that model's likelihood update. Neither finite-state nor general-state normalization proves that a frozen neural network is a calibrated Bayesian filter.

The physical experiments therefore test an implementation of the input-control idea, rather than directly verifying the analytic inverse. On a fixed example, a feasible teacher input with improved full-field loss shows that the frozen model admits at least that improvement. A failed finite optimization does not establish an unreachable target, and a teacher that uses the true field does not demonstrate that its control can be identified from deployment evidence. Comparing the student with the initialization and teacher measures how much of the found improvement is retained by an evidence-driven predictor. Full-field error and deployment-observation consistency assess complementary aspects of the resulting reconstruction.

For squared loss, Theorem 1 provides an output-level interpretation in terms of conditional-mean approximation, with the distillation limitation in Appendix B.1. Other losses, including relative ℓ_2, are evaluated according to their own objectives. None of these point-reconstruction metrics establishes full posterior matching; that requires a reliable conditional-distribution reference and distributional evaluation. Similarly, a correct common decoder pushes equal beliefs to equal field laws, but an inaccurate decoder can break that connection.

Finally, the causal construction in Theorem 2 uses only E_t for z_t. A blockwise translator that waits for several observations has a different availability time for the earlier inputs in the block. We therefore distinguish online reconstruction from strict per-measurement operation, and neither the offline nor blockwise results establish an active closed-loop sensing guarantee.

D PHYSICAL DATASETS AND SENSING PROTOCOLS

We evaluate five dataset–sensing configuration combinations built from four physical datasets: NS–Scan, NS–Uniform, airfoil flow, plate strain, and urban pollutant dispersion. Table 3 lists the available data; translator training subsets are distinct from these archive sizes.

Table 3: Target representations and physical-case splits. Spatial dimensions are $H \times W$ or $D \times H \times W$.

Dataset	Field type	Channels; frames; grid	Train/validation/test
NS	Dynamic, 2D	3; 50; 128×128	9,846/128/1,993
Pollutant	Time-averaged, 3D	3; 1; $11 \times 181 \times 141$	123/26/27
Plate strain	Dynamic, 2D shell surface	3; 51; 100×200	5,000/500/500 mixtures
Airfoil	Dynamic, 2D	6; 50; 96×192	256/16/16

For case n, sensor j, and sampling index k, observations under configuration L are

$$\mathbf{y}_{k,j}^{L,(n)} = \mathcal{H}_{\mathbf{r}_j^L(k)}[\mathbf{q}_k^{(n)}] + \boldsymbol{\eta}_{k,j}^{L,(n)}, \qquad L \in \{A, B\}. \tag{40}$$

Here $\mathbf{y}$ is the measured channel vector, $\mathbf{r}$ the sensor position, $\mathbf{q}$ the full field, $\mathcal{H}$ the spatial sampling operator, and $\boldsymbol{\eta}$ the measurement noise (zero when no noise is added). For transient datasets, $\mathbf{q}_k^{(n)} = \mathbf{q}^{(n)}(\cdot, t_k)$ at physical time t_k. For pollutant dispersion, $\mathbf{q}_k^{(n)}$ is the case's time-averaged statistics and k indexes successive UAV positions, not LES time. Sampling uses bilinear interpolation for NS/airfoil, circular-footprint averaging for the plate, and source-mesh barycentric interpolation for pollutant observations and Cartesian resampling. Figures show fields before training normalization, with one color scale per channel across the displayed frames or heights.

D.1 NS: INCOMPRESSIBLE FLOW PAST THREE CYLINDERS

Generation and quantities. We generate two-dimensional incompressible flow past three cylinders with Lilypad Weymouth (2015), using the boundary data immersion method Weymouth & Yue (2011) to solve

$$\nabla \cdot \mathbf{u} = 0, \qquad \partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\rho^{-1}\nabla p + \nu \nabla^2 \mathbf{u}, \qquad Re = U_{\mathrm{ref}} D_{\mathrm{ref}}/\nu. \tag{41}$$

Here $\mathbf{u} = (u, v)$ is velocity, p pressure, ρ density, and ν kinematic viscosity; U_{ref} and D_{ref} define the reference speed and length. Each case stores 50 (u, v, p) frames on a 128×128 grid and three cylinder contours, each represented by 40 coordinate pairs. These contour samples describe geometry, not sensor locations. The data generator uses unit inflow speed and $D_{\text{ref}} = 16$ grid cells, hence $\nu = 16/Re$ in solver units. Its Latin hypercube design (McKay et al., 2000) spans $Re \in [50, 500)$; diameter factors are $[0.5, 1.3]$ for the first cylinder and $[0.5, 1.1]$ for the others, multiplied by 16. Neighbor-separation parameters span $[1.2, 2.2]$, with full-turn angles and feasibility constraints. Each archived boundary defines the actual geometry. Figure 5 uses archived frame indices and solver-scale values; the archived dataset does not specify a conversion to seconds or pascals. The archive contains 9,974 training-pool cases, of which 128 are reserved for validation (random seed 20260916), plus 1,993 test cases; recorded Re ranges from 50 to 499.

Interfaces A/B. The uniform setting uses 16 stationary probes in grid coordinates:

$$\mathcal{R}_A = \{16, 48, 80, 112\}^2, \qquad \mathcal{R}_B = \{(x + 12, y) : (x, y) \in \mathcal{R}_A\}. \tag{42}$$

The separate scanning setting uses one probe reading three channels at $k = 0, \ldots, 49$:

$$x_A(k) = 12.8 + 102.4[1 - |2\,\mathrm{frac}(k/20) - 1|], \quad y_A(k) = 12.8 + 102.4k/50, \tag{43}$$

$$(x_B(k), y_B(k)) = (127 - x_A(k), 127 - y_A(k)). \tag{44}$$

The B route is obtained by rotating A by $180°$ about $(63.5, 63.5)$. Bilinear sampling follows `align_corners=True`. Prescribed probes are not rerouted around obstacles; solid validity is handled by the experiment's mask.

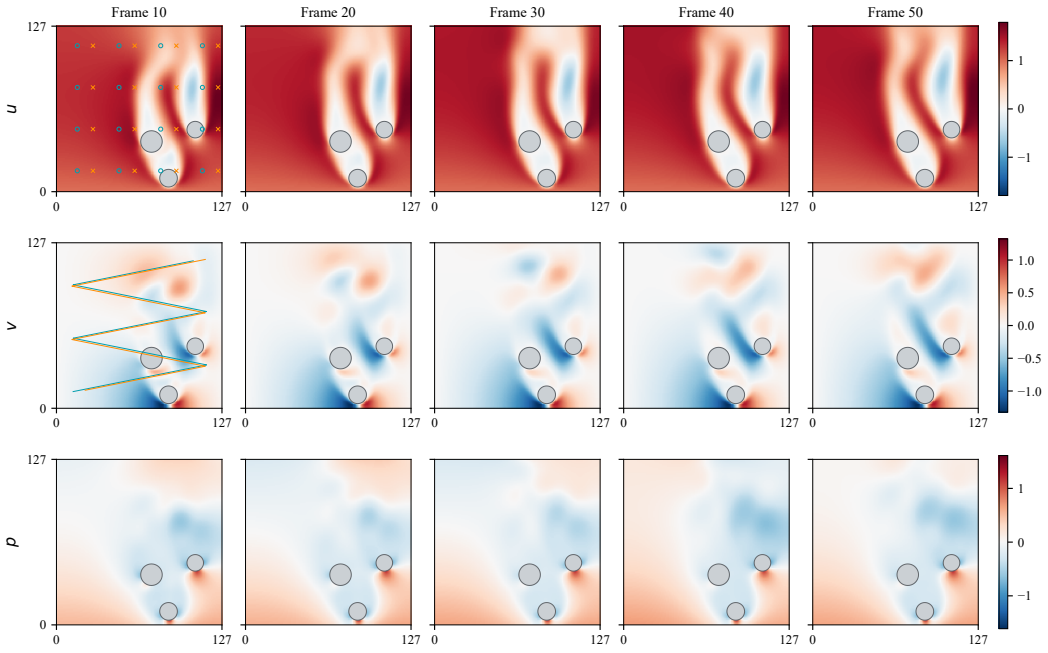

Figure 5: NS test case 11 ($Re = 68$), frames 10, 20, 30, 40, 50. Rows: u, v, p; axes: grid indices. First-column overlays show uniform A/B arrays on u and scanning A/B paths on v. Gray disks are the stored obstacles.

D.2 Pollutant: time-averaged fields in a statistically stationary regime

Source and physical quantities. Our urban pollutant benchmark is derived from the PPMLES dataset (Lumet et al., 2025; 2024a). Its underlying large-eddy simulation (LES) model was validated against the Mock Urban Setting Test (MUST) field experiment in Lumet et al. (2024b). We use the three-dimensional field statistics released in `ave_fields.h5`, accessed through a local node cache, and construct our own cropped Cartesian grids and simulated UAV observation trajectories.

The original dataset contains 200 AVBP simulations of neutral atmospheric flow and pollutant dispersion through an idealized urban environment comprising 120 MUST containers. The ensemble varies the inlet wind direction $\alpha \in [-90^\circ, 30^\circ]$ and friction velocity $u_* \in [0.07, 0.89]\,\mathrm{m\,s^{-1}}$. The simulations use the WALE subgrid-scale closure (Nicoud & Ducros, 1999) and a turbulent Schmidt number of 0.6. The inlet mean velocity follows

$$U(z) = \frac{u_*}{\kappa} \log\left(\frac{z + z_0}{z_0}\right), \tag{45}$$

where $\kappa = 0.4$, $z_0 = 0.045\,\mathrm{m}$, and the horizontal wind direction is $(\cos\alpha, \sin\alpha)$.

For physical interpretation, a constant-density filtered representation of the governing equations is

$$\partial_i \bar{u}_i = 0, \tag{46}$$

$$\partial_t \bar{u}_i + \partial_j(\bar{u}_i \bar{u}_j) = -\rho^{-1}\partial_i \bar{p} + \nu \partial_{jj} \bar{u}_i - \partial_j \tau_{ij}^{\mathrm{SGS}}, \tag{47}$$

$$\partial_t \widetilde{c} + \nabla \cdot (\widetilde{\mathbf{u}}\widetilde{c}) = \nabla \cdot [(D + D_t)\nabla \widetilde{c}] + S_c. \tag{48}$$

This representation explains the physical quantities rather than reproducing AVBP's compressible discretization. Here, the bars on velocity and pressure denote LES filtering.

Time-averaged reconstruction task. The released statistics are computed over a 200-s analysis window after spin-up to a statistically stationary regime. Statistical stationarity does not imply that the instantaneous flow ceases to fluctuate. Our three reconstruction targets are the mean concentration, concentration RMS fluctuation, and turbulent kinetic energy:

$$\bar{c} = \langle c \rangle_t, \qquad c_{\mathrm{rms}} = \sqrt{\langle (c - \bar{c})^2 \rangle_t}, \qquad k = \frac{1}{2}\sum_{i=1}^{3} \left\langle \left(u_i - \langle u_i \rangle_t\right)^2 \right\rangle_t. \tag{49}$$

Our task reconstructs spatial statistics of pollutant dispersion, rather than instantaneous turbulent dynamics. Using these statistical fields is therefore consistent with the reconstruction target. Within each case, the observation sequence indexes successive sensing locations in a fixed statistical field, not successive physical snapshots. The simulated UAV samples these fields; it does not represent an actual flight measuring the instantaneous pollutant plume.

The published concentration quantities use parts per million by volume (ppmv), and TKE uses $\mathrm{m^2\,s^{-2}}$. Because the full concentration-unit conversion is not documented in the local cache, concentration visualizations retain cache-value labels rather than asserting a verified ppmv scale.

Cartesian representation and data split. The released fields are defined on an unstructured tetrahedral mesh with 1,878,585 nodes. We resample the cached node values onto a regular Cartesian grid within

$$\Omega_{\mathrm{crop}} = [0, 70] \times [0, 90] \times [1.5, 2.5]\,\mathrm{m}, \tag{50}$$

using spacings $(\Delta x, \Delta y, \Delta z) = (0.5, 0.5, 0.1)\,\mathrm{m}$. The vertical interval selects a near-ground layer at heights of 1.5–2.5 m. This operation is spatial resampling, not an additional LES computation or a source of new physical information.

For a target point $\mathbf{x}$ inside a source tetrahedron, we use barycentric interpolation:

$$\mathcal{I}q(\mathbf{x}) = \sum_{a=1}^{4} w_a(\mathbf{x}) q(v_a), \qquad \sum_{a=1}^{4} w_a(\mathbf{x}) = 1, \tag{51}$$

where v_a are the tetrahedron vertices. The cropped grid contains 280,731 locations, of which 250,756 lie in the fluid domain. Requiring at least 97% node availability in the local cache retains 176 cases, which are split into 123 training, 26 validation, and 27 test cases using random seed 2026.

Negative interpolated mean concentrations are clipped to zero before training normalization:

$$\bar{c}_{\mathrm{used}}(\mathbf{x}) = \max\{\mathcal{I}\bar{c}(\mathbf{x}), 0\}. \tag{52}$$

Interpolation is considered valid only when all four source vertices are available. Solid regions and unavailable values are masked out during observation sampling, training, and evaluation; stored zero placeholders are not treated as measured zero concentration.

Interfaces A/B. The reference trajectory $\mathcal{A}$ is constructed from an obstacle-repaired tilted ellipse centered at $(35, 45)\,\mathrm{m}$, with semiaxes of 45 and 20 m and its long axis running from the upper-left to the lower-right of the horizontal domain. The deployment trajectory $\mathcal{B}$ is obtained by rotating the horizontal polyline of $\mathcal{A}$ by 90° and applying obstacle repair again. Route planning uses a 0.25 m grid, 1 m building inflation, and a 0.75 m domain-boundary margin.

The altitude along each trajectory is

$$z(s_{xy}) = 2 + 0.5 \sin\left(2\pi \frac{s_{xy}}{L_{xy}} + \phi_0\right) \quad \mathrm{m}, \tag{53}$$

where s_{xy}/L_{xy} is normalized horizontal arclength and ϕ_0 is the phase offset. The trajectories are resampled at equal three-dimensional arclength intervals. Their lengths are 239.487 m for $\mathcal{A}$ and 243.958 m for $\mathcal{B}$. Each contains 240 distinct sampling positions plus a repeated closure endpoint, with mean step lengths of 0.998 and 1.017 m, respectively, and approximately 1 m clearance from buildings.

Visualization. Figure 6 presents top-down views of horizontal slices at $z = 1.5$, 2.0, and $2.5\,\mathrm{m}$. These are slices of three-dimensional statistical fields, not averages over the vertical interval. The trajectories overlaid on the first panel are horizontal projections of the three-dimensional sensing paths.

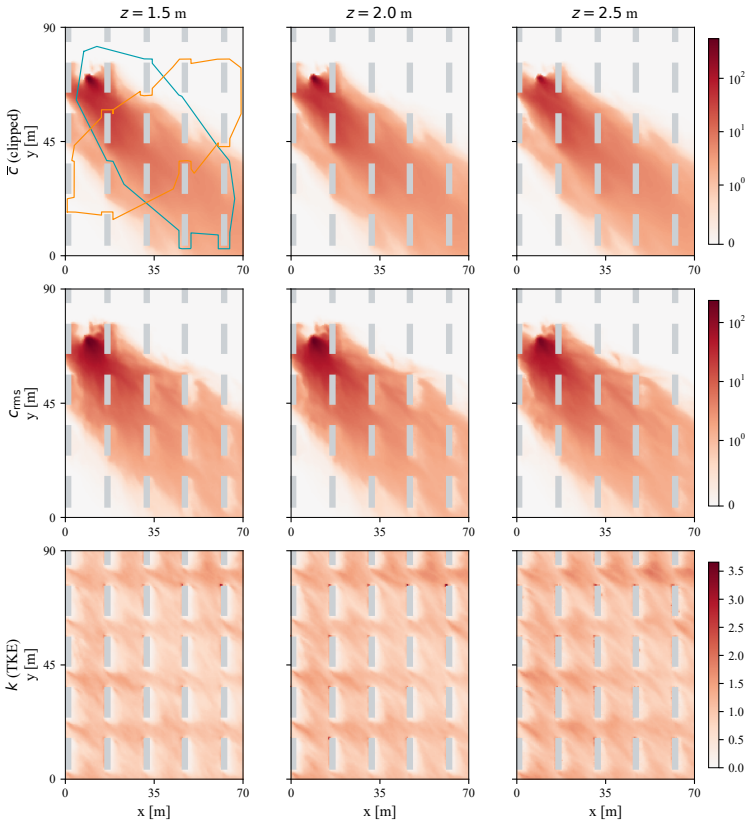

Figure 6: Top-down horizontal slices of pollutant case 000. Columns correspond to heights of $z = 1.5$, 2.0, and $2.5\,\mathrm{m}$; rows show clipped mean concentration, concentration RMS fluctuation, and TKE. Gray denotes solid geometry. The first panel overlays horizontal projections of the three-dimensional reference and deployment trajectories. Concentration color scales are linear below one cache unit and logarithmic above.

D.3 PLATE STRAIN FIELDS

Generation and quantities. Abaqus/Explicit (Dassault Systèmes SIMULIA, n.d.) simulates a $1000 \times 500 \times 10$ mm plate with 5 mm S4R shells (20,000 elements, 20,301 nodes), all six nodal degrees of freedom fixed at $x = 0$, and `NLGEOM=NO`. The linearly elastic, undamaged model obeys

$$\rho_s \ddot{\mathbf{d}} = \nabla \cdot \boldsymbol{\sigma} + \mathbf{b}, \quad \boldsymbol{\epsilon} = \tfrac{1}{2}(\nabla \mathbf{d} + \nabla \mathbf{d}^{\mathsf{T}}), \quad \boldsymbol{\sigma} = \mathsf{C}(E, \nu) : \boldsymbol{\epsilon}. \tag{54}$$

Targets are dimensionless $(\epsilon_{11}, \epsilon_{22}, \epsilon_{12})$ at the positive shell face, section point 5, in global X/Y, with tensor shear $\epsilon_{12} = E_{12}^{\text{Abaqus}}/2$. The reconstruction target is the transient two-dimensional surface strain field. Element centres form a 100×200 grid at $(2.5 + 5i, 2.5 + 5j)$ mm. A stationary normal load over a radius-12 mm region at $(500, 250)$ mm is a five-cycle Hann burst,

$$F(t) = \tfrac{F_0}{2} \sin(2\pi f t)[1 - \cos(2\pi t/T_b)], \quad T_b = 5/f, \quad 0 \le t \le T_b, \tag{55}$$

and zero afterwards. Outputs every $10\,\mu$s over 0.5 ms give 51 frames including $t = 0$. The 128 basis solves comprise three nominal 15/20/25 kHz cases and 125 Latin hypercube cases spanning $E = 190000$–215000 MPa, $\nu = 0.27$–0.33, $\rho_s = (7.65\text{–}8.05) \times 10^{-9}$ tonne/mm^3, and $f = 15$–25 kHz. We generate 6,000 synthetic strain sequences by interpolating the four nearest basis responses in normalized material/frequency parameter space, with normalized inverse-square-distance weights and linear force-amplitude scaling:

$$\boldsymbol{\epsilon}_{\text{mix}} = \frac{F_0}{5\,\text{N}} \sum_{b \in \mathcal{N}_4} \omega_b \boldsymbol{\epsilon}_b, \quad \omega_b \ge 0, \quad \sum_b \omega_b = 1, \quad F_0 \in [3, 7]\,\text{N}. \tag{56}$$

The interpolated sequences are split into 5,000/500/500 training/validation/test cases and share the 128-solve basis. Interpolation approximates the response at new parameter combinations. Reconstruction evaluation uses separately simulated held-out fields within the same parameter ranges.

Interfaces A/B. Twelve stationary footprints use $\mathcal{R}_A = \{200, 400, 600, 800\} \times \{125, 250, 375\}$ mm. The implemented deployment layout shifts each probe by 60 mm along x, keeping y fixed:

$$\mathcal{R}_B = \{(x + 60\,\text{mm}, y) : (x, y) \in \mathcal{R}_A\}. \tag{57}$$

Each averages all three components over element centres within a 10 mm radius. Gaussian noise is scaled by 0–1% of each sensor/channel's clean temporal RMS. The source remains fixed when probes move.

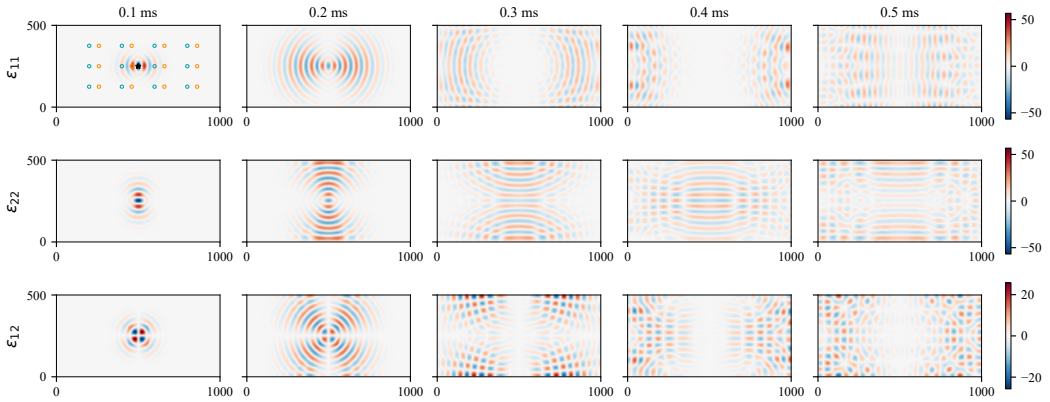

Figure 7: Independent plate case 000: $(\epsilon_{11}, \epsilon_{22}, \epsilon_{12})$ in 10^{-9} strain units; axes are mm. With initial index $k = 0$, columns are $k = 10, 20, 30, 40, 50$. A/B footprints and the fixed load (star) appear in the first panel. This is a fresh finite-element response, not a mixture.

D.4 AIRFOIL: TRANSIENT COMPRESSIBLE FLOW

Generation and quantities. The NACA0012 dataset contains 288 successful OpenFOAM v2012 `rhoCentralFoam` runs (Greenshields et al., 2010) with chord $c = 0.60105$ m, a no-slip air-

foil, freestream outer boundaries, and empty spanwise boundaries. Although its template uses SA-IDDES (Shur et al., 2008), the calculation is 2D and does not resolve 3D turbulence. Conservation equations are

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = 0, \tag{58}$$

$$\partial_t (\rho \mathbf{u}) + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + p\mathbf{I} - \boldsymbol{\tau}) = 0, \tag{59}$$

$$\partial_t (\rho E_t) + \nabla \cdot [(\rho E_t + p)\mathbf{u} - \boldsymbol{\tau} \cdot \mathbf{u} + \mathbf{q}_h] = 0, \tag{60}$$

with $E_t = e + \|\mathbf{u}\|^2/2$ and $p = \rho RT$. Targets (u, v, p, T, ρ, Ma) use m/s, m/s, Pa, K, kg/m^3, and dimensionless units. The exporter reads local $Ma = \|\mathbf{u}\|/a$, falling back to $\|\mathbf{u}\|/a_\infty$ if the solver field is absent; fallback usage is not verified across every run. Latin hypercube sampling varies attack angle 1°–7°, $Ma_\infty = 0.68$–0.80, and $\log_{10} Re = 6$–7.08. With $T_\infty = 293$ K, $p_\infty = 10^5$ Pa, $R = 287.6977$ J/(kg K), and design $\gamma = 1.4$,

$$\rho_\infty = p_\infty/(RT_\infty), \quad a_\infty = \sqrt{\gamma RT_\infty}, \quad U_\infty = Ma_\infty a_\infty, \quad \mu = \rho_\infty U_\infty c/Re. \tag{61}$$

The inflow direction follows attack angle. The perfect-gas setup uses $C_p = 1005$ J/(kg K), Prandtl number 0.71, and viscosity constant within each case. Cell-centred fields are linearly triangulated to 192×96 points in $[-0.4, 1.5] \times [-0.55, 0.55]$ m, with an STL-derived solid mask. Fifty frames span 0.0003125–0.015625 s at 0.0003125 s intervals: a short transient of roughly six to seven chord-convection times, not established stationary buffet. Splits are 256/16/16; normalization uses training fluid cells only.

Interfaces A/B. For $x(q) = x_{\mathrm{LE}} + q(x_{\mathrm{TE}} - x_{\mathrm{LE}})$ and upper/lower profiles $y_\pm(x)$, four stationary probes use

$$\mathbf{r}_\pm(q) = (x(q), y_\pm(x(q)) \pm \Delta y), \qquad \Delta y = 1.1/95 \text{ m}. \tag{62}$$

A has $q \in \{0.3, 0.7\}$ on both sides and B has $q \in \{0.4, 0.8\}$. The offset is one pixel along global y, not the surface normal. These virtual near-wall probes read six channels, rather than pressure alone at the solid surface.

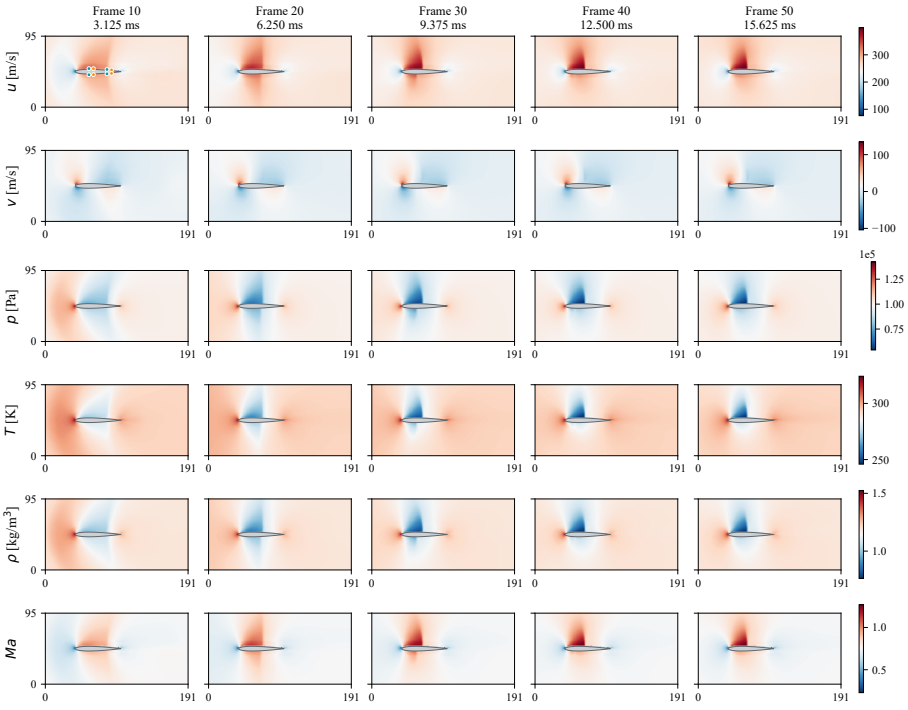

Figure 8: All six airfoil channels for test case 000 at frames 10, 20, 30, 40, 50. Axes are grid indices; colorbars retain physical units. First-panel markers show A/B near-wall probes. Each channel shares one range across its five frames.

D.5 Reconstruction losses and reported errors

In TT, the teacher optimizes virtual inputs through the frozen reconstructor using full-field targets; these targets are unavailable to the deployed student. For valid-domain mask M_n and training-normalized channel $\widetilde{q}_{n,c}$, an equal-case masked MSE is

$$\mathcal{L}_{\mathrm{MSE}} = \frac{1}{N}\sum_{n} \frac{\sum_{c,k,\mathbf{x}} M_n |\widehat{\widetilde{q}}_{n,c,k} - \widetilde{q}_{n,c,k}|^2}{\sum_{c,k,\mathbf{x}} M_n}, \qquad \widehat{\mathbf{q}} = F_A(\mathbf{z}). \tag{63}$$

The pollutant Transformer uses this equal-case reduction. Teacher objectives remain experiment-specific: the plate teacher averages physical-space relative L2 over 50 noninitial frames, whereas online TT aggregates full-field losses within each block using experiment-specific weights, with the preceding recurrent state held fixed (Equation 16). Define $\|v\|_M^2 = \sum_{\mathbf{x}} M(\mathbf{x})|v(\mathbf{x})|^2$. The channel-wise relative error, used for TT's pollutant results, is

$$E_{\mathrm{rel}} = \frac{1}{C}\sum_{c} \mathrm{mean}_{n,k} \frac{\|\widehat{q}_{n,c,k} - q_{n,c,k}\|_{M_n}}{\max(\|q_{n,c,k}\|_{M_n}, \varepsilon)}. \tag{64}$$

Evaluation undoes training normalization. Pollutant errors average its three statistical channels with no physical-time average. Channel-wise errors, joint-channel errors, and errors pooled over all frames use different reductions and are reported according to each experiment. The zero initial plate frame is excluded from per-frame relative errors.

E Finite-State Taylor–Green Experiment: Detailed Protocol

E.1 Physical fields and observations

The finite state is a trajectory identity $J \in \{0, 1, 2, 3\}$, sampled uniformly and held fixed throughout each sequence. It is not a spatial partition of one field. On the periodic domain $[0, 2\pi)^2$, set $\xi_j = x - \phi_{j,x}$ and $\eta_j = y - \phi_{j,y}$. The candidates are

$$u_j(x, y, t) = \sin\xi_j \cos\eta_j \, e^{-2\nu t}, \tag{65}$$

$$v_j(x, y, t) = -\cos\xi_j \sin\eta_j \, e^{-2\nu t}, \tag{66}$$

$$p_j(x, y, t) = \tfrac{1}{4}(\cos 2\xi_j + \cos 2\eta_j) e^{-4\nu t}. \tag{67}$$

We use $\nu = 0.04$, 20 equally spaced times in $[0, 4]$, and phases $(0, 0)$, $(0.42, 0.08)$, $(-0.12, 0.44)$, and $(0.34, -0.37)$. Fields for reconstruction evaluation are sampled on a 24×24 grid. Each measurement contains (u, v, p) with independent Gaussian noise of standard deviation 0.45 per channel, for both interfaces.

The fixed reference scan is selected before observing data from a periodic 25×25 candidate grid. The rule favors well-conditioned differences between candidate observations, with a maximum periodic displacement of 1.35 per update. Every reference evidence matrix is rank three. The displayed deployment trajectory rotates this scan by 0.72 radians about the domain center, with periodic wrapping. Additional active and random-walk policies are included in the offline student's training protocol; the displayed rotated setting is not an unseen trajectory family.

E.2 Exact posterior and translation references

Let $h^B_{t,j}$ denote the three-channel candidate observation at B_t. Starting from $\pi_0(j) = 1/4$, the target update is

$$\pi^B_t(j) = \frac{\pi^B_{t-1}(j) \exp[-\|b_t - h^B_{t,j}\|_2^2/(2\sigma^2)]}{\sum_k \pi^B_{t-1}(k) \exp[-\|b_t - h^B_{t,k}\|_2^2/(2\sigma^2)]}. \tag{68}$$

The reference filter applies the same rule at A_t to virtual input z_t. Analytic TT uses the exact evidence and the inverse in equation 34; its virtual inputs are unconstrained. The evidence-structured reference instead predicts the current likelihood ratios and applies the fixed inverse. These references use only the current observation prefix and are compared with the same prefix posterior. They are separate from the practical offline student below.

E.3 OFFLINE TEACHER–STUDENT PROTOCOL

The training set contains 1008 observation sequences (336 per policy), with 96 development and 300 fresh test sequences per policy. All splits share the same four physical candidates; they differ in sampled identities, noise, and observation sequences, not in unseen physical-field families. The teacher starts from the raw deployment measurements and uses 64 Adam steps at learning rate 0.06 to minimize mean physical-field relative ℓ_2 error over all 20 times. Each three-channel residual is bounded to RMS 2, and the best training-loss iterate, including initialization, is retained. Only training targets are used to generate teacher labels.

The student takes 11-dimensional observation, coordinate, initialization, and time features. A width-64 embedding and bidirectional GRU with 64 hidden units per direction feed a bounded residual head (59,395 parameters). It is trained for 2000 label-supervision updates, followed by development comparison of two 1000-update continuations: label-only and joint label plus field-loss training, both at learning rate 10^{-4} with matched sampling. The joint continuation was selected using development results before test evaluation. Three student seeds are 11, 29, and 47. This toy student protocol is reported explicitly rather than assumed identical to every physical benchmark.

E.4 METRICS AND VISUALIZATION

Figure 9 presents the candidate states, trajectories, and the archived results. Prefix KL evaluates $D_{\mathrm{KL}}(\pi_t^B \| \widetilde{\pi}_t)$ using identical available evidence. The analytic maximum absolute KL is 5.28×10^{-16}; the mean prefix KL is 2.4181 for No TT and 0.02232 for the structured reference. For the rotated setting, mean reconstruction relative ℓ_2 errors are 0.4346 (No TT), 0.1094 (offline student), and 0.1399 (structured prefix). Prefix and full-sequence conditional-mean references are 0.1333 and 0.06890, respectively; their evidence differs, and they are not universal relative-error lower bounds. Curves use 300 test sequences, with pointwise 95% Student-t intervals across three training seeds where applicable. The KL axis is linear below 0.05 and logarithmic above. Exact posterior matching is validated by the analytic reference; reconstruction improvement of the learned student is not claimed to establish exact posterior recovery.

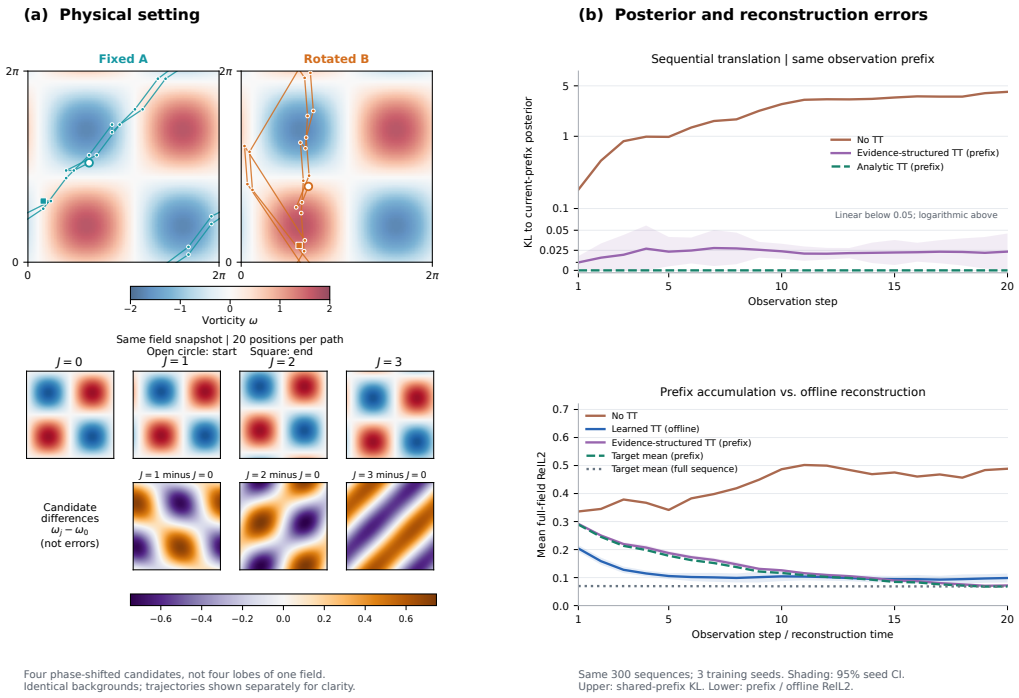

Four phase-shifted candidates, not four lobes of one field.
Identical backgrounds; trajectories shown separately for clarity.

Same 300 sequences; 3 training seeds. Shading: 95% seed CI.
Upper: shared-prefix KL. Lower: prefix / offline RelL2.

Figure 9: Finite-state experiment in detail. Left: reference/deployment trajectories and four phase-shifted candidate fields. Right: same-prefix posterior KL and reconstruction errors with explicitly distinguished prefix and full-sequence information. Candidate differences are not reconstruction errors.

F Additional Results

F.1 Full-field reconstruction visualizations

Plate strain. Figure 10 compares direct deployment observations with learned student TT using a frozen Transformer. We retain the previously fixed independent test sample 000 and final frame (zero-based index 50). The displayed component is ε_{22}, selected for the compact illustration after inspecting all three strain components. It illustrates spatial structure and residual errors. No test-time teacher optimization is used.

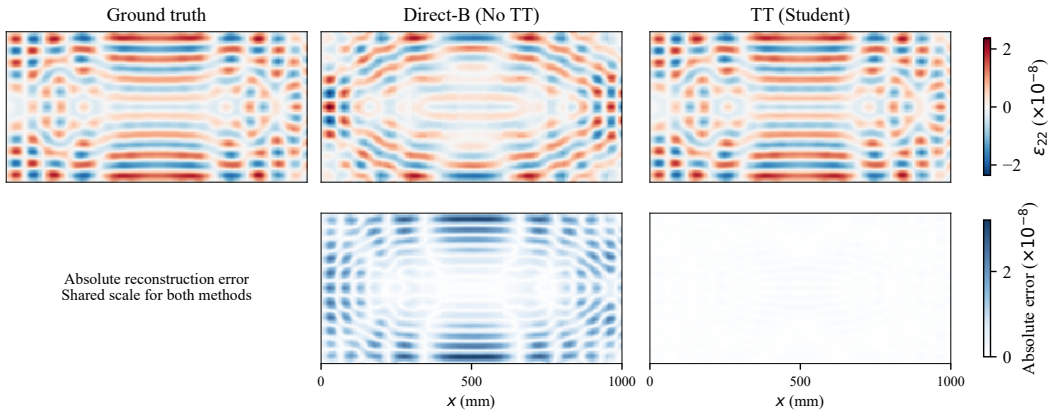

Figure 10: Plate ε_{22} reconstruction: ground truth, Direct-B (No TT), and Student TT. The lower row shows absolute errors with a common scale. Predictions are from the archived evaluation; no additional fitting or optimization is performed for visualization.

Airfoil flow. Figures 11 and 12 show the full spatial domain of the first Airfoil validation case at frames 1, 25, and 50, using the frozen Transformer and its learned offline TT. Ground truth, Native A, No TT, interpolation, and TT share a color scale within each field plot. Direct B denotes reconstruction without TT on the deployment interface $\mathcal{B}$. Native A uses genuine reference observations. The predictions come from the same evaluation as the Airfoil–Transformer main-table result.

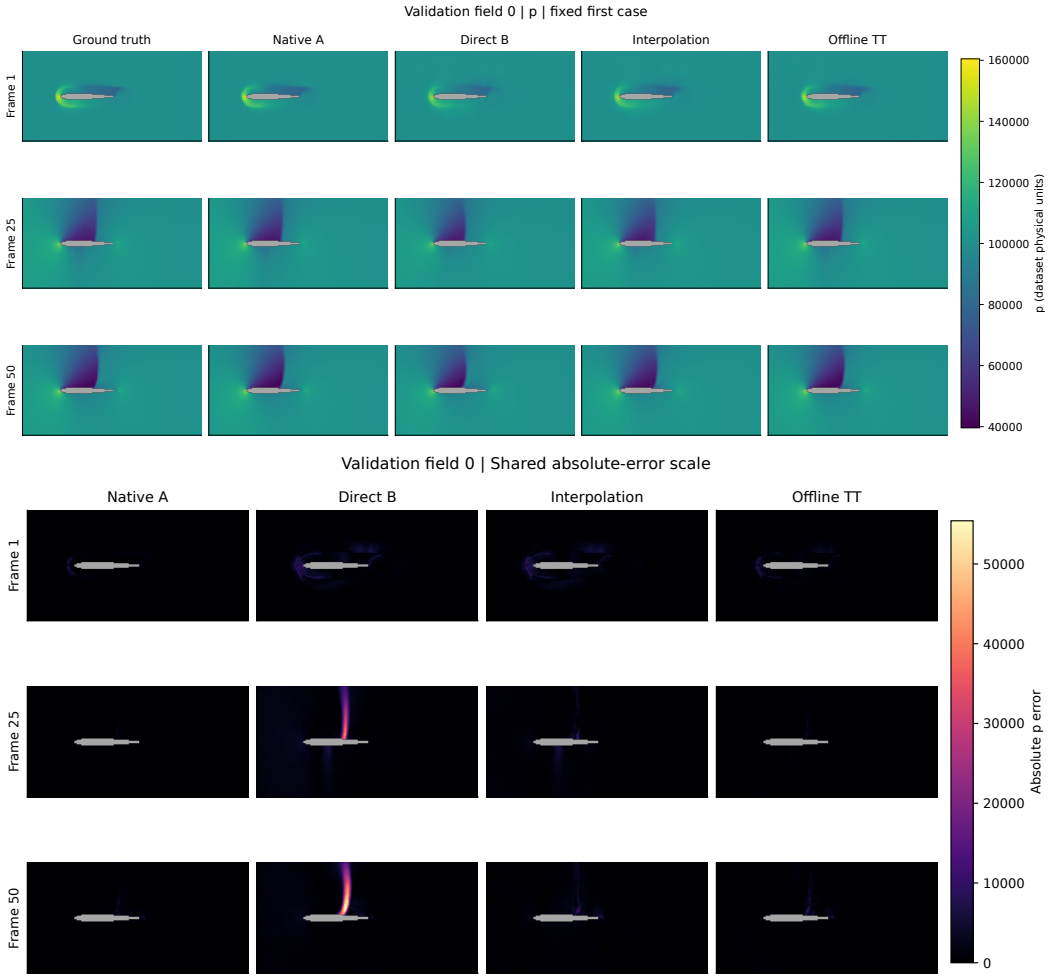

Figure 11: Airfoil pressure reconstruction and absolute error at three frames. Gray regions mask the solid airfoil. Field values share one scale across methods and frames; absolute errors share a separate scale.

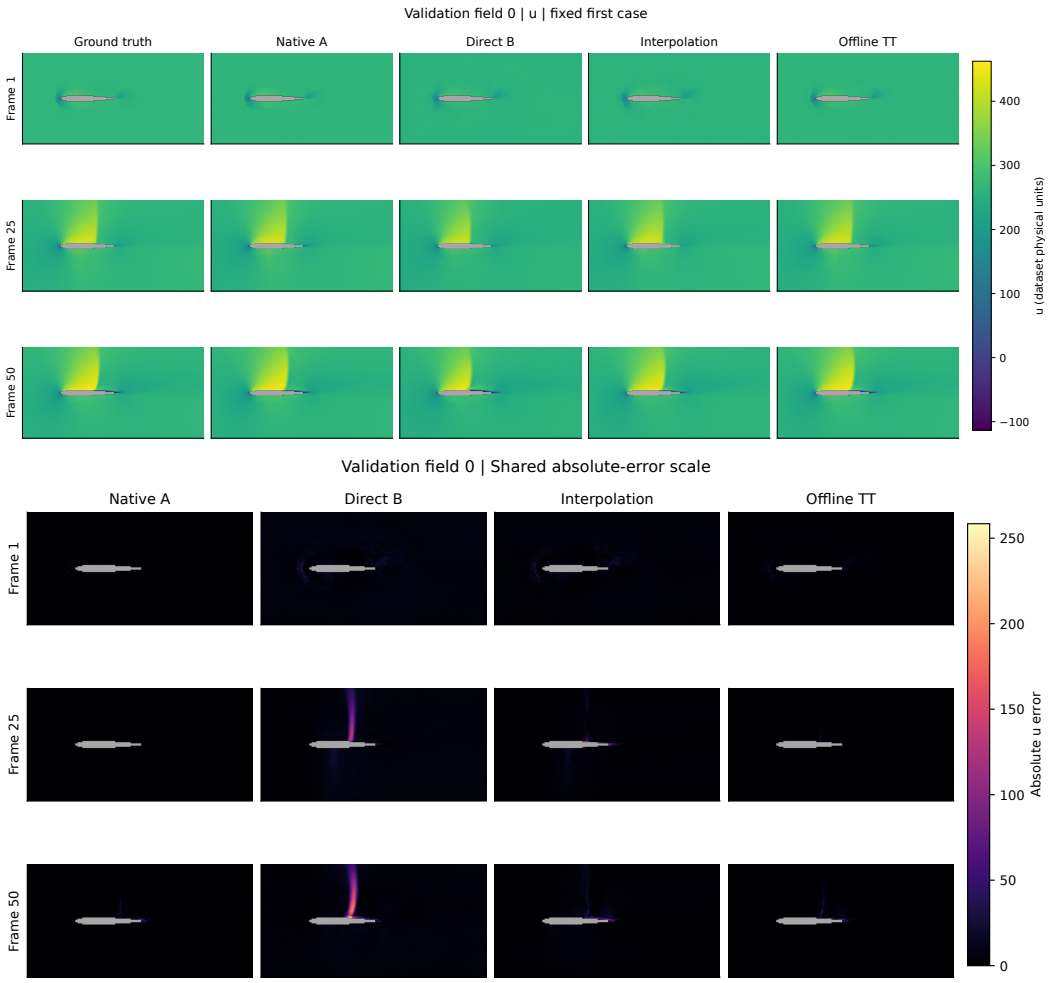

Figure 12: Airfoil streamwise velocity reconstruction and absolute error for the same case, frames, and methods as Figure 11.

Pollutant reconstruction. Figures 13 and 14 show reconstructed fields and absolute errors for the frozen Pollutant Transformer. We visualize concentration, concentration-fluctuation RMS, and turbulent kinetic energy on the horizontal slice $z = 2.0\,\mathrm{m}$ of case 048, whose TT channel-averaged relative ℓ_2 error is the median among the 27 evaluation cases. Ground truth, No TT, and TT share a color scale within each physical quantity; the two error maps likewise share a scale. Gray regions indicate invalid cells. The channel-averaged relative ℓ_2 error over the complete three-dimensional field decreases from 0.5028 without TT to 0.1646 with TT.

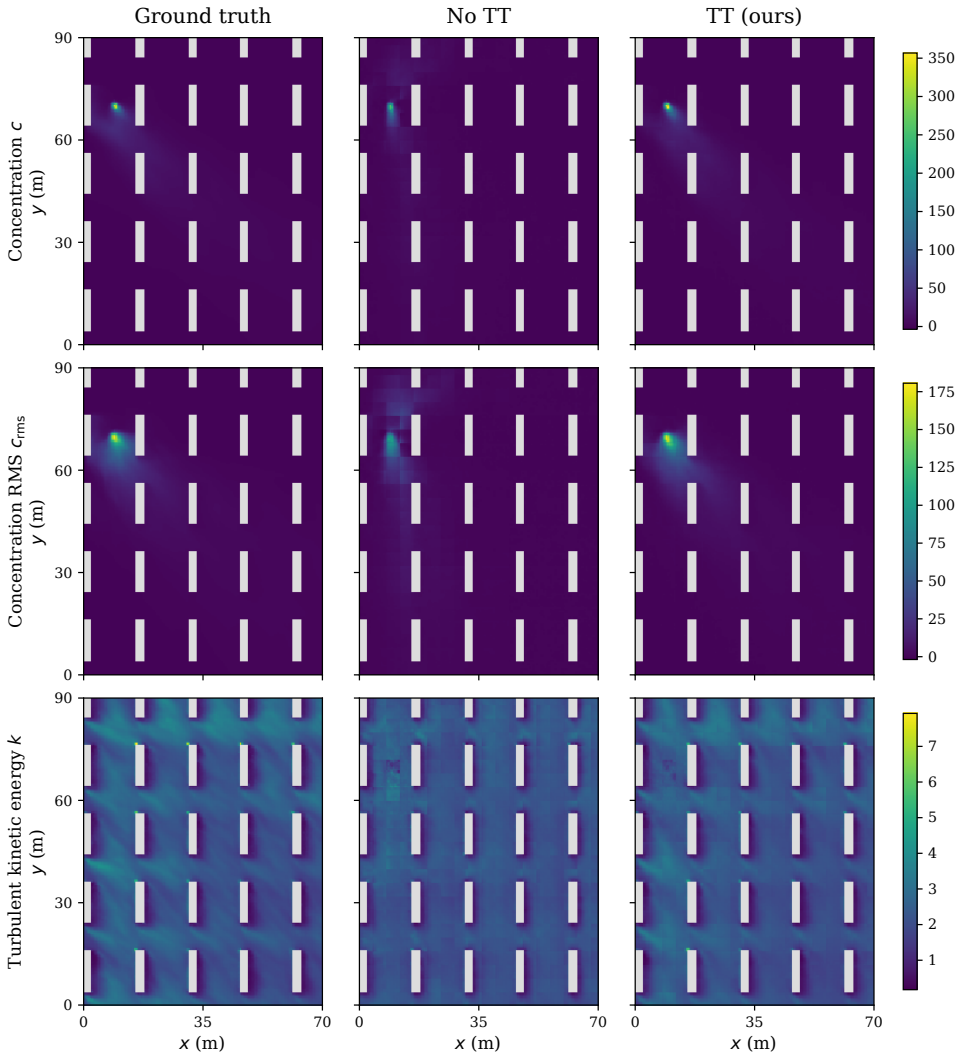

Figure 13: Pollutant reconstruction using the frozen Transformer for case 048 at $z = 2.0\,\mathrm{m}$. Columns show ground truth, No TT, and TT (ours). Rows show concentration, concentration-fluctuation RMS, and turbulent kinetic energy. Each row shares a color scale in the original data units; gray regions indicate invalid cells.

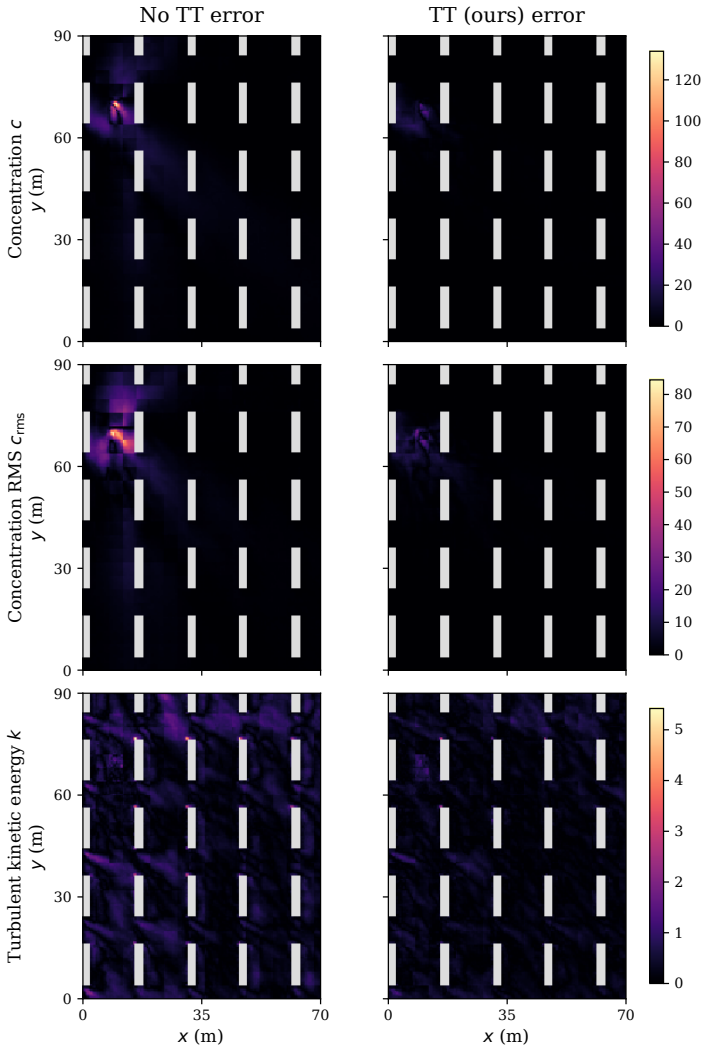

Figure 14: Absolute reconstruction errors for the same case and slice as Figure 13. Columns show No TT and TT (ours). Rows show concentration, concentration-fluctuation RMS, and turbulent kinetic energy. Errors are $|\widehat{X} - X|$ in the original data units. Each row shares a color scale across both methods; gray regions indicate invalid cells.

F.2 Ablations

F.2.1 Supervision on genuine reference observations

A-supervised trains a translator to predict genuine observations on the reference interface $\mathcal{A}$, using training fields to construct its targets. TT instead uses reconstruction-oriented teacher targets. Both feed virtual observations to the frozen reconstructor; genuine deployment-time reference observations are not supplied to either translator. Table 4 collects the corresponding results from the same evaluation settings as Table 1. TT has lower reconstruction error in 16 of 18 reporting rows; A-supervised is better in NS–Scan Transformer and Mobile-SHRED, NS–Uniform RU-CFM, and Airfoil Transformer. Thus, matching genuine reference observations is a useful alternative objective, but need not yield the lowest reconstruction error. This is an empirical objective comparison, not an assertion that all training curricula, checkpoint selection, and compute budgets are identical. In particular, the NS–Scan RU-CFM A-supervised run starts from scratch, and the NS FNO A-supervised runs do not reproduce TT's DAgger curriculum. The pollutant A-supervised translator is shared across the three frozen reconstructors. Plate FNO–128 and FNO–512 share baseline results and are not independent replications.

Table 4: Supervision-target comparison (relative ℓ_2 error; lower is better). Bold indicates the better of A-supervised and TT. Gain is $(\text{A-sup.} - \text{TT})/\text{A-sup.} \times 100\%$; negative values favor A-supervision. Metrics and splits follow Table 1.

Setting	Reconstructor	No TT	A-supervised	TT(ours)	Gain
NS–Scan	Transformer	0.6185	0.3805	**0.3735**	+1.8%
NS–Scan	Mobile-SHRED	0.6607	0.3801	**0.3755**	+1.2%
NS–Scan	FNO	0.5042	0.3625	**0.3449**	+4.9%
NS–Scan	RU-CFM	0.4299	0.3896	**0.3774**	+3.1%
NS–Uniform	Transformer	0.4416	0.2475	**0.2312**	+6.6%
NS–Uniform	SHRED	0.4213	0.3033	**0.2944**	+2.9%
NS–Uniform	FNO	0.3051	0.2927	**0.2882**	+1.5%
NS–Uniform	RU-CFM	0.4432	**0.2616**	0.2966	-13.4%
Airfoil ($\times 10^{-2}$)	Transformer	2.1350	**0.3880**	0.4030	-3.9%
Airfoil ($\times 10^{-2}$)	SHRED	2.1480	0.5160	**0.5070**	+1.7%
Airfoil ($\times 10^{-2}$)	FNO	2.8420	0.6170	**0.5510**	+10.7%
Plate	Transformer	1.3907	0.5804	**0.1538**	+73.5%
Plate	SHRED	1.2912	0.6320	**0.3751**	+40.6%
Plate	FNO–128	1.7894	0.8935	**0.7434**	+16.8%
Plate	FNO–512	1.7894	0.8935	**0.6897**	+22.8%
Pollutant	Transformer	0.5910	0.2786	**0.1703**	+38.9%
Pollutant	SHRED	1.0785	0.3352	**0.2303**	+31.3%
Pollutant	FNO	0.9744	0.3473	**0.2427**	+30.1%

Temporal supervision comparison. Figure 4(a) uses independent_000 from the Plate–Transformer independent test set. This setting has the largest aggregate relative improvement over A-supervision in Table 4. Among the three archived complete paired sequences (000, 007, and 037), we select the largest full-sequence, all-channel relative improvement for illustration. Its global relative ℓ_2 errors are 0.2771 (A-supervised) and 0.0730 (TT). The plot recomputes spatial-and-channel relative error separately at each frame from the original predictions, with identical ground truth and sensing interfaces verified for both methods. All 50 nonzero-target frames are shown without smoothing; the identically zero initial frame has undefined relative error and is omitted. The selected curve does not estimate average-case performance; the aggregate result is reported in the table. No additional training, teacher optimization, or prediction fitting is performed.

F.2.2 Teacher optimization budget

With student architecture, data, and training updates fixed, increasing the teacher budget from 8 to 16, 32, 64, and 128 Adam steps decreases mean student relative ℓ_2 error from 0.5515 to 0.5193, 0.4849, 0.4630, and 0.4591. This controlled NS–Scan Transformer study uses 256 training fields, four trajectory scales, 8000 student updates, and three seeds, distinct from the main training budget. Figure 15 retains the original numerical results and error bars.

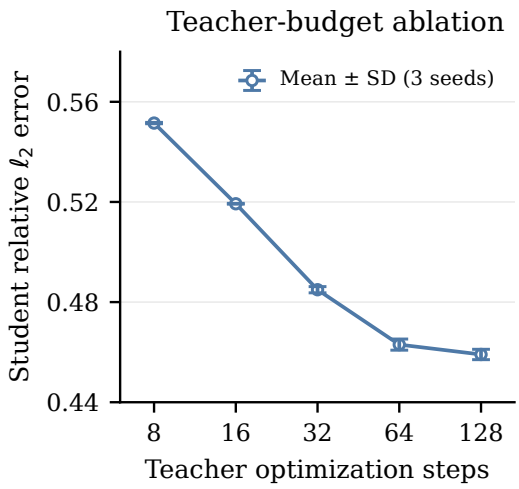

Figure 15: Teacher-budget ablation. Error bars show one standard deviation over three student-training seeds.

F.3 Mean Squared Reconstruction Error

F.3.1 Original-scale MSE across the main evaluation settings

Table 5 reports squared reconstruction errors for the same No TT and learned TT outputs evaluated by relative error in Table 1. For case i, channel c, frame t, and valid spatial domain Ω_{itc}, we compute

$$\widehat{\mathcal{R}}_{\mathrm{MSE}} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{T_i C}\sum_{t=1}^{T_i}\sum_{c=1}^{C}\frac{1}{|\Omega_{itc}|}\sum_{q\in\Omega_{itc}}(\widehat{X}_{itcq} - X_{itcq})^2. \tag{69}$$

The two methods share targets and valid domains. We average spatial points, channels, frames, and cases in that order, without dividing by each case's target energy. Original-scale values depend on channel units and amplitudes; their magnitudes are not comparable across datasets, and no cross-dataset MSE average is reported.

Evaluation scope and provenance. NS uses all 50 frames and the full grid; Scan RU-CFM has 13 nonoverlapping development fields, while the other NS rows have 128. Airfoil uses all 50 frames and the fluid mask: 16 validation fields for Transformer and 16 test fields for SHRED/FNO. Plate uses all 51 frames, including the zero initial field. Pollutant uses 27 test cases and the original three-dimensional valid masks, excluding solids and unavailable values. Its target is a static statistical field, so $T_i = 1$ here.

Sixteen rows are computed directly from complete saved predictions. Airfoil Transformer is replayed with the original frozen reconstructor and student; its original per-case metrics are reproduced within 10^{-6}, without training. Plate Transformer MSE is recovered from each case's saved global relative ℓ_2 error and target energy, divided by its element count; the identity is verified against all saved No TT predictions and the three available complete TT predictions. It does not recover channel-wise errors.

Table 5: Original-scale MSE for all 18 main-table reporting rows. Reduction is relative to No TT. Development and validation sets retain their original roles; they are not relabeled as independent tests. TT lowers sample-average MSE in all 18 rows.Original-scale MSE is sensitive to channel units and magnitudes, so channels with larger numerical scales can dominate the aggregate error. For Airfoil and Pollutant, we therefore also report MSE normalized by fixed training-set channel scales in Table 6.

Setting	Reconstructor	No TT	TT(ours)	Reduction
NS–Scan	Transformer	2.7312×10^{-1}	1.0876×10^{-1}	60.2%
NS–Scan	Mobile-SHRED	2.9462×10^{-1}	0.9310×10^{-1}	68.4%
NS–Scan	FNO	1.7784×10^{-1}	1.0140×10^{-1}	43.0%
NS–Scan	RU-CFM	1.3307×10^{-1}	9.9429×10^{-2}	25.3%
NS–Uniform	Transformer	1.1062×10^{-1}	3.4286×10^{-2}	69.0%
NS–Uniform	SHRED	1.0463×10^{-1}	5.0944×10^{-2}	51.3%
NS–Uniform	FNO	6.4841×10^{-2}	5.4813×10^{-2}	15.5%
NS–Uniform	RU-CFM	1.0782×10^{-1}	5.0490×10^{-2}	53.2%
Airfoil	Transformer	9.0778×10^{5}	4.0068×10^{4}	95.6%
Airfoil	SHRED	9.6004×10^{5}	6.6499×10^{4}	93.1%
Airfoil	FNO	1.4095×10^{6}	6.1411×10^{4}	95.6%
Plate	Transformer	1.1072×10^{-16}	1.3546×10^{-18}	98.8%
Plate	SHRED	9.5430×10^{-17}	8.0558×10^{-18}	91.6%
Plate	FNO–128	1.8330×10^{-16}	3.1631×10^{-17}	82.7%
Plate	FNO–512	1.8330×10^{-16}	2.7230×10^{-17}	85.1%
Pollutant	Transformer	5.0235×10^{1}	3.4234×10^{0}	93.2%
Pollutant	SHRED	2.4976×10^{2}	5.5666×10^{0}	97.8%
Pollutant	FNO	7.8044×10^{1}	5.6122×10^{0}	92.8%

F.3.2 Fixed-scale normalization and the risk interpretation

For Airfoil and Pollutant, we also report the verified fixed-training-scale results in Table 6. Replace each residual in equation 69 by $(\widehat{X}_{itcq} - X_{itcq})/s_c$, using the same training-derived $s_c > 0$ for both methods and all evaluation cases. Airfoil uses its original training standard deviations. Pollutant uses training RMS scales $(12.091879, 11.649949, 2.729463)$ for concentration, concentration-fluctuation RMS, and turbulent kinetic energy, respectively. This is fixed-scale MSE, not squared relative ℓ_2 error or a normalization by test-case target energy.

Table 6: MSE after fixed training-scale normalization, for the six settings with verified channel-wise normalized results. Evaluation cases and masks are identical to Table 5.

Setting	No TT	TT(ours)	Reduction
Airfoil Transformer	0.1083	0.0053	95.1%
Airfoil SHRED	0.1187	0.0075	93.6%
Airfoil FNO	0.1726	0.0076	95.6%
Pollutant Transformer	0.3806	0.0292	92.3%
Pollutant SHRED	1.8165	0.0462	97.5%
Pollutant FNO	0.8314	0.0493	94.1%

Connection to Theorem 1. At a fixed joint law of deployment evidence and target, the theorem gives

$$\mathcal{R}(f_0) - \mathcal{R}(f_{\mathrm{TT}}) = \mathbb{E}\|f_0(E) - m(E)\|_2^2 - \mathbb{E}\|f_{\mathrm{TT}}(E) - m(E)\|_2^2. \tag{70}$$

Thus a decrease in population squared risk equals a decrease in squared distance to the conditional mean, because the irreducible term cancels. For fixed diagonal scaling $D = \mathrm{diag}(s_c^{-1})$, the same orthogonality argument applies to DX:

$$\mathbb{E}\|D(f(E) - X)\|_2^2 = \mathbb{E}\|D(X - m(E))\|_2^2 + \mathbb{E}\|D(f(E) - m(E))\|_2^2.$$

Spatial and temporal averaging only adds the corresponding component risks; for masked domains the identity is conditional on the evaluation-domain metadata as well. Randomized decoders require the independence condition stated in Appendix B.1.

Tables 5 and 6 provide empirical squared-error results consistent with this interpretation.

F.4 CONTROLNET UNDER DIFFERENT PARAMETER BUDGETS

Table 7 compares TT with parameter-matched and full-size ControlNet adapters. Parameter counts refer to trainable adapters and exclude the frozen reconstructor. The matched adapter differs from TT by less than 0.3% in every setting. TT achieves lower error than this baseline in 14 of 18 rows and lower error than the full-size variant in seven rows. The two Plate FNO rows share the same backbone and ControlNet results.

In the seven rows where TT outperforms the full-size variant, the ControlNet adapter has 1,254,400 parameters for NS–Scan FNO, 17,499,648 for NS–Uniform Transformer, 2,214,400 for NS–Uniform FNO, 1,600,000 for Airfoil FNO, 1,817,216 for Plate FNO, and 2,514,048 for Pollutant FNO. These correspond to approximately 21–250 times the TT parameter counts. The full-size adapter improves many results but does not eliminate TT's advantage in all settings. Its copied-block architecture differs from the parameter-matched bottleneck adapter; implementation details are given in Appendix G.4.

Table 7: ControlNet capacity comparison. Errors follow the conventions of Table 1; lower is better. Bold and underline mark the lowest and second-lowest errors across the three adapters, respectively. Parameter counts exclude frozen backbones.

		Trainable parameters			Relative reconstruction error		
Physical field	Reconstructor	TT(ours)	Matched	Full-size	TT(ours)	Matched	Full-size
NS–Scan	Transformer	51,075	51,072	5,979,648	0.3735	<u>0.3676</u>	**0.2897**
	Mobile-SHRED	59,395	59,395	950,018	0.3755	<u>0.3649</u>	**0.2923**
	FNO	59,395	59,394	1,254,400	**0.3449**	0.4678	<u>0.3706</u>
	RU-CFM	21,967	21,903	68,987,952	<u>0.3774</u>	0.3842	**0.3353**
NS–Uniform	Transformer	69,872	69,842	17,499,648	**0.2312**	0.3105	<u>0.2377</u>
	SHRED	69,872	69,832	5,764,138	<u>0.2944</u>	0.3007	**0.2770**
	FNO	71,920	71,862	2,214,400	**0.2882**	0.3386	<u>0.3068</u>
	RU-CFM	119,024	118,962	101,243,952	0.2966	<u>0.2126</u>	**0.1640**
Airfoil ($\times 10^{-2}$)	Transformer	57,240	57,235	10,126,848	0.4030	<u>0.2795</u>	**0.2112**
	SHRED	57,240	57,233	3,153,194	<u>0.5070</u>	0.6866	**0.4453**
	FNO	57,240	57,236	1,600,000	**0.5510**	0.7724	<u>0.6008</u>
Plate	Transformer	67,044	66,997	14,254,592	<u>0.1538</u>	0.3857	**0.0591**
	SHRED	67,044	67,040	4,537,624	<u>0.3751</u>	0.5490	**0.2000**
	FNO–128	67,044	66,992	1,817,216	**0.7434**	0.9991	<u>0.9788</u>
	FNO–512	67,044	66,992	1,817,216	**0.6897**	0.9991	<u>0.9788</u>
Pollutant	Transformer	51,267	51,195	11,408,256	<u>0.1703</u>	0.2004	**0.1324**
	SHRED	51,267	51,252	2,456,942	<u>0.2303</u>	0.2397	**0.1409**
	FNO	51,267	51,255	2,514,048	**0.2427**	0.5626	<u>0.3235</u>

F.5 TRANSLATION ACROSS SENSOR LAYOUTS AND COUNTS

We further evaluate TT between the Scan and Uniform observation interfaces on NS, where the number of sensors changes between 1 and 16 per frame. We consider both Uniform16 $\rightarrow$ Scan1 and Scan1 $\rightarrow$ Uniform16, with the deployment interface mapped to the native interface of a frozen Transformer. As shown in Table 8, TT also supports translation between these interfaces, achieving relative ℓ_2 errors of 0.3032 and 0.3271, respectively, on 128 development fields.

F.6 OBSERVATION CONSISTENCY

We evaluate whether the reconstructed fields agree with the actual deployment observations on interface B. For both No TT and TT, we sample the reconstructed field at the deployment interface and compute the mean squared discrepancy from the observed values, using each dataset's fixed channel

Table 8: Relative ℓ_2 reconstruction error for cross-sensor interface translation on NS. Arrows indicate deployment B to native A. Native A is a reference using genuine native observations; direct No TT input is incompatible with the reconstructor's fixed input dimensions.

Setting ($B \to A$)	Reconstructor	Native A	No TT	TT (ours)
Uniform16 $\to$ Scan1	Transformer	0.2672	—	0.3032
Scan1 $\to$ Uniform16	Transformer	0.1715	—	0.3271

normalization scales. Table 9 reports the results for the 18 settings considered in Table 1. TT reduces observation error in 17 of the 18 settings; NS–Scan with RU-CFM exhibits a small increase.

Table 9: Observation consistency measured by normalized MSE at the deployment interface B. Lower is better. Bold marks the lower error in each row. Values are averaged over the corresponding evaluation cases.

Physical field	Reconstructor	No TT	TT (ours)
NS–Scan	Transformer	1.5141	**0.1027**
	Mobile-SHRED	1.6793	**0.0925**
	FNO	0.8658	**0.6223**
	RU-CFM	**0.5456**	0.5492
NS–Uniform	Transformer	0.5803	**0.1467**
	SHRED	0.5282	**0.2177**
	FNO	0.2667	**0.2315**
	RU-CFM	0.5814	**0.2157**
Airfoil	Transformer	1.0790	**0.0431**
	SHRED	0.4482	**0.0494**
	FNO	0.3996	**0.0588**
Plate	Transformer	2.8810	**0.0347**
	SHRED	2.8010	**0.2282**
	FNO–128	5.5480	**1.0360**
	FNO–512	5.5480	**0.7706**
Pollutant	Transformer	0.1597	**0.0114**
	SHRED	0.2606	**0.0245**
	FNO	0.3348	**0.0301**

G Implementation Details

Dataset, sensing, normalization, and metric definitions are in Appendix D. Adapter counts exclude frozen reconstructors. Plate FNO–128/512 share the reconstructor and baseline checkpoints.

G.1 Frozen reconstructors

Transformer. For both NS configurations, the observation embedding is a two-layer MLP with intermediate width 128 and output width 256. It receives the flattened sensor values and coordinates. A sinusoidal temporal embedding is followed by six causal Transformer encoder layers, each with width 256, eight attention heads, feed-forward width 1,024, and zero dropout. A linear projection and a transposed-convolution decoder with channel widths 256, 128, and 64 reconstruct the three-channel 128×128 field, using group normalization and a final convolutional projection. For Airfoil, four six-channel observations and their two-dimensional coordinates give a 32-dimensional input per frame. A linear embedding feeds six causal encoder layers with width 256, eight heads, and feed-forward width 1,024. The decoder projects to $128 \times 6 \times 12$, then applies four stride-two transposed convolutions with output channels 64, 32, 16, and 6 (kernel 4, padding 1), with GELU between stages, yielding a $6 \times 96 \times 192$ field.

The Plate Transformer uses width 512, four encoder layers, six decoder layers, eight heads, and feed-forward width 2,048. Its input is the 36-dimensional vector from 12 three-channel sensors. The model uses Fourier time and spatial embeddings with six and eight frequency bands, respectively. Spatial patch size is 10×10, giving 200 spatial query tokens on the 100×200 grid; each output token predicts 300 strain values. The checkpoint used for adaptation is the 100,000-update checkpoint. The Pollutant Transformer has width 576, six heads, two observation encoder layers, four field decoder layers, feed-forward width 2,304, and zero dropout. Observation tokens contain three values and three coordinates, with invalid tokens masked. Geometry-conditioned query tokens use $2 \times 8 \times 8$ patches; padding and cropping preserve the $11 \times 181 \times 141$ output grid. The selected checkpoint is at 40,000 updates. Its 240 observation indices represent a spatial route through a static field, not physical time.

SHRED and Mobile-SHRED. All SHRED-family models use a two-layer LSTM with hidden width 64 and a fully connected decoder with intermediate widths 350 and 400, ReLU activations, and zero dropout. The final linear layer outputs the full field. NS–Scan and NS–Uniform have input widths 3 and 48, respectively, and output width $3 \cdot 128 \cdot 128$. Airfoil uses input width 24 and output width $6 \cdot 96 \cdot 192$; Plate uses input width 36 and output width $3 \cdot 100 \cdot 200$. For these transient fields, each reconstructed frame is decoded from its causal LSTM prefix. The NS–Scan model uses the prescribed moving-probe sequence. The Pollutant model uses seven features per route position: three values, three normalized coordinates, and one validity indicator. Its final recurrent state is decoded to $3 \cdot 11 \cdot 181 \cdot 141$ values. Invalid measurements and their coordinates are zeroed before encoding.

Fourier neural operator. The two-dimensional FNOs use four spectral layers of width 32 with 12 retained modes along each spatial axis. Each layer adds a learned 1×1 local convolution; GELU follows the first three layers. A $32 \to 128 \to C$ pointwise projection outputs the field. Sparse observations are bilinearly splatted to a grid, with accumulated values divided by interpolation coverage. The input includes the value and coverage grids for the causal observation prefix, zero-filled future slots, two coordinate channels, and a time channel. This gives 203 input channels for NS, 353 for Airfoil, and 207 for Plate. The Pollutant FNO instead uses trilinear splatting and seven input channels (three values, coverage, and three coordinates). Its four three-dimensional spectral layers have width 16 and retained modes $(4, 12, 12)$, with four Fourier sign combinations. A $16 \to 32 \to 3$ pointwise head outputs the static field. Geometry masks are applied by the experiment wrapper.

G.2 Reconstructor training and checkpoint selection

The backbones are pretrained on their native observation interfaces before any adapter is trained. NS Transformer and FNO training uses AdamW (Loshchilov & Hutter, 2017) with learning rate 10^{-4} and weight decay 10^{-5}, with a 100,000-update training budget. NS SHRED training uses Adam with initial learning rate 10^{-3}, 500 warmup updates, and cosine decay, also for 100,000 updates. Airfoil models use AdamW with learning rate 10^{-4} and weight decay 10^{-5}; full-field training losses use normalized fluid-cell errors. Plate models use AdamW with initial learning rate 2×10^{-4}, weight decay 10^{-4}, 2,000 warmup updates, cosine decay, and gradient clipping at 1. Their training objective combines channel-relative squared error, cosine discrepancy, and normalized squared error, as implemented in the archived training code. Plate SHRED/FNO training batches contain four cases and four sampled frames per case. Pollutant SHRED and FNO use Adam with initial learning rate 3×10^{-4}, batch size 16, 500 warmup updates, and a 100,000-update budget; the learning rate decays to 3×10^{-5}. Their objective is equal-case masked MSE in the normalized channel space. Pollutant Transformer uses AdamW at 10^{-4} with zero weight decay, batch size 16, 1,000 warmup updates, cosine decay over 100,000 updates, and BF16 autocasting; validation is performed every 10,000 updates. Training budgets and selected checkpoints are distinct: for example, NS–Uniform Transformer uses the 90,000-update checkpoint and Pollutant Transformer uses the 40,000-update checkpoint.

Evaluation partitions. NS deterministic models and NS–Uniform RU-CFM use 128 development cases; NS–Scan RU-CFM uses 13 non-overlapping development cases. Airfoil Transformer uses 16 validation cases, and Airfoil SHRED/FNO use 16 test cases. Plate uses held-out physical solves, and

Pollutant uses 27 test cases. Comparisons preserve these partition roles. Archived configurations record case indices and checkpoint selection.

G.3 TRANSLATOR IMPLEMENTATION

Offline students encode the complete deployment sequence using a single bidirectional GRU with hidden width 64 per direction, preceded by a learned feature projection. Inputs contain deployment measurements, A/B coordinates, sequence position, and the deployable initialization; validity indicators are included where needed. Output heads predict bounded residuals around that initialization. Transient inputs contain 50 frames for NS and Airfoil and 51 for Plate; Pollutant contains 240 spatial samples. A complete offline forward pass is therefore not a per-frame call.

The Plate student has a $97 \to 64$ input projection, the bidirectional GRU, and a head with layer normalization followed by $128 \to 64 \to 36$ linear layers and GELU. Its last layer is initialized to zero. For each three-channel sensor vector, the residual is $2\sqrt{3}\, q/\sqrt{1+\|q\|_2^2}$. The Pollutant student uses a $14 \to 64$ SiLU embedding and a $128 \to 3$ output head, initialized to zero. Residual vectors are projected onto a ball of radius $2\sqrt{3}$ in normalized channel space, and invalid output slots remain zero. Both constructions bound the per-observation RMS adjustment by 2. NS–Scan online TT uses a $13 \to 32$ embedding, a unidirectional GRU of width 64, and a block output head; NS–Uniform online TT uses its separate 211-feature block student. Their parameter counts and all offline student counts are reported in Table 7.

Privileged supervision and optimization. Teachers optimize virtual observations through frozen reconstructors using training fields only. Students regress bounded controls from deployable inputs. For settings using a DAgger-style training procedure (Ross et al., 2011), teacher targets are regenerated under histories induced by the current student, and the student is subsequently trained on these targets. This exposes the student to supervision under its own induced histories, rather than only those from a fixed teacher pass. For Airfoil, the archived offline student training uses AdamW at 3×10^{-4}, zero weight decay, batch size 16, and 2,048 updates over 256 training cases; the teacher learning rate is 0.06. For Pollutant, all three students use 8,000 updates, effective batch size 16, learning rate 3×10^{-4}, and validation every 250 updates. The Pollutant teacher uses 64 Adam steps for Transformer and 128 for SHRED/FNO, learning rate 0.06, and RMS radius 2. Its full-field objective is the equal-case, equal-channel normalized RMSE. The teacher-budget ablation and supervision-target comparisons are reported in Appendix F.2; budgets are not assumed identical across backbones.

G.4 BASELINES AND SUPERVISION ABLATION

Native and unadapted reconstruction. Native A supplies genuine observations on the pretrained interface. No TT supplies deployment measurements under the model's existing input convention, including deployment coordinates when supported. Plate's fixed-slot models receive deployment values in their original slots. Native A measurements are unavailable to deployed translators.

Spatial interpolation. Interpolation maps deployment measurements to the original interface using the stored geometric interpolation operator, without learned weights or full-field targets. NS–Uniform and Plate use interpolation along the aligned sensor rows; Airfoil preserves the upper/lower probe grouping. For one scanning sensor, same-frame interpolation degenerates to copying its measured value to the reference slot. This is distinct from the causal spatial-temporal kernel used to initialize online Scan TT. Pollutant uses inverse-squared-distance weighting over four nearest observed locations and the experiment's validity masks.

Direct field refinement. The refinement baseline starts from the No-TT reconstructed field X_0 and optimizes field values directly, without a translator:

$$\min_{\widetilde{X}} \ \mathrm{MSE}_B(\mathcal{H}_B \widetilde{X}, b) + \mathrm{MSE}(\widetilde{X}, X_0). \tag{71}$$

Both terms use the experiment's normalized variables and valid entries. The coefficient on the anchoring term is 1. We apply 32 Adam steps with learning rate 0.03 and retain the minimum total objective, including the initialization. Deployment ground truth is not used to choose an iterate. For

online RU-CFM, refinement acts on the current output field, using current-frame evidence, and does not modify the recursively updated state. Pollutant refinement uses a geometry-normalized trilinear grid sampling operator; its observation data were originally obtained by tetrahedral interpolation, so the refinement operator is a grid surrogate. Timing includes both initial reconstruction and the 32 field updates.

ControlNet. We use task-specific conditional residual adapters under the ControlNet label. Both variants preserve the original frozen path through zero-initialized output projections, but their internal parameterizations differ.

Parameter-matched adapter. For a frozen-layer output h_l and an observation-only context c, the adapter adds

$$\Delta h_l = U_l \, \mathrm{SiLU}(D_l h_l + C_l \, \mathrm{SiLU}(Pc)) \,, \tag{72}$$

where affine biases are implicit. P is shared across injection sites, D_l is a low-rank down projection, C_l projects the context, and U_l is the zero-initialized up projection. This bottleneck variant does not copy pretrained blocks. The allocation routine searches shared context widths 1–128, greedily allocates layer ranks under the TT parameter budget, and selects the largest feasible count without using reconstruction scores. Shared weights are counted once; every counted parameter participates in the forward pass. The resulting counts differ from TT by less than 0.3% in all reported settings.

Full-size adapter. The larger variant trains copied source blocks and condition projections, with zero-initialized output connections back to the frozen network. In the generic runner, each copied block reads that frozen block's inputs, rather than forming a second fully independent encoder. Its correction is a zero-projected sum of the copied-block feature and the projected observation context. The dedicated Plate Transformer implementation uses its copied encoder pathway. Consequently, the capacity comparison in Appendix F.4 compares two parameterizations as well as two parameter budgets; it is not a width-only scaling experiment.

Injection sites and available evidence. Transformer adapters act at encoder layers; SHRED adapters act at the LSTM and the two intermediate decoder linear layers; FNO adapters act at its four spectral outputs. Offline context flattens the deployment observation features for the full sequence or route. For RU-CFM, entries beyond the end of the current five-frame block are zeroed. Context contains no target field. In the SHRED implementation the last-layer hidden state receives the corresponding final-sequence residual, while the cell state is retained.

Table 10: Parameter-matched ControlNet architecture. Context dimension is the flattened observation-only feature width, h is the shared projection width, and ranks follow injection-site order. Repeated ranks are written as rank $\times$ number of sites. Plate FNO–128/512 share one adapter.

Setting	Reconstructor	Context	h	Bottleneck ranks
NS–Scan	Transformer	550	36	9×6
NS–Scan	SHRED	550	35	8, 22, 26
NS–Scan	FNO	550	106	$2, 1 \times 3$
NS–Uniform	Transformer	8,050	1	$20 \times 3, 19 \times 3$
NS–Uniform	SHRED	6,450	2	7, 35, 38
NS–Uniform	FNO	8,050	8	$25 \times 3, 24$
Airfoil	Transformer	3,250	5	$13 \times 4, 12 \times 2$
Airfoil	SHRED	3,250	13	2, 8, 10
Airfoil	FNO	3,250	15	$26 \times 3, 25$
Plate	Transformer	4,947	5	$10 \times 3, 9$
Plate	SHRED	4,947	5	9, 25, 28
Plate	FNO	4,947	12	24×4
Pollutant	Transformer	2,400	15	6×2
Pollutant	SHRED	2,400	10	3, 16, 18
Pollutant	FNO	2,400	21	$4 \times 2, 3 \times 2$

Training protocol. Both capacity variants use 10,000 adapter updates, seed 20260925, AdamW at initial learning rate 10^{-4}, weight decay 10^{-4}, and gradient clipping at 1. The generic NS/Airfoil/Pollutant runner uses 100 warmup updates and cosine decay to one tenth of the peak learning rate. Checkpoints are selected at the fixed final budget. NS and Airfoil sample one case

and four frames per update; Plate samples four cases and four frames; Pollutant samples one static case; RU-CFM samples one case and one target frame. The losses are physical per-frame relative ℓ_2 for NS, normalized masked MSE for Airfoil, physical per-frame relative ℓ_2 for Plate, and equal-case/channel normalized RMSE for Pollutant. Adaptation training uses 9,000 cases for NS–Scan Transformer/SHRED and NS–Uniform Transformer, 1,000 for NS–Uniform SHRED, 9,846 for both NS FNO models, 256 for Airfoil, 5,000 for Plate, and 123 for Pollutant. Scan RU-CFM uses its 9,000-field training support; Uniform RU-CFM uses 5,500 training fields.

A-supervised ablation. A-supervised replaces optimized teacher controls with genuine training-time observations on A, while retaining deployable Student inputs. It is a supervision ablation, not a main-table baseline. Appendix F.2.1 records curriculum differences and the shared pollutant student.

G.5 Inference timing

Table 2 uses synchronized wall-clock timing on an NVIDIA H800 PCIe GPU, batch size 1, five warmups and 100 timed repetitions per case, randomized method order, and three cases per setting. Each call reconstructs all 50 Airfoil frames or 51 Plate frames. Inputs, geometric operators, and translator features are prepared on the GPU. Model + TT includes Student inference and reconstruction, excluding preprocessing, I/O, transfers, and teacher generation. Added time is the difference of mean pipeline times before rounding. The 0.39–0.57 ms range is per complete offline sequence, not per frame or online block.

For online TT with the NS–Scan RU-CFM reconstructor, we additionally measured inference over five-frame observation blocks. The mean runtime was 136.10 ms without TT and 138.15 ms with TT, corresponding to an additional 2.05 ms per block (approximately 1.5%). Each TT call generates one block of virtual observations, followed by five sequential reconstructor updates. These measurements include online token preprocessing, whereas the offline timings above use cached tokens. The additional time is the difference between the complete pipeline runtimes, rather than an isolated measurement of the translator's forward pass.