CoSGF: 4D Scientific Gaussian Fields for High-Resolution Reconstruction from Complementary Observations

Anonymous authors
Paper under double-blind review

Abstract

High-fidelity numerical simulations resolve the spatiotemporal structures needed to study complex physical systems, but generating and retaining long high-resolution trajectories can be computationally and storage intensive. Existing approaches reduce data requirements through augmentation, few-shot learning, flexible representations, or physical priors, but reconstruction is often organized around a dominant observation pathway; increasing model capacity or adding auxiliary constraints alone cannot recover information that is not identifiable from coarse measurements. We instead formulate scientific-field reconstruction as the joint inference of a shared physical field from heterogeneous but complementary scientific evidence. We instantiate this formulation with a 4D Scientific Gaussian Field whose primitives encode physical variables rather than appearance. Numerical measurements and retained scientific visualizations constrain the same field through their corresponding forward processes, while dataset-specific physical and geometric knowledge further restricts residual ambiguity. Rather than treating physics as an unrestricted third loss, physics-driven optimizer displacements are coordinated using the local sensitivity of active observations and finite forward verification, so refinement preferentially acts on weakly observed degrees of freedom while preserving evidence-supported content. The HR field remains inaccessible during reconstruction and is used only for evaluation, while the framework directly produces an analysis-ready high-resolution scientific field. In the reported comparison over four single-phase datasets at $4\times$ and $8\times$ reconstruction ratios, CoSGF achieves the highest mean PSNR among the compared methods.

1 Introduction

High-fidelity numerical simulations provide the detailed spatiotemporal fields required to study complex physical systems, but generating and retaining long high-resolution trajectories can incur substantial computational, storage, and I/O costs Brunton et al. (2020); Liang et al. (2018). Scientific-field reconstruction alleviates this burden by recovering fine-scale physical variables and structures from reduced data for subsequent visualization and quantitative analysis. Recent work studies hierarchical and probabilistic super-resolution of scientific data Wurster et al. (2023); Shen & Shen (2023); Wang et al. (2024), turbulence-field reconstruction Yousif et al. (2022), and super-resolution data assimilation Barthélémy et al. (2022). Related representation approaches use continuous or implicit fields Sitzmann et al. (2020); Tancik et al. (2020), neural operators Lu et al. (2021); Kovachki et al. (2023); Li et al. (2021), and physics-informed constraints Raissi et al. (2019); Karniadakis et al. (2021); Cuomo et al. (2022). Yet reconstruction is commonly centered on a primary numerical observation, with other information incorporated in task-specific or auxiliary forms. Inverse-volume and image-based methods show that rendered images can constrain latent volumetric or material properties Gkioulekas et al. (2013); Zhang et al. (2023), although they mainly target renderable representations and visual exploration rather than recovery of the underlying physical field. The central limitation is therefore not merely model capacity: numerical reduction discards fine-scale information, so distinct high-resolution states can remain consistent with the same coarse measurements. Scientific visualizations, generated from the same state through a different forward process,

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.

can provide complementary evidence about weakly determined structures. We therefore study how heterogeneous scientific evidence can be jointly reconciled in a shared physical representation to recover high-resolution field structures that reduced numerical data alone cannot reliably determine.

To address this problem, we introduce a 4D Scientific Gaussian Field as a shared physical representation for joint reconstruction from complementary scientific evidence. Unlike appearance-oriented Gaussian representations, each primitive parameterizes local physical variables and the collection defines a continuous spatiotemporal field. Coarse numerical measurements constrain this field through their actual observation process, preserving the relationship between the latent high-resolution state and the retained data. Retained scientific visualizations constrain the same field through differentiable volume rendering, allowing rendering-derived discrepancies to update the underlying physical variables rather than a separate appearance representation. Physical and geometric knowledge is incorporated through dataset-specific compatibility operators. To avoid overriding content already supported by active observations, physics-driven updates are preferentially directed toward weakly observed directions and accepted only after finite forward verification. Gaussian support is optimized jointly with field variables to represent nonuniform structures. Thus, different information sources retain their own scientific semantics while constraining the same latent physical state. The HR reference remains inaccessible during reconstruction and is used only for evaluation, while the optimized high-resolution field can be resampled, visualized, and quantitatively analyzed.

We evaluate the proposed method on four scientific datasets at $4\times$ and $8\times$ reconstruction ratios. Our evaluation examines field-level reconstruction accuracy together with scientifically meaningful structures and relevant physical quantities, and controlled ablations assess the contribution of numerical observations, visualization supervision, and physical information.

Our contributions are threefold: (1) we formulate high-resolution scientific-field reconstruction as joint inference from complementary scientific observations and instantiate it with a shared 4D Scientific Gaussian Field whose primitives encode physical variables; (2) we introduce dataset-specific physical compatibility with evidence-preserving refinement, so physics preferentially acts on degrees of freedom that remain weakly constrained by active observations; and (3) we provide controlled evaluation on four datasets at $4\times$ and $8\times$, characterizing the contribution of numerical, visualization, and physical information to reconstruction of the same spatiotemporal field.

2 Background

This section summarizes the technical background needed for our method. Section 2.1 reviews Gaussian primitives as localized continuous spatial representations, while Section 2.2 introduces differentiable Gaussian splatting, which connects a Gaussian representation to image-space observations. Their adaptation to scientific-field reconstruction is described in Section 4.

2.1 Gaussian Primitives

3D Gaussian Splatting (3DGS) represents a 3D scene using a collection of anisotropic Gaussian primitives Kerbl et al. (2023). The spatial support of the i-th primitive is defined as

$$g_i(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_i)^\top \boldsymbol{\Sigma}_i^{-1}(\mathbf{x}-\boldsymbol{\mu}_i)\right), \tag{1}$$

where $\boldsymbol{\mu}_i \in \mathbb{R}^3$ denotes its center and $\boldsymbol{\Sigma}_i \in \mathbb{R}^{3\times 3}$ determines its anisotropic spatial extent. The covariance is commonly parameterized as $\boldsymbol{\Sigma}_i = \mathbf{R}_i\mathbf{S}_i\mathbf{S}_i^\top\mathbf{R}_i^\top$, where $\mathbf{R}_i$ and $\mathbf{S}_i$ encode rotation and scale, respectively. Thus, each Gaussian provides an explicit and continuously evaluable local support whose position and spatial extent can be optimized. Conventional 3DGS further associates each primitive with opacity and appearance attributes for rendering. Our framework retains this spatial representation but assigns scientific state variables to the Gaussian primitives, as detailed in Section 4.1.

2.2 Differentiable Gaussian Splatting

Gaussian splatting renders 3D primitives by projecting them onto the image plane and compositing their screen-space contributions Kerbl et al. (2023). Dynamic extensions have also used explicit

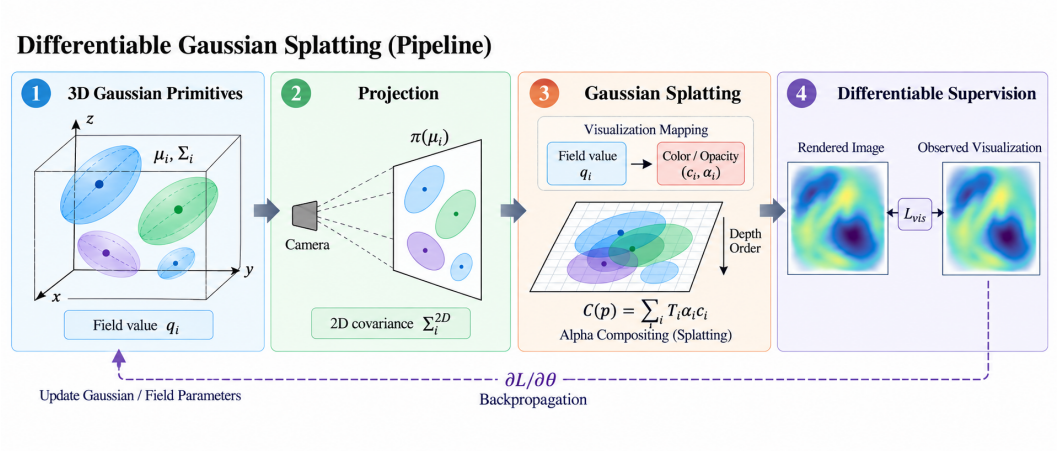


Figure 1: **Differentiable Gaussian splatting.** 3D Gaussians are projected and depth-composited into a rendered visualization, while image-space gradients propagate back to the Gaussian and field parameters.

spatiotemporal Gaussian representations Wu et al. (2024). Figure 1 illustrates both the forward rendering process and the backward gradient flow.

Given a rendering configuration ϕ, the Gaussian center is projected as $\mathbf{u}_i = \pi_\phi(\boldsymbol{\mu}_i)$. Under a local affine approximation, its covariance becomes

$$\boldsymbol{\Sigma}_i^{2D} = \left[\mathbf{J}_i \mathbf{W} \boldsymbol{\Sigma}_i \mathbf{W}^\top \mathbf{J}_i^\top\right]_{2\times 2}, \tag{2}$$

where $\mathbf{W}$ denotes the viewing transformation and $\mathbf{J}_i$ is the Jacobian of the projection. The corresponding screen-space footprint is

$$\rho_i(\mathbf{p}) = \exp\left(-\frac{1}{2}(\mathbf{p}-\mathbf{u}_i)^\top (\boldsymbol{\Sigma}_i^{2D})^{-1}(\mathbf{p}-\mathbf{u}_i)\right). \tag{3}$$

With $\alpha_i(\mathbf{p}) = o_i \rho_i(\mathbf{p})$, depth-sorted Gaussians are combined through front-to-back alpha compositing,

$$\mathbf{C}(\mathbf{p}) = \sum_i T_i(\mathbf{p})\, \alpha_i(\mathbf{p})\, \mathbf{c}_i, \qquad T_i(\mathbf{p}) = \prod_{j<i} (1-\alpha_j(\mathbf{p})). \tag{4}$$

For scientific visualization, the rendering attributes can be obtained from the scientific state $\mathbf{q}_i$ through a prescribed visualization mapping, $(\mathbf{c}_i, o_i) = \mathcal{V}_\phi(\mathbf{q}_i)$. Because projection, footprint evaluation, visualization mapping, and alpha compositing are differentiable, an image-space objective can propagate gradients to both Gaussian and field parameters. Section 4.3 uses this property to incorporate scientific visualizations into field reconstruction.

3 Problem Formulation and Motivation

Reconstruction setting. Let $\mathbf{q}^\star(\mathbf{x}, t) : \Omega \times \mathcal{T} \to \mathbb{R}^C$ denote the unknown high-resolution (HR) scientific field, where $\mathbf{x} \in \Omega \subset \mathbb{R}^3$ denotes space, $t \in \mathcal{T}$ time, and C the number of physical variables. The HR field itself is unavailable during reconstruction and is used only for evaluation. For dataset d, the available observations are generated from the same underlying field through different forward processes:

$$\underbrace{\mathbf{Y}_{\text{num}} = D_d(\mathbf{q}^\star)}_{\text{numerical observation}}, \qquad \underbrace{\mathbf{Y}_{\text{vis}}^{(m)} = R(\mathbf{q}^\star; \phi_m)}_{\text{visualization observation}}, \tag{5}$$

where D_d denotes the dataset-specific numerical observation process, R the scientific rendering process, and ϕ_m the known configuration of the m-th visualization. Available physical, geometric, and boundary knowledge is denoted by $\mathcal{K}_d$.

Why complementary observations? Numerical reduction removes fine-scale information, so distinct HR fields may satisfy

$$\mathbf{q}_1 \neq \mathbf{q}_2, \qquad D_d(\mathbf{q}_1) \approx D_d(\mathbf{q}_2), \tag{6}$$

while differing substantially in structures unresolved by the retained measurements. A second field-dependent observation need not share this ambiguity: $R(\mathbf{q}_1; \phi_m)$ and $R(\mathbf{q}_2; \phi_m)$ may remain distinguishable under the same visualization configuration. This observation-operator view is standard in nonlinear inverse problems and sparse reconstruction, where the local forward map and its sensitivity determine which state directions are observable Egbert & Kelbert (2012); Manohar et al. (2018). The retained observations therefore define an evidence-consistent set

$$\mathcal{S}_{\text{obs}} = \left\{ \mathbf{q} : D_d(\mathbf{q}) \approx \mathbf{Y}_{\text{num}},\ R(\mathbf{q}; \phi_m) \approx \mathbf{Y}_{\text{vis}}^{(m)}\ \forall m \right\}. \tag{7}$$

Physical knowledge can further reduce residual ambiguity in $\mathcal{S}_{\text{obs}}$ without invalidating content already supported by the observations.

Unified formulation. We therefore seek a single spatiotemporal physical representation $G_\Theta(\mathbf{x}, t)$ that explains the available evidence through its corresponding mechanism:

$$\begin{aligned} D_d(G_\Theta) &\approx \mathbf{Y}_{\text{num}}, \\ R(G_\Theta; \phi_m) &\approx \mathbf{Y}_{\text{vis}}^{(m)} \qquad \forall m, \\ G_\Theta &\models \mathcal{K}_d. \end{aligned} \tag{8}$$

These relations are not assigned to separate reconstruction levels or independent models; they constrain the same latent physical state. Observations determine what the reconstruction must preserve, while physical knowledge preferentially refines what remains weakly determined. Section 4 instantiates this formulation with a 4D Scientific Gaussian Field.

4 METHOD

Following Section 3, we reconstruct the unknown state with one shared spatiotemporal representation. Numerical and visualization observations constrain it through their forward processes, while dataset-specific physical knowledge refines residual ambiguity. The final output is the high-resolution physical field itself. Figure 2 summarizes this workflow. Heterogeneous numerical and visualization evidence is first mapped to the same 4D Scientific Gaussian Field. The field is then passed through the corresponding differentiable forward operators, whose discrepancies define observation losses. In parallel, physical and geometric compatibility proposes a candidate update; observation sensitivity projects this update toward weakly constrained directions, and finite forward verification accepts or backtracks the update according to the leakage budgets.

4.1 4D SCIENTIFIC GAUSSIAN FIELD

Shared physical representation. We represent the scientific state using N Gaussian primitives defined over the joint spatiotemporal coordinate $\mathbf{z} = [\mathbf{x}^\top, t]^\top \in \mathbb{R}^4$. The i-th primitive is parameterized by $\mathcal{G}_i = (\boldsymbol{\mu}_i, \boldsymbol{\ell}_i, m_i, \mathbf{s}_i)$, where $\boldsymbol{\mu}_i$ and $\boldsymbol{\ell}_i$ determine its 4D support, m_i denotes its mass, and $\mathbf{s}_i \in \mathbb{R}^C$ stores local scientific variables rather than RGB appearance. With $\boldsymbol{\sigma}_i = \exp(\boldsymbol{\ell}_i)$, its support is

$$g_i(\mathbf{z}) = \exp\left[-\frac{1}{2} \left\| (\mathbf{z} - \boldsymbol{\mu}_i) \oslash \boldsymbol{\sigma}_i \right\|_2^2 \right], \qquad G_\Theta(\mathbf{x}, t) = \mathcal{A}\big(\{ g_i, m_i, \mathbf{s}_i \}_{i=1}^N \big), \tag{9}$$

where $\mathcal{A}$ denotes the differentiable field aggregation and Θ collects the representation parameters. The same primitive set jointly represents all selected times, so $G_\Theta(\mathbf{x}, t)$ is one continuous 4D field rather than a collection of independent 3D reconstructions. To allocate capacity to nonuniform structures, Gaussian centers and log-scales are optimized jointly with their scientific attributes while primitive masses and topology remain fixed.

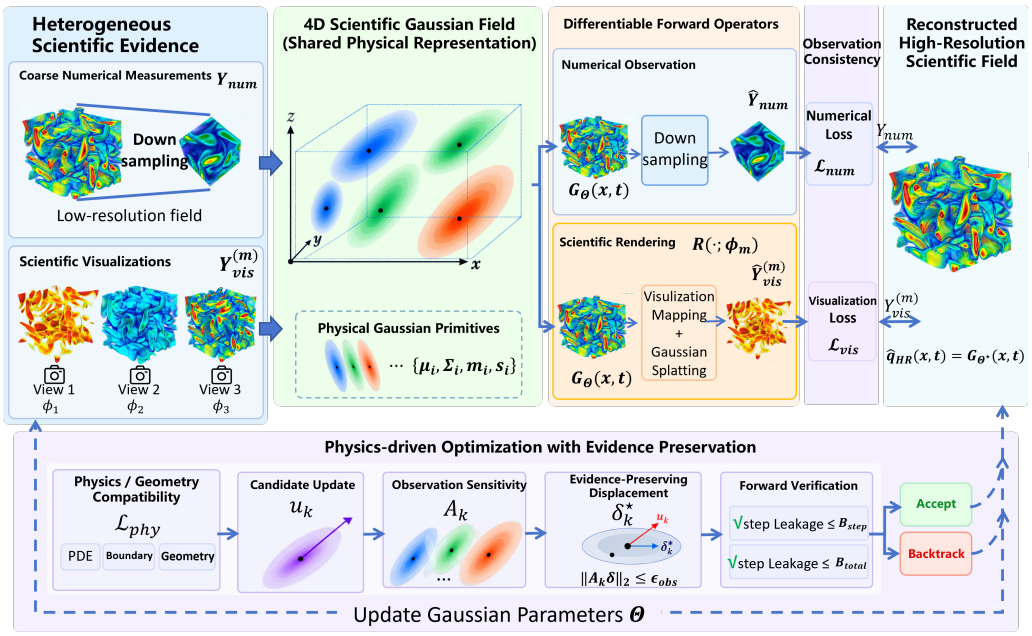


Figure 2: **Overview of the proposed framework.** Coarse numerical measurements and scientific visualizations provide heterogeneous evidence for one shared 4D Scientific Gaussian Field. Differentiable numerical and rendering operators produce observation losses, while physics and geometry propose a candidate update. The update is projected using observation sensitivity and accepted only after finite forward verification, yielding the reconstructed high-resolution scientific field.

4.2 Numerical Observation Consistency

The retained numerical data provide the primary quantitative anchor. Instead of interpreting coarse values as unavailable HR point samples, we reproduce the numerical observation process that generated them:

$$\widehat{\mathbf{Y}}_{\text{num}} = D_d(G_\Theta), \qquad \mathcal{L}_{\text{num}} = \ell_{\text{num}}\left(D_d(G_\Theta), \mathbf{Y}_{\text{num}}\right). \tag{10}$$

The operator D_d preserves the semantics of the retained measurement. For example, if a coarse value represents a cell or block average, $D_d(G_\Theta)$ computes the corresponding average over the predicted field rather than substituting a center-point sample. Thus, numerical supervision constrains the field only through information that is actually available.

4.3 Visualization-Guided Field Reconstruction

A retained scientific visualization supplies another observation of the same physical state through a different forward process. For visualization m with known configuration ϕ_m, the current scientific state is mapped to rendering attributes and rendered through the differentiable Gaussian splatting formulation of Section 2.2, following the broader use of differentiable volumetric rendering for image-space supervision Mildenhall et al. (2022):

$$\widehat{\mathbf{Y}}_{\text{vis}}^{(m)} = R(G_\Theta; \phi_m), \qquad \mathcal{L}_{\text{vis}} = \frac{1}{M}\sum_{m=1}^{M} \ell_{\text{vis}}\left(R(G_\Theta; \phi_m), \mathbf{Y}_{\text{vis}}^{(m)}\right). \tag{11}$$

Rendering attributes are derived from the current scientific state rather than learned as an independent RGB representation. Consequently, image-space discrepancies propagate through R to the same parameters already constrained by $\mathcal{L}_{\text{num}}$. The renderings are complementary measurements, not the reconstruction target. We use $\mathcal{L}_{\text{obs}} = \lambda_{\text{num}}\mathcal{L}_{\text{num}} + \lambda_{\text{vis}}\mathcal{L}_{\text{vis}}$ for observation fitting; physical refinement is coordinated separately.

Table 1: Dataset and observation protocol. HR fields are used only through the evaluation interface and are not available to reconstruction.

Dataset	Variables	HR grid	4× grid	8× grid	Eval. frames
Sphere Re300	u, v, w, p	128^3	32^3	16^3	231–240
Isotropic turbulence	u, v, w, p	128^3	32^3	16^3	291–300
Half-cylinder Re6400	u, v, w, p	$80 \times 240 \times 640$	$20 \times 60 \times 160$	$10 \times 30 \times 80$	141–150
Bubble Pressure Full	u, v, w, p, f	128^3	32^3	16^3	100–109

4.4 Physical Compatibility and Joint Optimization

Dataset-specific physical compatibility. Physical semantics differ across datasets in governing relations, retained variables, geometry, boundaries, and discretization. Physics-informed learning commonly incorporates governing knowledge through residual or compatibility terms Raissi et al. (2019); Karniadakis et al. (2021); Cuomo et al. (2022). For dataset d, we therefore define

$$\mathbf{q}_d^{\mathrm{phy}} = \mathcal{T}_d^{\mathrm{phy}}(G_\Theta), \qquad \mathbf{r}_d = \mathcal{P}_d\left(\mathbf{q}_d^{\mathrm{phy}}; \mathcal{K}_d\right), \tag{12}$$

where $\mathcal{T}_d^{\mathrm{phy}}$ converts the reconstructed field to the channels, coordinates, masks, and retained-grid semantics required by the physical model, and $\mathcal{P}_d$ evaluates a valid physical, boundary, or geometric relation. The physical objective is $\mathcal{L}_{\mathrm{phy}} = \mathbb{E}_{\mathbf{z}\sim\Omega_d^{\mathrm{phy}}}[\|\mathbf{r}_d(\mathbf{z})\|_2^2]$. Thus, datasets need not share an identical continuum PDE; the binding must instead remain valid for the retained scientific state.

Evidence-constrained joint optimization. All information sources act on the same parameters Θ, but physical compatibility is not introduced as an unrestricted third term in $\mathcal{L}_{\mathrm{obs}}$. At state Θ_k, let $\mathcal{O}_k(G_\Theta)$ collect the observation blocks protected during the current physical refinement and $A_k = J_\Theta[\mathcal{O}_k(G_\Theta)]_{\Theta_k}$ denote their local parameter sensitivity. Let $\mathbf{u}_k$ be the tentative optimizer displacement proposed by $\mathcal{L}_{\mathrm{phy}}$. We constrain the displacement that would actually be applied:

$$\delta_k^\star = \arg\min_\delta \frac{1}{2}\|\delta - \mathbf{u}_k\|_{M_k}^2 \quad \text{s.t.} \quad \|A_k\delta\|_2 \le \epsilon_{\mathrm{obs}}, \tag{13}$$

where M_k denotes the optimizer-space metric and ϵ_{obs} the permitted first-order observation change. This preferentially directs physics toward weakly observed degrees of freedom.

Because Eq. equation 13 is only a local condition, the candidate $\Theta_k' = \Theta_k + \alpha_k\delta_k^\star$ is re-evaluated through the actual nonlinear observation processes. Writing $\ell_k = \rho(\mathcal{O}_k(G_{\Theta_k'}), \mathcal{O}_k(G_{\Theta_k}))$, with ρ a normalized observation discrepancy, the update is accepted only if $\ell_k \le B_{\mathrm{step}}$ and $C_{k-1} + \ell_k \le B_{\mathrm{total}}$, where C_{k-1} is the cumulative accepted leakage; otherwise α_k is reduced, or the parameters and optimizer state are restored. This preserves evidence-supported content while using physical knowledge to refine residual ambiguity. Additional valid observation or compatibility operators can be incorporated without changing the shared reconstruction state. After optimization, $\widehat{\mathbf{q}}_{\mathrm{HR}}(\mathbf{x}, t) = G_{\Theta^\star}(\mathbf{x}, t)$. The HR reference never enters optimization and is used only for evaluation.

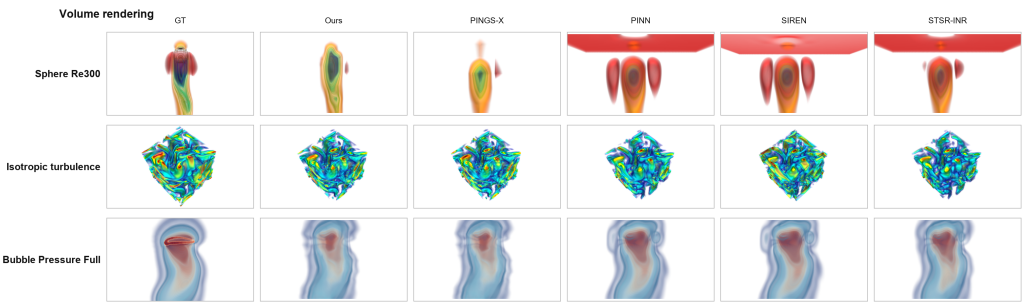


Figure 3: Qualitative volume-rendering comparison at the supplied 8× reconstruction setting. Rows are datasets and columns are the native reference (GT), Ours, and applicable baselines. The same rendering configuration is used within each row; the figure is intended to compare recovered structure and contrast rather than to replace the quantitative metrics in Table 2. Bubble Pressure Full is qualitative only and is excluded from the single-phase aggregate.

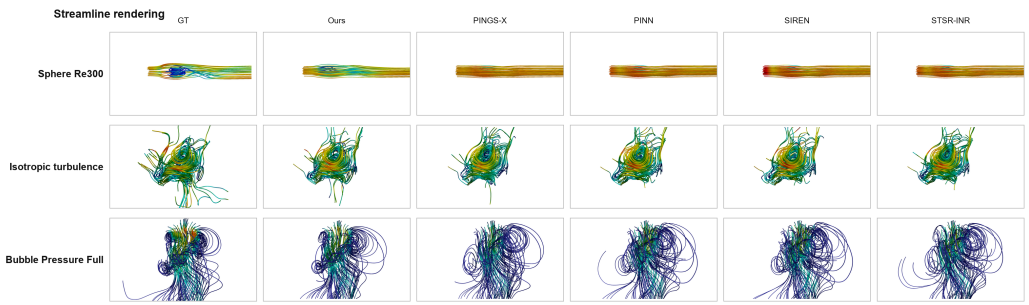


Figure 4: Qualitative streamline comparison at the supplied $8\times$ reconstruction setting. Rows are datasets and columns are the native reference (GT), Ours, and applicable baselines. Identical seed points and integration settings are used within each row, making topology, branch continuity, and recirculation directly comparable.

5 Experiments

We test whether retained visualizations complement reduced numerical observations, whether physical compatibility improves reconstruction without overwriting supported content, and whether the shared 4D field and protected refinement are necessary. We describe the datasets and evaluation protocol below, followed by quantitative, qualitative, and ablation results.

5.1 Datasets, Training, and Evaluation Protocol

Datasets and observations. We evaluate on four single-phase datasets—Sphere Re300, isotropic turbulence, Tangaroa, and Half-cylinder Re6400—and retain Bubble Pressure Full as an additional two-phase qualitative case. The HR and prescribed $4\times/8\times$ LR grids and evaluation frames are listed in Table 1; the Tangaroa entries remain provisional. Single-phase fields contain (u, v, w, p), while the bubble case additionally contains the liquid volume fraction f. Numerical observations are generated by the dataset-specific operator D_d, with the same cell/block averaging applied to reconstructed fields. Visualization observations are $\mathbf{Y}_{\text{vis}}^{(m)} = R(\mathbf{q}^\star; \phi_m)$, using recorded camera, sampling, scalar-mapping, and transfer-function settings.

Training and physical compatibility. The HR reference $\mathbf{q}^\star$ is used only for evaluation and never for initialization, optimization, refinement, hyperparameter selection, early stopping, or model selection. Controlled variants share the same LR observations, initialization, Gaussian budget, optimizer, observation schedule, stopping rule, and evaluation frames. Numerical and visualization fitting precede the physical refinement in Section 4.4. Physical operators follow dataset semantics: the single-phase cases use their corresponding incompressible-flow, geometry, boundary, and forcing conditions, whereas Bubble Pressure Full uses a variable-density/viscosity two-phase formulation with VOF transport, gravity, surface tension, and moving-frame correction. Reproducibility records include optimizer settings, loss weights, refinement intervals, random seeds, hardware, and runtime.

Baselines. We compare against methods compatible with the same LR numerical observations and reconstruction ratios, including tricubic interpolation, PINN-based reconstruction, PINGS-X, and an applicable scientific super-resolution baseline such as PSRFlow. Methods are evaluated only where their input channels and physical/rendering operators are compatible; in particular, a baseline is not applied to the two-phase case without support for volume fraction and interface dynamics. All visual comparisons use identical evaluation frames and rendering settings.

Evaluation metrics. Field fidelity is measured by PSNR between the reconstructed field $\widehat{F}$ and HR reference F,

$$\mathrm{PSNR}(F, \widehat{F}) = 10 \log_{10} \left(\frac{r^2}{\mathrm{MSE}(F, \widehat{F})} \right), \tag{14}$$

Table 2: Mean range-normalized PSNR (dB) over 20 runs for the four single-phase datasets. "All" averages u, v, w, p, velocity magnitude, and vorticity magnitude; "Vel." averages u, v, w. Higher is better.

		4×		8×	
Dataset	Method	All	Vel.	All	Vel.
Sphere Re300	PINGS-X	41.37	42.45	37.07	38.33
	STSR-INR	41.13	42.22	36.38	37.74
	PINN	40.82	41.91	36.04	37.39
	SIREN	40.72	41.81	34.56	35.99
	Ours	**41.39**	**42.47**	**39.14**	**39.40**
Isotropic turbulence	PINGS-X	**26.75**	28.28	20.99	22.14
	STSR-INR	26.70	28.19	20.86	21.88
	PINN	26.64	28.12	20.72	21.69
	SIREN	26.31	27.93	20.39	21.41
	Ours	**26.75**	**28.29**	**25.99**	**23.16**
Tangaroa	PINGS-X	26.61	25.05	26.90	25.29
	STSR-INR	26.35	24.82	26.56	24.90
	PINN	26.29	24.76	26.44	24.75
	SIREN	26.24	24.71	26.36	24.66
	Ours	**28.71**	**29.13**	**27.96**	**26.36**
Half-cylinder Re6400	PINGS-X	28.84	29.64	29.04	29.79
	STSR-INR	28.52	29.30	28.69	29.44
	PINN	28.40	29.16	28.51	29.22
	SIREN	28.33	29.09	28.40	29.10
	TV	29.14	29.95	29.32	**30.08**
	Ours	**31.46**	**31.19**	**30.63**	29.33

Table 3: Computational cost of CoSGF for the reported configuration. Time is measured in minutes and rendering throughput in frames per second. Hardware, measurement repetitions, and variability will be reported in the final version.

Method	Ratio	Peak memory (GB)	Time (min)	Render (FPS)
CoSGF	4×	8.4	14.6	34.7
CoSGF	8×	8.7	16.9	33.9

where r is the prescribed dynamic range after dataset normalization; scores are computed per physical channel and averaged over evaluation frames. Visual fidelity is measured by LPIPS on renderings produced with identical cameras, sampling, scalar mappings, and transfer functions, averaged over fixed views and frames. Physical consistency and observation preservation are reported separately in the method-validation experiments.

5.2 Qualitative and Quantitative Analysis

Quantitative reconstruction accuracy. Table 2 reports range-normalized PSNR over the four single-phase datasets. The "All" metric averages u, v, w, p, velocity magnitude, and vorticity magnitude, while "Vel." averages the three velocity components. CoSGF achieves consistently strong reconstruction across datasets and ratios, with particularly clear gains in several more challenging $8\times$ settings. For Sphere Re300, the methods are close at $4\times$, whereas at $8\times$ CoSGF reaches 39.14 dB on the All metric compared with 37.07 dB for PINGS-X. The separation is larger for isotropic turbulence, where CoSGF improves the $8\times$ All-field PSNR from 20.99 to 25.99 dB. On Half-cylinder Re6400, CoSGF obtains 31.46 and 30.63 dB on the All metric at $4\times$ and $8\times$, respectively.

These results indicate that the benefit is especially visible when coarse numerical observations leave more fine-scale structure unresolved. This is consistent with the motivation of CoSGF: visualization observations provide complementary constraints on structures weakly determined by the LR samples, while physical refinement further reduces residual ambiguity. The improvement is not uniform for every individual metric—for example, TV has a higher $8\times$ velocity score on Half-cylinder—but the overall results show that CoSGF remains effective under stronger numerical reduction.

Table 4: Ablation point estimates. PSNR is reported in dB; lower LPIPS, observation leakage, and physical residual are better. The physical residual is normalized to the full-model value. The final version will report the independent-seed count and uncertainty.

	PSNR ↑		LPIPS ↓		Leakage (%) ↓	Residual ↓
Variant	4×	8×	4×	8×		
Full CoSGF	**33.07**	**30.66**	**0.142**	**0.168**	0.7	**1.00**
−Vis	32.99	30.54	0.176	0.203	0.5	1.02
−Phy	33.01	30.57	0.145	0.171	0.6	1.63
−Shared	32.92	30.42	0.181	0.214	0.8	1.18
−Protect	33.00	30.53	0.149	0.177	3.2	1.02

Qualitative reconstruction. Figures 3 and 4 examine complementary structural properties using the supplied $8\times$ reconstructions with matched visualization settings. In the volume-rendering comparison, CoSGF preserves both global organization and localized structures: the Sphere Re300 wake remains compact and separated, isotropic turbulence retains connected vortical structures and local variation, and the bubble reconstruction preserves the interface geometry and local contrast. Bubble Pressure Full is used only as a qualitative two-phase demonstration and is excluded from the single-phase aggregate because equivalent two-phase baseline support is not established.

The streamline comparison further evaluates the reconstructed velocity field using identical seed points and integration settings. Because trajectory integration accumulates local velocity errors, it exposes differences in direction, connectivity, and recirculation that may be less apparent from pointwise metrics. CoSGF preserves the downstream recirculation pattern in the sphere case and more of the intertwined loops and outer trajectories in isotropic turbulence, while competing reconstructions more often smooth or terminate these structures.

Together, these evaluations characterize complementary aspects of the same reconstructed physical field: PSNR measures field-level agreement, volume rendering reveals localized spatial structure, and streamlines probe coherent transport geometry. Their consistent trends support the use of complementary numerical and visualization observations together with adaptive Gaussian support and evidence-preserving physical refinement, particularly when fine-scale structures are weakly constrained by coarse numerical measurements.

5.3 Computational Cost and Ablation Study

Table 3 reports the computational cost of CoSGF under the reported configuration, including peak memory, runtime, and rendering throughput. Runtime excludes data loading, while rendering throughput is measured using the fixed visualization settings adopted throughout evaluation.

For the ablation study, we remove visualization supervision (−Vis), physical compatibility (−Phy), the shared 4D field (−Shared), or evidence protection (−Protect), while keeping the initialization and optimization budgets fixed. As shown in Table 4, removing visualization supervision degrades LPIPS, removing physical compatibility substantially increases the physical residual, and removing evidence protection produces the largest observation leakage. Removing the shared 4D representation degrades both PSNR and LPIPS, indicating the importance of jointly representing the spatiotemporal field.

6 Conclusion

We presented CoSGF, a shared 4D Scientific Gaussian Field for reconstructing high-resolution physical fields from complementary numerical and visualization observations. Sensitivity-aware, forward-verified physical refinement improves weakly constrained structure while preserving observation-supported content. CoSGF achieves strong PSNR across datasets and ratios, with pronounced gains in challenging $8\times$ settings and improved structural fidelity in qualitative evaluations. The method assumes calibrated visualization metadata and valid dataset-specific operators; transfer under missing or misspecified information remains open.

REFERENCES

Sébastien Barthélémy, Julien Brajard, Laurent Bertino, and François Counillon. Super-resolution data assimilation. *Ocean Dynamics*, 72(8):661–678, 2022. doi: 10.1007/s10236-022-01523-x.

Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020. doi: 10.1146/annurev-fluid-010719-060214.

Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what's next. *Journal of Scientific Computing*, 92(3):88, 2022. doi: 10.1007/s10915-022-01939-z.

Gary D. Egbert and Anna Kelbert. Computational recipes for electromagnetic inverse problems. *Geophysical Journal International*, 189(1):251–267, 2012. doi: 10.1111/j.1365-246X.2011.05347.x.

Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics*, 32(6):1–13, 2013. doi: 10.1145/2508363.2508377.

George E. Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021. doi: 10.1038/s42254-021-00314-5.

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.

Nikola B. Kovachki, Zongyi Li, Kian Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operators: Learning maps between function spaces. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Zongyi Li, Nikola Kovachki, Kian 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.

Xin Liang, Sheng Di, Dingwen Tao, Sihuan Li, Shaomeng Li, Hanqi Guo, Zizhong Chen, and Franck Cappello. Error-controlled lossy compression optimized for high compression ratios of scientific datasets. In *2018 IEEE International Conference on Big Data*, pp. 438–447, 2018. doi: 10.1109/BigData.2018.8622520.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George E. Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021. doi: 10.1038/s42256-021-00302-5.

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*, 38(3):63–86, 2018. doi: 10.1109/MCS.2018.2810460.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2022. doi: 10.1145/3503250.

Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

Jingyi Shen and Han-Wei Shen. Psrflow: Probabilistic super resolution with flow-based models for scientific data. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2023. doi: 10.1109/TVCG.2023.3327171.

Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

Xinjie Wang, Maoquan Sun, Yundong Guo, Chunxin Yuan, Xiang Sun, Zhiqiang Wei, and Xiaogang Jin. Octree-based hierarchical sampling optimization for the volumetric super-resolution of scientific data. *Journal of Computational Physics*, 502:112804, 2024. doi: 10.1016/j.jcp.2024.112804.

Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20310–20320, 2024. doi: 10.1109/CVPR52733.2024.01920.

Skylar W. Wurster, Hanqi Guo, Han-Wei Shen, Tom Peterka, and Jiayi Xu. Deep hierarchical super resolution for scientific data. *IEEE Transactions on Visualization and Computer Graphics*, 29: 5483–5495, 2023. doi: 10.1109/TVCG.2022.3214420.

Mustafa Z. Yousif, Linqi Yu, and Hee-Chang Lim. Super-resolution reconstruction of turbulent flow fields at various reynolds numbers based on generative adversarial networks. *Physics of Fluids*, 34(1):015130, 2022. doi: 10.1063/5.0074724.

Youjia Zhang, Teng Xu, Junqing Yu, Yuteng Ye, Yanqing Jing, Junle Wang, Jingyi Yu, and Wei Yang. Nemf: Inverse volume rendering with neural microflake field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 22862–22872, 2023. doi: 10.1109/ICCV51070.2023.02095.

A Appendix

You may include other additional sections here.