

Spline-PINN: Approaching PDEs without Data Using Fast, Physics-Informed Hermite-Spline CNNs

Nils Wandel,¹ Michael Weinmann,² Michael Neidlin,³ Reinhard Klein¹

¹ University of Bonn

² Delft University of Technology

³ RWTH Aachen University

wandel@cs.uni-bonn.de, m.weinmann@tudelft.nl, neidlin@ame.rwth-aachen.de, rk@cs.uni-bonn.de

Abstract

Partial Differential Equations (PDEs) are notoriously difficult to solve. In general, closed-form solutions are not available and numerical approximation schemes are computationally expensive. In this paper, we propose to approach the solution of PDEs based on a novel technique that combines the advantages of two recently emerging machine learning based approaches. First, physics-informed neural networks (PINNs) learn continuous solutions of PDEs and can be trained with little to no ground truth data. However, PINNs do not generalize well to unseen domains. Second, convolutional neural networks provide fast inference and generalize but either require large amounts of training data or a physics-constrained loss based on finite differences that can lead to inaccuracies and discretization artifacts.

We leverage the advantages of both of these approaches by using Hermite spline kernels in order to continuously interpolate a grid-based state representation that can be handled by a CNN. This allows for training without any precomputed training data using a physics-informed loss function only and provides fast, continuous solutions that generalize to unseen domains. We demonstrate the potential of our method at the examples of the incompressible Navier-Stokes equation and the damped wave equation. Our models are able to learn several intriguing phenomena such as Karman vortex streets, the Magnus effect, Doppler effect, interference patterns and wave reflections. Our quantitative assessment and an interactive real-time demo show that we are narrowing the gap in accuracy of unsupervised ML based methods to industrial solvers for computational fluid dynamics (CFD) while being orders of magnitude faster.

Introduction

Partial differential equations (PDEs) are an important mathematical concept to describe for example the motion of fluids, the propagation of waves, the evolution of stock markets, gravity and more. However, solving partial differential equations is a hard problem since closed-form solutions are rarely available. Thus, developing fast and accurate numerical schemes in order to find approximate solutions is of great interest for applications such as e.g. physics engines in computer games, computer generated imagery (CGI) for movies,

or computational fluid dynamics (CFD) to help engineers with simulated wind tunnel experiments.

Recently, advances of machine learning (ML) based approaches have led to promising results coping with the high computational costs associated with classical numerical methods. Furthermore, some ML based approaches are by design differentiable. This means that they offer gradients that can be used for stability analysis, optimal control or reinforcement learning. However, ML based approaches often do not generalize to domains not seen during training [Kim and Lee 2020, Mohan et al. 2020, Thuerey et al. 2019, Um et al. 2020, Raissi, Perdikaris, and Karniadakis 2019] or rely on large amounts of training data which capture the variations to be expected in foreseen scenarios [Pfaff et al. 2021, Kim et al. 2019, Ladický et al. 2015]. Recent physics-constrained approaches based on finite differences [Tompson et al. 2017, Zhu et al. 2019, Wandel, Weinmann, and Klein 2021a,b] exhibit the potential to mitigate these problems but might lead to inaccuracies and discretization artifacts, especially at high Reynolds-Numbers.

Here, we propose to use a Hermite spline CNN to obtain continuous solutions that can be trained with a physics-informed loss only. This approach combines the advantages of i) leveraging physics-informed neural networks (PINNs) to overcome the need for large amounts of training data [Jin et al. 2021, Raissi, Yazdani, and Karniadakis 2018, Raissi, Perdikaris, and Karniadakis 2019] and ii) leveraging the faster inference and better generalization capabilities offered by convolutional neural networks [Tompson et al. 2017, Wandel, Weinmann, and Klein 2021a,b]. We demonstrate the effectiveness of our approach for the incompressible Navier-Stokes equations as well as the damped wave equation. The incompressible Navier-Stokes equations, which are particularly hard to solve due to the non-linear advection term and thus the main focus of this paper, are investigated for Reynolds numbers ranging from 2-10000. To assess the accuracy of our method, we compute drag and lift coefficients on a CFD benchmark domain [Schäfer and Turek 1996] and compare the results with official benchmark values. For both equations, we perform generalization experiments, evaluate the stability over long time horizons and present an interactive real-time demonstration. To ensure full reproducibility, our code is publicly available on github: https://github.com/aschethor/Spline_PINN.

Related Work

Recent developments indicate the potential of efficient surrogate models based on machine learning, and in particular deep learning, to approximate the dynamics of partial differential equations (PDEs).

Lagrangian methods like smoothed particle hydrodynamics (SPH) [Gingold and Monaghan 1977] represent fluids in the reference frame of individual particles that follow the fluid’s velocity field. Relying on this principle, learning-based Lagrangian approaches include the simulation of fluids based on regression forests [Ladický et al. 2015], graph neural networks [Mrowca et al. 2018, Li et al. 2019], continuous convolutions [Ummenhofer et al. 2020] and Smooth Particle Networks (SP-Nets) [Schenck and Fox 2018]. While Lagrangian methods are particularly suitable for fluid domains that exhibit large, dynamic surfaces such as waves and droplets, however, Eulerian methods typically allow a more accurate simulation of fluid dynamics within a certain fluid domain (see [Foster and Metaxas 1996]).

Eulerian methods model domain characteristics such as a velocity or pressure field on a fixed reference frame. Respective techniques leverage implicit neural representations, grid- and mesh-structures to describe the domain of a PDE.

Continuous Eulerian methods exploit the direct mapping of domain coordinates (i.e. positional coordinates and time) to field values (i.e. velocity, pressure, etc.) relying e.g. on implicit neural representations to obtain smooth and accurate simulations [Sirignano and Spiliopoulos 2018a, Grohs et al. 2018, Khoo, Lu, and Ying 2019, Raissi, Perdikaris, and Karniadakis 2019, Lu et al. 2021] and handle the curse of dimensionality faced by discrete techniques in the case of high-dimensional PDEs [Grohs et al. 2018]. Respective applications include the modeling of flow through porous media [Zhu and Zabarar 2018, Zhu et al. 2019, Tripathy and Bilonis 2018], fluid modeling [Yang, Yang, and Xiao 2016, Raissi, Yazdani, and Karniadakis 2018], turbulence modeling [Geneva and Zabarar 2019, Ling, Kurzawski, and Templeton 2016], wave propagation [Rasht-Behesht et al. 2021, Sitzmann et al. 2020] and the modeling of molecular dynamics [Schöberl, Zabarar, and Koutsourelakis 2019]. Training typically relies on penalizing residuals of the underlying PDEs based on physics-informed loss functions for a specific domain, which prevents the networks from generalizing to new domain geometries and being used in the scope of interactive scenarios. Recently, Wang et al. [Wang et al. 2021] proposed to stitch pretrained Deep Learning models together to speed up the learning process for boundary value problems. However, temporal dependencies are not embedded in their method yet preventing interactive applications.

Discrete Eulerian methods such as finite difference, lattice Boltzmann [Chen and Doolen 1998, Guo 2021] or finite element methods instead tackle the underlying PDEs on an explicit representation such as a discrete grid or a graph structure. Beyond early seminal work [Harlow and Welch 1965, Stam 1999], recent developments focus on exploiting the potential of deep learning techniques to achieve significant speed-ups while preserving accuracy. Learning fluid

simulations has been approached with deep generative models for efficient interpolations between various simulation settings [Kim et al. 2019] and turbulent flow fields within pipe domains [Kim and Lee 2020]. However, both of these approaches do not generalize to new domain geometries that have not been used during training. To speed up Eulerian fluid simulations, Tompson et al. [Tompson et al. 2017] proposed to learn a Helmholtz projection step with a CNN that generalizes to new domain geometries that have not been used during the training. However, this approach requires a path tracer to deal with the advection term of the Navier-Stokes equations. Furthermore, as viscosity is not considered in Eulerian fluid models, phenomena such as the Magnus effect and Karman vortex streets cannot be simulated. Ensuring incompressibility of fluids has been achieved with discretized vector potentials [Kim et al. 2019, Mohan et al. 2020], however, these techniques are not capable of generalizing to domain geometries not seen during the training. The approach by Geneva et al. [Geneva and Zabarar 2020] discards the pressure term in the Navier-Stokes equations which leads to the simpler Burgers’ equations, for which they learned the update step based on a physics-constrained framework. Thuerey et al. [Thuerey et al. 2019] learn to solve the Reynolds-averaged Navier-Stokes equations, but the specific approach prevents a generalization beyond air-foil flows and discards the temporal evolution of the fluid state. Um et al. [Um et al. 2020] proposed to learn a correction step which allows approximating solutions of a high-resolution fluid simulations in terms of a low-resolution differentiable fluid solver, but the generalization capability to domain geometries that have not been used during training has not been shown. In contrast to the aforementioned techniques, the approach by Wandel et al. [Wandel, Weinmann, and Klein 2021a] and its 3D extension [Wandel, Weinmann, and Klein 2021b] overcome the dependence on huge amounts of training data by using a physics-constrained loss and introducing a training cycle that recycles data generated by the network during training. This allows handling dynamic boundary conditions as required for interactions with the fluid [Wandel, Weinmann, and Klein 2021a,b] and, with an improved network [Wandel, Weinmann, and Klein 2021b], also dealing with temporally varying fluid parameters such as viscosity and density during simulation. However, their physics-constrained loss can lead to discretization artifacts and requires a pressure gradient regularizer for high Reynolds numbers. Further learning-based approaches also exploit graph representations in terms of graph neural networks [Harsch and Riedelbauch 2021, Sanchez-Gonzalez et al. 2020], graph convolutional networks [Gao, Zahr, and Wang 2021] as well as mesh representations [Pfaff et al. 2021] or subspace representations [Sirignano and Spiliopoulos 2018b, Ainsworth and Dong 2021]. Unfortunately, graph neural networks cannot make use of the highly efficient implementations for convolutional operations on grids and thus usually come with higher computational complexity per node compared to grid based CNNs.

To the best of our knowledge, the potential of Spline CNNs [Fey et al. 2017] has not yet been investigated for learning the dynamics of PDEs and, hence, the presented

physics-informed Hermite Spline CNN (Spline-PINN) approach is the first method that leverages both implicit and explicit characteristics to obtain fast, interactive, continuous surrogate models that generalize and can be trained without training data using a physics-informed loss based on the underlying PDEs.

Method

In the main part of our paper, we first briefly discuss the background regarding partial differential equations with a focus on the Navier-Stokes equations and the damped wave equation, before we introduce our physics-informed Hermite spline CNN approach (Spline-PINN), that combines the advantages of PINNs regarding physics-informed training without training data and Spline-CNNs regarding their generalization capability.

Partial Differential Equations

Partial Differential Equations describe the dependencies between partial derivatives of a multivariate function inside a domain Ω and usually need to be solved for given initial conditions at the beginning of the simulation and boundary conditions at the domain boundaries $\partial\Omega$.

The Incompressible Navier-Stokes Equations describe the dynamics of a fluid with a pressure-field p and a velocity-field \vec{v} inside a domain Ω by means of the incompressibility and momentum equation:

$$\nabla \cdot \vec{v} = 0 \quad \text{in } \Omega \quad (1)$$

$$\rho \dot{\vec{v}} = \rho(\partial_t \vec{v} + \vec{v} \cdot \nabla \vec{v}) = \mu \Delta \vec{v} - \nabla p + \vec{f} \quad \text{in } \Omega \quad (2)$$

Here, ρ represents the fluid's density and μ its viscosity. The external forces \vec{f} are neglected in our experiments. These two equations have to be solved for given initial conditions \vec{v}_0, p_0 and boundary conditions (BCs). Here, we consider the Dirichlet BC to set the velocity field $\vec{v} = \vec{v}_d$ at the domain boundaries $\partial\Omega$. In this work, we exploit the Helmholtz decomposition theorem and use a vector potential \vec{a} with $\vec{v} = \nabla \times \vec{a}$ to automatically ensure incompressibility. Furthermore, we restrict our considerations to 2D flows and, hence, only the z -component of \vec{a} , a_z , is needed. The Navier-Stokes equations are considered particularly hard to solve due to the non-linear advection term ($\vec{v} \cdot \nabla \vec{v}$) and, therefore, are the main focus of our investigations.

The Damped Wave Equation can be used to describe for example the dynamics of a thin membrane with height-field z and velocity-field v_z .

$$\partial_t z = v_z \quad \text{in } \Omega \quad (3)$$

$$\partial_t v_z = k \Delta z - \delta v_z \quad \text{in } \Omega \quad (4)$$

Here, k is the stiffness constant of the membrane and δ is a damping constant. As for the Navier-Stokes equations, we solve these equations for given initial conditions z_0, v_{z0} and Dirichlet boundary conditions: $z = z_d$ on $\partial\Omega$.

Although the damped wave equation contains only linear components, we present this additional example to demonstrate that our method might also work for different classes

of PDEs. The wave equation lays the foundation for many more complex equations such as the electromagnetic wave equation or the Schrödinger equation. Furthermore, for high damping constants δ , z converges towards the solution of the Laplace equation.

Spline-PINN

In this section, we introduce our Physics-Informed Hermite Spline CNN. First, we provide an overview on Hermite splines and then show how to incorporate them into a CNN for solving PDEs.

Hermite Splines In general, Hermite splines are piecewise polynomials that are defined by their values and their first n derivatives at given support points. To facilitate CNN-based processing, we arrange the support points on a regular grid. Figure 1 shows an example of Hermite spline kernel functions in 1D for $n = 0, 1, 2$. We define the kernel functions $h_i^n(x)$, such that the values and first n derivatives at the support points ($x = -1, 0, 1$) are set to 0 except the i th derivative at $x = 0$ which takes a value such that $h_i^n(x) \in [-1, 1]$. In contrast to B-Spline kernels, the support of Hermite spline kernels ranges only over the directly neighboring grid cells. This facilitates learning of high frequencies and allows for computationally more efficient interpolations with comparatively small transposed convolution kernels. By linearly combining Hermite spline kernel functions of order n at every grid-cell we obtain continuous piecewise polynomials with $(n + 1)$ -th derivatives of bounded variation. Choosing the right spline order n is important to be able to compute the physics informed loss for a given PDE as will be discussed below.

To obtain kernel functions in multiple dimensions, we use the tensor product of multiple 1D kernel functions:

$$h_{i,j,k}^{l,m,n}(x, y, t) = h_i^l(x) h_j^m(y) h_k^n(t) \quad (5)$$

In the supplementary, we show how basis flow fields can be obtained by taking the curl of $h_{i,j}^{l,m}(x, y)$. On a grid $\hat{x}, \hat{y}, \hat{t} \in \hat{X} \times \hat{Y} \times \hat{T}$ with discrete spline coefficients $c_{\hat{x}, \hat{y}, \hat{t}}^{i,j,k}$, we obtain a continuous Hermite spline $g(x, y, t)$ as follows:

$$g(x, y, t) = \sum_{\substack{i,j,k \in [0:l] \times [0:m] \times [0:n] \\ \hat{x}, \hat{y}, \hat{t} \in \hat{X} \times \hat{Y} \times \hat{T}}} c_{\hat{x}, \hat{y}, \hat{t}}^{i,j,k} h_{i,j,k}^{l,m,n}(x - \hat{x}, y - \hat{y}, t - \hat{t}) \quad (6)$$

Our goal is, to find spline coefficients $c_{\hat{x}, \hat{y}, \hat{t}}^{i,j,k}$ such that $g(x, y, t)$ matches the solution of PDEs as closely as possible. The partial derivatives of g with respect to x, y, t can be directly obtained by taking the corresponding derivatives of the spline kernel functions.

Pipeline Figure 2 depicts the pipeline of our Spline-PINN. We use a CNN (PDE Model) to map discrete Hermite-spline coefficients and boundary conditions from a timepoint \hat{t} to spline coefficients at a succeeding timepoint $\hat{t} + dt$. By recurrently applying the PDE Model on the spline coefficients, the simulation can be unrolled in time. From these discrete spline-coefficients, continuous Hermite splines can be efficiently obtained using transposed convolutions (see Equation 6) with interpolation kernels as depicted in Figure 1.

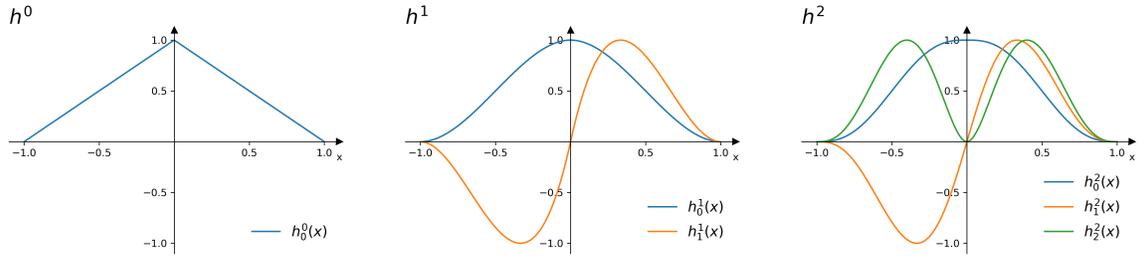


Figure 1: 1D Hermite spline kernels for $n = 0, 1, 2$ (scaled between -1 and 1). Note, that these kernel functions are in C^n and thus, the $(n + 1)$ -th derivatives are of bounded variation. The case $n = 0$ could be considered as a linear interpolation between spline coefficients.

This way, training samples as well as evaluation samples can be taken at any point in space and time.

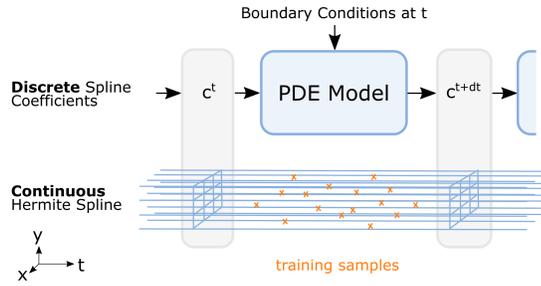


Figure 2: Pipeline of PDE Model with Hermite spline interpolation. Since the solution is continuous, training samples and evaluation samples can be obtained at any point in space and time. More detailed views of the PDE Models used for the Navier-Stokes equations and damped wave equation are provided in the supplementary material.

Physics Informed Loss When training neural networks on PDEs, the information provided by the PDEs themselves can be used to save or even completely spare out training-data. In literature, two major approaches have been established:

Physics-*constrained* approaches [Zhu et al. 2019] compute losses *explicitly* - usually on a grid topology using finite differences in order to approximate the derivatives of a PDE. This approach is suitable to train CNNs and allows neural networks to learn e.g. the incompressible Navier-Stokes equations without any ground truth training-data and generalize to new domains [Wandel, Weinmann, and Klein 2021a]. However, relying on finite differences may lead to inaccuracies and strong discretization artifacts - especially at high Reynolds-numbers (see Figure 3).

In contrast, physics-*informed* approaches [Raissi, Perdikaris, and Karniadakis 2019] compute losses *implicitly* - usually by taking derivatives of an implicit field description. This approach enables efficient training of implicit neural networks and yields continuous solutions that can be evaluated at any point in space and time. However, implicit neural networks do not generalize well to novel domains but usually require retraining of the network.

Here, we can combine the advantages of both approaches:

By using a convolutional neural network that processes spline coefficients which can be considered as a discrete hidden latent description for a continuous implicit field description based on Hermite splines, our Spline-PINN approach is capable to generalize to new domain geometries and yields continuous solutions that avoid the detrimental effects of a discretized loss function based on finite differences.

We aim to optimize the spline coefficients such that the integrals of the squared residuals of the PDEs over the domain / domain-boundaries and time steps are minimized. To compute these integrals, we uniformly randomly sample points within the given integration domains.

For the Navier-Stokes equation, we consider a momentum loss term L_p defined as:

$$L_p = \int_{\Omega} \int_{\hat{t}}^{\hat{t}+dt} \|\rho(\partial_t \vec{v} + \vec{v} \cdot \nabla \vec{v}) - \mu \Delta \vec{v} + \nabla p\|^2 \quad (7)$$

It is important that the residuals are of bounded variation, otherwise we can not compute the integral. L_p contains third order derivatives for a_z in space (see viscosity term: $\mu \Delta(\nabla \times a_z)$) and thus, following the argument in Figure 1, we have to choose at least $l, m = 2$ for the Hermite spline kernels of a_z . The boundary loss term L_b is given by:

$$L_b = \int_{\partial\Omega} \int_{\hat{t}}^{\hat{t}+dt} \|\vec{v} - \vec{v}_d\|^2 \quad (8)$$

L_p and L_b are then combined in the final loss term $L_{\text{tot}}^{\text{flow}}$ according to:

$$L_{\text{tot}}^{\text{flow}} = \alpha L_p + \beta L_b \quad (9)$$

Here, we set hyperparameters $\alpha = 10$ and $\beta = 20$. β was chosen higher compared to α to prevent the flow field from leaking through boundaries. A more thorough discussion regarding the choice of hyperparameters is provided in the supplementary material.

For the damped wave equation, we need two loss terms inside the domain Ω :

$$L_z = \int_{\Omega} \int_{\hat{t}}^{\hat{t}+dt} \|\partial_t z - v_z\|^2 \quad (10)$$

$$L_v = \int_{\Omega} \int_{\hat{t}}^{\hat{t}+dt} \|\partial_t v_z - k \Delta z + \delta v_z\|^2 \quad (11)$$

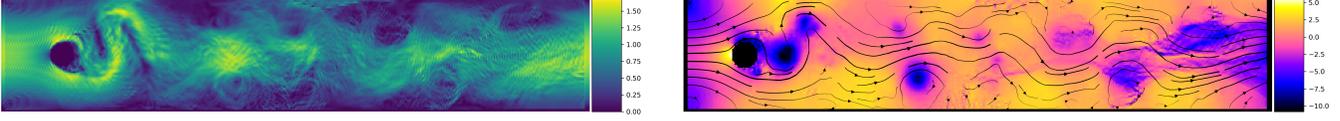


Figure 3: Severe discretization artifacts appear using a physics-constrained loss based on a finite differences Marker and Cell grid at $Re = 10000$ - even at twice the resolution compared to our Hermite spline approach.

Similar to the incompressible Navier-Stokes equation, we have second order derivatives of z in space (see $k\Delta z$ -Term) and thus have to choose at least $l, m = 1$ for the Hermite spline kernels of z . The boundary loss term is:

$$L_b = \int_{\partial\Omega} \int_{\hat{t}}^{\hat{t}+dt} \|z - z_d\|^2 \quad (12)$$

L_z, L_v and L_b are then combined in the final loss term $L_{\text{tot}}^{\text{wave}}$:

$$L_{\text{tot}}^{\text{wave}} = \alpha L_z + \beta L_v + \gamma L_b \quad (13)$$

For the wave equation, we set the hyperparameters $\alpha = 1, \beta = 0.1, \gamma = 10$.

Training Procedure As presented in Figure 4, we use a training procedure similar to [Wandel, Weinmann, and Klein 2021a] but replace the physics-constrained loss based on finite differences by the just introduced physics-informed loss on the spline coefficients. First, we initialize a training pool of randomized training domains and spline coefficients. At the beginning, all spline coefficients can be set to 0, thus no ground truth data is required. Then, we draw a random minibatch (batchsize = 50) from the training pool and feed it into the PDE Model, which predicts the spline coefficients for the next time step. Then, we compute a physics-informed loss inside the volume spanned by the spline coefficients of the minibatch and the predicted spline coefficients in order to optimize the weights of the PDE Model with gradient descent. To this end, we use the Adam optimizer (learning rate = 0.0001). Finally, we update the training pool with the just predicted spline coefficients in order to fill the pool with more and more realistic training data over the course of training. From time to time, we reset spline coefficients to 0 in the training pool in order to also learn the warm-up phase from 0 spline coefficients. Training took 1-2 days on a NVidia GeForce RTX 2080 Ti.

Results

In the following, we provide qualitative and quantitative evaluations of our approach for fluid and wave simulations. Additional results as well as demonstrations where the user can interact with the PDE by dynamically changing boundary conditions, are contained in our supplemental.

Fluid Simulation

Qualitative Evaluation The dynamics of a fluid depends strongly on the Reynolds number, which is a dimensionless quantity that relates the fluid density ρ , mean velocity $\|\vec{v}\|$, obstacle diameter L and viscosity μ in the following way:

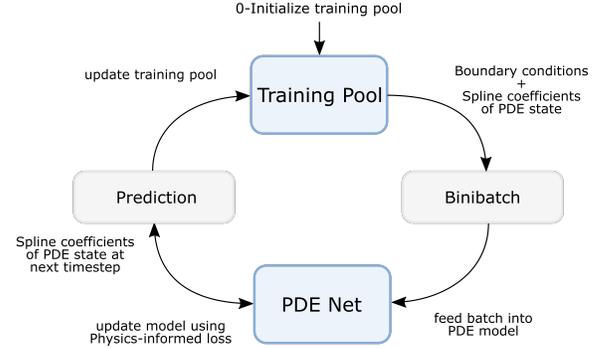


Figure 4: Training cycle similar to [Wandel, Weinmann, and Klein 2021a]

$$Re = \frac{\rho \|\vec{v}\| L}{\mu} \quad (14)$$

To validate our method, we reconstructed the DFG benchmark setup [Schäfer and Turek 1996] and scaled its size and viscosity by a factor of 100. This way, we obtained a grid size of 41×220 cells for the fluid domain and the obstacle diameter is 10 cells.

For very small Reynolds numbers, the flow field becomes basically time-reversible which can be recognized by the symmetry of the stream lines in Figure 5 a). At $Re = 20$ a steady wake is forming behind the obstacle (see Figure 5 b). For higher Reynolds numbers such as $Re = 100$, this wake becomes unstable and a von Karman vortex street is forming (see Figure 5 c). Figure 5 d) and e) show turbulent flow fields at very high Reynolds numbers ($Re = 1000$ and $Re = 10000$). Notice the high level of details in the velocity field learned by our method without any ground truth data.

Quantitative Evaluation To evaluate our method quantitatively, we computed the forces exerted by the fluid on a cylinder (see Figure 6) and compared the resulting drag and lift coefficients to an implicit physics-informed neural network similar to [Raissi, Yazdani, and Karniadakis 2018], a MAC grid based physics-constrained neural network [Wandel, Weinmann, and Klein 2021a], an industrial CFD solver (Ansys) and official results reported on a CFD benchmark [Schäfer and Turek 1996]. To the best of our knowledge, this is the first time such forces were obtained by a fully unsupervised machine learning based approach. The forces can be separated into two terms stemming from viscous friction

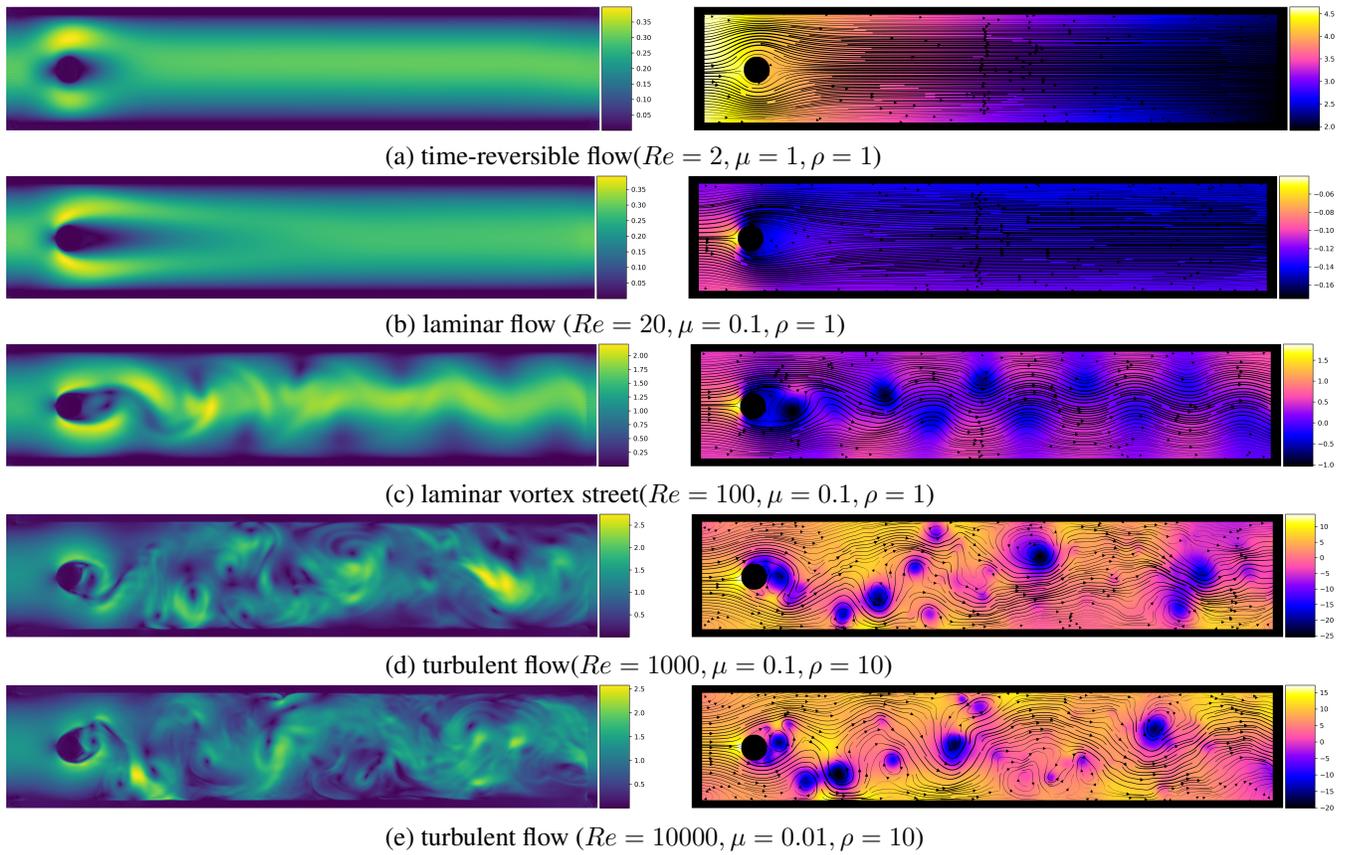


Figure 5: Flow and pressure fields around a cylinder obtained by our method at different Reynolds numbers. Left side: velocity magnitude; Right side: pressure field and stream lines of velocity field. An animated real-time visualization of these experiments is provided in the supplementary video (www.youtube.com/watch?v=QC98LCtZn0).

(\vec{F}_μ) and the pressure field (\vec{F}_p):

$$\vec{F}_\mu = \int_S \mu(\nabla\vec{v})\vec{n}ds; \vec{F}_p = \int_S -p\vec{n}ds \quad (15)$$

$$\vec{F}_{\text{tot}} = \vec{F}_\mu + \vec{F}_p \quad (16)$$

Figure 6 shows the distribution of such viscous drag and pressure forces along the cylinder surface S with surface normals \vec{n} .

The drag-force F_D is the parallel force-component of \vec{F}_{tot} to the flow direction ($F_D = \vec{F}_{\text{tot},x}$) while the lift force F_L is its orthogonal component ($F_L = \vec{F}_{\text{tot},y}$). From these forces, it is possible to compute drag and lift coefficients as follows:

$$C_D = \frac{2F_D}{U_{\text{mean}}^2 L}, C_L = \frac{2F_L}{U_{\text{mean}}^2 L} \quad (17)$$

Tables 1 and 2 compare the drag and lift coefficients obtained by our method to an implicit PINN, a finite-difference based MAC grid approach [Wandel, Weinmann, and Klein 2021a], an industrial CFD solver (Ansys) and the official DFG benchmark [Schäfer and Turek 1996] for Reynolds numbers 2, 20 and 100. For all Reynolds numbers, the Spline-PINN approach returned significantly improved drag coefficients compared to the implicit PINN and

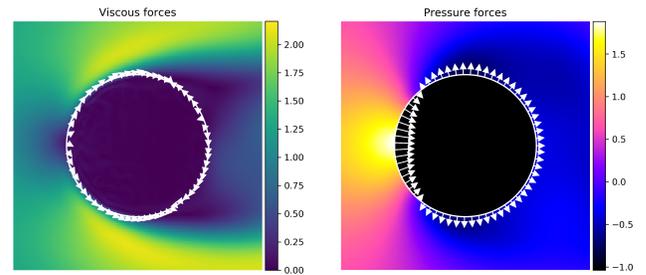


Figure 6: White arrows indicate viscous and pressure forces (see Equation 15) acting on a cylinder at $Re = 100$ (see Figure 5 c). While forces from viscous friction are parallel to the obstacle's surface, pressure forces are always perpendicular to the surface.

MAC-approach. For $Re = 100$, the implicit PINN fails to capture the dynamics of the von Karman vortex street as it is only trained on the domain boundaries without any additional information stemming for example from a moving die [Raissi, Yazdani, and Karniadakis 2018] or dynamic boundary conditions [Jin et al. 2021]. In contrast, our method is

able to reproduce such oscillations as can be seen in Figure 7. The implicit PINN required retraining for every setting which took about 1 day while performing the computations with the MAC grid based approach or our Spline-PINN was achieved basically in real-time at 60 and 30 iterations per second respectively. The Ansys solver took 9 sec for $Re = 2$, 13 sec for $Re = 20$ and 37 min for $Re = 100$. A comparison of errors in $E[|\nabla \cdot \vec{v}|]$ and L_p as well as a stability analysis can be found in the supplemental.

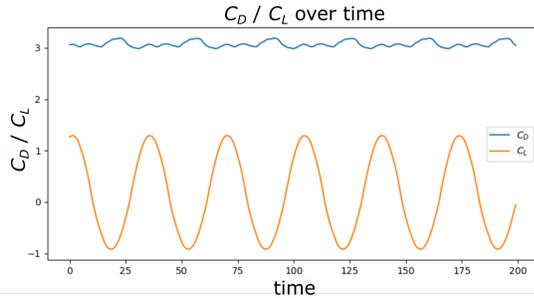


Figure 7: Oscillating drag and lift coefficients over time obtained by our Spline-PINN at $Re = 100$.

Wave Simulation

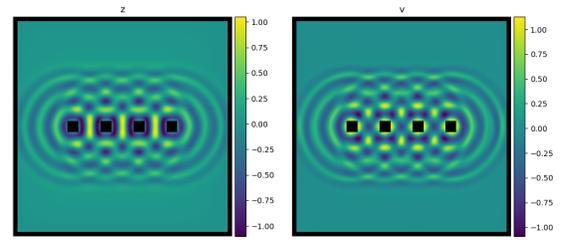
Qualitative Evaluation Figure 8 shows several experimental results that were obtained by our method for the damped wave equation on a 200×200 domain. First, we investigated interference patterns that arise when waves from different directions are superimposed. In Figure 8 a), the waves of 4 oscillators interfere with each other. Then, in Figure 8 b), we investigate whether our method is able to learn the Doppler effect for moving oscillators. This effect is well-known e.g. for changing the pitch of an ambulance-car that drives by. Finally, in Figure 8 c), we show the reflection and interference behavior of waves hitting Dirichlet boundaries.

As for the fluid simulations, all of these results were obtained without relying on any ground truth data. Furthermore, the domain in Figure 8 c) was not contained in the randomized set of training domains indicating good generalization performance.

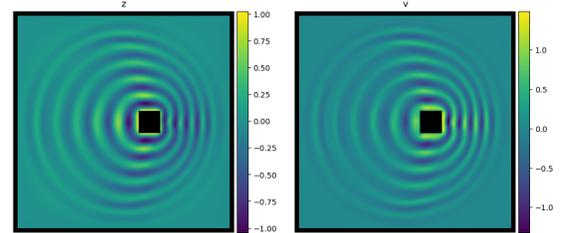
Quantitative Evaluation Table 3 compares the losses of our method on the oscillator domain (see Figure 8 a). We trained two versions with two different spline orders in the spatial dimensions ($l, m = 1$ and $l, m = 2$) and observed significantly better performance for $l, m = 2$. The stability of our approach for the wave equation is examined in the supplementary material.

Conclusion

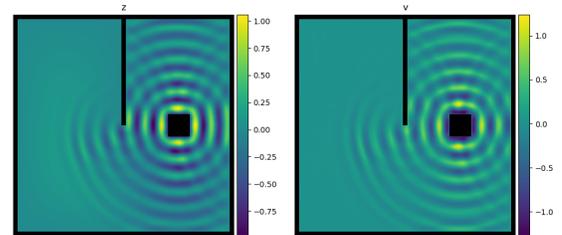
In this work, we approach the incompressible Navier-Stokes equation and the damped wave equation by training a continuous Hermite Spline CNN using physics-informed loss functions only. While finite-difference based methods break down at high Reynolds-numbers due to discretization artifacts, our method still returns visually appealing results at



(a) Interference patterns forming around 4 oscillators.



(b) Doppler effect of an oscillator moving to the right.



(c) Wave reflections on domain boundaries.

Figure 8: Results of our method for the wave equation ($k = 10, \delta = 0.1$). Left side: height field z ; Right side: velocity field v_z . The domain boundaries are marked in black. For better performance, second-order splines were used in x and y . An animated real-time visualization of these experiments can be found at www.youtube.com/watch?v=QC98LCtZn0.

$Re = 10000$. Furthermore, we investigated drag and lift coefficients on a CFD benchmark setup and observed reasonable accordance with officially reported values which is remarkable given the fact that our method does not rely on any ground truth data.

In the future, further boundary conditions (e.g. von Neumann BC) could be incorporated into the PDE model as well. To further refine the solutions at the boundary layers, a multigrid Hermite spline CNN could be considered. The fully differentiable nature of our method may also help in reinforcement learning scenarios, optimal control, sensitivity analysis or gradient based shape optimization.

We believe that our method could have applications in physics engines of computer games or in computer-generated imagery as it provides fast and visually plausible solutions. The obtained drag and lift coefficients indicate that in the future, unsupervised ML based methods could reach levels of accuracy that are sufficiently close to traditional industrial CFD solvers making them suitable for fast prototyping in engineering applications. We firmly believe

Method	Re=2		Re=20	
	C_D	C_L	C_D	C_L
implicit PINN	25.3	0.478	3.299	0.0744
MAC grid [Wandel, Weinmann, and Klein 2021a]	25.76	-0.824	4.414	-0.597
Spline-PINN (ours)	29.7	-0.456	4.7	5.64e-04
Ansys	32.035	0.774	5.57020	0.00786
DFG-Benchmark [Schäfer and Turek 1996]	-	-	5.58	0.0106

Table 1: Drag and lift coefficients obtained by an implicit PINN, our Hermite spline approach, an industrial CFD solver (Ansys) and official results from the DFG-benchmark [Schäfer and Turek 1996] for $Re = 2, 20$.

Method	Re=100		
	C_D	C_L	time
implicit PINN	1.853	-0.02445	~ 1 day
MAC grid [Wandel, Weinmann, and Klein 2021a]	(2.655 / 2.693 / 2.725)*	(-0.757 / 0.0184 / 0.86)*	~ 15 sec
Spline-PINN (ours)	(2.985 / 3.068 / 3.188)*	(-0.926 / 0.179 / 1.295)*	~ 10 sec
Ansys	(3.234 / 3.273 / 3.31)*	(-1.14 / -0.059 / 1.07)*	37 min
DFG-Benchmark [Schäfer and Turek 1996]	(3.1569 / 3.1884 / 3.220)*	(-1.0206 / -0.0173 / 0.9859)*	-

Table 2: Drag and lift coefficients obtained by an implicit PINN, our Hermite spline approach, an industrial CFD solver (Ansys) and official results from the DFG-benchmark [Schäfer and Turek 1996] for $Re = 100$. *:(minimum/average/maximum)-values for oscillating coefficients

Spline order	L_z	L_v	L_b
$l, m = 1$	8.511e-02	1.127e-02	1.425e-03
$l, m = 2$	5.294e-02	6.756e-03	1.356e-03

Table 3: Quantitative results of wave equation.

that moving from explicit physics-constrained losses to implicit physics-informed losses on continuous fields based on discrete latent descriptions such as spline coefficients will positively influence the performance of future ML based PDE solvers that generalize.

Acknowledgements

This work has been supported by the German Research Foundation (DFG), FOR 2535 Anticipating Human Behavior.

References

Ainsworth, M.; and Dong, J. 2021. Galerkin Neural Networks: A Framework for Approximating Variational Equations with Error Control. *arXiv preprint arXiv:2105.14094*.

Chen, S.; and Doolen, G. D. 1998. Lattice Boltzmann method for fluid flows. *Annual review of fluid mechanics*, 30(1): 329–364.

Fey, M.; Lenssen, J. E.; Weichert, F.; and Müller, H. 2017. SplineCNN: Fast Geometric Deep Learning with Continuous B-Spline Kernels. *CoRR*, abs/1711.08920.

Foster, N.; and Metaxas, D. 1996. Realistic Animation of Liquids. *Graphical Models and Image Processing*, 58(5): 471 – 483.

Gao, H.; Zahr, M. J.; and Wang, J. 2021. Physics-informed graph neural Galerkin networks: A unified framework for solving PDE-governed forward and inverse problems. *CoRR*, abs/2107.12146.

Geneva, N.; and Zabarar, N. 2019. Quantifying model form uncertainty in Reynolds-averaged turbulence models with Bayesian deep neural networks. *Journal of Computational Physics*, 383: 125 – 147.

Geneva, N.; and Zabarar, N. 2020. Modeling the dynamics of PDE systems with physics-constrained deep auto-regressive networks. *Journal of Computational Physics*, 403: 109056.

Gingold, R. A.; and Monaghan, J. J. 1977. Smoothed particle hydrodynamics: theory and application to non-spherical stars. *Monthly notices of the royal astronomical society*, 181(3): 375–389.

Grohs, P.; Hornung, F.; Jentzen, A.; and Von Wurstemberger, P. 2018. A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of Black-Scholes partial differential equations. *arXiv preprint arXiv:1809.02362*.

Guo, Z. 2021. Well-balanced lattice Boltzmann model for two-phase systems. *Physics of Fluids*, 33(3): 031709.

Harlow, F. H.; and Welch, J. E. 1965. Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *The physics of fluids*, 8(12): 2182–2189.

- Harsch, L.; and Riedelbauch, S. 2021. Direct Prediction of Steady-State Flow Fields in Meshed Domain with Graph Networks. arXiv:2105.02575.
- Jin, X.; Cai, S.; Li, H.; and Karniadakis, G. E. 2021. NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 426: 109951.
- Khoo, Y.; Lu, J.; and Ying, L. 2019. Solving for high-dimensional committor functions using artificial neural networks. *Research in the Mathematical Sciences*, 6(1): 1.
- Kim, B.; Azevedo, V. C.; Thuerey, N.; Kim, T.; Gross, M.; and Solenthaler, B. 2019. Deep fluids: A generative network for parameterized fluid simulations. In *Computer Graphics Forum*, volume 38, 59–70. Wiley Online Library.
- Kim, J.; and Lee, C. 2020. Deep unsupervised learning of turbulence for inflow generation at various Reynolds numbers. *Journal of Computational Physics*, 406: 109216.
- Ladický, L.; Jeong, S.; Solenthaler, B.; Pollefeys, M.; and Gross, M. 2015. Data-Driven Fluid Simulations Using Regression Forests. *ACM Trans. Graph.*, 34(6).
- Li, Y.; Wu, J.; Tedrake, R.; Tenenbaum, J. B.; and Torralba, A. 2019. Learning Particle Dynamics for Manipulating Rigid Bodies, Deformable Objects, and Fluids. In *ICLR*.
- Ling, J.; Kurzawski, A.; and Templeton, J. 2016. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *Journal of Fluid Mechanics*, 807: 155–166.
- Lu, L.; Meng, X.; Mao, Z.; and Karniadakis, G. E. 2021. DeepXDE: A Deep Learning Library for Solving Differential Equations. *SIAM Review*, 63(1): 208–228.
- Mohan, A. T.; Lubbers, N.; Livescu, D.; and Chertkov, M. 2020. Embedding Hard Physical Constraints in Neural Network Coarse-Graining of 3D Turbulence. arXiv:2002.00021.
- Mrowca, D.; Zhuang, C.; Wang, E.; Haber, N.; Fei-Fei, L.; Tenenbaum, J. B.; and Yamins, D. L. K. 2018. Flexible Neural Representation for Physics Prediction. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, 8813–8824. Red Hook, NY, USA: Curran Associates Inc.
- Pfaff, T.; Fortunato, M.; Sanchez-Gonzalez, A.; and Battaglia, P. W. 2021. Learning Mesh-Based Simulation with Graph Networks. *ICLR 2021*:2010.03409.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. 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.; Yazdani, A.; and Karniadakis, G. E. 2018. Hidden Fluid Mechanics: A Navier-Stokes Informed Deep Learning Framework for Assimilating Flow Visualization Data. arXiv preprint arXiv:1808.04327.
- Rasht-Behesht, M.; Huber, C.; Shukla, K.; and Karniadakis, G. E. 2021. Physics-informed Neural Networks (PINNs) for Wave Propagation and Full Waveform Inversions. arXiv:2108.12035.
- Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; and Battaglia, P. W. 2020. Learning to Simulate Complex Physics with Graph Networks. arXiv:2002.09405.
- Schenck, C.; and Fox, D. 2018. SPNets: Differentiable Fluid Dynamics for Deep Neural Networks. In *Conference on Robot Learning*, 317–335.
- Schäfer, M.; and Turek, S. 1996. Benchmark Computations of Laminar Flow Around a Cylinder (The CFD Benchmarking Project). <http://www.mathematik.tu-dortmund.de/feat-flow/en/benchmarks/cfdbenchmarking.html>. (accessed at 08-Sep-2021).
- Schöberl, M.; Zabaras, N.; and Koutsourelakis, P.-S. 2019. Predictive collective variable discovery with deep Bayesian models. *The Journal of Chemical Physics*, 150(2): 024109.
- Sirignano, J.; and Spiliopoulos, K. 2018a. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375: 1339 – 1364.
- Sirignano, J.; and Spiliopoulos, K. 2018b. DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375: 1339–1364.
- 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.
- Stam, J. 1999. Stable fluids. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 121–128.
- Thuerey, N.; Weissenow, K.; Prantl, L.; and Hu, X. 2019. Deep Learning Methods for Reynolds-Averaged Navier-Stokes Simulations of Airfoil Flows. *AIAA Journal*, 1–12.
- Tompson, J.; Schlachter, K.; Sprechmann, P.; and Perlin, K. 2017. Accelerating eulerian fluid simulation with convolutional networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3424–3433. JMLR. org.
- Tripathy, R. K.; and Bilonis, I. 2018. Deep UQ: Learning deep neural network surrogate models for high dimensional uncertainty quantification. *Journal of Computational Physics*, 375: 565 – 588.
- Um, K.; Fei, R.; Holl, P.; Brand, R.; and Thuerey, N. 2020. Solver-in-the-Loop: Learning from Differentiable Physics to Interact with Iterative PDE-Solvers. arXiv:2007.00016.
- Ummenhofer, B.; Prantl, L.; Thuerey, N.; and Koltun, V. 2020. Lagrangian Fluid Simulation with Continuous Convolutions. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Wandel, N.; Weinmann, M.; and Klein, R. 2021a. Learning Incompressible Fluid Dynamics from Scratch - Towards Fast, Differentiable Fluid Models that Generalize. *9th International Conference on Learning Representations, ICLR 2021*.
- Wandel, N.; Weinmann, M.; and Klein, R. 2021b. Teaching the Incompressible Navier-Stokes Equations to Fast Neural Surrogate Models in 3D. *Physics of Fluids (AIP)*.

Wang, H.; Planas, R.; Chandramowlishwaran, A.; and Bostanabad, R. 2021. Train Once and Use Forever: Solving Boundary Value Problems in Unseen Domains with Pre-trained Deep Learning Models. *CoRR*, abs/2104.10873.

Yang, C.; Yang, X.; and Xiao, X. 2016. Data-driven projection method in fluid simulation. *Computer Animation and Virtual Worlds*, 27(3-4): 415–424.

Zhu, Y.; and Zabaras, N. 2018. Bayesian deep convolutional encoder–decoder networks for surrogate modeling and uncertainty quantification. *Journal of Computational Physics*, 366: 415 – 447.

Zhu, Y.; Zabaras, N.; Koutsourelakis, P.-S.; and Perdikaris, P. 2019. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394: 56 – 81.