



Can neural networks learn finite elements?

Julia Novo ^{a,*}, Eduardo Terrés

^a Departamento de Matemáticas, Universidad Autónoma de Madrid, Spain

ARTICLE INFO

Dataset link: <https://github.com/EduardoTerres/Can-Neural-Networks-learn-Finite-Elements.git>

Keywords:

Neural networks
Finite elements
Partial differential equations

ABSTRACT

The aim of this note is to construct a neural network for which the linear finite element approximation of a simple one dimensional boundary value problem is a minimum of the cost function to find out if the neural network is able to reproduce the finite element approximation. The deepest goal is to shed some light on the problems one encounters when trying to use neural networks to approximate partial differential equations.

1. Introduction

For a brief introduction to neural networks and their applications to the numerical approximation of partial differential equations we refer the reader to reviews [1,2]. For the origins of the finite element method and its influence in modern computing we refer to review [3].

The literature concerning the use of neural networks to approach partial differential equations has increased significantly in recent years. Most of the results are heuristic and the ability of neural networks to perform better than classical methods is not always clear. With this short note we would like to shed some light on the difficulties one encounters when trying to define a neural network to approach a given problem better than classical finite element methods. To this end, we choose a very simple model problem, that nevertheless allows us to understand some general characteristics of neural networks.

When choosing a neural network to approach a problem one has to decide its structure and activation function. The structure of the network, a priori unknown, has a great impact on the obtained results. The use of different activation functions also leads to different approaches. Another important question is the theoretical ability of a neural network to approximate a function arbitrarily well, see [4].

In this note, following [5], we choose the structure of the neural network in such a way that for certain values of the weights and biases the output of the network is the linear finite element approximation to the model problem. The activation function is the so called ReLU. With this structure, certain theoretical capabilities are guaranteed, given that piecewise linear approximations over a certain partition can be exactly represented by the neural network. In particular, the response of the network could be the finite element approximation itself, with well-known approximation properties. However, we observed a first problem. Even in this case in which the structure of the network is known, the result of the optimization problem that finds the values of the weights and biases can lead to an approximation far away from the linear finite element approximation that a priori gives a minimum of the optimization problem. The idea behind this fact is very simple: the problem we are trying to solve is highly over-determined, thus it is difficult to find the optimal value. On a second try, we add some initial information to the neural network concerning

* Corresponding author.

E-mail addresses: julia.novo@uam.es (J. Novo), eduardo.terres@estudiante.uam.es (E. Terrés).

¹ Research is supported by grant PID2022-136550NB-I00 funded by MCIN/AEI/10.13039/501100011033 and by ERDF A way of making Europe, by the European Union.

the values of weights and biases to lead somehow the answer toward the output we know is a minimum of the problem. As we show in the numerical experiments, we consider several levels in the procedure of adding information. However, if we know how to compute the linear finite element approximation in a deterministic way it probably does not make sense to use neural networks to get this approximation. So, the question is, can we find a balance between the deterministic information we supply and the free values we keep (output values of the optimization problem) so that the resulting approximation has the same or, if possible, better approximation capabilities than the standard finite element approximation? This question can perhaps be re-formulated in the following way: can physics-informed, or similarly defined, neural networks, [6], give better approximations than standard methods? Our conclusion is that treating neural networks as black boxes, without providing sufficient domain-specific insights does not offer guarantees for success. At least with the procedure we followed, in our simple model problem we do not obtain, in general, better approximations than the finite element method. Moreover, we only obtain neural network approximations close to those of the finite element method in our last try, in which we have fixed so many weights and biases, that the response of the neural network, regardless of the output values of the optimization problem, is always a continuous piecewise linear approximation with respect to the partition in which the linear finite element approximation is defined.

This is not the first reference to pose the above question. In [7], the authors make a deep comparison between PINNs and finite elements. Their study suggests that for certain classes of PDEs for which classical methods are applicable, PINNs do not produce better results. As the authors of [7] state, PINNs could however be efficient in high-dimensional problems for which classical techniques are prohibitively expensive and when combining PDEs and data. We also want to mention the recent paper [8] in which the author considers a comparison of fast multigrid solvers for partial differential equations and corresponding attempts with neural networks.

The outline of the paper is as follows. In Section 2 we introduce the model problem (based on a one dimensional convection–diffusion equation) and describe the neural networks that mimic finite elements. In Section 3 we present a series of numerical experiments in which we distinguish two cases, corresponding to diffusion dominated and convection dominated examples. In the convection dominated case, instead of the standard linear finite element approximation we try to reproduce the stabilized streamline-upwind Petrov–Galerkin (SUPG) linear finite element approximation [9] for which a different cost function from that of the standard method is used.

2. Model problem. Neural networks that mimic finite elements

As a model problem we consider the following steady convection–diffusion problem in one dimension. Find $u \in C^2(0, 1) \cap C[0, 1]$ such that

$$\begin{aligned} -\epsilon u'' + u' &= 1, \quad x \in (0, 1), \\ u(0) &= u(1) = 0, \end{aligned} \quad (1)$$

where ϵ is a positive parameter. The weak solution of (1) is: find $u \in H_0^1(0, 1)$ such that

$$\epsilon(u', v') + (u', v) = (1, v), \quad \forall v \in H_0^1(0, 1), \quad (2)$$

where (\cdot, \cdot) represents the $L^2(0, 1)$ inner product and we have used standard notation for the Sobolev space $H_0^1(0, 1)$. For any integer $N > 0$ we consider a partition

$$\tau_N = \{0 = x_0 < x_1 < \dots < x_N = 1\}, \quad (3)$$

of $[0, 1]$ and denote by $h_i = x_{i+1} - x_i$. We denote by V_h the linear finite element space of continuous functions that are piecewise linear with respect to the partition τ_N and satisfy the homogeneous Dirichlet boundary conditions of the problem. Then, the linear finite element approximation to the model problem (1) satisfies the weak form (2) over V_h : find $u_h \in V_h$ such that

$$\epsilon(u_h', v_h') + (u_h', v_h) = (1, v_h), \quad \forall v_h \in V_h. \quad (4)$$

With the standard nodal Lagrange basis functions, $\phi_i \in V_h$, defined by $\phi_i(x_j) = \delta_{ij}$ we can write $u_h(x) = \sum_{i=1}^{N-1} u_h(x_i) \phi_i(x)$. Then, u_h is unique and its vector of coefficients $\mathbf{u} = [u_h(x_1), \dots, u_h(x_{N-1})]^T$ is the unique solution of the linear system $\mathbf{A}\mathbf{u} = \mathbf{f}$, with $\mathbf{A} = \epsilon\mathbf{K} + \mathbf{C}$ and $\mathbf{f} = [f_1, \dots, f_{N-1}]^T$ with $f_i = (h_{i-1} + h_i)/2$. The matrices \mathbf{K} and \mathbf{C} are tridiagonal and its elements are:

$$\begin{aligned} k_{i,i-1} &= -\frac{1}{h_{i-1}}, \quad k_{i,i} = \frac{1}{h_{i-1}} + \frac{1}{h_i}, \quad k_{i,i+1} = -\frac{1}{h_i} \\ c_{i,i-1} &= -\frac{1}{2}, \quad c_{i,i} = 0, \quad c_{i,i+1} = \frac{1}{2}. \end{aligned} \quad (5)$$

Let us denote by σ the ReLU activation function defined as $\sigma(x) = x$ for $x \geq 0$ and $\sigma(x) = 0$ for $x < 0$. Then, following [5], we can write

$$\phi_i(x) = \sigma\left(\frac{x - x_{i-1}}{h_{i-1}}\right) - \sigma\left((x - x_i)\left(\frac{1}{h_{i-1}} + \frac{1}{h_i}\right)\right) + \sigma\left(\frac{x - x_{i+1}}{h_i}\right).$$

Following the notation in [1], any of the basis functions can be written as a neural network of the form

$$\phi_i(x) = W_i^{[3]} \sigma(W_i^{[2]} x + b_i^{[2]}), \quad i = 1, \dots, N-1, \quad (6)$$

where

$$W_i^{[2]} = \begin{pmatrix} \frac{1}{h_{i-1}} \\ \frac{1}{h_{i-1}} + \frac{1}{h_i} \\ \frac{1}{h_i} \end{pmatrix}, \quad b_i^{[2]} = \begin{pmatrix} \frac{-x_{i-1}}{h_{i-1}} \\ \frac{1}{h_{i-1}} + \frac{1}{h_i} \\ \frac{-x_{i+1}}{h_i} \end{pmatrix}, \quad (7)$$

$$W_i^{[3]} = (1, -1, 1). \quad (8)$$

We now define a neural network that mimics the finite element approximation.

Definition 1. Let $F : \mathbb{R} \rightarrow \mathbb{R}$, be a neural network with three layers where layer 1 is the Input layer and layer 3 is the Output layer, and layers 1, 2 and 3 have 1, $3(N-1)$ and 1 neurons, respectively. More precisely, for $W^{[2]} \in \mathbb{R}^{3(N-1) \times 1}$, $b^{[2]} \in \mathbb{R}^{3(N-1)}$ and $W^{[3]} \in \mathbb{R}^{1 \times 3(N-1)}$,

$$F(x) = W^{[3]} \sigma(W^{[2]}x + b^{[2]}).$$

Proposition 1. Let us define $W^{[2]}$, $b^{[2]}$ and $W^{[3]}$ by:

$$W^{[2]} = \begin{pmatrix} W_1^{[2]} \\ \vdots \\ W_i^{[2]} \\ \vdots \\ W_{N-1}^{[2]} \end{pmatrix}, \quad b^{[2]} = \begin{pmatrix} b_1^{[2]} \\ \vdots \\ b_i^{[2]} \\ \vdots \\ b_{N-1}^{[2]} \end{pmatrix}, \quad W^{[3]} = \begin{pmatrix} u_h(x_1) W_1^{[3]} \\ \vdots \\ u_h(x_i) W_i^{[3]} \\ \vdots \\ u_h(x_{N-1}) W_{N-1}^{[3]} \end{pmatrix}^T$$

where $W_i^{[2]}$, $b_i^{[2]}$ and $W_i^{[3]}$ are defined in . Then $F \equiv u_h$.

Proof. The proof is immediate using the definition of $u_h(x) = \sum_{i=1}^{N-1} u_h(x_i) \phi_i(x)$, the expression of the basis functions as neural networks (6) and Definition 1.

Proposition 1 says that the neural network F mimics the finite element approximation in the sense that there exist values of the weights and biases for which we recover the finite element approximation. We observe that there are different values of weights and biases that satisfy the property $F \equiv u_h$ since $\sigma(\alpha x) = \alpha \sigma(x)$ for any $\alpha > 0$.

Once the structure of the neural network is chosen (depending only on the value N), we still need to define a cost function. Since we want the network to mimic the finite element approximation, we define the cost function as follows

$$\begin{aligned} \text{Cost} = \sum_{i=1}^{N-1} & \left(-\epsilon \left[-\frac{1}{h_{i-1}} F(x_{i-1}) + \left(\frac{1}{h_{i-1}} + \frac{1}{h_i} \right) F(x_i) - \frac{1}{h_i} F(x_{i+1}) \right] \right. \\ & \left. + \frac{1}{2} [F(x_{i+1}) - F(x_{i-1})] - \frac{1}{2} (h_{i-1} + h_i) \right)^2 + F(0)^2 + F(1)^2. \end{aligned} \quad (9)$$

It is easy to check that choosing $F(x_i) = u_h(x_i)$, $i = 0, \dots, N$ then $\text{Cost} = 0$. Then, according to Proposition 1, there exist weights and biases for which there is a minimum of the optimization problem.

It is well-known that in the convection dominated regime (in our case ϵ small in (1)) the standard finite element approximation produces spurious oscillations unless the mesh size is sufficiently small. Stabilized approximations are usually applied to mitigate this problem. The so called SUPG method is one of the most popular finite element stabilizations, [9]. Then, in the convection dominated regime instead of (9) we consider

$$\begin{aligned} \text{Cost} = \sum_{i=1}^{N-1} & \left(-\epsilon \left[-\frac{1}{h_{i-1}} F(x_{i-1}) + \left(\frac{1}{h_{i-1}} + \frac{1}{h_i} \right) F(x_i) - \frac{1}{h_i} F(x_{i+1}) \right] \right. \\ & \left. + [F(x_i) - F(x_{i-1})] - \frac{1}{2} (h_{i-1} + h_i) \right)^2 + F(0)^2 + F(1)^2. \end{aligned} \quad (10)$$

It is easy to check that $F(x_i) = u_h^s(x_i)$, $i = 0, \dots, N$, with u_h^s the linear finite element SUPG approximation to problem (1) with stabilization parameter $\delta_i = h_{i-1}/2$ in (x_{i-1}, x_i) , satisfies $\text{Cost} = 0$.

3. Numerical experiments

We consider the neural network of Definition 1. In the numerical experiments we use the gradient descent method with back propagation. In the sequel for simplicity we consider a uniform partition (3) for different values of N . We denote by η , the learning rate, and N_{iter} the maximum number of iterations in the gradient descent method. We have tried to add a regularization term to the Cost function of the form $\beta \|w\|^2$, where $\|w\|^2$ represents the L^2 norm of the vector that contains all the weights and biases. However, most of the experiments are done with $\beta = 0$, since a small value of β does not affect the results and by increasing the value of β we get worse results. We have carried out two types of experiments. In the diffusion dominated regime we consider the cost function (9). In the convection dominated regime we consider the cost function (10).




Fig. 1. $N = 40$, $N_{\text{iter}} = 10^4$, $\eta = 10^{-4}$, $\beta = 10^{-4}$.




Fig. 2. $N = 20$, $N_{\text{iter}} = 3 \times 10^5$, $\eta = 10^{-6}$, $\beta = 0$.

3.1. Diffusion dominated regime

We fix the value of $\epsilon = 0.1$ along this section. In a first experiment all the parameters, weights and biases, in the neural network are free and we initialize them with random values. We observe that, while the finite element approximation has $N - 1$ free parameters, our neural network has $9(N - 1)$. In Fig. 1 we have plotted, on the left, the neural network approximation in blue and the exact solution in red and, on the right, the values of the cost function in blue and the L^2 error between the exact solution and the neural network approximation in black. Let us observe that we use different scales on the left and right axis for the picture on the right. The left axis is the scale for the cost function while the right axis is the scale for the error. After less than 1000 iterations the value of the cost function stabilizes at around 0.5. With this procedure the neural network is not able to reach a minimum of the cost function. The main issue the network faces here is solving a highly over-determined problem. We have tried different values of N , N_{iter} , η and β without improving (even worsening) the results shown in Fig. 1.

In a second experiment, we try to add some initial information to the neural network with the aim of improving the results obtained in the previous failed attempt. Along this line, we initialize $W^{[2]}$ and $b^{[2]}$ with the values and, as before, we use random values for $W^{[3]}$. After initialization, the neural network is free to adjust the values of $W^{[2]}$ and $b^{[2]}$ in the pursuit of a configuration that minimizes the cost function and, potentially, outperforms the finite element solution. Let us observe that we can write the neural network as $F(x) = \sum_{i=1}^{3(N-1)} w_i^3 \sigma(w_i^2 x + b_i^2)$ which means that $F(x)$ is a piecewise linear function over a mesh with nodes $-b_i^2/w_i^2$. If all the nodes are positive then $F(0) = 0$ while $F(1) = \sum_{i=1}^{3(N-1)} w_i^3 \sigma(w_i^2 + b_i^2)$. In the experiments we have observed that $F(0) = 0$ while $F(1)$ is small (not strongly imposed). In Fig. 2 we have plotted on the left the approximation we have obtained and on the right the values of the cost function and the error in the neural network, as the number of iterations increases. The experiment is done for $N = 20$ and we can observe (see the vertical lines on the left picture corresponding to values of the nodes $x_j = j/20$, $j = 1, \dots, 19$) that the method gives an approximation with small error at the nodes of the partition $\tau_{20} = \{x_j\}_{j=1}^{19}$. As expected, the approximation is piecewise linear but over a mesh different from τ_{20} , which means that the neural network has modified (at least some) of the initial given values for $W^{[2]}$ and $b^{[2]}$. The fact that the neural network is accurate over τ_{20} but has a big error in the extra nodes makes sense since the cost function is defined over the nodes of τ_{20} and does not have any added information about other extra nodes. Let us observe that the cost function (9) is blind for the nodes not belonging to τ_{20} which means that the spurious oscillations of the neural network approximation at the extra nodes cannot be avoided with this algorithm. At this point it would be interesting to be able to add some information to the neural network (cost function) in order to produce a better approximation than u_h . Taking into account that the response $F(x)$ is a piecewise linear approximation over a mesh with $3(N - 1)$ nodes while u_h is based on $N - 1$ this could be, in principle, possible. However, our aim in this experiment is to check whether a neural network for which the linear finite element approximation is a minimum of the cost function is able by itself to reproduce (learn) the linear finite element approximation.




Fig. 3. $N = 40$, $N_{\text{iter}} = 5 \times 10^5$, $\eta = 10^{-7}$ and $\beta = 0$ on the left and $N = 100$, $N_{\text{iter}} = 3 \times 10^5$, $\eta = 10^{-8}$ and $\beta = 0$ on the right.




Fig. 4. $N = 20$, $N_{\text{iter}} = 2 \times 10^5$, $\eta = 10^{-6}$, $\beta = 0$.

Table 1

Errors in the finite element and neural network approximations.

	L^2 FEM	H^1 FEM	L^2 NN	H^1 NN
$N = 20$	3.87e-3	3.20e-1	3.87e-3	3.20e-1
$N = 40$	9.72e-4	1.61e-1	9.72e-4	1.61e-1
$N = 100$	1.56e-4	6.45e-2	3.30e-4	2.60e-1

In Fig. 3 we have represented on the left the approximation corresponding to $N = 40$ with a smaller value of the learning rate $\eta = 10^{-7}$. As before, we mark with vertical lines the nodes of the uniform partition τ_{40} , at which the neural network gives a small error. On the right of Fig. 3 we have represented the approximation for $N = 100$. Spurious oscillations appear in both approximations with smaller local maximum and minimum values in the last case: $N = 100$.

In our third experiment we try to lead the network to the linear finite element approximation in a stronger way. To this end, we fix the values of $W^{[2]}$ and $b^{[2]}$ to those in and keep only $3(N - 1)$ free parameters, those in $W^{[3]}$. We observe that in this case our neural network can be written as

$$F(x) = \sum_{i=1}^{N-1} w_i^{3,1} \sigma\left(\frac{x - x_{i-1}}{h_{i-1}}\right) + w_i^{3,2} \sigma\left((x - x_i)\left(\frac{1}{h_{i-1}} + \frac{1}{h_i}\right)\right) + w_i^{3,3} \sigma\left(\frac{x - x_{i+1}}{h_i}\right). \quad (11)$$

This means that the response of the network is a continuous piecewise linear function with respect to the original partition τ_N that satisfies in a strong way the boundary condition $F(0)$. In our numerical experiments we have plotted the resulting functions $F(x)$ at different stages of the iterative algorithm (different values of the number of iterations) and we have observed that the main effort of the neural network seems to be to enforce the condition $F(1) = 0$.

In Fig. 4 we can see the neural network approximation on the left and the cost function and error of the approximation on the right for $N = 20$. In this case the final approximation generated by the neural network is a good approximation and essentially indistinguishable from the linear finite element approximation u_h . We have obtained similar results for $N = 40$ and $N = 100$. In Fig. 5 we plot the absolute value of the difference between the linear finite element and the neural network approximations for $N = 20, 40$ and 100 . We observe that in all cases the error between both approximations increases from a point around $x = 0.6$ and has a maximum at $x = 1$. The errors at $x = 1$ are around 10^{-11} , 10^{-8} and 10^{-3} for $N = 20, 40$ and 100 , respectively.

As we can also observe in Table 1, in which we have represented the errors in the finite element and neural network approximations, both approximations give the same errors for $N = 20$ and $N = 40$, both in L^2 and H^1 norms, while for $N = 100$ the finite element approximation produces smaller errors. In our experiments we observed that for $N = 100$, both the cost function and the error are decreasing functions (they do not stabilize). We stopped the algorithm after 5×10^5 iterations. We guess that by




Fig. 5. Absolute value of errors between finite element and neural network approximations for $N = 20, 40$ and 100 .

Table 2

Errors in the SUPG finite element and neural network approximations.

	L^2 SUPG	H^1 SUPG	L^2 NN	H^1 NN
$N = 20$	1.25e-1	21.91	1.25e-1	21.91
$N = 40$	8.52e-2	21.45	5.79e-2	42.67
$N = 100$	4.87e-2	20.04	2.18e-2	134.37




Fig. 6. $N = 20$, $N_{\text{iter}} = 10^6$, $\eta = 10^{-6}$, $\beta = 0$.

increasing the number of iterations, the network could have finally reached the error of the finite element approximation as in the previous cases: $N = 20$ and $N = 40$. Our aim is, however, not to reproduce the linear finite element approximation that can be obtained solving the linear system described in (4)–(5) but find out how much information needs an a priori blind neural network with a cost function for which the linear finite element is a minimum to learn (or reproduce) this approximation. Since, in general, one does not even know the structure of the neural network, with this simple experiment we just want to stand out the problems that Physical Informed Neural Networks (or similarly based networks) may have in obtaining good approximations to partial differential equations.

3.2. Convection dominated regime

We have reproduced the previous experiments in the convection dominated regime. In this section we take $\epsilon = 0.001$ and, as before, we consider a uniform partition (3) for different values of N . In this case, instead of (9) we use the cost function (10). In view of the previous results we did not expect to get much different results (probably worse) in this more difficult problem.

As before, we do not obtain a reasonable approximation for the case in which all the weights and biases are free. In the second case, in which we initialize $W^{[2]}$ and $b^{[2]}$ as in (9), but then we allow the neural network to change those values, we obtained similar results with some peculiarities. The case $N = 20$ is indeed very similar. As we can observe in Fig. 6, the neural network approximation is a piecewise linear approximation over a refined partition of τ_{20} . Again, the neural network approximation is only a good approximation to the true solution in the nodes of τ_{20} .

In Fig. 7 we have represented on the left the neural network approximation for $N = 40$ and on the right the one corresponding to $N = 100$. Here, we find some differences with respect to the diffusion dominated regime. We observe that in this case the errors at the nodes of the underlying uniform partition are larger than in the diffusion dominated case and also that the error at $x = 1$ is very large. This means that the neural network fails at imposing the condition $F(1) = 0$. As it is known, there is a boundary layer at $x = 1$ in this extremely simple convection dominated problem.

To conclude, we fix $W^{[2]}$ and $b^{[2]}$ as in (9) and leave $W^{[3]}$ free to check if the network is able to reproduce the SUPG linear finite element approximation. In Table 2 we have collected the errors of the SUPG and neural network approximations. As expected, the errors are larger than those in Table 1 for the diffusion dominated case, in particular all the errors in H^1 are considerably large.




Fig. 7. $N = 40$, $N_{\text{iter}} = 10^6$, $\eta = 10^{-7}$ and $\beta = 0$ on the left and $N = 100$, $N_{\text{iter}} = 10^6$, $\eta = 10^{-8}$ and $\beta = 0$ on the right.




Fig. 8. $N = 40$, $N_{\text{iter}} = 2 \times 10^6$, $\eta = 10^{-7}$, $\beta = 0$.




Fig. 9. $N = 100$, $N_{\text{iter}} = 4 \times 10^5$, $\eta = 10^{-9}$, $\beta = 0$.

While for $N = 20$ the network is able to reproduce the SUPG approximation, similarly to how it reproduced the FE approximation in the diffusion dominated case (compare with Table 1), the situation changes for $N = 40$ and $N = 100$. For these values of N , the L^2 errors are smaller in the neural network approximation, while the H^1 errors are considerably larger. We believe that the explanation for this is the weak imposition of the right homogeneous Dirichlet boundary condition at $x = 1$. While in the SUPG method this condition is strongly imposed, in the neural network is not. In fact, the neural network approximations fail to impose the condition $F(1) = 0$ within the specified number of iterations in both cases. To study this in more detail, we have represented the results of the experiments in Figs. 8 and 9. It is surprising to see that, although the approximations are quite accurate (except for the values at $x = 1$ being far from 0 in both cases), the values of the cost function are large. For $N = 40$ (Fig. 8) the cost has not yet stabilized, hence we believe that with enough iterations the neural network would have eventually approached the SUPG approximation, leading to a further increase in the L^2 error. For $N = 100$ (Fig. 9), the value of the cost function stabilizes at around 1, meaning that the neural network is not able to reproduce the SUPG approximation. However, it yields a better approximation in terms of the L^2 error than SUPG, even though the value of the cost function is large. Since the limit solution ($\epsilon = 0$) of problem (1) is the function $G(x) = x$, which satisfies only the left boundary condition, the reason for having a boundary layer at $x = 1$ is the strong imposition of the right boundary condition. As we can see in the left picture of Fig. 9, the neural network approximation for $N = 100$ is close to the function $G(x) = x$. On the other hand, $G(x) = x$ is also close in L^2 to the solution of problem (1) for the selected value of ϵ . Observe that function $G(x) = x$ can be expressed as a neural network of the form (11) by taking all parameters equal to zero except for $w_1^{3,1} = h$. Actually, after the training process, the neural network presents a value of $w_1^{3,1} = 0.0098 \approx h = 0.01$, and

most of the weights in $W^{[3]}$ are close to zero ($< 10^{-10}$). Thus, in this particular example, the weak imposition of the right boundary condition together with the particular form of the output of the neural network seems to make the neural network approach the limit solution of (1).

4. Conclusions

In this note we have constructed a neural network for which the linear finite element approximation of a simple one dimensional boundary value problem is a minimum of the cost function to find out if the neural network is able to reproduce the finite element approximation. In the first attempt, the optimization problem entails too many free parameters, rendering it highly over-determined and making it difficult for the network to find the optimal values. In this case, the approximation is far away from the linear finite element approximation and the theoretical solution. In subsequent steps, we shift the focus into finding out how much information an a priori blind neural network with a cost function for which the linear finite element is a minimum needs to correctly learn this approximation. With this simple example we just want to stand out the problems that PINNs (or similarly based networks) may have in obtaining good approximations to partial differential equations since, in general, one does not even know the structure of the neural network. Our conclusion is that in our simple model problem we do not obtain, in general, better approximations than the finite element method. In fact, we only obtain neural network approximations close to those of the finite element method in our last try, where we have fixed so many weights and biases that the response of the neural network, regardless of the output values of the optimization problem, is always a continuous piecewise linear approximation with respect to the partition in which the linear finite element approximation is defined.

Our conclusions are in agreement with the deeper study of [7] in which the authors compare PINNs and finite elements through various linear and nonlinear partial differential equations. Their study suggests that for certain classes of PDEs for which classical methods are applicable, PINNs are not able to outperform these. However, we agree with the authors of [7] that PINNs could be efficient in high-dimensional problems for which classical techniques are prohibitively expensive and when combining both PDEs and data.

Data availability

The MATLAB code for the training algorithms can be found in the following GitHub repository: <https://github.com/EduardoTerrés/Can-Neural-Networks-learn-Finite-Elements.git>.

References

- [1] C.F. Higham, D.J. Higham, Deep learning: an introduction for applied mathematicians, *SIAM Rev.* 61 (2019) 860–891.
- [2] L. Lu, X. Meng, Z. Mao, G.E. Karniadakis, DeepXDE: A deep learning library for solving differential equations, *SIAM Rev.* 63 (2021) 208–228.
- [3] M.J. Gander, G. Wanner, From Euler, Ritz, and Galerkin to modern computing, *SIAM Rev.* 54 (2012) 627–666.
- [4] A. Pinkus, Approximation theory of the MLP model in neural networks, *Acta Numer.* 8 (1999) 143–195.
- [5] J. He, L. Li, J. Xu, Ch. Zheng, Relu deep neural networks and linear finite elements, *J. Comput. Math.* 38 (2020) 502–527.
- [6] S. Mishra, R. Molinaro, Estimates of the generalization error of physics-informed neural networks for approximating PDEs, *IMA J. Numer. Anal.* 43 (2023) 1–43.
- [7] T.G. Grossmann, U.J. Komorowska, J. Latz, C.-B. Schonlieb, Can physics-informed neural networks beat the finite element method?, 2023, <http://dx.doi.org/10.48550/arXiv.2302.04107>.
- [8] U. Rüde, Fast solvers! when natural intelligence counts, in: *Conference Trends in Scientific Computing, Dortmund, 2024*.
- [9] A.N. Brooks, T.J.R. Hughes, Streamline upwind/Petrov–Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier–Stokes equations, *Comput. Methods Appl. Mech. Engrg.* 32 (1982) 199–259.