

# PIGS-X: Physics-Informed Normalized Gaussian Splatting with Axes Alignment for Efficient Super-Resolution of 4D Flow MRI

Sun Jo<sup>1,5\*</sup>, Seok Young Hong<sup>4\*</sup>, Jinhyun Kim<sup>2,5</sup>, Seungmin Kang<sup>3,5</sup>, Ahjin Choi<sup>2,5</sup>,  
Don-Gwan An<sup>3,5</sup>, Simon Song<sup>3,5</sup> †, Je Hyeong Hong<sup>1,2,5</sup> †

<sup>1</sup>Department of Artificial Intelligence, Hanyang University, Seoul, Republic of Korea

<sup>2</sup>Department of Electronic Engineering, Hanyang University, Seoul, Republic of Korea

<sup>3</sup>Department of Mechanical Engineering, Hanyang University, Seoul, Republic of Korea

<sup>4</sup>School of Social Sciences and School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore

<sup>5</sup>Center for Precision Medicine Platform Based on Smart Hemo-Dynamic Index, Hanyang University

{choyw5, jhkim02, tmdals1213, caj0328, adg0324, simonsong, jhh37}@hanyang.ac.kr, seokyoung.hong@ntu.edu.sg

## Abstract

4D flow magnetic resonance imaging (MRI) is a reliable, non-invasive approach for estimating blood flow velocities, vital for cardiovascular diagnostics. Unlike conventional MRI focused on anatomical structures, 4D flow MRI requires high spatiotemporal resolution for early detection of critical conditions such as stenosis or aneurysms. However, achieving such resolution typically results in prolonged scan times, creating a trade-off between acquisition speed and prediction accuracy. Recent studies have leveraged physics-informed neural networks (PINNs) for super-resolution of MRI data, but their practical applicability is limited as the prohibitively slow training process must be performed for each patient. To overcome this limitation, we propose PINGS-X, a novel framework modeling high-resolution flow velocities using axes-aligned spatiotemporal Gaussian representations. Inspired by the effectiveness of 3D Gaussian splatting (3DGS) in novel view synthesis, PINGS-X extends this concept through several non-trivial novel innovations: (i) normalized Gaussian splatting with a formal convergence guarantee, (ii) axes-aligned Gaussians that simplify training for high-dimensional data while preserving accuracy and the convergence guarantee, and (iii) a Gaussian merging procedure to prevent degenerate solutions and boost computational efficiency. Experimental results on computational fluid dynamics (CFD) and real 4D flow MRI datasets demonstrate that PINGS-X substantially reduces training time while achieving superior super-resolution accuracy.

**Code** — <https://github.com/SpatialAILab/PIGS-X>

**Extended version** — <https://arxiv.org/abs/2511.11048>

## Introduction

4D flow MRI is a powerful, non-invasive method for quantifying blood flow velocities (Markl et al. 2012) without involving ionizing radiation. It captures time-resolved, three-dimensional velocity fields using phase-contrast MRI, which

\*These authors contributed equally.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Super-resolution of a carotid artery velocity field. From a low-resolution (LR) input, PINGS-X optimizes a set of 4D spatiotemporal Gaussians (visualized here in 3D at a single time frame). This explicit representation allows us to reconstruct a high-resolution (HR) output that faithfully recovers the complex flow patterns of the ground truth.

extends conventional MRI by encoding velocity-sensitive phase shifts into the MR signal. Crucially, the detailed hemodynamic insights from high-resolution 4D flow MRI are vital for predicting the progression of critical conditions such as stenosis and aneurysms (Ferdian et al. 2023).

Nevertheless, achieving high spatiotemporal resolution is challenging, as the multi-dimensional acquisition process involving four velocity encodings for each 3D volume and time frame is inherently slow (Markl et al. 2012). While modern acceleration techniques like parallel imaging (Deshmane et al. 2012) and compressed sensing (Lustig, Donoho, and Pauly 2007; Kontogiannis and Juniper 2022) are used, they have inherent limitations. Parallel imaging suffers from noise amplification at high acceleration rates, while compressed sensing relies on sparsity assumptions that fail in complex flow, often blurring structures (Gottwald et al. 2020) and underestimating peak velocities (Pathrose et al. 2021). This creates a clinical dilemma: accept rapid, low-resolution scans that risk obscuring information, or perform long, high-fidelity scans that are impractical due to patient discomfort and motion artifacts (Zaitsev, Maclaren, and Herbst 2015). This motivates a complementary strategy of post-hoc spatial super-resolution to enhance detail from

rapidly acquired data without increasing scan duration.

This potential has motivated various super-resolution strategies for 4D flow MRI (Ferdian et al. 2023; Saitta et al. 2024). Data-driven deep learning methods can learn effective low-to-high-resolution mappings (Shit et al. 2022; Ferdian et al. 2023) but require large patient datasets and often struggle with generalization to unseen cardiovascular domains (Ericsson et al. 2024). The alternative Physics-Informed Neural Network (PINN) paradigm (Raissi, Perdikaris, and Karniadakis 2019a) enforces physical laws on a per-scan basis, circumventing the need for large datasets. However, PINNs must be re-trained for each patient (Callmer et al. 2025), and their reliance on slow implicit neural representations creates a critical computational bottleneck for high-dimensional 4D flow MRI (Saitta et al. 2024; Callmer et al. 2025).

To address these challenges, we draw inspiration from novel view synthesis, where explicit representations have recently surpassed slower, implicit models. Notably, 3D Gaussian splatting (3DGS) (Kerbl et al. 2023) achieved state-of-the-art results with training times orders of magnitude faster than neural radiance fields (NeRFs) (Mildenhall et al. 2020). This motivates our key question: *Can a bespoke Gaussian-based representation overcome the computational bottlenecks of PINNs for 4D flow MRI super-resolution while retaining theoretical guarantee of convergence?*

In this paper, we address above question by introducing PINGS-X: Physics-Informed Normalized Gaussian Splatting with  $\alpha$ Xes alignment—a novel framework that adapts the principles of 3DGS to model 4D velocity fields using a set of optimized spatiotemporal Gaussians. As direct adoption of 3DGS is infeasible in physics-informed super-resolution, we propose several key innovations as follows:

- **Normalized Gaussian splatting with a formal convergence guarantee:** We propose a normalized Gaussian splatting scheme which enables faithful representation of continuous flow profiles. We provide a formal proof of convergence with a highly favorable convergence rate.
- **Axes-aligned Gaussians for efficient training:** We introduce axes-aligned Gaussians that simplify computation and improve training efficiency for high-dimensional data such as 4D flow MRI while preserving the convergence guarantee with minimal impact on accuracy.
- **Gaussian merging for stability and scalability:** To prevent potential degenerate solutions induced by the normalization scheme and enhance computational scalability, we present a density control strategy that merges nearby Gaussians with similar characteristics.

Each contribution is evaluated on synthetic CFD and real 4D flow MRI datasets with spatially averaged data, creating a more realistic challenging testbed than the sampling-based settings used in previous work (Dirix et al. 2022).

## Related Work

**Super-resolution for 4D Flow MRI.** Methods to enhance the resolution of 4D flow MRI fall into two main categories.

First, data-driven deep learning, uses architectures like CNNs to learn direct low-to-high-resolution mappings from

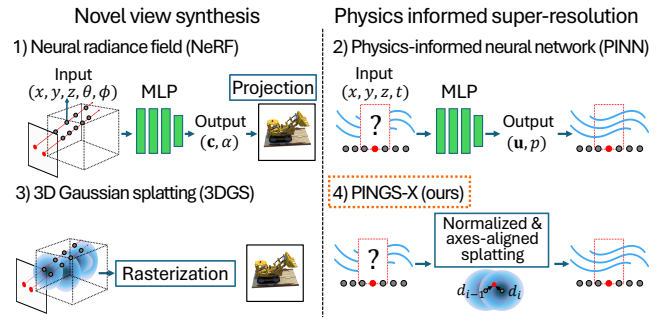


Figure 2: Illustration of the architectural parallel between neural rendering and physics-informed learning. Our work is motivated by the shared use of slow implicit representations (NeRFs and PINNs) and aims to transfer the efficiency of the explicit 3DGS (Kerbl et al. 2023) framework to our domain.

paired training data (Shit et al. 2022; Ferdian et al. 2023; Long et al. 2023). While powerful, their reliance on large and diverse training datasets is a significant practical challenge, leading to concerns about generalization to unseen cardiovascular domains (Ericsson et al. 2024).

Second, Physics-Informed Neural Networks (PINNs) offer a compelling alternative by embedding governing laws like the Navier-Stokes equations into the optimization process (Raissi, Perdikaris, and Karniadakis 2019a; Fathi et al. 2020). This allows them to be trained on a single 4D flow MRI scan without requiring large datasets (Fathi et al. 2020; Saitta et al. 2024). A known synergy exists with neural rendering; for example, the *Siren* activation function is now often used in PINNs to better represent complex physical fields. Despite this, the fundamental reliance of PINNs on a slow-to-train implicit MLP creates a major computational bottleneck for high-dimensional 4D flow MRI data, limiting their clinical feasibility (Saitta et al. 2024; Callmer et al. 2025; Cho et al. 2024).

**The shift toward explicit representations in rendering.** The computational bottleneck in PINNs mirrors a similar challenge recently overcome in novel view synthesis. Early methods like neural radiance fields (NeRFs) were also limited by slow, implicit MLP-based models (Mildenhall et al. 2020). This limitation was largely solved by 3D Gaussian splatting, which achieved state-of-the-art results with orders-of-magnitude faster training by using an explicit set of 3D Gaussians (Kerbl et al. 2023). This successful transition from a slow implicit to a fast explicit representation provides the primary motivation for our work: to bring similar efficiency gains to physics-informed super-resolution.

**Physics-informed Gaussian representations.** A few recent studies have explored the intersection of physics and Gaussian representations, but each has key limitations that our work aims to address. Physics-Informed Gaussian Splatting (PIGS) (Max Rensen and Eisemann 2024) models time-resolved physical fields by moving 3D Gaussians through the time domain. However, its use of unnormalized Gaussian weights means it lacks theoretical convergence guarantees,

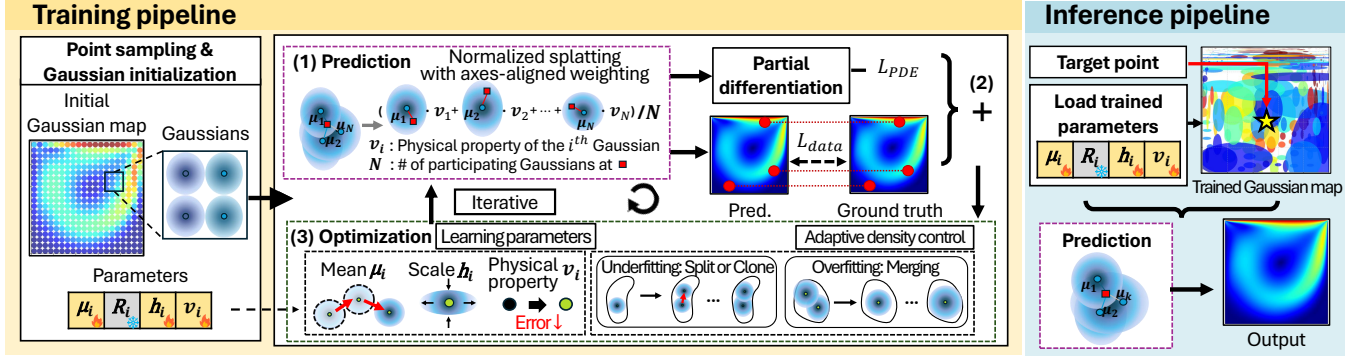


Figure 3: Our PINGS-X framework. For training, we initialize (axes-aligned) spatiotemporal Gaussians from low-resolution data and iteratively optimize them. In each step, we (1) predict values using Normalized Gaussian Splatting (NGS), (2) compute a combined data ( $L_{data}$ ) and physics ( $L_{PDE}$ ) loss, and (3) update the Gaussians parameters. The Gaussian density is adjusted via splitting, cloning, and our proposed *merging*. At inference, we use the trained Gaussians for prediction via NGS.

which, as our experiments will show, can lead to inferior prediction accuracy. In contrast, Physics-Informed Gaussians (PIG) (Kang et al. 2024) uses a lightweight MLP to process a Gaussian feature embedding. However, by retaining a neural network in its prediction path, it remains an implicit method and, as we found in our experiments, fundamentally shares the same computational overhead as standard PINNs.

**Summary and positioning.** A need remains for a theoretically grounded Gaussian splatting approach for efficient physical-field super-resolution. Our work addresses this gap.

## Preliminaries

### 4D flow MRI and governing equations

**Acquisition principle.** 4D flow MRI is an extension of phase-contrast MRI that encodes velocity-sensitive phase shifts into the MR signal along three orthogonal directions. For each cardiac phase, the scanner acquires four complex k-space volumes—one reference and three flow-encoded scans. Typical exams capture 20–30 cardiac frames at 1–3 mm isotropic spatial resolution. After inverse Fourier reconstruction, phase unwrapping, and velocity scaling, these volumes provide a time-resolved velocity field  $\mathbf{v}(x, y, z, t) = [v_x, v_y, v_z]^\top$  defined on the voxel lattice.

**Governing partial differential equations (PDEs).** Blood flow in large arteries and veins is often approximated as an incompressible Newtonian fluid (Lynch, Nama, and Figueroa 2022). Subsequently, the velocity field obeys the incompressible Navier–Stokes equation in Eq. (1), which describe the motion of a viscous fluid, and the continuity equation, which describes mass conservation in Eq. (2).

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{v} + \mathbf{g}, \quad (1)$$

$$\nabla \cdot \mathbf{v} = 0, \quad (2)$$

where  $p(x, y, z, t)$  is pressure,  $\rho$  the fluid density,  $\nu$  the kinematic viscosity, and  $\mathbf{g}$  is the gravity. Eq. (1) and Eq. (2) supply the physics-informed loss terms later used in training.

**Spatial averaging.** Each voxel in the acquired velocity field represents a spatial average of all sub-voxel velocities, leading to a loss of fine-scale details. This motivates the need for super-resolution methods that recover fine-grained velocities consistent with Eq. (1) and Eq. (2).

### 3D Gaussian splatting in novel view synthesis

As shown in Fig. 2, 3DGS represents a scene with a sparse  $N$  set of anisotropic 3D Gaussians  $\{\mathcal{G}_i\}$ . Each Gaussian has center (mean)  $\boldsymbol{\mu}_i \in \mathbb{R}^3$ , covariance  $\Sigma_i = \mathbf{R}_i \text{diag}(h_{i1}^2, h_{i2}^2, h_{i3}^2) \mathbf{R}_i^\top$ , color  $\mathbf{c}_i \in \mathbb{R}^3$ , and opacity  $0 \leq \alpha_i \leq 1$ , where  $\{h_{ij}\}$  represent the scales of  $\mathcal{G}_i$  along its principal axes and  $\mathbf{R}_i \in SO(3)$  represent its rotation.

**Rasterization.** After perspective projection, the  $i$ -th Gaussian appears on the image plane with 2D mean  $\boldsymbol{\mu}'_i$  and covariance  $\Sigma'_i$ . For a pixel at screen coordinate  $\mathbf{x} \in \mathbb{R}^2$ , Gaussians are sorted front-to-back and blended by closed-form alpha-compositing as follows:

$$\tilde{\alpha}_i(\mathbf{x}) = \alpha_i \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}'_i)^\top \Sigma'_i (\mathbf{x} - \boldsymbol{\mu}'_i)\right), \quad (3)$$

$$\hat{\mathbf{c}}(\mathbf{x}) = \sum_{i=1}^N \mathbf{c}_i \tilde{\alpha}_i(\mathbf{x}) \prod_{j < i} (1 - \tilde{\alpha}_j(\mathbf{x})), \quad (4)$$

where  $\hat{\mathbf{c}} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$  is the rendered RGB color at pixel  $\mathbf{x}$ .

**Optimization pipeline.** Training alternates between two stages, namely (i) joint gradient-based optimization of Gaussian variables  $\{\boldsymbol{\mu}_i, \Sigma_i, \mathbf{c}_i, \alpha_i\}$ , and (ii) adaptive density control to *split/clone* Gaussians in high-error regions and *prune* low-contribution Gaussians. This yields NeRF-level quality with training times orders-of-magnitude shorter.

## Proposed Method

We introduce PINGS-X, a framework for physics-informed super-resolution that leverages the explicit representation of 3D Gaussian Splatting. As illustrated in Fig. 3, our approach is built upon three key innovations detailed in this section: Normalized Gaussian Splatting (NGS), the use of axes-aligned Gaussians, and a Gaussian merging procedure.

## Normalized Gaussian splatting and its convergence

**Motivation.** A straightforward application of 3DGS to model a physical field (e.g. velocity and pressure)  $\hat{\mathbf{v}} : \mathbb{R}^q \rightarrow \mathbb{R}^p$  at any spatiotemporal point  $\mathbf{x} \in \mathbb{R}^q$  is to use an (unnormalized) weighted sum of  $N$  Gaussians  $\{\mathcal{G}_i\}$  each with physical property  $\mathbf{v}_i \in \mathbb{R}^p$ , center  $\boldsymbol{\mu}_i \in \mathbb{R}^q$  and covariance  $\Sigma_i := \mathbf{R}_i \text{diag}(h_{i1}^2, \dots, h_{iq}^2) \mathbf{R}_i^\top$  ( $\mathbf{R}_i \in SO(q)$  is the rotation and  $h_{ij}$  is the Gaussian scale about the  $j$ -th axis) as follows:

$$\hat{\mathbf{v}}(\mathbf{x}) = \sum_{i=1}^N z_i(\mathbf{x}) \mathbf{v}_i, \quad (5)$$

where  $z_i(\mathbf{x}) = \exp(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^\top \Sigma_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i))$  can be viewed as the (unnormalized) influence of  $\mathcal{G}_i$ . This formulation similarly used in PIGS (Max Rensen and Eisemann 2024) is simpler than the original 3DGS rendering equation as it is a direct evaluation in the spatiotemporal domain and thus omits the projection and alpha compositing steps. However, the unnormalized sum above has a fundamental limitation: it is not a *well-behaved* prediction scheme.

In regions far from any Gaussian center, the prediction  $\hat{\mathbf{v}}(\mathbf{x})$  collapses to zero (Fig. 4). To compensate, the model must adopt inefficient strategies, such as learning overly large Gaussians that cause over-smoothing, or constantly adding new Gaussians which can create oscillatory behavior. This instability is further evidenced in our Rosenbrock experiment (Table 1 and Fig. 5), and we also provide a theory in the extended version confirming that convergence of  $\hat{\mathbf{v}}$  is not guaranteed, as a distant  $\mathbf{x}$  makes the prediction collapse to zero unless the number of Gaussians is adjusted.

**Our solution.** To address this issue, we introduce Normalized Gaussian splatting (NGS), which redefines the prediction by normalizing the Gaussian weights:

$$\hat{\mathbf{v}}(\mathbf{x}) = \sum_{i=1}^N w_i(\mathbf{x}) \mathbf{v}_i, \quad \text{where } w_i(\mathbf{x}) = \frac{z_i(\mathbf{x})}{\sum_{j=1}^N z_j(\mathbf{x})}. \quad (6)$$

This ensures the prediction is a convex combination of the Gaussian properties, providing inherent stability as the predicted value is always bounded by the minimum and maximum of the values from the influencing Gaussians. Consequently, adding new Gaussians serves as a principled refinement—placing new “control points” to capture finer details instead of a corrective measure (Fig. 4). Such stability is the basis for our method’s universal approximation capability, which we formalize in the following theorem.

**Universal approximation property: consistency.** We formally establish the asymptotic convergence in a statistical sense. The convergence rate is explicitly derived in Theorem 1. We provide the proof in the extended version.

**Theorem 1** *Suppose  $(\boldsymbol{\mu}_i, \mathbf{v}_i) \in \mathbb{R}^{q+p}$ ,  $i = 1, \dots, N$ ,  $q, p \in \mathbb{N}$ , are samples from the random vector  $(\boldsymbol{\mu}, \mathbf{v})$ , whose cumulative distribution function is absolutely continuous with respect to the Lebesgue measure with density  $f$ . We assume:*

- (i) (i.i.d) *The samples  $(\boldsymbol{\mu}_1, \mathbf{v}_1), (\boldsymbol{\mu}_2, \mathbf{v}_2), \dots, (\boldsymbol{\mu}_N, \mathbf{v}_N)$  are independent and identically distributed.*




Figure 4: A 1D visualization comparing unnormalized and normalized sums of Gaussians ( $\sigma = 12$ ). (a) Unnormalized sum: with only two Gaussians ( $x = \{5, 25\}$ ,  $v = \{1, 3\}$ ), the prediction (dotted red) decays to zero in the middle. Adding a third Gaussian at  $x = 15$ ,  $v = 1.3$  (orange) to fill this gap results in an oscillatory profile (solid blue). (b) Normalized sum: we observe a smooth prediction curve between the two Gaussians. Adding the third Gaussian adjusts the curve while maintaining a stable, monotonic result.

- (ii) (*Smoothness condition*) *Both  $\mathbf{v}$  and  $f$  are  $\beta$ -times continuously differentiable in some neighborhood of  $\mathbf{x} \in \mathbb{R}^q$ .*  
(iii) (*Bandwidth conditions*) *The matrix  $\Sigma_i = \Sigma_i(N)$  is a sequence of bandwidth matrices such that  $[\text{tr}(\Sigma_i)]^{\beta/2} \rightarrow 0$  and  $N \cdot \det(\Sigma_i)^{1/2} \rightarrow \infty$  as  $N \rightarrow \infty$  for all  $i$ .*

Then, for any  $\mathbf{x} \in \mathbb{R}^q$ , we have the following convergence:

$$\hat{\mathbf{v}}(\mathbf{x}) - \mathbf{v}(\mathbf{x}) \rightarrow 0 \quad (7)$$

in probability, as the number of Gaussians  $N \rightarrow \infty$ , where  $\mathbf{v}(\mathbf{x}) = \mathbb{E}[\mathbf{v} | \boldsymbol{\mu} = \mathbf{x}]$ . Specifically, we have

$$\hat{\mathbf{v}}(\mathbf{x}) - \mathbf{v}(\mathbf{x}) = O_p \left( \frac{1}{N} \sum_{i=1}^N [\text{tr}(\Sigma_i)]^{\beta/2} + \sqrt{\frac{1}{N^2} \sum_{i=1}^N \frac{1}{\det(\Sigma_i)^{1/2}}} \right) \quad (8)$$

where  $O_p(\cdot)$  denotes the stochastic Big-O notation.

**Remark.** NGS can be viewed as transforming the classic (static) Nadaraya-Watson (NW) estimator (Nadaraya 1964; Watson 1964) into a dynamic, learnable framework. Unlike the static nature of NW, which uses fixed points and a global bandwidth, our approach optimizes the Gaussian centers ( $\boldsymbol{\mu}_i$ ) and learns an individual anisotropic covariance/bandwidth ( $\Sigma_i$ ) for each kernel. Coupled with adaptive density control to add or remove Gaussians, this creates a notably more flexible representation for complex physical fields.

## Axes-aligned splatting for efficient super-resolution

**Motivation.** Optimizing the full anisotropic covariance of NGS becomes prohibitive in high-dimensional domains like 4D spatiotemporal space. The primary challenge is parameterizing the higher-dimensional (4D) rotation of each Gaussian, which is non-trivial to optimize with gradient descent.

**Our solution.** To circumvent this representational complexity, we constrain the Gaussians to be axes-aligned, meaning their covariance matrices are diagonal, i.e.  $\Sigma_i = \text{diag}(h_{i1}^2, \dots, h_{iq}^2)$ . This provides a simpler and more direct representation that avoids the need to parameterize complex 4D rotations. As a secondary benefit, this also reduces

# Gauss.	Method	$t$ (s) ↓	Rel. err. (%) ↓	RMSE ↓
16	PIGS	61.1	3.38	7.47e-03
	PIGS (ours)	66.5	<b>1.00</b>	<b>2.22e-03</b>
	PIGS-X (ours)	<b>48.4</b>	3.29	7.28e-03
100	PIGS	58.2	2.45	5.42e-03
	PIGS (ours)	68.7	<b>0.10</b>	<b>2.17e-04</b>
	PIGS-X (ours)	<b>50.1</b>	0.14	3.01e-04
400	PIGS	68.2	3.14	6.94e-03
	PIGS (ours)	74.5	0.26	5.83e-04
	PIGS-X (ours)	<b>54.2</b>	<b>0.09</b>	<b>1.92e-04</b>
600	PIGS	72.6	3.53	7.82e-03
	PIGS (ours)	80.1	<b>0.04</b>	<b>7.95e-05</b>
	PIGS-X (ours)	<b>62.9</b>	0.08	1.73e-04

Table 1: Results for fitting the Rosenbrock function  $f(x, y) = (1 - x)^2 + 100(x - y^2)^2$ , analyzing only the effect of the number of Gaussians ( $N$ ) without position updates and adaptive density control of Gaussians. The error for the unnormalized sum (PIGS) remains high and unstable, while our normalized methods (PIGS/PIGS-X) demonstrate clear convergence as  $N$  increases.

the number of covariance parameters from 10 to 4 for each Gaussian, improving computational efficiency.

This approach, which we term Normalized and Axes-aligned Gaussian Splatting (NGS-X), preserves the favorable convergence properties of our general method—a claim we validate theoretically below and empirically in Fig. 5.

**Corollary 1** *In addition to the same assumptions used in Theorem 1, if we further suppose  $\Sigma_i$  is diagonal with entries  $h_{i1}^2, h_{i2}^2, \dots, h_{iq}^2$  for each  $i$ , the expression (8) simplifies to*

$$\hat{\mathbf{v}}(\mathbf{x}) - \mathbf{v}(\mathbf{x}) = O_p \left( \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^q h_{ij}^\beta + \sqrt{\frac{1}{N^2} \sum_{i=1}^N \frac{1}{h_{i1} \dots h_{iq}}} \right). \quad (9)$$

**Remark.** In the case where each bandwidth  $h_{ij}$  is of same order of magnitude, i.e.  $h_{ij} = O(h) \forall i, j$ , using the bias-variance balancing bandwidth, we attain the best possible convergence rate in the minimax sense (Stone 1980, 1982).

### Gaussian merging for scalable training

**Motivation.** In 3DGS, the training loss is computed on 2D renderings from multiple camera poses, meaning overlapping Gaussians are implicitly disambiguated by multi-view geometric constraints. Our framework is different, as we directly approximate a high-dimensional physical field without projection or alpha-compositing. With distance-based kernels, this can lead to *degenerate solutions*, where multiple Gaussians accumulate in the same region to represent a single smooth profile. Such duplicates inflate memory and slow optimization without improving accuracy.

Note, pruning is ineffective in our scenario as these redundant Gaussians can be equally impactful. Hence, the appropriate strategy is to *merge* them into a single representative.

**Our solution.** Our merging procedure identifies Gaussians with similar influence on the prediction. For each Gaussian  $\mathcal{G}_i$ , we first compute its normalized influence vector  $\mathbf{w}_i = [w_i(\mathbf{x}_1), \dots, w_i(\mathbf{x}_K)]^\top$  at all  $K$  training points.




Figure 5: Absolute error maps for the Rosenbrock function (yellow indicates higher error). Our PINGS/PIGS-X converges as the number of Gaussians (in parentheses) increases, while (unnormalized) PIGS does not.

We then build an undirected graph where nodes are Gaussians and an edge connects any pair  $(i, j)$  whose influence vectors have a cosine similarity exceeding a threshold (0.9).

The connected components of this graph identify clusters of redundant Gaussians, which are then merged in a single pass. The new Gaussian’s mean  $\boldsymbol{\mu}$  is the average of the merged means, and its physical property is the predicted field value at that new mean,  $\hat{\mathbf{v}}(\boldsymbol{\mu})$ . This merge step is invoked periodically during training (e.g., every 100 epochs), improving training stability and efficiency (Table 4).

### Training procedure

**Initialization of Gaussians.** We initialize the Gaussians  $\{\mathcal{G}_i\}$  using a simple grid-based sampling strategy to ensure uniform coverage of the spatiotemporal domain. Each physical property  $\mathbf{v}_i \in \mathbb{R}^p$  is set to the value of the nearest low-resolution data point to the initial mean  $\boldsymbol{\mu}_i$ . This straightforward initialization was chosen to keep the setup simple and avoid dependence on complex, data-specific heuristics.

**Loss function.** We train our model with a composite loss function  $L$  that combines a data fidelity term,  $L_{data}$ , with a physics-informed regularizer,  $L_{PDE}$  to yield PINGS-X:

$$L = L_{data} + \lambda L_{PDE}, \quad (10)$$

$$L_{data} = \frac{1}{K} \sum_{k=1}^K \|\boldsymbol{\omega} \odot (\hat{\mathbf{v}}(\mathbf{x}_k) - \mathbf{v}_k)\|^2 \quad (11)$$

$$L_{PDE} = \frac{1}{K} \sum_{k=1}^K \sum_{m=1}^M \|\mathbf{g}_m(\hat{\mathbf{v}}(\mathbf{x}_k))\|^2, \quad (12)$$

where  $L_{data}$  is essentially the mean squared error between the prediction  $\hat{\mathbf{v}}$  and the low-resolution observation  $\mathbf{v}_k$ . The binary vector  $\boldsymbol{\omega} \in \mathbb{R}^p$  masks unobserved components of the physical properties; for 4D flow MRI, we use it to exclude the pressure term which is not directly measured.  $L_{PDE}$  is the mean squared error of  $M$  physical constraints,  $\{\mathbf{g}_m\}$ . For 4D flow MRI, these are based on the dimensionless form of Eq. (1) and Eq. (2). from the dimensionless Navier-Stokes equations (see extended version) for which we set  $\lambda = 1.0$ .

Note, unlike PINNs, which compute derivatives for  $L_{PDE}$  via computationally expensive backpropagation, our

Methods	Lid-driven		Y-shaped		L-shaped	
	$t \downarrow$ (min)	Rel. $\downarrow$ err.(%)	$t \downarrow$ (min)	Rel. $\downarrow$ err.(%)	$t \downarrow$ (min)	Rel. $\downarrow$ err.(%)
Nada.-Watson	<1	3.03	<1	5.11	<1	3.99
PINN	51.4	12.20	45.0	11.54	43.6	11.84
Siren	59.8	3.20	54.0	4.34	50.3	2.49
XPINN (tanh)	90.2	8.47	165.0	8.07	132.6	5.54
XPINN (sin)	105.8	2.71	195.5	3.89	158.0	2.15
PIG	820.1	2.33	292.9	7.06	297.8	1.48
PIGS-X	<b>21.9</b>	<b>1.13</b>	<b>4.8</b>	<b>2.62</b>	<b>6.3</b>	<b>1.15</b>

Table 2: Quantitative comparison on synthetic 2D CFD datasets. PINGS-X consistently outperforms other methods.

# Initial Gaussians	Lid-driven		Y-shaped		L-shaped	
	$t \downarrow$ (min)	Rel. $\downarrow$ err.(%)	$t \downarrow$ (min)	Rel. $\downarrow$ err.(%)	$t \downarrow$ (min)	Rel. $\downarrow$ err.(%)
200	35.5	1.13	5.6	2.60	6.7	1.16
400	22.2	1.13	5.2	2.62	6.7	1.15
800	9.4	1.36	5.0	2.74	6.5	1.12
1200	15.7	1.15	4.7	2.51	7.0	1.17
1600	10.6	1.33	4.8	2.59	5.1	1.12

Table 3: Sensitivity analysis for the number of initial Gaussians. The performance of PINGS-X is robust across a wide range of initial counts, highlighting the effectiveness of its adaptive density control (split/clone/merge) during training.

explicit formulation allows for their efficient, analytical calculation. This avoids the compounding cost of automatic differentiation, making our training process significantly faster.

**Boundary conditions.** Boundary conditions are implicitly enforced by sampling training data on boundaries.

## Experimental Results and Discussions

This section experimentally evaluates PINGS-X, detailing the setup, presenting results on synthetic 2D CFD and real 4D flow MRI datasets, and including an ablation study.

### Datasets

We evaluate our method on both synthetic CFD and real 4D flow MRI datasets. For both cases, low-resolution data was generated via spatial averaging of the original high-resolution fields. Further details on the datasets and the spatial averaging process are in the extended version.

**Synthetic CFD datasets.** We used three steady-state, incompressible 2D flow simulations: a *lid-driven cavity*, *L-shaped* and *Y-shaped* channels with the last mimicking vascular geometries. High-resolution ground truth velocity and pressure fields were generated using standard CFD solvers (OpenFOAM and StarCCM+) with appropriate boundary conditions. Low-resolution data was created by spatially averaging these fields, providing a more challenging setup than simple point-sampling. For these cases, the model maps 2D coordinates ( $q = 2$ ) to a 3D physical property vector containing velocity and pressure  $(v_x, v_y, p)$ , hence  $p = 3$ .

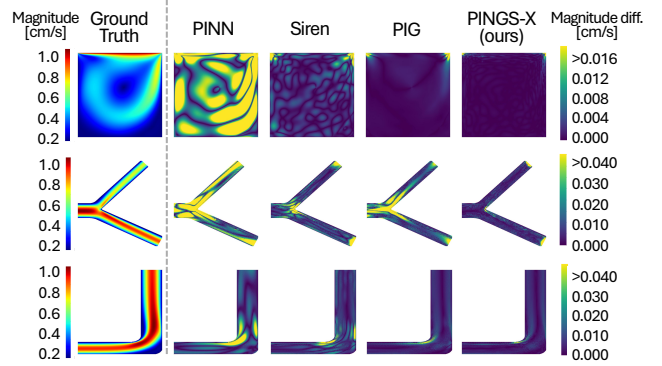


Figure 6: Qualitative comparison of ground truth velocity and absolute error maps for the synthetic 2D CFD datasets (Lid-driven, Y-shape, and L-shape). For the error maps (yellow = high error), the color scale is clipped to the mean error of the PINN baseline to aid visual comparison.

Method	Lid-driven		Y-shaped		L-shaped	
	$t \downarrow$ (min)	Rel. $\downarrow$ err. (%)	$t \downarrow$ (min)	Rel. $\downarrow$ err. (%)	$t \downarrow$ (hr)	Rel. $\downarrow$ err. (%)
PIGS-X	21.9	1.13	4.8	2.62	6.3	1.15
w/o norm.	23.3	NC	3.4	NC	4.3	NC
w/o axes-align.	70.3	1.01	23.3	2.27	22.7	1.13
w/o merging	*12.1	*1.38	*35.1	*1.66	*32.8	*1.12

Table 4: Ablation study of PINGS-X components. NC denotes non-convergence. \* indicates experiments terminated early ( $< 2000$  epochs) due to out-of-memory (OOM) issues.

**Real 4D flow MRI dataset.** We used a high-resolution 4D Flow MRI dataset of a carotid artery phantom that replicates patient vasculature (Ko et al. 2019), acquired at 0.35 mm spatial and 25 ms temporal resolution. From 34 time frames per cardiac cycle, five frames around the peak systolic phase were used for constructing the 4D dataset. From this data, low-resolution inputs for our  $\times 8$  and  $\times 64$  super-resolution tasks were generated via spatial averaging. For this task, the model maps 4D spatiotemporal coordinates ( $q = 4$ ) to a 4D physical property vector containing the 3D velocity and pressure  $(v_x, v_y, v_z, p)$ , hence  $p = 4$ .

### Experimental settings

**Baselines and general configuration.** For the task of physics-informed super-resolution, we compared our method, PINGS-X, against several representative baselines: a standard PINN (Raissi, Perdikaris, and Karniadakis 2019b), PINN with a Siren activation function (Sitzmann et al. 2020), XPINN (Jagtap and Karniadakis 2021) (using both *tanh* and *Siren* activations), and the recent Physics-Informed Gaussians (PIG) (Kang et al. 2024), all of which were run for 100,000 epochs. To evaluate against its theoretical foundation, we also included the classic Nadaraya-Watson (NW) estimator as a non-learnable reference. All models were trained using an Adam optimizer. More details are in the extended version.

Method	Carotid ( $\times 8 = \times 2^3$ )			Carotid ( $\times 64 = \times 4^3$ )		
	$t \downarrow$ (hr)	Rel. $\downarrow$ err. (%)	RMSE $\downarrow$ (cm/s)	$t \downarrow$ (min)	Rel. $\downarrow$ err. (%)	RMSE $\downarrow$ (cm/s)
PINN	30.1	25.93	2.75	259.0	39.70	4.21
Siren	30.8	10.63	1.13	299.2	18.49	1.96
PIGS-X	<b>2.6</b>	<b>8.98</b>	<b>0.95</b>	<b>3.9</b>	<b>17.59</b>	<b>1.87</b>

Table 5: Quantitative comparison on the 4D flow MRI (carotid) dataset for two spatial averaging settings ( $\times 8$  and  $\times 64$ ). PINGS-X shows superior performance in both training time and prediction error.

**Implementation details.** For synthetic CFD experiments, PINN, Siren, and each XPINN sub-network shared an architecture of 5 hidden layers with 128 neurons and were trained using a full-batch approach. Models with Tanh activation (PINN, XPINN (tanh)) used a learning rate of  $10^{-4}$  (Wang et al. 2023), while those with Siren activation (Siren, XPINN (sin)) used  $5 \times 10^{-6}$  and  $\omega_0=10$ . For PIG, we used the official implementation with default settings, while PINGS-X was initialized with 400 Gaussians, trained for 10,000 epochs with a learning rate of  $10^{-2}$ , and applied splitting/cloning and merging every 100 epochs.

For the real 4D flow MRI dataset, we omitted PIG and XPINN due to high training time and NW as it showed much inferior performance. PINN and Siren baselines utilized larger MLP with 8 hidden layers and 256 neurons, trained with a batch size of 10,000. The learning rates were identical to the synthetic case, with Siren using  $\omega_0=20$  for improved stability. For PINGS-X framework, Gaussians were initialized by downsampling input points to 1/10, yielding 4,863 Gaussians for the  $\times 8$  task and 753 for  $\times 64$ . Training was conducted for 1,000 epochs with a learning rate of  $10^{-2}$  and batch size 10,000, applying split/clone and merge every 100 epochs. More details are in the extended version.

**Evaluation metrics** We evaluate models based on wall-clock training time ( $t$ ) and two error metrics computed over the  $K$  high-resolution test points. Our primary metric is the relative  $L^2$  error, defined as  $\sqrt{\sum_{k=1}^K \|\hat{\mathbf{v}}(\mathbf{x}_k) - \mathbf{v}_k\|_2^2 / \sum_{k=1}^K \|\mathbf{v}_k\|_2^2}$ . For the 4D flow MRI experiments, we also report the root mean squared error (RMSE),  $\sqrt{\frac{1}{K} \sum_{k=1}^K \|\hat{\mathbf{v}}(\mathbf{x}_k) - \mathbf{v}_k\|_2^2}$ , in cm/s. In both metrics,  $\hat{\mathbf{v}}(\mathbf{x}_k)$  is the prediction at  $\mathbf{x}_k$  and  $\mathbf{v}_k$  is the corresponding ground truth. For these calculations, we exclude the unknown pressure component so the metrics directly reflect the error in the predicted velocity.

## Comparative analysis

**Results on synthetic CFD datasets.** As shown in Table 2, PINGS-X consistently achieves the lowest relative  $L^2$  error across all synthetic datasets. Notably, this superior accuracy is achieved with a dramatic reduction in training time. For instance, on the Y-shaped dataset, PINGS-X converges in just 4.8 minutes, making it approximately 40 times faster than the next most accurate model (XPINN (sin)). Fig. 6 also qualitatively demonstrates that PING-X achieved the lowest




Figure 7: Visualizations of ground truth velocity field averaged over five time frames and prediction error maps (magnitude and direction) achieved by different methods on the 4D flow MRI dataset for the  $\times 8$  spatially averaged data.

error among all the baseline methods.

**Results on real 4D Flow MRI dataset.** On the real-world 4D flow MRI dataset, PINGS-X again demonstrates superior performance (Table 5). For both the  $\times 8$  and the more challenging  $\times 64$  super-resolution tasks, our method converged to a solution with lower error significantly faster than the PINN and Siren baselines. The qualitative results in Fig. 7 corroborate this, showing error maps for both velocity magnitude and angle that are visibly closer to the ground truth.

## Ablation study

Our ablation studies (Table 4) validate our key design choices. Removing normalization prevents the model from converging, confirming its necessity for stability. Disabling axes-alignment significantly increases training time for negligible accuracy gain, validating its efficiency benefits. Finally, removing Gaussian merging generally increases both training time and error. The effectiveness of our complete adaptive density control (with merging) is further highlighted in Table 3, which shows that the full PINGS-X framework is robust to the initial number of Gaussians.

## Conclusion

In this work, we introduced *PIGS-X*, a framework for the physics-informed super-resolution of 4D flow MRI data, designed to address the computational expense of PINN-based methods. Inspired by successes in computer graphics, we developed an explicit, spatiotemporal Gaussian representation with three key innovations: a Normalized Gaussian splatting scheme with theoretical convergence guarantees, an axes-aligned representation for efficient training in (high-dimensional) spatiotemporal space, and a Gaussian merging procedure for robust, scalable training. Experiments on synthetic and real-world spatially averaged data show that PINGS-X achieves high accuracy with significantly reduced training time compared to competing methods. By bridging the gap between explicit representations and physics-informed learning, PINGS-X offers a promising direction for efficient and accurate scientific modeling.

## Acknowledgments

This work was partly supported by National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (No.2021R1A2B5B03002103 and No.RS-2022-NR070832), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2020-II201373, Artificial Intelligence Graduate School Program(Hanyang University)), Hanyang University (HY-202500000003991) and the CoHASS Research Support Grant (No.025637-00001), funded by Nanyang Technological University, Singapore.

## References

- Callmer, P.; Bonini, M.; Ferdian, E.; Nordsletten, D.; Giese, D.; Young, A. A.; and Marlevi, D. 2025. Deep learning for temporal super-resolution 4D Flow MRI. *arXiv preprint arXiv:2501.08780*.
- Cho, J.; Nam, S.; Yang, H.; Yun, S.-B.; Hong, Y.; and Park, E. 2024. Separable Physics-Informed Neural Networks. *Advances in Neural Information Processing Systems*, 36: 56213–56236.
- Deshmane, A.; Gulani, V.; Griswold, M. A.; and Seiberlich, N. 2012. Parallel MR imaging. *Journal of Magnetic Resonance Imaging*, 36(1): 55–72.
- Dirix, P.; Buoso, S.; Peper, E. S.; and Kozerke, S. 2022. Synthesis of patient-specific multipoint 4D flow MRI data of turbulent aortic flow downstream of stenotic valves. *Scientific Reports*, 12(1): 16004.
- Ericsson, L.; Hjalmarsson, A.; Akbar, M. U.; Ferdian, E.; Bonini, M.; Hardy, B.; Schollenberger, J.; Aristova, M.; Winter, P.; Burris, N.; et al. 2024. Generalized Super-Resolution 4D Flow MRI—Using Ensemble Learning to Extend Across the Cardiovascular System. *IEEE Journal of Biomedical and Health Informatics*.
- Fathi, M. F.; Perez-Raya, I.; Baghaie, A.; Berg, P.; Janiga, G.; Arzani, A.; and D’Souza, R. M. 2020. Super-resolution and denoising of 4D-flow MRI using physics-informed deep neural nets. *Computer Methods and Programs in Biomedicine*, 197: 105729.
- Ferdian, E.; Marlevi, D.; Schollenberger, J.; Aristova, M.; Edelman, E. R.; Schnell, S.; Figueroa, C. A.; Nordsletten, D.; and Young, A. A. 2023. Cerebrovascular super-resolution 4D flow MRI—sequential combination of resolution enhancement by deep learning and physics-informed image processing to non-invasively quantify intracranial velocity, flow, and relative pressure. *Medical Image Analysis*, 88: 102831.
- Gottwald, L. M.; Peper, E. S.; Zhang, Q.; Coolen, B. F.; Strijkers, G. J.; Nederveen, A. J.; and van Ooij, P. 2020. Pseudospiral sampling and compressed sensing reconstruction provides flexibility of temporal resolution in accelerated aortic 4D flow MRI: A comparison with k-t principal component analysis. *NMR in Biomedicine*, 33(4): e4255.
- Jagtap, A. D.; and Karniadakis, G. E. 2021. Extended Physics-informed Neural Networks (XPINNs): A Generalized Space-Time Domain Decomposition based Deep Learning Framework for Nonlinear Partial Differential Equations. In *AAAI Spring Symposium: MLPS*, volume 10.
- Kang, N.; Oh, J.; Hong, Y.; and Park, E. 2024. PIG: Physics-Informed Gaussians as Adaptive Parametric Mesh Representations. *arXiv preprint arXiv:2412.05994*.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Ko, S.; Lee, J.; Song, S.; Kim, D.; Lee, S. H.; and Cho, J.-H. 2019. Patient-specific hemodynamics of severe carotid artery stenosis before and after endarterectomy examined by 4D flow MRI. *Scientific Reports*, 9(1): 18554.
- Kontogiannis, A.; and Juniper, M. P. 2022. Physics-informed compressed sensing for PC-MRI: an inverse Navier-Stokes problem. *IEEE Transactions on Image Processing*, 32: 281–294.
- Long, D.; McMurdo, C.; Ferdian, E.; Mauger, C. A.; Marlevi, D.; Nash, M. P.; and Young, A. A. 2023. Super-resolution 4D flow MRI to quantify aortic regurgitation using computational fluid dynamics and deep learning. *The International Journal of Cardiovascular Imaging*, 39(6): 1189–1202.
- Lustig, M.; Donoho, D. L.; and Pauly, J. M. 2007. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6): 1182–1195.
- Lynch, S.; Nama, N.; and Figueroa, C. A. 2022. Effects of non-Newtonian viscosity on arterial and venous flow and transport. *Scientific Reports*, 12(1): 20568.
- Markl, M.; Frydrychowicz, A.; Kozerke, S.; Hope, M.; and Wieben, O. 2012. 4D flow MRI. *Journal of Magnetic Resonance Imaging*, 36(5): 1015–1036.
- Max Rensen, B. B., Michael Weinmann; and Eisemann, E. 2024. *Physics-Informed Gaussian Splatting*. Master’s thesis, Delft University of Technology Computer Graphics and Visualization Group. Available at <https://repository.tudelft.nl/record/uuid:7069ec01-9b5e-4957-9af4-850c4cf32d20>.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Nadaraya, E. A. 1964. On estimating regression. *Theory of Probability & Its Applications*, 9(1): 141–142.
- Pathrose, A.; Ma, L.; Berhane, H.; Scott, M. B.; Chow, K.; Forman, C.; and Markl, M. 2021. Highly accelerated aortic 4D flow MRI using compressed sensing: Performance at different acceleration factors in patients with aortic disease. *Magnetic Resonance in Medicine*, 85(4): 2174–2187.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019a. 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.

- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019b. 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.
- Saitta, S.; Carioni, M.; Mukherjee, S.; Schönlieb, C.-B.; and Redaelli, A. 2024. Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI. *Computer Methods and Programs in Biomedicine*, 246: 108057.
- Shit, S.; Zimmermann, J.; Ezhov, I.; Paetzold, J. C.; Sanches, A. F.; Pirkel, C.; and Menze, B. H. 2022. SRflow: Deep learning based super-resolution of 4D-flow MRI data. *Frontiers in Artificial Intelligence*, 5: 928181.
- Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33: 7462–7473.
- Stone, C. 1980. Optimal Rates of Convergence for Nonparametric Estimators. *The Annals of Statistics*, 8: 1348–1360.
- Stone, C. 1982. Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics*, 10: 1040–1053.
- Wang, S.; Sankaran, S.; Wang, H.; and Perdikaris, P. 2023. An expert’s guide to training physics-informed neural networks. *arXiv preprint arXiv:2308.08468*.
- Watson, G. S. 1964. Smooth regression analysis. *Sankhya: The Indian Journal of Statistics, Series A*, 26(4): 359–372.
- Zaitsev, M.; Maclaren, J.; and Herbst, M. 2015. Motion artifacts in MRI: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging*, 42(4): 887–901.