Collocation polynomial neural forms and domain fragmentation for solving initial value problems

Toni Schneidereit¹ · Michael Breuß¹

Received: 7 April 2021 / Accepted: 12 December 2021 / Published online: 27 December 2021
© The Author(s) 2021

Abstract

Several neural network approaches for solving differential equations employ trial solutions with a feedforward neural network. There are different means to incorporate the trial solution in the construction, for instance, one may include them directly in the cost function. Used within the corresponding neural network, the trial solutions define the so-called neural form. Such neural forms represent general, flexible tools by which one may solve various differential equations. In this article, we consider time-dependent initial value problems, which require to set up the neural form framework adequately. The neural forms presented up to now in the literature for such a setting can be considered as first-order polynomials. In this work, we propose to extend the polynomial order of the neural forms. The novel collocation-type construction includes several feedforward neural networks, one for each order. Additionally, we propose the fragmentation of the computational domain into subdomains. The neural forms are solved on each subdomain, whereas the interfacing grid points overlap in order to provide initial values over the whole fragmentation. We illustrate in experiments that the combination of collocation neural forms of higher order and the domain fragmentation allows to solve initial value problems over large domains with high accuracy and reliability.

Keywords Collocation neural forms · Polynomial neural forms · Trial solution · Initial value problems · Domain fragmentation

1 Introduction

Over the last decades, several neural network approaches for solving differential equations have been developed [1–3]. The application and extension of these approaches are a topic of recent research, including work on different network architectures like Legendre [4] and polynomial neural networks [5] as well as computational studies [6, 7].

One of the early proposed methods [8] introduced a trial solution (TS) in order to define a cost function using one feedforward neural network. The TS is supposed to satisfy given initial or boundary values by construction. It is also

referred to as neural form (NF) in this context [8, 9] which we will adopt from here on. Let us note that such NFs represent a general tool that enable to solve ordinary differential equations (ODEs), partial differential equations (PDEs) and systems of ODEs/PDEs alike. We will refer here to this approach as the trial solution method (TSM). Later, the initial method from [8] has been extended by a NF with two feedforward neural networks, which allows to deal with boundary value problems for irregular boundaries [10] and yields broader possibilities for constructing the TS [9]. In the latter context, let us also mention [11] where an algorithm is proposed in order to create a TS based on grammatical evolution. Focusing on initial value problems (IVPs), one approach employs to learn solution bundles [12], making the trained neural forms reusable for enquired initial values.

A technique related to TSM that avoids the explicit construction of trial solutions has been proposed in [13]. The given initial or boundary values from the underlying differential equation are included in the cost function as

✉ Toni Schneidereit
Toni.Schneidereit@b-tu.de

Michael Breuß
breuss@b-tu.de

¹ Applied Mathematics Group, Brandenburg University of Technology Cottbus-Senftenberg, Platz der Deutschen Einheit 1, 03046 Cottbus, Germany

additional terms, so that the NF can be set to equal the neural network output. We will refer to this approach as modified trial solution method (mTSM).

The fact that the neural network output computation resembles a linear combination of basis functions leads to a network architecture as presented in [14] (for PDEs). In that work, one hidden layer incorporates two sets of activation functions, one of which is supposed to satisfy the PDE and the second dealing with boundary conditions. The basis function coefficients are set to be the connecting weights from the hidden layer to the output neuron, and the sum over all basis functions and coefficients makes up the NF.

Motivated by the construction principle of collocation methods in numerical analysis, we propose in this paper a novel extension of the NF approach. Our neural form extension is based on the observation that the NF using one feedforward neural network as employed in [8] may be interpreted as a first-order collocation polynomial. We propose to extend the corresponding polynomial order of the neural form. The novel construction includes several feedforward neural networks, one for each order. Compared to a collocation method from standard numerics, the networks take on the role of coefficients in the collocation polynomial expansion.

Furthermore, we aim to approximate initial value problems on fairly large domains. Therefore, and based on the NF structures, we also propose a fragmentation of the computational domain into subdomains. In each subdomain, we solve the initial value problem with a collocation neural form. This is done proceeding in time from one domain fragment to the adjacent subdomain. The interfacing grid points in any subdomain provide the initial value for the next subdomain. On a first glance, one may think of similarities to domain decomposition methods for PDEs in numerical analysis, cf. [15, 16]. We also show how to combine the domain fragmentation with the newly developed collocation polynomial neural forms.

Let us note that this article is a substantial extension of our conference paper [17]. There we briefly introduced just the principle of the collocation-based polynomial NF. Here, we present a much more detailed investigation of the collocation-based NF, and as a substantial novelty the expansion and combination with our novel domain fragmentation method for solving initial value problems. That enables us to discuss both methods in a merged context, which we find highly suitable. Moreover, we updated the notation to a more suitable form.

Related work and problem statement In our previous work [18], we have shown that even simple feedforward neural networks (5 hidden layer neurons) are capable of solving a stiff initial value problem. The investigated and studied neural forms approaches are based on [8, 13].

Please find the additional related work to be addressed in the previous introduction paragraphs. The best results in the abovementioned computational study were provided by random weight initialisation, but with the drawback of a spread between the results with different random initialisation for unchanged computational parameters. There we now want to improve constant weight initialisation. The main advantage of the latter is that results are exactly reproducible with unchanged computational parameters. Based on the neural forms approach [8], we have already shown that a polynomial extension leads to a significant increase in numerical accuracy for constant weight initialisation [17]. In the present work, we now combine the polynomial neural forms and employ domain fragmentation to split the solution domain into smaller subdomains. This technique solves the neural forms on such subdomains with the initial values provided by previous subdomains. The inherent flexibility enables the approach to achieve useful results, even on fairly large domains.

2 Setting up the neural form (NF)

In this section, we first recall the TSM and its modified version mTSM, respectively, compare [8, 13]. Then, we proceed with details on the feedforward neural networks that we employ, followed by a description of the novel collocation-based neural form and the subdomain approach.

2.1 Construction of the neural form

Consider an initial value problem written in a general form as

$$G(t, u(t), \dot{u}(t)) = 0, \quad u(t_0) = u_0, \quad t \in D \subset \mathbb{R} \quad (1)$$

with given initial value $u(t_0) = u_0$. In order to connect G with a neural network, several approaches introduce a NF as a differentiable function $\tilde{u}(t, \mathbf{p})$, where the vector $\mathbf{p}$ contains the network weights. With the collocation method, we discretise the domain D by a uniform grid with $n + 1$ grid points t_i ($t_0 < t_1 < \dots < t_n$), so that the initial value problem (1) leads to the formulation

$$G(t_i, \tilde{u}(t_i, \mathbf{p}), \dot{\tilde{u}}(t_i, \mathbf{p})) = 0 \quad (2)$$

Let us note that, in a slight abuse of notation, we identify $G(t_i, \tilde{u}(t_i, \mathbf{p}), \dot{\tilde{u}}(t_i, \mathbf{p}))$ with the vector of corresponding entries (grid points), since this enables to give many formula a more elegant, compact notation.

In order to satisfy the given initial value, TSM [8] employs the NF as a sum of two terms

$$\tilde{u}(t, \mathbf{p}) = A(t) + F(t, N(t, \mathbf{p})) \quad (3)$$

where $A(t)$ is supposed to match the initial condition (with the simplest choice to be $A(t) = u(t_0)$), while $F(t, N(t, \mathbf{p}))$ is constructed to eliminate the impact of $N(t, \mathbf{p})$ at t_0 . The choice of $F(t, N(t, \mathbf{p}))$ determines the influence of $N(t, \mathbf{p})$ over the domain.

Since the NF as used in this work satisfies given initial values by construction, we define the corresponding cost function incorporating Eq. (3) as

$$E[\mathbf{p}] = \frac{1}{2} \left\| G(t_i, \tilde{u}(t_i, \mathbf{p}), \dot{\tilde{u}}(t_i, \mathbf{p})) \right\|_2^2 \quad (4)$$

which is subject to minimisation. Although Eq. (2) denotes a system of equations which may be solvable w.r.t. $\mathbf{p}$, the actual equation of interest is Eq. (4) and its optimisation will return suitable neural network weights.

Let us now turn to the mTSM approach after [13]. The mTSM approach chooses the NF to be equivalent to the neural network output directly

$$\tilde{u}(t, \mathbf{p}) = N(t, \mathbf{p}) \quad (5)$$

Since no condition is imposed by the initial value on the NF in this way, the conditions are added to the cost function when relying on Eq. (5):

$$E[\mathbf{p}] = \frac{1}{2} \left\| G(t_i, \tilde{u}(t_i, \mathbf{p}), \dot{\tilde{u}}(t_i, \mathbf{p})) \right\|_2^2 + \frac{1}{2} \left\| N(t_0, \mathbf{p}) - u(t_0) \right\|_2^2 \quad (6)$$

2.2 Neural network architecture and optimisation

In this section, we will describe how a feedforward neural network with one hidden layer operates in our setting. Specific variants will be addressed in the corresponding sections.

Fig. 1 The architecture of our incorporated feedforward neural network with one hidden layer bias

As depicted in Fig. 1, we employ one hidden layer, with H hidden layer neurons supplemented by one bias neuron. Having in addition one bias neuron in the input layer and a linear output layer neuron, the neural network output reads

$$N(t, \mathbf{p}) = \sum_{j=1}^H \rho_j \sigma_j + \gamma \quad (7)$$

Thereby, $\sigma_j = \sigma(z_j) = 1/(1 + e^{-z_j})$ represents the sigmoid activation function with the weighted sum $z_j = v_j t + \eta_j$. Here, v_j (input layer neuron), η_j (input layer bias neuron), ρ_j (hidden layer neurons) and γ (hidden layer bias neuron) denote the weights which are stored in the weight vector $\mathbf{p}$. The input layer passes the domain data t (that is in practice t_i), weighted by v_j and η_j , to the hidden layer for processing. The neural network output $N(t, \mathbf{p})$ is again a weighted sum of the values $\rho_j \sigma_j$ with γ added. With $N(t, \mathbf{p})$ given, the neural forms and cost functions in Eqs. (4),(6) are obtained.

As usual, the cost function gradient is used to update $\mathbf{p}$ in order to find a (local) minimum in the weight space. One training cycle is called an epoch and consists of a full iteration over all training points. The gradient, with respect to the network weights, determines their influence on the network output. With this information, each incorporated weight can be adjusted individually to lead the network into a local minimum during the training process. For optimising the cost function, we consider here Adam (adaptive moment estimation) [19] which is a stochastic gradient descent method, using adaptive learning for every weight.

If a weight update is performed after the gradient computation for a single grid point, we call this method single batch training (SBtraining) here. An alternative proceeding, performing the weight update after a complete iteration over all grid points, averaging the cost function gradient and training error, is denoted here as full batch training (FBtraining).

Let us comment in some detail on the relation between grid points and training points. Our setting is an unsupervised learning framework, where grid points are used for domain discretisation, and where the unknowns are values of the ODE solution at exactly these grid points. Thus, in our setting, the grid points are identical with the training points. Let us stress in this context, that the grid points by themselves stay fixed during network optimisation.

3 The novel collocation neural form (CNF)

Making Eq. (3) precise for our grid-based setting, a suitable choice for the neural form of TSM with given $u(0) = u_0$ is

$$\tilde{u}(t, \mathbf{p}) = u_0 + N(t, \mathbf{p})t \quad (8)$$

where t will be evaluated at grid points t_i . Please note, if the initial point is different from $t_0 = 0$, this results in a shift of $t \rightarrow (t - t_0)$ in Eq. (8). Compared to a first-order polynomial $q_1(t) = a_0 + a_1t$, one may find similarities in the structure. Motivated by the expansion of an m -th order collocation function polynomial [20]

$$q_m(t) = a_0 + \sum_{k=1}^m a_k t^k \quad (9)$$

we are lead to set up our collocation-based NF (CNF) approach for TSM:

$$\tilde{u}_C(t, \mathbf{P}_m) = u_0 + \sum_{k=1}^m N_k(t, \mathbf{p}_k) t^k \quad (10)$$

The weight vector is denoted by $\mathbf{p}_k$ and we define the matrix $\mathbf{P}_m$ of m weight vectors $\mathbf{P}_m = (\mathbf{p}_1, \dots, \mathbf{p}_m)$.

The use of higher-order monomial powers t^k as in Eq. (10) not only generalises previous methods, but may also enable better stability and accuracy properties, as we show in this paper. Let us also observe, that the neural networks take on the roles of coefficient functions for the values of t^k . We conjecture at this point that this construction makes sense since in this way, several possible multipliers (not only t as in Eq. (8)) are included for neural form construction. It is important to mention that the new neural form construction Eq. (10) fulfils the initial condition.

Let us stress that the proposed ansatz (Eq. (10)) includes m neural networks, where $N_k(t, \mathbf{p}_k)$ represents the k -th neural network

$$N_k(t, \mathbf{p}_k) = \sum_{j=1}^H \rho_{j,k} \sigma(v_{j,k}t + \eta_{j,k}) + \gamma_k \quad (11)$$

The corresponding cost function is then given as in Eq. (4).

We extend the mTSM method in a similar way as we obtained the TSM extension in Eq. (10):

$$\tilde{u}_C(t, \mathbf{P}_m) = N_1(t, \mathbf{p}_1) + \sum_{k=2}^m N_k(t, \mathbf{p}_k) t^{k-1} \quad (12)$$

Thereby, the first neural network $N_1(t, \mathbf{p}_1)$ is set to learn the initial condition in the same way as stated in Eq. (6).

From now on, we will refer to the number of neural networks in the neural form as the collocation neural form order, denoted by m .

4 The novel subdomain collocation neural form (SCNF)

The previous described TSM and mTSM approaches use the IVP structure together with the given initial value in order to train the neural networks on a certain domain. In a prior experimental study [7], we figured out that especially TSM tends to struggle with approximating the solution on larger domains. However, on small domains, the numerical error tends to remain small. Since the domain variable t effectively acts as a scaling of $N(t, \mathbf{p})$, we conjecture that a large domain size variation may introduce the need for a higher amount of training points or the use of a more complex neural network architecture.

These circumstances motivate us to introduce a second stage of discretising the domain. That is, we split the solution domain D in subdomains $D_l, l = 1, \dots, h$, with $n + 1$ grid points $t_{i,l}$ in each subdomain. Now, the CNF is solved separately in each subdomain. The interfacing grid points overlap, i.e. the computed value $\tilde{u}_C(t_{n,l-1}, \mathbf{P}_{m,l-1})$ at the last grid point of any subdomain D_{l-1} is set to be the new initial value $\tilde{u}_C(t_{0,l}, \mathbf{P}_{m,l})$ for the next subdomain D_l .

Since the CNF for TSM is constructed in order to satisfy the given initial values, we force the subdomain CNF (SCNF) to also hold that characteristic. Therefore, the SCNF is constructed to satisfy the new initial values in each subdomain, namely

$$\begin{aligned} \tilde{u}_C(t_{i,l}, \mathbf{P}_{m,l}) &= \tilde{u}_C(t_{0,l}, \mathbf{P}_{m,l}) \\ &+ \sum_{k=1}^m N_k(t_{i,l}, \mathbf{p}_{k,l})(t_{i,l} - t_{0,l})^k \end{aligned} \quad (13)$$

The neural networks are now scaled by $(t_{i,l} - t_{0,l})^k$, which in fact may avoid higher scaling factors, depending on the subdomain size. The arising cost function, similar to Eq. (4), is

$$E_l[\mathbf{P}_{m,l}] = \frac{1}{2} \left\| G(t_{i,l}, \tilde{u}_C(t_{i,l}, \mathbf{P}_{m,l}), \dot{\tilde{u}}_C(t_{i,l}, \mathbf{P}_{m,l})) \right\|_2^2 \quad (14)$$

Proceeding to mTSM, we also adopt the CNF approach and set the first neural network to learn the new initial values in each subdomain. That is, the SCNF reads

$$\begin{aligned} \tilde{u}_C(t_{i,l}, \mathbf{P}_{m,l}) &= N_1(t_{i,l}, \mathbf{p}_{1,l}) \\ &+ \sum_{k=2}^m N_k(t_{i,l}, \mathbf{p}_{k,l})(t_{i,l} - t_{0,l})^{k-1} \end{aligned} \quad (15)$$

and the corresponding cost function

$$E_l[\mathbf{P}_{m,l}] = \frac{1}{2} \left\| G(t_{i,l}, \tilde{u}_C(t_{i,l}, \mathbf{P}_{m,l}), \dot{\tilde{u}}_C(t_{i,l}, \mathbf{P}_{m,l})) \right\|_2^2 + \frac{1}{2} \left\| N(t_{0,l}, \mathbf{p}_{1,l}) - \tilde{u}_C(t_{0,l}, \mathbf{P}_{m,l}) \right\|_2^2 \quad (16)$$

Let us note at this point, that G in Eqs. (14) and (16) shares the same structure as the general problem in Eq. (1). However, the original solution function $u(t)$ has been replaced by the SCNF $\tilde{u}_C(t_{i,l}, \mathbf{P}_{m,l})$. Therefore, G involved in the cost function now relies on one or more neural networks, depending on the neural forms order. Once trained, each subdomain has its unique learned weight matrix $\mathbf{P}_{m,l}$ which can later be used to recreate the solution or evaluate the solution at grid points intermediate to the training points.

In order to keep the overview of all terms and indices, we sum them up again: The i -th grid point in the l -th subdomain is denoted by $t_{i,l}$, while $t_{0,l}$ is the initial point in the subdomain D_l with the initial value $\tilde{u}_C(t_{0,l}, \mathbf{P}_{m,l})$. That is, $t_{n,l-1}$ and $t_{0,l}$ are overlapping grid points. In D_1 , $\tilde{u}_C(t_{0,1}, \mathbf{P}_{m,1}) = u(t_0)$ holds. The matrix $\mathbf{P}_{m,l}$ contains the set of the m neural network weights in the corresponding subdomain l . Finally, $N_k(t_{i,l}, \mathbf{p}_{k,l})$ denotes the k -th neural network in D_l .

5 Experiments and results

This section is divided into experiments on the collocation neural form (CNF), followed by experiments on the subdomain collocation neural form (SCNF). Prior to this, we will provide detailed information about how the weight initialisation for the different neural networks is realised. The discussion of constant weight initialisation is also one of the main subjects in the experimental section. As stated before, the specific neural network configurations will be addressed in the subsequent experiments.

Weight initialisation with $\mathbf{p}_{const}^{init}$ applies to all corresponding neural networks so that they use the same initial values. Increasing the m for the initialisation with $\mathbf{p}_{rnd}^{init}$ works systematically. For $m = 1$, a set of random weights for the neural network is generated. For $m = 2$ (now with two neural networks), the first neural network is again initialised with the generated weights from $m = 1$, while for neural network number two, a new set of weights is generated. This holds for all m for higher orders, subsequently, in all experiments. To achieve comparability, the same random initialised weights are used in all experiments.

For optimisation, we use Adam, which parameters are fixed with, as employed in [19], $\alpha = 1e-3$, $\beta_1 = 9e-1$, $\beta_2 = 9.99e-1$ and $\epsilon = 1e-8$.

5.1 Experiments on the collocation neural form (CNF)

In this section, we want to test our novel CNF approach with the initial value problem

$$\dot{u}(t) + 5u(t) = 0, \quad u(0) = 1 \quad (17)$$

which has the analytical solution $u(t) = e^{-5t}$ and is solved over the entire domain $D = [0, 2]$ (without domain fragmentation). Equation (17) involves a damping mechanism, making this a simple model for stiff phenomena [21].

The numerical error Δu shown in subsequent diagrams in this section is defined as the l_1 -norm of the difference between the exact solution and the corresponding CNF

$$\Delta u = \frac{1}{n+1} \|u(t_i) - \tilde{u}_C(t_i, \mathbf{P}_m)\|_1 \quad (18)$$

If we do not say otherwise, the fixed computational parameters in the subsequent experiments are: 1 input layer bias, 1 hidden layer with 5 sigmoid neurons, 1e5 training epochs, 10 training points, $D = [0, 2]$ and the weight initialisation values which are $\mathbf{p}_{const}^{init} = -10$ and $\mathbf{p}_{rnd}^{init} \in [-10.5, -9.5]$. These values may seem arbitrarily chosen, but we found them to work well for both TSM and mTSM so that a useful comparison is possible.

Weight initialisation Let us comment in some more detail on weight initialisation. The weight initialisation plays an important role and determines the starting point for gradient descent. Poorly chosen, the optimisation method may fail to find a suitable local minimum. The initial neural network weights are commonly chosen as small random values [22]. Let us note that this is sometimes considered as a computational characteristic of the stochastic gradient descent optimisation. Another option is to choose the initialisation to be constant. This method is not commonly used for the optimisation of neural networks since random weight initialisation may lead to better results. However, constant initialisation returns reliably results of reasonable quality if the computational parameters in the network remain unchanged.

As previous experiments have documented [7, 8, 13], both TSM and mTSM are able to solve differential equations up to a certain degree of accuracy. However, an example illustrating the accuracy of five computations with random weights $\mathbf{p}_{rnd}^{init}$, respectively, constant weights $\mathbf{p}_{const}^{init}$ shows that the quality of approximations may vary considerably, see Table 1. As observed in many experiments, even a small discrepancy in the initialisation with several sets of random weights in the same range may lead to a significant difference in accuracy. On the other hand, the network initialisation with constant values very often gives reliable results by the proposed novel approach. This

Table 1 Results for five different realisations during optimisation (mTSM, $m = 2$)

No.	$\Delta u(\mathbf{p}_{rnd}^{init})$	$\Delta u(\mathbf{p}_{const}^{init})$
1	5.7148e-6	2.6653e-6
2	7.5397e-6	2.6653e-6
3	3.7249e-5	2.6653e-6
4	1.1894e-5	2.6653e-6
5	7.7956e-6	2.6653e-6

motivates us to study in detail the effects of constant network initialisation.

5.1.1 CNF experiment: number of training epochs

The first experiment shows for different m how the numerical error Δu behaves depending on the number of training epochs. The diagrams only display every hundredth data point.

In Fig. 2a with TSM and $\mathbf{p}_{const}^{init}$ results for $m = 1$ (blue) do not provide any useful approximation, independent of the batch training method selected. With a second neural network for $m = 2$ (orange) in the neural form, Δu approximately lowers by one order of magnitude so that we now obtain a solution which can be considered to rank at the lower end of reliability. However, the most interesting result in Fig. 2a is $m = 5$ (green) with the best accuracy at the end of the optimisation process but with the drawback of occurring oscillations. These may arise by the chosen optimisation method.

Table 2 shows the numeric error $\Delta u(t_i)$ at individual grid points t_i . For both CNF and Runge-Kutta 4 (RK4), the results were computed with ten grid points, resulting in the CNF approach with $\mathbf{p}_{const}^{init}$ performing better over Runge-Kutta 4. However, further refining the grid for Runge-Kutta 4 will result in significantly lower $\Delta u(t_i)$.

Table 2 Numerical error comparison at individual grid points with $m = 5$ and $\mathbf{p}_{const}^{init}$

t_i	$\Delta u(t_i)$ (TSM)	$\Delta u(t_i)$ (RK4)
0.00	0.0000e0	0.0000e0
0.22	3.8158e-5	0.1769e0
0.44	3.5101e-5	0.1478e0
0.66	1.4318e-5	9.4013e-2
0.88	1.2001e-5	5.3900e-2
1.11	4.5407e-5	2.9361e-2
1.33	5.2069e-6	1.5546e-2
1.55	6.6105e-5	8.0942e-3
1.77	1.2052e-5	4.1712e-3
2.00	7.9787e-5	2.1357e-3

For mTSM with SBtraining and $\mathbf{p}_{const}^{init}$, already $m = 1$ converges to a solution accuracy that can be considered reliable. However, we observe within Fig. 2b that only the transition from $m = 1$ (blue) to $m = 2$ (orange) affects Δu with increasing accuracy, while heavy oscillations start to occur.

In not documented results with $\mathbf{p}_{rnd}^{init}$, m has only minor influence on the accuracy. In particular, FBtraining for mTSM shows the same trend for both initialisation methods with only minor differences in the last epochs.

Let us note that the displayed results show the best approximations using constant or random initialisation. This means, we obtain the best results for TSM with FBtraining, $m = 5$ (green) and for mTSM with SBtraining, $m \geq 2$, respectively.

Concluding this experiment, we were able to get better results with $\mathbf{p}_{const}^{init}$ over $\mathbf{p}_{rnd}^{init}$. Increasing the m to at least order five seems to be a good option for TSM and

Fig. 2 Experiment in 5.1.1 Number of training epochs, (blue) $m = 1$, (orange) $m = 2$, (yellow) $m = 3$, (purple) $m = 4$, (green) $m = 5$ (color figure online)

FBtraining, whereas further m may provide even better approximations. For mTSM, we can not observe benefits for m above order 2.

Moreover, we see especially that the increase in the order of the neural form in (10) appears to have a similar impact on solution accuracy as the discretisation order in classical numerical analysis.

5.1.2 CNF experiment: domain size variation

Investigating the methods concerning different domain sizes provides information on the reliability of computations on larger domains. The domains in this experiment read as $D = [0, t_{end}]$ and we directly compare in this experiment $\mathbf{p}_{const}^{init}$ with $\mathbf{p}_{rnd}^{init}$.

In Fig. 3a, b, we observe TSM from around $t_{end} = 3.5$ to incrementally plateau to unreliable approximations. Increasing m improves Δu on small domains and shifts the observable step-like accuracy degeneration towards larger domains.

However, even with $m = 5$ (green), the results starting from domain size $t_{end} = 3.5$ towards larger sizes are unreliable. Previous to the first plateau higher m provide significant better Δu for $\mathbf{p}_{const}^{init}$, while there are only minor changes for $\mathbf{p}_{rnd}^{init}$ for the TSM method. This holds for both SBtraining and FBtraining, and one can say that in this experiment, TSM works better with $\mathbf{p}_{rnd}^{init}$, even without increasing m .

Turning to the mTSM extension, we observe in Fig. 3c with SBtraining the existence of a certain point from where different m return equal values, whereas FBtraining returns (close to) equal results for all the investigated domain sizes. However, we see some evidence for the use of $m = 2$ (orange) over $m = 1$ (blue) to show an overall good performance. A further increase of m is not necessary with this approach, confirming results from Experiment 5.1.1.

Let us also note that, with mTSM, we find that a small domain seems to favour $\mathbf{p}_{const}^{init}$ which then provides better results than $\mathbf{p}_{rnd}^{init}$.

Fig. 3 Experiment in 5.1.2 Domain size variation, (blue) $m = 1$, (orange) $m = 2$, (yellow) $m = 3$, (purple) $m = 4$, (green) $m = 5$ (color figure online)

5.1.3 CNF experiment: number of training points variation

The behaviour of numerical methods highly depends on the chosen amount of grid points, so that in this experiment, we analogously investigate the influence of the number of training points (nTP). In every computation, the domain D is discretised by equidistant grid points.

As in the previous experiments, the m shows a major influence on the results with TSM, and the best approximations are provided by $\mathbf{p}_{const}^{init}$ with $m = 5$ (green) as shown in Fig. 4a. An interesting behaviour (observed also in a different context in Fig. 2a) is the equivalence between $m = 3$ (yellow) and $m = 4$ (purple). Both converge to almost exactly the same Δu , where one may assume a saturation for the m . However, another increase in the order decreases the numerical error again by one order of accuracy.

Turning to mTSM in Fig. 4b, we again find a major increase in accuracy after a transition from $m = 1$ (blue) to $m = 2$. For nTP = 50, values for $m \geq 2$ converge to the same results as provided by TSM with $m = 5$.

Concluding this experiment, we again find evidence that increasing m in the proposed approach provides an improved accuracy for $\mathbf{p}_{const}^{init}$. However, increasing nTP seems not to improve the accuracy from a certain point on, unlike for numerical methods. But one could argue that the analogy between the number of grid points for numerical methods here is the number of epochs.

5.2 Experiments on the subdomain collocation neural form (SCNF)

In Sect. 5.1, while the test equation is stiff, its solution is at the same time very smooth and the equation is solved on a small domain. However, Fig. 3 in Experiment 5.1.2 shows

that TSM does not provide reliable solutions on larger domains. Hence, we want to show that the novel SCNF approach is able to work even on a fairly large domain with a different initial value problem. Therefore, we use the following test equation

$$\dot{u}(t) - t \sin(10t) + u(t) = 0, \quad u(0) = -1 \quad (19)$$

with the analytical solution

$$u(t) = \sin(10t) \left(\frac{99}{10201} + \frac{t}{101} \right) + \cos(10t) \left(\frac{20}{10201} - \frac{10t}{101} \right) - \frac{10221}{10201} e^{-t} \quad (20)$$

The solution is shown in Fig. 5 for $t \in [0, 15]$ and incorporates heavily oscillating and increasing characteristics, similar to instabilities. Although our approach is not limited to certain types of IVPs, we find Eq. (19) to represent possible real-world behaviour and find it suitable to serve as an example IVP.

The numerical error Δu_l is now defined as the averaged l_1 -norm of the difference between the exact solution and the corresponding SCNF in each subdomain

$$\Delta u_l = \frac{1}{n+1} \|u(t_{i,l}) - \tilde{u}_C(t_{i,l}, \mathbf{p}_{m,l})\|_1 \quad (21)$$

whereas Δu averages the numerical error of the subdomains

$$\Delta u = \frac{1}{h} \sum_{l=1}^h \Delta u_l \quad (22)$$

The weight initialisation works as employed in Sect. 5, and the values are fixed to $\mathbf{p}_{const}^{init} = 0$ and $\mathbf{p}_{md}^{init} \in [-0.5, 0.5]$. In the subsequent experiments, the solution domain is kept constant to $D = [0, 15]$, and the neural networks are training with $1e5$ epochs.

Fig. 4 Experiment in 5.1.3 Number of training points variation, (blue) $m = 1$, (orange) $m = 2$, (yellow) $m = 3$, (purple) $m = 4$, (green) $m = 5$ (color figure online)

Fig. 5 Analytical solution for initial value problem (19)

In addition, we use the method of training the neural networks incrementally which has been employed in [13]. That is, we initially train the neural networks for the first grid point, afterwards for the first two grid points. We continue the procedure up to a FBtraining of all grid points in each subdomain. The initial weight initialisation is the same in each subdomain.

Please note at this point, that we provide an explicit comparison to the Runge-Kutta 4 method only in the last experiment of this section. We find that our approach provides by construction a very high flexibility, to deal with many types of initial value problems that may require specific constructions in classic numerical analysis (e.g. by symplectic integration). However, in simple settings, we will usually find the network-based approach at the moment not competitive to numerical state-of-the-art solvers (w.r.t. computational efficiency) which have been developed and refined over decades. We think that because of the much higher flexibility of the network-based tool, this comparison would not be entirely adequate.

As an example for the flexibility, a recent neural network approach [12] makes it superficial to restart the computation of numerical solutions when considering multiple, different initial conditions. We also demonstrate the flexibility in this paper in Sect. 5.2.5 and show how invariants can be simply added to the cost function. Regarding numerical methods for handling those problems, special constructions are often needed.

In our test example in Eq. (19), we find a graphical comparison in the subsequent experiments to not provide further information since the visual differences between analytical solution and solution with Runge-Kutta-4 with adaptive time stepping are minor.

A scaling experiment The original TSM neural form (Eq. (3)) is theoretically capable of approximating every continuous function, according to the universal approximation theorem [23]. However, Table 3 shows results for a TSM neural form with a single neural network. For

Table 3 TSM neural form, $\mathbf{p}_{md}^{init}$

domain D	Δu	Neurons	nTP
[0, 1]	8.4228e-4	5	10
[0, 2]	9.2191e-4	10	20
[0, 3]	1.9448e-3	15	30
[0, 4]	1.6751e-2	20	40

different domains, we scaled the number of hidden layer neurons linearly and averaged ten computations for each domain with the same computational parameters.

The results in Table 3 provide the following message. Increasing the domain size forces the neural network to incorporate more hidden layer neurons and grid points. Indeed, to reach, e.g. (averaged) $\Delta u = 9.5873e-4$ for $D = [0, 3]$, learning the neural network required 50 hidden layer neurons and 75 grid points. In general, determining a suitable architecture in terms of the number of hidden layer neurons and training points is a challenging task.

In subsequent experiments, we find the SCNF to be able to solve the initial value problem with neural networks including a small fixed amount of hidden layer neurons and training points in each subdomain. At the same time, this allows to define various important parameters in a simple and straightforward way.

5.2.1 SCNF experiment: CNF versus SCNF

In the first experiment, we compare results of a SCNF with a CNF that is solved over the entire domain. For comparability the total number of training points is constant, namely $nTP = 1000$ for the red line and $nTP = 10$ with 100 subdomains for the black/dotted line. However, the comparison of two CNFs with the same architecture would not be meaningful because the domain size has a significant influence. Therefore, we decided to realise the CNF (red) with a neural network incorporating 1 input layer bias and 100 sigmoid neurons with 1 bias. The SCNF (black/dotted) features neural networks with 1 input layer bias and 5 sigmoid neurons with 1 bias per subdomain. Both CNF and SCNF incorporate $m = 3$. In addition, we did not increase the domain size incrementally for this experiment, to reduce the number of parameters that prevent comparability.

The CNF solution (red) shows throughout all experiments in Fig. 6 no useful approximation. In total, the number of hidden layer neurons and training points that would be needed to obtain a useful approximation seems to be much higher. Nonetheless, the SCNF approach (black/dotted) working with the same number of training

Fig. 6 Experiment in 5.2.1 CNF versus SCNF, (orange) analytical solution, (red) CNF solution, (black/dotted) SCNF solution (color figure online)

points was able to solve the initial value problem in a satisfactory way. From a qualitative perspective, both TSM and mTSM together with $\mathbf{p}_{const}^{init}$ and $\mathbf{p}_{rnd}^{init}$ provide similar results.

Concluding this experiment, we see that the SCNF method provides a useful solution to the initial value problem. In addition, the incorporated small number of hidden layer neurons enables a much more effective training of the neural networks.

5.2.2 SCNF experiment: variation of the polynomial neural forms order (m)

The ability to approximate the initial value problem with SCNF, depending on different m , is subject to this experiment. Here, the SCNFs include 1 input layer bias and 5 sigmoid neurons with 1 bias. The solution domain is split into 60 subdomains with 10 grid points in each subdomain. Here, we employ incremental learning in the subdomains.

Results for TSM with $m = 1$ in Fig. 7a and mTSM with $m = 1$ in Fig. 7b indicate that the original TSM and mTSM methods are not useful over larger domains, even when

employing domain fragmentation. However, the SCNF of first order is able to get back on the solution trend, although several subdomains do not provide correct approximations. In total, both solutions for $m = 1$ (especially mTSM) cannot be considered to be reliable.

That changes for $m = 2$, at least for TSM in Fig. 7c. Here, we find, with the exception of some local extreme points, the SCNF to be a reasonable approximation of the initial value problem. This statement, however, does not hold for mTSM. Although the general trend now is much closer to the analytical solution, there are still subdomains which do not approximate the solution well.

Results shown in Table 4 represent the l_∞ -norm of differences between analytical and computed solution, for mTSM as displayed Fig. 7, measured at the last grid points in D_{l-1} , namely $t_{n,l-1}$, and the corresponding initial points in D_l , $t_{0,l}$. We propose to consider this measure, since it indicates how well the solution can be met over the subdomains. We find that increasing m has a major influence on the accuracy.

We conjecture that learning the subdomain initial values becomes easier for mTSM, the more neural networks are

Fig. 7 Experiment 5.2.2 m variation, $\mathbf{p}_{const}^{init}$, (orange) analytical solution, (black/dotted) SCNF solution (color figure online)

Table 4 l_∞ -norm for mTSM interface grid points, $\mathbf{p}_{const}^{init}$

m	$t_{n,l-1}$	$t_{0,l}$
1	2.7834	4.8758
2	0.5763	0.7478
3	0.0846	0.0848

incorporated. That is mainly because the first neural network can so to say focus on learning the initial values, while the other networks are more engaged with the IVP

structure. We think that this conjecture can be confirmed by the decreasing discrepancy between the overlapping at the interfaces for higher orders of m .

The overall best solutions here are provided by $m = 3$ (Fig. 7e, f) for both TSM and mTSM in this experiment.

We tend to favour TSM over mTSM, since the initial value in each subdomain is satisfied by the corresponding SCNF (where the learned value at $t_{n,l-1}$ is set to be the initial value for $t_{0,l}$) and does not have to be learned again.

Fig. 8 Experiment in 5.2.3 Number of subdomain variation, (blue) $\mathbf{p}_{rnd}^{init}$, (yellow) $\mathbf{p}_{const}^{init}$ (color figure online)

5.2.3 SCNF experiment: number of subdomain variation

In this experiment, we investigate the influence of the total number of subdomains on the numeric error Δu . Figure 8 shows the behaviour for $\mathbf{p}_{rnd}^{init}$ (blue) and $\mathbf{p}_{const}^{init}$ (yellow). The SCNF incorporate $m = 3$, 1 input layer bias, 5 sigmoid neurons with 1 bias and 10 grid points in each subdomain. We again employ incremental learning in the subdomains.

Let us first comment on the SCNF for TSM in Fig. 8a. Despite minor differences between the solutions corresponding to $\mathbf{p}_{rnd}^{init}$ and $\mathbf{p}_{const}^{init}$ for smaller numbers of subdomains, both initialisation methods show a very similar trend. A saturation regime seems to appear for around 350 subdomains with $\Delta u \approx 1e-5$.

Turning to mTSM in Fig. 8b, we again observe a similar behaviour between the methods with $\mathbf{p}_{rnd}^{init}$ and $\mathbf{p}_{const}^{init}$. Although the differences almost disappear for larger numbers of subdomains, we find that even at 400 subdomains, the numerical error Δu cannot compete with TSM here.

Let us note again, that the chosen weight initialisation approach for $\mathbf{p}_{rnd}^{init}$ (see Sect. 5) means that the random weights are initialised in the same way in each subdomain. In undocumented tests, we observed that the results may show slight to significant variations, when the random weights are generated independently for each network over the subdomains. However, the results we have shown here using $\mathbf{p}_{rnd}^{init}$ represent a rather typical trend observed in the results.

In conclusion, one can obtain very good approximations with the TSM SCNF approach for both weight initialisation methods. That means, choosing $\mathbf{p}_{const}^{init}$ over $\mathbf{p}_{rnd}^{init}$ has no downsides, which leads us to again support the use of constant weight initialisation.

5.2.4 SCNF experiment: numerical error in the subdomains

The last experiment investigates the numeric error Δu_l in each subdomain D_l , depending on different m . Again, the SCNFs feature 1 input layer bias and 5 sigmoid neurons with 1 bias. The solution is computed with 100 subdomains together with 10 grid points each and incremental learning in the subdomains.

Throughout Fig. 9, $m = 1$ shows the least good results. Although, if we compare mTSM with $\mathbf{p}_{const}^{init}$ in Fig. 9b to results for 60 subdomains in Fig. 7b, increasing the domain fragments by 40 subdomains seems to prevent the solution from diverging. Random weight initialisation works better for $m = 1$, especially with TSM.

Solutions provided by $m = 3$ and $m = 5$ are much better than for $m = 1$, and increasing the order clearly tends to increase the accuracy. For TSM with both $m = 3$ and $m = 5$, as well as for mTSM with $m = 5$ from a certain subdomain on, the numerical error saturates. Let us note, that for both $\mathbf{p}_{const}^{init}$ and $\mathbf{p}_{rnd}^{init}$, the differences in the overall numerical error Δu are not significant in these cases.

In this experiment, we again tend to favour TSM with $\mathbf{p}_{const}^{init}$. Although $m = 1$ does not work well, the other shown higher orders provide good approximations with saturation regimes. The results confirm our preference of constant initialisation, because $\mathbf{p}_{const}^{init}$ does not depend on a good generation of random weights by chance.

5.2.5 SCNF experiment: system of initial value problems

In this section, we study a non-linear system of initial value problems. The example system we consider reads

$$\dot{u}(t) = a_u v(t) w(t) \quad (23)$$

$$\dot{v}(t) = a_v w(t) u(t) \quad (24)$$

Fig. 9 Experiment in 5.2.4 Numerical error in the subdomains, (blue) $m = 1$, (orange) $m = 3$, (yellow) $m = 5$

$$\dot{w}(t) = a_w u(t) v(t) \quad (25)$$

with

$$a_u = \frac{I_v - I_w}{I_v I_w}, \quad a_v = \frac{I_w - I_u}{I_w I_u}, \quad a_w = \frac{I_u - I_v}{I_u I_v} \quad (26)$$

where I_u, I_v, I_w are non-zero real numbers, and given initial values for u, v, w . The equations describe the angular momentum of a free rigid body [24, 25] with the centre of mass at the origin. These coupled initial value problems are often denoted as Euler equations and feature time invariant characteristics since the independent variable t (time) does not explicitly appear on the right-hand side. The quadratic invariant expression

$$R^2 = u^2(t) + v^2(t) + w^2(t) \quad (27)$$

conserves the so-called magnitude and describes a sphere, while another quadratic invariant

$$H = \frac{1}{2} \left(\frac{u^2(t)}{I_u} + \frac{v^2(t)}{I_v} + \frac{w^2(t)}{I_w} \right) \quad (28)$$

conserves the kinetic energy and describes an ellipsoid.

Both invariant quantities force the solution to stay on the intersection formed by the sphere and the ellipsoid. Since Eqs. (27) and (28) remain unchanged over time along the solutions, we have

$$H = \frac{1}{2} \left(\frac{u^2(0)}{I_u} + \frac{v^2(0)}{I_v} + \frac{w^2(0)}{I_w} \right) \quad (29)$$

$$R^2 = u^2(0) + v^2(0) + w^2(0) \quad (30)$$

The corresponding initial values $u(0), v(0), w(0)$ are given as in Fig. 10, and the fixed principle moments of interior (see [24], Section 14.3) have the values

$$I_u = 2, \quad I_v = 1, \quad I_w = \frac{2}{3} \quad (31)$$

assigned. In general, the neural network approach allows given invariant expressions to be directly added to the cost function, due to its flexibility. For systems of initial value problems, the cost function can be obtained by assigning each solution function its own SCNF and sum up the retrieved l_2 -norms, here together with the invariant quantities:

Fig. 10 Experiment in 5.2.5 Runge-Kutta 4 (coloured/solid) and SCNF (black/dots) solution of the free rigid body problem in Eq. (23) for different initial values, (blue) $\{u(0) = \cos(1.1); v(0) = 0.6; w(0) = \sin(1.1)\}$, (orange) $\{u(0) = \cos(1); v(0) = 0.7; w(0) = \sin(1)\}$, (yellow) $\{u(0) = \cos(1.2); v(0) = 0.5; w(0) = \sin(1.2)\}$

$$\begin{aligned}
 E[\mathbf{p}] = & \frac{1}{2} \left\| \dot{\tilde{u}}_C - a_u \tilde{v}_C \tilde{w}_C \right\|_2^2 + \frac{1}{2} \left\| \dot{\tilde{v}}_C - a_v \tilde{w}_C \tilde{u}_C \right\|_2^2 \\
 & + \frac{1}{2} \left\| \dot{\tilde{w}}_C - a_w \tilde{u}_C \tilde{v}_C \right\|_2^2 \\
 & + \frac{1}{2} \left\| \tilde{u}_C^2 + \tilde{v}_C^2 + \tilde{w}_C^2 - R^2 \right\|_2^2 \\
 & + \frac{1}{2} \left\| \frac{\tilde{u}_C^2}{2I_u} + \frac{\tilde{v}_C^2}{2I_v} + \frac{\tilde{w}_C^2}{2I_w} - H \right\|_2^2
 \end{aligned} \quad (32)$$

In order to visualise the invariant behaviour, the computed results in Fig. 10 are obtained for $t \in [0, 30]$, which allows the solution to pass its own initial points more than once. The solution domain is fragmented into 40 subdomains with ten training points in each subdomain. Due to overlapping grid points at the subdomain intersections, the total number of unique training points is $nTP = 361$, and the same amount was used to obtain computational result with Runge-Kutta 4. Each SCNF has $m = 3$, and therefore, features three neural networks involved. The latter are initialised with zeros and other training parameters and methods remain unchanged (see Sect. 5.2).

In Fig. 10, we display both the Runge-Kutta 4 (coloured/solid) and the SCNF (black/dots) solution. As mentioned above, the curves lay on intersections formed between a sphere and an ellipsoid. Changing the point of view in Fig. 10 reveals the contours of intersections between the corresponding spheres and ellipsoids for the different initial conditions (Fig. 11).

Figure 12 shows the training error arising from the cost function in Eq. (32) over the incorporated subdomains. Since the SCNF solution in each subdomain is computed independently (except for the provided initial value), the training error can differ significantly, even between two adjacent subdomains. Additionally, the flexibility of those independent computations enables this decreasing

Fig. 11 Experiment in 5.2.5 Top view of Fig. 10 to visualise the spherical and ellipsoidal intersections, colours are adopted from Fig. 10

Fig. 12 Experiment in 5.2.5 Training error over the subdomains for $\{u(0) = \cos(1.2); v(0) = 0.5; w(0) = \sin(1.2)\}$

behaviour of the training error, although the previous subdomain may have shown large errors.

In conclusion, the SCNF approach can be easily extended to handle systems of initial value problems, even with nonlinear and time invariant characteristics. A qualitative comparison between the SCNF solution and the Runge-Kutta 4 solution shows only very minor visual differences. As mentioned in Sect. 5.2, a quantitative comparison may favour Runge-Kutta 4 in terms of the numerical error in this experiment. However, the Runge-Kutta 4 method is not a symplectic integrator and may not preserve quadratic invariants over time.

6 Conclusion and future work

The proposed CNF and SCNF approaches merging collocation polynomial basis functions with neural networks and domain fragmentation show clear benefits over the previous neural form constructions. We have studied in detail the constant weight initialisation for our novel CNF approach with a basic stiff initial value problem. Depending on the batch learning methods, the collocation-based extension seems to have some benefits for both TSM and mTSM. For the TSM CNF, this effect is more significant than observed for the mTSM extension.

Focussing on mTSM and the CNF approach, using two neural networks, one for learning the initial value and one multiplied by t , seems to have some advantages over other possible mTSM settings. Considering approximation quality as most imperative, we find mTSM with $m = 2$ to provide the overall best results for the investigated initial value problem.

We find that the proposed SCNF approach combines many advantages of the new developments. Employing higher-order CNF methods, it is possible to solve initial value problems over large domains with very high accuracy, and at the same time with reasonable optimisation effort. Moreover, many computational parameters can be fixed easily for this setting, which is a significant issue with other TSM and mTSM variations.

As another important conclusion, in the experiments, we were able to show that we can favour constant weight initialisation over random weight initialisation. Nonetheless, the complexity of the IVP has an impact on the architecture and on the values of the initial constant weights. As we have pointed out in a computational study [18], DE and network-related parameters may not be independent of each other. However, the underlying relation between problem complexity and necessary neural network architecture is yet part of future work.

When focussing on constant weight initialisation, we find a further investigation on how to find suitable (constant) initial weights to be of interest. The same holds for the sensitivity of the neural network parameters.

Future research may also include work on other possible collocation functions and on combining the networks with other discretisation methods. In addition to that, let us note that we find optimal control problems to be a possible and relevant potential field of applications for our method, see, for instance, [26] for recent research in that area.

Acknowledgements This publication was funded by the Graduate Research School (GRS) of the Brandenburg University of Technology Cottbus-Senftenberg. This work is part of the Research Cluster Cognitive Dependable Cyber Physical Systems.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declaration

Conflict of interest The authors have no conflicts of interest to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Yadav N, Yadav A, Kumar M (2015) An Introduction to Neural Network Methods for Differential Equations. SpringerBriefs in Applied Sciences and Technology, Netherlands. <https://doi.org/10.1007/978-94-017-9816-7>
2. Maede AJ Jr, Fernandez AA (1994) The numerical solution of linear ordinary differential equations by feedforward neural networks. *Math Comput Model* 19(12):1–25. [https://doi.org/10.1016/0895-7177\(94\)90095-7](https://doi.org/10.1016/0895-7177(94)90095-7)
3. Dissanayake MWMG, Phan-Thien N (1994) Neural-network-based approximations for solving partial differential equations. *Commun Numer Methods Eng* 10(3):195–201. <https://doi.org/10.1002/cnm.1640100303>
4. Mall S, Chakraverty S (2016) Application of Legendre neural network for solving ordinary differential equations. *Appl Soft Comput* 43:347–356. <https://doi.org/10.1016/j.asoc.2015.10.069>
5. Zjavka L, Pedrycz W (2016) Constructing general partial differential equations using polynomial and neural networks. *Neural Netw* 73:58–69. <https://doi.org/10.1016/j.neunet.2015.10.001>
6. Famelis IT, Kaloutsas V (2020) Parameterized neural network training for the solution of a class of stiff initial value systems. *Neural Comput Appl*. <https://doi.org/10.1007/s00521-020-05201-1>
7. Schneidereit T, Breuß M (2020) Solving Ordinary Differential Equations using Artificial Neural Networks—A study on the solution variance. In: *Proceedings of the Conference Algorithm*, pp 21–30

8. Lagaris IE, Likas A, Fotiadis DI (1998) Artificial neural networks for solving ordinary and partial differential equations. *IEEE Trans Neural Netw* 9(5):987–1000. <https://doi.org/10.1109/72.712178>
9. Lagari PL, Tsoukalas LH, Safarkhani S, Lagaris IE (2020) Systematic construction of neural forms for solving partial differential equations inside rectangular domains, subject to initial, boundary and interface conditions. *Int J Artif Intell Tools* 29(5):2050009. <https://doi.org/10.1142/S0218213020500098>
10. Lagaris IE, Likas A, Papageorgiou DG (2000) Neural-network methods for boundary value problems with irregular boundaries. *IEEE Trans Neural Netw* 11(5):1041–1049. <https://doi.org/10.1109/72.870037>
11. Tsoulos IG, Gavrilis D, Glavas E (2009) Solving differential equations with constructed neural networks. *Neurocomputing* 72(10):2385–2391. <https://doi.org/10.1016/j.neucom.2008.12.004>
12. Flamant C, Protopapas P, Sondak D (2020) Solving Differential Equations Using Neural Network Solution Bundles. [arXiv:2006.14372](https://arxiv.org/abs/2006.14372)
13. Piscopo ML, Spannowsky M, Waite P (2019) Solving differential equations with neural networks: applications to the calculation of cosmological phase transitions. *Phys Rev D* 100(1):016002. <https://doi.org/10.1103/PhysRevD.100.016002>
14. Rudd K, Ferrari S (2015) A constrained integration (CINT) approach to solving partial differential equations using artificial neural networks. *Neurocomputing* 155:277–285. <https://doi.org/10.1016/j.neucom.2014.11.058>
15. Golovashkin DL, Kazanskiy NL (2009) Mesh domain decomposition in the finite-difference solution of Maxwell's equations. *Opt Memory Neural Netw* 18(3):203–211. <https://doi.org/10.3103/S1060992X09030102>
16. Jagtap AD, Kharazmi E, Karniadakis GE (2020) Conservative physics-informed neural networks on discrete domains for conservation laws: applications to forward and inverse problems. *Comput Methods Appl Mech Eng* 365:113028. <https://doi.org/10.1016/j.cma.2020.113028>
17. Schneidereit T, Breuß M (2021) Polynomial Neural Forms Using Feedforward Neural Networks for Solving Differential Equations. *Artificial Intelligence and Soft Computing, ICAISC 2021. Lecture Notes in Computer Science*, vol 12854, pp 236–245. https://doi.org/10.1007/978-3-030-87986-0_21
18. Schneidereit T, Breuß M (2020) Computational characteristics of feedforward neural networks for solving a stiff differential equation. *arXiv preprint: arXiv:2012.01867*
19. Kingma DP, Ba J (2017) Adam: A Method for Stochastic Optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
20. Antia HM (2012) *Numerical methods for scientists and engineers*, 1st edn. Hindustan Book Agency, New Delhi
21. Dahlquist GG (1978) G-stability is equivalent to A-stability. *BIT Numer Math* 18(4):384–401. <https://doi.org/10.1007/BF01932018>
22. Fernández-Redondo M, Hernández-Espinosa C (2001) Weight initialization methods for multilayer feedforward. *ESANN* 119–124
23. Cybenko G (1989) Approximation by superpositions of a sigmoidal function. *Math Control Signal Syst* 2(4):303–314. <https://doi.org/10.1007/BF02551274>
24. Griffiths DF, Higham DJ (2010) *Numerical Methods for Ordinary Differential Equations*. Springer Undergraduate Mathematics Series, London. <https://doi.org/10.1007/978-0-85729-148-6>
25. Hairer E, Lubich C, Wanner G (2006) *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer series in computational mathematics. Springer, Berlin. <https://doi.org/10.1007/3-540-30666-8>
26. Woźniak M, Połap D (2017) Hybrid neuro-heuristic methodology for simulation and control of dynamic systems over time interval. *Neural Netw* 93:45–56. <https://doi.org/10.1016/j.neunet.2017.04.013>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.