

LagrangianSplats: Divergence-Free Transport of Gaussian Primitives for Fluid Reconstruction

NINGXIAO TAO, School of Intelligence Science and Technology, Peking University, China

BAOQUAN CHEN*, Peking University, China

MENGYU CHU*, State Key Laboratory of General Artificial Intelligence, Peking University, China

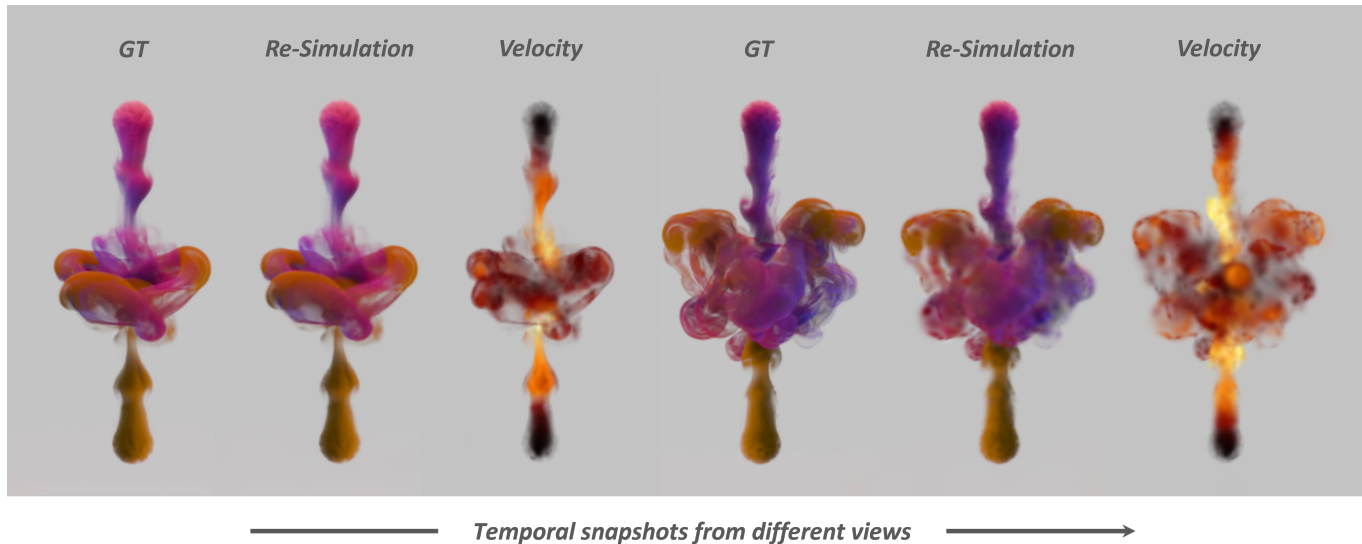


Fig. 1. Reconstruction results on the Biplume dataset, featuring highly turbulent flow dynamics. We compare our re-simulation results against the Ground Truth from two distinct viewpoints at different timestamps. Additionally, we visualize the predicted velocity magnitude using a heatmap, demonstrating our method’s ability to capture intricate flow details. Quantitatively, our re-simulation achieves high fidelity, with sequence-averaged metrics of **37.70 PSNR**, **0.9764 SSIM**, and **0.0758 LPIPS** over the full re-simulation sequence, indicating rigorous transport consistency enabled by a physically plausible velocity reconstruction that is divergence-free by construction.

Reconstructing 3D fluid velocity fields from sparse 2D video observations is a highly ill-posed inverse problem, demanding both transport consistency with observed motion and physical validity under fluid laws. Existing methods typically impose these constraints through soft penalties, often leading to compromised accuracy and convergence issues. We introduce a reconstruction framework that structurally enforces both constraints. Specifically, we parameterize the reconstructed velocity using a continuous Divergence-Free Kernel representation, driving the advection of a Lagrangian 3D Gaussian Splatting representation. This formulation intrinsically guarantees both flow incompressibility and long-range transport coherence by construction. To enable the efficient optimization of such a constrained system, we introduce a novel Sliding Window scheme that propagates gradients over meaningful temporal horizons while maintaining tractable training costs. Experiments on synthetic and real-world datasets demonstrate that our method outperforms state-of-the-art baselines in both transport consistency and physical accuracy, enabling applications such as

high-quality re-simulation and flow analysis. Our implementation is released at <https://github.com/taoningxiao/LagrangianSplats.git>.

CCS Concepts: • **Computing methodologies** → **Physical simulation; Reconstruction; Volumetric models.**

Additional Key Words and Phrases: Fluid Reconstruction, Divergence-Free Representation, 3D Gaussian Splatting


ACM Reference Format:

Ningxiao Tao, Baoquan Chen, and Mengyu Chu. 2026. LagrangianSplats: Divergence-Free Transport of Gaussian Primitives for Fluid Reconstruction. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26)*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3799902.3811188>

1 INTRODUCTION

Fluid phenomena are ubiquitous and captivating. Reconstructing their underlying dynamics, specifically the 3D velocity field, from sparse 2D videos is a critical pursuit in computer vision and graphics. This capability unlocks broad applications ranging from special effects [Gregson et al. 2014] and aerodynamics analysis [Rosset et al. 2023] to high-fidelity forecasting [Bi et al. 2023]. However, this task presents a formidable ill-posed inverse problem due to the inherent

*corresponding authors

Please use nonacm option or ACM Engage class to enable CC licenses. 
This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
SIGGRAPH Conference Papers '26, July 19–23, 2026, Los Angeles, CA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2554-8/2026/07
<https://doi.org/10.1145/3799902.3811188>

ambiguity of inferring 3D flow from 2D visual inputs. A successful reconstruction of the velocity field must simultaneously satisfy two rigorous criteria. First, transport consistency requires that the reconstructed velocity field correctly advects the smoke density to reproduce the observed temporal evolution. Second, physical validity demands that the recovered flow obeys fundamental fluid laws, such as incompressibility and momentum conservation. Achieving both objectives within a unified framework remains a significant challenge.

Existing methods struggle to enforce strict transport consistency. Eulerian approaches like PINF [Chu et al. 2022] and HyFluid [Yu et al. 2023] model density and velocity as neural fields. While demonstrating the potential of joint optimization, their Eulerian nature lacks explicit long-term temporal correspondence. Their short-range advection supervision often traps velocity optimization in local minima, preventing efficient convergence to globally accurate transport.

Several methods attempt to improve long-term transport consistency. PICT [Wang et al. 2024] enforces consistency using neural trajectory representations. However, it relies on aggressive regularization to ensure trajectory validity, forcing the reconstruction toward overly smoothed dynamics and suppressing turbulent transport. A theoretically more robust solution is GlobalTransport [Franz et al. 2021], which defines the smoke reconstruction by explicitly advecting a canonical initial volume throughout the entire sequence. However, its reliance on differentiable rendering and advection makes it computationally prohibitive, restricting its use to scenarios with known lighting conditions and obstacle geometry. As a scene representation, 3D Gaussian Splatting (3DGS) [Kerbl et al. 2023] offers a promising Lagrangian perspective. FluidNexus [Gao et al. 2025] adapts this paradigm for single-view smoke reconstruction by advecting Gaussians through time. Despite its Lagrangian nature, it relies on greedy frame-to-frame velocity estimation, which discards valuable long-range gradients needed to disambiguate complex flow dynamics.

In parallel to transport issues, ensuring physical validity remains a critical bottleneck. Whether employing implicit neural fields [Chu et al. 2022; Yu et al. 2023; Wang et al. 2024] or Position-Based Fluid representations [Gao et al. 2025], these approaches typically supervise the velocity field using soft PINN [Raissi et al. 2019] penalties derived from fluid laws. However, optimization based on soft constraints is often difficult to converge and cannot strictly guarantee physical plausibility. In particular, failure to strictly enforce the divergence-free condition introduces degeneracies: the reconstruction may explain temporal changes through artificial sources and sinks rather than through material transport. This ambiguity corrupts the velocity recovery and limits physical interpretability.

We present a unified reconstruction framework that enforces both constraints by construction. To ensure physical validity, we parameterize the velocity field using the Divergence-Free Kernel (DFK) representation [Ni et al. 2025]. Eliminating the need for soft penalties, this approach inherently guarantees incompressibility and effectively avoids nonphysical density hallucinations. To achieve transport consistency, we employ a Lagrangian approach where dynamic smoke is represented by 3D Gaussian primitives advected from the initial frame using the DFK velocity field. While optimizing

the full dependency chain offers a strong global constraint, it is computationally prohibitive. Therefore, we introduce a Sliding Window optimization strategy. This mechanism restricts gradient propagation to effective temporal horizons, preserving essential long-range transport logic while ensuring computational tractability.

Our main contributions are summarized as follows:

- We propose a strictly divergence-free fluid velocity field reconstruction framework by integrating 3DGS with the DFK representation, achieving results that are both transport-consistent and physically accurate.
- We introduce a Sliding Window optimization mechanism that enforces Lagrangian transport consistency over meaningful temporal horizons, enabling effective long-range gradient propagation for accurate flow dynamics recovery.
- We demonstrate state-of-the-art performance in velocity field reconstruction on both synthetic and real datasets, facilitating downstream applications such as high-fidelity re-simulation and quantitative flow analysis.

2 RELATED WORK

Fluid Reconstruction from Sparse Views. Reconstructing 3D fluid dynamics is a problem of significant interest across both computer graphics and scientific research. Early scientific approaches like structured light [Gu et al. 2012; Ji et al. 2013; Atcheson et al. 2009] and PIV [Grant 1997; Elsinga et al. 2006; Xiong et al. 2017] relied on specialized hardware, limiting them to laboratory settings. In graphics, efforts shifted towards reconstructing fluids from multi-view RGB videos input. ScalarFlow [Eckert et al. 2019] demonstrated this by integrating differentiable fluid solvers into the reconstruction loop. Notably, GlobalTransport [Franz et al. 2021] introduced a strict Lagrangian formulation, defining the smoke state by explicitly advecting a canonical initial volume. While this guarantees long-term transport consistency, its reliance on traditional differentiable volume rendering incurs computational bottlenecks.

Implicit Neural Representations (INRs) offer a promising continuous alternative to traditional discretizations. This paradigm has achieved remarkable success across geometry modeling [Park et al. 2019; Saito et al. 2019; Mescheder et al. 2019; Peng et al. 2020], neural rendering [Mildenhall et al. 2020; Barron et al. 2021; Müller et al. 2022], and physics simulation [Chen et al. 2023]. In the context of fluid reconstruction, recent works have leveraged INRs to jointly model density and velocity fields. PINF [Chu et al. 2022] and HyFluid [Yu et al. 2023] incorporate physics-informed losses to enhance physical plausibility, while NeuSmoke [Qiu et al. 2024] combines 3D neural fields with 2D refinement for efficiency. To enforce temporal coherence, PICT [Wang et al. 2024] adopts a trajectory-based representation for long-term supervision. However, its over-constrained formulation often limits its ability to capture fine-grained turbulent details. More broadly, INRs generally depend on PINN-based objectives [Raissi et al. 2019], whose ill-conditioned optimization landscape often hinders effective training and limits reconstruction fidelity [Rathore et al. 2024; Cao and Zhang 2025].

3D Gaussian Splatting for Dynamic and Fluid Scenes. The advent of 3D Gaussian Splatting [Kerbl et al. 2023] has catalyzed a shift toward dynamic scene reconstruction. Numerous approaches [Luiten et al.

2024; Yang et al. 2024; Wu et al. 2024; Duan et al. 2024; Hong et al. 2025] extend 3DGS to the temporal domain, primarily prioritizing visual fidelity in general scenes.

To enable physically plausible fluid reconstruction, recent works have integrated physics-informed supervision into the 3DGS framework. GaussFluids [Du et al. 2025] incorporates a density-based soft constraint to guide particle distribution and encourage incompressibility, yet it lacks explicit momentum supervision governed by the Navier-Stokes equations. FluidGS [Xie et al. 2025] transfers Gaussian density fields to Eulerian grids to apply PINN constraints; however, this supervision remains a soft regularization restricted to a localized temporal horizon. Most closely related to our work is FluidNexus [Gao et al. 2025], which sharing a similar Lagrangian motivation to advect 3DGS using Position-Based Fluid representation [Macklin and Müller 2013]. However, FluidNexus adopts a sequential, frame-by-frame estimation strategy, which discards valuable long-range gradients from future frames. Our approach diverges by addressing the problem through a global transport perspective, capturing long-range temporal dependencies often overlooked by greedy, sequential updates.

Divergence-Free Velocity Representation. Divergence-free motion is critical for realistic fluid simulation. Traditional solvers typically enforce incompressibility via pressure projection [Zhu and Bridson 2005; Jiang et al. 2015; Bridson 2015; Stam 2023]. Recent research has explored representations that are divergence-free by construction, such as maintaining vector potentials [Chang et al. 2021; Lyu et al. 2024] or employing specialized grid interpolation schemes [Nabizadeh et al. 2024; Roy-Chowdhury et al. 2024]. In the data-driven domain, efforts have been made to encode these physical constraints into neural architectures. For instance, Kim et al. [2019] trained Convolutional Neural Networks to predict velocity potentials, while Richter-Powell et al. [2022] proposed neural parameterizations that inherently satisfy the divergence-free condition. Unlike these grid-based or implicit approaches, we adopt the Divergence-Free Kernel [Ni et al. 2025]. Its mesh-free, analytically incompressible nature naturally complements the Lagrangian 3DGS, ensuring continuous fields suitable for sparse particle advection.

3 BACKGROUND

In this section, we briefly review the foundational representations that underpin our framework: 3D Gaussian Splatting (Sec. 3.1) and Divergence-Free Kernels (Sec. 3.2).

3.1 3D Gaussian Splatting

3D Gaussian Splatting [Kerbl et al. 2023] represents a scene as a set of 3D anisotropic Gaussians. Each Gaussian k is parameterized by its mean position $\boldsymbol{\mu}_k \in \mathbb{R}^3$, covariance matrix Σ_k , opacity $\alpha_k \in [0, 1]$, and view-dependent color represented by spherical harmonics. To ensure positive semi-definiteness, the covariance matrix Σ_k is further decomposed into a scaling matrix \mathbf{S}_k and a rotation matrix \mathbf{R}_k , such that $\Sigma_k = \mathbf{R}_k \mathbf{S}_k \mathbf{S}_k^\top \mathbf{R}_k^\top$.

To render an image from a specific viewpoint, the 3D Gaussians are projected onto the 2D image plane. The pixel color $C(\mathbf{p})$ is computed using volumetric α -blending, sorting the Gaussians from

front to back:

$$C(\mathbf{p}) = \sum_{i \in \mathcal{N}} c_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j), \quad (1)$$

where c_i is the color of the i -th Gaussian, and σ_i is the effective opacity computed by multiplying the learned opacity α_i with the projected 2D Gaussian’s probability density. This explicit representation allows for real-time rasterization and efficient differentiability, which is crucial for our long-range transport optimization.

3.2 Divergence-Free Kernel

DFK [Ni et al. 2025] represent velocity field $\mathbf{v}(\mathbf{x})$ over a spatial domain as a linear combination of matrix-valued kernels $\boldsymbol{\psi}_i$:

$$\mathbf{v}(\mathbf{x}) = \sum_i \boldsymbol{\psi}_i(\mathbf{x}) \boldsymbol{\omega}_i, \quad (2)$$

where $\boldsymbol{\omega}_i \in \mathbb{R}^3$ is the trainable weight vector associated with the i -th kernel centered at \mathbf{x}_i . The matrix-valued kernel $\boldsymbol{\psi}_i : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 3}$ is derived from a scalar radial basis function via a second-order differential operator:

$$\boldsymbol{\psi}_i(\mathbf{x}) = (-\mathbf{I}\nabla \cdot \nabla + \nabla\nabla^\top) \phi_i(\mathbf{x}), \quad (3)$$

where \mathbf{I} is the identity matrix, $\nabla\nabla^\top$ denotes the Hessian operator, and $\phi_i(\mathbf{x})$ is a scalar kernel defined as:

$$\phi_i(\mathbf{x}) = \phi \left(\frac{\|\mathbf{x} - \mathbf{x}_i\|}{h_i} \right), \quad (4)$$

with h_i being the support radius. Following Ni et al. [2025], we employ the Wendland C^4 -continuous piecewise-polynomial function [Wendland 1995] as the radial basis function $\phi(r)$.

By construction, the divergence of the resulting field $\mathbf{v}(\mathbf{x})$ in Eq. (2) is guaranteed to be identically zero for any choice of weights $\boldsymbol{\omega}_i$. Beyond ensuring incompressibility, this representation offers several desirable properties, including compact support, positive definiteness, and second-order differentiability, making it an ideal parameterization for physically plausible fluid reconstruction.

4 METHOD

This section describes our reconstruction framework. Given sparse multi-view videos of smoke, our goal is to jointly recover both the evolving smoke and the driving velocity field. As illustrated in Fig. 2, our physics-integrated framework seamlessly combines the efficient differentiable rendering of 3DGS with the physical guarantees of DFK-parameterized velocity. At its core is a Lagrangian-based optimization scheme that drives advection via a strictly divergence-free velocity field, inherently ensuring both transport consistency and physical plausibility by construction.

We first define our hybrid representation in Sec. 4.1, coupling 3DGS geometry with the DFK velocity parameterization. Building on this model, Sec. 4.2 introduces our Sliding Window strategy, an algorithm designed to capture long-range dynamics effectively by propagating gradients over manageable temporal horizons. Finally, Sec. 4.3 outlines the optimization objectives, comprising reconstruction terms and physics-informed regularizers to guide the joint recovery process.




Fig. 2. **Overview of our proposed framework.** We reconstruct smoke dynamics from sparse-view videos using a hybrid representation of 3DGS (purple splats) for density and DFK (orange nodes) for velocity. To ensure transport consistency, we employ a Sliding Window strategy. Within the active window, the smoke state at intermediate frames is explicitly defined by advecting the Gaussians from the window’s start. As indicated by the icons, the Gaussian parameters at the window start and the velocity fields are optimizable, while the advected states are derived but allow gradients to pass through. Consequently, gradients from the image loss \mathcal{L}_{img} back-propagate through the advection chain (red dashed arrows), jointly refining the initial Gaussian attributes and the velocity field with long-term temporal foresight. Physical validity is further enhanced via vorticity and regularization losses \mathcal{L}_{vor} and \mathcal{L}_{reg} .

4.1 Modeling Smoke with Advected Gaussians

We represent the evolving smoke at time t as a set of 3D Gaussians,

$$\mathcal{G}_t = \{G_i^t\}_{i=1}^N, \quad (5)$$

Crucially, rather than optimizing their positions as free variables, we treat them as passive tracers carried by the DFK-parameterized velocity field $\mathbf{v}(\mathbf{x}, t)$. Specifically, the position of each Gaussian at time t is determined by advecting its initial state from the first frame:

$$\frac{d\boldsymbol{\mu}_i}{dt} = \mathbf{v}(\boldsymbol{\mu}_i, t), \quad \boldsymbol{\mu}_i^t = \mathcal{A}_{0 \rightarrow t}(\boldsymbol{\mu}_i^0, \mathbf{v}), \quad (6)$$

where $\mathcal{A}_{0 \rightarrow t}(\cdot)$ denotes the Lagrangian advection operator. In this way, the movement of Gaussians is explicitly constrained by the advection process, providing a physically grounded and differentiable link between the reconstructed smoke and the underlying flow dynamics.

4.2 Sliding Window Optimization

We propose a sliding window strategy to jointly optimize the smoke and the underlying velocity field from sparse video inputs. Given a sequence of N frames and a window size w , the optimization proceeds in two phases: an initial warm-up phase and a subsequent sliding phase.

Initial Warm-up Phase. We begin by reconstructing the smoke at the first frame ($t = 0$) using standard 3D Gaussian Splatting to

provide an initial guess. The optimization then focuses on the initial window spanning frames $[0, w]$, where the first frame ($t = 0$) acts as the *anchor frame*. Within this window, the intrinsic attributes (e.g., opacity) of the anchor frame are actively optimized as trainable parameters to refine the initial guess. For any subsequent frame $t \in (0, w]$, the positions of Gaussians are derived via the advection operator $\mathcal{A}_{0 \rightarrow t}$ defined in Eq. (6). To enforce short-term physical consistency, the intrinsic attributes of these advected Gaussians strictly inherit the updated values from the anchor and remain constant along the intra-window trajectory.

To ensure training stability, we adopt a progressive strategy. Starting from a short horizon, we gradually extend the advection range from $t = 1$ to $t = w$. At each step, the advected Gaussians are rendered and supervised by an image reconstruction loss \mathcal{L}_{img} , which combines L1 and D-SSIM terms [Kerbl et al. 2023]. Gradients from this loss are back-propagated analytically to update the time-varying velocity field across the window and the parameters of the Gaussians at the initial frame.

Crucially, when aggregating gradients over the window, we account for the diminishing influence of distant physical interactions by applying a temporal discount factor $\gamma \in (0, 1)$. Specifically, to update the parameters at frame t using future supervision, the objective is formulated as a weighted sum:

$$\mathcal{L}_{\text{total}} = \sum_j \gamma^j \mathcal{L}_{\text{img}}(t + j), \quad (7)$$

where j represents the temporal distance from the current frame t to the supervised frame $t + j$. In our implementation, we set $\gamma = 0.9$ as the default value. This discounting strategy effectively suppresses noise arising from long-range dependencies, thereby stabilizing the joint optimization of the smoke and the underlying velocity field.

Subsequent Sliding Phase. After the warm-up, we transition to the sliding phase. Consider the window shifting from the span $[s - 1, s - 1 + w]$ to $[s, s + w]$, where the start frame index s progresses from 1 to $N - 1 - w$. At this stage, the Gaussians at frame $s - 1$ have been optimized by the previous windows sufficiently. Since the temporal discount factor γ renders the impact of gradients from future frames negligible on this distant past state, we fix its parameters to truncate the computational graph.

In the new window, optimization follows a similar advection logic as the initial warm-up phase, but with a key difference in handling the anchor frame s . To strictly enforce global Lagrangian consistency, the positions of Gaussians at frame s are not optimized directly. Instead, they are constrained to be the result of advecting fixed positions from frame $s - 1$ using the velocity field \mathbf{v}_{s-1} . Consequently, position-related gradients back-propagated to frame s will flow further back, optimizing velocity \mathbf{v}_{s-1} . This ensures the advection trajectory connecting the windows remains physically continuous.

While positions are strictly constrained by the advection from the previous window, the intrinsic attributes (e.g., opacity) of the new anchor frame s are actively optimized. We initialize these attributes using values from frame $s - 1$ to promote temporal smoothness, but allow them to be optimized as free variables. This enables the model to capture progressive appearance changes like the diffusion of the smoke. Crucially, the remaining frames in the window ($t \in (s, s + w)$) are advected from this new anchor, strictly inheriting its updated intrinsic attributes just as in the warm-up phase. Finally, as the majority of frames in the new window are already pre-optimized, we bypass the progressive schedule and directly optimize over the full window width w . This process repeats until the end of the sequence, efficiently scaling our algorithm to arbitrary video lengths while enforcing a continuous physical trajectory.

Additionally, to model continuous smoke emission, we designate a specific inflow region. When a new frame first enters the sliding window, we augment the model by injecting new Gaussians into this region. Once added, these new Gaussians are treated identically to the other advected ones, fully participating in the standard velocity-driven optimization process.

4.3 Optimization Objectives

The overall optimization objective is designed to jointly ensure high-fidelity visual reconstruction and physical plausibility. It comprises two primary components: reconstruction losses that supervise the smoke’s geometry and appearance against observed videos, and physics-informed velocity constraints that ensure the underlying velocity field to adhere to fluid dynamics principles.

Reconstruction Losses. As defined in Eq. (7), our primary supervision comes from the discounted image reconstruction loss. Given a

rendered frame and ground truth, the per-frame loss is a combination of L1 distance and D-SSIM [Kerbl et al. 2023]:

$$\mathcal{L}_{\text{img}} = (1 - \lambda_{\text{ssim}})\mathcal{L}_1 + \lambda_{\text{ssim}}\mathcal{L}_{\text{D-SSIM}}, \quad (8)$$

where we set $\lambda_{\text{ssim}} = 0.2$.

Additionally, to prevent visual artifacts caused by overly skinny Gaussian particles, we also apply the anisotropic regularization term $\mathcal{L}_{\text{aniso}}$ proposed by Xie et al. [2025]:

$$\mathcal{L}_{\text{aniso}} = \sum (||s_x - s_y||^2 + ||s_y - s_z||^2 + ||s_x - s_z||^2), \quad (9)$$

where s_x, s_y, s_z are Gaussian particles’ scaling factor in each axes.

Physics-Informed Velocity Constraints. While the DFK representation guarantees the velocity field is divergence-free, we further enforce physical plausibility through two loss terms.

First, to penalize spurious motion in empty regions where no smoke density exists to constrain the flow, we impose a regularization based on the L_1 norm of the velocity magnitude:

$$\mathcal{L}_{\text{reg}} = \sum_{\mathbf{x}} |\mathbf{v}(\mathbf{x})|_1. \quad (10)$$

This term encourages the velocity field to remain zero in the absence of smoke, guiding the optimization toward a minimal-energy solution.

Second, to capture realistic fluid dynamics, we enforce the momentum conservation described by the Navier-Stokes equations. Since our field is analytically incompressible, we can work directly with the Vorticity Transport Equation [Cottet 2001]. We minimize the residual of the vorticity evolution:

$$\mathcal{L}_{\text{vor}} = \left| \frac{\partial \boldsymbol{\omega}}{\partial t} + (\mathbf{v} \cdot \nabla) \boldsymbol{\omega} - (\boldsymbol{\omega} \cdot \nabla) \mathbf{v} \right|_1, \quad (11)$$

where $\boldsymbol{\omega} = \nabla \times \mathbf{v}$ is the vorticity. By combining the hard divergence-free constraint of DFK with this soft vorticity transport constraint, our method effectively reconstructs a velocity field that adheres to the full incompressible Navier-Stokes dynamics.

Total Objective. Finally, we integrate the aforementioned terms into a unified objective function to guide the joint optimization process. The total loss is formulated as a weighted sum:

$$\mathcal{L} = \mathcal{L}_{\text{total}} + \lambda_{\text{aniso}}\mathcal{L}_{\text{aniso}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}} + \lambda_{\text{vor}}\mathcal{L}_{\text{vor}}, \quad (12)$$

where $\mathcal{L}_{\text{total}}$ is the temporally discounted objective from Eq. (7). The hyperparameters λ_{aniso} , λ_{reg} and λ_{vor} control the relative importance of the shape regularization, kinetic energy sparsity, and vorticity transport consistency, respectively. By minimizing this composite objective, our method effectively balances the trade-off between transport consistency and physical accuracy.

5 EXPERIMENTS

We evaluate our method against state-of-the-art approaches on a diverse set of synthetic and real-world scenarios. Experiments demonstrate that our framework outperforms existing methods in terms of both transport consistency and physical validity. In the following, we will first detail our experimental setup.

Table 1. Quantitative comparison on the ScalarSyn (synthetic) and ScalarReal (real-world) benchmarks. We evaluate transport consistency using re-simulation metrics (PSNR, SSIM, LPIPS) and assess physical validity via divergence (Div), velocity error (MSE- v), and directional consistency (Cos- v). Note that both MSE- v and Cos- v are computed only within the smoke volume, masked by regions where GT density is greater than 0, to focus on the relevant dynamics. Additionally, while our method guarantees zero divergence analytically, we report the numerical divergence computed via finite differences on the grid for fair comparison with baselines. Our method consistently outperforms all baselines, achieving the highest visual fidelity and the lowest physical errors.

Model	ScalarSyn						ScalarReal					
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Div \downarrow	MSE- $v\downarrow$	Cos- $v\uparrow$	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Div \downarrow	MSE- $v\downarrow$	Cos- $v\uparrow$
PINF	30.06	0.9225	0.08284	0.004024	0.1465	0.03890	31.24	0.9635	0.1020	0.003482	–	–
PICT	31.00	0.8998	0.1000	0.002380	0.1166	0.06472	31.97	0.9490	0.09932	0.002177	–	–
HyFluid	32.41	0.9340	0.08536	0.3268	0.4033	0.1367	32.62	0.9555	0.08186	0.003058	–	–
FluidNexus	31.93	0.9501	0.1087	0.01774	0.1430	0.05362	35.83	0.9781	0.07847	0.04148	–	–
Ours	46.50	0.9930	0.02581	0.000099	0.1116	0.2731	40.00	0.9782	0.06440	0.000008	–	–

Datasets. We conduct quantitative benchmarking on the ScalarSyn dataset [Eckert et al. 2019], which provides ground-truth (GT) velocity fields, and the ScalarReal dataset [Eckert et al. 2019], a real-world capture. Additionally, to assess robustness in challenging scenarios, we introduce three new datasets: Suzanne for smoke interacting with complex obstacle geometries over long sequences, Sphere for boundary interactions and Biplume for high turbulence.

Baselines. We compare our framework with leading fluid reconstruction methods, focusing on representative implicit neural representations: PINF [Chu et al. 2022], HyFluid [Yu et al. 2023], and PICT [Wang et al. 2024]. We also include the recent 3DGS-based approach, FluidNexus [Gao et al. 2025], for a comprehensive comparison.

Evaluation Metrics. We assess reconstruction quality from two complementary perspectives.

First, for transport consistency, we perform the re-simulation task proposed by Yu et al. [2023]. Specifically, we advect the initial reconstructed smoke solely using the recovered velocity field over the entire sequence. The re-simulated results are then rendered and compared with the corresponding ground-truth images. Given the highly sparse input views in our standard setting, we primarily evaluate these re-simulation metrics under the training views. However, to further assess our method’s performance from unobserved angles, we additionally include test-view re-simulation evaluations in the Suzanne scene. For both settings, we report sequence-averaged PSNR, SSIM [Wang et al. 2004], and LPIPS [Zhang et al. 2018], obtained by first computing the metrics for each frame and then averaging them over the full re-simulation sequence.

Second, for physical validity, we measure the average divergence (Div) of the reconstructed velocity field across all datasets to quantify violation of incompressibility. For synthetic datasets where GT velocity is available, we additionally evaluate the flow accuracy using two metrics: the Mean Squared Error (MSE- v) to measure value deviation, and the average Cosine Similarity (Cos- v) to assess directional alignment between the reconstructed and ground-truth velocities. Both metrics are computed within the smoke volume, masked by regions where GT density is greater than zero.

Visualization. For qualitative analysis, we visualize velocity fields via middle slices of the xy , zy , and xz planes arranged side-by-side, following standard protocols [Chu et al. 2022; Wang et al. 2024]. In

this visualization, velocity direction is encoded as color hue and magnitude as brightness. The image border acts as a circular legend to identify directions; for instance, colors appearing at the top of the border (purplish-red) denote upward motion, whereas colors at the bottom (green) denote downward motion.

With the experimental setup established, the remainder of this section is organized as follows: Sec. 5.1 presents a comprehensive comparison against baselines on standard benchmarks. To demonstrate the scalability and generalization of our approach, Sec. 5.2 provides a challenging experiment (the Suzanne scene) featuring a long sequence of smoke interacting with complex obstacles. There, we comprehensively evaluate our method against baselines using quantitative velocity metrics as well as re-simulation tasks in both training and novel testing views. Finally, Sec. 5.3 presents ablation studies. Furthermore, additional evaluations regarding obstacle interactions and robustness in highly turbulent scenarios, along with implementation details, are deferred to the supplementary.

5.1 Comparisons on Standard Benchmarks

We evaluate our method against all baselines on the ScalarSyn and ScalarReal datasets.

We first evaluate transport consistency through re-simulation. As shown in Fig. 3 and Fig. 4, baseline methods often suffer from noticeable density drift and dissipation over long sequences, failing to preserve the smoke’s shape. Our method, conversely, maintains sharp structural details and correct advection trajectories even at late timestamps. As shown in Table 1, this qualitative observation is also mirrored by our quantitative metrics: we consistently outperform all baselines in PSNR, SSIM, and LPIPS. We attribute this robustness to our Sliding Window strategy, which propagates gradients from future observations back to earlier frames. This mechanism leverages temporal foresight to lock the velocity field to the smoke’s true Lagrangian trajectory, ensuring long-term consistency.

Next, we examine the physical validity of the reconstructed velocity fields. As illustrated in Fig. 5 and Fig. 10, our method produces clean flow structures that closely align with the Ground Truth. In comparison, PINF and PICT tend to produce overly smoothed fields lacking fine details, whereas HyFluid exhibits significant chaotic noise. FluidNexus, limited by its discrete PBF representation, often generates discontinuous, spotty velocity fields. Quantitatively, As shown in Table 1, our method achieves the lowest Velocity MSE and




Fig. 3. Comparison of re-simulation results on the ScalarSyn dataset. We visualize a late-time snapshot of the smoke density advected by the reconstructed velocity fields of different methods. Baseline methods suffer from severe density drift and dissipation, failing to maintain the smoke’s shape over long durations. Our method preserves sharp structural details and correct advection trajectories, achieving superior visual fidelity that is visually indistinguishable from the Ground Truth.




Fig. 4. Comparison of re-simulation results on the ScalarReal dataset. Despite the complexity of real-world capture, our method maintains long-term transport consistency, preserving the plume’s structure significantly better than baselines. This confirms the efficacy of our Sliding Window strategy in capturing accurate Lagrangian dynamics from sparse real-world observations.




Fig. 5. Visualization of reconstructed velocity fields on the ScalarSyn dataset. It can be seen that both PINF and PICT produce overly smoothed fields, while HyFluid and FluidNexus exhibit significant noise or discontinuities. In contrast, our method reconstructs clean, coherent flow structures that closely align with the Ground Truth.

highest Cosine Similarity on ScalarSyn, indicating superior accuracy in recovered dynamics. Furthermore, thanks to the DFK representation, we maintain a near-zero divergence across all datasets, which is often orders of magnitude lower than baselines. This confirms that integrating a strictly divergence-free basis by construction is far more effective than optimizing soft physical penalties.

5.2 Evaluation on Suzanne Scene

To comprehensively evaluate our method’s robustness over extended periods and its capability to handle intricate boundary conditions, we introduce the Suzanne scene. This scene simulates a rising plume of smoke interacting with the complex geometry of a monkey head model (Suzanne in Blender [Blender 2011]), spanning a long sequence of 10 seconds.




Fig. 6. Visualization of reconstructed velocity fields on the Suzanne scene. While PINF and PICT can reconstruct plausible overall flow patterns, our method recovers flow structures that more closely align with the Ground Truth. Notably, the velocity field produced by FluidNexus collapses to near-zero, as its extensively hard-coded settings severely limit its generalization to this novel setup.




Fig. 7. Ablation study on the Sliding Window strategy. We compare the re-simulation quality of our method under different temporal discount factors: $\gamma = 0.9$ (our default long-range setting) versus $\gamma = 0.1$ (approximating short-sighted optimization). While the short-sighted model captures the general motion, it suffers from degradation in fine structural details and transport accuracy, as reflected by the lower PSNR, SSIM, and higher LPIPS scores. In contrast, our full model with $\gamma = 0.9$ effectively utilizes long-range temporal foresight to maintain high-fidelity consistency with the Ground Truth.

We benchmark our framework against PINF, PICT, and FluidNexus. Note that HyFluid is excluded from this evaluation because it does not support the reconstruction of colored scene. Additionally, we found that FluidNexus struggles to generalize to this new scenario. Due to its extensively hard-coded settings, such as fixed priors on initial smoke colors, FluidNexus fails entirely to drive the flow optimization, producing a degenerate, near-zero velocity field (as clearly evident in Fig. 6).

Consistent with previous results, our method demonstrates superior performance in both physical validity and transport consistency on this scene, as summarized in Table 2. Quantitatively, our reconstructed velocity fields achieve the best results across divergence (Div), mean squared error (MSE- v), and cosine similarity (Cos- v). As visualized in Figure 6, our method reconstructs a velocity field that most closely resembles the Ground Truth, even in the presence of complex obstacles. For the training-view re-simulation task, our approach consistently outperforms all evaluated baselines in PSNR, SSIM, and LPIPS. As illustrated in Figure 8, our method robustly preserves sharp structural details over the long duration, whereas other approaches suffer from severe density drift.

We also evaluate the re-simulation performance on novel test views to assess the generalization of the reconstructed physics and representations, as quantitatively reported in Table 2 and visually compared in Figure 9. In the novel test view, our method continues to surpass existing approaches in terms of PSNR and SSIM. We note that our method yields slightly higher LPIPS scores on the test views. We attribute this to the Gaussian representation’s inherent tendency to overfit training views. This can compromise high-frequency textural details when rendered from unseen perspectives, presenting an interesting direction for future improvement.

5.3 Ablation Study

We conduct ablation studies on the ScalarSyn dataset to validate the effect of our Sliding Window strategy for transport consistency and the impact of physical regularization terms on velocity accuracy.

Effect of Sliding Window Strategy. To assess the importance of long-range gradient propagation, we compare our default setting where the temporal discount factor γ is 0.9 against a variant with a reduced factor of $\gamma = 0.1$. This lower value severely limits the temporal horizon, approximating a short-sighted optimization. As shown in Fig. 7, reducing the temporal foresight leads to a consistent degradation in re-simulation quality across all metrics, including PSNR, SSIM, and LPIPS. This confirms that propagating gradients from future observations is essential for maintaining high-fidelity transport consistency and recovering accurate flow trajectories.

Effect of Physical Regularizations. While the DFK representation guarantees divergence-free motion by construction, we incorporate additional loss terms to further refine the solution. We evaluate the impact of the kinetic energy regularization \mathcal{L}_{reg} and the vorticity transport loss \mathcal{L}_{vor} by training variants without them. As illustrated in Fig. 11, removing both terms results in a higher velocity Mean Squared Error. Adding the kinetic energy regularization \mathcal{L}_{reg} effectively suppresses noise in empty regions. Incorporating the vorticity transport loss \mathcal{L}_{vor} further improves the flow accuracy, yielding the



Fig. 8. Comparison of training-view re-simulation results on the Suzanne scene. While baseline methods struggle with severe density drift and structural degradation over the long duration, our method robustly handles the obstacle interaction, maintaining sharp structural details and accurate advection trajectories that closely match the Ground Truth.




Fig. 9. Comparison of novel test-view re-simulation results on the Suzanne scene. By rendering the advected smoke from completely unseen camera angles, we evaluate the 3D physical and structural consistency of the reconstructed fields. While baseline methods exhibit clear volumetric artifacts and incorrect density distributions, our framework synthesizes structurally plausible novel views with superior PSNR and SSIM. We also observe that our method yields slightly higher LPIPS scores. We attribute this finding to a degree of overfitting to the training views, which can occasionally compromise fine textural details when rendered from unseen perspectives.




Fig. 10. Visualization of reconstructed velocity fields on the ScalarReal dataset. Compared to baselines which often produce noisy or chaotic motion artifacts in background regions, our method generates a smooth and physically plausible velocity field. The clean flow structure and clear boundaries demonstrate the robustness of our DFK representation on real-world data.

Table 2. Quantitative comparison on the long-duration Suzanne scene. We evaluate across three aspects: physical validity (Velocity Metrics), training-view visual fidelity (Train View Re-Simulation), and novel-view generalization (Test View Re-Simulation). Note that HyFluid is excluded as it does not support colored scene reconstruction. FluidNexus struggles to reconstruct a physically meaningful velocity field and collapses to near zero velocity in this scenario due to its extensively hard-coded settings (e.g., fixed assumptions about initial smoke color). Our method achieves the best performance across all velocity and training-view metrics, and maintains superior PSNR and SSIM on novel test views, demonstrating robust transport consistency and generalization.

Model	Velocity Metrics			Train View Re-Simulation			Test View Re-Simulation		
	Div↓	MSE- v ↓	Cos- v ↑	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
PINF	0.000260	0.02136	0.1947	27.80	0.9241	0.1199	26.91	0.9221	0.1236
PICT	0.000665	0.05257	0.1679	24.34	0.9213	0.1115	23.63	0.9168	0.1171
FluidNexus	Collapses to zero			25.18	0.9231	0.1292	24.21	0.9178	0.1331
Ours	0.000014	0.02120	0.2087	39.54	0.9731	0.1066	29.36	0.9383	0.1392

best performance. This demonstrates that while DFK ensures kinematic validity through mass conservation, these physics-informed losses are valuable for recovering the correct dynamics governed by momentum conservation.

6 CONCLUSION

In this paper, we presented a novel framework for reconstructing high-fidelity fluid velocity fields from sparse multi-view videos. By integrating the efficiency of 3D Gaussian Splatting with the physical rigor of Divergence-Free Kernels, we enforce strict incompressibility by construction. Furthermore, our Sliding Window optimization strategy successfully enforces long-term Lagrangian consistency while maintaining computational tractability. This mechanism enables the propagation of gradients over meaningful temporal horizons, recovering accurate flow dynamics that are both visually consistent with the input observations and physically plausible.

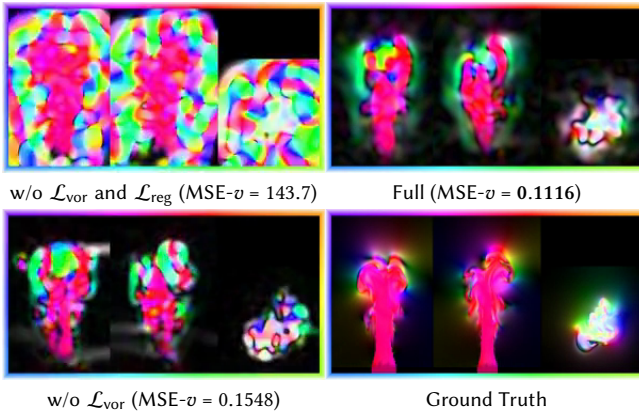


Fig. 11. Ablation study on physical regularization terms. We compare velocity field reconstructions under different loss configurations: (Top-Left) without any physical regularization, (Bottom-Left) with only kinetic energy regularization \mathcal{L}_{reg} , and (Top-Right) our full model including vorticity transport loss \mathcal{L}_{vor} . As shown, removing all regularizers leads to significant background noise and high Velocity MSE. Introducing \mathcal{L}_{reg} suppresses spurious motion in empty regions, drastically reducing the error. Finally, incorporating \mathcal{L}_{vor} further refines the flow structure to match the Ground Truth (Bottom-Right), demonstrating that these physics-informed losses are crucial for recovering accurate flow dynamics beyond basic divergence-free representation.

Despite these advancements, our method is subject to several limitations that point toward exciting avenues for future research.

First, our current formulation is primarily designed for divergence-free, incompressible fluids (e.g., smoke). Because it relies on DFKs, the method does not readily extend to compressible flows, and it can be challenging to accurately capture highly nonlinear and chaotic fluid dynamics, such as shocks. Other phenomena like free-surface liquids, elastic bodies, or cloth are also not covered in the present work. Extending our Lagrangian-based sliding window framework to recover the dynamics and physical parameters of these diverse materials and complex flow regimes remains a promising direction. Second, while DFK guarantees mass conservation (divergence-free), the momentum conservation is only encouraged via soft physics-informed losses. Consequently, discrepancies between the reconstructed velocity and the ground truth may persist in complex scenarios. Developing a velocity parameterization that naturally satisfies momentum transport by construction could lead to improvements in reconstruction accuracy. Third, although our method achieves a substantial speedup compared to implicit neural representations, it is not yet capable of real-time reconstruction. Specifically, our total training time typically ranges from 30 minutes to 3 hours depending on scene complexity, whereas implicit representations often require over ten hours. Detailed per-scene timings are provided in the supplementary. Bridging the gap between offline optimization and real-time inference is crucial for deploying such techniques in time-sensitive applications like interactive aerodynamic analysis. We believe that further optimizing the computational graph or exploring feed-forward prediction models could be key to unlocking this potential.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (Project No. 62595772).

REFERENCES

- Bradley Atcheson, Wolfgang Heidrich, and Ivo Ihrke. 2009. An evaluation of optical flow algorithms for background oriented schlieren imaging. *Experiments in Fluids* 46, 3 (01 Mar 2009), 467–476. <https://doi.org/10.1007/s00348-008-0572-7>
- Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5855–5864.
- K Bi, L Xie, H Zhang, X Chen, X Gu, and Q Tian. 2023. Accurate medium-range global weather forecasting with 3D neural networks. *Nature* 619, 7970 (2023), 533–538.
- Blender. 2011. (CC) Blender Foundation. <https://www.blender.org/>. Online.

- Robert Bridson. 2015. *Fluid simulation for computer graphics*. CRC press.
- Wenbo Cao and Weiwei Zhang. 2025. An analysis and solution of ill-conditioning in physics-informed neural networks. *J. Comput. Phys.* 520 (2025), 113494.
- Jumyung Chang, Ruben Partono, Vinicius C Azevedo, and Christopher Batty. 2021. Curl-flow: Boundary-respecting pointwise incompressible velocity interpolation for grid-based fluids. *arXiv preprint arXiv:2104.00867* (2021).
- Honglin Chen, Rundi Wu, Eitan Grinspun, Changxi Zheng, and Peter Yichen Chen. 2023. Implicit neural spatial representations for time-dependent pdes. In *International Conference on Machine Learning*. PMLR, 5162–5177.
- Mengyu Chu, Lingjie Liu, Quan Zheng, Erik Franz, Hans-Peter Seidel, Christian Theobalt, and Rhaleb Zayer. 2022. Physics informed neural fields for smoke reconstruction with sparse data. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–14.
- Georges-Henri Cottet. 2001. Vortex methods: theory and practice. *Measurement Science and Technology* (2001).
- F. Du, Y. Zhang, Y. Ji, X. Wang, C. Yao, J. Kosinka, S. Frey, A. Telea, and X. Ban. 2025. GaussFluids: Reconstructing Lagrangian Fluid Particles from Videos via Gaussian Splatting. In *Pacific Graphics*.
- Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 2024. 4d-rotor gaussian splatting: towards efficient novel view synthesis for dynamic scenes. In *ACM SIGGRAPH 2024 Conference Papers*. 1–11.
- Marie-Lena Eckert, Kiwon Um, and Nils Thuerey. 2019. ScalarFlow: a large-scale volumetric data set of real-world scalar transport flows for computer animation and machine learning. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–16.
- Gerrit E. Elsinga, Fulvio Scarano, Bernhard Wieneke, and Bas W van Oudheusden. 2006. Tomographic particle image velocimetry. *Experiments in fluids* 41, 6 (2006), 933–947.
- Erik Franz, Barbara Solenthaler, and Nils Thuerey. 2021. Global Transport for Fluid Reconstruction with Learned Self-Supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1632–1642.
- Yue Gao, Hong-Xing Yu, Bo Zhu, and Jiajun Wu. 2025. FluidNexus: 3D fluid reconstruction and prediction from a single video. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 26091–26101.
- Ian Grant. 1997. Particle image velocimetry: a review. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 211, 1 (1997), 55–76.
- James Gregson, Ivo Ihrke, Nils Thuerey, and Wolfgang Heidrich. 2014. From capture to simulation: connecting forward and inverse problems in fluids. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Jinwei Gu, Shree K Nayar, Eitan Grinspun, Peter N Belhumeur, and Ravi Ramamoorthi. 2012. Compressive structured light for recovering inhomogeneous participating media. *IEEE transactions on pattern analysis and machine intelligence* 35, 3 (2012), 1–1.
- Haoqin Hong, Ding Fan, Fubin Dou, Zhi-Li Zhou, Haoran Sun, Congcong Zhu, and Jinrun Chen. 2025. Physics-Informed Deformable Gaussian Splatting: Towards Unified Constitutive Laws for Time-Evolving Material Field. *arXiv preprint arXiv:2511.06299* (2025).
- Yu Ji, Jinwei Ye, and Jingyi Yu. 2013. Reconstructing gas flows using light-path approximation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2507–2514.
- Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. 2015. The affine particle-in-cell method. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–10.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)* 42, 4 (2023), 1–14.
- Byungsoo Kim, Vinicius C Azevedo, Nils Thuerey, Theodore Kim, Markus Gross, and Barbara Solenthaler. 2019. Deep fluids: A generative network for parameterized fluid simulations. In *Computer graphics forum*, Vol. 38. Wiley Online Library, 59–70.
- Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. 2024. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *2024 International Conference on 3D Vision (3DV)*. IEEE, 800–809.
- Luan Lyu, Xiaohua Ren, Wei Cao, Jian Zhu, Enhua Wu, and Zhi-Xin Yang. 2024. Wavelet potentials: An efficient potential recovery technique for pointwise incompressible fluids. In *Computer Graphics Forum*, Vol. 43. Wiley Online Library, e15023.
- Miles Macklin and Matthias Müller. 2013. Position based fluids. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–12.
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4460–4470.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* 41, 4 (2022), 1–15.
- Mohammad Sina Nabizadeh, Ritoban Roy-Chowdhury, Hang Yin, Ravi Ramamoorthi, and Albert Chern. 2024. Fluid Implicit Particles on Coadjoint Orbits. *arXiv preprint arXiv:2406.01936* (2024).
- Xingyu Ni, Jingrui Xing, Xingqiao Li, Bin Wang, and Baoquan Chen. 2025. Representing Flow Fields with Divergence-Free Kernels for Reconstruction. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 8, 4 (2025), 1–21.
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 165–174.
- Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. 2020. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 523–540.
- Jiaxiang Qiu, Ruihong Cen, Zhong Li, Han Yan, Ming-Ming Cheng, and Bo Ren. 2024. NeuSmoke: Efficient Smoke Reconstruction and View Synthesis with Neural Transportation Fields. In *SIGGRAPH Asia 2024 Conference Papers*. 1–12.
- M. Raissi, P. Perdikaris, and G.E. Karniadakis. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378 (2019), 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- Pratik Rathore, Weimu Lei, Zachary Frangella, Lu Lu, and Madeleine Udell. 2024. Challenges in training PINNs: a loss landscape perspective. In *Proceedings of the 41st International Conference on Machine Learning*. 42159–42191.
- Jack Richter-Powell, Yaron Lipman, and Ricky TQ Chen. 2022. Neural conservation laws: A divergence-free perspective. *Advances in Neural Information Processing Systems* 35 (2022), 38075–38088.
- Nicolas Rosset, Guillaume Cordonnier, Regis Duvigneau, and Adrien Bousseau. 2023. Interactive design of 2D car profiles with aerodynamic feedback. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, 427–437.
- Ritoban Roy-Chowdhury, Tamar Shinar, and Craig Schroeder. 2024. Higher order divergence-free and curl-free interpolation on MAC grids. *J. Comput. Phys.* 503 (2024), 112831.
- Shunске Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. 2019. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2304–2314.
- Jos Stam. 2023. Stable fluids. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 779–786.
- Yiming Wang, Siyu Tang, and Mengyu Chu. 2024. Physics-Informed Learning of Characteristic Trajectories for Smoke Reconstruction. In *ACM SIGGRAPH 2024 Conference Papers*. 1–11.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- Holger Wendland. 1995. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in computational Mathematics* 4, 1 (1995), 389–396.
- Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 2024. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 20310–20320.
- Youchen Xie, Chen Li, Sheng Qiu, Zhi-Jun Wang, Chenhui Li, Yibo Zhao, Zan Gao, and Changbo Wang. 2025. FluidGS: Physics Informed Gaussian Splatting for Dynamic Fluid Reconstruction from Sparse Views. In *Proceedings of the 33rd ACM International Conference on Multimedia*. 8438–8447.
- Jinhui Xiong, Ramzi Idoughi, Andres A Aguirre-Pablo, Abdulrahman B Aljedaani, Xiong Dun, Qiang Fu, Sigurdur T Thoroddsen, and Wolfgang Heidrich. 2017. Rainbow particle imaging velocimetry for dense 3D fluid velocity imaging. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–14.
- Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. 2024. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 20331–20341.
- Hong-Xing Yu, Yang Zheng, Yuan Gao, Yitong Deng, Bo Zhu, and Jiajun Wu. 2023. Inferring Hybrid Neural Fluid Fields from Videos. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
- Yongning Zhu and Robert Bridson. 2005. Animating sand as a fluid. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 965–972.