Extremization to fine tune physics informed neural networks for solving boundary value problems

Thiruthummal, A. A., Shelyag, S. & Kim, E.

Published PDF deposited in Coventry University's Repository

Original citation:

Thiruthummal, AA, Shelyag, S & Kim, E 2024, 'Extremization to fine tune physics informed neural networks for solving boundary value problems', Communications in Nonlinear Science and Numerical Simulation, vol. 137, 108129. https://dx.doi.org/10.1016/j.cnsns.2024.108129

DOI 10.1016/j.cnsns.2024.108129 ISSN 1007-5704 ESSN 1878-7274

Publisher: Elsevier

This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/)

Contents lists available at ScienceDirect



Communications in Nonlinear Science and Numerical Simulation



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

Research paper

Extremization to fine tune physics informed neural networks for solving boundary value problems

Abhiram Anand Thiruthummal^{a,c,*}, Sergiy Shelyag^{b,c}, Eun-jin Kim^a

^a Centre for Fluids and Complex Systems, Coventry University, Coventry, United Kingdom

^b College of Science and Engineering, Flinders University, Adelaide, Australia

^c School of Information Technology, Deakin University, Melbourne, Australia

ARTICLE INFO

Keywords: Physics informed neural networks Theory of functional connections Boundary value problems PDEs

ABSTRACT

We propose a novel method for fast and accurate training of physics-informed neural networks (PINNs) to find solutions to boundary value problems (BVPs) and initial boundary value problems (IBVPs). By combining the methods of training deep neural networks (DNNs) and Extreme Learning Machines (ELMs), we develop a model which has the expressivity of DNNs with the fine-tuning ability of ELMs. We showcase the superiority of our proposed method by solving several BVPs and IBVPs which include linear and non-linear ordinary differential equations (ODEs), partial differential equations (PDEs) and coupled PDEs. The examples we consider include a stiff coupled ODE system where traditional numerical methods fail, a 3+1D non-linear PDE, Kovasznay flow and Taylor–Green vortex solutions to incompressible Navier–Stokes equations and pure advection solution of 1+1 D compressible Euler equation.

The Theory of Functional Connections (TFC) is used to exactly impose initial and boundary conditions (IBCs) of (I)BVPs on PINNs. We propose a modification to the TFC framework named Reduced TFC and show a significant improvement in the training and inference time of PINNs compared to IBCs imposed using TFC. Furthermore, Reduced TFC is shown to be able to generalize to more complex boundary geometries which is not possible with TFC. We also introduce a method of applying boundary conditions at infinity for BVPs and numerically solve the pure advection in 1+1 D Euler equations using these boundary conditions.

1. Introduction

Differential equations are used to mathematically model various phenomena in the fields of engineering, physics, chemistry, economics and biology. A tiny fraction of such differential equations admit an analytic closed-form solution. Therefore the study of solutions of differential equations requires the use of a variety of computational methods. There exist several numerical methods to find the solutions to ordinary differential equations (ODEs) and partial differential equations (PDEs). Traditionally, such methods are based on discretization of geometric or frequency spatial and temporal domains. Also, some numerical approximation of the derivatives involved in the differential equations is used on the grid defined by the discretization. Recently, however, a novel family of numerical methods for ODE and PDE solution began to emerge based on utilization of artificial neural networks. These methods exploit the universal function approximation capabilities of the neural networks and have several properties, unusual for numerical methods, such as differentiability of the provided solution.

* Corresponding author at: Centre for Fluids and Complex Systems, Coventry University, Coventry, United Kingdom.

E-mail addresses: thiruthuma@uni.coventry.ac.uk (A.A. Thiruthummal), sergiy.shelyag@flinders.edu.au (S. Shelyag), ad3116@coventry.ac.uk (E.-j. Kim).

https://doi.org/10.1016/j.cnsns.2024.108129

Received 1 August 2023; Received in revised form 23 May 2024; Accepted 3 June 2024

Available online 10 June 2024

^{1007-5704/© 2024} The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

A.A. Thiruthummal et al.

Physics-informed neural networks (PINNs) utilize physical laws written in the form of differential equations to facilitate the learning process to find the network weights, which can then be combined with other information to achieve various objectives. One of the first uses of PINN can be traced to [1] where the authors used neural networks along with unsupervised learning to solve boundary value problems (BVPs) of ODEs and PDEs. The term 'physics informed neural network' was coined much later in [2] where a semi-supervised learning paradigm was introduced to find solutions to PDEs in a data-driven manner from sparse and noisy experimental data. The idea of PINN was later extended for data-driven discovery of PDEs [3]. For data-driven discovery of PDEs and PDE solutions, PINN outperforms traditional methods, which requires computing the solution of PDEs multiple times iteratively to estimate the parameters.

In the case of solving IBVPs PINNs have some advantages over traditional methods like Finite Difference, Finite Volume or Finite Element Methods (FD, FVM and FEM). A PINN uses neural networks as an analytic approximation of the solution whereas FD/FVM/FEM discretize the spatial domain into elements, within which the solution is approximated using finite differences instead of derivatives, averaging and interpolation over a control volume, or fitting an arbitrary function within a control volume and minimizing the error. This generally works satisfactorily, however the problems represented by hyperbolic PDEs are often numerically unstable and require some additional numerical stabilization techniques, inevitably leading to unphysically large and uncontrollable numerical diffusion. Spatial discretization necessitates the need for a numerically stable interpolation scheme to compute the value of the solution between two elements, whereas it is straightforward to compute the values from the analytic expression of the neural network. Furthermore, the number of elements scales exponentially with the number of dimensions.

Nevertheless, for BVPs of low-dimensional PDEs PINN was still significantly slower and less accurate than an efficient spatial grid-based solver in computing solutions. To address this problem in [4] the authors make use of the Extreme Learning Machine (ELM) [5] algorithm, which significantly reduces training time of PINNs. ELM used a shallow neural network to approximate the solution and used Gauss–Newton Extremization (GNE) (Section 4.1) to train only the weights in the final layer of the neural network. In [6] by combining ELM with domain decomposition, PINNs are shown to outperform FEM in terms of solution accuracy and computational time on a wide variety of BVPs.

Another drawback of PINNs was the approximate loss function based method of imposing initial conditions (ICs) and boundary conditions (BCs). This required additional hyperparameter choice in determining the relative importance of loss functions. In [1] a trial function was constructed from the neural network output which exactly satisfy the ICs and BCs. But this method was difficult to generalize due to a lack of a general prescription to generate the trial functions for arbitrary ICs and BCs. Recently, based on the Theory of Functional Connections (TFC) [7,8] a neural network method of solving IBVPs while exactly imposing the ICs and BCs was introduced [9]. This method was further developed to use ELM [10] producing solutions orders of magnitude more accurate than grid-based methods in some instances. A fundamental limitation of shallow networks used in ELM is its lack of expressivity compared to a deep neural network (DNN) [11]. In [10] this results in DNN outperforming ELM in terms of solution accuracy in the context of Burgers equation and Navier–Stokes equations. In [6] the authors could improve the expressivity by using multiple ELMs to represent the solution.

In this work, by combining the ideas of ELM and deep networks, we propose a framework wherein a DNN is initially trained using standard methods and then GNE is used to significantly improve the solution accuracy. By combining these methods we get the expressive power of a DNN [11] combined with the fine tuning ability of shallow networks trained with ELMs. We also present a modification to the TFC framework used in [9,10] to significantly reduce the computational time required for training and inference.

The rest of this manuscript is structured as follows: Section 2 provides a concise introduction to PINN and discusses the need for TFC framework to impose constraints in BVPs. Section 3 introduces the TFC framework and demonstrates the application of TFC to impose constraints. In Section 3.3 a modification to the TFC framework is proposed to speed up computation. In Section 3.4 we discuss the drawbacks of TFC and how to generalize to complex boundary geometries. Section 4 details the neural network architectures and the training methods used for the examples provided in the section that follows. Section 5 uses various known analytic solutions to linear and non-linear ODEs (Section 5.1), PDEs (Section 5.2) and Coupled PDEs (Section 5.3) to show the superiority of our proposed method and discuss its drawbacks. Section 6 contains discussion. Some supplementary information is also provided in Appendix A-Appendix C.

2. Physics informed neural networks

Neural networks are considered as universal function approximators [12,13]. Let $\mathcal{N}_{\{\theta\}}$ be an arbitrary function represented by a neural network, where $\{\theta\}$ denotes the parameters of the neural network. The idea of PINN is to make this neural network 'physics informed' by imposing $\mathcal{N}_{\{\theta\}}$ to satisfy some given differential equation(s).

Let $\mathcal{D}(\mathbf{x}, f(\mathbf{x}))$ denote a general differential operator of an *m*-dimensional variable \mathbf{x} acting on function $f(\mathbf{x})$. Any differential equation can be written as $\mathcal{D}(\mathbf{x}, f(\mathbf{x})) = 0$. In order to impose that $\mathcal{N}_{\{\theta\}}$ satisfy this differential equation, the differential equation itself can be used as a loss function while training. The derivatives in the differential equation can be accurately computed to an arbitrary precision using auto-differentiation [14]. Note that since neural networks approximate a function, in practice this differential equation is never exactly satisfied by $\mathcal{N}_{\{\theta\}}$. This deviation of the differential operator acting on $\mathcal{N}_{\{\theta\}}$ from identically being equal to zero is called residual, denoted by \mathcal{R} :

$$\mathcal{R}(x; \{\theta\}) = \mathcal{D}\left(\mathbf{x}, \mathcal{N}_{\{\theta\}}(\mathbf{x})\right).$$

(1)

For a neural network training batch with N samples x_1, \ldots, x_n , the loss function can now be defined as a mean squared residual (MSR) of the differential equation:

$$\mathcal{L}_{DE}(\{\theta\}) = \frac{1}{N} \sum_{i=1}^{N} \|\mathcal{R}(\mathbf{x}_i, \{\theta\})\|^2.$$
⁽²⁾

Minimizing the value of \mathcal{L}_{DE} results in $\mathcal{N}_{\{\theta\}}$ more closely satisfying the differential equation $\mathcal{D}(\mathbf{x}, f(\mathbf{x})) = 0$. Note that, in the most general case we can have neural networks with multiple outputs $\mathcal{N}_1, \dots, \mathcal{N}_l$ informed by a system of differential equations $\mathcal{D} \equiv (\mathcal{D}_1, \dots, \mathcal{D}_k) = 0$. Then the same Eq. (2) holds with $\mathcal{R} \equiv (\mathcal{R}_1, \dots, \mathcal{R}_k)$.

In literature [2–4,15–17] \mathcal{L}_{DE} can be found paired with other loss functions during training to achieve mainly three objectives:

· Data Driven Solution Discovery

The objective of this framework is to find the solution of a differential equation based on sparse experimental data. Let $v_E(\mathbf{y})$ denote some noisy experimental measurement of the state of the system at point \mathbf{y} . We can approximate this state using the neural network $\mathcal{N}_{I_{0}}^{v}$. For M experimental data points $\mathbf{y}_1, \dots, \mathbf{y}_M$, the data driven loss can be defined as:

$$\mathcal{L}_{DD}(\{\theta\}) = \frac{1}{M} \sum_{i=1}^{M} \|\mathcal{N}_{\{\theta\}}^{v}(\mathbf{y}_{i}) - v_{E}(\mathbf{y}_{i})\|^{2}$$
(3)

The solution can be discovered by minimizing the combined loss function, which in this case will be a weighted sum of \mathcal{L}_{DE} and \mathcal{L}_{DD} .

$$\mathcal{L}(\{\theta\}) = \omega_{DE} \mathcal{L}_{DE}(\{\theta\}) + \omega_{DD} \mathcal{L}_{DD}(\{\theta\}) \tag{4}$$

Here ω_{DE} and ω_{DD} are weights which needs to be carefully chosen for the optimal training of the neural network.

Data Driven Parameter Discovery

For a differential equation $\mathcal{D}_{\{\gamma\}}(\mathbf{x}, f(\mathbf{x})) = 0$ parameterized by a set of parameters $\{\gamma\}$. Experimental measurement of the state $v_E(\mathbf{y}_i)$ at points $\{\mathbf{y}_i\}$ can be used to find the parameters of the differential equation. The loss function to minimize in this case will be the same as in Eq. (4) but with an additional set of parameters $\{\gamma\}$ to optimize.

$$\mathcal{L}(\{\theta\},\{\gamma\}) = \omega_{DE}\mathcal{L}_{DE}(\{\theta\},\{\gamma\}) + \omega_{DD}\mathcal{L}_{DD}(\{\theta\},\{\gamma\})$$
(5)

· Solving Boundary Value Problems

PINN framework can also be used to numerically solve the differential equation based on ICs and BCs. If f(x) denotes a solution to the differential equation $\mathcal{D}(\mathbf{x}, f(\mathbf{x})) = 0$, we can denote its *n* ICs and BCs applied on the respective boundaries $\partial \Omega_1, \dots, \partial \Omega_n$ as a set of constraints:

$$B_1(f(\mathbf{x})) = 0, \quad \forall \mathbf{x} \in \partial \Omega_1$$
...
$$B_n(f(\mathbf{x})) = 0, \quad \forall \mathbf{x} \in \partial \Omega_n$$
(6)

These constraints can be of any form including Dirichlet, Neumann and Robin BCs. The loss function for a single constraint can be written as

$$\mathcal{L}_{BC,k}(\{\theta\}) = \frac{1}{|\Omega_k|} \sum_{\mathbf{x}_i \in \Omega_k} \|\mathcal{B}_i(f(\mathbf{x}_i))\|^2$$
(7)

where $|\Omega_k|$ denotes the number of datapoints used in the summation. The total loss for the ICs and BCs can then be the weighted sum of the individual losses, where the weights $\omega_{BC,i}, \ldots, \omega_{BC,n}$ need to be carefully chosen for optimal training.

$$\mathcal{L}_{BC}(\{\theta\}) = \sum_{i=1}^{n} \omega_{BC,i} \mathcal{L}_{BC,i}(\{\theta\})$$
(8)

The solution of a BVP can then be found by minimizing the loss

$$\mathcal{L}(\{\theta\}) = \mathcal{L}_{BC}(\{\theta\}) + \omega_{DE}\mathcal{L}_{DE}(\{\theta\}) \tag{9}$$

In this work we will focus our attention solely on BVP. Since this is a multi-objective optimization problem (Eqs. (8) & (9)), the accuracy of the solutions will be sensitive to the hyperparameters ω_{DE} , $\omega_{BC,1}$, ..., $\omega_{BC,n}$. Tuning of these hyperparameters becomes challenging and time consuming especially if a large number of boundary conditions is given. In [9,10] this issue was overcome by using a mathematical framework called Theory of Functional Connections (TFC) to exactly satisfy the ICs and BCs without using a loss function. We provide a concise introduction to TFC in Section 3. Using TFC the neural network output $\mathcal{N}(x)$ is converted to a constrained function $f^c(x)$ which always satisfy the ICs and BCs. The training then becomes a single objective optimization problem as shown in Fig. 1.



Fig. 1. Structure of a fully connect neural network solving a BVP with *k* coupled differential equations D_1, \ldots, D_k and *n* initial and boundary conditions B_1, \ldots, B_n acting on *l* functions dependent on *m* variables x_1, \ldots, x_m .

3. Theory of functional connections

TFC is a mathematical framework for converting constrained problems into unconstrained problems. The idea of TFC is to write down the most general function which satisfies the constraints of the problem. In the case of PINNs, the constraints are in the form of initial conditions and boundary conditions of the differential equations. In the following subsections we will provide a brief self-contained exposition of TFC. See [7,8] for a rigorous introduction.

3.1. Univariate case

Let us consider an arbitrary univariate function f(x) subject to the constraints $C_1[f] = c_1, C_2[f] = c_2, ..., C_n[f] = c_n$, where $C_1, ..., C_n$ are linear functionals. Functional in this context refers to a mapping from the space of all functions to the real/complex numbers. The term linear implies that $C[\alpha f + \beta g] = \alpha C[f] + \beta C[g]$ for any two functions f and g, and any two real/complex numbers α and β . Some examples of linear functionals are: $C[f] = f(x_1), C[f] = f'(x_1), C[f] = \int_a^b f(x)dx$. The idea of TFC is to write down a constrained expression $f^c(x)$ for f(x), such that $f^c(x)$ always satisfy the constraints:

$$f^{c}(x) = \mathcal{N}(x) + \sum_{i=1}^{n} \eta_{i} s_{i}(x).$$
(10)

Here $\mathcal{N}(x)$ is some arbitrary function known as the free function, which in the context of PINN will be represented by a neural network. η_1, \ldots, η_n are some unknown coefficients and $\{s_1(x), \ldots, s_n(x)\}$ is a set of linearly independent functions we choose, called the support functions. An example of a set of support functions is the polynomial basis $\{1, x, \ldots, x^{n-1}\}$. Since by definition $f^c(x)$ satisfy all the constraints, Eq. (10) can be written as the following set of equations:

$$c_{1} = C_{1}[f] = C_{1}[\mathcal{N}] + \sum_{i=1}^{n} \eta_{i}C_{1}[s_{i}],$$
:
$$c_{n} = C_{n}[f] = C_{n}[\mathcal{N}] + \sum_{i=1}^{n} \eta_{i}C_{n}[s_{i}].$$
(11)

Once a set of support functions is chosen, we can solve for η_i to get an exact expression for $f^c(x)$ in terms of the free function $\mathcal{N}(x)$:

$$\begin{bmatrix} \eta_1 \\ \vdots \\ \eta_n \end{bmatrix} = \begin{bmatrix} C_1[s_1] & \cdots & C_1[s_n] \\ \vdots & \ddots & \vdots \\ C_n[s_1] & \cdots & C_n[s_n] \end{bmatrix}^{-1} \begin{bmatrix} c_1 - C_1[\mathcal{N}] \\ \vdots \\ c_n - C_n[\mathcal{N}] \end{bmatrix}.$$
(12)

The linearity of the functionals was used to split the RHS of Eq. (11) into multiple terms. Note that in some rare cases the support functions will need to be carefully chosen to make the matrix in Eq. (12) invertible. For most commonly occurring ICs and BCs the polynomial basis can be used as the support functions.

When solving a differential equation with PINN, $\mathcal{N}(x)$ represents the neural network output, C_1, \ldots, C_n the initial and boundary conditions and $f^c(x)$ the solution of the differential equation which satisfy the initial and boundary conditions.

3.2. Multivariate case

For a multi-variate function $f(x_1, x_2, ..., x_m)$, this procedure of finding the constrained expression can be applied iteratively on one variable at a time to get the general constrained expression. To do so we first define the operator C^i which is a mapping from the space of multivariate functions which depend on $x_1, x_2, ..., x_m$, to a space of multivariate functions which depend on $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_m$. In short the operator C^i acts only on the ith variable of the multivariate function. Now we can represent n_i constraints on each of the *i*th dimension as follows:

$$C_{1}^{i}[f] = c_{1}^{i}(x_{1}, \dots, x_{i-1}, x_{i+1}, \dots, x_{m}),$$

$$C_{2}^{i}[f] = c_{2}^{i}(x_{1}, \dots, x_{i-1}, x_{i+1}, \dots, x_{m}),$$

$$\vdots$$

$$C_{n}^{i}[f] = c_{n}^{i}(x_{1}, \dots, x_{i-1}, x_{i+1}, \dots, x_{m}),$$
(13)

where $C_1^i, \ldots, C_{n_i}^i$ are linear operators acting only on the variable x_i . For example the Dirichlet BC $f(x_1, 5, x_3, \ldots, x_n) = \sin(x_1 + x_3 + x_4 + \cdots + x_n)$ has the form in Eq. (13) with $C^2 = f(x_1, 5, x_3, \ldots, x_n)$ and $c^2(x_1, x_3, x_4, \ldots, x_n) = \sin(x_1 + x_3 + x_4 + \cdots + x_n)$. Similarly we can represent Neumann and Robin BCs. Note that this form of constraints which are represented by linear operators which act on a single variable limits the shape of domain boundaries on which BCs can be applied. BCs can only be applied on boundaries of the form $x_i = constant$. This is a fundamental limitation of the TFC framework. See Section 3.4 for further discussion and generalization.

Now starting with a free function $\mathcal{N}(x_1, \dots, x_n)$, represented by the neural network, we can iteratively apply the constraints to derive the general constrained expression as follows:

$$f_{1}^{c}(x_{1},...,x_{m}) = \mathcal{N}(x_{1},...,x_{n}) + \sum_{i=1}^{n_{1}} \eta_{i}^{1} s_{i}(x_{1})$$

$$f_{2}^{c}(x_{1},...,x_{m}) = f_{1}^{c}(x_{1},...,x_{n}) + \sum_{i=1}^{n_{2}} \eta_{i}^{2} s_{i}(x_{2})$$

$$\vdots$$

$$f_{m}^{c}(x_{1},...,x_{m}) = f_{m-1}^{c}(x_{1},...,x_{n}) + \sum_{i=1}^{n_{m}} \eta_{i}^{m} s_{i}(x_{m})$$
(14)

Here f_1^c satisfies the C^1 constraints, f_2^c satisfies the C^1 and C^2 constraints, and so on. Using the same method for the one-dimensional case (Eq. (10) & Eq. (11)), at each iteration we can compute the η_i^j . $f^c \equiv f_m^c$ now satisfy all the constraints. Note that at each step of the iteration we can use different set of support functions. In this work we exclusively use polynomial basis functions as support functions. For high-dimensional functions, the algebra of this iterative application of constraints becomes cumbersome, therefore we make use of symbolic computation capabilities of Mathematica to get the constrained expression.

3.3. Reduced TFC

For the univariate case of TFC, in order to get the exact expression for $f^c(x)$ (Eq. (10)), we need to derive the expression for η_i using Eq. (12). This solution contains the quantities $C_1[\mathcal{N}], \ldots, C_n[\mathcal{N}]$. In the case of PINN, since $\mathcal{N}(x)$ is represented by a neural network, computing these quantities requires additional evaluation of the neural network. The number of additional evaluations depends on the type of constraints. If *C* is a value constraint like $C[\mathcal{N}] = \mathcal{N}(\bar{x})$ where \bar{x} is some constant, computing $C_n[\mathcal{N}]$ requires an evaluation of the neural network for the value \bar{x} , a derivative constraint $C[\mathcal{N}] = \mathcal{N}'(\bar{x})$ requires a more expensive auto-differentiation to evaluate, and an integral constraint $C_n[\mathcal{N}] = \int \mathcal{N}(x) dx$ requires multiple evaluations. For *n* constraints, computing of $f^c(x)$ in the univariate case therefore requires at least n + 1 evaluations of the neural network.

For the case of the multivariate functions, for the sake of simplicity we will assume all the constraints are value constraints. In Eq. (14), following similar arguments for the univariate case, evaluation of f_1^c requires $n_1 + 1$ evaluation of the neural network. The evaluation of f_2^c requires $n_2 + 1$ evaluations of f_1^c , hence $(n_2 + 1) \times (n_1 + 1)$ evaluations of the neural network. Using inductive logic we get that evaluation of f_m^c requires $(n_1 + 1) \times \cdots \times (n_m + 1)$ evaluations of the neural network.

In this work, we propose the idea of Reduced TFC to reduce the number of evaluations of the neural network, which can be computationally expensive. To achieve this we note that majority of the boundary value problems we encounter in practice have value or derivative constraints.

Suppose for the univariate case we have a *k*th derivative constraint $C[f] := f^{(k)}(\tilde{x}) = c$ where \tilde{x} and c are some constants, we can modify the free function in Eq. (10), $\mathcal{N}(x) \to \mathcal{N}(x)(x - \tilde{x})^{k+1}$, so that $C[\mathcal{N}(x)(x - \tilde{x})^{k+1}] = 0$. Therefore an additional evaluation of the $\mathcal{N}(x)$ represented by the neural network is not required. To generalize this, if we have *n* derivative constraints $f^{(k_1)}(\tilde{x}_1) = c_1, \ldots, f^{(k_n)}(\tilde{x}_n) = c_n$, the Reduced TFC constraint expression can be written as

$$f^{c}(x) = \mathcal{N}(x)(x - \tilde{x}_{1})^{k_{1}+1} \cdots (x - \tilde{x}_{n})^{k_{n}+1} + \sum_{i=0}^{n} \eta_{i} s_{i}(x).$$
(15)

The solution for η_i will then be given by a simplified version of Eq. (12):

$$\begin{bmatrix} \eta_1 \\ \vdots \\ \eta_n \end{bmatrix} = \begin{bmatrix} s_1^{(k_1)}(x_1) & \cdots & s_n^{(k_1)}(x_1) \\ \vdots & \ddots & \vdots \\ s_1^{(k_n)}(x_n) & \cdots & s_n^{(k_n)}(x_n) \end{bmatrix}^{-1} \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix}$$
(16)

In the multivariate case, for the set of constraints

$$C_{j}^{i}[f] := \partial_{x_{i}}^{k_{j}} f(x_{1}, \dots, x_{m}) \Big|_{x_{i} = \bar{x}_{j}^{i}} = c_{j}^{i}$$
(17)

the following modification is required to the free function:

$$\mathcal{N}(x_1, \dots, x_m) \to \mathcal{N}(x_1, \dots, x_m) \prod_{i,j} (x_i - \tilde{x}_j^i)^{k_j + 1}.$$
(18)

This can then be substituted in Eq. (14) and all η_j^i can be computed. The resulting $f^c(x_1, \ldots, x_m)$ requires only a single evaluation of $\mathcal{N}(x_1, \ldots, x_m)$ which is represented by the neural network. As we will see in Sections 5.2.3, 5.2.4 and 5.3.1, compared to TFC, Reduced TFC will lead to a significant speedup in computing the solution of PDEs.

3.4. Remarks

For the multivariate TFC while deriving the expression for the constrained expression f^c (Eq. (14)), we assume that the constraints are of the form $C_j^i[f] = c_j^i(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_m)$, where C_j^i is an operator acting only on the variable x_j . This limits the geometry of the hypersurface on which ICs and BCs can be applied. The hypersurface should always be of the form $\partial \Omega = \{x \mid x_j = \text{const.}\}$ for some coordinate x_j or some combination of it. The class of BVPs where ICs and BCs are applied to the boundary of the domain $[a_1, b_1] \times \cdots \times [a_m, b_m]$ satisfy this condition. In real world applications a large class of BVPs are defined with BCs defined on more complex geometries. The unit circle $\partial \Omega = \{(x, y) \mid x^2 + y^2 = 1\}$ is an example of a boundary on which the TFC framework fails since it cannot be expressed in the form $\{x \mid x_j = \text{const.}\}$. While for this specific case we can still use TFC after a coordinate change to the polar coordinates (r, θ) and the boundary can be represented by the equation r = 1, such coordinate changes are not possible in general.

For complex boundary geometries in 2D an extension of TFC using bijective mapping can be used [18]. For more complicated geometries in arbitrary dimensions, a method based on approximate distance functions has been proposed in [19]. The form of Reduced TFC expression (Eq. (15)) can also be generalized to complex geometries. Note that the form of Eq. (15) is such that the 1st term in the RHS is identically zero on the boundaries and the 2nd term satisfies the values at the boundary. This can in principle be extended to more complex boundaries where Dirichlet, Neumann or mixed boundary conditions are imposed. Let $b_1(\mathbf{x}) = 0, b_2(\mathbf{x}) = 0, \dots, b_n(\mathbf{x}) = 0$ be the *n* boundaries where BCs are applied and k_1, k_2, \dots, k_n be the order of the derivative constraint applied at the boundary. Note that k = 0 denotes a Dirichlet BC. Then the constrained expression can be written as

$$f^{c}(\mathbf{x}) = \mathcal{N}(\mathbf{x})b_{1}(\mathbf{x})^{k_{1}+1} \cdots b_{n}(\mathbf{x})^{k_{n}+1} + \mathcal{G}(\mathbf{x}).$$
⁽¹⁹⁾

Here the 1st term in the RHS is identically zero at the boundaries after applying the appropriate number of derivative operations and the 2nd term $\mathcal{G}(\mathbf{x})$ satisfies the values at the boundaries. For a combination of boundary conditions it is not straightforward to derive a function $\mathcal{G}(\mathbf{x})$ to satisfy the values at the boundaries for all boundary conditions. For rectangular boundaries this is solved using TFC. For more complicated BCs a neural network can be used to represent $\mathcal{G}(\mathbf{x})$ and the neural network can be trained to learn the values at the boundaries using the loss function in Eq. (8). Once trained, the neural network $\mathcal{G}(\mathbf{x})$ can be used to impose the boundary conditions and only the neural network $\mathcal{N}(\mathbf{x})$ needs to be trained to solve the differential equation. Further investigation is needed in this direction to find better ways to derive $\mathcal{G}(\mathbf{x})$ for the case of multiple BCs with complex boundary geometries.

For the sake of simplicity, in this work we mostly use ICs and BCs on rectangular boundaries where TFC is applicable. Note that TFC and Reduced TFC cannot be used to impose periodic BCs. For PINNs periodic BCs can be imposed by changing the neural network architecture [20]. This is explored in Section 5.3.2. Also see Section 5.3.3 for a case where BCs are imposed at infinity.

4. Neural network & training

In this work, we use a fully connected neural network (FCNN) to represent the function(s) which are solution(s) of differential equation(s). FCNNs have been previously used in the context of PINNs [1,6,9,20]. In this work, we mainly use two different network architectures: a single hidden layer (shallow) FCNN [5] and a multi-layer (deep) FCNN. Based on previous studies [4,10], a shallow FCNN, trained with extreme learning machine (ELM) [5] algorithm is found to outperform DNNs in terms of solution accuracy and training speed in different scenarios. The ELM algorithm uses Gauss–Newton Extremization (GNE) to train the weights in the final layer of the neural network.

4.1. Gauss-Newton extremization

GNE is a popular algorithm for non-linear least squares optimization. Let $\mathcal{N}_{\{\lambda\}}(\mathbf{x})$ be a non-linear fitting function parameterized by a set of parameters $\{\lambda\}$. If \mathcal{N} represents a neural network $\{\lambda\}$ can be a subset of parameters of the neural network. Now for an input variable x_i and its corresponding output y_i , the residual of the fitting function $\mathcal{R}_i(\{\lambda\})$ is defined as

$$\mathcal{R}_i(\{\lambda\}) := y_i - \mathcal{N}_{\{\lambda\}}(\mathbf{x}) \tag{20}$$

Note that for solving BVPs using PINNs the residual \mathcal{R} will be given by Eq. (1). The Gauss–Newton algorithm is then used to iteratively minimize the sum of squares of the residual, $\sum_i \mathcal{R}_i(\{\lambda\})^2$. Starting from an initial guess λ_0 for the parameter vector, GNE prescribes the following iteration to minimize the sum of squares of \mathcal{R}_i :

$$\lambda_{t+1} - \lambda_t = -\left(\mathbf{J}_t^{\mathsf{T}} \mathbf{J}_t\right)^{-1} \mathbf{J}_t^{\mathsf{T}} \mathcal{R}(\lambda_t).$$
⁽²¹⁾

Here $\mathcal{R} \equiv (\mathcal{R}_i, ...)$ is the residual vector and **J** is the Jacobian matrix defined as

$$\left(\mathbf{J}_{t}\right)_{ij} := \frac{\partial \mathcal{R}_{i}\left(\lambda_{t}\right)}{\partial\left(\lambda_{t}\right)_{j}}.$$
(22)

Note that Eq. (21) involves a matrix inversion, which is numerically unstable. An alternate way to compute the quantity on the RHS of Eq. (21) is to note that for the linear least squares (LLS) problem of minimizing $\|\mathbf{J}_t - \mathcal{R}(\lambda_t)\beta\|^2$ w.r.t. β , the analytic solution is given by $\beta = (\mathbf{J}_t^{\mathsf{T}} \mathbf{J}_t)^{-1} \mathbf{J}_t^{\mathsf{T}} \mathcal{R}(\lambda_t)$. Therefore the RHS of Eq. (21) can be computed in a numerically stable way using LLS. In this work we use the PyTorch function *torch.linalg.lstsq* for LLS which utilizes the standard implementation in LAPACK library [21]. LAPACK provides several different algorithms for computing the LLS solution. Out of these methods, singular value decomposition (SVD) is found to provide consistent results throughout different problems we have considered in this work. See Appendix A for details.

4.2. Neural network architecture

The ELM architecture consists of a single hidden layer FCNN, with randomly initialized constant input weights and biases, and trainable weights and no bias in the final output layer. The ELM architecture with *m* inputs (x_1, \ldots, x_m) , *l* outputs $(\mathcal{N}_1, \ldots, \mathcal{N}_l)$ and *H* nodes in the hidden layer can be written as

$$\mathcal{N}_k = \sum_{j=1}^H \beta_{kj} \sigma \left(\sum_{i=1}^m w_{ji} x_i + b_j \right), \quad 1 \le k \le l.$$
(23)

Here w_{ji} and b_j are the untrainable weights and biases respectively, and β_{kj} are the trainable weights. σ is a non-linear function known as the activation function. We use tanh activation throughout this work. The values of β_{kj} are computed using GNE during training.

The deep FCNN architecture we use in this study consists of M hidden layers. All hidden layers except the last has H nodes with biases and the last hidden layer has H_E nodes without bias. This choice was made to replicate the structure of ELM for the last layer. The mathematical description of the neural network can be written as:

$$h_{j}^{(1)} = \sigma \left(\sum_{i=1}^{m} w_{ji}^{(1)} x_{i} + b_{j}^{(1)} \right),$$

$$h_{j}^{(2)} = \sigma \left(\sum_{i=1}^{H} w_{ji}^{(2)} h_{i}^{(1)} + b_{j}^{(2)} \right),$$

$$\vdots$$

$$\mathcal{N}_{k} = \sum_{j=1}^{H_{E}} \beta_{kj} h_{j}^{(M)}, \quad 1 \le k \le l.$$
(24)

Here $h_j^{(k)}$ are the values of nodes in the *k*th hidden layer and $w_{ji}^{(k)}$ and $b_j^{(1)}$ the corresponding weights and biases respectively. β_{kj} are the weights of the final output layer. If we choose M = 1 we get the ELM architecture.

See Fig. 1 for the complete layout of the PINN.

4.3. Training

The deep FCNN used in the context of PINNs are typically trained with stochastic gradient descent method like Adam [22] or quasi-Newton method like BFGS [23]. In most cases BFGS optimizer is found to provide the best solution accuracy for PINN. An exception can be found in Section 5.1.1, where we used Adam. In all other examples considered in this work a limited memory variant of the BFGS algorithm (L-BFGS) [24] is used for optimizing the parameters of the deep network. The specific implementation of the optimizers we use are from the PyTorch library [25]. In this work we propose that after training the deep FCNN with the desired optimizer, the β_{kj} parameters be further optimized using GNE. In Section 5 we show that combining optimizers in this manner results in orders of magnitude improvement in the solution accuracy.

Note that any linear differential equation $D(\mathbf{x}, f(\mathbf{x}))$ can be written as $\tilde{D}(\mathbf{x}, f(\mathbf{x})) + g(\mathbf{x})$, where \tilde{D} is the linear homogeneous part. If one of the outputs \mathcal{N}_k of the neural network satisfies a linear differential equation, using Eq. (24) we can write

$$D\left(\mathbf{x}, \mathcal{N}_{k}(\mathbf{x})\right) = \sum_{j=1}^{H_{E}} \beta_{kj} \tilde{D}\left(\mathbf{x}, h_{j}^{(M)}(\mathbf{x})\right) + g(\mathbf{x}).$$
(25)

Finding the minimum of the residual of $D(\mathbf{x}, \mathcal{N}_k(\mathbf{x}))$ w.r.t. β_{kj} now becomes a LLS problem which can be minimized with a single GNE iteration. The same argument applies to a coupled linear system of differential equations and the weights β_{kj} can be computed using a single GNE iteration. The arguments still hold when using TFC constrained expression, since TFC and Reduced TFC use linear transformations to generate the constrained expressions. For a non-linear differential equation, we need to perform multiple iterations in GNE and the iteration is stopped when the MSR (Eq. (2)) stops decreasing.

We train the neural network by uniformly sampling points from the domain and compute the loss function (Eq. (2)). For ELM the loss function is minimized w.r.t. β_{kj} using GNE. For the deep network, the loss function is first minimized w.r.t. all the parameters

Table 1

A summary of different problems considered in the Results section. The Remarks column outlines important observations for each of the considered problems.

Section	Problem	Remarks
ODE		
5.1.1	$y'(t) = -\sin t$	Showcase limitation of ELM when learning complex solutions and necessity of deep networks. GNE solution is demonstrated to improve with increase in number of neurons in the final layer. Adam used since L-BFGS fails to converge.
5.1.2	Non-linear coupled ODEs	An example where traditional numerical methods fail. Incremental training method introduced. PINN fails without incremental training.
PDE		
5.2.1	2D linear & non-linear PDEs with same solution	ELM & L-BFGS + GNE give similar solution accuracy for linear PDE. L-BFGS + GNE outperforms ELM for non-linear PDE.
5.2.2	1+1 D Burgers Equation	GNE fails if large gradients are present. L-BFGS + GNE shown to work over a wider range of gradients compared to ELM. Drawback of TFC discussed.
5.2.3	2+1 D Heat Equation	Smaller solution error with L-BFGS compared to ELM but significantly slower. L-BFGS + GNE improves upon L-BFGS by more than 2 orders of magnitude with a small increase in computation time over ELM. 4x speedup with Reduced TFC compared to TFC.
5.2.4	3+1 D Non-linear PDE	20x - 40x speedup with Reduced TFC compared to TFC.
Coupled PDEs		
5.3.1	Kovasznay flow	Steady state solution to 2D incompressible stationary Navier–Stokes equations. L-BFGS + GNE significantly outperforms traditional as well as other neural network based methods.
5.3.2	Taylor–Green vortex	Unsteady solution to 2D incompressible Navier–Stokes equations. Periodic BCs applied by changing the neural network architecture and using sinusoidal activation function. Showcase benefits of domain decomposition on large domains.
5.3.3	Pure advection in 1+1 D Euler Equation	ELM fails to solve. Incremental training method used. Imposing vanishing BCs at infinity.

using Adam or L-BFGS. Then after the desired number of iterations is reached, the loss function is minimized w.r.t. β_{kj} using GNE. New random samples are generated for each iteration of Adam or L-BFGS. This prevents overfitting even if the number of samples is small. For GNE we use a larger number of sample points to prevent overfitting. Throughout this study the number of sample points for Adam or L-BFGS and GNE is chosen after a crude manual hyperparameter search.

Performing GNE requires computing the Jacobian. In this work for the ease of implementation, Jacobian is computed using forward mode auto-differentiation implemented in PyTorch as *torch.autograd.functional.jacobian()*. The specific vectorized implementation of Jacobian in PyTorch is memory intensive and while performing GNE using a given set of sample points, the Jacobian needs to be computed in batches and the batch size depends on the neural network architecture and the operations performed to compute the residual. Note that it is possible to significantly speed up the computation of Jacobian by deriving an analytic expression for it. A drawback is that it needs to be done on a case by case basis. This process can be automated using symbolic computation packages like Mathematica but is not pursued in this work.

During training, the progress is quantified using the root mean squared residual (RMSR) of the differential equation or the system of differential equations. Since the actual training curves of the neural network will have significant fluctuations, the training curves shown in Section 5 represent the values of the best model which has the lowest RMSR obtained until that particular training step. Therefore the training curves will be monotonically decreasing. Throughout this work, unless stated otherwise, all the parameters of the neural networks are initialized using Xavier uniform initialization [26]. We use tanh activation function in this study for simplicity. See [27,28] for trainable activation functions.

5. Results

In this section by solving different BVPs we show the superiority of the proposed extremization method for DNNs. The ICs and BCs are applied using TFC or Reduced TFC. The exact expressions can be found in Appendix C. Some examples considered in this section are taken from [1,10]. Note that ELMs and DNNs trained after applying boundary conditions using TFC is referred to as X-TFC and Deep-TFC respectively in [10]. All computations in this section were performed using double-precision arithmetic on a workstation with AMD EPYC 7552 CPU, Nvidia Titan RTX GPU and 64 GB of RAM. The general outline of this section is given in Table 1:



Fig. 2. Root mean square of the residual of the ODE for neural networks with (left) 1 hidden layer with 400 neurons and (right) 4 hidden layers with (32, 32, 32, 400) neurons, trained on the interval $t \in [0, t_{max}]$.



Fig. 3. Comparison of the exact solution with the neural network approximation computed using a different training method on the domain $t \in [0, 100]$. The neural network had 4 hidden layers with (32, 32, 32, 400) neurons. Note that the error in the Adam solution is apparent when we zoom into a small section of the domain (right panel).

5.1. ODE

5.1.1. $y'(t) = -\sin t$

In this subsection we will use a simple example $y'(t) = -\sin t$ to show the limitation of ELM in learning complicated functions. Since this is a linear ODE, a single Gauss–Newton iteration is required to find the solution using ELM. In this example we adjust the complexity of the solution by changing the size of the solution domain $[0, t_{max}]$. Throughout this example neural networks are trained using 1000 random uniformly sampled points from the domain for Adam and 2000 random uniformly sampled points for GNE. For uniformity all networks using Adam were trained for 100,000 steps. Note that this training method is extremely slow and the incremental training introduced in Section 5.1.2 will significantly speed up the process, but is not used here for simplicity.

In Fig. 2 (left) it can be seen that the single hidden layer neural network fails to learn the solution irrespective of the training algorithm, when the domain size becomes large. Up to 1200 neurons in the hidden layer was tried but yielded similar results. For the deep network in Fig. 2 (right), Adam starts performing better than ELM at around $t_{max} = 20$. This is because ELM is only optimizing the weights of the final layer of the neural network whereas with Adam we are optimizing the entire network. Now by combining Adam with GNE we create a training method which can efficiently search the large parameter space of the DNN and extremize w.r.t. the final layer to fine-tune the solution. This results in orders of magnitude improvement in the RMSR of the ODE as seen in Fig. 2 (right). This improvement in RMSR is reflected in the error of the solution as can be seen in Fig. 3 and Fig. 4.

The learning curve in Fig. 5 shows the benefit of using Adam before extremization w.r.t. the final layer. The plot was generated from the same training process where a copy of the current best model from Adam training was made at regular intervals and GNE was done on this copy. Note that the sudden drop in RMSR value for Adam at around 35,000 steps corresponds to a similar but more significant drop in RMSR for the case of Adam + GNE. The error in the numerical solution before and after using GNE is shown in Fig. 4. Adam in this case helped us to efficiently search the parameter space of the DNN. After Adam learned the essential features of the solution, using GNE we were able to fine-tune the solution to significantly improve the accuracy. There is nothing special about Adam in this training process. Any optimizing algorithm that can efficiently search a high dimensional parameter space can be paired with the GNE.

Fig. 6 shows the effect of the configuration of the neural network on its ability to learn the solution of the ODE. A single hidden layer network is unable to learn the solution of the ODE irrespective of the number of neurons in the hidden layer. However, adding



Fig. 4. The absolute error in the solution after training a 4 hidden layer neural network with (32, 32, 32, 400) neurons for 100,000 steps on the domain $t \in [0, 100]$.



Fig. 5. Learning curve of a DNN with 4 hidden layers with (32, 32, 32, 400) neurons trained on the domain $t \in [0, 100]$.

an extra hidden layer Adam is more capable to learn the solution and GNE further refines and improves this solution. Beyond a threshold of around 100 neurons, the number of neurons in the final layer had no significant effect on the accuracy of the solution when Adam was used. As for GNE an increase in the number of neurons was associated with an increase in the accuracy of the solution. This is mainly due to the fact that with an increase in the number of neurons the extremization step has more degrees of freedom to fine-tune the solution.

Note that for PINNs, L-BFGS optimizer is generally found to be faster. It also leads to the lower RMSR compared to Adam. However, for this example L-BFGS fails to converge when the domain size becomes too large. Even though incremental training described in the next subsection can be used to mitigate this, we use Adam in this example for simplicity. Throughout the rest of this study we will exclusively use L-BFGS.

5.1.2. Stiff coupled non-linear ODEs

A system of ODEs

$$\frac{d}{dt}u = \cos(t) + u^2 + v - \left(1 + t^2 + \sin^2(t)\right),$$
(26)

$$\frac{d}{dt}v = 2t - \left(1 + t^2\right)\sin(t) + uv\tag{27}$$

has an analytical solution for the initial conditions u(0) = 0 and v(0) = 1, given by

$$u(t) = \sin(t)$$

(28)



Fig. 6. Root mean square of residual of ODE trained in the interval $t \in [0, 100]$ with different neural networks using (left) Adam and (right) Adam + GNE. All hidden layers of neural network except the last had 32 neurons.



Fig. 7. Comparison of exact solution with numerical solution found using BDF method for the coupled ODEs Eqs. (26) and (27). The relative and absolute tolerance for BDF was set at 10^{-15} .

 $v(t) = 1 + t^2$.

(29)

This system of ODEs was numerically solved in [1] using a neural network on the domain [0, 3]. On a larger domain this system exhibits stiffness and numerical instability. We tried solving this system of ODEs on the domain [0, 10] using several standard and commonly used ODE solvers like LSODA [29], RK45 [30], DOP853 [31] and BDF [32] implemented in *solve_ivp* function of *scipy* package. All these solvers failed to solve the system of ODEs in different ways. For instance in the case of LSODA the step size became too small irrespective of the prescribed tolerance values. RK45 finds an incorrect solution when absolute and relative tolerances are greater than 10^{-10} . For smaller tolerance values RK45 step size effectively becomes 0 and integration fails. A similar effect, although with smaller tolerance values, is observed with DOP853 method. BDF, which is capable of providing numerical solutions for some stiff problems, with relative and absolute tolerance set at 10^{-15} gave an incorrect solution as shown in Fig. 7.

Different training algorithms for PINNs also failed to converge when directly trained on the full domain [0, 10]. Instead we start with a smaller domain [0, 0.5] and increase the domain size during training when certain conditions are met. For ELM, since the equation under consideration is non-linear, GNE iterations are repeated until the RMSR stops decreasing, then we increase the domain size by 0.5. As for L-BFGS we start with [0, 0.5] domain and increase the domain size by 0.5 when the RMSR becomes smaller than 5×10^{-2} . This process is repeated until the full domain size [0, 10] is reached. Then the training continues without any domain size increments. We call this incremental training. 1000 random uniformly sampled points were chosen from the domain for each L-BFGS training step irrespective of the domain size. 4000 random uniformly sampled points were used for GNE in ELM and after L-BFGS.

Fig. 8 shows the training curve for incremental training with different training methods. The initial fluctuation in the RMSR values are due to change in domain size. The L-BFGS method reaches full domain size faster than ELM, but ELM solution has significantly lower RMSR. The L-BFGS + GNE method takes about twice the amount of time as compared to ELM to achieve marginally lower RMSR value. The corresponding absolute errors in the solutions are shown in Fig. 9. Note that the increase in error towards the end



Fig. 8. Evolution of RMSR with incremental training for different training algorithms. The black vertical dotted lines denote the point when L-BFGS reaches full domain size [0,10] with incremental training. This line appears at different locations on the I-LBFGS and I-LBFGS + GNE curves since the latter requires additional time to perform GNE after I-LBFGS reaches full domain size. ELM had a single hidden layer with 400 neurons. L-BFGS optimized a neural network with 2 hidden layers with (32,400) neurons.



Fig. 9. The absolute error in the solution found using different incremental training methods on the domain $t \in [0, 10]$. ELM had a single hidden layer with 400 neurons. L-BFGS optimized a neural network with 2 hidden layers with (32, 400) neurons.

of the domain is not a trend that would extrapolate if we increase the domain size. For this specific system of ODEs this trend of increasing error towards the end of the domain is seen irrespective of the domain $[0, t_{max}]$, where $t_{max} \gtrsim 8$. Fig. 10 shows the error in the solution for neural networks trained on a larger domain [0, 20].

Similar to the example given in Section 5.1.1, the advantage of deep network trained with incremental L-BFGS + GNE is seen when we further increase the domain size to [0, 20] as shown in Fig. 10. Using more neurons in the hidden layer resulted in higher RMSR for incremental ELM.

5.2. PDE

5.2.1. 2D linear and non-linear PDE

The linear PDE we use in this subsection is given by

$$\nabla^2 u(x, y) = \left(2 - \pi^2 y^2\right) \sin(\pi x)$$
(30)

and the non-linear PDE is

۲

$$\nabla^2 u(x, y) + u(x, y) \frac{\partial}{\partial y} u(x, y) = \sin(\pi x) \left(2 - \pi^2 y^2 + 2y^3 \sin(\pi x) \right)$$
(31)

We use the same set of mixed boundary conditions on both equations

$$u(0, y) = 0$$

(32)



Fig. 10. The absolute error in the solution found using different incremental training methods on the domain $t \in [0, 20]$. ELM had a single hidden layer with 400 neurons. L-BFGS optimized a neural network with 3 hidden layers with (32, 32, 400) neurons.

Table 2

Statistics of numerical solution for linear and non-linear PDE found using ELM and L-BFGS + GNE. ELM had a single hidden layer with 200 neurons. L-BFGS optimized a neural network with 3 hidden layers with (32, 32, 400) neurons.

	ELM			L-BFGS + GNE		
	RMSR	Mean Abs. Err.	Max. Abs. Err.	RMSR	Mean Abs. Err.	Max. Abs. Err.
Linear	9.2×10^{-10}	2.1×10^{-12}	1.2×10^{-11}	1.0×10^{-9}	1.6×10^{-12}	8.1×10^{-12}
Non-Linear	3.7×10^{-7}	6.0×10^{-10}	4.2×10^{-9}	1.5×10^{-9}	4.2×10^{-12}	2.3×10^{-11}

u(1, y) = 0	(33)
u(x,0) = 0	(34)
$\frac{\partial}{\partial y}u(x,1) = 2\sin(\pi x)$	(35)

and both PDEs have the same analytic solution

$$u(x,y) = y^2 \sin(\pi x). \tag{36}$$

For ELM we use a single hidden layer with 200 neurons and for the deep network we use 2 hidden layers with (32, 200) neurons. Decreasing the number of neurons in the last hidden layer decreased the RMSR value while increasing only leads to marginal improvement in RMSR. Increasing the number of hidden layers seemed not to affect the RMSR but increased the computational time. We use 1000 random uniformly sampled points from the domain for each L-BFGS training step and 2000 random uniformly sampled points for the GNE in ELM and after L-BFGS training.

For the linear equation ELM required a single iteration of GNE and the solution took 1.5 s to compute. In the case of L-BFGS + GNE, 2 steps of L-BFGS took around 2 s and an additional 1.5 s for GNE taking a total of 3.5 s to compute the solution. In the case of non-linear equation GNE involves multiple iterations. The ELM solution took around 8 s. For the hybrid training 3 L-BFGS steps took around 3.5 s and the GNE took 19 s taking a total of 22.5 s to compute. For both these cases pure L-BFGS RMSR and mean absolute error were around 10^{-1} . We need to train for significantly longer duration with pure L-BFGS to attain an absolute error which is orders of magnitude worse than ELM and L-BFGS + GNE.

Even though the analytic solution which the neural network has to learn is the same for both the PDEs, ELM performed worse with the non-linear PDE compared to the linear PDE according to Table 2. In comparison, for L-BFGS + GNE the solution of linear and non-linear PDEs had similar statistics. In the case of the linear PDE, a LLS problem w.r.t. weights in the last layer of the neural network can be extremized in a single iteration of GNE. However for the non-linear PDE a more complicated optimization landscape is found where combining L-BFGS with GNE is advantageous compared to just using GNE in the case of ELM. Note that in the case of the non-linear PDE solved with ELM the LLS algorithm used to perform the Gauss–Newton iteration is seen to significantly affect the numerical solution. This is discussed further in Appendix A.

5.2.2. 1+1 D Burgers' equation

The (viscous) Burgers' equation is a non-linear PDE given by the following equation:

$$\frac{\partial u}{\partial t} + \alpha u \frac{\partial u}{\partial x} = v \frac{\partial^2 u}{\partial x^2}.$$
(37)

A previous study [10] addressed solving Burgers' equation using ELM and DNN. In order to compare with their results, we solve the PDE on the same domain $(x, t) \in [-3, 3] \times [0, 1]$ using the same boundary conditions:

$$u(-3,t) = \frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}(-3-ct)\right)$$
(38)



Fig. 11. Plot of analytic solution (Eq. (41)) of Burgers' equation for (left) v = 1, (center) v = 0.1 and (right) v = 0.01.



Fig. 12. Learning curve of neural networks using L-BFGS (dashed curves) and L-BFGS + GNE (solid curves) for the Burgers' equation with different values of ν . All the neural networks were trained for 15,000 L-BFGS steps. ELM took between 15 s to 25 s depending on the value of ν .

$$u(3,t) = \frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}(3-ct)\right)$$

$$u(x,0) = \frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}x\right).$$
(39)
(40)

The exact analytic solution for this boundary value problem is given by

$$u(x,t) = \frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}(x-ct)\right).$$
(41)

In this case the viscosity parameter v plays a role in determining the sharpness of the gradients in the exact solution:

$$\nabla u(x,t) = \left(\frac{c^2}{\alpha v + \alpha v \cosh\left(\frac{c}{2v}(x-ct)\right)}, -\frac{c^3}{\alpha v + \alpha v \cosh\left(\frac{c}{2v}(x-ct)\right)}\right)$$
(42)

$$\Rightarrow \nabla u(x,t)|_{x=ct} = \left(\frac{c^2}{2\alpha\nu}, -\frac{c^2}{2\alpha\nu}\right).$$
(43)
The magnitude of the gradient at $x = ct$ is inversely proportional to ν and as $\nu \to 0$ the analytic solution becomes discontinuous.

The magnitude of the gradient at x = ct is inversely proportional to v and as $v \rightarrow 0$ the analytic solution becomes discontinuous. This can be readily seen in Fig. 11. Since all the optimization methods we use in this work rely on computing the gradient terms in Eq. (37) and further computing gradients of the residual of Eq. (37) w.r.t. the weights in the neural network, we expect PINN to perform worse as v becomes small. In this subsection we will look at how robust the proposed method is to large gradients when training.

We use similar architecture as used in [10] and set $\alpha = c = 1$. ELM had a single hidden layer with 601 neurons. The deep network had 4 hidden layers with (32, 32, 32, 200) neurons. Each L-BFGS step was trained on 1000 random uniformly sampled points from the domain and GNE in ELM and after L-BFGS was done with 2000 random uniformly sampled points from the domain. All L-BFGS training in this subsection with and without GNE was done for 15,000 steps for consistency. Most of the neural networks converged to the lowest RMSR well before this, as shown in Fig. 12.



Fig. 13. (left) Mean absolute error and (right) maximum absolute error for Burgers' equation solved with different training algorithms.



Fig. 14. Absolute value of residual of Burgers' equation for (left) v = 1, (center) v = 0.1 and (right) v = 0.01 after training with L-BFGS for 15,000 steps.

From Fig. 13 it can be seen that for $v \leq 0.2$ ELM starts performing worse than DNN trained with L-BFGS. In comparison L-BFGS + GNE consistently performs better than both ELM and L-BFGS. For $v \leq 0.01$ it was observed that applying GNE on a model trained with L-BFGS led to similar or worse solution error. This is likely due to GNE being sensitive to large gradients.

Fig. 14 shows the residual of the Burgers equation for DNNs trained with L-BFGS for different values of v. For large values of v the residuals are small and the values are evenly distributed throughout the domain. As v gets smaller the residuals start to assume large values around the lines x = 0 and x = ct. Here x = ct is the location of the discontinuity when $v \rightarrow 0$ and the large value of residual at x = 0 line is due to the presence of $\tanh\left(\frac{c}{2v}x\right)$ term in the TFC constrained expression Eq. (C.6). This is a fundamental limitation of TFC when imposing discontinuous boundary conditions or boundary conditions with sharp gradients. In this case, if we impose the boundary condition as a loss function (Eq. (9)), we can get rid of the large residual values along the x = 0 line, but we will still be left with large residual value along x = ct. During GNE, these larger residuals localized to a small area of the domain will result in a Jacobian matrix mostly populated by small values with sparsely occurring large values. This in turn will result in large updates (Eq. (21)) for the weights. This can be seen in Fig. 15. All the statistics of the proposed updates scales approximately as v^{-1} . For very small values of v (v < 0.01), these large updates result in the solution becoming worse after GNE. For small enough values of v all gradient based methods will fail due to exploding gradients, but GNE fails before methods like Adam and L-BFGS since it is more sensitive to large gradients.

The method described in [33] was found to be helpful for extremely small values of v, even up to v = 0. Nevertheless, the results were not always consistent, occasionally producing solutions with the discontinuity on a line other than x = ct.

5.2.3. 2+1 D heat equation

In this subsection we will solve the 2D heat equation

$$\frac{\partial^2}{\partial x^2}u(x,y,t) + \frac{\partial^2}{\partial y^2}u(x,y,t) = \kappa \frac{\partial}{\partial t}u(x,y,t)$$
(44)

subject to the following Dirichlet boundary conditions:

u(0, y, t) = 0 u(L, y, t) = 0 u(x, 0, t) = 0 u(x, H, t) = 0(45)
(46)
(47)
(47)
(47)



Fig. 15. Statistics of the updates proposed to the weights in the last layer of the DNN after the first iteration of GNE. Note that the statistics scales approximately as v^{-1} .



Fig. 16. Statistics of the solution to heat equation computed with different methods and using Reduced TFC to impose constraints. Solid lines denote L-BFGS + GNE, dotted lines denote L-BFGS and x marks ELM. The time taken by L-BFGS is not shown in the figure since it takes significantly longer than other methods. See Table 3 for exact values.

$$u(x, y, 0) = \sin\left(\frac{\pi x}{L}\right) \sin\left(\frac{\pi y}{H}\right)$$
(49)

The analytic solution is given by:

$$u(x, y, t) = \sin\left(\frac{\pi x}{L}\right) \sin\left(\frac{\pi y}{H}\right) e^{-\left(\frac{\pi^2}{L^2} + \frac{\pi^2}{H^2}\right)t}$$
(50)

The equation is solved in the domain $x, y, t \in [0, L] \times [0, H] \times [0, 1]$. In this study we use the values L = 2, H = 1, $\kappa = 1$. The ELM network had a single hidden layer with 400 neurons. Further increase in the number of neurons did not improve the numerical solution. The deep network had 3 hidden layers with (32, 32, 400) neurons. At each L-BFGS step the network was trained on 2000 random uniformly sampled points from the domain and GNE in ELM and after L-BFGS was done with 2000 random uniformly sampled points from the domain.

Fig. 16 shows that the solution error from L-BFGS is slightly better than ELM but L-BFGS takes significantly longer to compute. L-BFGS + GNE on the other hand takes slightly longer than ELM to compute the solution but achieves lower error by more than 2 orders of magnitude. From Table 3 it can be seen that using Reduced TFC we are able to compute solutions around 4 times

Table 3

Statistics for different training algorithms. The shallow ELM had a single hidden layer with 400 neurons. The deep network had 3 hidden layers with (32, 32, 400) neurons. The number of L-BFGS training steps used were 400 for pure L-BFGS and 30 for L-BFGS + GNE.

	RMSR	Mean Abs. Err.	Max. Abs. Err.	Time
Reduced TFC				
ELM	4.0×10^{-3}	4.6×10^{-5}	3.1×10^{-4}	4.9 s
L-BFGS	2.9×10^{-3}	2.8×10^{-5}	2.0×10^{-4}	264 s
L-BFGS + GNE	1.1×10^{-5}	6.2×10^{-8}	5.9×10^{-7}	32 s
TFC				
ELM	4.4×10^{-3}	6.0×10^{-5}	3.6×10^{-4}	19.5 s
L-BFGS	4.2×10^{-3}	3.8×10^{-5}	2.5×10^{-4}	955 s
L-BFGS + GNE	3.2×10^{-5}	1.3×10^{-7}	1.6×10^{-6}	136 s

faster with also a marginal improvement in error. This improvement is due to fact that with Reduced TFC we require only a single evaluation of the neural network compared to multiple evaluations in the case of TFC.

5.2.4. 3+1 D non-linear PDE

In this subsection we will showcase an example where Reduced TFC provides significant advantage over TFC. We will solve the following 3+1 dimensional non-linear PDE.

$$\partial_{x}u(x, y, z, t) \partial_{y}u(x, y, z, t) \partial_{z}u(x, y, z, t) + \partial_{t}^{2}u(x, y, z, t) = \left((t-1)tx(z-1) + x^{2}\cos\left(x^{2}y\right) + \frac{3}{2}x\sqrt{y}z\right) \\ \left((t-1)ty(z-1) + 2xy\cos\left(x^{2}y\right) + y^{3/2}z\right) \\ \left(2\pi t^{2}\cos(2\pi z) + (t-1)txy + xy^{3/2}\right) \\ + 2xy(z-1) + 2\sin(2\pi z)$$
(51)

subject to the following Dirichlet boundary conditions:

$$u(0, y, z, t) = t^{2} \sin(2\pi z)$$
(52)

$$u(x, 0, z, t) = t^2 \sin(2\pi z)$$
 (53)

$$u(x, y, 1, t) = \sin(x^2 y) + x y^{3/2}$$
(54)

$$u(x, y, z, 0) = \sin(x^2 y) + x y^{3/2} z$$
(55)

$$u(x, y, z, 1) = \sin\left(x^2y\right) + xy^{3/2}z + \sin(2\pi z).$$
(56)

The exact analytic solution for this boundary value problem is given by:

$$u(x, y, z, t) = t^{2} \sin(2\pi z) + \sin(x^{2}y) + xy^{3/2}z + xyt(z-1)(t-1).$$
(57)

The PDE was solved in the domain $(x, y, z, t) \in [0, 1] \times [0, 1] \times [0, 1] \times [0, 1]$. The ELM used in this case had 1 single hidden layer with 400 neurons. The deep network had 3 hidden layers with (32, 32, 400) neurons. At each L-BFGS step the network was trained on 2000 random uniformly sampled points from the domain and GNE in ELM and after L-BFGS was done with 2000 random uniformly sampled points from the domain.

Table 4 shows that L-BFGS + GNE provides only a marginal improvement in accuracy over ELM. The important thing to note here is that, using Reduced TFC we are able to achieve more than 20 times speedup in the case of ELM and L-BFGS + GNE, and more than 40 times speedup for L-BFGS.

5.3. Coupled PDE

5.3.1. Kovasznay flow solution

The 2D incompressible stationary Navier-Stokes equation given by

$$u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} = -\frac{1}{\rho}\frac{\partial p}{\partial x} + v\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right)$$
(58)
$$u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} = -\frac{1}{\rho}\frac{\partial p}{\partial y} + v\left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2}\right)$$
(59)
$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0$$
(60)

$$\frac{\partial x}{\partial x} + \frac{\partial y}{\partial y} = 0,$$
(60)

1)

Table 4

Statistics for different training algorithms. The ELM had 1 single hidden layer with 400 neurons. The deep network had 3 hidden layers with (32,32,400) neurons. The number of training steps used were 800 for pure L-BFGS and 10 for L-BFGS + GNE.

	RMSR	Mean Abs. Err.	Max. Abs. Err.	Time
Reduced TFC				
ELM	2.8×10^{-9}	5.9×10^{-11}	5.4×10^{-9}	7 s
L-BFGS	1.5×10^{-4}	3.9×10^{-6}	6.2×10^{-5}	47 s
L-BFGS + GNE	6.6×10^{-10}	1.8×10^{-11}	2.0×10^{-9}	10 s
TFC				
ELM	3.3×10^{-6}	1.0×10^{-7}	1.2×10^{-5}	150 s
L-BFGS	3.4×10^{-4}	7.3×10^{-6}	2.0×10^{-4}	1899 s
L-BFGS + GNE	9.0×10^{-7}	2.2×10^{-8}	1.9×10^{-6}	245 s



Fig. 17. Plot of analytic Kovasznay flow solution with v = 0.025, $\rho = 1$ and $p_0 = 1$.

where *u* and *v* are *x* and *y* components of the velocity respectively, *p* the pressure, *v* the viscosity and ρ the density. The Kovasznay flow [34] is an exact solution to the above coupled system of PDEs given as:

$$u(x, y) = 1 - e^{\lambda x} \cos(2\pi y) \tag{61}$$

$$(x, y) = \frac{\lambda}{2\pi} e^{\lambda x} \sin(2\pi y)$$
(62)

$$p(x, y) = p_0 - \frac{1}{2}e^{2\lambda x}$$
, where p_0 is an arbitrary constant (63)

$$\lambda = \frac{1}{2\nu} - \sqrt{\frac{1}{4\nu^2} + 4\pi^2}.$$
(64)

The Kovasznay flow solution shown in Fig. 17 is often used to benchmark traditional [35,36] as well as neural networkbased [37,38] numerical solvers. Compared to these methods, in this subsection we show that using L-BFGS + GNE we are able to achieve orders of magnitude more accurate numerical solutions in a fraction of the time. For this test we set the values v = 0.025, $\rho = 1$ and $p_0 = 1$. The domain is chosen as $(x, y) \in [0, 2] \times [0, 2]$ and we apply Dirichlet boundary conditions on u(x, y) and v(x, y) at all 4 boundaries of the domain based on the analytic solution. In principle without applying any boundary condition on p(x, y), we can still solve for u(x, y) and v(x, y). In practice when no boundary condition is applied on p(x, y) we have an extra degree of freedom, related to the pressure p_0 which is found to become extremely large (~ 10⁷) when solved using neural networks. This in turn leads to a loss in precision as the computations are limited by double precision arithmetic. Therefore we apply Dirichlet boundary condition for the pressure p(x, y) on any one of the boundaries.

The shallow network used for ELM had 400 neurons in the single hidden layer. Increasing the number of neurons further did not decrease the error. The deep network had 3 hidden layers with (32, 32, 400) neurons. At each L-BFGS step the network was trained on 2000 random uniformly sampled points from the domain and GNE in ELM and after L-BFGS was done with 3000 random uniformly sampled points from the domain.

As illustrated in Fig. 18, even though L-BFGS + GNE takes longer to compute the solution compared to ELM, the LBGS + GNE solution has several orders of magnitude lower error. L-BFGS, even after running for 600 training steps, still lagged behind ELM in terms of accuracy while taking significantly longer to compute. Table 5 shows the advantage of Reduced TFC over TFC. Reduced TFC provided a 50% improvement in runtime. This pales in comparison to Sections 5.2.3 and 5.2.4 where we achieved a speedup of 4 times and 20 times respectively. This is mostly due to the fact that for this boundary value problem, the computational time



Fig. 18. Statistics of the Kovasznay flow solution computed with different methods using Reduced TFC to impose constraints for (**left**) u(x, y), (**center**) v(x, y) and (**right**) p(x, y). Solid lines denote L-BFGS + GNE, dotted lines denote L-BFGS and x marks ELM. The time taken by L-BFGS is not shown in the figure since it takes significantly longer than other methods. See Table 5 for exact values.

Table 5

Statistics of the Kovasznay flow solution computed with different training methods using TFC and Reduced TFC. The number of L-BFGS training steps used were 600 for pure L-BFGS and 40 for L-BFGS + GNE.

	RMSR	RMSR		Mean Abs. Err.			Max. Abs. Err.			Time
	Eq. (58)	Eq. (59)	Eq. (60)	u(x, y)	v(x, y)	p(x, y)	u(x, y)	v(x, y)	p(x, y)	
Reduced TFC										
ELM L-BFGS L-BFGS + GNE	2.8×10^{-4} 1.4×10^{-3} 2.9×10^{-8}	2.6×10^{-4} 7.9×10^{-4} 1.9×10^{-8}	2.4×10^{-4} 6.9×10^{-4} 2.1×10^{-8}	$\begin{array}{c} 1.8 \times 10^{-5} \\ 7.5 \times 10^{-5} \\ 9.0 \times 10^{-10} \end{array}$	$\begin{array}{c} 1.3 \times 10^{-5} \\ 5.3 \times 10^{-5} \\ 6.0 \times 10^{-10} \end{array}$	1.9×10^{-5} 1.5×10^{-4} 1.2×10^{-9}	$\begin{array}{c} 6.4 \times 10^{-5} \\ 3.2 \times 10^{-4} \\ 5.3 \times 10^{-9} \end{array}$	7.3×10^{-5} 2.6×10^{-4} 3.2×10^{-9}	1.1×10^{-4} 8.9×10^{-4} 1.1×10^{-8}	95 s 857 s 182 s
TFC										
ELM L-BFGS L-BFGS + GNE	$\begin{array}{c} 1.4 \times 10^{-4} \\ 1.3 \times 10^{-3} \\ 6.4 \times 10^{-8} \end{array}$	$\begin{array}{c} 1.3 \times 10^{-4} \\ 9.3 \times 10^{-4} \\ 2.5 \times 10^{-8} \end{array}$	1.2×10^{-4} 1.0×10^{-3} 3.8×10^{-8}	7.0×10^{-6} 8.8×10^{-5} 2.4×10^{-9}	5.1×10^{-6} 8.2×10^{-5} 1.4×10^{-9}	$\begin{array}{c} 1.5 \times 10^{-5} \\ 1.1 \times 10^{-4} \\ 5.2 \times 10^{-9} \end{array}$	3.7×10^{-5} 3.1×10^{-4} 1.6×10^{-8}	2.8×10^{-5} 4.7×10^{-4} 7.2×10^{-9}	$\begin{array}{c} 1.0 \times 10^{-4} \\ 5.9 \times 10^{-4} \\ 4.5 \times 10^{-8} \end{array}$	147 s 1307 s 318 s

is dominated by computing the various derivatives in the PDE using auto-differentiation rather than evaluating the neural network multiple times to impose constraints using TFC.

Note that the slowest step in the GNE process is the computation of the Jacobian matrix. Since GPUs are equipped with a limited amount of VRAM, this computation is often done in batches. For instance in the case of the neural network used in this subsection, for each Gauss–Newton iteration, Jacobian matrix of 3000 sample outputs w.r.t. 1200 weights was computed in batches of 100 and took around 25 s. Once the Jacobian matrix was computed, the LLS took on average 560 ms per Gauss–Newton iteration. Since we use random sampling from the domain to train, multiple GPUs can readily be used to parallelize the batch-wise computation of the Jacobian. For multi-dimensional PDEs the synchronization time between all the GPUs and time for computing the LLS will be negligible compared to the time it takes to compute the Jacobian matrix, resulting in an almost linear speedup. Another option discussed in Section 4.3 is to derive an analytic expression for Jacobian.

5.3.2. Taylor-Green vortex solution

The 2D incompressible Navier-Stokes equations are written as

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} = -\frac{1}{\rho} \frac{\partial p}{\partial x} + v \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right)$$

$$\frac{\partial v}{\partial x} + u \frac{\partial v}{\partial y} = -\frac{1}{\rho} \frac{\partial p}{\partial x} + v \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} \right)$$
(65)

$$\frac{\partial t}{\partial x} = \frac{\partial y}{\partial y} + \frac{\partial y}{\partial x^2} + \frac{\partial y^2}{\partial y^2}$$

$$\frac{\partial v}{\partial y} = 0$$
(65)

$$\frac{\partial x}{\partial x} + \frac{\partial y}{\partial y} = 0 \tag{67}$$

The Taylor–Green vortex solution [39] to the Navier–Stokes equations describes an unsteady flow with decaying vortex and can be written as follows:

$$u(x, y, t) = \sin x \cos y e^{-2vt}$$
 (68)

$$v(x, y, t) = -\cos x \sin y e^{-2vt}$$
(69)

$$p(x, y, t) = \frac{\rho}{4} (\cos 2x + \sin 2y)e^{-4\nu t}.$$
(70)

In this case we will solve the Navier–Stokes equations by applying periodic boundary conditions and initial conditions which match the analytic solution for u(x, y, t), v(x, y, t) and p(x, y, t) at t = 0. As we discussed in Section 5.3.1, the initial condition on p(x, y, t) is not required in principle but necessary in practice.



Fig. 19. Structure of the neural network imposing periodic boundary conditions along the x and y axes. Raw outputs are the neural network outputs before applying TFC constraints.

Based on [20], the periodic boundary conditions were exactly imposed using an additional layer with a periodic activation function right after the input layer as shown in Fig. 19. The other constraints are imposed using Reduced TFC. The variables x and y of the periodic dimensions are first passed to a set of m periodic activation functions defined as follows:

$$S_{1i}(x) = A_{1i} \sin\left(\frac{2\pi}{L_x}x + \phi_{1i}\right), \quad 1 \le i \le m,$$
(71)

$$S_{2i}(y) = A_{2i} \sin\left(\frac{2\pi}{L_y}y + \phi_{2i}\right), \quad 1 \le i \le m.$$
(72)

Here L_x and L_y are the periods of the x and y variables respectively, which have the values $L_x = L_y = 2\pi$ in this example. $A_{1i}, A_{2i}, \phi_{1i}$ and ϕ_{2i} are weights which are initialized randomly with Xavier normal initialization [26] and remain constant for ELM but are trainable parameters in the case of DNN trained with L-BFGS. $S_{1i}(x), S_{2i}(y)$ and the non-periodic variable *t* now form the new input variables which are then passed to a FCNN. Note here that Eqs. (71) and (72) do not exactly follow the prescription in [20]. The current form of Eqs. (71) and (72) was found to be optimal for ELM. L-BFGS and L-BFGS + GNE was found to be agnostic to the specific form, as long as the periodic activation was present, since they have the additional freedom to train the parameters.

The shallow network used for ELM has 15 neurons (m = 15) each for x and y variables in the periodic activation layer and 400 neurons in the single hidden layer. Increasing the number of neurons in the hidden layer was found to decrease the RMSR but took significantly longer to train due to computational constraints. The deep network had the same configuration for the periodic activation layer and had 3 hidden layers with (32, 32, 400) neurons. At each L-BFGS step the network was trained on 2000 random uniformly sampled points from the domain and GNE in ELM and after L-BFGS was done with 3000 random uniformly sampled points from the domain.

Fig. 20 shows the learning curve for the various training methods used. ELM and L-BFGS + GNE outperform pure L-BFGS. L-BFGS + GNE can achieve the same RMSR value as ELM slightly faster, and being trained for a longer duration is able to achieve a better RMSR. This corresponds to decreased error in the L-BFGS + GNE solution compared to ELM as shown in Fig. 21. Note that in Fig. 21, the absolute error in p(x, y, t) is extremely large for all the training methods, even though we applied boundary conditions as we did in Section 5.3.1. This is because p(x, y, t) in this section has an additional gauge freedom where $p(x, y, t) \rightarrow p(x, y, t) + g(t)$ will also satisfy the Navier–Stokes equations since p(x, y, t) only appears as a gradient w.r.t. x and y variables in the equation. Imposing any kind of boundary condition at t = 0 will only constrain the value of g(t) at t = 0. Even though it is not possible to unambiguously learn the function p(x, y, t), the value of its derivatives can be determined. This is evident in Fig. 21 where the errors in $\partial_x p(x, y, t)$ $\partial_y p(x, y, t)$ are small.

An important thing to consider here is that the error in u, v and ∇p are almost a constant with a slight increase towards the end of the t domain. However, the exact solutions given by Eqs. (68)–(70) are exponentially decaying in time. This means that the



Fig. 20. Learning curve of the different training algorithms used for solution of the Taylor-Green vortex problem. For L-BFGS with and without GNE 700 training steps were used.



Fig. 21. Mean absolute error at different time slices t for the Taylor–Green vortex solution computed using different training methods.



Fig. 22. Heat map of numerical solution u(x, y, t) to Taylor–Green vortex problem at t = 10 found using (left) ELM and (right) L-BFGS + GNE.

relative error in the solution will be exponentially increasing. This increase in relative error translates to visible artifacts in the ELM solution at t = 10 as shown in Fig. 22. L-BFGS + GNE on the other hand yields a solution which is visibly indistinguishable from the true solution. This is even more evident in Fig. 23 which shows that the maximum relative error in the ELM solution is around 30% whereas for L-BFGS + GNE the maximum relative error is around 3%.



Fig. 23. Heat map of relative error of numerical solution u(x, y, t) to Taylor–Green vortex problem at t = 10 found using (left) ELM and (right) L-BFGS + GNE.



Fig. 24. Heat map of relative error of numerical solution u(x, y, t) to Taylor–Green vortex problem at t = 10 found using a combination of BTM with(**left)** ELM and (**right**) L-BFGS + GNE.

In order to better learn long-time solutions to time-dependent differential equations, a method, named block time marching (BTM), was introduced in [6]. With BTM, the time domain is decomposed into sub-domains represented by independent neural networks. Starting from the sub-domain containing the initial condition, the neural networks are trained one at a time and successively, with value from the previous neural network being used as the initial condition for the current network. For this example we decompose the domain into two sub-domains [0, 5.25] and [5, 10.25]. The 5% extension of the domain was to take into account the slight increase in error towards the end of the domain boundary. The first neural network is trained with the initial condition at t = 0 and the second network is trained with initial condition at t = 5, the values of which are evaluated from the previously trained neural network.

Fig. 24 shows the relative error in the solution found using BTM. The accuracy of both ELM and L-BFGS + GNE improved with BTM, resulting in a 0.8% maximum relative error in the case of ELM and 0.06% in the case of L-BFGS + GNE at t = 10. Note that it is in principle possible to further improve the solution accuracy by decomposing the domain into smaller sub-domains and applying BTM. BTM is but an example of a general class of domain decomposition PINN methods [6,17] which may be necessary to find accurate solutions on large domains due to computational hardware limitations. This example serves as an illustration of the applicability of extremization applied on top of domain decomposition to improve solution accuracy. Further in-depth study is required into this topic and is not undertaken at this point.

5.3.3. Pure advection in compressible flow

The compressible 1-dimensional Euler equations are given as

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x} (\rho u) = 0$$

$$\frac{\partial}{\partial t} (\rho u) + \frac{\partial}{\partial x} (\rho u^2 + p) = 0$$
(73)
(74)



Fig. 25. Learning curves of incremental L-BFGS and incremental L-BFGS + GNE. The incremental training reached the full domain size in 27 steps. The curves are not plotted against training time to show the overlap in RMSR values. The 200 L-BFGS steps took 152 s. The GNE took on average 50 s.

$$\frac{\partial E}{\partial t} + \frac{\partial}{\partial x} \left((E+p)u \right) = 0, \tag{75}$$

where the total (kinetic + internal) energy E is defined as:

$$E = \frac{1}{2}\rho u^2 + \frac{p}{\gamma - 1}.$$
 (76)

We also assume $\gamma = 1.4$. For this system of PDEs, if we start with a constant pressure $p(x, 0) = p_0$ and constant velocity $u(x, 0) = u_0$, the result is pure advection in density, where we have $\rho(x, t) = \rho(x - u_0 t, 0)$. The pressure and velocity remain constant in time. In this example we start with the ICs

$$p(x,0) = 1,$$
 (77)

$$u(x,0) = 1,$$
 (78)

$$\rho(x,0) = \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad \mu = 0.5, \, \sigma = 0.1.$$
(79)

We solve this system of PDEs on the domain $(x,t) \in [0,1] \times [0,1]$. We use the BC that the density vanishes at infinity, $\lim_{x\to\pm\infty} \rho(x,t) = 0$. The ICs on p and u are imposed using Reduced TFC. In order to impose the BCs on ρ we modify the Reduced TFC constrained expression as follows:

$$\int_{p}^{c} (\mathbf{x}, t) = \mathcal{N}_{p}(\mathbf{x}, t) t + 1$$
(80)

$$f_{u}^{c}(x,t) = \mathcal{N}_{u}(x,t)t + 1$$
(81)

$$f_{\rho}^{c}(x,t) = \exp\left(\mathcal{N}_{\rho}(x,t)t - \frac{(x-\mu)^{2}}{2\sigma^{2}}\right).$$
(82)

Here \mathcal{N}_{ρ} , \mathcal{N}_{u} , \mathcal{N}_{ρ} are the neural network outputs before applying constraints. The proof that Eq. (82) satisfies the BC at infinity is given in Appendix B.

In order to train the neural network we use the incremental training approach we used in Section 5.1.2. The neural network was first trained on the [0, 0.01] time domain. After a threshold RMSR of 10^{-3} was reached, the domain size was increased by 50%. This process was repeated until the full time domain size of [0, 1.0] was reached. Then the training is continued till the desired RMSR value is reached. For this example it is possible to train the neural network without incremental training, but it was found to be significantly slower. Note that the PyTorch implementation of L-BFGS while training, occasionally encountered large gradients and resulted in *NaN* values in the neural network parameters. When this happens the neural network is reverted back to the previous best parameters (lowest RMSR) and training was resumed.

We tried different number of output neurons $\{100, 200, 300, 400, 500, 600\}$ for the ELM, with and without incremental training. Directly training on the full domain always resulted in *NaN* values of the parameters. With incremental training ELM initially learns well when domain size is small but later fails and network parameters become *NaN* when the domain size is increased. A similar behavior can be seen in Fig. 25 when GNE initially yields better RMSR, but as the domain size increases, it fails and overlaps the learning curve of pure L-BFGS. After reaching the full domain size, it takes further L-BFGS training for GNE to start improving the RMSR. The DNN we used in this example had 3 hidden layers with (32, 32, 400) neurons.



Fig. 26. Errors in the solution of pure advection computed using (top) incremental L-BFGS and (bottom) incremental L-BFGS + GNE.

Fig. 26 shows the errors for the solutions of u, p and ρ . The use of GNE has resulted in a 5 orders of magnitude improvement in the solution accuracy. However it is to be noted that similar to the case of the Burgers' equation in Section 5.2.2, the improvement with GNE decreases with sharper gradients, which in this case means smaller σ in Eq. (79). This results in difficulties in solving shockx solutions to the Euler equations [40]. Even though crude solutions can be computed using PINN trained with Adam and L-BFGS [33,41,42], GNE always fails due to exploding gradients as discussed in Section 5.2.2.

6. Discussion

In this work we proposed a novel extremization method for fast and accurate training of PINNs for solving (I)BVPs by combining SGD or L-BFGS method for training DNNs, with GNE. By combining the ideas from training DNNs and ELMs, our proposed method could retain the expressive power of DNNs while benefiting from the fine tuning ability of ELMs. Using TFC to exactly satisfy ICs and BCs, in Section 5 we showed the superiority of our proposed method compared to traditional training methods for DNNs and ELMs in solving BVPs of various ODEs, PDEs and coupled PDEs. In all cases the extremization method was found to produce solutions which are orders of magnitude better than DNNs trained with SGD or L-BFGS method alone. It also performed better than ELMs and depending on the complexity of the solution resulted in marginal to significant improvement in solution accuracy. The extremization method was shown to work in cases where ELM failed (Section 5.1.1, Section 5.2.2, Section 5.3.3) and further improve the accuracy of solutions computed using domain decomposition methods (Section 5.3.2). The utility of this method is yet to be investigated in the case of data-driven solutions and data-driven parameter discovery of PDEs. The simplicity of this algorithm and the ease of implementation suggest that this method can be readily used in other contexts where overfitting is not a concern or adequate measures can be taken to prevent overfitting.

The extremization method is in general significantly faster than SGD or L-BFGS training methods since a lesser number of SGD or L-BFGS training iterations need to be performed before applying GNE, compared to using SGD or L-BFGS alone, to achieve similar or better solution accuracy. Our proposed method is slower than ELM since it requires additional SGD or L-BFGS training iterations before performing GNE, whereas ELM only uses GNE.

We also proposed a modification to the TFC framework called Reduced TFC (Section 3.3). Compared to TFC, Reduced TFC was shown in Sections 5.2.3, 5.2.4 & 5.3.1 to provide up to 40x speed-up in computational time. In Section 3.4 we discussed the limitation of TFC in imposing boundary conditions on complex boundary geometries and how Reduced TFC can be used to solve this in principle. More work is needed in this direction to derive analytic expression for $\mathcal{G}(\mathbf{x})$ in Eq. (19).

In Section 5.3.1 a significant amount of computational time was spend on GNE and we suggested a multi-GPU implementation of GNE to speed up this computation. Such an implementation is expected to provide an almost linear speed-up of GNE computation with the number of GPUs used. Another problem-specific method of speeding up the computation is to derive an analytic expression for the Jacobian. This can in principle be done by first deriving a TFC or Reduced TFC-constrained expression f^c of Eq. (24), plugging it into the differential equation to get the expression for the residual \mathcal{R} and then using Eq. (22) to compute the analytic expression for the Jacobian. The expression for the Jacobian will often be complicated and the use of a symbolic computation package like Mathematica is recommended. The analytic expression for the Jacobian will enable us to compute the Jacobian in a single forward pass with little memory overhead to evaluate the analytic expression.

Table A.6	
LAPACK [21] algorithm names and t	the specific operation it performs.

0	
Algorithm name	Operation
gels	solve LLS using QR or LQ factorization
gelsd	solve LLS using divide-and-conquer SVD
gelsy	solve LLS using complete orthogonal factorization
gelss	solve LLS using SVD

In Section 5.2.2 through the example of Burgers' equation we showed the failure of GNE when large gradients are present in the solution. GNE is not a unique choice in performing extremization, and other methods like the Levenberg–Marquardt [43] and robust Gauss–Newton [44] need to be investigated. For large enough gradients, since all gradient-based methods fail, there is a need to explore non-gradient-based optimization methods [45–47] for PINN training. Even with non-gradient-based optimization, the residuals of the differential equations still contain derivatives which can lead to convergence failure of the optimization algorithm when large enough gradients or discontinuities are present. Some training algorithms have been proposed [16,48] that partly solve this problem by using residuals of the differential equations while training the neural network. Such methods require careful treatment of the problem on a case-by-case basis. Another fundamental limitation of commonly used neural network architectures is that they can only represent continuous functions and approximating discontinuous solutions leads to exploding gradients during training. This issue is addressed in [49] where a neural network architecture with trainable discontinuous activation functions is proposed, which can represent arbitrary discontinuous functions in solutions. Even though such a network can represent an arbitrary discontinuous function, new training methods need to be developed to train it in the context of PINNs.

Software availability

The Python source code for the methods introduced in this work is available from [50].

CRediT authorship contribution statement

Abhiram Anand Thiruthummal: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. Sergiy Shelyag: Supervision, Validation, Writing – review & editing, Investigation. Eun-jin Kim: Supervision, Validation, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

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

Acknowledgments

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. Abhiram Anand Thiruthummal thanks Dr. Abhishek Kumar for the valuable discussion on exact solutions of Navier–Stokes equations.

Appendix A. Least squares

In this section we look at the effect of specific LLS algorithm on the solution accuracy achieved from GNE. We use the non-linear equation (Eq. (31)) in Section 5.2.1 as a case study. We use the same neural network architecture and sample size as in Section 5.2.1. ELM was trained with GNE. L-BFGS + GNE was trained with 3 steps of L-BFGS and then used GNE. The training process was repeated 100 times for each network architecture with each LLS algorithm to generate the statistics shown in Fig. A.27.

In Fig. A.27, it can be seen that *gelsd* and *gelss* which use SVD decomposition (Table A.6) consistently perform better than the other methods. An anomalous behavior can be seen when using *gels* to perform GNE after L-BFGS. In this case, even though the RMSR value was better than for the other algorithms, the mean and maximum absolute errors were in line with the other algorithms. The explanation for these trends is beyond the scope of this work and requires further investigation. Throughout this work we use *gelsd* to perform LLS.

Communications in Nonlinear Science and Numerical Simulation 137 (2024) 108129



Fig. A.27. Statistics of the solution to the non-linear PDE given by Eq. (31) computed by (Top) ELM and (bottom) L-BFGS + GNE using different LLS algorithms provided in the LAPACK library [21]. See Table A.6 for details on the algorithms.

Appendix B. Boundary condition at infinity

Consider a neural network with multiple inputs x and a bounded activation function in at least one of the layers, say, layer k. Then the outputs from the layer k, represented by $h_i^{(k)}(x)$ will be bounded:

$$\left|h_{j}^{(k)}(\mathbf{x})\right| < C_{1} \ \forall \mathbf{x}, \ 0 < C_{1} < \infty.$$
(B.1)

Now let Φ denote the output(s) of the neural network after the forward pass through the rest of the layers. Then the output(s) will also be bounded if there are no singularities in the activation functions in the rest of the layers:

$$\left|\boldsymbol{\Phi}\left(\boldsymbol{h}_{j}^{(k)}(\boldsymbol{\mathbf{x}})\right)\right| < C_{2} \; \forall \boldsymbol{\mathbf{x}}, \; 0 < C_{2} < \infty.$$
(B.2)

In order to apply the BC at infinity consider a function g(x) which vanishes at infinity, $\lim_{|x|\to\infty} g(x) = 0$. Then the constrained expression for vanishing BC at infinity can be written as:

$$f^c = \tilde{g}(\boldsymbol{\Phi})g(x)$$
 or (B.3)

$$f^c = g\left(\tilde{g}(\boldsymbol{\Phi}) + x\right),\tag{B.4}$$

where $\tilde{g}(x)$ represents any function without singularities in the domain $|\Phi| < C_2$. By considering the asymptotic behavior of g and \tilde{g} similar expressions can be derived for the case of unbounded activation functions. A BC with finite value at infinity can also be imposed by adding constants to the constrained expression Eq. (B.3) or Eq. (B.4).

In Section 5.3.3 the constrained expression we used (Eq. (79)) can be derived using both Eqs. (B.3) and (B.4). In terms of Eq. (B.3), g(x) is the initial condition $\rho(x, 0)$ and $\tilde{g}(x) = \exp(t x)$. The specific form of \tilde{g} was chosen so that the density $\rho \ge 0$ and $\tilde{g} \to 1$ when $t \to 0$ to satisfy the initial condition.

Appendix C. Constrained expressions

The explicit form of the TFC and Reduced TFC constrained expressions we use in this work are as follows:

• Section 5.1.1

$$f_{c}^{*}(t) = \mathcal{N}_{v}(t) - \mathcal{N}_{v}(0) + 1 \tag{C.1}$$

• Section 5.1.2

$$f_{u}^{c}(t) = \mathcal{N}_{u}(t) - \mathcal{N}_{u}(0)$$
(C.2)
$$f_{v}^{c}(t) = \mathcal{N}_{v}(t) - \mathcal{N}_{v}(0) + 1$$
(C.3)

• Section 5.2.1

$$f_{u1}^{c}(x,y) = \mathcal{N}_{u}(x,y) - (1-x)\mathcal{N}_{u}(0,y) + x\mathcal{N}_{u}(1,y)$$
(C.4)

$$f_{u}^{c}(x,y) = f_{u1}^{c}(x,y) - f_{u1}^{c}(x,0) + y \left(2\sin(\pi x) - \partial_{y}f_{u1}^{c}(x,y)\Big|_{y=1}\right)$$
(C.5)

• Section 5.2.2

$$f_{u1}^{c}(x,t) = \mathcal{N}_{u}(x,t) - \mathcal{N}_{u}(x,0) + \frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}x\right) + \frac{3+x}{2}\left(\frac{c}{2\nu} - \frac{c}{2} \tanh\left(\frac{c}{2\nu}(3-ct)\right) - f^{c}(3,t)\right)$$
(C.6)

$$f_{u}^{c}(x,y) = f_{u1}^{c}(x,t) + \frac{3+x}{6} \left(\frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}(-3-ct)\right) - f_{u1}^{c}(-3,t)\right) + \frac{3+x}{6} \left(\frac{c}{\alpha} - \frac{c}{\alpha} \tanh\left(\frac{c}{2\nu}(3-ct)\right) - f_{u1}^{c}(3,t)\right)$$
(C.7)

• Section 5.2.3 TFC constraints:

$$f_{u1}^{c}(x, y, t) = \mathcal{N}_{u}(x, y, t) - \frac{1}{L} \left((L - x)\mathcal{N}_{u}(0, y, t) + x\mathcal{N}_{u}(L, y, t) \right)$$
(C.8)

$$f_{u2}^{c}(x,y,t) = f_{u1}^{c}(x,y,t) - \frac{1}{H} \left((H-y)f_{u1}^{c}(x,0,t) + yf_{u1}^{c}(x,H,t) \right)$$
(C.9)

$$f_{u}^{c}(x, y, t) = f_{u2}^{c}(x, y, t) - f_{u2}^{c}(x, y, 0) + \sin\left(\frac{\pi x}{L}\right) \sin\left(\frac{\pi y}{H}\right)$$
(C.10)

Reduced TFC constraints:

$$f_u^c(x, y, t) = \mathcal{N}_u(x, y, t) \left(x - L\right) x \left(y - H\right) y t + \sin\left(\frac{\pi x}{L}\right) \sin\left(\frac{\pi y}{H}\right)$$
(C.11)

• Section 5.2.4 TFC constraints:

$$f_{u}^{c}(x, y, z, t) = \mathcal{N}_{u}(x, y, z, t) - \mathcal{N}_{u}(0, y, z, t)t^{2} + \sin(2\pi z)$$
(C.12)

$$f_{u_2}^c(x, y, z, t) = f_{u_1}^c(x, y, z, t) - f_{u_1}^c(x, 0, z, t)t^2 + \sin(2\pi z)$$
(C.13)

$$f_{u_{3}}^{c}(x, y, z, t) = f_{u_{2}}^{c}(x, y, z, t) - f_{u_{2}}^{c}(x, y, 1, t) + \sin(x^{2}y) + xy^{3/2}$$

$$+ (1 - t)(\sin(x^{2}y) + xy^{3/2}z - f_{u_{3}}^{c}(x, y, z, 0))$$

$$(C.14)$$

+
$$t\left(\sin\left(x^2y\right) + xy^{3/2}z + \sin(2\pi z) - f_{u3}^c(x, y, z, 0)\right)$$
 (C.15)

Reduced TFC constraints:

$$f_u^c(x, y, z, t) = \mathcal{N}_u(x, y, z, t) x y(z-1) t(t-1) + x y^{3/2} z + \sin\left(x^2 y\right) + t^2 \sin\left(2\pi z\right)$$
(C.16)

• Section 5.3.1 TFC constraints:

$$\begin{aligned} f_{u1}^{c}(x,y) = \mathcal{N}_{u}(x,y) &+ \frac{x_{2} - x_{1}}{x_{2} - x_{1}} \left(1 - e^{\lambda x_{2}} \cos(2\pi y) - \mathcal{N}_{u}(x_{2},y) \right) &+ \frac{x_{2} - y_{1}}{x_{2} - x_{1}} \left(1 - e^{\lambda x_{1}} \cos(2\pi y) - \mathcal{N}_{u}(x_{1},y) \right) \\ f_{u}^{c}(x,y) = f_{u1}^{c}(x,y) &+ \frac{y_{2} - y_{1}}{y_{2} - y_{1}} \left(1 - e^{\lambda x} \cos(2\pi y_{1}) - f_{u1}^{c}(x,y_{1}) \right) \\ &+ \frac{y - y_{1}}{y_{2} - y_{1}} \left(1 - e^{\lambda x} \cos(2\pi y_{2}) - f_{u1}^{c}(x,y_{2}) \right) \end{aligned}$$
(C.17)
(C.17)

$$\begin{aligned} f_{v1}^{c}(x,y) &= \mathcal{N}_{v}(x,y) + \frac{x_{2} - x}{x_{2} - x_{1}} \left(\frac{\lambda}{2\pi} e^{\lambda x_{1}} \sin(2\pi y) - \mathcal{N}_{v}(x_{1},y) \right) \\ &+ \frac{x - x_{1}}{x_{2} - x_{1}} \left(\frac{\lambda}{2\pi} e^{\lambda x_{2}} \sin(2\pi y) - \mathcal{N}_{v}(x_{2},y) \right) \\ f_{v}^{c}(x,y) &= f_{v1}^{c}(x,y) + \frac{y_{2} - y}{x_{2} - x_{1}} \left(\frac{\lambda}{2\pi} e^{\lambda x} \sin(2\pi y_{1}) - f_{v1}^{c}(x,y_{1}) \right) \end{aligned}$$
(C.19)

$$\left(\begin{array}{c} (x,y) = f_{v_1}(x,y) + \frac{y-y_1}{y_2 - y_1} \left(\frac{\lambda}{2\pi} e^{-x} \sin(2\pi y_1) - f_{v_1}(x,y_1) \right) \\ + \frac{y-y_1}{y_2 - y_1} \left(\frac{\lambda}{2\pi} e^{\lambda x} \sin(2\pi y_2) - f_{v_1}^c(x,y_2) \right) \end{array} \right)$$
(C.20)

$$f_p^c(x,y) = \mathcal{N}_p(x,y) - \mathcal{N}_p(x_1,y_1) + p_0 - \frac{1}{2}e^{2\lambda x_1}$$
(C.21)

Reduced TFC constraints:

$$\begin{aligned} f_{u1}^{c}(x,y) &= \frac{x_{2} - x}{x_{2} - x_{1}} \left(1 - e^{\lambda x_{1}} \cos(2\pi y) \right) + \frac{x - x_{1}}{x_{2} - x_{1}} \left(1 - e^{\lambda x_{2}} \cos(2\pi y) \right) \\ f_{u}^{c}(x,y) &= \mathcal{N}_{u}(x,y) \left(x - x_{1} \right) \left(x - x_{2} \right) \left(y - y_{1} \right) \left(y - y_{2} \right) + f_{u1}^{c}(x,y) \\ &+ \frac{y_{2} - y}{y_{2} - y_{1}} \left(1 - e^{\lambda x} \cos(2\pi y_{1}) - f_{u1}^{c}(x,y_{1}) \right) \end{aligned}$$
(C.22)

Communications in Nonlinear Science and Numerical Simulation 137 (2024) 108129

$$+\frac{y-y_1}{y_2-y_1}\left(1-e^{\lambda x}\cos(2\pi y_2)-f_{u_1}^c(x,y_2)\right)$$
(C.23)

$$f_{v_1}^c(x,y) = + \frac{x_2 - x}{x_2 - x_1} \left(\frac{\lambda}{2\pi} e^{\lambda x_1} \sin(2\pi y) \right) + \frac{x - x_1}{x_2 - x_1} \left(\frac{\lambda}{2\pi} e^{\lambda x_2} \sin(2\pi y) \right)$$

$$f_v^c(x,y) = \mathcal{N}_v(x,y) (x - x_1) (x - x_2) (y - y_1) (y - y_2) + f_v^c(x,y)$$
(C.24)

$$+ \frac{y_2 - y}{y_2 - y_1} \left(\frac{\lambda}{2\pi} e^{\lambda x} \sin(2\pi y_1) - f_{v_1}^c(x, y_1) \right) \\ + \frac{y - y_1}{2\pi} \left(\frac{\lambda}{2\pi} e^{\lambda x} \sin(2\pi y_2) - f_{v_1}^c(x, y_2) \right)$$
(C.25)

$$y_2 - y_1 \left(2\pi \right)$$

$$f_p^c(x, y) = \mathcal{N}_p(x, y) \left(x - x1 \right) \left(y - y1 \right) + p_0 - \frac{1}{2} e^{2\lambda x_1}$$
(C.26)

• Section 5.3.2 Reduced TFC constraints:

 $f_{u}^{c}(x, y, t) = \mathcal{N}_{u}(x, y, t) t + \sin x \cos y$ (C.27)

$$f_v^c(x, y, t) = \mathcal{N}_v(x, y, t)t - \cos x \sin y \tag{C.28}$$

$$f_{p}^{c}(x, y, t) = \mathcal{N}_{p}(x, y, t) t + \frac{\mu}{4} (\cos 2x + \sin 2y)$$
(C.29)

• Section 5.3.3 Reduced TFC constraints:

$$f_p^c(x,t) = \mathcal{N}_p(x,t)t + 1$$
 (C.30)

$$f_{u}^{c}(x,t) = \mathcal{N}_{u}(x,t)t + 1$$
 (C.31)

$$f_{\rho}^{c}(x,t) = \exp\left(\mathcal{N}_{\rho}(x,t)t - \frac{(x-\mu)^{2}}{2\sigma^{2}}\right)$$
(C.32)

Note that the form of $f_a^c(x,t)$ is not based on TFC. It imposes boundary condition at infinity as shown in Appendix B.

References

- Lagaris Isaac E, Likas Aristidis, Fotiadis Dimitrios I. Artificial neural networks for solving ordinary and partial differential equations. IEEE Trans Neural Netw 1998;9(5):987–1000.
- [2] Raissi Maziar, Perdikaris Paris, Karniadakis George Em. Physics informed deep learning (Part I): Data-driven solutions of nonlinear partial differential equations. 2017, arXiv preprint arXiv:1711.10561.
- [3] Raissi Maziar, Perdikaris Paris, Karniadakis George E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J Comput Phys 2019;378:686–707.
- [4] Dwivedi Vikas, Srinivasan Balaji. Physics informed extreme learning machine (PIELM)-A rapid method for the numerical solution of partial differential equations. Neurocomputing 2020;391:96–118.
- [5] Huang Guang-Bin, Zhu Qin-Yu, Siew Chee-Kheong. Extreme learning machine: Theory and applications. Neurocomputing 2006;70(1-3):489-501.
- [6] Dong Suchuan, Li Zongwei. Local extreme learning machines and domain decomposition for solving linear and nonlinear partial differential equations. Comput Methods Appl Mech Engrg 2021;387:114129.
- [7] Mortari Daniele, Leake Carl. The multivariate theory of connections. Mathematics 2019;7(3):296.
- [8] Leake Carl, Johnston Hunter, Mortari Daniele. The multivariate theory of functional connections: Theory, proofs, and application in partial differential equations. Mathematics 2020;8(8):1303.
- [9] Leake Carl, Mortari Daniele. Deep theory of functional connections: A new method for estimating the solutions of partial differential equations. Mach Learn Knowl Extract 2020;2(1):37–55.
- [10] Schiassi Enrico, Furfaro Roberto, Leake Carl, De Florio Mario, Johnston Hunter, Mortari Daniele. Extreme theory of functional connections: A fast physics-informed neural network method for solving ordinary and partial differential equations. Neurocomputing 2021;457:334–56.
- [11] Poole Ben, Lahiri Subhaneil, Raghu Maithra, Sohl-Dickstein Jascha, Ganguli Surya. Exponential expressivity in deep neural networks through transient chaos. In: Advances in neural information processing systems, vol. 29, 2016.
- [12] Hornik Kurt, Stinchcombe Maxwell, White Halbert. Multilayer feedforward networks are universal approximators. Neural Netw 1989;2(5):359-66.
- [13] Kidger Patrick, Lyons Terry. Universal approximation with deep narrow networks. In: Conference on learning theory. PMLR; 2020, p. 2306–27.
- [14] Margossian Charles C. A review of automatic differentiation and its efficient implementation. Wiley Interdiscip Rev: Data Min Knowl Discov 2019;9(4):e1305.
- [15] Cai Shengze, Mao Zhiping, Wang Zhicheng, Yin Minglang, Karniadakis George Em. Physics-Informed Neural Networks (PINNs) for fluid mechanics: A review. Acta Mech Sinica 2022;1–12.
- [16] Lv Chunyue, Wang Lei, Xie Chenming. A hybrid physics-informed neural network for nonlinear partial differential equation. 2021, arXiv preprint arXiv:2112.01696.
- [17] Jagtap Ameya D, Karniadakis George E. Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations. In: AAAI spring symposium: MLPS. 2021.
- [18] Mortari Daniele, Arnas David. Bijective mapping analysis to extend the theory of functional connections to non-rectangular 2-dimensional domains. Mathematics 2020;8(9):1593.
- [19] Sukumar N, Srivastava Ankit. Exact imposition of boundary conditions with distance functions in physics-informed deep neural networks. Comput Methods Appl Mech Engrg 2022;389:114333.
- [20] Dong Suchuan, Ni Naxian. A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks. J Comput Phys 2021;435:110242.
- [21] Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J, et al. LAPACK users' guide. 3rd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics; 1999.
- [22] Kingma Diederik P, Ba Jimmy. Adam: A method for stochastic optimization. 2014, arXiv preprint arXiv:1412.6980.
- [23] Fletcher Roger. Practical methods of optimization. John Wiley & Sons; 2013.

- [24] Liu Dong C, Nocedal Jorge. On the limited memory BFGS method for large scale optimization. Math Program 1989;45(1):503-28.
- [25] Paszke Adam, Gross Sam, Massa Francisco, Lerer Adam, Bradbury James, Chanan Gregory, et al. Pytorch: An imperative style, high-performance deep learning library. In: Advances in neural information processing systems, vol. 32, 2019, PyTorch Version 1.12.0. CUDA Version 10.1.
- [26] Glorot Xavier, Bengio Yoshua. 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, p. 249–56.
- [27] Jagtap Ameya D, Kawaguchi Kenji, Em Karniadakis George. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. Proc R Soc A 2020;476(2239):20200334.
- [28] Gnanasambandam Raghav, Shen Bo, Chung Jihoon, Yue Xubo, et al. Self-scalable tanh (stan): Faster convergence and better generalization in physics-informed neural networks. 2022, arXiv preprint arXiv:2204.12589.
- [29] Petzold Linda. Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations. SIAM J Sci Stat Comput 1983;4(1):136–48.
- [30] Dormand John R, Prince Peter J. A family of embedded Runge-Kutta formulae. J Comput Appl Math 1980;6(1):19-26.
- [31] Hairer Ernst, Nørsett Syvert P, Wanner Gerhard. Solving ordinary differential equations. Nonstiff problems, vol. 1, Springer-Vlg; 1993.
- [32] Shampine Lawrence F, Reichelt Mark W. The matlab ode suite. SIAM J Sci Comput 1997;18(1):1-22.
- [33] Liu Li, Liu Shengping, Yong Heng, Xiong Fansheng, Yu Tengchao. Discontinuity computing with physics-informed neural network. 2022, arXiv preprint arXiv:2206.03864.
- [34] Kovasznay Leslie I George. Laminar flow behind a two-dimensional grid. In: Mathematical proceedings of the cambridge philosophical society, vol. 44, (no. 1):Cambridge University Press; 1948, p. 58–62.
- [35] Lee Euntaek, Ahn Hyung Taek. A reconstruction-based cell-centered high-order finite volume method for incompressible viscous flow simulation on unstructured meshes. Comput & Fluids 2018;170:187–96.
- [36] Wijayanta Agung Tri, et al. Numerical solution strategy for natural convection problems in a triangular cavity using a direct meshless local Petrov-Galerkin method combined with an implicit artificial-compressibility model. Eng Anal Bound Elem 2021;126:13–29.
- [37] Wang Bo, Zhang Wenzhong, Cai Wei. Multi-scale deep neural network (mscalednn) methods for oscillatory stokes flows in complex domains. 2020, arXiv preprint arXiv:2009.12729.
- [38] Lou Qin, Meng Xuhui, Karniadakis George Em. Physics-informed neural networks for solving forward and inverse flow problems via the Boltzmann-BGK formulation. J Comput Phys 2021;447:110676.
- [39] Taylor Geoffrey Ingram, Green Albert Edward. Mechanism of the production of small eddies from large ones. Proc R Soc Lond Ser A 1937;158(895):499–521.
- [40] Sod Gary A. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. J Comput Phys 1978;27(1):1-31.
- [41] Mao Zhiping, Jagtap Ameya D, Karniadakis George Em. Physics-informed neural networks for high-speed flows. Comput Methods Appl Mech Engrg 2020;360:112789.
- [42] Papados Alexandros. Solving hydrodynamic shock-tube problems using weighted physics-informed neural networks with domain extension.
- [43] Moré Jorge J. The Levenberg-Marquardt algorithm: Implementation and theory. In: Numerical analysis. Springer; 1978, p. 105–16.
- [44] Qin Youwei, Kavetski Dmitri, Kuczera George. A robust Gauss-Newton algorithm for the optimization of hydrological models: From standard Gauss-Newton to robust Gauss-Newton. Water Resour Res 2018;54(11):9655–83.
- [45] Larson Jeffrey, Menickelly Matt, Wild Stefan M. Derivative-free optimization methods. Acta Numer 2019;28:287-404.
- [46] Aly Ahmed, Guadagni Gianluca, Dugan Joanne B. Derivative-free optimization of neural networks using local search. In: 2019 IEEE 10th annual ubiquitous computing, electronics & mobile communication conference. IEEE; 2019, p. 0293–9.
- [47] Chen Xiangyi, Liu Sijia, Xu Kaidi, Li Xingguo, Lin Xue, Hong Mingyi, et al. Zo-adamm: Zeroth-order adaptive momentum method for black-box optimization. Adv Neural Inf Process Syst 2019;32.
- [48] Han Jihun, Nica Mihai, Stinchcombe Adam R. A derivative-free method for solving elliptic partial differential equations with deep neural networks. J Comput Phys 2020;419:109672.
- [49] Della Santa Francesco, Pieraccini Sandra. Discontinuous neural networks and discontinuity learning. J Comput Appl Math 2023;419:114678.
- [50] PINN-extremization. 2024, https://github.com/keygenx/PINN-Extremization. [Accessed: 23 May 2024].