

ODE \mathcal{N} : A Framework to Solve Ordinary Differential Equations using Artificial Neural Networks

Liam L. H. Lau^{§ 1, *} and Denis Werth^{§ 2, †}

¹*Gonville & Caius College, Trinity Street, Cambridge, CB2 1TA, UK*

²*École Normale Supérieure Paris-Saclay, Department of Physics, France*

We explore in detail a method to solve ordinary differential equations using feedforward neural networks. We prove a specific loss function, which does not require knowledge of the exact solution, to be a suitable standard metric to evaluate neural networks' performance. Neural networks are shown to be proficient at approximating continuous solutions within their training domains. We illustrate neural networks' ability to outperform traditional standard numerical techniques. Training is thoroughly examined and three universal phases are found: (i) a prior tangent adjustment, (ii) a curvature fitting, and (iii) a fine-tuning stage. The main limitation of the method is the nontrivial task of finding the appropriate neural network architecture and the choice of neural network hyperparameters for efficient optimization. However, we observe an optimal architecture that matches the complexity of the differential equation. A user-friendly and adaptable open-source code (ODE \mathcal{N}) is provided on GitHub.

I. INTRODUCTION

Neural networks (NNs) are known to be powerful tools due to their role as universal continuous function approximators [1]. Rewarding NNs for their improving performances in specific tasks can be done by minimizing a specially designed loss function. In fact, the last ten years have seen a surge in the capabilities of NNs, specifically deep neural networks, due to the increase in supervised data and compute. For example, one can cite AlexNet [2], a convolutional neural network for image recognition and GPT-2 [3], a powerful Transformer language model.

Although NNs and machine learning are used in physics [4–13], the use of NNs has still not been widely accepted in the scientific community due to their black-box nature, the unclear and often unexplored effects of hyperparameter choice for a particular problem, and a large computational cost for training.

Usually, neural networks are used for problems involving dimensional reduction [14], data visualization, clustering [15], and classification [16]. Here, we use the potential of NNs to solve differential equations. Differential equations are prevalent in many disciplines including Physics, Chemistry, Biology, Economics, and Engineering. When an analytical solution cannot be found, numerical methods are employed and are very successful [17–19]. However, some systems exhibit differential equations that are not efficiently solved by usual numerical methods. Such differential equations can be numerically solved to the required accuracy by novel specific codes, such as `oscode` [20] which efficiently solves one-

dimensional, second-order, ordinary differential equations (ODEs) with rapidly oscillating solutions which appear in cosmology and condensed matter physics [21] among other systems. Yet, specifically written algorithms [22] and codes require a lot of time and resources.

In this paper, we ask whether the unsupervised method of solving ODEs using NNs, first developed by Lagaris et al. [23] and improved in [24–31], is robust. In previous investigations, there has been no clear consensus on what exact method should be used. Multiple loss functions were used, sometimes requiring the exact solution to be known [32] [33], and the boundary/initial conditions were either included in the model as an extra term in the loss function or treated independently by means of trial functions. No explicit benchmark to evaluate the NN performances was transparently introduced. Training and extrapolation performances were not meticulously studied, making NNs appear as black-boxes.

This work aims at filling the mentioned gaps and demystifying a few aspects. We provide a clear step-by-step method to solve ODEs using NNs. We prove a specific loss function from the literature [24] to be an appropriate benchmark metric to evaluate the NN performances, as it does not require the exact solution to be known. The training process is investigated in detail and is found to exhibit three phases that we believe are the general method that the NN learns to solve a given ODE with boundary/initial conditions. We explore the effect of training domain sampling, extrapolation performances and NN architectures. Contrary to numerical integration, approximating the ODE solution by a NN gives a function that can be continuously evaluated over the entire training domain. However, finding an appropriate NN architecture to reach the desired accuracy requires extensive testing and intuition. An open-source and user-

[§] These authors contributed equally to this work.

* lhl2@cam.ac.uk

† denis.werth@ens-paris-saclay.fr

friendly code (ODE \mathcal{N}) that accompanies this manuscript² is provided, along with animations that help to visualize the training process.

In the next section, we briefly present the basics of NNs, followed by the method used to solve ODEs with NNs in Section III. In Section IV, three different ODEs are solved as illustration. A specific loss function is proved to be a universal metric to assess the accuracy and different simple NN models are tested in Section V, particularly highlighting the role of complexity when choosing the NN. Finally, conclusion and outlook are provided in Sections VI.

II. BACKGROUND

The basic unit of a NN is a neuron i that takes a vector of input features $\mathbf{x} = (x_1, x_2, \dots)$ and produces a scalar output $a_i(\mathbf{x})$. In almost all cases, a_i can be decomposed into a linear operation taking the form of a dot product with a set of weights $\mathbf{w}^{(i)} = (w_1^{(i)}, w_2^{(i)}, \dots)$ followed by re-centering with an offset called the bias $b^{(i)}$, and a non-linear transformation *i.e.* an activation function $\sigma_i : \mathbb{R} \rightarrow [0, 1]$ which is usually the same for all neurons. One can write the full input-output formula for one neuron as follows

$$a_i(\mathbf{x}) = \sigma_i(\mathbf{w}^{(i)} \cdot \mathbf{x} + b^{(i)}). \quad (1)$$

A NN consists of many such neurons stacked into layers, with output of one layer serving as input for the next. Thus, the whole NN can be thought of as a complicated non-linear transformation of inputs x into an output \mathcal{N} that depends on the weights and biases $\boldsymbol{\theta}$ of all the neurons in the input, hidden, and output layers [34]:

$$\text{Neural Network} = \mathcal{N}(x, \boldsymbol{\theta}). \quad (2)$$

The NN is trained by finding the value of $\boldsymbol{\theta}$ that minimizes the loss function $\mathcal{L}(\boldsymbol{\theta})$, a function that judges how well the model performs on corresponding the NN inputs to the NN output. Minimization is usually done using a gradient descent method [34]. These methods iteratively adjust the NN parameters $\boldsymbol{\theta}$ in the direction (on the parameter surface) where the gradient \mathcal{L} is large and negative. In this way, the training procedure ensures the parameters $\boldsymbol{\theta}$ flow towards a local minimum of the loss function. The main advantage of NNs is to exploit their layered structure to compute the gradients of \mathcal{L} in a very efficient way using backpropagation [34].

III. METHOD

The starting point of solving ODEs using NNs is the universal approximation theorem. The theorem states that a feedforward NN with hidden layers containing a finite number of neurons can approximate any continuous functions at any level of accuracy [1]. Thus, one may expect NNs to perform well on solving ODEs. Despite the existence of such NNs, the theorem does not provide us with a recipe to find these NNs nor touch upon their learnabilities.

Differential equations with an unknown solution f can be written as a function \mathcal{F} of f and their derivatives in the form

$$\mathcal{F}[x, f(x), \nabla f(x), \dots, \nabla^p f(x)] = 0 \quad \text{for } x \in \mathcal{D}, \quad (3)$$

with some Dirichlet/Neumann boundary conditions or initial conditions, where \mathcal{D} is the differential equation domain and $\nabla^p f(x)$ is the p -th gradient of f . The notations and the approach can easily be generalized to coupled and partial differential equations. Following classic procedures, we then discretize the domain \mathcal{D} in N points $\mathbf{x} = (x_1, x_2, \dots, x_N)$. For each point $x_i \in \mathcal{D}$, Eq. (3) must hold so that we have

$$\mathcal{F}[x_i, f(x_i), \nabla f(x_i), \dots, \nabla^p f(x_i)] = 0. \quad (4)$$

Writing the differential equation in such a way allows us to convert the problem of finding solutions to an optimization problem. Indeed, an approximated solution is a function that minimizes the square of the left-hand side of Eq. (3). In our approach, we identify the approximated solution of the differential equation (3) with a NN $\mathcal{N}(x, \boldsymbol{\theta})$ and express the loss function as follows [24]:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) = & \frac{1}{N} \sum_{i=1}^N \mathcal{F}[x_i, \mathcal{N}(x_i, \boldsymbol{\theta}), \nabla \mathcal{N}(x_i, \boldsymbol{\theta}), \dots, \nabla^p \mathcal{N}(x_i, \boldsymbol{\theta})]^2 \\ & + \sum_j [\nabla^k \mathcal{N}(x_j, \boldsymbol{\theta}) - K_j]^2, \end{aligned} \quad (5)$$

with K_j being some chosen constants. The loss function contains two terms: the first one forces the differential equation (3) to hold, and the second one encodes the boundary/initial conditions³. Note that $\mathcal{L}(\boldsymbol{\theta})$ does not depend on the exact solution⁴. The approximated solution $\mathcal{N}(x, \boldsymbol{\theta})$ is found by finding the set of weights and biases $\boldsymbol{\theta}$ such that the loss function $\mathcal{L}(\boldsymbol{\theta})$ is minimized on the training points x_i for $i = 1, 2, \dots, N$. Solving ODEs using NNs can then be expressed in terms of the following steps:

² <https://github.com/deniswerth/ODEN>

³ Usually, $k = 0$ and/or $k = 1$.

⁴ In some works, the mean squared error [32] [33] is used.

1. Write the differential equation in the form shown in Eq. (3);
2. Discretize the differential equation domain in N points $\mathbf{x} = (x_1, x_2, \dots, x_N)$;
3. Construct a neural network $\mathcal{N}(x, \theta)$ by choosing its architecture *i.e.* the number of hidden layers and the number of neurons in each layer, and the activation functions;
4. Write the loss function $\mathcal{L}(\theta)$ in the form shown in Eq. (5) including boundary/initial conditions;
5. Minimize $\mathcal{L}(\theta)$ to find the optimal NN parameters θ^* *i.e.* train the NN;
6. Once the NN has been trained on the discrete domain, $\mathcal{N}(x, \theta^*)$ is an approximation of the exact solution that can be evaluated continuously within the training domain.

For our setup, as ODE solutions are scalar functions of a scalar independent variable, the NN has one input neuron and one output neuron. The first term of $\mathcal{L}(\theta)$ in Eq. (5) requires computing the gradients of the NN output with respect to the NN input. Gradient computation, as well as the network implementation and the training procedure, is performed using the Keras [35] framework with a Tensorflow 2 [36] backend. The general method can be easily extended to coupled and partial differential equations⁵.

IV. APPLICATIONS

The method is verified against the exact solution of first and second order ODEs. We then use the method to solve an ODE that exhibits rapid oscillations. Training processes were computed using 12 CPU cores and took between a few seconds to a few minutes, depending on the number of epochs and the NN complexity.

A. First order ordinary differential equation

We first demonstrate the accuracy and efficiency of the NN solver when applied to the following simple first order differential equation with $x \in [0, 2]$, subject to the initial value $f(0) = 0$:

$$\frac{df}{dx}(x) + f(x) = e^{-x} \cos(x). \quad (6)$$

⁵ For example, coupled differential equations can be solved by corresponding the number of output neurons to the number of equations in the system one wants to solve.

The analytic solution is $f(x) = e^{-x} \sin(x)$. The equation being simple, the NN is expected to reach an acceptable accuracy for relatively few epochs. The accuracy and efficiency of the solver are illustrated in **FIG. 1** where the NN has been trained for 10^4 epochs reaching 1.6×10^{-4} as mean relative error. The bottom panel in **FIG. 1** exhibits clear wavy patterns showing that the NN solution wiggles around the exact solution. Better accuracy can be found by increasing the number of epochs.




FIG. 1: The NN solution (dots) of Eq. (6) subject to the initial value $f(0) = 0$ overlaid on the exact solution (solid line). The lower panel shows the relative error. The network was trained for 10^4 epochs using Adam optimizer with 100 uniformly spaced training points in $[0,2]$. One hidden layer of 10 neurons with a sigmoid activation function was used.

B. Stationary Schrödinger equation

The one dimensional time-independent Schrödinger equation, with potential $V(x)$, has the following differential equation⁶:

$$\frac{d^2\psi}{dx^2}(x) + 2m(E_n - V(x))\psi(x) = 0. \quad (7)$$

The energy E_n is quantized with an integer n . For a harmonic potential $V(x) = x^2$, the n^{th} level has energy $n + 1/2$ with a corresponding analytical solution to the n^{th} energy eigenfunction $\psi_n(x)$ given in terms of the Hermite polynomials. **FIG. 3** shows the NN evaluation of the energy eigenfunction $\psi(x)$ for $n = 1$. In this example, Dirichlet boundary conditions are imposed at $x = \pm 2$. Note that the relative error is maintained at $\sim 10^{-4}$ in between the boundary conditions and degrades outside.

⁶ We set $\hbar = 1$.







(a) 10^3 epochs(b) 6×10^3 epochs(c) 7×10^3 epochs(d) 8×10^3 epochs(e) 10^4 epochs(f) 1.8×10^4 epochs

FIG. 2: The NN $n = 2$ energy eigenfunction (dots) as a solution of Eq. (7) along with the exact solution (solid line), shown for comparative purposes, at different epochs during the training. The same boundary conditions, architecture, activation function, and domain sampling as FIG. 3 were used.

The NN is found to perform well on solving a boundary-value problem. As the exact solution rapidly vanishes outside the boundary conditions, the relative error dramatically increases so that it is no longer a valid measure of accuracy.




FIG. 3: NN solution (dots) of Eq. (7) for $n = 1$ overlaid on the exact solution (solid line). Dirichlet boundary conditions are imposed at $x = \pm 2$. The lower plot shows the relative error. The network was trained for 5×10^4 epochs using Adam optimizer with 100 uniformly spaced training points in $[-5, 5]$ (but displayed for $[-4, 4]$). One hidden layer of 50 neurons with a sigmoid activation function was used.

The required number of epochs during the training can only be found by iterative testing. **FIG. 2** shows the neural network prediction for the $n = 2$ energy eigenfunction for $x \in [-4, 4]$ at different epochs. One can see a general tendency in the training process that was found for other equations as well: (i) first, the NN fits the general trend of the solution by adjusting its prediction along a tangent of the curve (**FIGs. 2a** and **2b**), (ii) then, the NN fits the curvature to reveal the true shape of the solution (**FIGs. 2c-2e**), (iii) finally, it fine-tunes to adjust the prediction to the exact solution, decreasing the relative error (**FIG. 2f**). These three phases of the training can be used to guess the number of epochs required to reach the desired relative error. Generally, the more complex the solution is, the more epochs are necessary to let the NN fit the curvature. Once the right curve is found, increasing the number of epochs increases the accuracy.

C. Burst equation

To illustrate the NN's ability to fit a complex shape, we solve the following second-order differential equation

$$\frac{d^2 f}{dx^2}(x) + \frac{n^2 - 1}{(1 + x^2)^2} f(x) = 0. \quad (8)$$

The solution of Eq. (8) is characterized by a burst of

approximately $n/2$ oscillations in the region $|x| < n$ [20]. An analytical solution for the equation is

$$f(x) = \frac{\sqrt{1+x^2}}{n} \cos(n \arctan x). \quad (9)$$

The exact solution and the NN prediction is shown in **FIG. 4** for $n = 10$. With usual numerical solvers such as Runge-Kutta based approaches, the relative error grows exponentially during the step by step integration, failing to reproduce the oscillations of Eq. (9) (see **FIG. 4**) [20]. With an optimization-based approach, the trained NN is able to capture the behavior of the rapid oscillations, albeit requiring the training of a more complicated architecture over a greater number epochs when compared to the first order differential equation example in **FIG. 1** and the Schrödinger equation example in **FIG. 3**. Training on minibatches of the discretized domain was found to outperform training using the whole domain as a batch. Even for boundary conditions being imposed only at positive x values, the network is able to retain the symmetry of the solution for negative x values without losing accuracy.




FIG. 4: NN solution (dots) of Eq. (8) for $n = 10$ overlaid on the exact solution Eq. (9) (solid line). Dirichlet boundary conditions are imposed at $x = 1.5$ and $x = 3$. The lower plot shows the relative error. The network was trained for 2×10^5 epochs with minibatches of size 30 using Adamax optimizer with 300 uniformly spaced training points in $[-7, 7]$. Three hidden layers, each having 30 neurons with a tanh activation function were used.

V. PERFORMANCES

We find that even if the NNs are not susceptible to overfitting within the training domain, the main issue of the NN approach arises from the arbitrariness in the choice of NN hyperparameters and the choice of the NN

architecture. The flexibility of the method enables us to obtain a thorough understanding of how well different NN models perform on solving differential equations.

Traditional non-linear activation functions such as sigmoid and tanh are found to reproduce the ODE solution with less training than modern activation functions such as ReLU, ELU, and SELU. Rather than using a random uniform distribution or constant initialization, we find that the properly scaled normal distribution for weight and bias initialization called Xavier initialization⁷ [37] [38], makes the training process converge faster. Xavier initialization is used for all figures. Also, Adam based approaches including Adamax and Nadam, with a learning rate set to 10^{-3} , are found to perform better than stochastic gradient descent.

A. Identifying the loss function with the mean squared error

To evaluate the accuracy of the NN solution and the training performance, we suggest to use the loss function \mathcal{L} in Eq. (5). Indeed, we show that the loss function as written in Eq. (5) can be identified with the mean absolute error computed with the L^2 norm, with some scaling. The identification makes the loss function an excellent indicator of how accurate the found solution is without the need of an exact solution. Furthermore, the numerical stability and the computational cost of the training process can be assessed without computing the relative error, as it is automatically encrypted in the loss function (5).

The main idea of this identification is to write the loss function (5) as a continuous functional of the unknown solution f and Taylor expand it around an exact solution f^0 to the second order. Using the same notations as in Section III, let us take the continuous limit of the first term in Eq. (5), turning the sum into an integral, and defining the loss \mathcal{L} as a functional $\mathcal{L}[f]$

$$\mathcal{L} \rightarrow \mathcal{L}[f], \quad (10)$$

with

$$\begin{aligned} \mathcal{L}[f] = & \int_{\mathcal{D}} dx \mathcal{F}[x, f(x), \nabla f(x), \dots, \nabla^p f(x)]^2 \\ & + \sum_j [\nabla^k f(x_j) - K_j]^2. \end{aligned} \quad (11)$$

The loss functional has to be minimized on the entire differential domain \mathcal{D} rather than on a set of discrete points $x_i \in \mathcal{D}$. Using the physics language, $\mathcal{L}[f]$ is to be understood as the action with some Lagrangian

$\mathcal{F}[x, f(x), \nabla f(x), \dots, \nabla^p f(x)]^2$. For clarity, we define $\bar{\mathcal{F}} = \mathcal{F}^2$. The Cauchy-Peano theorem guarantees the existence of an exact solution f^0 to differential equations subject to boundary/initial conditions. The solution f^0 minimizes the loss function making it vanish $\mathcal{L}[f^0] = 0$. We now Taylor expand $\mathcal{L}[f]$ around f^0 to the second order

$$\begin{aligned} \mathcal{L}[f] = & \int_{\mathcal{D}} dx \bar{\mathcal{F}}[x, f^0(x), \nabla f^0(x), \dots, \nabla^p f^0(x)] \\ & + \int_{\mathcal{D}} dx \frac{\delta \bar{\mathcal{F}}}{\delta f}[x, f^0(x), \nabla f^0(x), \dots, \nabla^p f^0(x)] (f(x) - f^0(x)) \\ & + \int_{\mathcal{D}} dx \frac{\delta^2 \bar{\mathcal{F}}}{\delta f^2}[x, f^0(x), \nabla f^0(x), \dots, \nabla^p f^0(x)] (f(x) - f^0(x))^2 \\ & + \mathcal{O}([f(x) - f^0(x)]^3). \end{aligned} \quad (12)$$

The second term in Eq. (11) vanishes because f^0 satisfies the boundary/initial conditions. The zero order term in Eq. (12) vanishes because f^0 is an exact solution *i.e.* $\bar{\mathcal{F}}[x, f^0(x), \nabla f^0(x), \dots, \nabla^p f^0(x)] = 0$ for all $x \in \mathcal{D}$. The first order term in Eq. (12) vanishes because f^0 minimizes the loss function so that the functional derivative is zero⁸ *i.e.* $\frac{\delta \bar{\mathcal{F}}}{\delta f}[x, f^0(x), \nabla f^0(x), \dots, \nabla^p f^0(x)] = 0$ for all $x \in \mathcal{D}$. The remaining non-vanishing term is the second order term. As the $\frac{\delta^2 \bar{\mathcal{F}}}{\delta f^2}[x, f^0(x), \nabla f^0(x), \dots, \nabla^p f^0(x)]$ term does not depend on the NN solution and so does not depend on the parameters θ , it is a function of x only and can be seen as a scaling $\alpha(x)$

$$\mathcal{L}[f] = \int_{\mathcal{D}} dx \alpha(x) [f(x) - f^0(x)]^2 + \mathcal{O}([f(x) - f^0(x)]^3). \quad (13)$$

Written in this form, the loss function appears to be the mean absolute error using the L^2 norm integrated over the entire domain \mathcal{D} , with some scaling function $\alpha(x)$. After discretizing \mathcal{D} and identifying $f(x)$ with the neural network $\mathcal{N}(x, \theta)$, the loss function (5) takes the final form

$$\begin{aligned} \mathcal{L}(\theta) \approx & \frac{1}{N} \sum_{i=1}^N \alpha(x_i) [\mathcal{N}(x_i, \theta) - f^0(x_i)]^2 \\ & + \sum_j [\nabla^k \mathcal{N}(x_j, \theta) - K_j]^2. \end{aligned} \quad (14)$$

Note that taking the continuous limit requires enough discrete points to make it valid. Here, NNs are usually trained with $\sim \mathcal{O}(100)$ points within the considered domains. Through this identification, we better understand the three training phases shown in section IV B. Indeed, as the loss function (5) can be Taylor expanded around

⁷ GlorotNormal in Tensorflow 2.

⁸ By analogy with physics, this term would be the Euler-Lagrange equation after an integration by parts.

the exact solution, the NN first minimizes the loss leading order term making it learn the general tendency of the solution first by finding the tangents. Then, the NN minimizes higher order terms, making the NN learn specific local patterns, like curvature. Finally, the NN adjusts the solution by considering even higher orders of the loss. These universal phases enable the method to converge locally with a stable accuracy that does not decrease over the entire training domain.

B. Domain sampling

Our hypothesis is that using fine sampling for regions with more features, such as oscillations, would lead to greater accuracy *i.e.* lower loss. Moreover, the computational cost of training the NN could be reduced by avoiding dense sampling of featureless regions in the domain.

We probe different domain samplings and expose their impact on the NN performance. **FIG. 5** illustrates the effect of different domain samplings for the $n = 5$ energy eigenfunction of Eq. (7): evenly spaced points, random uniformly distributed points, and points sampled from a Gaussian distribution centered at $x = 0$ with standard deviation of 1, to accentuate the oscillatory region.

Surprisingly, a NN trained on evenly spaced points performs better than the other distributions. Every point in the region contributes equally in the loss function (5) and the NN learns the ODE solution uniformly over the entire domain. The uniform behavior can be seen in the lower panel of **FIG. 5** as the relative error remains constant over the entire domain. Moreover, increasing the number of training points is found to have no influence on the NN performances but requires more training to reach the same accuracy. Randomizing the points is found to increase the relative error by an order of magnitude. Choosing a sampling distribution that accentuates a certain region fails to reproduce the ODE solution's shape for the same training time. Further training is then necessary to reproduce the same accuracy.

C. Extrapolation performances

Once the NN is trained, it is expected to be a continuous function within the finite, bounded training domain.

Here, we illustrate the continuous nature of the NN solution. **FIG. 6** shows the extrapolation performances for points inside the training domain and outside the training domain. The $n = 5$ energy eigenfunction of Eq. (7) is solved within $[-4, 4]$ using 100 evenly spaced points. The NN prediction is then evaluated in $[0, 2]$, and in $[-6, -4] \cup [4, 6]$.

The NN is able to reproduce the ODE solution with the same accuracy within the training domain but fails to reproduce the solution outside the training domain, even with more training. Note that the mean relative error is smaller than in **FIG. 5** as more epochs were used for




FIG. 5: NN solution of Eq. (7) for $n = 5$ overlaid on the exact solution (solid line) for (i) an evenly spaced training domain (blue dots), (ii) a random uniform training domain (green dots), and (iii) a random Gaussian training domain centered at $x = 0$ with width 1 (red dots). The lower panel shows the relative error. All training domains contain 100 points in $[-5, 5]$. Dirichlet boundary conditions at $x = \pm 4$ were imposed. The network was trained for 4×10^5 epochs using Adamax optimizer and tanh activation, and two hidden layers with 20 neurons in each layer were used.




FIG. 6: NN prediction of Eq. (7) for $n = 5$ overlaid on the exact solution (solid line) for (i) the training points, 100 evenly spaced points in $[-4, 4]$ (blue dots), (ii) 40 evenly spaced points in $[0, 2]$, and (iii) 30 evenly spaced points in $[-6, -4]$ and $[4, 6]$. The lower panel shows the relative error. Dirichlet boundary conditions were imposed at $x = \pm 3$. The network was trained for 5×10^5 epochs using three hidden layers with 30 neurons in each layer. The same optimizer and activation as in **FIG. 5** were used.

the training. In classical numerical analysis, the sample spacing *i.e.* the integration step is usually small to avoid numerical instabilities [22]. Here, the optimization approach enables us to train the NN on a reduced number

of points, hence reducing the training time, and allows then to evaluate the prediction on a fine mesh.

D. Neural network architectures

Choosing the exact architecture for a NN remains an art that requires extensive numerical experimentation and intuition, and is often times problem-specific. Both the number of hidden layers and the number of neurons in each layer can affect the performance of the NN.

A detailed systematic study to quantify the effect of the NN architecture is presented in **FIG. 7**. We vary both the number of neurons in a single hidden layer and the number of hidden layers for all three equations (6), (7) and (8). An unsatisfactory value of the loss after plateauing can be viewed as a poor NN architecture performance, as further training in general does not reduce its value. Here, the loss function is used to evaluate the NN performances during the training process as it can be identified with the scaled mean squared error (see section V A).

Intuitively, one would expect an increasing NN complexity to enhance the accuracy. We found that some NN architectures, not necessarily the more complex ones, perform better than others in reaching the desired accuracy for a certain number of epochs. In general, increasing the number of hidden layers make the loss decrease more rapidly in the beginning of the training until roughly 10^4 epochs, but some simpler NN architectures prove to reach a better accuracy after more training.

A clear link between the ODE complexity and the number of NN parameters is shown. The first order ODE (6) being simple, a NN with one hidden layer containing 20 neurons is able to outperform a more complex NN with one hidden layer containing 1000 neurons (**FIG. 7a**). The difference affects the loss value by a few orders of magnitude. In this case, the number of hidden layers does not show any significant enhancement for the same amount of training. The Schrödinger equation (7) being more complex because it is a second order ODE, a two hidden layer NN performs better than a single hidden layer NN. However, the loss function for more than two hidden layers rapidly reaches a plateau and does not decrease after more training, making too complex NNs inefficient (**FIG. 7b**). Increasing the number of hidden layers for the Burst equation (8) to three is shown to both decrease the loss function faster and reach a smaller value than simple NNs (**FIG. 7c**). This tendency reveals the ODE complexity.

Indeed, for all three equations, an empiric optimal NN architecture can be found and its complexity increases with the ODE complexity. The main limitation is that we find no clear recipe of how to guess the optimal architecture, as it has to be done by testing. However, we suggest that too complex NNs (more than three hidden layers and more than 100 neurons in each layer) should not be used as the problem of solving ODEs is in general too simple.

VI. CONCLUSION

We have explored and critiqued the method of using NNs to find solutions to differential equations provided by Piscopo et al. [24]. We found that the NN was able to perform better than typical numerical solvers for some differential equations, such as for highly oscillating solutions. We must note that the proposed method should not be seen as a replacement of numerical solvers as, in most cases, such methods meet the stability and performance required in practice.


Our main message here is that the NN approach no longer appears as a black-box but a rather intuitive way of constructing accurate solutions. This approach was studied in detail. Three phases of fitting were characterized: first finding the general trend, secondly adjusting the curvature of the solution, and finally making small adjustments to improve the accuracy of the solution. Within the training domain, the NN was found to provide a continuous approximate function that matches the analytical solution to arbitrary accuracy depending on training time. However, extrapolation outside the training domain fails. We found that training in smaller mini-batches rather than the whole discretized domain used in Lagaris et al. [23] and Piscopo et al. [24] gives a greatly reduced loss.

A specifically designed loss function from the literature [24] was proved to be the appropriate metric for evaluating the solution accuracy and the NN performances without the need of the exact solution, which is usually not known. Finally, we found the limitation of the method is finding a suitable architecture. There is no trivial relationship between the NN architecture and the accuracy of the NN approximate solution for a general differential equation, though a general tendency to increase the number of NN parameters to solve more complex differential equations was highlighted.

A range of questions can be immediately explored. The performances of more sophisticated NN structures with dropouts and recurrent loops can be studied. Other sampling schemes can also be tested. Another question is whether the convergence of the NN solution to a certain accuracy can be achieved with fewer epochs. With the method we described, convergence comes locally, similar to a Taylor series. One might be able to reformulate the NN such that convergence comes globally, via a Fourier series representation, or using a different complete basis. Such reformulations might help the NN to learn general representations about the ODEs [39]. Global convergence may give better extrapolation results. Nevertheless, NNs show great potential as a support for standard numerical techniques to solve differential equations.

ACKNOWLEDGMENTS

D.W. thanks the ENS Paris-Saclay for its continuing support via the normalien civil servant grant.



(a) First order differential equation (6)



(b) Schrödinger equation (7) for $n = 2$ (c) Burst equation (8) for $n = 10$

FIG. 7: Loss function \mathcal{L} , proved to be the scaled mean squared error, during the training for different NN architectures. The first order ODE, the Schrödinger equation for $n = 2$ and the Burst equation for $n = 10$ are shown in 7a, 7b and 7c respectively. For each Figure, the upper panel is generated by varying the number of neurons in a single hidden layer, and the lower panel is generated by varying the number of hidden layers with 20 neurons in each layer. The same optimizers, activation functions and training domains as in FIG. 1, 3 and 4 for each equation were used. The learning rate was set 10^{-4} for each optimizer.

-
- [1] K. Hornik, M. Stinchcombe, and H. White, *Neural Networks* **2**, 359 (1989).
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., 2012) pp. 1097–1105.
- [3] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, (2019).
- [4] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [5] A. Tanaka and A. Tomiya, *Journal of the Physical Society of Japan* **86**, 063001 (2017), <https://doi.org/10.7566/JPSJ.86.063001>.
- [6] C.-D. Li, D.-R. Tan, and F.-J. Jiang, *Annals of Physics* **391** (2017), 10.1016/j.aop.2018.02.018.
- [7] S. J. Wetzel and M. Scherzer, *Phys. Rev. B* **96**, 184410 (2017).
- [8] Q. Wei, R. G. Melko, and J. Z. Y. Chen, *Phys. Rev. E* **95**, 032504 (2017).
- [9] H. Huang, B. Xiao, H. Xiong, Z. Wu, Y. Mu, and H. Song, *Nuclear Physics A* **982**, 927 (2019), the 27th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions: Quark Matter 2018.
- [10] J. Caldeira, W. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. Story, *Astronomy and Computing* **28**, 100307 (2019).
- [11] H. Shen, J. Liu, and L. Fu, *Phys. Rev. B* **97**, 205140 (2018).
- [12] S. Arai, M. Ohzeki, and K. Tanaka, *Journal of the Physical Society of Japan* **87**, 033001 (2018), <https://doi.org/10.7566/JPSJ.87.033001>.
- [13] N. Yoshioka, Y. Akagi, and H. Katsura, *Phys. Rev. B* **97**, 205110 (2018).
- [14] W. Wang, Y. Huang, Y. Wang, and L. Wang, in *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014* (IEEE Computer Society, 2014).
- [15] J. Herrero, A. Valencia, and J. Dopazo, *Bioinformatics* **17** (2001), 10.1093/bioinformatics/17.2.126.
- [16] E. Byvatov, U. Fechner, J. Sadowski, and G. Schneider, *Journal of Chemical Information and Computer Sciences* **43**, 1882 (2003), PMID: 14632437, <https://doi.org/10.1021/ci0341161>.
- [17] H.-P. Breuer, U. Dorner, and F. Petruccione, *Computer Physics Communications* **132**, 30 (2000).
- [18] S. Kang, B. Lee, J. Yang, S. Kim, and S. Kim, in *Frontiers in the Convergence of Bioscience and Information Technologies 2007, FBIT 2007, Jeju Island, Korea, October 11-13, 2007*, edited by D. Howard and P. Rhee (IEEE Computer Society, 2007) pp. 26 – 31.
- [19] Y. Yan, *Chin. J. Chem. Phys.* **30** (2017), 10.1063/1674-0068/30/cjcp1703025.
- [20] F. J. Agocs, W. J. Handley, A. N. Lasenby, and M. P. Hobson, *Phys. Rev. Research* **2**, 013030 (2020).
- [21] J. Wensch, M. Dne, W. Hergert, and A. Ernst, *Computer Physics Communications* **160**, 129 (2004).
- [22] J. Bamber and W. Handley, *Phys. Rev. D* **101**, 043517 (2020).
- [23] I. E. Lagaris, A. Likas, and D. I. Fotiadis, *IEEE Transactions on Neural Networks* **9**, 987 (1998).
- [24] M. L. Piscopo, M. Spannowsky, and P. Waite, *Phys. Rev. D* **100**, 016002 (2019).
- [25] A. Koryagin, R. Khudorozhkov, and S. Tsimfer, “Pydens: a python framework for solving differential equations with neural networks,” 1909.11544.
- [26] Z. Liu, Y. Yang, and Q.-D. Cai, “Solving differential equation with constrained multilayer feedforward network,” 1904.06619v1.
- [27] Y. Lu, A. Zhong, Q. Li, and B. Dong, in *Beyond Finite Layer Neural Networks: Bridging Deep Architectures and Numerical Differential Equations*, Proceedings of Machine Learning Research, Vol. 80, edited by J. G. Dy and A. Krause (PMLR, 2017).
- [28] S. Mall and S. Chakraverty, *Advances in Artificial Neural Systems* **2013**, 1 (2013).
- [29] C. Michoski, M. Milosavljevic, T. Oliver, and D. Hatch, *Neurocomputing* (2020), 10.1016/j.neucom.2020.02.015.
- [30] I. Bolodurina, D. Parfenov, and L. Zabrodina, “Algorithm for solving ordinary differential equations using neural network technologies,” (2020) pp. 79–92.
- [31] J. Sirignano and K. Spiliopoulos, *Journal of Computational Physics* **375** (2017), 10.1016/j.jcp.2018.08.029.
- [32] Y. Bekele, “Deep learning for one-dimensional consolidation,” 2004.11689.
- [33] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations,” 1711.10561.
- [34] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, 1st ed. (Springer, 2007).
- [35] F. Chollet *et al.*, “Keras,” <https://keras.io> (2015).
- [36] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, in *OSDI16: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI16 (USENIX Association, USA, 2016) p. 265283.
- [37] X. Glorot and Y. Bengio, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, Vol. 9, edited by Y. W. Teh and M. Titterton (PMLR, Chia Laguna Resort, Sardinia, Italy, 2010) pp. 249–256.
- [38] K. He, X. Zhang, S. Ren, and J. Sun, in *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV 15 (IEEE Computer Society, USA, 2015) p. 10261034.
- [39] M. Magill, F. Qureshi, and H. de Haan, in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018) pp. 4071–4081.