FLUIDWEAVE: RECONSTRUCTING PHYSICAL FIELDS FROM SPARSE SENSORS VIA GAUSSIAN PRIMITIVES

Anonymous authors
Paper under double-blind review

ABSTRACT

Reconstructing continuous physical fields from sparse surface-mounted sensors is central to scientific monitoring, engineering design, and digital-twin systems. Many neural reconstructors predict queried field values through latent features that do not themselves define a field over the target domain. To make this spatial structure explicit, we introduce FLUIDWEAVE, an end-to-end framework built around a shared spatial state that couples a continuously queryable field proposal with its residual correction. Inspired by 3D Gaussian Splatting, we instantiate this state as K anisotropic Gaussian primitives predicted from the unordered sensor set, whose centers, covariances, and amplitudes encode location, support, and field value, respectively. Their partition-of-unity aggregation forms the normalized proposal, while the full primitive state conditions the correction decoder. To analyze this proposal–correction mechanism, we derive a coverage–amplitude error bound, decompose local sensor sensitivity across the two pathways, and establish an exact identity for the residual's error reduction. Across four complementary benchmarks, from 2D cylinder flow and AirfRANS to 3D FlowBench LDC-3D and PhySense-Car, FLUIDWEAVE outperforms several strong reconstruction baselines.

1 INTRODUCTION

Reconstructing high-dimensional physical fields from sparse measurements has been studied extensively in data-driven sensing and reconstruction (Manohar et al., 2018; Callaham et al., 2019; Liu et al., 2025). We focus on the setting in which measurements are collected by sensors mounted on accessible solid surfaces, while monitoring, design, and control require continuous fields throughout the surrounding or enclosed domain. Reconstructing these domain-wide fields from sparse surface measurements is an underdetermined inverse problem typically addressed by incorporating prior knowledge about physically plausible fields.

Data-driven methods learn such priors from examples and have advanced physical-field reconstruction from sparse observations (Callaham et al., 2019; Fukami et al., 2021; Santos et al., 2023; Zhao et al., 2024; Nguyen et al., 2026). A common formulation encodes the observations into a global or query-conditioned latent representation and directly decodes field values at target coordinates (Santos et al., 2023; Luo et al., 2024). Although effective, the latent state in this formulation need not itself define a field over the target domain; the spatial allocation of sensor-derived information is instead deferred to the final decoder. Prior work on structured functional embeddings, relational representations, and local continuous features has shown that aligning intermediate representations with target structure can encode useful inductive biases (Gordon et al., 2020; Battaglia et al., 2018; Peng et al., 2020). For sparse field reconstruction, such alignment can make the spatial allocation of sensor-derived information explicit before final decoding, while producing an intermediate field that is directly queryable, independently evaluable, and reusable by the final predictor. This motivates expressing the learned prior through an explicit spatial state defined over the target domain.

Accordingly, we introduce FLUIDWEAVE, an end-to-end framework whose shared spatial state defines a continuously queryable field proposal and conditions its residual correction. We instantiate this state with K anisotropic Gaussians predicted jointly from the unordered sensor set, adopting the explicit primitive representation of 3D Gaussian Splatting (3DGS) for coordinate-based physical-field reconstruction. Each primitive carries a location, anisotropic support, field amplitude, and aggregation weight. Partition-of-unity aggregation forms the normalized proposal; the full primitive

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.

state conditions the correction decoder, and both pathways are optimized jointly end to end. To our knowledge, FLUIDWEAVE is the first framework for reconstructing domain-wide physical fields from sparse surface-mounted sensors in which an amortized, per-case Gaussian primitive state serves jointly as a normalized field proposal and as conditioning for continuous correction.

To characterize how this shared state supports reconstruction, we provide an architecture-aligned theoretical analysis of the proposal–correction mechanism. We first characterize the local dimensionality of the sensor-conditioned primitive state and bound proposal error through spatial coverage and amplitude mismatch. We then decompose first-order sensor sensitivity across the proposal and correction pathways. Finally, an exact identity relates residual alignment and magnitude to error reduction.

We evaluate FLUIDWEAVE on four public flow-reconstruction benchmarks: Cy_Taira Cylinder (Santos et al., 2023) on a regular 2D grid, AirfRANS (Bonnet et al., 2022) on irregular 2D meshes, FlowBench LDC-3D (Tali et al., 2025) on a volumetric 3D grid, and PhySense-Car (Ma et al., 2025) on unstructured 3D surfaces. Across these settings, FLUIDWEAVE achieves lower primary reconstruction error than the compared reconstruction methods under the corresponding evaluation protocols. Detailed ablation studies validate the contributions of the major architectural components, while complementary diagnostics relate the proposed analysis to observed model behavior. In summary, our contributions are threefold:

1. We introduce FLUIDWEAVE, an end-to-end framework for reconstructing domain-wide physical fields from sparse surface-mounted sensors. Its shared spatial state, instantiated with sensor-conditioned Gaussian primitives, forms a normalized field proposal and conditions continuous correction.
2. We develop an architecture-aligned analysis of the proposal–correction mechanism, characterizing proposal error, local sensor sensitivity, and correction-induced error reduction.
3. We evaluate FLUIDWEAVE on four public 2D/3D flow benchmarks spanning regular grids, irregular meshes, volumetric grids, and unstructured surfaces. It achieves the lowest primary reconstruction error among the compared methods under each benchmark protocol, while controlled ablations and diagnostics verify the contributions of its principal components.

2 RELATED WORK

Sparse-sensor physical-field reconstruction. Classical methods recover high-dimensional fields using POD or sparse representations, often together with data-driven sensor placement (Manohar et al., 2018; Callaham et al., 2019). Learning-based approaches replace fixed bases with shallow networks and Voronoi-assisted CNNs (Erichson et al., 2020; Fukami et al., 2021), attention-based set encoders (Santos et al., 2023), and spectral or spatiotemporal operator models (Zhao et al., 2024; Nguyen et al., 2026; Dang & Nguyen, 2025). Conditional and attentive neural processes cast reconstruction as set-to-function prediction (Garnelo et al., 2018a;b; Kim et al., 2019); ConvCNP constructs a translation-equivariant functional embedding through kernel-based SetConv (Gordon et al., 2020), while conditional INRs provide continuous coordinate-based reconstruction (Luo et al., 2024). More broadly, neural operators and geometry-aware PDE architectures learn field mappings over regular and irregular domains (Lu et al., 2021; Li et al., 2021; Azizzadenesheli et al., 2024; Li et al., 2023a;b; Hao et al., 2023; Wu et al., 2024). Within this landscape, FLUIDWEAVE addresses reconstruction under a prescribed sparse surface-sensor layout through a distinct intermediate representation: it amortizes a finite, per-case anisotropic Gaussian state whose normalized aggregate contributes an explicit field proposal and whose complete primitive tokens condition continuous correction. This differs from direct latent or grid decoders and from ConvCNP's kernel-smoothed functional embedding while retaining continuous queries on irregular coordinates. Data assimilation (Law et al., 2015; Carrassi et al., 2018) and differentiable sensor placement (Liu et al., 2025; Ma et al., 2025) address the complementary problems of sequential state estimation and sensor acquisition design, respectively.

Gaussian spatial representations. EWA splatting and 3D Gaussian Splatting represent visual scenes with explicit anisotropic Gaussian primitives (Zwicker et al., 2002; Kerbl et al., 2023); subsequent work improves their geometric and scale-aware treatment (Huang et al., 2024; Yu et al., 2024). For scattered spatial data, radial basis functions and partition-of-unity methods construct normalized local approximants (Wendland, 2004; Schaback & Wendland, 2006; Babuška & Melenk,

1997), while meshfree CFD uses renormalized kernels and weighted particle discretizations to improve consistency and conservation (Lanson & Vila, 2008; Liu & Liu, 2006; Gaburov & Nitadori, 2011; Hopkins, 2015). More recently, Gaussian functions have been adopted as physical-field representations: Gaussian Fluids represents velocity fields for physics-based fluid simulation (Xing et al., 2025), whereas Gaussian Particle Operator learns explicit Gaussian bases with Petrov–Galerkin modal coupling for forward PDE operator learning (Li et al., 2026). FLUIDWEAVE differs in both problem setting and computational role: it reconstructs domain-wide fields from sparse surface sensors by amortizing a per-case Gaussian state whose normalized aggregate forms an additive proposal and whose complete primitive state conditions continuous correction.

Approximation and sensitivity analysis. Classical theories of RBF approximation and partition-of-unity methods provide error estimates for approximants constructed from prescribed local bases (Wendland, 2004; Schaback & Wendland, 2006; Babuška & Melenk, 1997), while neural approximation theory characterizes the expressive power and convergence rates of learned function classes (Mhaskar, 1996; Yarotsky, 2017). Our analysis complements these perspectives by addressing the coupling specific to FLUIDWEAVE: a sensor-conditioned spatial state contributes to the prediction through both normalized aggregation and learned correction. We bound proposal error in terms of spatial coverage and amplitude discrepancy, decompose local sensor sensitivity across the proposal and correction pathways, and derive an exact identity for correction-induced error reduction.

3 FLUIDWEAVE: METHOD

3.1 PROBLEM SETTING

Let $\mathcal{D} \subset \mathbb{R}^d$ be the computational domain, $\Omega \subseteq \mathcal{D}$ the valid target domain, and $\boldsymbol{f} : \Omega \to \mathbb{R}^C$ a C-channel physical field. For each case, N sensors provide the unordered observation set

$$\mathcal{O} = \{(\boldsymbol{x}_i^{\text{obs}}, \boldsymbol{y}_i)\}_{i=1}^N, \qquad \boldsymbol{y}_i = \mathcal{H}_i[\boldsymbol{f}] + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{y}_i, \boldsymbol{\varepsilon}_i \in \mathbb{R}^{C_{\text{obs}}}, \tag{1}$$

where $\boldsymbol{x}_i^{\text{obs}} \in \mathcal{D}$ is the sensor coordinate, $\mathcal{H}_i$ is its benchmark-specific C_{obs}-component measurement functional (point evaluation or a boundary/surface trace), and $\boldsymbol{\varepsilon}_i$ is measurement noise. Let $\boldsymbol{c}$ collect any benchmark-provided geometry and operating context (and be empty for sensor-only tasks). From paired sensor–field cases, we learn the amortized inverse map $(\boldsymbol{x}, \mathcal{O}, \boldsymbol{c}) \mapsto \hat{\boldsymbol{f}}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c})$ for arbitrary $\boldsymbol{x} \in \Omega$; inference uses no target-field values.

3.2 OVERVIEW

FLUIDWEAVE is organized around a shared per-case spatial state. A permutation-invariant encoder maps $\mathcal{O}$ and the case-level part of $\boldsymbol{c}$ to a fused context $\boldsymbol{z} \in \mathbb{R}^D$, and a joint head predicts a bank $\mathcal{S}$ of K Gaussian primitives. At each query $\boldsymbol{x}$, normalized aggregation produces a field proposal $\boldsymbol{S}(\boldsymbol{x})$, while a query token attends to $\mathcal{S}$ and a coordinate decoder predicts a correction $\boldsymbol{R}(\boldsymbol{x})$. The two paths are optimized jointly:

$$\hat{\boldsymbol{f}}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c}) = \boldsymbol{S}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c}) + \boldsymbol{R}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c}). \tag{2}$$

We denote the primitive state by $\mathcal{S}$ and its field-valued proposal by $\boldsymbol{S}$; component analyses refer to the latter as the *Gaussian scaffold*. Sections 4 and 5.3 characterize the corresponding proposal, correction, and sensitivity behavior.

3.3 SENSOR-CONDITIONED GAUSSIAN STATE AND FIELD PROPOSAL

Primitive generation. The encoder applies a shared token MLP and mean/max pooling to obtain $\boldsymbol{z}_{\text{obs}} \in \mathbb{R}^{D_{\text{obs}}}$. Conditioning layers fuse it with the case-level part of $\boldsymbol{c}$ into $\boldsymbol{z} \in \mathbb{R}^D$; without additional context, $\boldsymbol{z} = \boldsymbol{z}_{\text{obs}}$ and $D = D_{\text{obs}}$. A joint head maps this context to the complete primitive state

$$\mathcal{S}(\mathcal{O}, \boldsymbol{c}) = \{(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, w_k, \boldsymbol{a}_k)\}_{k=1}^K, \tag{3}$$

where $\boldsymbol{\mu}_k \in \mathcal{D} \subset \mathbb{R}^d$ is a center, $\boldsymbol{\Sigma}_k \in \mathbb{R}^{d \times d}$ is a positive-definite covariance, $w_k \in \mathbb{R}_{>0}$ is a weight, and $\boldsymbol{a}_k \in \mathbb{R}^C$ is a field amplitude. Centers and scales are bounded to the normalized computational domain, with dataset-specific covariance parameterizations in Appendix F. All primitive parameters are generated jointly in one pass and learned end-to-end from the reconstruction objective, rather than fitting a dictionary at inference.

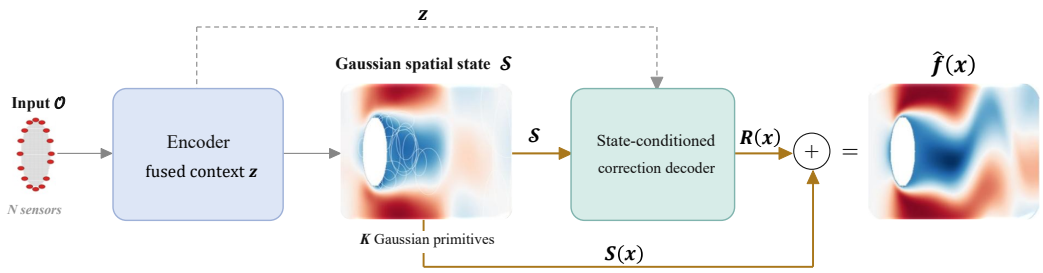

Figure 1: FluidWeave architecture. Sparse surface observations generate a per-case Gaussian spatial state. Its normalized aggregation forms the field proposal, while the same state conditions a residual correction; their sum yields the continuously queryable prediction $\hat{\boldsymbol{f}}(\boldsymbol{x})$.

Normalized field proposal. For query $\boldsymbol{x}$, suppressing the fixed dependence of the primitive parameters on $(\mathcal{O}, \boldsymbol{c})$, define

$$\varphi_k(\boldsymbol{x}) = \exp\left[-\tfrac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu}_k)^\top \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{x}-\boldsymbol{\mu}_k)\right], \qquad \psi_k(\boldsymbol{x}) = \frac{w_k\varphi_k(\boldsymbol{x})}{\sum_{j=1}^{K} w_j\varphi_j(\boldsymbol{x})}. \tag{4}$$

Writing the raw basis mass as $m(\boldsymbol{x}) = \sum_j w_j\varphi_j(\boldsymbol{x})$, the implementation uses the guarded denominator $m_{\text{num}}(\boldsymbol{x}) = \max\{m(\boldsymbol{x}), \epsilon_{\text{num}}\}$ for numerical stability; Appendix A gives its relation to the analytic normalization. Using the analytic weights in Equation (4), the proposal is

$$\boldsymbol{S}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c}) = \sum_{k=1}^{K} \psi_k(\boldsymbol{x})\boldsymbol{a}_k, \qquad \psi_k(\boldsymbol{x}) \geq 0, \quad \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) = 1. \tag{5}$$

Thus $\boldsymbol{S}(\boldsymbol{x}) \in \mathbb{R}^C$ is a normalized spatial aggregation of sensor-conditioned amplitudes. The bank can be evaluated and intervened on through its proposal, coverage, routing, and sensor response.

3.4 State-Conditioned Residual Decoder

The correction path receives both the proposal and its generating state. To make the correction query-dependent, each query attends to primitive tokens that encode the state's location, support, weight, and amplitude. Let $\gamma(\boldsymbol{x})$ collect raw and Fourier coordinate features (Tancik et al., 2020), $\boldsymbol{c}_{\boldsymbol{x}}$ be the query-local part of $\boldsymbol{c}$ (empty when unavailable), and $\xi(\boldsymbol{\Sigma}_k)$ denote the implemented covariance parameters. Define

$$\boldsymbol{q}_{\boldsymbol{x}} = E_q([\gamma(\boldsymbol{x}), \boldsymbol{z}, \boldsymbol{c}_{\boldsymbol{x}}]), \qquad \boldsymbol{t}_k = E_p([\boldsymbol{\mu}_k, \xi(\boldsymbol{\Sigma}_k), w_k, \boldsymbol{a}_k]). \tag{6}$$

Here $[\cdot]$ is concatenation; learned embeddings E_q and E_p produce $\boldsymbol{q}_{\boldsymbol{x}}, \boldsymbol{t}_k \in \mathbb{R}^{D_t}$. With the row-stacked matrix $\boldsymbol{T} = [\boldsymbol{t}_1^\top; \ldots; \boldsymbol{t}_K^\top] \in \mathbb{R}^{K\times D_t}$, the implemented cross-attention block is

$$\boldsymbol{u}_{\boldsymbol{x}} = \mathrm{LN}(\boldsymbol{q}_{\boldsymbol{x}} + \mathrm{MHA}(\boldsymbol{q}_{\boldsymbol{x}}, \boldsymbol{T}, \boldsymbol{T})), \qquad \boldsymbol{h}_{\boldsymbol{x}} = \boldsymbol{u}_{\boldsymbol{x}} + \mathrm{FFN}(\boldsymbol{u}_{\boldsymbol{x}}), \tag{7}$$

where LN, MHA, and FFN are layer normalization, multi-head attention, and a feed-forward network; $\boldsymbol{u}_{\boldsymbol{x}}, \boldsymbol{h}_{\boldsymbol{x}} \in \mathbb{R}^{D_t}$. The residual MLP combines this state-conditioned feature with the context, proposal, and guarded basis mass $m_{\text{num}}(\boldsymbol{x})$:

$$\boldsymbol{R}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c}) = D_R([\gamma(\boldsymbol{x}), \boldsymbol{z}, \boldsymbol{c}_{\boldsymbol{x}}, \boldsymbol{S}(\boldsymbol{x}), m_{\text{num}}(\boldsymbol{x}), \boldsymbol{h}_{\boldsymbol{x}}]). \tag{8}$$

Here D_R is a learned MLP with output in $\mathbb{R}^C$; every input on the right is determined by $(\boldsymbol{x}, \mathcal{O}, \boldsymbol{c})$. Appendix F specifies benchmark-dependent geometry features and the additional PhySense-Car sensor-token attention. Thus each query receives primitive location, support, and amplitude while retaining the additive proposal in Equation (2).

3.5 Training and Inference

Training. We train the complete model end-to-end with field supervision at sampled query points, a soft observation-consistency term, and, when available, an additional surface-target term,

$$\mathcal{L} = \mathcal{L}_{\text{field}} + \lambda_{\text{obs}}\mathcal{L}_{\text{obs}} + \lambda_{\text{surf}}\mathcal{L}_{\text{surf}}. \tag{9}$$

Here $\mathcal{L}_{\text{field}}$ is mean squared error at sampled domain queries, $\mathcal{L}_{\text{obs}}$ applies the operators $\mathcal{H}_i$ to the prediction, and $\mathcal{L}_{\text{surf}}$ uses separate surface queries on AirfRANS and is omitted on datasets without such targets. The nonnegative scalars λ_{obs} and λ_{surf} weight the two auxiliary losses. Dataset-specific query sampling, conditioning, and optimization are reported in Appendix F. The main models are trained on clean sensor values; controlled test-time input perturbations are evaluated separately.

Inference. One sensor encoding and primitive prediction support any number of query coordinates. The same coordinate-query architecture can therefore be evaluated on a regular grid, an irregular mesh, a volumetric grid, or an unstructured surface without changing the coordinate-query interface.

4 Architecture-Aligned Analysis

We analyze the implemented sensor-value path $\boldsymbol{y} \mapsto \boldsymbol{z} \mapsto \mathcal{S}$ and $(\boldsymbol{y}, \boldsymbol{z}, \mathcal{S}) \mapsto (\boldsymbol{S}, \boldsymbol{R}) \mapsto \hat{\boldsymbol{f}}$, with sensor coordinates, geometry, and operating metadata fixed. Here $\boldsymbol{y} \in \mathbb{R}^M$ stacks the $M = NC_{\text{obs}}$ scalar sensor values and $\boldsymbol{z} \in \mathbb{R}^D$ is the fused context; the explicit $\boldsymbol{y}$ argument includes any direct sensor-token conditioning of the residual. The analysis characterizes proposal error, perturbation propagation, and residual correction.

Shared state and proposal error. For a fixed case, suppress $(\mathcal{O}, \boldsymbol{c})$ and consider $\boldsymbol{x} \in \Omega$ with $m(\boldsymbol{x}) \geq \epsilon_{\text{num}} > 0$. Define the local coverage radius

$$\rho(\boldsymbol{x}) = \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{x} - \boldsymbol{\mu}_k\|_2. \tag{10}$$

Proposition 1 (Deterministic proposal error bound)**.** *If $\boldsymbol{f}$ admits an L-Lipschitz extension $\bar{\boldsymbol{f}} : \mathcal{D} \to \mathbb{R}^C$ satisfying $\bar{\boldsymbol{f}}|_\Omega = \boldsymbol{f}$, then*

$$\|\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{S}(\boldsymbol{x})\|_2 \leq L\rho(\boldsymbol{x}) + \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{a}_k - \bar{\boldsymbol{f}}(\boldsymbol{\mu}_k)\|_2. \tag{11}$$

The terms separate spatial coverage from amplitude–target discrepancy with respect to $\bar{\boldsymbol{f}}$; the proof and an integrated form are in Appendix B.

Proposal and full-output sensitivity. For clean and perturbed values at the same case and query, let $\boldsymbol{S} = \sum_k \psi_k \boldsymbol{a}_k$, $\boldsymbol{S}' = \sum_k \psi'_k \boldsymbol{a}'_k$, and $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_K)^\top$.

Proposition 2 (Finite perturbation bound)**.** *If both weight sets are exactly normalized and $A_{\max} = \max_{1 \leq k \leq K} \|\boldsymbol{a}_k\|_2$, then*

$$\|\boldsymbol{S}'(\boldsymbol{x}) - \boldsymbol{S}(\boldsymbol{x})\|_2 \leq \max_{1 \leq k \leq K} \|\boldsymbol{a}'_k - \boldsymbol{a}_k\|_2 + A_{\max} \|\boldsymbol{\psi}'(\boldsymbol{x}) - \boldsymbol{\psi}(\boldsymbol{x})\|_1. \tag{12}$$

The terms separate amplitude variation from changes in geometry and normalized routing. At almost every differentiable input,

$$\frac{\partial \boldsymbol{S}}{\partial \boldsymbol{y}} = \sum_{k=1}^{K} \psi_k \frac{\partial \boldsymbol{a}_k}{\partial \boldsymbol{y}} + \sum_{k=1}^{K} (\boldsymbol{a}_k - \boldsymbol{S})(\nabla_{\boldsymbol{y}} \psi_k)^\top. \tag{13}$$

Let $\boldsymbol{J}_S = \partial \boldsymbol{S}/\partial \boldsymbol{y}$, $\boldsymbol{J}_R = \partial \boldsymbol{R}/\partial \boldsymbol{y}$, and $\boldsymbol{J}_{\hat{f}} = \partial \hat{\boldsymbol{f}}/\partial \boldsymbol{y}$ denote the sensor-value Jacobians. At one query each lies in $\mathbb{R}^{C \times M}$; over a fixed query set, stack outputs and Jacobians vertically. Because the residual includes all direct and state-mediated sensor dependence,

$$\boldsymbol{J}_{\hat{f}} = \boldsymbol{J}_S + \boldsymbol{J}_R. \tag{14}$$

After stacking the fixed evaluation queries into $\mathbb{R}^{M_{\text{out}}}$, let $\boldsymbol{W} \in \mathbb{R}^{M_{\text{out}} \times M_{\text{out}}}$ satisfy $\boldsymbol{W} \succeq 0$ and define $\langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\boldsymbol{W}} = \boldsymbol{u}^\top \boldsymbol{W} \boldsymbol{v}$ and $\|\boldsymbol{v}\|^2_{\boldsymbol{W}} = \langle \boldsymbol{v}, \boldsymbol{v} \rangle_{\boldsymbol{W}}$ for $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^{M_{\text{out}}}$. For a small zero-mean perturbation $\boldsymbol{\varepsilon} \in \mathbb{R}^M$ with covariance $\boldsymbol{\Sigma}_\varepsilon = \mathbb{E}[\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^\top]$, first-order propagation gives

$$\mathbb{E}\|\hat{\boldsymbol{f}}(\boldsymbol{y} + \boldsymbol{\varepsilon}) - \hat{\boldsymbol{f}}(\boldsymbol{y})\|^2_{\boldsymbol{W}} \approx Q_{\hat{f}} = Q_S + Q_R + 2C_{SR}, \tag{15}$$

where

$$Q_A = \mathrm{tr}(\boldsymbol{W}\boldsymbol{J}_A\boldsymbol{\Sigma}_\varepsilon\boldsymbol{J}_A^\top), \quad A \in \{S, R, \hat{f}\}, \qquad C_{SR} = \mathrm{tr}(\boldsymbol{W}\boldsymbol{J}_S\boldsymbol{\Sigma}_\varepsilon\boldsymbol{J}_R^\top). \tag{16}$$

Thus Q_S and Q_R are pathway contributions and C_{SR} records whether their local sensitivities reinforce or offset one another; finite-noise behavior is evaluated in Section 5.3.

For the stacked outputs in $\mathbb{R}^{M_{\mathrm{out}}}$, reuse $\boldsymbol{f}$, $\boldsymbol{S}$, and $\boldsymbol{R}$ and set $\boldsymbol{e}_S = \boldsymbol{f} - \boldsymbol{S}$.

Proposition 3 (Residual correction identity)**.**

$$\begin{aligned}\Delta_{\mathrm{corr}} &= \|\boldsymbol{f} - \boldsymbol{S}\|_{\boldsymbol{W}}^2 - \|\boldsymbol{f} - (\boldsymbol{S} + \boldsymbol{R})\|_{\boldsymbol{W}}^2 \\ &= 2\langle \boldsymbol{e}_S, \boldsymbol{R}\rangle_{\boldsymbol{W}} - \|\boldsymbol{R}\|_{\boldsymbol{W}}^2.\end{aligned} \tag{17}$$

If $\|\boldsymbol{e}_S\|_{\boldsymbol{W}} > 0$ and $\|\boldsymbol{R}\|_{\boldsymbol{W}} > 0$, let $\beta = \|\boldsymbol{R}\|_{\boldsymbol{W}}/\|\boldsymbol{e}_S\|_{\boldsymbol{W}}$ denote relative correction magnitude. Let $\cos\vartheta = \langle \boldsymbol{e}_S, \boldsymbol{R}\rangle_{\boldsymbol{W}}/(\|\boldsymbol{e}_S\|_{\boldsymbol{W}}\|\boldsymbol{R}\|_{\boldsymbol{W}})$ denote alignment. Then

$$\frac{\Delta_{\mathrm{corr}}}{\|\boldsymbol{e}_S\|_{\boldsymbol{W}}^2} = 2\beta\cos\vartheta - \beta^2. \tag{18}$$

Positive Δ_{corr} means that the residual reduces proposal error; Section 5.3 reports these quantities and Appendix C gives the derivations.

5 Experiments

Experimental setup. We evaluate the same coordinate-query framework on four complementary spatial representations: Cy_Taira Cylinder (Santos et al., 2023) uses 4/8/16 wall sensors on a regular grid; AirfRANS (Bonnet et al., 2022) uses eight surface-pressure sensors on irregular meshes; FlowBench LDC-3D (Tali et al., 2025) uses 5/10/20/40 surface-mounted sensor layouts on 64^3 volumes; and PhySense-Car (Ma et al., 2025) reconstructs pressure on unstructured car surfaces. Metrics are physical-space relative L^2 (Cylinder), the AirfRANS three-channel aggregate of per-field normalized MSEs, global and channel-macro relative L^2 (LDC-3D), and released normalized-pressure relative L^2 (PhySense-Car). Main benchmark summaries use five runs; Table 5 reports three common paired runs for controlled ablations, while PhySense uses three. Full protocols and 3D settings appear in Appendix F.

Baselines. We compare POD–Ridge, DeepONet (Lu et al., 2021), Senseiver (Santos et al., 2023), RecFNO (Zhao et al., 2024), FLRNet (Nguyen et al., 2026), and Voronoi-CNN (Fukami et al., 2021). Public implementations and published architectures are evaluated using the benchmark-specific splits, sensor layouts, and evaluators summarized in Appendix G. Published PhySense values and evaluations in this study are reported in separate groups in Table 4. Together, the four tasks test whether the same proposal–correction computation applies across output discretizations and spatial dimensions; the controlled studies in Section 5.3 assess the roles of its shared state.

5.1 Main Results

Cylinder (regular grid). Table 1 reports three wall-sensor layouts under the disjoint 50/200/4,750 protocol. FluidWeave yields the lowest five-seed mean among the compared methods at all three sensor counts, supporting the sensor-to-state construction across observation densities. The full protocol and the separate Interior-8 study are in Appendices G and H.

AirfRANS (irregular meshes). Table 2 reports native-mesh reconstruction and two operating-condition shifts. FluidWeave yields the lowest mean on Full-8 and retains this ordering under the tested angle-of-attack and Reynolds-number shifts, supporting coordinate-query reconstruction on irregular domains beyond the training conditions considered. FluidWeave also has the lowest error in all three target channels on the common three-seed diagnostic subset (Appendix H).

FlowBench LDC-3D (volumetric grid). Table 3 reports four sensor counts on 64^3 four-channel volumes, with FluidWeave attaining the lowest global and channel-macro mean among six methods at every count under both the joint raw-channel and channel-balanced norms.

Table 1: Cylinder relative L^2 on the disjoint 50/200/4,750 protocol (mean $\pm$ SE over five training seeds). Protocol details are in Appendix F.2. Lower is better.

Method	Surface-4	Surface-8	Surface-16
FLUIDWEAVE (ours)	**0.0642 $\pm$ 0.0018**	**0.0358 $\pm$ 0.0033**	**0.0160 $\pm$ 0.0003**
Senseiver	0.0732 $\pm$ 0.0051	0.0563 $\pm$ 0.0029	0.0545 $\pm$ 0.0048
DeepONet	0.0925 $\pm$ 0.0015	0.1301 $\pm$ 0.0061	0.0946 $\pm$ 0.0010
RecFNO	0.0745 $\pm$ 0.0072	0.1025 $\pm$ 0.0050	0.0705 $\pm$ 0.0025
FLRNet	0.0965 $\pm$ 0.0039	0.0917 $\pm$ 0.0030	0.0880 $\pm$ 0.0036
Voronoi-CNN	0.0731 $\pm$ 0.0008	0.0419 $\pm$ 0.0045	0.0206 $\pm$ 0.0017

Table 2: AirfRANS aggregate normalized-field MSE with eight surface-pressure sensors (mean $\pm$ SE over five training seeds; POD–Ridge is deterministic). Implementation settings are reported in Appendix G. Lower is better; the best mean in each column is bold.

Method	Full-8	AoA-shift-8	Re-shift-8
FLUIDWEAVE (ours)	**0.2526 $\pm$ 0.0186**	**0.5280 $\pm$ 0.0452**	**0.3129 $\pm$ 0.0102**
FLRNet	0.2931 $\pm$ 0.0042	0.8181 $\pm$ 0.0194	0.3142 $\pm$ 0.0018
Voronoi-CNN	0.3085 $\pm$ 0.0135	0.8084 $\pm$ 0.0284	0.3366 $\pm$ 0.0160
DeepONet	0.3289 $\pm$ 0.0148	0.5937 $\pm$ 0.0443	0.3859 $\pm$ 0.0082
Senseiver	0.3891 $\pm$ 0.0139	0.7546 $\pm$ 0.0404	0.4647 $\pm$ 0.0301
RecFNO	0.3939 $\pm$ 0.0225	0.8078 $\pm$ 0.0580	0.4400 $\pm$ 0.0181
POD–Ridge	0.5849	1.2982	0.6024

Table 3: FlowBench LDC-3D global and channel-macro relative L^2 over (u, v, w, p), using Reynolds conditioning statistics computed from the 700 training cases only. Entries are means over five training seeds; SE summaries and per-channel values are in Appendix H.10. Implementation settings are provided in Appendix G. Lower is better.

	Global relative L^2				Channel-macro relative L^2			
Method	$N=5$	10	20	40	$N=5$	10	20	40
FLUIDWEAVE (ours)	**.0303**	**.0284**	**.0281**	**.0273**	**.2246**	**.2229**	**.2247**	**.2225**
DeepONet	.0590	.0465	.0493	.0483	.3211	.2889	.2942	.2956
Senseiver	.1009	.0760	.1499	.0906	.4480	.4025	.5530	.4490
RecFNO	.1144	.1219	.1328	.1238	.4522	.4535	.5096	.4821
FLRNet	.0670	.0533	.0503	.0475	.3610	.3558	.3530	.3526
Voronoi-CNN	.0820	.0691	.0691	.0489	.5653	.4709	.4867	.4635

PhySense-Car (unstructured surface). Table 4 reports 15–200 sensors on approximately 95k surface points. Among the methods evaluated in this study, FLUIDWEAVE has the lowest three-seed mean at every sensor count, extending the same sensor-to-state and coordinate-query interface to an unstructured surface.

5.2 QUALITATIVE RESULTS

Figures 2 and 3 show held-out reconstructions across all four spatial representations. On AirfRANS and Cylinder, FLUIDWEAVE recovers the dominant pressure variation and wake organization, with errors concentrated near sharp gradients. LDC streamlines reproduce the principal three-dimensional recirculation, while the PhySense examples show lower relative L^2 and pointwise errors than the released reference.

All outputs are evaluated on their benchmark's native spatial representation; Appendix Figure 5 visualizes the primitive-state support for Cylinder.

Table 4: PhySense-Car normalized-pressure relative L^2 under the released sensor-only evaluator (lower is better). Published rows are point estimates from Table 3 of Ma et al. (2025); other rows are three-seed means ± SE.

Method	$N=15$	$N=30$	$N=50$	$N=100$	$N=200$
Published results					
Senseiver	0.1053	0.1022	0.1018	0.1012	0.1009
DiffusionPDE	0.2095	0.1055	0.0987	0.0966	0.0967
PhySense	0.0465	0.0416	0.0395	0.0382	0.0375
PhySense-opt	**0.0386**	**0.0372**	**0.0370**	**0.0370**	**0.0369**
Evaluations in this study					
FLRNet	0.0196 ± 0.00013	0.0202 ± 0.00016	0.01864 ± 0.00002	0.0162 ± 0.00014	0.0155 ± 0.00008
Voronoi-CNN	0.0413 ± 0.00056	0.0378 ± 0.00069	0.0353 ± 0.00041	0.0340 ± 0.00025	0.0338 ± 0.00029
FLUIDWEAVE (ours)	$\mathbf{0.0112 \pm 0.00123}$	$\mathbf{0.0093 \pm 0.00039}$	$\mathbf{0.0085 \pm 0.00004}$	$\mathbf{0.0079 \pm 0.00020}$	$\mathbf{0.0072 \pm 0.00018}$

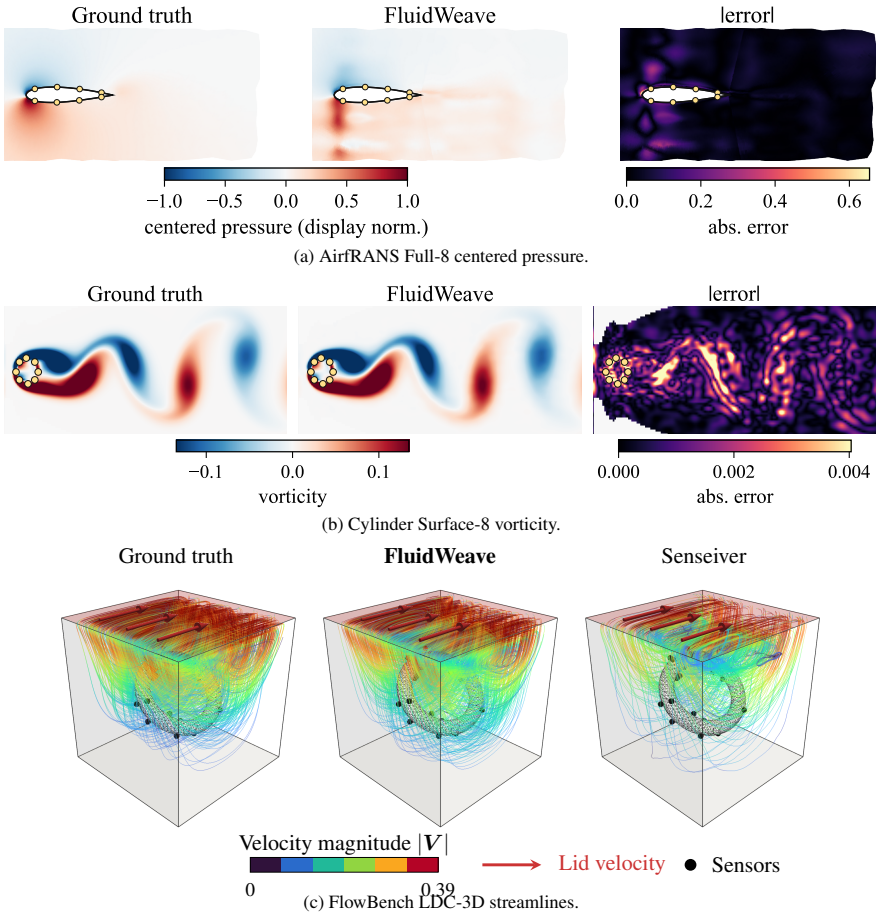

(a) AirfRANS Full-8 centered pressure.

(b) Cylinder Surface-8 vorticity.

(c) FlowBench LDC-3D streamlines.

Figure 2: Held-out AirfRANS (a), Cylinder (b), and LDC-3D (c) reconstructions; (a,b) show ground truth, FLUIDWEAVE, and absolute error, while (c) compares FLUIDWEAVE with Senseiver.

5.3 ABLATION STUDIES

Shared-state contribution. Table 5 compares representations, state interfaces, and prediction pathways on AirfRANS Full-8. Removing state features or attention raises mean MSE, supporting state-conditioned correction; removing the additive proposal raises it to 0.2682. Cross-case state intervention increases mean global relative L^2 by 0.0276, showing that predictions depend on the generated state.

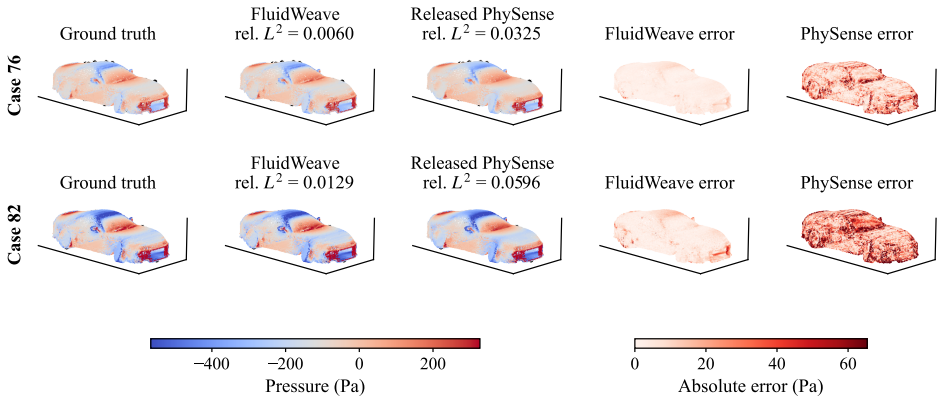

Figure 3: PhySense-Car ($N = 50$) reconstructions and errors for held-out cases 76/82.

Table 5: AirfRANS Full-8 ablations (3 paired runs, mean ± SE); 5-run aggregate in Table 2.

Reference: **FluidWeave (full)**			0.2367 ± 0.0261
Representation controls		*State-interface controls*	
Variant	MSE	Variant	MSE
Direct decoder (no spatial state)	0.3227 ± 0.0122	No global token	0.2759 ± 0.0207
Generic token routing	0.2627 ± 0.0127	No state features	0.2992 ± 0.0099
		No state attention	0.3446 ± 0.0083
Pathway controls		*Objective control*	
No additive proposal	0.2682 ± 0.0086	No obs.-consistency loss	0.2620 ± 0.0070
Proposal only	0.6524 ± 0.0091		
Residual only (broader)	0.3530 ± 0.0075		

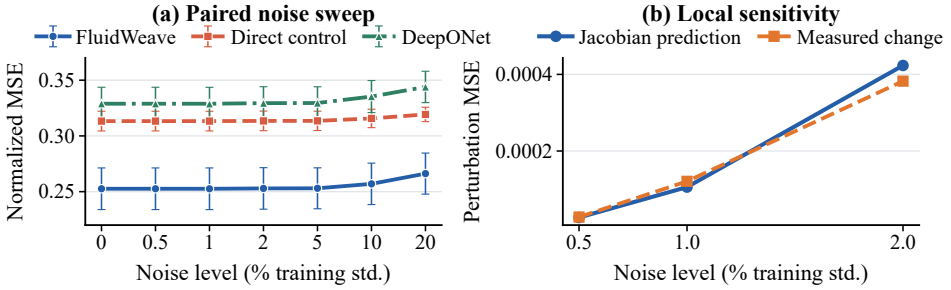

Figure 4: AirfRANS Full-8 five-seed noise sweep (mean ± SE; left) and Jacobian agreement (right).

Correction. $\Delta_{\text{corr}} > 0$ for every test case, with mean magnitude $\beta = 0.8923$ and mean alignment $\cos\vartheta = 0.9127$. **State instantiation.** Generic-token and Wendland C2 controls raise mean MSE from 0.2367 to 0.2627 and 0.2745, supporting the Gaussian instantiation. **Sensor perturbations.** At 0.5–2% noise, Jacobian predictions are within 4–14% of measured changes; FluidWeave retains the lowest absolute error through the five-seed 0–20% sweep (Figure 4; Appendix H).

6 Conclusion

We introduced FluidWeave for domain-wide field reconstruction from sparse surface-mounted sensors. Its amortized Gaussian state forms a continuous proposal and conditions a residual decoder, making the learned spatial prior explicit before decoding. Across four 2D/3D flow benchmarks, FluidWeave attains the lowest primary error among the evaluated methods; ablations support complementary proposal and correction roles. In future work, we will extend the state–proposal–correction design to additional physical fields and sensing operators.

AI USE STATEMENT

Generative-AI assistants were used for language editing and manuscript polishing. The authors reviewed all suggested edits and take full responsibility for the manuscript and its claims.

REPRODUCIBILITY STATEMENT

Reproducibility details, including dataset splits, preprocessing, key architectural and optimization settings, numbers of independent training runs, and evaluation protocols, are provided in the appendix. Complete theoretical assumptions and proofs are also included. Code and configurations will be released upon publication.

REFERENCES

Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 6(5):320–328, 2024. doi: 10.1038/s42254-024-00712-5.

Ivo Babuška and Jens M Melenk. The partition of unity method. *International journal for numerical methods in engineering*, 40(4):727–758, 1997.

Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matthew Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. URL `https://arxiv.org/abs/1806.01261`.

Florent Bonnet, Jocelyn Mazari, Paola Cinnella, and Patrick Gallinari. Airfrans: High fidelity computational fluid dynamics dataset for approximating reynolds-averaged navier–stokes solutions. *Advances in Neural Information Processing Systems*, 35:23463–23478, 2022.

Jared L Callaham, Kazuki Maeda, and Steven L Brunton. Robust flow reconstruction from limited measurements via sparse representation. *Physical Review Fluids*, 4(10):103907, 2019.

Alberto Carrassi, Marc Bocquet, Laurent Bertino, and Geir Evensen. Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *Wiley Interdisciplinary Reviews: Climate Change*, 9(5):e535, 2018. doi: 10.1002/wcc.535.

Hiep Vo Dang and Phong C. H. 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, 11 2025. ISSN 1530-9827. doi: 10.1115/1.4070332.

N Benjamin Erichson, Lionel Mathelin, Zhewei Yao, Steven L Brunton, Michael W Mahoney, and J Nathan Kutz. Shallow neural networks for fluid flow reconstruction with limited sensors. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2238): 20200097, 2020.

Kai Fukami, Romit Maulik, Nesar Ramachandra, Koji Fukagata, and Kunihiko Taira. Global field reconstruction from sparse sensors with voronoi tessellation-assisted deep learning. *Nature Machine Intelligence*, 3(11):945–951, 2021.

Evghenii Gaburov and Keigo Nitadori. Astrophysical weighted particle magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 414(1):129–154, 2011. doi: 10.1111/j.1365-2966.2011.18313.x.

Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and S. M. Ali Eslami. Conditional neural processes. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1704–1713, 2018a.

Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018b. URL https://arxiv.org/abs/1807.01622.

Vivette Girault and Pierre-Arnaud Raviart. Finite element methods for navier-stokes equations: theory and algorithms. *NASA STI/Recon Technical Report A*, 87:52227, 1986.

Jonathan Gordon, Wessel P. Bruinsma, Andrew Y. K. Foong, James Requeima, Yann Dubois, and Richard E. Turner. Convolutional conditional neural processes. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=Skey4eBYPS.

Pierre Grisvard. *Elliptic problems in nonsmooth domains*. SIAM, 2011.

Zhongkai Hao, Zhengyi Wang, Hang Su, Chengyang Ying, Yinpeng Dong, Songming Liu, Ze Cheng, Jian Song, and Jun Zhu. Gnot: A general neural operator transformer for operator learning. In *International conference on machine learning*, pp. 12556–12569. PMLR, 2023.

Philip F. Hopkins. A new class of accurate, mesh-free hydrodynamic simulation methods. *Monthly Notices of the Royal Astronomical Society*, 450(1):53–110, 2015. doi: 10.1093/mnras/stv195.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pp. 1–11, 2024.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. doi: 10.1145/3592433.

Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, S. M. Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=SkE6PjC9KX.

V Kozlov, V Maz'ya, and J Rossmann. *Spectral Problems Associated with Corner Singularities of Solutions to Elliptic Equations*, volume 85 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2000. doi: 10.1090/surv/085.

Nathalie Lanson and Jean-Paul Vila. Renormalized meshfree schemes i: Consistency, stability, and hybrid methods for conservation laws. *SIAM Journal on Numerical Analysis*, 46(4):1912–1934, 2008. doi: 10.1137/S0036142903427718.

Kody Law, Andrew Stuart, and Konstantinos Zygalakis. *Data Assimilation: A Mathematical Introduction*. Springer, 2015. doi: 10.1007/978-3-319-20325-6.

Zhihao Li, Yu Feng, Zhilu Lai, and Wei Wang. From basis to basis: Gaussian particle representation for interpretable PDE operators. *arXiv preprint arXiv:2602.21551*, 2026. doi: 10.48550/arXiv.2602.21551. URL https://arxiv.org/abs/2602.21551.

Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=c8P9NQVtmnO.

Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023a.

Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36:35836–35854, 2023b.

M. B. Liu and G. R. Liu. Restoring particle consistency in smoothed particle hydrodynamics. *Applied Numerical Mathematics*, 56(1):19–36, 2006. doi: 10.1016/j.apnum.2005.02.012.

Xu Liu, Wei Peng, Xiaoya Zhang, Xiaoyu Zhao, Weien Zhou, Wen Yao, and Xiaoqian Chen. Enhancing deep learning-based field reconstruction with a differentiable learning framework. *Nature Machine Intelligence*, 7(7):1129–1140, 2025.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.

Xihaier Luo, Wei Xu, Balasubramanya Nadiga, Yihui Ren, and Shinjae Yoo. Continuous field reconstruction from sparse observations with implicit neural networks. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=kuTZMZdCPZ.

Yuezhou Ma, Haixu Wu, Hang Zhou, Huikun Weng, Jianmin Wang, and Mingsheng Long. Physense: Sensor placement optimization for accurate physics sensing. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=zIzZxDsNNP.

Krithika Manohar, Bingni W Brunton, J Nathan Kutz, and Steven L Brunton. Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems Magazine*, 38(3):63–86, 2018.

Hrushikesh N Mhaskar. Neural networks for optimal approximation of smooth and analytic functions. *Neural computation*, 8(1):164–177, 1996.

Phong CH Nguyen, Joseph B Choi, and Quang-Trung Luu. Deep learning with fourier features for regressive flow field reconstruction from sparse sensor measurements. *Scientific Reports*, 16:5980, 2026.

Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision*, pp. 523–540. Springer, 2020. doi: 10.1007/978-3-030-58580-8_31.

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.

Robert Schaback and Holger Wendland. Kernel techniques: from machine learning to meshless methods. *Acta numerica*, 15:543–639, 2006.

Ronak Tali, Ali Rabeh, Cheng-Hau Yang, Mehdi Shadkhah, Samundra Karki, Abhisek Upadhyaya, Suriya Dhakshinamoorthy, Marjan Saadati, Soumik Sarkar, Adarsh Krishnamurthy, Chinmay Hegde, Aditya Balu, and Baskar Ganapathysubramanian. Flowbench: A large scale benchmark for flow simulation over complex geometries. *Journal of Data-centric Machine Learning Research*, 2025. URL https://openreview.net/forum?id=0xrOmmB2KQ.

Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020.

Roger Temam. *Navier–Stokes equations: theory and numerical analysis*, volume 343. American Mathematical Society, 2024.

Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.

Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast transformer solver for pdes on general geometries. In *International Conference on Machine Learning*, pp. 53681–53705. PMLR, 2024.

Jingrui Xing, Bin Wang, Mengyu Chu, and Baoquan Chen. Gaussian Fluids: A grid-free fluid solver based on Gaussian Spatial Representation. In *ACM SIGGRAPH 2025 Conference Papers*, SIGGRAPH '25, New York, NY, USA, 2025. Association for Computing Machinery. doi: 10.1145/3721238.3730620. URL https://doi.org/10.1145/3721238.3730620.

Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural networks*, 94: 103–114, 2017.

Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19447–19456, 2024.

Xiaoyu Zhao, Xiaoqian Chen, Zhiqiang Gong, Weien Zhou, Wen Yao, and Yunyang Zhang. Recfno: A resolution-invariant flow and heat field reconstruction method from sparse observations via fourier neural operator. *International Journal of Thermal Sciences*, 195:108619, 2024.

Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. EWA splatting. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):223–238, 2002. doi: 10.1109/TVCG.2002.1021576.

APPENDIX

A NOTATION AND METRIC CONVENTIONS

For one case, $\mathcal{O} = \{(\boldsymbol{x}_i^{\mathrm{obs}}, \boldsymbol{y}_i)\}_{i=1}^{N}$ denotes the observation set and $\boldsymbol{y} \in \mathbb{R}^M$, with $M = NC_{\mathrm{obs}}$, stacks all sensor values. The possibly empty $\boldsymbol{c}$ denotes fixed geometry or operating context. The sensor-set encoding is $\boldsymbol{z}_{\mathrm{obs}} \in \mathbb{R}^{D_{\mathrm{obs}}}$; conditioning layers produce $\boldsymbol{z} \in \mathbb{R}^D$, with $\boldsymbol{z} = \boldsymbol{z}_{\mathrm{obs}}$ and $D = D_{\mathrm{obs}}$ when $\boldsymbol{c}$ is empty. We reserve $\mathcal{S}$ for the primitive bank and $\boldsymbol{S}$ for its field proposal. The vectorized state is $\boldsymbol{\theta} = \mathrm{vec}(\mathcal{S}) \in \mathbb{R}^{P_\theta}$, where P_θ is its total scalar dimension. At query $\boldsymbol{x}$, $\varphi_k(\boldsymbol{x})$ is the unnormalized Gaussian response, $\psi_k(\boldsymbol{x})$ is its normalized routing weight, and $\boldsymbol{a}_k$ is its C-channel field amplitude. Consequently, $\boldsymbol{S}(\boldsymbol{x}) = \sum_{k=1}^{K} \psi_k(\boldsymbol{x})\boldsymbol{a}_k$ and $\hat{\boldsymbol{f}}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c}) = \boldsymbol{S}(\boldsymbol{x}) + \boldsymbol{R}(\boldsymbol{x})$. Dependence on $(\mathcal{O}, \boldsymbol{c})$ is suppressed when fixed.

Weighted field seminorms. Let M_{out} be the total number of scalar outputs after stacking a fixed evaluation-query set. For $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^{M_{\mathrm{out}}}$ and a positive semidefinite evaluation matrix $\boldsymbol{W} \in \mathbb{R}^{M_{\mathrm{out}} \times M_{\mathrm{out}}}$, define the weighted bilinear form and seminorm

$$\langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\boldsymbol{W}} = \boldsymbol{u}^\top \boldsymbol{W} \boldsymbol{v}, \qquad \|\boldsymbol{v}\|_{\boldsymbol{W}}^2 = \langle \boldsymbol{v}, \boldsymbol{v} \rangle_{\boldsymbol{W}}. \tag{19}$$

The residual-correction diagnostics use the same channel scaling and spatial quadrature as the destandardized squared-relative metric. Aggregate normalized-field MSE instead averages squared error over the three train-standardized target channels and then over query points and cases. The two metrics therefore quantify complementary aspects of reconstruction behavior.

Numerical normalization. Analytically, positive Gaussian responses and primitive weights give $m(\boldsymbol{x}) = \sum_{k=1}^{K} w_k \varphi_k(\boldsymbol{x}) > 0$. On all AirfRANS Full-8 diagnostic queries, $m(\boldsymbol{x}) \geq 2.5249 \times 10^{-4} > 10^{-6}$; hence the guarded implementation coincides with the analytic partition-of-unity normalization.

Differentiability. The rank and Jacobian statements apply at inputs where the implemented network is differentiable. With standard piecewise-smooth activations this holds almost everywhere. The finite-perturbation bound and correction identity are deterministic and allow dependent primitive variations.

B PROPOSAL APPROXIMATION PROPERTIES

This section establishes the representational properties of the implemented normalized proposal used in Section 4.

B.1 JOINT LATENT GENERATION

Proposition 4 (Local rank of primitive variation)**.** *Let $\boldsymbol{y} \in \mathbb{R}^M$ stack the sensor values at fixed sensor coordinates and case context, let $\boldsymbol{z} = g_{\mathrm{ctx}}(\boldsymbol{y}) \in \mathbb{R}^D$ be the fused context, and let the vectorized primitive state be $\boldsymbol{\theta} = g_{\mathrm{state}}(\boldsymbol{z}) \in \mathbb{R}^{P_\theta}$. At every point where g_{ctx} and g_{state} are differentiable,*

$$\mathrm{rank}\left(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{y}}\right) \leq \min\{M, D, P_\theta\}. \tag{20}$$

Proof of Proposition 4. At a differentiable input, the chain rule gives

$$\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{y}} = \frac{\partial g_{\mathrm{state}}}{\partial \boldsymbol{z}} \frac{\partial g_{\mathrm{ctx}}}{\partial \boldsymbol{y}}.$$

The rank of this product is bounded by the rank of either factor and by the intermediate dimension D, which proves Equation (20). The statement characterizes local per-instance primitive variation through the amortized state generator. □

B.2 Coverage–Amplitude Error Bound

Proof of Proposition 1. Because $\sum_{k=1}^{K} \psi_k(\boldsymbol{x}) = 1$ and $\psi_k(\boldsymbol{x}) \geq 0$,

$$\begin{aligned} \boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{S}(\boldsymbol{x}) &= \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \big(\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{a}_k\big) \\ &= \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \big[\bar{\boldsymbol{f}}(\boldsymbol{x}) - \bar{\boldsymbol{f}}(\boldsymbol{\mu}_k) + \bar{\boldsymbol{f}}(\boldsymbol{\mu}_k) - \boldsymbol{a}_k\big]. \end{aligned} \tag{21}$$

Here $\bar{\boldsymbol{f}}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x})$ on Ω. The triangle inequality and L-Lipschitz continuity imply

$$\begin{aligned} \|\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{S}(\boldsymbol{x})\|_2 &\leq \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\bar{\boldsymbol{f}}(\boldsymbol{x}) - \bar{\boldsymbol{f}}(\boldsymbol{\mu}_k)\|_2 + \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\bar{\boldsymbol{f}}(\boldsymbol{\mu}_k) - \boldsymbol{a}_k\|_2 \\ &\leq L \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{x} - \boldsymbol{\mu}_k\|_2 + \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{a}_k - \bar{\boldsymbol{f}}(\boldsymbol{\mu}_k)\|_2, \end{aligned} \tag{22}$$

which is Equation (11). □

The same argument yields an integrated statement for any finite measure ν on Ω. For a vector field $\boldsymbol{g} : \Omega \to \mathbb{R}^C$, we use $\|\rho\|_{L^1(\nu)} = \int_\Omega \rho(\boldsymbol{x})\, \mathrm{d}\nu(\boldsymbol{x})$ and $\|\boldsymbol{g}\|_{L^1(\nu)} = \int_\Omega \|\boldsymbol{g}(\boldsymbol{x})\|_2\, \mathrm{d}\nu(\boldsymbol{x})$:

$$\|\boldsymbol{f} - \boldsymbol{S}\|_{L^1(\nu)} \leq L \|\rho\|_{L^1(\nu)} + \int_\Omega \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{a}_k - \bar{\boldsymbol{f}}(\boldsymbol{\mu}_k)\|_2\, \mathrm{d}\nu(\boldsymbol{x}). \tag{23}$$

The bound applies directly to jointly learned centers and amplitudes and separates upper-bound terms associated with spatial coverage and amplitude prediction.

B.3 Connection to State-Conditioned Correction

Proposition 1 separates two upper-bound terms associated with proposal error. The correction $\boldsymbol{R}(\boldsymbol{x}; \mathcal{O}, \boldsymbol{c})$ conditions on the fused context $\boldsymbol{z}$ and primitive state $\mathcal{S}$ that determine coverage and amplitudes. Its realized effect is quantified by Δ_{corr} in Proposition 3 and the measurements in Section 5.3.

C Perturbation and Correction Diagnostics

Because the model jointly generates the primitive state through amortized inference, we analyze finite and local perturbations of the implemented proposal and use an exact identity to evaluate residual correction.

C.1 Finite Proposal Perturbations

Proof of Proposition 2. At a fixed query, suppress $\boldsymbol{x}$ and add and subtract $\sum_{k=1}^{K} \psi'_k \boldsymbol{a}_k$:

$$\boldsymbol{S}' - \boldsymbol{S} = \sum_{k=1}^{K} \psi'_k (\boldsymbol{a}'_k - \boldsymbol{a}_k) + \sum_{k=1}^{K} (\psi'_k - \psi_k) \boldsymbol{a}_k. \tag{24}$$

Normalization and nonnegativity give

$$\left\| \sum_{k=1}^{K} \psi'_k (\boldsymbol{a}'_k - \boldsymbol{a}_k) \right\|_2 \leq \max_{1 \leq k \leq K} \|\boldsymbol{a}'_k - \boldsymbol{a}_k\|_2,$$

while the second term is bounded by

$$\left\| \sum_{k=1}^{K} (\psi'_k - \psi_k) \boldsymbol{a}_k \right\|_2 \leq A_{\max} \|\boldsymbol{\psi}' - \boldsymbol{\psi}\|_1.$$

Combining the inequalities proves Equation (12). □

The bound applies to correlated changes in amplitudes, centers, scales, and weights generated by the shared context.

C.2 Centered Proposal Jacobian

Differentiating $\boldsymbol{S} = \sum_{k=1}^{K} \psi_k \boldsymbol{a}_k$ with respect to the stacked sensor-value vector $\boldsymbol{y} \in \mathbb{R}^M$ gives

$$\frac{\partial \boldsymbol{S}}{\partial \boldsymbol{y}} = \sum_{k=1}^{K} \psi_k \frac{\partial \boldsymbol{a}_k}{\partial \boldsymbol{y}} + \sum_{k=1}^{K} \boldsymbol{a}_k (\nabla_{\boldsymbol{y}} \psi_k)^\top. \tag{25}$$

Since $\sum_{k=1}^{K} \psi_k = 1$, $\sum_{k=1}^{K} \nabla_{\boldsymbol{y}} \psi_k = 0$, and therefore

$$\sum_{k=1}^{K} \boldsymbol{a}_k (\nabla_{\boldsymbol{y}} \psi_k)^\top = \sum_{k=1}^{K} (\boldsymbol{a}_k - \boldsymbol{S})(\nabla_{\boldsymbol{y}} \psi_k)^\top.$$

This proves Equation (13). The decomposition allows statistical dependence between the two pathways. Since $\nabla_{\boldsymbol{y}} \psi_k \in \mathbb{R}^M$, the second term is an outer product in $\mathbb{R}^{C \times M}$.

For sufficiently small $\boldsymbol{\varepsilon}$, $\boldsymbol{S}(\boldsymbol{y} + \boldsymbol{\varepsilon}) - \boldsymbol{S}(\boldsymbol{y}) \approx \boldsymbol{J}_S \boldsymbol{\varepsilon}$. Using $\mathbb{E}[\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^\top] = \boldsymbol{\Sigma}_\varepsilon$,

$$\begin{aligned} \mathbb{E}\|\boldsymbol{J}_S \boldsymbol{\varepsilon}\|_{\boldsymbol{W}}^2 &= \mathbb{E}\ \mathrm{tr}\big(\boldsymbol{W} \boldsymbol{J}_S \boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^\top \boldsymbol{J}_S^\top\big) \\ &= \mathrm{tr}\big(\boldsymbol{W} \boldsymbol{J}_S \boldsymbol{\Sigma}_\varepsilon \boldsymbol{J}_S^\top\big). \end{aligned} \tag{26}$$

For the complete prediction, define the output change over the fixed evaluation queries as $\delta \hat{\boldsymbol{f}} = \hat{\boldsymbol{f}}(\boldsymbol{y} + \boldsymbol{\varepsilon}) - \hat{\boldsymbol{f}}(\boldsymbol{y})$. Its first-order approximation is

$$\delta \hat{\boldsymbol{f}} \approx (\boldsymbol{J}_S + \boldsymbol{J}_R)\boldsymbol{\varepsilon}.$$

Expanding its weighted second moment and using symmetry of $\boldsymbol{W}$ and $\boldsymbol{\Sigma}_\varepsilon$ gives

$$\mathbb{E}\|\delta \hat{\boldsymbol{f}}\|_{\boldsymbol{W}}^2 \approx \mathrm{tr}\big[\boldsymbol{W}(\boldsymbol{J}_S + \boldsymbol{J}_R)\boldsymbol{\Sigma}_\varepsilon(\boldsymbol{J}_S + \boldsymbol{J}_R)^\top\big] \tag{27}$$

$$= Q_S + Q_R + 2C_{SR}. \tag{28}$$

This proves Equations (14)–(16). The finite-noise sweep evaluates the complete prediction beyond this local first-order regime.

C.3 Residual Correction

Proof of Proposition 3. With $\boldsymbol{e}_S = \boldsymbol{f} - \boldsymbol{S}$, expand

$$\|\boldsymbol{e}_S - \boldsymbol{R}\|_{\boldsymbol{W}}^2 = \|\boldsymbol{e}_S\|_{\boldsymbol{W}}^2 - 2\langle \boldsymbol{e}_S, \boldsymbol{R}\rangle_{\boldsymbol{W}} + \|\boldsymbol{R}\|_{\boldsymbol{W}}^2. \tag{29}$$

Subtracting this expression from $\|\boldsymbol{e}_S\|_{\boldsymbol{W}}^2$ proves Equation (17). Dividing by $\|\boldsymbol{e}_S\|_{\boldsymbol{W}}^2$ and substituting β and $\cos\vartheta$ proves Equation (18). □

The identity holds for every proposal and residual. Section 5.3 reports Δ_{corr} across test cases and training seeds.

D Dimension-Specific Analysis

The algebraic statements in Section 4 retain the same form in two and three spatial dimensions, while dimension determines the primitive state size and evaluation cost.

Table 6: Intrinsic per-primitive degrees of freedom in d dimensions; the two covariance rows are alternative parameterizations.

Quantity or covariance option	General d	$d = 2$	$d = 3$
Center	d	2	3
Diagonal scales	d	2	3
Fully anisotropic covariance	$d(d+1)/2$	3	6
Weight	1	1	1
Amplitude	C	C	C

The fully anisotropic implementations use positive diagonal scales together with a 2D angle or a normalized 3D quaternion. A unit quaternion is stored in four coordinates but has three rotational degrees of freedom, yielding the six intrinsic covariance degrees of freedom shown above.

With diagonal or axis-aligned covariances, evaluating K primitives at Q queries costs $O(KQd)$ before decoder operations; a dense full-covariance quadratic form costs $O(KQd^2)$. The experiments therefore use dataset-appropriate covariance parameterizations to balance expressivity and evaluation cost.

The coverage radius retains the same form in either dimension: $\rho(\boldsymbol{x}) = \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{x} - \boldsymbol{\mu}_k\|_2$. Its learned value reflects primitive placement and bank size. The local rank bound remains governed by the sensor-value dimension and context dimension D, while d and K determine the state parameterization and evaluation cost.

E Regularity Interpretation for Physical Fields

Proposition 1 applies to physical fields that admit a Lipschitz extension from the fluid domain Ω to the computational domain $\mathcal{D}$ containing the learned centers. On a convex subdomain $U \subset \Omega$, a field component in $W^{1,\infty}(U)$ admits a Lipschitz representative with a constant controlled by its essential gradient norm (Temam, 2024; Girault & Raviart, 1986). More generally, if $\boldsymbol{f}$ admits an α-Hölder extension $\bar{\boldsymbol{f}} : \mathcal{D} \to \mathbb{R}^C$ satisfying $\bar{\boldsymbol{f}}|_\Omega = \boldsymbol{f}$, with $\alpha \in (0, 1]$ and Hölder constant L_α, the same proof replaces the coverage term by

$$L_\alpha \sum_{k=1}^{K} \psi_k(\boldsymbol{x}) \|\boldsymbol{x} - \boldsymbol{\mu}_k\|_2^\alpha. \tag{30}$$

Near corners, sharp trailing edges, boundary layers, and wakes, larger local gradients (Grisvard, 2011; Kozlov et al., 2000) produce spatially varying bounds, motivating the spatially resolved error analysis reported with the reconstruction results.

F Experimental Setup

F.1 AirfRANS

AirfRANS (Bonnet et al., 2022) provides steady two-dimensional RANS solutions around NACA-series airfoils on irregular meshes. Each case contains velocity (u_x, u_y), pressure p, signed distance to the airfoil, nearest-surface normal, and operating metadata. We predict (u_x, u_y, p) from sparse surface pressure.

Surface-pressure sensor layout. The Full-8 protocol uses eight sensors at normalized surface-arclength positions

$$\{0.0625, 0.1875, 0.3125, 0.4375, 0.5625, 0.6875, 0.8125, 0.9375\}.$$

The 4-, 12-, and 16-sensor studies use analogous uniform arclength schedules. Sensor locations are fixed across cases and methods.

Splits. We use the released AirfRANS train/test partitions, reserving approximately 10% of each training pool for validation. The train/validation/test counts are 720/80/200 for Full, 724/80/196 for angle-of-attack shift, and 454/50/496 for Reynolds-number shift. Case identifiers and source files are disjoint across the training, validation, and test partitions. All normalization statistics are estimated from the corresponding training partition.

Query and evaluation. Training samples 2,048 domain query points per case, with 60% concentrated near the airfoil, plus 256 surface queries. Targets are standardized per channel using training-set statistics. Query conditioning contains signed distance, nearest-surface normal, normalized arclength, and flow-frame coordinates. Reynolds number and angle of attack are standardized using training statistics.

Following the normalized-field MSE evaluation of AirfRANS (Bonnet et al., 2022), we compute squared error in standardized target space. Our primary *aggregate normalized-field MSE* averages

this error over query points and the three target channels, followed by an average over cases. The secondary metric first destandardizes the three outputs and computes

$$\mathcal{E}_{\mathrm{rel}} = \left(\frac{\sum_{q=1}^{Q} \sum_{\ell=1}^{C} (\hat{f}_{q,\ell} - f_{q,\ell})^2}{\sum_{q=1}^{Q} \sum_{\ell=1}^{C} f_{q,\ell}^2} \right)^{1/2} \tag{31}$$

per case, where $\hat{f}_{q,\ell}$ and $f_{q,\ell}$ are the predicted and target values in destandardized dataset units, $q \in \{1, \ldots, Q\}$ indexes the Q valid evaluation queries, and $\ell \in \{1, \ldots, C\}$ indexes the $C = 3$ target channels. Main benchmark and controlled-ablation tables report standard error (SE) across the independent training runs stated in their captions. Here SE denotes the standard error of the mean across those training runs.

The regular-grid baselines use a shared 96×192 nonuniform grid derived from marginal quantiles of training-case coordinates only. Model selection uses validation cases, and predictions from the selected models are sampled back to every native mesh point before applying the common Table 2 evaluator.

F.2 Senseiver Cylinder (Cy_Taira)

We use the Cy_Taira dataset distributed with the Senseiver public repository (Santos et al., 2023). It contains 5,000 vorticity snapshots for flow past a cylinder at $Re = 100$ on a 112×192 grid. The fixed training set contains 50 frames, a fixed 200-frame subset of the remaining data is used for validation, and the other 4,750 frames form the disjoint evaluation set in Table 1. Field mean and standard deviation are computed from the 50 training frames only.

Interior-8 protocol. Sensor positions follow the published Senseiver wake layout, and Table 10 reports the corresponding Interior-8 evaluation. The main Cylinder study uses the wall-mounted Surface-4/8/16 layouts below.

Surface-4/8/16 protocol. Sensors are placed on the cylinder wall and snapped to valid integer grid locations. The layouts are fixed across cases and methods:

Layout	Grid coordinates [row, column]
Surface-4	[55, 20], [49, 14], [55, 7], [62, 13]
Surface-8	[55, 20], [51, 18], [49, 14], [51, 9], [55, 7], [60, 9], [62, 13], [60, 18]
Surface-16	[55, 20], [53, 19], [51, 18], [50, 16], [49, 14], [50, 11], [51, 9], [53, 8], [55, 7], [58, 8], [60, 9], [61, 11], [62, 13], [61, 16], [60, 18], [58, 19]

All learned rows use five independent training runs; uncertainty is the standard error (SE) across these runs. Configurations are selected on the 200-frame validation set and evaluated on the disjoint 4,750-frame test set.

F.3 FluidWeave Implementation Details

Shared core. All experiments use a per-sensor MLP encoder with mean/max pooling, a joint primitive head, query-to-primitive cross-attention with a residual feed-forward layer, and a Fourier-feature residual MLP conditioned on the proposal, basis mass, global context, and cross-attention output. Query tokens combine coordinates, sensor context, and benchmark-specific geometry; primitive tokens contain the predicted center, covariance, weight, and amplitude. PhySense-Car additionally uses query-to-sensor attention, whereas LDC-3D includes the query signed-distance feature.

Cylinder Surface-4/8/16. The primitive head predicts $K = 256/384/448$ rotated anisotropic Gaussians. The residual path has four linear layers with hidden width $320/448/512$ and six Fourier-frequency bands. No case metadata or per-query geometry features are used.

Primitive parameterization (cylinder). Centers lie in the normalized domain. Scales use $[0.01, 0.35]$ for Surface-4/8 and $[0.005, 0.25]$ for Surface-16. Positive weights are produced by a sigmoid, amplitudes are unconstrained, and rotation is enabled.

AirfRANS (all splits). The model uses $K = 96$ axis-aligned anisotropic Gaussians, so rotation is disabled. Primitive scales lie in $[0.01, 0.35]$. The residual MLP has three linear layers, six Fourier-frequency bands, and hidden width 192/128/192 for Full/AoA/Reynolds. Full and AoA use a learned global token; the Reynolds configuration uses the pooled context directly. Query features comprise signed distance, nearest-surface normal, normalized arclength, and flow-frame coordinates; metadata modulation uses gain 0.35.

Training (cylinder). Surface-4 uses learning rate 8×10^{-4} and Surface-8/16 use 5×10^{-4}. Training uses Adam, weight decay 10^{-6}, batch size 16, gradient clipping, validation-based learning-rate reduction, and early stopping. Observation-consistency weights are $0.02/0.01/0.005$ for the three sensor counts.

Training (AirfRANS). Training uses Adam at 10^{-3}, weight decay 10^{-6}, batch size 2, gradient clipping at 1.0, and validation-based learning-rate reduction and early stopping. Observation-consistency and surface-loss weights are 0.05 and 0.2. The main accuracy and sensor-noise comparisons use five independent training runs; controlled branch, representation, and correction analyses use a common three-run subset, as stated with each result.

Compute. Experiments were run on a single NVIDIA RTX 5060 Ti with 16 GB memory. The fixed Full-8 configuration typically completes in about 50 minutes per training run.

Three-dimensional FLUIDWEAVE configurations. Table 7 summarizes the architecture and optimization settings used for the two 3D studies. Both use the shared primitive-state and correction design; the table lists dataset-specific inputs, outputs, conditioning, and model-selection settings.

Table 7: FluidWeave configurations for the 3D experiments. "Full volume" and "full surface" mean that inference covers every valid evaluation query.

Configuration	FlowBench LDC-3D	PhySense-Car
Output / conditioning	(u, v, w, p); SDF and Reynolds	pressure; sensor only
Primitives K	128	128
Sensor/context width	192 / 192	192 / 192
Residual MLP	width 384, 5 layers	width 384, 5 layers
State attention	192 dim, 4 heads	192 dim, 4 heads
Fourier bands / global token	8 / 64 dim	8 / 64 dim
Trainable parameters	1,808,900	2,309,537
Training queries per case	8,192 fluid voxels	8,192 surface points
Batch size / optimizer	8 / AdamW	8 / AdamW
Learning rate / weight decay	3×10^{-4} / 10^{-5}	3×10^{-4} / 10^{-5}
Schedule	cosine, at most 200 epochs	cosine, 301 fixed epochs
Model selection	lowest validation global rel. L^2	fixed final epoch
Evaluation queries	full volume	full surface
Independent training runs	5	3

F.4 FLOWBENCH LDC-3D DATASET AND EVALUATION PROTOCOL

FlowBench LDC-3D (Tali et al., 2025) contains 1,000 steady lid-driven-cavity simulations on 64^3 grids. Each case contains random internal obstacles, an obstacle signed-distance field, Reynolds number, and four targets (u, v, w, p). The split is 700/100/200 train/validation/test cases. For $N \in \{5, 10, 20, 40\}$, all methods use the same nested obstacle-mounted probe layouts, with values sampled from the adjacent fluid shell two processed-grid voxels from the boundary. This fixed sampling stand-off represents surface-mounted probes measuring the surrounding flow while keeping every layout tied to the obstacle geometry. Each probe observes (u, v, w, p); all methods receive the same readings and coordinates, signed-distance geometry, and Reynolds conditioning. Training uses 8,192 sampled valid fluid voxels per case; validation and test cover all valid fluid voxels.

Reynolds number is normalized using the 700 training cases, with mean 289.7757 and standard deviation 301.3884. All reported methods are trained with this normalization and evaluation protocol. Model selection uses validation performance; reported metrics are computed over the full held-out test set. Appendix H.10 reports global, channel-macro, per-channel, sensor-value, and metadata-intervention results.

F.5 PhySense-Car Dataset and Evaluation Protocol

PhySense-Car (Ma et al., 2025) reconstructs pressure on unstructured ShapeNet car surfaces with roughly 95k points per case. We retain the released split, pressure normalization constants, fixed sensor indices, and sensor-only evaluator. Cases 1–75 are used for training and 76–100 for testing. Three independently trained fixed-epoch FluidWeave models use sensor counts sampled from $[10, 200]$ and are evaluated at $N \in \{15, 30, 50, 100, 200\}$. The locally evaluated baselines use three independent training runs with the same split, sensor indices, normalization, and released evaluator. Appendix I provides the mask, sampling, conditioning, and training settings.

F.6 Sensor-Noise Protocol

For a standardized pressure-sensor value $\widetilde{y}$, test-time noise is

$$\widetilde{y}' = \widetilde{y} + \sigma_{\text{noise}}\zeta, \qquad \zeta \sim \mathcal{N}(0, 1), \tag{32}$$

with dimensionless noise level $\sigma_{\text{noise}} \in \{0, 0.005, 0.01, 0.02, 0.05, 0.10, 0.20\}$ and an independent standard normal draw ζ for each perturbed sensor value. Each level is a fraction of the training-set sensor-pressure standard deviation. Sensor coordinates, geometry, metadata, queries, and trained models remain fixed. At every nonzero level, each trained model is evaluated on all 200 test cases with five independent noise draws per case, shared across methods. Metrics are first averaged over cases and noise draws for each model; tables then report mean $\pm$ SE across five independent training runs, while Figure 4 reports the corresponding mean $\pm$ SE.

G Baseline Implementations

Evaluation protocol. Within each benchmark, all methods trained in our study share the data split, sensor layout, target fields, and evaluator. Results are summarized over the stated number of independent training runs. Hyperparameter and checkpoint selection follow the benchmark-specific protocols summarized below. We use released implementations or published architectures under these protocols.

Senseiver. For Cylinder, we run the public Senseiver implementation under the surface protocol in Appendix F.2. For AirfRANS, the same set-to-field attention architecture encodes sensor values and coordinates and is queried at native-mesh points. Input and output dimensions follow the corresponding benchmark fields (Santos et al., 2023).

DeepONet. The branch network encodes sparse observations and case context, while the trunk network encodes query coordinates and geometry features. Their paired features predict the target channels at each query (Lu et al., 2021).

RecFNO. Following RecFNO (Zhao et al., 2024), sparse observations are lifted to a regular latent grid, processed by Fourier operator blocks, and sampled at the evaluation coordinates.

FLRNet and Voronoi-CNN. For AirfRANS, we retain FLRNet's released sensor-mapping and decoder modules (Nguyen et al., 2026), adapted for end-to-end three-field reconstruction. For Voronoi-CNN (Fukami et al., 2021), we follow the original Voronoi input construction using eight pressure measurements and a sensor-location mask, with a three-channel output layer. Both methods use the common data splits and evaluator.

POD–Ridge. Training fields are compressed with principal components, and ridge regression maps sensor observations and case descriptors to modal coefficients. The number of modes and regularization strength are selected on validation data.

G.1 Matched Direct Conditional Decoder

The matched direct decoder is the closest non-scaffold control. It receives the same sparse values and coordinates, sensor encoder, case context, per-query geometry, and Fourier coordinate features as FluidWeave, but predicts the field directly without constructing a primitive state. It uses a four-layer, width-248 MLP with learning rate 10^{-3} and is trained with five independent runs under the fixed Full-8 protocol. Its training objective combines field reconstruction and surface supervision with weights 1 and 0.2, respectively; the full FluidWeave model additionally uses observation consistency with weight 0.05.

Table 8: Capacity and traced compute for the matched direct control. FLOPs use the same representative forward batch.

Method	Parameters	Traced forward FLOPs
FluidWeave	236,283	745,895,864
Matched direct decoder	228,563	1,212,295,872

The comparison is parameter-comparable: the two models differ by 3.27% in parameter count, while the direct decoder uses higher traced compute.

G.2 Three-Dimensional Baselines

For FlowBench LDC-3D, the DeepONet, Senseiver, and RecFNO implementations evaluated in this study contain 1.79M, 3.63M, and 2.24M parameters, respectively. All methods share the fixed sensor layouts, Reynolds conditioning, 8,192-query objective, five independent training runs, and four-channel evaluation protocol.

For PhySense-Car, Table 4 places the point estimates reported by Ma et al. (2025) alongside three-run evaluations conducted in this study. These methods use the released sensor indices, normalization, split, and evaluator. Full protocol details are in Appendix I.

H Additional Results

H.1 Cylinder Per-Run Results

Table 9: Per-run physical-space relative L^2 on the disjoint 4,750-frame Cylinder evaluation set.

Setting	Method	Run 1	Run 2	Run 3	Run 4	Run 5
Surface-4	FluidWeave	0.067908	0.064002	0.057227	0.066225	0.065453
	Senseiver	0.068654	0.063124	0.089461	0.080278	0.064379
	DeepONet	0.098148	0.092515	0.091591	0.090081	0.089978
	RecFNO	0.078314	0.054161	0.097464	0.077267	0.065517
	FLRNet	0.100319	0.090855	0.103355	0.104128	0.083838
	Voronoi-CNN	0.072977	0.072722	0.073565	0.075514	0.070793
Surface-8	FluidWeave	0.047905	0.029811	0.034845	0.036137	0.030438
	Senseiver	0.063805	0.059187	0.047721	0.059072	0.051516
	DeepONet	0.127184	0.141928	0.108602	0.130783	0.141973
	RecFNO	0.113815	0.093549	0.087638	0.108059	0.109204
	FLRNet	0.091979	0.085362	0.099578	0.096775	0.084909
	Voronoi-CNN	0.031685	0.057459	0.035543	0.045574	0.039141
Surface-16	FluidWeave	0.016616	0.015417	0.016989	0.015632	0.015234
	Senseiver	0.056016	0.070619	0.045151	0.056497	0.044304
	DeepONet	0.095878	0.095706	0.096024	0.094957	0.090500
	RecFNO	0.065817	0.074231	0.063508	0.072289	0.076757
	FLRNet	0.087521	0.089401	0.092864	0.095344	0.074811
	Voronoi-CNN	0.021066	0.019574	0.018513	0.026808	0.016856

All rows use the same disjoint 4,750-frame evaluation set as Table 1. Across the 15 paired layout–run comparisons with Voronoi-CNN, FLUIDWEAVE has lower error in 14, including every Surface-4 and Surface-16 run, complementing the five-run aggregate results in the main table.

H.2 CYLINDER INTERIOR-8 RESULTS

Table 10 reports the established Interior-8 sensor layout as a complementary Cylinder evaluation.

Table 10: Cylinder Interior-8 relative L^2. The Senseiver point estimate is reported from its publication; the remaining entries are evaluated in this study using the same sensor layout.

Method	Relative L^2
FLUIDWEAVE (ours)	**0.0097**
Senseiver	0.0390
RecFNO	0.0875
DeepONet	0.3847

H.3 MECHANISM DIAGNOSTICS

We measure correction, coverage, and local sensitivity on AirfRANS Full-8. For the correction analysis, each of three independently trained models is evaluated on all 200 AirfRANS Full-8 test cases. The state-conditioned correction reduces mean global relative L^2 from 0.7827 for the proposal to 0.3151 for the complete prediction. Across all three models, the measured Δ_{corr} from Proposition 3 is positive for every test case. The normalized correction magnitude and alignment are $\beta = 0.8923$ and $\cos\vartheta = 0.9127$. We compute β and $\cos\vartheta$ separately for each test case and trained model, then report their equal-weight averages over 200 test cases and three trained models.

For the cross-case state intervention, ten fixed-point-free donor permutations replace the complete primitive state of each recipient case. Recipient sensor context, coordinates, geometry features, metadata, and queries remain unchanged. Across these interventions, global relative L^2 increases by 0.0276 on average (95% bootstrap confidence interval $[0.0260, 0.0293]$); 77.85% of recipient–donor clusters show increased error. For each recipient–donor pair, the effect is first averaged over the three trained models. The confidence interval is obtained by resampling the ten donor-pairing summaries, each computed over the 200 recipient cases.

The intervention provides direct evidence that the complete prediction depends on the primitive state. Controlled decoder ablations reinforce this result: removing state features raises aggregate normalized-field MSE from 0.2367 to 0.2992, and removing query-to-state attention raises it to 0.3446.

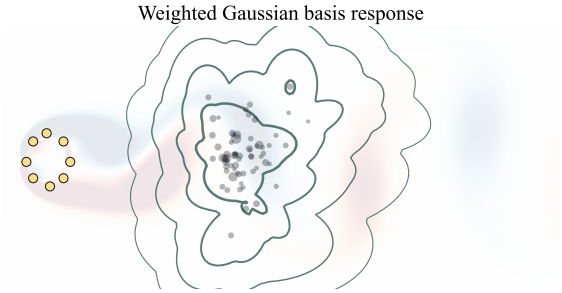

Figure 5: Evaluated primitive-support state on Cylinder Surface-8. Yellow and gray markers denote sensors and highest-weight primitive centers (marker size: weight). Teal contours show maximum weighted Gaussian basis response; the background shows display-normalized ground-truth vorticity (red–blue) for spatial reference.

Coverage and local sensitivity. We evaluate the normalized coverage radius $\rho(\boldsymbol{x})$ from Proposition 1 on 4,096 deterministic native-mesh queries for every test case and each of the same three models. Across the evaluated queries, the Pearson and Spearman associations between $\rho(\boldsymbol{x})$ and scaffold point error are 0.1173 and 0.2398, respectively, consistent with the spatial-coverage term in Proposition 1.

For local sensitivity, we hold sensor coordinates, geometry, and metadata fixed and differentiate with respect to the eight sensor values. Exact input directions are evaluated on 30 fixed test cases per model, using 512 deterministic queries per case. Table 11 reports averages over the resulting 90 case–model units. The correction path dominates the local quadratic sensitivity, while the positive mean cross-term indicates aligned scaffold/correction perturbations on average. For these diagnostics, $\boldsymbol{\Sigma}_\varepsilon = \boldsymbol{I}_8$ and $\boldsymbol{W} = \boldsymbol{I}/M_{\mathrm{out}}$, so the reported Q terms use unit-variance normalized sensor directions and mean squared normalized-field output change.

Table 11: AirfRANS Full-8 coverage and Jacobian diagnostics. The finite-noise ratio is empirical divided by first-order predicted mean squared output change. The local Jacobian comparison covers 0.5–2% perturbations; the complete noise sweep reports finite-perturbation behavior.

Coverage association	Pearson	Spearman	Samples
$\rho(\boldsymbol{x})$ vs. scaffold point error	0.1173	0.2398	2,457,600
Jacobian terms	Q_S	Q_R	C_{SR}
Mean	0.00196	1.03936	0.00842
Noise std.	0.5%	1%	2%
Empirical/predicted	1.045	1.143	0.903

For $Q_S Q_R > 0$, define the scale-free interaction $\eta_{SR} = C_{SR}/\sqrt{Q_S Q_R}$. Its mean is 0.1818. Across 0.5–2% perturbations, the local prediction is within 4–14% of the empirical mean, supporting the local first-order sensitivity model over the evaluated range.

H.4 FluidWeave K-ablation

The main Surface-8 configuration uses $K = 384$, selected jointly with the remaining hyperparameters on the validation split. To examine sensitivity to primitive-bank resolution, we vary only K while retaining the residual hidden dimension (448), attention dimension (160), four heads, scale range $[0.01, 0.35]$, and $\lambda_{\mathrm{obs}} = 0.01$ at their $K = 384$ values.

Table 12: Cylinder Surface-8 primitive-count study. Values are mean ± SE over five independent training runs on the disjoint 4,750-frame evaluation set.

K	FluidWeave	Scaffold-only
64	**0.0282 ± 0.0008**	0.3471 ± 0.0067
128	0.0287 ± 0.0007	0.2933 ± 0.0113
256	0.0298 ± 0.0015	0.2501 ± 0.0056
384	0.0358 ± 0.0033	**0.2356 ± 0.0089**

In this sweep, increasing K lowers mean scaffold-only error but raises mean full-model error. Thus, better standalone proposal accuracy need not improve the final reconstruction under the fixed remaining configuration.

H.5 Branch Ablations and Sensor-Count Scaling

The main benchmark aggregate in Table 2 summarizes five independent training runs. The controlled-ablation table reports the component and representation controls over a common three-run subset shared by all variants; its reference row is the corresponding three-run summary. The two reference values therefore summarize different run sets for the same Full-8 model.

Table 13: AirfRANS Full-8 component and representation controls: aggregate normalized-field MSE (mean ± SE over three common paired training runs; lower is better).

Variant	Parameters	MSE
Reference and representation		
FLUIDWEAVE (full)	0.236M	**0.2367 ± 0.0261**
Direct decoder (no spatial state)	0.229M	0.3227 ± 0.0122
Generic token routing	0.236M	0.2627 ± 0.0127
State and merge controls		
No additive proposal (state retained)	0.236M	0.2682 ± 0.0086
No state features	0.235M	0.2992 ± 0.0099
No state attention	0.210M	0.3446 ± 0.0083
No global token	0.212M	0.2759 ± 0.0207
Pathway and objective controls		
Proposal only	0.114M	0.6524 ± 0.0091
Residual only (broader control)	0.209M	0.3530 ± 0.0075
No observation-consistency loss	0.236M	0.2620 ± 0.0070

With the same loss terms and weights as the direct decoder, FLUIDWEAVE without observation-consistency loss achieves a lower mean MSE over the three runs (0.2620 versus 0.3227; Table 13).

Generic-token routing control. The generic-token control preserves the per-case token budget, normalized proposal aggregation, correction decoder, and training objective, while removing the geometric interpretation of the generated tokens. Let $\boldsymbol{g}_k$ denote the k-th raw token emitted by the joint head. For the reported AirfRANS control, $\boldsymbol{g}_k$ has dimension $2 + 2 + C + 1$; its C-dimensional amplitude block is denoted by $\boldsymbol{a}_k$, and the full vector is projected to a token feature $\tilde{\boldsymbol{t}}_k = E_p^{\mathrm{tok}}(\boldsymbol{g}_k)$. Thus, the entries of $\boldsymbol{g}_k$ are not interpreted as centers, scales, or Gaussian weights.

Let $\boldsymbol{q}_{\boldsymbol{x}}$ be the query token constructed with the same coordinate, context, global-token, and benchmark-specific query features as in the full model. For attention head h with key dimension d_h, the query-to-token routing weights are

$$\alpha_k^{(h)}(\boldsymbol{x}) = \frac{\exp\big((Q_h \boldsymbol{q}_{\boldsymbol{x}})^\top (K_h \tilde{\boldsymbol{t}}_k)/\sqrt{d_h}\big)}{\sum_{j=1}^{K} \exp\Big((Q_h \boldsymbol{q}_{\boldsymbol{x}})^\top (K_h \tilde{\boldsymbol{t}}_j)/\sqrt{d_h}\Big)}, \qquad \bar{\alpha}_k(\boldsymbol{x}) = \frac{1}{H}\sum_{h=1}^{H} \alpha_k^{(h)}(\boldsymbol{x}), \tag{33}$$

where $k, j \in \{1, \dots, K\}$, $h \in \{1, \dots, H\}$, Q_h and K_h are the learned query and key projections for head h, and d_h is its key dimension. The reported control uses $H = 4$ heads, and the implementation averages the per-head weights. Consequently, $\bar{\alpha}_k(\boldsymbol{x}) \geq 0$ and $\sum_k \bar{\alpha}_k(\boldsymbol{x}) = 1$, so the control's normalized proposal is

$$\boldsymbol{S}_{\mathrm{tok}}(\boldsymbol{x}) = \sum_{k=1}^{K} \bar{\alpha}_k(\boldsymbol{x}) \boldsymbol{a}_k, \qquad \kappa_{\mathrm{tok}}(\boldsymbol{x}) = \sum_{k=1}^{K} \bar{\alpha}_k(\boldsymbol{x})^2. \tag{34}$$

The residual decoder uses the same query/context inputs, cross-attention block, loss, and additive merge as the full model; its proposal-related state features are $\boldsymbol{S}_{\mathrm{tok}}(\boldsymbol{x})$ and the routing concentration $\kappa_{\mathrm{tok}}(\boldsymbol{x})$. The resulting prediction is therefore $\hat{\boldsymbol{f}}_{\mathrm{tok}}(\boldsymbol{x}) = \boldsymbol{S}_{\mathrm{tok}}(\boldsymbol{x}) + \boldsymbol{R}_{\mathrm{tok}}(\boldsymbol{x})$. This provides a controlled comparison between geometry-defined Gaussian routing and generic content-based token routing with the same token budget, residual-decoder design, and training objective. Its mean MSE is 0.2627 ± 0.0127, compared with 0.2367 ± 0.0261 for the full model.

The no-additive-proposal control retains the Gaussian state, proposal-derived state features, global token, query-to-state attention, and residual decoder. It changes only the final merge from $\boldsymbol{S}(\boldsymbol{x}) + \boldsymbol{R}(\boldsymbol{x})$ to $\boldsymbol{R}(\boldsymbol{x})$, giving a mean MSE of 0.2682 ± 0.0086 versus 0.2367 ± 0.0261 for the full model. This isolates the additive proposal pathway while keeping state-conditioned correction available.

Residual-only is a broader residual-path control: it removes the additive proposal together with the proposal-value and basis-mass features and query-to-state attention, while retaining the global context token. Proposal only removes the correction pathway. The independently tuned direct decoder predicts the field without constructing a primitive state.

State-interface controls. *No state features* removes $\boldsymbol{S}(\boldsymbol{x})$ and $m_{\mathrm{num}}(\boldsymbol{x})$ from Equation (8), retaining attention and the additive proposal. *No state attention* removes E_q, E_p, and the block in Equation (7), omitting $\boldsymbol{h}_{\boldsymbol{x}}$ from the residual input but retaining $\boldsymbol{S}(\boldsymbol{x})$ and $m_{\mathrm{num}}(\boldsymbol{x})$. Full-8 augments the inputs to E_q and D_R with an extra learned projection of $\boldsymbol{z}$. *No global token* removes this projection, retaining $\boldsymbol{z}$ itself and the primitive-state pathways.

For 4/8/12/16 sensors, FLUIDWEAVE gives aggregate MSEs 0.2997, 0.2526, 0.2833, and 0.2813, while DeepONet gives 0.3108, 0.3289, 0.3320, and 0.3171. FLUIDWEAVE is lower at every tested sensor count.

H.6 MATCHED DIRECT DECODER: PER-RUN RESULTS

Table 14: Full-8 aggregate normalized-field MSE for paired training runs. The lower value in each run column is bold.

Method	Run 1	Run 2	Run 3	Run 4	Run 5
FLUIDWEAVE	**0.195864**	**0.229128**	**0.285197**	**0.298828**	**0.254034**
Matched direct	0.341863	0.326172	0.300002	0.301066	0.297250

Under the common Full-8 protocol, FLUIDWEAVE is lower in all five paired runs, with five-run means of 0.252610 and 0.313271 for FLUIDWEAVE and the direct decoder, respectively.

H.7 MATCHED NON-GAUSSIAN KERNEL CONTROL

We isolate the radial response while retaining the same sensor-conditioned centers, scales, weights, amplitudes, normalized aggregation, state-aware correction, losses, 10-epoch budget, and 236,283 parameters. The Gaussian $\exp(-r^2/2)$ is replaced by the compact Wendland C2 response $[1 - r/3]_+^4(4r/3 + 1)$. Its support radius is fixed to three learned scales before training, matching the conventional effective 3σ Gaussian neighborhood.

Table 15: Matched radial-response control on AirfRANS Full-8. Entries are aggregate normalized-field MSE; mean $\pm$ SE is over three paired training runs.

Kernel	Mean $\pm$ SE
Gaussian	**0.2367 $\pm$ 0.0261**
Wendland C2	0.2745 $\pm$ 0.0304

Gaussian achieves the lower mean error under this matched control, supporting its use in the FluidWeave proposal–correction architecture.

H.8 AIRFRANS FULL-8 PER-FIELD DIAGNOSTICS

Table 16 resolves the aggregate Full-8 score into the three standardized target channels on the common three-run diagnostic subset.

Table 16: AirfRANS Full-8 per-field normalized MSE on the common three-run diagnostic subset (mean $\pm$ SE across independent training runs). Lower is better.

Method	u_x	u_y	p
FLUIDWEAVE (ours)	**0.4227 $\pm$ 0.0602**	**0.2056 $\pm$ 0.0255**	**0.0819 $\pm$ 0.0036**
Matched direct decoder	0.5711 $\pm$ 0.0118	0.2892 $\pm$ 0.0148	0.1078 $\pm$ 0.0137
DeepONet	0.5393 $\pm$ 0.0346	0.3313 $\pm$ 0.0266	0.1175 $\pm$ 0.0201

FLUIDWEAVE improves all three output fields on this common subset, showing that the aggregate Full-8 gain is distributed across velocity and pressure.

H.9 COMPLETE NOISE-SWEEP MSE

Table 17: AirfRANS aggregate normalized-field MSE under paired sensor noise. Entries report absolute MSE in train-standardized field space as mean ± SE over five independent training runs. All noise levels are paired within the five-model ensemble trained with the main Full-8 configuration.

Method	0%	0.5%	1%	2%	5%	10%	20%
FLUIDWEAVE	**0.2526 ± 0.0186**	**0.2526 ± 0.0187**	**0.2526 ± 0.0187**	**0.2529 ± 0.0186**	**0.2530 ± 0.0184**	**0.2570 ± 0.0186**	**0.2662 ± 0.0184**
Matched direct	0.3133 ± 0.0089	0.3133 ± 0.0089	0.3133 ± 0.0088	0.3135 ± 0.0088	0.3135 ± 0.0087	0.3157 ± 0.0082	0.3193 ± 0.0065
DeepONet	0.3289 ± 0.0148	0.3289 ± 0.0148	0.3289 ± 0.0147	0.3294 ± 0.0148	0.3295 ± 0.0145	0.3353 ± 0.0144	0.3439 ± 0.0141

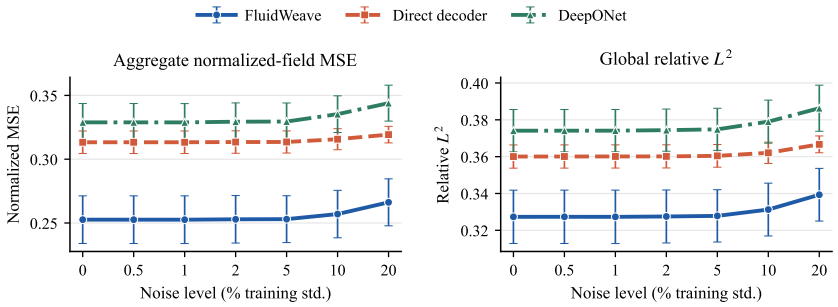

Figure 6: AirfRANS Full-8 sensor-noise sweep over five training runs (mean ± SE). FLUIDWEAVE retains the lowest absolute error under both metrics throughout the tested range.

Together with Section 5.3, these values show that FLUIDWEAVE has the lowest MSE across the tested range.

H.10 LDC-3D CHANNEL-BALANCED RESULTS AND CONTROLS

The main LDC-3D metric is a joint relative norm over the four raw channels and therefore depends on their numerical scales. We complement it with per-case relative L^2 for u, v, w, and p and a channel-balanced macro average. Tables below report mean ± SE over the same five independent training runs as Table 3.

Table 18: LDC-3D channel-balanced and per-channel relative L^2 for $N = 5$.

Method	Global	Macro	u	v	w	p
FLUIDWEAVE	**0.0303±0.0003**	**0.2246±0.0032**	0.1065±0.0015	0.1173±0.0018	0.6540±0.0095	0.0207±0.0003
DeepONet	0.0590±0.0086	0.3211±0.0265	0.2107±0.0193	0.2300±0.0252	0.8022±0.0498	0.0415±0.0118
Senseiver	0.1009±0.0055	0.4480±0.0111	0.3285±0.0153	0.3619±0.0336	1.0001±0.0009	0.1017±0.0189
RecFNO	0.1144±0.0192	0.4522±0.0206	0.3463±0.0165	0.3574±0.0111	0.9786±0.0221	0.1264±0.0400
FLRNet	0.0670±0.0043	0.3610±0.0046	0.1700±0.0069	0.2061±0.0088	0.9967±0.0012	0.0713±0.0039
Voronoi-CNN	0.0820±0.0170	0.5653±0.1307	0.3595±0.0901	0.3703±0.0787	1.4482±0.3390	0.0833±0.0238

Table 19: LDC-3D channel-balanced and per-channel relative L^2 for $N = 10$.

Method	Global	Macro	u	v	w	p
FLUIDWEAVE	**0.0284±0.0003**	**0.2229±0.0035**	0.1055±0.0014	0.1164±0.0015	0.6510±0.0113	0.0186±0.0003
DeepONet	0.0465±0.0022	0.2889±0.0048	0.1782±0.0081	0.1960±0.0067	0.7544±0.0044	0.0269±0.0007
Senseiver	0.0760±0.0159	0.4025±0.0253	0.2566±0.0353	0.2836±0.0386	1.0011±0.0015	0.0686±0.0266
RecFNO	0.1219±0.0267	0.4535±0.0273	0.3419±0.0208	0.3553±0.0151	0.9773±0.0255	0.1396±0.0529
FLRNet	0.0533±0.0016	0.3558±0.0044	0.1664±0.0066	0.2020±0.0097	0.9965±0.0014	0.0584±0.0022
Voronoi-CNN	0.0691±0.0100	0.4709±0.0708	0.2987±0.0472	0.4089±0.1013	1.1107±0.1353	0.0652±0.0098

Table 20: LDC-3D channel-balanced and per-channel relative L^2 for $N = 20$.

Method	Global	Macro	u	v	w	p
FLUIDWEAVE	**0.0281**±**0.0004**	**0.2247**±**0.0034**	0.1074±0.0017	0.1183±0.0016	0.6562±0.0104	0.0166±0.0004
DeepONet	0.0493±0.0020	0.2942±0.0031	0.1886±0.0060	0.2031±0.0042	0.7560±0.0021	0.0291±0.0023
Senseiver	0.1499±0.0257	0.5530±0.0513	0.5086±0.1058	0.5255±0.0777	1.0196±0.0111	0.1582±0.0390
RecFNO	0.1328±0.0280	0.5096±0.0515	0.3933±0.0437	0.4005±0.0377	1.1098±0.0885	0.1349±0.0398
FLRNet	0.0503±0.0015	0.3530±0.0036	0.1633±0.0053	0.1981±0.0078	0.9964±0.0014	0.0542±0.0019
Voronoi-CNN	0.0691±0.0104	0.4867±0.0867	0.3225±0.0650	0.3858±0.0974	1.1706±0.2036	0.0677±0.0126

Table 21: LDC-3D channel-balanced and per-channel relative L^2 for $N = 40$.

Method	Global	Macro	u	v	w	p
FLUIDWEAVE	**0.0273**±**0.0004**	**0.2225**±**0.0037**	0.1053±0.0017	0.1165±0.0017	0.6534±0.0116	0.0147±0.0003
DeepONet	0.0483±0.0024	0.2956±0.0050	0.1864±0.0058	0.2075±0.0073	0.7607±0.0044	0.0279±0.0044
Senseiver	0.0906±0.0032	0.4490±0.0094	0.3139±0.0045	0.3796±0.0378	1.0296±0.0294	0.0726±0.0153
RecFNO	0.1238±0.0321	0.4821±0.0525	0.3537±0.0311	0.3663±0.0203	1.0689±0.1023	0.1398±0.0601
FLRNet	0.0475±0.0012	0.3526±0.0035	0.1657±0.0059	0.1991±0.0076	0.9967±0.0013	0.0488±0.0013
Voronoi-CNN	0.0489±0.0089	0.4635±0.1660	0.2131±0.0474	0.2705±0.0838	1.3250±0.5223	0.0456±0.0108

Sensor and metadata interventions. At $N = 20$, we retain sensor coordinates and change only the named input. Zeroing or cross-case shuffling sensor values sharply increases global and channel-macro error for both FLUIDWEAVE and DeepONet, while assigning Reynolds metadata from another case also degrades both models (Table 22). Per-channel results show strong sensor dependence in pressure reconstruction and sensitivity of velocity predictions to Reynolds metadata.

Table 22: LDC-3D $N = 20$ global and channel-macro relative L^2 under input interventions. Entries are five-run means; per-channel means for FLUIDWEAVE are shown below.

Method	Metric	Clean	Zero values	Shuffled values	Cross-case Re
FLUIDWEAVE	Global	0.0281	0.9458	8.2684	0.1593
FLUIDWEAVE	Macro	0.2247	0.4855	3.1101	0.4309
DeepONet	Global	0.0493	0.9508	8.2742	0.1486
DeepONet	Macro	0.2942	0.5475	3.1537	0.4077
FLUIDWEAVE	u	0.1074	0.1213	0.1699	0.3565
FLUIDWEAVE	v	0.1183	0.1278	0.1673	0.3647
FLUIDWEAVE	w	0.6562	0.7021	0.7959	0.8090
FLUIDWEAVE	p	0.0166	0.9909	11.3075	0.1936

I PHYSENSE-CAR: PROTOCOL DETAILS

Data, sensing, and evaluation. We use the released raw files, pressure normalization constants, and split: cases 1–75 for training and 76–100 for testing, with approximately 95k surface points per case. Following the released sensor protocol, all methods evaluated in this study use the same fixed sensor indices for $N \in \{15, 30, 50, 100, 200\}$. Speed and yaw are set to zero to match the sensor-only setting, so models receive only sensor coordinates and pressure readings. We report normalized relative-L^2 using the released 3σ pressure mask and normalization constants. Table 4 separates the published PhySense (Ma et al., 2025) Table 3 values from the three-run evaluations conducted in this study.

Training and evaluation. Three independently trained FLUIDWEAVE models use 301 epochs and $N \sim \mathrm{Uniform}[10, 200]$ per sample; each model is evaluated at all five sensor counts. For the locally evaluated baselines, settings are selected using cases 1–70/71–75 for training/validation, and final models use cases 1–75, three independent training runs, 8,192 targets per case, and $N \sim \mathrm{Uniform}[10, 200]$. We report standard error (SE) over training runs. Figure 3 shows two held-out cases at $N = 50$ with identical sensor draws across methods.