Contents lists available at ScienceDirect



Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Preconditioned FEM-based neural networks for solving incompressible fluid flows and related inverse problems

Franziska Griese^(D), Fabian Hoppe^(D), Alexander Rüttgers^(D), Philipp Knechtges^{(D)*}

German Aerospace Center (DLR), Institute of Software Technology, High-Performance Computing Department, Linder Höhe, Cologne, 51147, Germany

ARTICLE INFO

Keywords: FEM-based neural network Machine learning Preconditioning Finite elements Parametric PDEs Stokes Navier–Stokes

ABSTRACT

The numerical simulation and optimization of technical systems described by partial differential equations is expensive, especially in multi-query scenarios in which the underlying equations have to be solved for different parameters. A comparatively new approach in this context is to combine the good approximation properties of neural networks (for parameter dependence) with the classical finite element method (for discretization). However, instead of considering the solution mapping of the PDE from the parameter space into the FEM-discretized solution space as a purely data-driven regression problem, so-called physically informed regression problems have proven to be useful. In these, the equation residual is minimized during the training of the neural network, i.e., the neural network "learns" the physics underlying the problem. In this paper, we extend this approach to saddle-point and non-linear fluid dynamics problems, respectively, namely stationary Stokes and stationary Navier-Stokes equations. In particular, we propose a modification of the existing approach: Instead of minimizing the plain vanilla equation residual during training, we minimize the equation residual modified by a preconditioner. By analogy with the linear case, this also improves the condition in the present non-linear case. Our numerical examples demonstrate that this approach significantly reduces the training effort and greatly increases accuracy and generalizability. Finally, we show the application of the resulting parameterized model to a related inverse problem.

1. Introduction

Understanding physical systems is crucial in many applications in science and engineering. Many important physical systems are modeled by partial differential equations (PDEs) and then solved using numerical methods. A more recent approach to simulate physical systems is to use neural networks (NNs) as surrogate models, which reduce complexity and speed up evaluation compared to numerical methods. These deep neural networks are universal function approximators [1] and are typically trained with data from simulations or measurements. However, generating data from experiments or simulations is expensive. In addition, a data-driven approach does not or only insufficiently take into account natural laws such as the conservation of energy, mass and momentum, thus generalizes poorly to unseen cases. To cope with these and similar issues, the field of physics-informed machine learning has emerged at the intersection of (numerical) mathematics, machine learning/artificial intelligence, and computational engineering; see, e.g., the surveys [2–5].

* Corresponding author.

https://doi.org/10.1016/j.cam.2025.116663

Received 21 August 2024; Received in revised form 28 February 2025

Available online 5 April 2025

E-mail addresses: franziska.griese@dlr.de (F. Griese), fabian.hoppe@dlr.de (F. Hoppe), alexander.ruettgers@dlr.de (A. Rüttgers), philipp.knechtges@dlr.de (P. Knechtges).

^{0377-0427/© 2025} The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Among the numerous approaches to utilize the power of neural networks for solving differential equations, in particular so-called physics-informed neural networks (PINNs), first introduced in [6,7] and brought to broad attention by [8], gained a lot of attention in the last years. PINNs embed the residual of the respective differential equations into the loss function of a neural network and compute the required derivatives exactly by automatic differentiation [9]. Hence, this approach is often considered to be able to avoid direct discretization; in particular, it is clearly a mesh-free approach. Since training exploits the underlying physics, encoded in the loss function, the PINN approach is able to solve a differential equation without the need for simulated or measured data. However, as higher order derivatives of the network in the loss function lead to a high learning complexity, and additional loss terms that incorporate boundary conditions lead to a multi-objective optimization problem (at least in the vanilla PINN approach), training PINNs is known to be potentially challenging; see, e.g., [10]. For further literature dealing with PINNs, extensions thereof, or similar techniques such as the Deep Ritz method, the Deep Nitsche method, and weak adversarial networks — all of these methods have in common that they do not explicitly apply discretization — we refer to the aforementioned survey papers.

In the following, we focus our literature overview on physics-informed neural networks that make explicit use of discretization, at least to a certain degree. The so-called variational PINNs (vPINNs), a Petrov-Galerkin-like approach, in which a neural network serves as trial function that is tested with a set of finitely many test functions, are introduced in [11]. In the VarNet approach [12], specifically finite element functions are used as test functions, and in h_p -vPINNs [13] essentially h_p -finite element functions are used as test functions. In contrast, the Galerkin neural network approach [14] extends the classical Galerkin procedure, that is usually bound to trial and test functions coming from linear spaces, to trial and test functions both coming from non-linear sets of functions parameterized by neural networks. PhyGeoNet [15] maps the underlying problem to a rectangular reference domain and then employs PINNs loss evaluated on a rectangular grid of points together with convolutional neural networks. Extending this idea to less structured meshes, physics-informed graph neural Galerkin networks [16] combine graph convolutional NNs with a Galerkin finite element method (FEM) discretization. In [17], finite element discretization is considered and the proposed network architecture mimics subsequent refinement steps of the underlying mesh; moreover, a block Jacobi preconditioner is applied. In [18] a somehow hybrid combination of PINNs and FEM-discretization is combined with preconditioning and applied to several model problems, including stationary Navier-Stokes equations (with homogeneous boundary conditions): while the network still approximates the non-discretized solution, it is only evaluated at grid points of a Lagrange FEM-discretization, and the loss functions for the training is constituted in terms of the left-preconditioned FEM-residual. Nonlinear problems are treated by iteratively solving linear problems. A model problem very similar to the one under consideration in the present paper is dealt with in [19] using classical PINNs again. A so-called time-stepping-oriented training procedure, inspired by an control-theoretic ansatz to improve condition number of Jacobi matrices, is employed and allows to deal with 3d problems up to Reynolds number 5,000. We see these results as an encouraging sign that investigating the use of classical numerical techniques and "tricks" is promising in the context of physics-informed ML.

Recently, the use of PINNs has also been proposed to speed up certain element-local computations within a finite element method [20], and neural networks have been used to correct coarse level FEM-solutions on a finer level; see, e.g., [21–23]. Finally, a combination of Isogeometric Analysis with classical PINNs has been proposed in [24]. At the end of this overview, let us briefly comment also on two methods with a similar title that are only slightly related to the present context. In [25], the differential equation is discretized using classical FEM, while a neural network is used to parameterize an unknown parameter function in the differential equation. In [26], the finite element method is cast into a neural network. The domain-decomposition-based (nonlinear) preconditioning of the training of physics-informed neural networks developed in [27] refers to layer-parallel decomposition of the respective neural networks.

Finally, [28,29] combine neural networks and finite elements directly: the parametric differential equation under consideration is discretized a priori by finite elements, and a neural network is trained to map the parameters (of the differential equation) to the coefficients of the corresponding discrete solution (w.r.t. the chosen finite element basis). The loss function is given by the residual of the FEM-discretized differential equation, making this approach data-free and physics-informed. In [28] the applicability to 1D-convection–diffusion problems, as well as to a truss problem, and in [29] to further elasticity problems, such as the Euler–Bernoulli beam or the Reissner–Mindlin plate, has been demonstrated. For a discussion of the advantages of this approach compared to PINNs, in particular the simplified handling of boundary conditions or irregularly shaped domains, we refer to [28]. While [28,29] are concerned with numerical parameters, an extension to a still fully discretized, but more operator-learning-like scenario has been presented in [30] in the context of a time-dependent linear heat equation. Among the vast amount of literature on operator learning [31], another popular method to address parametric problems, we mention exemplarily the recent contribution on hypersonic flows [32].

In our opinion, combining the mathematically proven and quantified ability of FEMs to solve differential equations with the strong approximation capabilities of neural networks to deal with parameter dependencies, should be considered as a simple, but highly appealing idea. In the present work, we thus stick to this approach and extend it in two main directions:

- 1. First, we consider stationary 2D Stokes and Navier–Stokes equations, i.e., equations of type and characteristics that, given their inherent saddle-point structure, are even for classical FEM quite challenging.
- Second, we modify the physics-informed loss function by adding a preconditioner, from left and right, directly to the (nonlinear) residual. In conjunction with L-BFGS, this significantly speeds up training and finally improves accuracy of the obtained solutions.

It should be noted that preconditioning from *both* sides and solving the nonlinear problem all-at-once instead of solving a sequence of linear problems distinguishes our approach from the preconditioning strategy proposed in [18]. Moreover, although the model problem in [19] is very similar to ours, their ansatz for preconditioning is completely different from ours.

Regarding the existing literature on physics-informed machine learning for fluid dynamics problems, we refer to the aforementioned survey papers on discretization-free approaches, and only comment on techniques that explicitly rely on discretization. The neural network multigrid solver [21] already mentioned above has been applied to Navier–Stokes equations. DiscretizationNet [33] combines finite volume discretization for the instationary Navier–Stokes equations together with an iterative application of an encoder–decoder convolutional neural network, while [34] employ physics-informed loss evaluated on a grid together with a 3D U-Net architecture for the neural network. Finally, we mention that beside these machine learning-based approaches, a large number of classical model reduction techniques have been developed over the last decades, especially for applications in fluid dynamics; see, e.g., [35].

Structure of the paper

In Section 2 we introduce the proposed method; first, we describe the underlying setting and the idea in general, and then we discuss this approach in more detail on behalf of our two concrete model problems in Sections 2.3 and 2.4. Numerical examples illustrating the advantages of our method are given in Section 3, and a prototypical application to probabilistic inverse problems is presented in Section 4. We conclude the paper with a brief summary and outlook on future work in Section 5. Finally, the Appendix contains the proof of a fact that can be understood as a heuristic motivation for the chosen approach.

2. Preconditioned loss functions for physics-informed, FEM-based neural networks

In this section, we introduce our proposed method and briefly discuss its realization, both first on an general, rather abstract level. After that, we provide some more concrete details regarding its realization and implementation on the basis of two exemplary, but non-trivial, model problems.

2.1. Problem statement

In the following, we will consider problems that can be written in the following abstract variational form:

Seek
$$u \in V$$
 s.t. $R(u, v) = 0 \quad \forall v \in V$, (1)

where *V* denotes a suitable vector space¹ and $R: V \times V \to \mathbb{R}$ is continuous and linear in the second argument, but not necessarily bi- or trilinear. In fact, we will even consider the following parameterized version of (1):

Seek
$$u(\lambda) \in V$$
 s.t. $R_1(u(\lambda), v) = 0 \quad \forall v \in V$, (2)

where $\lambda \in \Lambda \subset \mathbb{R}^p$ denotes a vector of parameters from a parameter range Λ and each $R_{\lambda} : V \times V \to \mathbb{R}$ satisfies the assumptions as stated for (1). As the notation " $u(\lambda)$ " indicates, we will *assume* in the following that for each $\lambda \in \Lambda$ there is a unique $u_{\lambda} \in V$ that satisfies (2) which makes the map $\lambda \mapsto u(\lambda)$ a well-defined map $\Lambda \to V$.

Let now $V_h \subset V$ be a finite dimensional subspace, e.g., resulting from a conforming finite element discretization of V; the index h indicates the discretization parameter, e.g., the mesh size used for finite element discretization. The so-called Galerkin discretization of our problem (2) is given by:

Seek
$$u_h(\lambda) \in V_h$$
 s.t. $R_{\lambda}(u_h(\lambda), v_h) = 0 \quad \forall v_h \in V_h.$ (3)

Again, we will assume without further notice that also this discrete problem is well-defined in the sense that for each $\lambda \in \Lambda$ there is a unique $u_h(\lambda) \in V_h$ satisfying (2). To obtain a more familiar notation, one may choose a basis $v_h^i \in V_h$, $i = 1, ..., N_h$, of V_h such that $u_h(\lambda) = \sum_{i=1}^{N_h} \mathbf{u}_h^i(\lambda) v_h^i$ is satisfied with a (parameter-dependent) coefficient vector $\mathbf{u}_h(\lambda) \in \mathbb{R}^{N_h}$. Then, it suffices to test (3) with $v_h = v_h^i$ for $i = 1, ..., N_h$ and one obtains a system of N_h equations for $\mathbf{u}_h(\lambda)$:

Seek
$$\mathbf{u}_h(\lambda) \in \mathbb{R}^{N_h}$$
 s.t. $\mathbf{R}_{\lambda,h}(\mathbf{u}_h(\lambda)) := \left(R_\lambda\left(u_h(\lambda), v_h^i\right)\right)_{i=1,\dots,N_h} = \mathbf{0}.$ (4)

2.2. Proposed method

Since (4) is a possibly nonlinear system of equations, usually consisting of a high number of variables, solving (4) for several choices of λ may become costly. To deal with this issue, we propose a method that makes use of the strong approximation capabilities of neural networks. More precisely, we want to approximate the map $\lambda \mapsto \mathbf{u}_h(\lambda)$ solving (4) by a neural network $\mathbf{u}_{h,NN}^{\theta} : \mathbb{R}^{\rho} \to \mathbb{R}^{N_h}$; hereby, θ denotes the network parameters of the neural network. Training this neural network — or in other words: finding the right parameters θ that make $\mathbf{u}_{h,NN}^{\theta}$ a hopefully good approximation of $\lambda \mapsto \mathbf{u}_h(\lambda)$ — requires to determine a suitable so-called loss function. Previous work [28] has suggested to use the ℓ^2 -norm of the discrete equation residual, i.e., to determine θ = $\operatorname{argmin}_{\theta} \mathcal{L}$ of

$$\mathcal{L} := \frac{1}{|A_{\text{train}}|} \sum_{\lambda \in A_{\text{train}}} \|\mathbf{R}_{\lambda,h}(\mathbf{u}_{h,\text{NN}}^{\theta}(\lambda))\|_{\ell^{2}}^{2},$$
(5)

 $^{^{1}}$ V is usually a space of functions of a certain regularity, e.g., a Sobolev space.



Fig. 1. Schematic of FEM-based neural networks.

where $\Lambda_{\text{train}} \subset \Lambda$ denotes a finite subset used as training data. In this paper we propose a slightly different approach, denoting our NN output as $\tilde{\mathbf{u}}_{\theta,NN}^{\theta}(\lambda)$, we determine $\theta = \operatorname{argmin}_{\theta} \mathcal{L}_{\mathbf{P}_{h}}$ via

$$\mathcal{L}_{\mathbf{P}_{\mathbf{h}}} := \frac{1}{|\Lambda_{\text{train}}|} \sum_{\lambda \in \Lambda_{\text{train}}} \|\mathbf{P}_{\lambda,h}^{L} \mathbf{R}_{\lambda,h} (\mathbf{P}_{\lambda,h}^{R} \tilde{\mathbf{u}}_{h,\text{NN}}^{\theta}(\lambda))\|_{\ell^{2}}^{2}, \tag{6}$$

with $\mathbf{P}_{\lambda,h}^{L}, \mathbf{P}_{\lambda,h}^{R} \in \mathbb{R}^{N_{h} \times N_{h}}$ being the left, and respectively right, preconditioners. The desired $\mathbf{u}_{h,NN}^{\theta}(\lambda)$ can then be reobtained via $\mathbf{u}_{h,NN}^{\theta}(\lambda) = \mathbf{P}_{\lambda,h}^{R} \tilde{\mathbf{u}}_{h,NN}^{\theta}(\lambda)$. Of course, such a preconditioner must be chosen carefully taking into account the problem characteristics. As already mentioned in the introduction, we directly apply preconditioning the possibly nonlinear residual, which is in contrast to [18] where nonlinear problems are treated by solving linear problems in iterative manner. Moreover, unlike in [18] we precondition from both sides.

Fig. 1 shows a schematic illustration of how FEM-based neural networks can be implemented in practice on a general level. The neural network takes the parameter vector λ as input and produces $\tilde{\mathbf{u}}_{h,NN}^{\theta}(\lambda)$ as output. During training, the FEM module needs to carry out the calculation of the preconditioned residual $\mathbf{P}_{\lambda,h}^{L}\mathbf{R}_{\lambda,h}\circ\mathbf{P}_{\lambda,h}^{R}$, either directly and all-at-once, or by calculating/applying the residual and the preconditioners separately; from these quantities the preconditioned loss function $\mathcal{L}_{\mathbf{P}_{h}}$ is computed. During inference, the FEM module only needs to provide the right-preconditioner $\mathbf{P}_{\lambda,h}^{R}$ in order to allow the computation of $\mathbf{u}_{h,NN}^{\theta}(\lambda)$ from the NN output $\tilde{\mathbf{u}}_{h,NN}^{\theta}(\lambda)$.

To optimize the loss function w.r.t. the neural network parameters θ , usually variants of the gradient descent method, such as the Adam optimizer [36], or quasi-Newton methods, such as the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm [37] are employed. These methods require the gradient of the loss function with respect to the network parameters, which is calculated by using automatic differentiation (AD) [9]; thus all computations in the FEM Module, such as, e.g., setting up stiffness and/or mass matrices etc., must be implemented automatically differentiable. In this paper we will solely employ L-BFGS. Like most optimizers popular in deep learning, L-BFGS only requires gradients of the objective function, i.e., first-order information, and thus avoids the expensive computation of second-order derivatives by AD, which would be necessary for second-order methods as, e.g., Newton's method. Nevertheless, as a quasi-Newton method, L-BFGS lies somewhere between first- and second-order methods as it is able to exploit curvature information by successively building an approximation to the inverse Hessian matrix during the iteration. The associated hopes for faster convergence have been confirmed, e.g., in the case of classic PINNs [8].

Our desire of preconditioning the problem is also largely driven by this choice of optimizer. Since [38] it is known that BFGS for a quadratic loss function and using perfect line search yields the same iterates as a CG algorithm. Since the CG algorithm is known to have a convergence rate that is highly dependent on the condition number of the Hessian, it seems desirable to precondition the optimization problem. In fact, if $\mathbf{R}_{\lambda,h}(\mathbf{P}_{\lambda,h}^{P}\tilde{\mathbf{u}}_{h,NN}^{O}(\lambda)) = 0$ (or ≈ 0 , respectively) for $\lambda \in \Lambda_{\text{train}}$, it holds that

$$\nabla^2_{\theta} \mathcal{L}_{\mathbf{P}_h}(\theta) = \frac{2}{|\Lambda_{\text{train}}|} \sum_{\lambda} H_{\lambda}^T H_{\lambda}$$

(or that $\nabla^2_{\theta} \mathcal{L}_{\mathbf{P}_h}(\theta)$ is dominated by this term, respectively) with

$$H_{\lambda} = \mathbf{P}_{\lambda,h}^{L} \cdot [D_{\mathbf{u}} \mathbf{R}_{\lambda,h}] (\mathbf{P}_{\lambda,h}^{R} \tilde{\mathbf{u}}_{h,\text{NN}}^{\theta}(\lambda)) \cdot \mathbf{P}_{\lambda,h}^{R} \cdot D_{\theta} \tilde{\mathbf{u}}_{h,\text{NN}}^{\theta}(\lambda).$$

Hence, we hope to improve the conditioning of the Hessian of $\mathcal{L}_{\mathbf{P}_h}$ at some minimizer θ by choosing $\mathbf{P}_{\lambda,h}^L$, $\mathbf{P}_{\lambda,h}^R$ in such a way that the condition number of $\mathbf{P}_{\lambda,h}^L \cdot [D_{\mathbf{u}} \mathbf{R}_{\lambda,h}](\mathbf{P}_{\lambda,h}^R \tilde{\mathbf{u}}_{h,NN}^\theta(\lambda)) \cdot \mathbf{P}_{\lambda,h}^R$ becomes relatively small. Even if these considerations by no means constitute a strict proof, we consider them to be a conclusive motivation for our approach, and our numerical observations in Section 3 seem to support this.



Fig. 2. Coarse triangular mesh around the NACA 0012 airfoil.

2.3. Example 1: Stokes equations in 2D

In the following, we show some more concrete details regarding the realization and implementation of the proposed FEM-based neural network for solving the Stokes flow around an airfoil with a parameterizable angle of attack. We consider the problem in a 2-dimensional Lipschitz-bounded domain Ω , where the Stokes problem consists of finding two functions, the velocity $u(\lambda) : \Omega \to \mathbb{R}^2$ and the pressure $p(\lambda) : \Omega \to \mathbb{R}$ such that

$$-\Delta u(\lambda) + \nabla p(\lambda) = f_{\lambda},\tag{7}$$

$$\nabla \cdot u(\lambda) = g_{\lambda},\tag{8}$$

with functions $f_{\lambda} : \Omega \to \mathbb{R}^2$, $g_{\lambda} : \Omega \to \mathbb{R}$ and suitable boundary conditions, which will be introduced precisely below. Eq. (7) represents the conservation of momentum, Eq. (8) the conservation of mass. Of course, depending on the regularity of the formulation one seeks, these equations can be understood in a strong sense, almost everywhere, or, if f, g become distributions, in a distributional sense.

For the sake of simplicity we will focus in the following on $f \in L^2(\Omega)^2$ and $g \in L^2(\Omega)$, and choose function spaces $X = \{x \in H^1(\Omega)^2 | x|_{\partial \Omega_D} = 0\}$ for velocity and $M = L^2(\Omega)$ for pressure. Here, $\partial \Omega_D$ is supposed to signify the Dirichlet boundary, which shall not comprise the entire boundary $\partial \Omega$. For the time being, we impose zero as the Dirichlet boundary condition for the velocity, which will be alleviated later on.

A weak formulation of the system (7)–(8) with a conforming Galerkin ansatz space $V_h := X_h \times M_h \subset X \times M$ then reads as follows:

Seek
$$u_h \in X_h$$
 and $p_h \in M_h$ such that
 $a(u_h, v_h) + b(v_h, p_h) = f(v_h) \quad \forall v_h \in X_h,$
 $b(u_h, q_h) = g(q_h) \quad \forall q_h \in M_h,$
(9)

with the bilinear forms $a(u_h, v_h) := \int_{\Omega} \nabla u_h : \nabla v_h$ and $b(v_h, p_h) := -\int_{\Omega} p_h \nabla \cdot v_h$. The two functionals f and g are defined as $f(v_h) := \int_{\Omega} v_h f_\lambda$ and $g(q_h) := -\int_{\Omega} q_h g_\lambda$, respectively. R_λ is accordingly defined as the residual of Eq. (9). Note that the fully continuous variant, that is, setting $X_h = X$ and $M_h = M$, defines a well-posed problem [39, Prop. 4.7]. However, to obtain a well-posed problem in the discrete setting, the discrete problem must satisfy the Ladyzhenskaya–Babuška–Brezzi (LBB) condition [40–42] and hence, the choice of X_h , M_h must be handled with care. We use Taylor–Hood mixed finite elements, i.e., continuous, piecewise linear ansatz functions for the pressure and continuous, piecewise quadratic ansatz functions for the velocity. With this choice of V_h , the LBB condition for (9) is satisfied and we obtain a well-posed discrete problem [39, Lemma 4.23].

 $\begin{aligned} & \text{LBB condition for (9) is satisfied and we obtain a well-posed discrete problem [39, Lemma 4.23].} \\ & \text{Let } \{v_h^i\}_{1 \le i \le N_u}, \{q_h^i\}_{1 \le i \le N_p} \text{ be a basis for } X_h, M_h, \text{ with } N_u, N_p \text{ the respective dimension of the corresponding subspace. When defining } \mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_{N_u})^T \text{ with } u_h = \sum_{i=1}^{N_u} \mathbf{u}_i v_h^i \in X_h \text{ and } \mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_{N_p}) \text{ with } p_h = \sum_{i=1}^{N_p} \mathbf{p}_i q_h^i \in M_h, \text{ we obtain the residual} \end{aligned}$

$$\mathbf{R}_{\lambda,h}(\mathbf{u},\mathbf{p}) = \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix},\tag{10}$$

with $\mathbf{A}_{i,j} = a(v_h^j, v_h^i) \in \mathbb{R}^{N_u, N_u}$, $\mathbf{B}_{k,i} = b(v_h^i, q_h^k) \in \mathbb{R}^{N_p, N_u}$, $\mathbf{f}_i = f(v_h^i) \in \mathbb{R}^{N_u}$ and $\mathbf{g}_k = g(q_h^k) \in \mathbb{R}^{N_p}$. The matrix **A** is symmetric positive definite and the system matrix in (10) symmetric, indefinite, and exhibits a typical saddle point structure. For such kind of linear systems, a basic block diagonal preconditioner is given by $\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & -\mathbf{S} \end{bmatrix}^{-1}$, where $\mathbf{S} = -\mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T$ is the so-called Schur complement; see, e.g., [43]. Since **A** and $-\mathbf{S}$ are symmetric positive definite, we can use their Cholesky decompositions $\mathbf{A} = \mathbf{L}\mathbf{L}^T$ and $-\mathbf{S} = \mathbf{M}\mathbf{M}^T$ to define our two preconditioners by

$$\mathbf{P}_{\lambda,h}^{L} = \begin{bmatrix} \mathbf{L}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^{-1} \end{bmatrix}, \qquad \mathbf{P}_{\lambda,h}^{R} = \begin{bmatrix} \mathbf{L}^{T} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^{T} \end{bmatrix}.$$
(11)

Hence, we obtain the preconditioned residual introduced in Section 2 as follows:

F 1 3

1 77 5 7

$$\mathbf{P}_{\lambda,h}^{L}\mathbf{R}_{\lambda,h}(\mathbf{u},\mathbf{p}) = \begin{bmatrix} \mathbf{L}^{T} & \mathbf{L}^{-1}\mathbf{B}^{T} \\ \mathbf{M}^{-1}\mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} - \begin{bmatrix} \mathbf{L}^{-1}\mathbf{f} \\ \mathbf{M}^{-1}\mathbf{g} \end{bmatrix} \\
= \begin{bmatrix} \mathbf{I} & \mathbf{L}^{-1}\mathbf{B}^{T}\mathbf{M}^{-T} \\ \mathbf{M}^{-1}\mathbf{B}\mathbf{L}^{-T} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{L}^{T}\mathbf{u} \\ \mathbf{M}^{T}\mathbf{p} \end{bmatrix} - \begin{bmatrix} \mathbf{L}^{-1}\mathbf{f} \\ \mathbf{M}^{-1}\mathbf{g} \end{bmatrix}, \\
= \begin{bmatrix} \mathbf{I} & \mathbf{L}^{-1}\mathbf{B}^{T}\mathbf{M}^{-T} \\ \mathbf{M}^{-1}\mathbf{B}\mathbf{L}^{-T} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{p}} \end{bmatrix} - \begin{bmatrix} \mathbf{L}^{-1}\mathbf{f} \\ \mathbf{M}^{-1}\mathbf{g} \end{bmatrix}, \tag{12}$$

where **I** is the identity matrix and $\tilde{\mathbf{u}} := \mathbf{L}^T \mathbf{u}$, $\tilde{\mathbf{p}} := \mathbf{M}^T \mathbf{p}$, i.e., $\begin{bmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{p}} \end{bmatrix} = \mathbf{P}_{\lambda,h}^R \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}$. Lemma 2.1 in [44] shows that the conditioning of the plain saddle point problem (10) is bounded from below by the conditioning

Lemma 2.1 in [44] shows that the conditioning of the plain saddle point problem (10) is bounded from below by the conditioning of the discretization of the Laplacian. As such, classical FEM theory [39, Theorem 9.14] shows that for a Lagrange finite element, as we have chosen for the velocity discretization, the conditioning scales with h^{-2} . Therefore, it is usually necessary to apply preconditioning to the simple saddle-point problem. In fact, the condition number of the square of the preconditioned stiffness matrix in (12) is bounded independently of the discretization, as shown in the Appendix. There, it is even shown that the number of distinct eigenvalues of this matrix is three. Considering again that BFGS with perfect line search is equivalent to CG for quadratic optimization problems [38], BFGS applied to such an quadratic optimization problem with the squared preconditioned matrix as its Hessian, would converge in at most three iterations². Of course, these properties do not directly transfer to our setting, as the problem under consideration is highly non-linear and non-convex due to the presence of a neural network. Nevertheless, it does not seem unreasonable to assume that applying the preconditioner to the linear part of the problem will lead to at least a gradual improvement of the overall properties in this case as well. In fact, our numerical examples in Section 3 support this.

Let us now turn to the explicit numerical example, where we want to solve the Stokes Eqs. (7)–(8) with source terms $f_{\lambda} = 0$, $g_{\lambda} = 0$, and non-homogeneous Dirichlet boundary conditions. We define the domain Ω as $(0, 5)^2$ with a NACA 0012 airfoil [46] of length one cut out in the center (see Fig. 2). A Dirichlet boundary condition dependent on the angle of attack λ is used at the inflow boundary $\partial \Omega^+ := \{(x, y) \mid x = 0 \text{ or } y = 0\}$, a Neumann boundary condition at the outflow boundary $\partial \Omega^- := \{(x, y) \mid x = 5 \text{ or } y = 5\}$ and a no-slip boundary condition at the airfoil boundary $\partial \Omega^A$:

$$u = [\cos(\lambda \pi / 180), \sin(\lambda \pi / 180)]^T \qquad \text{on } \partial \Omega^+,$$

$$p - \partial_{\nu} u = 0 \qquad \text{on } \partial \Omega^-,$$

$$u = 0 \qquad \text{on } \partial \Omega^A.$$

Here ∂_{ν} signifies the partial derivative in normal direction on the boundary. Note that the homogeneous Neumann boundary condition on the outflow boundary effectively sets the mean outflow pressure to zero.

Fig. 2 shows the triangular mesh used in the experiments. As we expect larger changes closer to the airfoil, the mesh closer to the airfoil is finer than at the outer edge. We use 2×3193 degrees of freedom for the velocity and 839 degrees of freedom for the pressure, which corresponds to 7225 degrees of freedom in total.

Homogeneous Neumann boundary conditions are known to be so-called natural conditions and thus included in (10) by construction, and do not need any special treatment. However, our aforementioned theory assumes homogeneous Dirichlet boundary conditions for the velocity as well. In order to enforce the inhomogeneous boundary conditions, we will implicitly construct a lifting function: Firstly, we assemble all the matrices and residuals oblivious to the Dirichlet boundary conditions. Secondly, we exploit that the chosen basis functions are Lagrangian, i.e., each basis function v_h^i has an associated Lagrangian node $x_i \in \partial \Omega$ such that

$$v_h^i(\mathbf{x}_j) = \delta_{ij} \quad \forall i, j.$$
⁽¹³⁾

The Dirichlet boundary conditions are now enforced on all such $x_j \in \partial \Omega^+ \cup \partial \Omega^A$ by adjusting the right-hand side vector **f** and the stiffness matrix in (10) accordingly: First, we multiply the columns of the stiffness matrix corresponding to the boundary nodes by the specified boundary values and then subtract the result from **f**. Next, we set the corresponding values of **f** to the boundary values. In the stiffness matrix we modify the rows and columns corresponding to the boundary nodes so that the diagonal value is unity and the off-diagonal entries are set to zero. Note that this is done before preconditioning, which is possible because the modifications do not affect the symmetry properties of the system (10).

In the present case, the (preconditioned) stiffness matrix and right hand side vector of the residual (10) are independent of the prediction of the neural network. Thus they need to be calculated only once and can be reused for every loss calculation in the training of the neural network.

It is also noteworthy that, by the aforementioned approach of incorporating the Dirichlet boundary conditions, we obtain in principle a single objective, avoiding the issues of weakly imposing the Dirichlet conditions via an additional objective term. However, we believe that avoiding the multi-objective scenario is not per se an issue, as one could also rewrite our objective with multiple terms. In our view, the more important part is that these potentially additional terms that impose the Dirichlet conditions must be mutually non-conflicting. We explicitly ensure this by erasing parts of the assembled stiffness matrix and adequately modifying the source terms, as mentioned above.

² The situation is analogous to the Schur-preconditioned stable FEM solutions of the Stokes problem, when solved with Krylov subspace methods, such as MINRES, cf. [45, p. 292].



Fig. 3. Loss values per L-BFGS iteration when training FEM-based NNs with preconditioning (pre) and without preconditioning (nopre) for Stokes (S) and Navier–Stokes (NS) equations with an angle of attack $\lambda = 1^{\circ}$ and different viscosities η .

2.4. Example 2: Stationary Navier-Stokes equations in 2D

To avoid redundancies, we focus on the changes that are necessary to extend our approach from the Stokes to the Navier–Stokes equations. Thus, details not discussed further remain the same as in the previous section. The Navier–Stokes equations are as follows:

$$\eta \Delta u(\lambda) + u(\lambda) \cdot \nabla u(\lambda) + \nabla p(\lambda) = f_{\lambda}, \tag{14}$$

$$\nabla \cdot u(\lambda) = g_{\lambda},\tag{15}$$

i.e., compared to the Stokes Eqs. (7)–(8) in particular the non-linear convection term $u(\lambda) \cdot \nabla u(\lambda)$ has been added. Moreover, we have introduced the so-called kinematic viscosity $\eta > 0$ as additional parameter. Discretization leads to

Seek
$$u_h \in X_h$$
 and $p_h \in M_h$ such that
 $a(u_h, v_h) + c(z_h; u_h, v_h) + b(v_h, p_h) = f(v_h) \quad \forall v_h \in X_h,$
 $b(u_h, q_h) = g(q_h) \quad \forall q_h \in M_h,$
(16)

with *b* as before, *a* as before but multiplied by η , and the trilinear form $c(z_h; u_h, v_h) = \int_{\Omega} (z_h \cdot \nabla u_h) \cdot v_h$ originating from the nonlinear term. We again use Taylor–Hood elements for the discretization in the non-linear case, as these are already necessary for a stable solution in the linear case. To formulate the corresponding nonlinear algebraic equation, we utilize Lagrangian basis functions as in (10) and obtain the discrete residual

$$\mathbf{R}_{\lambda,h}(\mathbf{u},\mathbf{p}) = \begin{bmatrix} \mathbf{A} + \mathbf{C}(\mathbf{u}) & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix},\tag{17}$$

with $\mathbf{C}(\mathbf{u})_{i,j} = c(u_h; v_h^j, v_h^i) \in \mathbb{R}^{N_u, N_u}$. We use the same preconditioning as for the Stokes problem, which yields the preconditioned system

$$\mathbf{P}_{\lambda,h}\mathbf{R}_{\lambda,h}(\mathbf{u},\mathbf{p}) = \begin{bmatrix} \mathbf{I} + \mathbf{L}^{-1}\mathbf{C}(\mathbf{u})\mathbf{L}^{-T} & \mathbf{L}^{-1}\mathbf{B}^{T}\mathbf{M}^{-T} \\ \mathbf{M}^{-1}\mathbf{B}\mathbf{L}^{-T} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{L}^{T}\mathbf{u} \\ \mathbf{M}^{T}\mathbf{p} \end{bmatrix} - \begin{bmatrix} \mathbf{L}^{-1}\mathbf{f} \\ \mathbf{M}^{-1}\mathbf{g} \end{bmatrix}.$$
(18)

Again we use the NN to directly predict $\tilde{\mathbf{u}} := \mathbf{L}^T \mathbf{u}$ and $\tilde{\mathbf{p}} := \mathbf{M}^T \mathbf{p}$. In our numerical examples we will solve the Navier–Stokes Eqs. (14)–(15) on the same geometry, mesh, and boundary conditions as in the previous section.

Since the Navier–Stokes problem is nonlinear, the stiffness matrix depends on the solution, i.e., the prediction of the NN in our case. This means that unlike in (12), where the system matrix and the right hand side could be calculated once in advance, we now have to recalculate parts of the residual (18) in every L-BFGS iteration during the training of the network. More precisely, this applies to the convection term C(u) and the right hand side vector **f** due to the inclusion of Dirichlet boundary conditions as we will describe below. We note that

$$\mathbf{C}_{ij}(\mathbf{u}) = \int_{\Omega} \sum_{k} \mathbf{u}_{k} v_{h}^{k} \cdot \nabla v_{h}^{j} \cdot v_{h}^{i} = \sum_{k} \mathbf{u}_{k} \underbrace{\int_{\Omega} v_{h}^{k} \cdot \nabla v_{h}^{j} \cdot v_{h}^{i}}_{=:\tilde{\mathbf{C}}_{i,i,k}},$$

i.e., $C(\mathbf{u})$ can be calculated by a contraction of the tensor $\mathbf{u} = \mathbf{L}^{-T} \tilde{\mathbf{u}}$ (or $\mathbf{u} = \tilde{\mathbf{u}}$ in the case of no preconditioning) with a 3D-tensor $\tilde{\mathbf{C}}$ that is independent of \mathbf{u} and thus can be precomputed and reused. Since $\tilde{\mathbf{C}}$ is sparse, this can be done in a memory-efficient manner.

Let us now briefly comment on enforcing Dirichlet boundary conditions in the residual (17). Except for the nonlinear part, $C(\mathbf{u})$, this can be handled in the same way as for the Stokes equations. The only difference is that instead of $C(\mathbf{u})$ we actually have to compute $C(\hat{\mathbf{u}})$ with the slightly modified argument $\hat{\mathbf{u}} = \hat{\mathbf{u}}(\mathbf{u})$ in which the respective boundary conditions are already enforced at all



Fig. 4. Eigenvalues of the Hessian of the loss function, evaluated after training the FEM-based NNs with preconditioning (pre) and without preconditioning (nopre) for Stokes equations with an angle of attack $\lambda = 1^{\circ}$. Eigenvalues are shown in gray, a kernel density estimate of the respective distribution functions in blue and orange. Note that the *y*-axis employs symmetric log axis scaling, i.e. logarithmic scaling with a linear scaling exception for values within [-2,2].

Lagrange nodes on the Dirichlet boundary, i.e.,

$$\hat{\mathbf{u}}_{i} := \begin{cases} \mathbf{u}_{i} & \text{if } x_{i} \notin \partial \Omega^{+} \cup \partial \Omega^{A}, \\ u_{\text{BC}}(x_{i}) & \text{if } x_{i} \in \partial \Omega^{+} \cup \partial \Omega^{A}, \end{cases}$$
(19)

where $u_{BC}(x_i)$ denotes the value of *u* prescribed at the boundary node x_i . As for the Stokes equations, the boundary conditions are included into (17) before preconditioning. With these changes, the training of the FEM-based NN can be carried out as described above.

3. Numerical results

In the following, we show results of the proposed FEM-based neural networks for solving the Stokes and Navier–Stokes flow around an airfoil with a parameterizable angle of attack as described in Section 2. In particular, we will focus on the comparison between using a preconditioner and not using a preconditioner.

The respective neural networks and an FEM module as described in Sections 2, 2.3, and 2.4 have been implemented in PyTorch [47]. Unless stated otherwise, the neural networks are plain vanilla feed forward networks consisting of five hidden layers with 50 neurons each and an output layer whose size is equal to the number of degrees of freedom of the discretization. When testing various hyperparameters, we found that small changes in the number of layers and number of neurons per layer did not have a major effect on the results. However, it turned out that a non-linear activation function which is not bounded is of great advantage to avoid local minima during optimization, which is why we use SELU [48] as the activation function except for the last, linear layer. To facilitate efficient training, we set the network weights with Xavier's normal initialization [49] scaled by 3/4. For the training, we use PyTorch's implementation³ of L-BFGS. To make the results more comparable with "classical" implementations, we choose max_iter=1 to ensure that only one L-BFGS step is performed per iteration. Furthermore, we choose a line search within L-BFGS that seeks a solution that fulfills the strong Wolfe condition.

In the following, the term "error" does not refer to the discretization error, but to the error between our NN-based solution and a pure FEM-based reference solution that uses the same mesh and Taylor–Hood discretization. The latter was computed with FEniCSx [50,51].

3.1. The non-parametric case

We start with experiments with only a single angle of attack as training data and use a linear NN consisting of only one input and one output layer. Considering this non-parametric and thus highly underdetermined case allows us to evaluate the benefits of preconditioning separately from other effects.

Both the preconditioned and the non-preconditioned NNs are trained and evaluated for a single angle of attack $\lambda = 1^{\circ}$ and 40000 L-BFGS iterations, for each: the Stokes problem and the Navier–Stokes equation. In the latter case, the solutions are computed with viscosities $\eta = 1$ and $\eta = 0.1$. Fig. 3 shows that in all these cases preconditioning significantly speeds up the training in terms of a drastically reduced number of iterations. For the preconditioned cases we can even observe a stagnation of the residual at levels that are probably related to numerical accuracy; this is probably related to a similar reason as the effect described in the paragraph below (12) in Section 2.3. Fig. 5 shows for the Navier–Stokes equation with $\eta = 1.0$ that training is also much faster in terms of a single iteration that is due to the few additional operations related to preconditioning.

³ https://pytorch.org/docs/stable/generated/torch.optim.LBFGS.html [Accessed August 21, 2024].



Fig. 5. Loss values over time, in logarithmic scaling, when training FEM-based NNs with preconditioning (pre) and without preconditioning (nopre) for the Navier–Stokes (NS) equations with an angle of attack $\lambda = 1^{\circ}$ and viscosity $\eta = 1.0$. The initial construction times of the system matrices in (17) and (18) are shown in cyan, the training times of the iterations in purple.

Table 1

Relative errors of the FEM-based neural network trained for Stokes (S) and Navier–Stokes (NS) using the non-preconditioned loss (5) (nopre) and the preconditioned loss (6) (pre) with an angle of attack of 1° and viscosity η .

	Pre.	η	L ² -Error			L^{∞} -Error		
			u _x	u _y	р	u _x	u _y	р
S	nopre	-	6.0801e-03	2.5899e-02	8.6753e-02	7.6672e-02	1.1719e-01	9.9965e-01
	pre	-	7.3656e-08	1.1760e-07	1.3784e-07	1.5618e-07	1.6593e-07	4.1410e-06
NS	nopre	1.0	7.3695e-03	4.4370e-02	1.5622e-01	8.3582e-02	1.5751e-01	9.9945e-01
	nopre	0.1	8.8126e-01	15.1430e-00	24.3703e-00	1.7591e-00	11.3491e-00	2.2046e-00
	pre	1.0	1.4740e-06	8.7947e-06	5.9979e-06	3.1987e-06	9.7846e-06	7.3386e-06
	pre	0.1	1.7492e-05	1.2692e-04	2.0684e-04	4.7161e-05	1.4152e-04	9.0761e-05

Table 1 shows the achieved relative errors in L^2 - and L^{∞} -norms. Again, the benefits of preconditioning are clearly visible in the reduction of the final error by several orders of magnitude. In particular, note that for $\eta = 0.1$ (which corresponds to an increased Reynolds number) a satisfying accuracy could only be achieved with preconditioning. As Fig. 6 illustrates, there is almost no visible difference between the NN-based solutions and the pure FEM-based ones in the end.

To further illustrate the effectiveness of the preconditioning, we computed the eigenvalues of the Hessian at the last iteration of the training of the Stokes equations, once with preconditioning and once without preconditioning. Note that computing the Hessian using the automatic differentiation machinery in PyTorch is possible, though possibly quite costly, and can only be employed for small problem sizes. The eigenvalues, and a kernel density estimate, for the two cases, with and without preconditioning, is plotted in Fig. 4. It is clearly visible that even though our argumentation in Section 2 and the Appendix totally neglected the non-linearity introduced by the neural networks, the conclusion of clustered eigenvalues still holds. As noted earlier, we assume L-BFGS coupled with perfect line search to share the same characteristics as the CG method for a nearly quadratic optimization problem. Given that in our least-squares setting the cluster of zero eigenvalues can be neglected [52], the convergence of the CG algorithm for a quadratic optimization problem is determined by the number of remaining eigenvalue clusters [45]. Of course, this analysis is not exact, given that our problem is made non-linear by the neural network and potentially the advection term in the Navier–Stokes equations, and thus potentially deviates a lot from the aforementioned assumptions. Nonetheless, it clearly demonstrates the mechanisms at work, indicating the effectiveness of classical preconditioning techniques.

That the assumptions made in the design of the preconditioner are only approximately fulfilled, could not be more visible than in the high Reynolds number case, as illustrated in Fig. 3. For $\eta = 0.01$ the advection term clearly dominates, such that the chosen preconditioner, which takes only the diffusive part into account, performs poorly. The latter is clearly indicated by the increased number of iterations. Moreover, for higher Reynolds numbers techniques like, e.g., Streamline Upwind Petrov–Galerkin (SUPG) should be employed to further enhance the numerical properties [45,53].

3.2. The parametric case

In this section, we consider the true parametric problem and train with multiple angles of attack, namely 3, 5, 9, and 17 different λ distributed equidistantly in the range [1°, 45°]. A maximum of 4000 iterations of L-BFGS is used to train the respective networks.

Fig. 7 shows the loss values during training for the different configurations under consideration. Even in this case, preconditioning is able to accelerate the decay of the loss by more than two orders of magnitude for both the Stokes and the Navier–Stokes equations. For the Stokes problem, we also performed the experiments on a finer mesh, which is obtained from the original one by a single uniform refinement and results in 28085 degrees of freedom. The behavior of the losses on the original and the finer mesh can be



Fig. 6. Results for Navier–Stokes equation trained with an angle of attack $\lambda = 1^{\circ}$ and viscosity $\eta = 0.1$. The first row shows the prediction of the FEM-based neural network with preconditioning, the second row the reference solution obtained with the FEM, the third row the absolute error of the prediction.

taken as an indication of mesh independence, which is consistent with our expectations from the Appendix and our discussion in the previous section.

In the following, we will in particular consider the training errors, i.e., the error on the seen training data, as well as generalization errors: more specifically the errors seen on interpolating input data as well as on extrapolating input data. To evaluate the interpolation and extrapolation qualities, we test with unseen angles of attack within the training range, $\lambda \in \{5^\circ, 16.5^\circ, 30^\circ, 40^\circ\}$, and outside the training range, $\lambda \in \{47.5^\circ, 50^\circ, 55^\circ\}$, respectively. Fig. 8 shows the average errors on the training data ("training") as well as the generalization errors ("interpolation" and "extrapolation") in L^2 - and L^∞ -norm for the different sizes of the training set. For clarity, we only report errors of the velocity in *x*-direction as the errors of velocity in *y*-direction and pressure behave similarly. It is clearly visible that preconditioning does not only heavily improve quality of the surrogate model on the training set itself, but also its inter- and extrapolation capabilities; in fact, the errors obtained without preconditioning are far beyond an acceptable range. The increase of the training errors for a growing size of the training set for the preconditioned problems is likely due to the fact that the same network architecture (i.e. the same number of degrees of freedom) is used to solve a increasingly complex problem. As expected, the generalization errors decrease with increasing training data.

Our findings can be summarized as follows: The use of preconditioned loss functions speeds up the training procedure, both for linear and nonlinear problems and, moreover, improves the quality of the resulting surrogate models both in terms of training and generalization errors.

4. Application: Solving an inverse problem related to stationary 2D Navier-Stokes equations

In the following, we demonstrate the relevance of our method by applying it to a prototypical inverse problem related to the stationary 2D Navier–Stokes equations. This problem corresponds to an application scenario in which pressure measurements from two sensors on the airfoil are used to determine the angle of attack. Since measurements in a real-world setting are always noisy, it is crucial to be able to quantify the aleatoric uncertainty coming from this noise. Our goal is to estimate a probability density for the unknown attack angle λ associated with these observations. Therefore, we pursue a probabilistic approach and use a Markov Chain Monte Carlo (MCMC) method in combination with our fully differentiable FEM-based NNs.

Let $\mathbf{y}_1, \dots, \mathbf{y}_m \in \mathbb{R}^2$ be noisy observations of the pressure at two fixed points of the airfoil. We assume that the two pressure sensors are located at 8.18% and 99.94% of the chord on the top of the airfoil; see Fig. 9 for an illustration. Moreover, we assume that the



Fig. 7. Loss values per L-BFGS iteration when training FEM-based NNs for multiple angles of attack $\#\lambda_{\text{train}}$ with preconditioning (pre) and without preconditioning (nopre). For the Stokes problem, in addition to the loss behavior for the original mesh, we also show the behavior for a finer mesh.



Fig. 8. Averaged relative L^2 - and L^{∞} -errors of the variable u_x predicted from FEM-based NNs trained with and without preconditioning for Stokes and Navier–Stokes ($\eta = 1$) with several numbers of angles of attack $\#\lambda_{\text{train}}$ and tested with angles of attack inside (interpolation) and outside (extrapolation) the training range. The dots represent the mean errors, the shadings correspond to the range in which the errors occurred for the given set of testing angles.



Fig. 9. Left: Location of sensors on airfoil. Right: Pressure p at sensor locations for different angles of attack λ predicted by FEM-based NN.



Fig. 10. Left: Observed pressure measurements y_i of two sensors. Right: Corresponding posterior distribution of the angle of attack λ of the airfoil, as generated with the MCMC algorithm.

measurements are independent and randomly distributed according to a normal distribution with standard deviation $\sigma = 0.5$. We thus obtain a log-likelihood that modulo some λ -independent constant *C* is given by

$$\log p(\boldsymbol{y}|\boldsymbol{\lambda}) = C - \sum_{i=1}^{m} \frac{1}{\sigma^2} \|\hat{\mathbf{p}}_{\boldsymbol{\lambda}} - \mathbf{y}_i\|^2.$$
(20)

We deliberately use $\hat{\mathbf{p}}_{\lambda}$ to signify the result of the preconditioned neural network evaluated at just these two measurement points. Encompassing this likelihood with a uniform prior $p(\lambda) \sim \mathcal{U}([1^\circ, 45^\circ])$, we are interested in the resulting posterior distribution $p(\lambda|y) \propto p(y|\lambda)p(\lambda)$.

Given that computations of $\hat{\mathbf{p}}_{\lambda}$ exhibit a nonlinear dependence on λ , we employ a Hamiltonian Monte Carlo (HMC) method [54, 55], more precisely the No-U-Turn Sampler (NUTS). The latter is a gradient-based MCMC method that approximates the desired density function. HMC allows to sample from this posterior by using an approximate Hamiltonian dynamics simulation, which is then corrected by a Metropolis acceptance step. To solve the Hamiltonian differential equation, derivatives of the target posterior density function with respect to λ must be computed. Here, an important advantage of neural network-based surrogate modeling comes into play: the surrogate model is not only differentiable (which may theoretically also be the case with other techniques), but in fact these derivatives can be computed efficiently using AD. We use the implementation of NUTS that comes with the PyTorch-based library Pyro [56].

In order to test this method, we choose the following setup: Fig. 9 (right) shows the forward problem of predicting the pressure at the sensor locations for 17 angles of attack λ that are equidistantly distributed in the range [1°, 45°]. Using this FEM-based NN as surrogate model, we generate 10 artificial and noisy pressure "measurements" by adding Gaussian noise with a standard deviation of 0.5 to the respective pressure predictions evaluated at a reference angle of attack of $\lambda = 5^{\circ}$. The "measurements" are shown in Fig. 10.

Fig. 10 displays the 1000 samples of the posterior distribution, generated by MCMC with 1000 warm-up steps. The generated distribution has an empirical mean of 5.03 and a standard deviation of 0.32. Hence, the combination of MCMC and FEM-based NNs was able to deliver plausible results for this test case.

5. Conclusion

FEM-based/FEM-inspired neural networks are known to combine the best of two worlds: finite elements and neural networks. On the one hand, the (parametric) PDE under consideration is discretized by a method that is well-understood and reliable, the

incorporation of boundary conditions and complex geometries is relatively straightforward. On the other hand, the discretization benefits from the power of neural networks to handle parameter dependence, e.g., by obtaining a cheap-to-evaluate and fully differentiable parameter-to-solution map, which can be advantageous in the context of inverse problems.

In this paper, we have advanced this type of technique by applying it for the first time to PDEs from fluid dynamics that result in a saddle point problem after discretization. In particular, we have introduced a modification of this approach in which preconditioning applied to the residual loss allows to improve both the speed of training and the resulting accuracy. We have demonstrated the relevance of our approach with numerical examples on behalf of a prototypical but realistic model problem from fluid dynamics and a related inverse problem.

Having provided a proof-of-concept for our method in this paper, an interesting goal for future research would be to investigate the possibilities for the inclusion of, e.g., distributed parallel methods or low-cost preconditioners, as they are typical for largescale finite element implementations. Another goal would be to consider other PDEs or PDE systems with different characteristics (including suitable preconditioners). In particular, the extension to time-dependent problems or to flow problems with higher Reynolds numbers that require stabilization shall be mentioned.

Acknowledgments

This research was carried out in part during the project PISA (Physics Inspired AI) by the German Aerospace Center (DLR). The DLR is supported by the Federal Ministry for Economic Affairs and Climate Action (BMWK) on the basis of a decision by the German Bundestag. Last but not least, the authors want to thank the anonymous reviewers for their valuable comments.

Appendix. Mesh-independent conditioning of the FEM layer for the preconditioned Stokes case

In this appendix, we show that contribution of the matrix in (12) to the loss in (6) is well-conditioned, as far as the last layer of the NN is concerned. More precisely, we will show that the square of

$$\mathbf{Z} = \begin{bmatrix} \mathbf{I} & \mathbf{L}^{-1}\mathbf{B}^{T}\mathbf{M}^{-T} \\ \mathbf{M}^{-1}\mathbf{B}\mathbf{L}^{-T} & \mathbf{0} \end{bmatrix}$$
(A.1)

has an eigenvalue distribution that is independent of the mesh resolution. The following derivation is based on [57].

Lemma 1. Consider $\mathbf{Y} = \mathbf{Z}^2$ with \mathbf{Z} as given in (A.1), then \mathbf{Y} has the following annihilating polynomial

$$(\mathbf{Y} - \mathbf{I})\left(\left(\mathbf{Y} - \frac{3}{2}\mathbf{I}\right)^2 - \frac{5}{4}\mathbf{I}\right) = 0.$$

Proof. Introducing $C = M^{-1}BL^{-T}$, first note that Y is given by

$$\mathbf{Y} = \begin{bmatrix} \mathbf{I} + \mathbf{C}^T \mathbf{C} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{C} \mathbf{C}^T \end{bmatrix} \,.$$

The latter can, using the definitions of L, M, S, be simplified to

$$\mathbf{C} \mathbf{C}^{T} = \mathbf{M}^{-1} \mathbf{B} \underbrace{\mathbf{L}^{-T} \mathbf{L}^{-1}}_{=\mathbf{A}^{-1}} \mathbf{B}^{T} \mathbf{M}^{-T}$$
$$= \mathbf{M}^{-1} \underbrace{\mathbf{B} \mathbf{A}^{-1} \mathbf{B}^{T}}_{-\mathbf{S}} \mathbf{M}^{-T}$$
$$= \mathbf{I}.$$

Furthermore, $\mathbf{P} = \mathbf{C}^T \mathbf{C}$ is a projection operator

$$\mathbf{C} \mathbf{P} = \mathbf{C} \mathbf{C}^T \mathbf{C} = \mathbf{C}$$
$$\mathbf{P}^2 = \mathbf{C}^T \mathbf{C} \mathbf{P} = \mathbf{C}^T \mathbf{C} = \mathbf{P}.$$

As such, we obtain

$$\left(\mathbf{Y} - \frac{3}{2}\mathbf{I}\right)^2 = \begin{bmatrix} \left(\mathbf{P} - \frac{1}{2}\mathbf{I}\right)^2 + \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \frac{5}{4}\mathbf{I} \end{bmatrix},$$

which yields

$$\left(\mathbf{Y} - \frac{3}{2}\mathbf{I}\right)^2 - \frac{5}{4}\mathbf{I} = \begin{bmatrix} \mathbf{P} - \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$$

The conclusion then follows easily from the observation that

$$\mathbf{Y} - \mathbf{I} = \begin{bmatrix} \mathbf{P} & \mathbf{C}^T \\ \mathbf{C} & 0 \end{bmatrix} \cdot \Box$$

As an immediate consequence of the last Lemma, the distinct eigenvalues of **Y** are $\left\{1, \frac{3}{2} \pm \frac{\sqrt{5}}{2}\right\}$, thus clearly bounding the condition number of the matrix **Y**.

Data availability

No data was used for the research described in the article.

References

- [1] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, Neural Netw. 2 (5) (1989) 359-366.
- [2] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nat. Rev. Phys. 3 (6) (2021) 422-440.
- [3] S. Cuomo, V.S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics–Informed neural networks: Where we are and what's next, J. Sci. Comput. 92 (3) (2022) 88.
- [4] Z. Hao, S. Liu, Y. Zhang, C. Ying, Y. Feng, H. Su, J. Zhu, Physics-informed machine learning: A survey on problems, methods and applications, 2023.
- [5] D.N. Tanyu, J. Ning, T. Freudenberg, N. Heilenkötter, A. Rademacher, U. Iben, P. Maass, Deep learning methods for partial differential equations and related parameter identification problems, Inverse Problems 39 (10) (2023) 103001.
- [6] M.W.M.G. Dissanayake, N. Phan-Thien, Neural-network-based approximations for solving partial differential equations, Commun. Numer. Methods Eng. 10 (3) (1994) 195–201.
- [7] I.E. Lagaris, A. Likas, D.I. Fotiadis, Artificial neural networks for solving ordinary and partial differential equations, IEEE Trans. Neural Netw. 9 (5) (1998) 987–1000.
- [8] M. Raissi, P. Perdikaris, G.E. Karniadakis, 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.
- [9] A. Griewank, On automatic differentiation and algorithmic linearization, Pesqui. Oper. 34 (2014) 621-645.
- [10] S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, SIAM J. Sci. Comput. 43 (5) (2021) A3055–A3081.
- [11] E. Kharazmi, Z. Zhang, G.E. Karniadakis, Variational physics-informed neural networks for solving partial differential equations, 2019, arXiv preprint arXiv:1912.00873.
- [12] R. Khodayi-Mehr, M. Zavlanos, VarNet: Variational neural networks for the solution of partial differential equations, in: Learning for Dynamics and Control, PMLR, 2020, pp. 298–307.
- [13] E. Kharazmi, Z. Zhang, G.E. Karniadakis, hp-VPINNs: Variational physics-informed neural networks with domain decomposition, Comput. Methods Appl. Mech. Engrg. 374 (2021) 113547.
- [14] M. Ainsworth, J. Dong, Galerkin neural networks: A framework for approximating variational equations with error control, SIAM J. Sci. Comput. 43 (4) (2021) A2474–A2501.
- [15] H. Gao, L. Sun, J. Wang, PhyGeoNet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state PDEs on irregular domain, J. Comput. Phys. 428 (2021) 110079.
- [16] H. Gao, M.J. Zahr, J. Wang, Physics-informed graph neural Galerkin networks: A unified framework for solving PDE-governed forward and inverse problems, Comput. Methods Appl. Mech. Engrg. 390 (2022) 114502.
- [17] C. Uriarte, D. Pardo, Á.J. Omella, A finite element based deep learning solver for parametric PDEs, Comput. Methods Appl. Mech. Engrg. 391 (2022) 114562.
- [18] S. Liu, C. Su, J. Yao, Z. Hao, H. Su, Y. Wu, J. Zhu, Preconditioning for physics-informed neural networks, 2024.
- [19] W. Cao, W. Zhang, An analysis and solution of ill-conditioning in physics-informed neural networks, J. Comput. Phys. 520 (2025) 113494.
- [20] P. Pantidis, M.E. Mobasher, Integrated finite element neural network (I-FENN) for non-local continuum damage mechanics, Comput. Methods Appl. Mech. Engrg. 404 (2023) 115766.
- [21] N. Margenberg, D. Hartmann, C. Lessig, T. Richter, A neural network multigrid solver for the Navier–Stokes equations, J. Comput. Phys. 460 (2022) 110983.
- [22] U. Kapustsin, U. Kaya, T. Richter, A hybrid finite element/neural network solver and its application to the Poisson problem, PAMM 23 (3) (2023) e202300135.
- [23] M. Hintermüller, D. Korolev, A hybrid physics-informed neural network based multiscale solver as a partial differential equation constrained optimization problem, 2024.
- [24] M. Möller, D. Toshniwal, F. van Ruiten, Physics-informed machine learning embedded into isogeometric analysis, in: Mathematics: Key Enabling Technology for Scientific Machine Learning, Platform Wiskunde, Amsterdam, 2021, pp. 57–59.
- [25] S.K. Mitusch, S.W. Funke, M. Kuchta, Hybrid FEM-NN models: Combining artificial neural networks with the finite element method, J. Comput. Phys. 446 (2021) 110651.
- [26] P. Ramuhalli, L. Udpa, S.S. Udpa, Finite-element neural networks for solving differential equations, IEEE Trans. Neural Netw. 16 (6) (2005) 1381–1392.
- [27] A. Kopaničáková, H. Kothari, G.E. Karniadakis, R. Krause, Enhancing training of physics-informed neural networks using domain decomposition-based preconditioning strategies, SIAM J. Sci. Comput. 46 (5) (2024) S46–S67.
- [28] R.E. Meethal, A. Kodakkal, M. Khalil, A. Ghantasala, B. Obst, K. Bletzinger, R. Wüchner, Finite element method-enhanced neural network for forward and inverse problems, Adv. Model. Simul. Eng. Sci. 10 (1) (2023) 6.
- [29] T. Le-Duc, H. Nguyen-Xuan, J. Lee, A finite-element-informed neural network for parametric simulation in structural mechanics, Finite Elem. Anal. Des. 217 (2023) 103904.
- [30] Y. Yamazaki, A. Harandi, M. Muramatsu, A. Viardin, M. Apel, T. Brepols, S. Reese, S. Rezaei, A finite element-based physics-informed operator learning framework for spatiotemporal partial differential equations on arbitrary domains, 2024.
- [31] L. Lu, P. Jin, G. Pang, Z. Zhang, G.E. Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, Nat. Mach. Intell. 3 (3) (2021) 218–229.
- [32] A. Peyvan, V. Kumar, Fusion DeepONet: A data-efficient neural operator for geometry-dependent hypersonic flows on arbitrary grids, 2025.
- [33] R. Ranade, C. Hill, J. Pathak, DiscretizationNet: A machine-learning based solver for Navier–Stokes equations using finite volume discretization, Comput. Methods Appl. Mech. Engrg. 378 (2021) 113722.
- [34] N. Wandel, M. Weinmann, R. Klein, Teaching the incompressible Navier–Stokes equations to fast neural surrogate models in three dimensions, Phys. Fluids 33 (4) (2021) 047117.
- [35] P. Benner, et al. (Eds.), Model Order Reduction, Vols 1-3, De Gruyter, 2022.
- [36] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Yoshua Bengio, Yann LeCun (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings, 2015.

- [37] D.C. Liu, J. Nocedal, On the limited memory BFGS method for large scale optimization, Math. Program. 45 (1) (1989) 503-528.
- [38] L. Nazareth, A relationship between the BFGS and conjugate gradient algorithms and its implications for new algorithms, SIAM J. Numer. Anal. 16 (5) (1979) 794-800.
- [39] A. Ern, J. Guermond, Theory and Practice of Finite Elements, vol. 159, Springer, 2004.
- [40] O.A. Ladyzhenskaya, The Mathematical Theory of Viscous Incompressible Flow, Gordon and Breach Science, 1969.
- [41] I. Babuška, The finite element method with Lagrangian multipliers, Numer. Math. 20 (3) (1973) 179–192.
- [42] F. Brezzi, On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers, R.A.I.R.O. Anal. Numérique 8 (R-2) (1974) 129–151.
- [43] M. Benzi, G. Golub, J. Liesen, Numerical solution of saddle point problems, Acta Numer. 14 (2005) 1–137.
- [44] T. Rusten, R. Winther, A preconditioned iterative method for saddlepoint problems, SIAM J. Matrix Anal. Appl. 13 (3) (1992) 887-904.
- [45] H.C. Elman, D.J. Silvester, A.J. Wathen, Finite Elements and Fast Iterative Solvers, in: Numerical Mathematics and Scientific Computation, Oxford University Press, 2005.
- [46] E.N. Jacobs, K.E. Ward, R.M. Pinkerton, The Characteristics of 78 Related Airfoil Section from Tests in the Variable-Density Wind Tunnel, Technical Report 460, National Advisory Committee for Aeronautics, 1933.
- [47] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems 32, Curran Associates, Inc., 2019, pp. 8024–8035.
- [48] G. Klambauer, T. Unterthiner, A. Mayr, S. Hochreiter, Self-normalizing neural networks, Adv. Neural Inf. Process. Syst. 30 (2017).
- [49] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.
- [50] I.A. Baratta, J.P. Dean, J.S. Dokken, M. Habera, J.S. Hale, C.N. Richardson, M.E. Rognes, M.W. Scroggs, N. Sime, G.N. Wells, DOLFINx: the next generation FEniCS problem solving environment, 2023, preprint.
- [51] M.S. Alnaes, A. Logg, K.B. Ølgaard, M.E. Rognes, G.N. Wells, Unified form language: A domain-specific language for weak formulations of partial differential equations, ACM Trans. Math. Softw. 40 (2014).
- [52] W.J. Kammerer, M.Z. Nashed, On the convergence of the conjugate gradient method for singular linear operator equations, SIAM J. Numer. Anal. 9 (1) (1972) 165–181.
- [53] A.N. Brooks, T.J.R. Hughes, Streamline upwind/Petrov–Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier–Stokes equations, Comput. Methods Appl. Mech. Engrg. 32 (1) (1982) 199–259.
- [54] M. Betancourt, M. Girolami, Hamiltonian Monte Carlo for hierarchical models, Curr. Trends Bayesian Methodol. Appl. 79 (30) (2015) 2-4.
- [55] R. Neal, MCMC using Hamiltonian dynamics, in: Handbook of Markov Chain Monte Carlo, 2012.
- [56] E. Bingham, J.P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, N.D. Goodman, Pyro: Deep universal probabilistic programming, J. Mach. Learn. Res. (2018).
- [57] M.F. Murphy, G.H. Golub, A.J. Wathen, A note on preconditioning for indefinite linear systems, SIAM J. Sci. Comput. 21 (6) (2000) 1969–1972.