# Weierstraß-Institut für Angewandte Analysis und Stochastik Leibniz-Institut im Forschungsverbund Berlin e. V.

Preprint

ISSN 2198-5855

# Using deep neural networks for detecting spurious oscillations in discontinuous Galerkin solutions of convection-dominated convection-diffusion equations

Derk Frerichs-Mihov<sup>1</sup>, Linus Henning<sup>2</sup>, Volker John<sup>1,2</sup>

submitted: December 20, 2022

 <sup>1</sup> Weierstrass Institute Mohrenstr. 39
 10117 Berlin Germany
 E-Mail: derk.frerichs-mihov@wias-berlin.de volker.john@wias-berlin.de  <sup>2</sup> Freie Universität Berlin Department of Mathematics and Computer Science Arnimallee 6 14195 Berlin Germany E-Mail: linus.henning@fu-berlin.de

No. 2986 Berlin 2022



2020 Mathematics Subject Classification. 65N30, 68T07.

*Key words and phrases.* Convection-diffusion equations, discontinuous Galerkin methods, spurious oscillations, deep neural networks, slope limiter.

Edited by Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS) Leibniz-Institut im Forschungsverbund Berlin e. V. Mohrenstraße 39 10117 Berlin Germany

Fax:+493020372-303E-Mail:preprint@wias-berlin.deWorld Wide Web:http://www.wias-berlin.de/

# Using deep neural networks for detecting spurious oscillations in discontinuous Galerkin solutions of convection-dominated convection-diffusion equations

Derk Frerichs-Mihov, Linus Henning, Volker John

#### Abstract

Standard discontinuous Galerkin (DG) finite element solutions to convection-dominated convection-diffusion equations usually possess sharp layers but also exhibit large spurious oscillations. Slope limiters are known as a post-processing technique to reduce these unphysical values. This paper studies the application of deep neural networks for detecting mesh cells on which slope limiters should be applied. The networks are trained with data obtained from simulations of a standard benchmark problem with linear finite elements. It is investigated how they perform when applied to discrete solutions obtained with higher order finite elements and to solutions for a different benchmark problem.

# 1 Introduction

Convection-diffusion equations are a basic model to describe the distribution of a scalar quantity in fluids. Besides modeling the heat distribution in a room (energy balance), they can describe the concentration of drugs in blood and the propagation of chemical substances in water (mass balance) to name just a few. Mathematically speaking, they are given in a bounded domain  $\Omega \subset \mathbb{R}^d$ ,  $d \in \{2, 3\}$ , with polyhedral Lipschitz boundary  $\Gamma = \Gamma_D \cup \Gamma_N$  with  $\Gamma_D \cap \Gamma_N = \emptyset$ . The steady-state convection-diffusion-reaction problem with homogeneous Neumann boundary conditions on  $\Gamma_N$  then reads as follows: Find a sufficiently smooth function u such that

$$\begin{aligned} -\varepsilon \Delta u + \boldsymbol{b} \cdot \nabla u + cu &= f & \text{in } \Omega, \\ u &= g & \text{on } \Gamma_{\mathrm{D}}, \\ \varepsilon \nabla u \cdot \boldsymbol{n} &= 0 & \text{on } \Gamma_{\mathrm{N}}, \end{aligned}$$
(1)

where the coefficient  $\varepsilon \in \mathbb{R}$ ,  $\varepsilon > 0$ , is the diffusion coefficient, the convection field is denoted by  $\boldsymbol{b} \in [W^{1,\infty}(\Omega)]^d$ ,  $c \in L^{\infty}(\Omega)$  describes the reaction coefficient, and  $f \in L^2(\Omega)$  models sources. On the Dirichlet boundary  $\Gamma_D$  Dirichlet conditions g are prescribed and the outer unit normal vector on the boundary of  $\Omega$  is denoted by  $\boldsymbol{n}$ . At the inflow boundary  $\Gamma_- = \{ \boldsymbol{x} \in \Gamma : \boldsymbol{b}(\boldsymbol{x}) \cdot \boldsymbol{n}(\boldsymbol{x}) < 0 \}$ , Dirichlet boundary conditions have to be prescribed, i.e.,  $\Gamma_- \subset \Gamma_D$ .

Especially in practical applications, the so-called convection-dominated regime is of particular interest that is mathematically characterized by  $\varepsilon \ll L \| \mathbf{b} \|_{L^{\infty}(\Omega)}$ , where L is a characteristic length scale of the problem. In this regime, usually the size of occurring layers is much smaller than the mesh width on computationally feasible grids. In this sense, the convection-dominated situation is a multiscale problem, where the layers are the small scales that cannot be resolved. It is well known that in this case the discrete solution to Equation (1) obtained by classical numerical schemes, like the central finite difference method and Galerkin finite element method, exhibits huge so-called spurious oscillations,

i.e., unphysical values such as negative concentrations or an unreasonable amount of energy, e.g., see [46, 31, 32, 3].

Besides conforming finite element methods also the attention for discontinuous Galerkin (DG) methods to discretize second order elliptic equations has risen during the last decades [44, 14, 15], even though they already have been invented in 1973 [43]. They allow, for instance, easily, compared conforming finite elements, to perform hp-refinement on both simplicial and also polygonal and polyhedral meshes [18, 8]. With respect to convection-diffusion equations, DG methods with a standard upwind flux are stable discretizations in the convection-dominated regime. It was shown in [25, 4, 34, 14] that they even control the streamline derivative without needing an additional term, as it is the case for conforming finite elements. Furthermore, DG methods are known to produce very sharp layers in the convection-dominated regime. But on the other hand these methods also have the flaw of producing large overand undershoots [3, 19, 20].

A computational cheap way to significantly reduce spurious oscillations in a post-processing step are so-called slope limiters. In a first step, they identify so-called troubled-cells where over- and undershoots occur and, in a second step, replace the solution locally by a polynomial of lower degree. The solution is usually replaced by a constant approximation [17, 16] or a (at most) linear one [10, 44]. This approach is typical for appropriate numerical methods for multiscale problems in the sense that different schemes are applied for the different scales. Using low order finite elements in a vicinity of layers is fine since error bounds for higher order elements contain the norm of the solution in a higher order Sobolev space and these norms scale (locally at layers) with inverse powers of the diffusion coefficient. The power increases with the order of the Sobolev space and finally there is a huge constant in the error bound such that it cannot be expected to obtain a better accuracy in a neighborhood of layers when using higher order elements there. In [19, 20] several of these methods have been numerically investigated for convection-dominated convection-diffusion equations. These methods share the advantage that they are computationally cheap, keep the higher order approximation away from layers, and most important that most of the methods also reduce the oscillations significantly, but not completely [19, 20].

Within the last decades the interest in deep neural networks as a form of deep learning has risen sharply. Thanks to their ability to be universal function approximators [28, 11, 24, Chapter 6.4.1] and classifiers [45], they have also been used to detect troubled-cells. In 2018, Ray and Hesthaven have trained a multilayer perceptron (MLP) to identify troubled-cells for one-dimensional scalar and systems of conservation laws [41]. They have observed that their MLP detector can mimic a classical limiter but without the need of fine-tuning a parameter. Their results have been extended by the authors to twodimensional problems in [42]. Liu et al. have trained a convolutional neural network (CNN) based shock detector for Euler's Gas equations and saw that their detector was significantly faster than classical ones [37]. In 2018 and 2020, Han Veiga and Abgrall have trained a MLP detector based on data from a Runge-Kutta DG scheme and tested how well it can then be transferred to a residual distribution (RD) scheme without retraining it [47, 2]. Again, they have used scalar transport equations and Euler's Gas equations. Morgan et al. have trained and tested a MLP detector with a Lagrangian RD method to detect troubled-cells in the two-dimensional Taylor-Green vortex and Triple-point vortex [40]. Beck et al. have trained a CNN based limiter in 2020 for a DG spectral element method to approximate the solution of Euler's Gas equations [7]. Their limiter is able to both detect cells and also locate the position of the shock inside the cell. As it can be seen, all these results are for different numerical schemes for the discretization of scalar or hyperbolic systems of equations, mainly Euler's Gas equations. Neural networks have been applied also with respect to other aspects of the numerical solution of partial differential equations, e.g., see [33, 38, 48, 39, 6].

The goal of this paper is to create a MLP based limiter for convection-dominated convection-diffusion

problems and hence extending the aforementioned results. To this end, a limiter is going to be trained with data that is obtained by applying classical limiters to the lowest order discrete solution of a standard benchmark problem defined in [29]. Several architectures are tested and it will be investigated how well the resulting limiter can be applied to higher order solutions and to another benchmark problem, the so-called Hemker example from [26]. We like to emphasize that it is not the goal to speed up simulations with the MPL based limiter, since standard limiters are already very efficient. It is investigated whether or not a MLP based limiter can be constructed at all that works equally well (or even better) with respect to the reduction of spurious oscillations.

The remainder of the paper is structured as follows: In Section 2, both the standard DG method for discretizing Equation (1) and relevant classical slope limiters are introduced. Section 3 follows with explaining the multilayer perceptron model and how the data is created with which the MLP limiter is trained. Several architectures are trained in Section 4 and are tested numerically. The paper concludes with a short summary and outlook. All data and code created and used for this work can be found at www.doi.org/10.20347/40vd-f944 [21].

In what follows the usual notation is used for Lebesgue and Sobolev spaces and their respective norms. The inner product in  $L^2(\Omega)$  is denoted by  $(\cdot, \cdot)$ , a norm of a space X is denoted by  $\|\cdot\|_X$  and a seminorm by  $|\cdot|_X$ .

# 2 Discontinuous Galerkin methods and Slope Limiter for convection-diffusion equations

### 2.1 Discontinuous Galerkin methods

Equation (1) can be transformed to its weak formulation in a standard way which then reads as follows: Find  $u \in H^1(\Omega)$  such that u = g on  $\Gamma_D$  and

$$(\varepsilon \nabla u, \nabla v) + (\boldsymbol{b} \cdot \nabla u + cu, v) = (f, v) \qquad \forall v \in H^{1}_{\mathrm{D},0}(\Omega),$$
(2)

where  $H^1_{\mathrm{D},0}(\Omega) \coloneqq \{ v \in H^1(\Omega) \ : \ v = 0 \text{ on } \Gamma_{\mathrm{D}} \}.$  Under the assumptions that

$$c - \frac{1}{2} \nabla \cdot \boldsymbol{b} \ge 0, \qquad \qquad \Gamma_{\mathrm{D}} \neq \emptyset, \qquad \qquad \boldsymbol{b} \cdot \boldsymbol{n} \ge 0 \text{ on } \Gamma_{\mathrm{N}},$$

by applying the Lax–Milgram Lemma it can be proven that problem (2) possesses a unique weak solution, e.g., see [46, Section III.1.1].

To introduce the DG discretization of (2), some notation needs to be fixed. Let  $\{\mathcal{T}_h\}, 0 < h$ , be a quasi-uniform family of decompositions of  $\overline{\Omega}$  into simplicial or quadrilateral/hexahedral meshes such that for any h it holds  $\overline{\Omega} = \bigcup_{K \in \mathcal{T}_h} K$  and the cells have pairwise disjoint interiors. As usual the triangulations should be admissible, see [9, p. 38, p. 51], i.e., among other properties, each facet of a mesh cell that lies on  $\Gamma$  is either contained in  $\Gamma_D$  or  $\Gamma_N$ . The set of all facets is denoted by  $\mathcal{E}_h := \bigcup_{K \in \mathcal{T}_h} \mathcal{E}_h(K)$ , where  $\mathcal{E}_h(K)$  is the set of all facets  $E \subset \partial K$  of a cell K. Furthermore, this set can be decomposed into the set of all interior facets  $\mathcal{E}_h^I$  and boundary facets  $\partial \mathcal{E}_h := \mathcal{E}_h \cap \partial \Omega$ . The inflow boundary facets are called  $\mathcal{E}_h^- := \Gamma_- \cap \mathcal{E}_h$ , the set of Dirichlet boundary facets is denoted by  $\mathcal{E}_h^D := \partial \mathcal{E}_h \cap \Gamma_D$  and the notation  $\mathcal{E}_h^{ID} := \mathcal{E}_h^I \cup \mathcal{E}_h^D$  is set. In addition to that, |K| denotes the d-volume and  $h_K$  the diameter of a cell  $K \in \mathcal{T}_h$  and  $h := \max_{K \in \mathcal{T}_h} h_K$  is defined. Due to the regularity of the family of triangulations there exists a constant C > 0 such that for all  $\mathcal{T}_h$  and  $K \in \mathcal{T}_h$  it holds  $h_E \leq h_K \leq Ch_E$ , where  $h_E$  is the diameter of a facet  $E \in \mathcal{E}_h(K)$ .

If there exists a facet  $E \in \mathcal{E}_h(K_i) \cap \mathcal{E}_h(K_j)$  that is shared by the cells  $K_i, K_j \in \mathcal{T}_h$ , then the cells are called neighbors. Under the assumption that there is a fixed numbering of the mesh cells  $K_0, K_1, \ldots \in \mathcal{T}_h$ , the unit normal vector  $\mathbf{n}_E$  on a facet  $E \in \mathcal{E}_h$  is defined by

$$\boldsymbol{n}_E \coloneqq egin{cases} \boldsymbol{n}_K, & ext{if } E \in \partial \mathcal{E}_h \cap \mathcal{E}_h(K) ext{ for a } K \in \mathcal{T}_h, \\ \boldsymbol{n}_{K_i}, & ext{if } K_i ext{ and } K_j ext{ are neighbors along facet } E ext{ and } i < j, \end{cases}$$

where  $n_K$  denotes the outward unit normal vector of the cell  $K \in \mathcal{T}_h$ .

The below defined DG space is a subspace of the broken Sobolev space

$$H^{k}(\mathcal{T}_{h}) = \{ v \in L^{2}(\Omega) : v|_{K} \in H^{k}(K) \text{ for any } K \in \mathcal{T}_{h} \} \supset H^{k}(\Omega), \qquad k \in \mathbb{N},$$

that is equipped with its piecewise-defined norm and semi-norm

$$\|v\|_{H^{k}(\mathcal{T}_{h})}^{2} \coloneqq \sum_{K \in \mathcal{T}_{h}} \|v\|_{H^{k}(K)}^{2}, \quad |v|_{H^{k}(\mathcal{T}_{h})}^{2} \coloneqq \sum_{K \in \mathcal{T}_{h}} |v|_{H^{k}(K)}^{2}.$$

Given a fixed  $p \in \mathbb{N}$ , the finite element space is defined by

$$V_{h,p} \coloneqq \left\{ v_h \in L^2(\Omega) : v_h |_K \in \mathcal{R}_p(K) \text{ for any } K \in \mathcal{T}_h \right\} \subset H^k(\mathcal{T}_h),$$

where  $\mathcal{R}_p(K) \coloneqq P_p(K)$  is the space of polynomials of at most degree p on simplicial mesh cells and  $\mathcal{R}_p(K) \coloneqq Q_p(K)$  is the tensor product space of polynomials of at most degree p on quadrilateral/hexahedral cells.

Both  $H^k(\mathcal{T}_h)$  and  $V_{h,p}$  contain functions that are discontinuous along interior facets  $E \in \mathcal{E}_h$ . Hence, a given function  $v \in V_{h,p}$  itself is not well-defined on E but its jump  $\llbracket v \rrbracket$  and average  $\langle v \rangle$  can be defined for any  $\boldsymbol{x} \in E$  by

$$\llbracket v \rrbracket(\boldsymbol{x}) \coloneqq \begin{cases} v|_{K_i}(\boldsymbol{x}) - v|_{K_j}(\boldsymbol{x}), & \text{if } K_i \text{ and } K_j \text{ are neighbors along facet } E \text{ and} \\ & i < j, \\ v|_K(\boldsymbol{x}), & \text{if } E \in \partial \mathcal{E}_h \cap \mathcal{E}_h(K) \text{ for a } K \in \mathcal{T}_h, \end{cases}$$

and

$$\langle v \rangle \left( \boldsymbol{x} \right) \coloneqq \begin{cases} \frac{1}{2} (v|_{K_i}(\boldsymbol{x}) + v|_{K_j}(\boldsymbol{x})), & \text{if } K_i \text{ and } K_j \text{ are neighbors along facet } E \\ & \text{and } i \neq j, \\ v|_K(\boldsymbol{x}), & \text{if } E \in \partial \mathcal{E}_h \cap \mathcal{E}_h(K) \text{ for a } K \in \mathcal{T}_h. \end{cases}$$

Finally, the DG discretization of (1) reads as follows: Find  $u_h \in V_{h,p}$  such that

$$a_{\mathrm{DG}}(u_h, v_h) = f_{\mathrm{DG}}(v_h) \quad \forall v_h \in V_{h,p},$$
(3)

where the bilinear form  $a_{\mathrm{DG}} : H^1(\mathcal{T}_h) \times H^1(\mathcal{T}_h) \to \mathbb{R}$  is defined as  $a_{\mathrm{DG}}(v, w) \coloneqq a_{\varepsilon}(v, w) + a_{bc}(v, w)$ , where  $v, w \in H^1(\mathcal{T}_h)$ , with

$$a_{\varepsilon}(v,w) = \sum_{K\in\mathcal{T}_{h}} \int_{K} \varepsilon \nabla v \cdot \nabla w \, \mathrm{d}\boldsymbol{x} - \sum_{E\in\mathcal{E}_{h}^{\mathrm{ID}}} \varepsilon \int_{E} \left( \left\langle \nabla v \cdot \boldsymbol{n}_{E} \right\rangle \llbracket w \rrbracket + \kappa \left\langle \nabla w \cdot \boldsymbol{n}_{E} \right\rangle \llbracket v \rrbracket \right) \, \mathrm{d}\boldsymbol{s} + \sum_{E\in\mathcal{E}_{h}^{\mathrm{I}}} \frac{\sigma}{h_{E}} \int_{E} \llbracket v \rrbracket \llbracket w \rrbracket \, \mathrm{d}\boldsymbol{s} + \sum_{E\in\mathcal{E}_{h}^{\mathrm{D}}} \frac{2\sigma}{h_{E}} \int_{E} vw \, \mathrm{d}\boldsymbol{s}$$
(4)

Berlin 2022

and

$$a_{bc}(v,w) = \sum_{K\in\mathcal{T}_h} \int_K \left( \boldsymbol{b}\cdot\nabla vw + cvw \right) d\boldsymbol{x} - \sum_{E\in\mathcal{E}_h^{\mathrm{I}}} \int_E \boldsymbol{b}\cdot\boldsymbol{n}_E \left[\!\left[v\right]\!\right] \langle w \rangle ds + \sum_{E\in\mathcal{E}_h^{\mathrm{I}}} \int_E \frac{\eta}{2} |\boldsymbol{b}\cdot\boldsymbol{n}_E| \left[\!\left[v\right]\!\right] \left[\!\left[w\right]\!\right] ds - \sum_{E\in\mathcal{E}_h^{-}} \int_E \boldsymbol{b}\cdot\boldsymbol{n}_E vw ds.$$
(5)

The discrete right-hand side  $f_{\rm DG}$  :  $H^1(\mathcal{T}_h) \to \mathbb{R}$  of (3) is defined by

$$f_{\rm DG}(w) = \sum_{K \in \mathcal{T}_h} \int_K f w \, \mathrm{d}\boldsymbol{x} - \sum_{E \in \mathcal{E}_h^-} \int_E \boldsymbol{b} \cdot \boldsymbol{n}_E g w \, \mathrm{d}s \\ - \sum_{E \in \mathcal{E}_h^{\rm D}} \varepsilon \kappa \int_E \nabla w \cdot \boldsymbol{n}_E g \, \mathrm{d}s + \sum_{E \in \mathcal{E}_h^{\rm D}} \frac{2\sigma}{h_E} \int_E g w \, \mathrm{d}s.$$
(6)

The discrete scheme (3) contains three user-chosen parameters. The parameter  $\kappa$  controls the symmetry of (4) where  $\kappa = 1$  corresponds to the symmetric interior penalty Galerkin (SIPG),  $\kappa = 0$  to the incomplete (IIPG), and  $\kappa = -1$  to the non-symmetric (NIPG) discretization of the Laplacian. The stability parameter  $\sigma$ , also called penalty parameter, in (4) and (6) that is incorporated as in [34, Section 2.2] influences the coercivity of  $a_{\varepsilon}$ : The bilinear form for the SIPG and IIPG method is coercive if  $\sigma$  is sufficiently large, where  $\sigma$  is proportional to  $\varepsilon$ , and for the NIPG method it is coercive for any  $\sigma > 0$ , e.g., see [44, Chapter 2.7.1]. Last but not least, the stabilization parameter  $\eta \ge 0$  appearing in (5) has to be chosen by the user. The choice  $\eta = 0$  corresponds to a centered flux and  $\eta = 1$  to an upwind flux discretization across the facet E, e.g., see [14, p. 55, p. 65]. It can be proven that DG methods converge asymptotically with an optimal rate in the DG norm, with an optimal rate in the  $L^2$ -norm only for the SIPG variant and suboptimally for the IIPG and NIPG method, e.g., see [34, 44, 14, 15] and the references therein.

#### 2.2 Slope Limiters

Slope limiters are a cheap post-processing technique to reduce spurious oscillations in the discrete solution. After the solution  $u_h$  of (3) is computed, the solution is adapted as follows:

- 1 Identify and mark cells in which the function might show spurious oscillations by
  - 1.1 computing (cell wise) a set of *features* of the solution and
  - 1.2 based on these features deciding whether to mark or not to mark a cell.
- 2 Approximate the solution locally on the marked cells by a polynomial of lower degree.

These steps can be translated into mathematics by introducing some mappings. Let  $\mathcal{F}_l: V_{h,p} \times \mathcal{T}_h \to \mathbb{R}^{n_l}$  be a function that maps locally a discrete function to  $n_l \in \mathbb{N}$  features, and  $\mathcal{M}_l: \mathbb{R}^{n_l} \to \{0, 1\}$  be a decision maker function. The post-processing techniques can then be seen as mappings  $l: V_{h,p} \to V_{h,p}$  defined cell wise on a cell  $K \in \mathcal{T}_h$  for  $u_h \in V_{h,p}$  by

$$l(u_h)|_K \coloneqq \begin{cases} u_h|_K, & \text{if } \mathcal{M}_l(\mathcal{F}_l(u_h, K)) = 0, \\ \Pi_{l,K}(u_h), & \text{else}, \end{cases}$$

where  $\Pi_{l,K}: V_{h,p}|_K \to \mathcal{R}_p(K)$  reconstructs the solution in marked cells.

Different methods differ only in their respective functions  $\mathcal{F}_l$ ,  $\mathcal{M}_l$ ,  $\Pi_{l,K}$ . Several of these methods are described in detail and were tested numerically in [19, 20] and they will be briefly recalled here. Since for what comes later only  $\mathcal{M}_l$  and  $\mathcal{F}_l$  are important,  $\Pi_{l,K}$  is only mentioned in passing. For the sake of presentation, the methods are described in two dimensions on triangles, but they can be easily extended to three dimensions or to quadrilateral/hexahedral meshes.

It is worth to emphasize that for all the methods presented below, the mappings  $\mathcal{F}_l$  and  $\mathcal{M}_l$  act locally, i.e., they are cell wise defined. Especially  $\mathcal{F}_l$  can be computed using only information of the discrete solution on a cell and possibly neighbors, and globally defined quantities like a tolerance or reference values. The mapping  $\mathcal{M}_l$  then takes *cell wise features* and returns locally the decision whether to mark a cell or not.

LinTriaReco This method was proposed in [10] and described again in [44, pp. 103-104] and [19].

Let  $K \in \mathcal{T}_h$  be a simplicial cell with facets  $E_i \in \mathcal{E}_h(K)$ , i = 0, 1, 2, and neighbors  $K_i \in \mathcal{T}_h$  along these edges. Using the notation  $m_i$ , i = 0, 1, 2, for the edge mid points and  $\overline{u}_{h,K} \coloneqq \int_K u_h \, \mathrm{d}\boldsymbol{x}/|K|$ for the integral mean of  $u_h$  in K, the feature mapping of *LinTriaReco* is defined by

$$\mathcal{F}_{\text{LTR}}(u_h, K) \coloneqq \{ \overline{u}_{h, K_0}, \ u_h|_K(m_0), \ \overline{u}_{h, K_1}, \ u_h|_K(m_1), \ \overline{u}_{h, K_2}, \ u_h|_K(m_2), \ \overline{u}_{h, K}, \ tol \},$$
(7)

where  $tol \in \mathbb{R}$ ,  $tol \ll 1$  is a fixed positive tolerance. Hence,  $n_{\text{LTR}} = 8$ . Let  $\lfloor a, b \rfloor := \min\{a, b\}$  and  $\lceil a, b \rceil := \max\{a, b\}$  for  $a, b \in \mathbb{R}$ . The decision maker  $\mathcal{M}_{\text{LTR}}$  is given by

$$\mathcal{M}_{\mathrm{LTR}}(\mathcal{F}_{\mathrm{LTR}}(u_h, K)) \coloneqq \begin{cases} 1, & \text{if } \mathcal{E}_h(K) \cap \partial \mathcal{E}_h = \emptyset \land \\ & \left(u_h|_K(m_0) \notin \left[\lfloor \overline{u}_{h,K_0}, \overline{u}_{h,K} \rfloor - tol, \lceil \overline{u}_{h,K_0}, \overline{u}_{h,K} \rceil + tol \right] \lor \\ & u_h|_K(m_1) \notin \left[\lfloor \overline{u}_{h,K_1}, \overline{u}_{h,K} \rfloor - tol, \lceil \overline{u}_{h,K_1}, \overline{u}_{h,K} \rceil + tol \right] \lor \\ & u_h|_K(m_2) \notin \left[\lfloor \overline{u}_{h,K_2}, \overline{u}_{h,K} \rfloor - tol, \lceil \overline{u}_{h,K_2}, \overline{u}_{h,K} \rceil + tol \right] \biggr) \\ 0, & \text{else.} \end{cases}$$
(8)

The tolerance tol is introduced to prevent marking of cells due to numerical round-off errors. Hence roughly speaking, *LinTriaReco* marks an interior cell K if for at least one edge the value of the solution at the edge midpoint is not between the cell averages of the function in the cell and the corresponding neighbor.

For the reconstruction  $\Pi_{LTR,K}$ , three affine functions are constructed based on the cell averages of the discrete solution in the cell and its neighbors of which one is chosen, e.g., see [19, 44, p. 104].

<u>ConstTriaReco</u> This method was proposed in [19] and is a modification of *LinTriaReco*. Instead of evaluating the function at the edge midpoint the integral mean  $\overline{u}_{h,K}^E := \int_E u_h |_K ds/h_E$  is used. Furthermore, for boundary edges  $E \in \mathcal{E}_h \cap \partial \mathcal{E}_h$ , i.e., edges along which the cell has no neighbor, a virtual neighbor is constructed by mirroring the opposite vertex along the edge E. Then, on this virtual neighbor the discrete function is defined to be the continuation of  $u_h|_K$  which exists and is well-defined since  $u_h|_K$  is a polynomial of degree at most p. In this way every triangle has three neighbors and a cell average in each neighbor can be computed.

The feature mapping is then given by

$$\mathcal{F}_{\text{CTR}}(u_h, K) \coloneqq \{ \overline{u}_{h,K_0}, \ \overline{u}_{h,K}^{E_0}, \ \overline{u}_{h,K_1}, \ \overline{u}_{h,K}^{E_1}, \ \overline{u}_{h,K_2}, \ \overline{u}_{h,K}^{E_2}, \ \overline{u}_{h,K}, \ tol \},$$
(9)

and hence  $n_{\rm CTR} = 8$ .

ConstTriaReco's decision maker is then analogously defined by

$$\mathcal{M}_{\mathrm{CTR}}(\mathcal{F}_{\mathrm{CTR}}(u_h, K)) \coloneqq \begin{cases} 1, & \text{if } \overline{u}_{h,K}^{E_0} \notin [\lfloor \overline{u}_{h,K_0}, \overline{u}_{h,K} \rfloor - tol, \lceil \overline{u}_{h,K_0}, \overline{u}_{h,K} \rceil + tol] \lor \\ & \overline{u}_{h,K}^{E_1} \notin [\lfloor \overline{u}_{h,K_1}, \overline{u}_{h,K} \rfloor - tol, \lceil \overline{u}_{h,K_1}, \overline{u}_{h,K} \rceil + tol] \lor \\ & \overline{u}_{h,K}^{E_2} \notin [\lfloor \overline{u}_{h,K_2}, \overline{u}_{h,K} \rfloor - tol, \lceil \overline{u}_{h,K_2}, \overline{u}_{h,K} \rceil + tol] \end{cases}$$
(10)  
0, else.

To reconstruct the solution,  $\Pi_{\text{CTR},K}(u_h) \coloneqq \overline{u}_{h,K}$  is used, which led often to good results in the numerical studies of [19].

**ConstJumpMod** A different approach is taken by *ConstJumpMod* that was proposed in [19] and improved in [20]. Based on the marking criterion of [17, 16] *ConstJumpMod* tries to approximate the local order of convergence along each edge and marks a cell if this order is smaller than some reference value.

Let  $0 < C_0 \in \mathbb{R}$  be a positive constant, L be a characteristic length scale of the problem and  $u_0$  a characteristic scale of the solution. For each edge  $E \in \mathcal{E}_h$  the quantity

$$\alpha_E \coloneqq \begin{cases} \ln\left(\frac{1}{C_0 L u_0^2} \int_E \left[\!\left[u_h\right]\!\right]^2 \, \mathrm{d}s \right) \middle/ \ln\left(\frac{h_E}{L}\right), & \text{if } E \in \mathcal{E}_h^\mathrm{I}, \\ \alpha_{\mathrm{ref}}, & \text{else,} \end{cases}$$

can be computed, where  $\alpha_{ref} \in \mathbb{R}$  is a fixed positive reference value. These values are used to define the feature mapping that is given by

$$\mathcal{F}_{\text{CJM}}(u_h, K) \coloneqq \{ \alpha_{E_0}, \alpha_{E_1}, \alpha_{E_2}, \alpha_{\text{ref}} \}$$
(11)

and it follows that  $n_{\rm CJM} = 4$ .

To mark a cell K, the decision maker

$$\mathcal{M}_{\mathrm{CJM}}(\mathcal{F}_{\mathrm{CJM}}(u_h, K)) \coloneqq \begin{cases} 1, & \text{if } \min_{i=0,1,2} \alpha_{E_i} < \alpha_{\mathrm{ref}}, \\ 0, & \text{else,} \end{cases}$$
(12)

is used. Note, to prevent having infinite values in the feature set before computing  $\mathcal{M}_{\rm CJM}$ , these values can be replaced by  $\alpha_{\rm ref}$  without changing the result of  $\mathcal{M}_{\rm CJM}$ , which might be beneficial for the implementation.

The solution in the marked cells is again replaced by the cell integral mean, i.e.,  $\Pi_{\text{CJM},K}(u_h) \coloneqq \overline{u}_{h,K}$ .

**ConstJumpNorm** Based on the previous approach, the method *ConstJumpNorm* was introduced in [20] that depends on the mean  $L^{\infty}(E)$ -norm of the jump of the function  $u_h$ . The  $L^1$ - and  $L^2$ -norms have been investigated as well but significant differences could not be observed [20]. If this jump is larger than a fixed positive reference value  $0 < \beta_{\text{ref}} \in \mathbb{R}$  the cell is marked.

To be precise, for each edge  $E \in \mathcal{E}_h$  the quantity

$$\beta_E \coloneqq \begin{cases} \|\llbracket u_h \rrbracket\|_{L^{\infty}(E)} , & \text{if } E \in \mathcal{E}_h^{\mathrm{I}}, \\ 0, & \text{else,} \end{cases}$$

can be defined. Based on this quantity, the feature mapping

$$\mathcal{F}_{\text{CJN}}(u_h, K) \coloneqq \{\beta_{E_0}, \beta_{E_1}, \beta_{E_2}, \beta_{\text{ref}}\}$$
(13)

can be computed, so that  $n_{\rm CJN} = 4$ .

The decision maker function of *ConstJumpNorm* is given by

$$\mathcal{M}_{\text{CJN}}(\mathcal{F}_{\text{CJN}}(u_h, K)) \coloneqq \begin{cases} 1, & \text{if } \max_{i=0,1,2} \beta_{E_i} \ge \beta_{\text{ref}}, \\ 0, & \text{else,} \end{cases}$$
(14)

and the solution is approximated by  $\Pi_{\text{CJN},K}(u_h) \coloneqq \overline{u}_{h,K}$ .

### 3 Deep neural networks as spurious oscillations detector

Deep learning techniques such as deep (neural) networks are a subpart of machine learning which try to approximate a possibly unknown function by learning it from data [24, p. 1-8]. In the following, multilayer perceptrons (MLPs) also known as feed forward neural networks are briefly introduced; see also [27, 24, Chapter 6] for detailed information.

Mathematically speaking, MLPs can be seen as functions that map an input domain  $\mathcal{X}$  to some output domain  $\mathcal{Y}$  by composing a sequence of functions  $g_1, g_2, \ldots, g_\ell$ , i.e.,

$$x \mapsto g_{\ell}(g_{\ell-1}(\dots g_1(x))\dots) \in \mathcal{Y} \quad (x \in \mathcal{X}).$$

Here each  $g_i$ ,  $i = 1, 2, ..., \ell$ , also called *i*-th layer has the form  $g_i(\bullet) = \sigma_i(W_i \bullet +b_i)$ , where  $W_i$  is a rectangular matrix called weight matrix,  $b_i$  is a vector called bias vector,  $\sigma_i$  is a component wise defined nonlinear function called activation function. The first layer is called input layer, the last layer is called output layer and the layers in between hidden layers. In other words, starting with x as value(s) of the input layer each following layer takes as input all the output of the previous layer, also called nodes, performs an affine transformation and applies component wise an activation function. MLPs can be therefore characterized or rather parametrized by their corresponding weights, biases and activation functions, which is why they are often referred to as *parameters*. Different activation functions can be used, but what they all have in common is that they are nonlinear, which is crucial to approximate nonlinear functions [24, p. 168]. Possible choices are the sigmoid function  $\sigma(x) = 1/(1 + e^{-x})$ , the rectified linear unit (ReLU) function  $\sigma(x) = \max\{0, x\}$  or the hyperbolic tangent  $\sigma(x) = \tanh(x) = (e^x - e^{-x})/(e^x + e^{-x})$ . To reach the goal that a MLP approximates a given function, the parameters need to be chosen accordingly. They are chosen in an optimization process that is called *training*.

Let  $F : \mathcal{X} \to \mathcal{Y}$  be the function that will be approximated by a MLP. During the training the parameters are optimized to minimize a given loss function  $\mathcal{L}$  over a given finite data set  $\mathcal{D} \subset \mathcal{X} \times \mathcal{Y}$  which consists of pairs  $(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}$  of features  $x_i$  and labels  $y_i = F(x_i)^1$ .

<sup>&</sup>lt;sup>1</sup>This is so-called supervised learning. See [24, p. 103-105] for unsupervised and reinforcement learning.

In this work, different approaches are investigated to approximate (combinations of) the decision maker functions  $\mathcal{M}_{LTR}$ ,  $\mathcal{M}_{CTR}$ ,  $\mathcal{M}_{CJM}$ , and  $\mathcal{M}_{CJN}$  by MLPs, see below. The concrete choice  $F = \mathcal{M}_{LTR}$ ,  $\mathcal{X} = \mathcal{F}_{LTR}$ ,  $\mathcal{Y} = \{0, 1\}$ ,

$$\mathcal{L}(D) \coloneqq -\frac{1}{N} \sum_{i=1}^{N} y_i \ln(\hat{y}_i) + (1 - y_i) \ln(1 - \hat{y}_i), \tag{15}$$

where N is the number of training data in  $\mathcal{D}$  and  $\hat{y}_i$  is the prediction of the MLP, may serve as a simple example and is used in Section 4.2.

During the training the parameters p are updated by

$$p \to p - \eta \nabla_p \mathcal{L}(p),$$

where  $0 < \eta \in \mathbb{R}$  is a positive step width, also called learning rate, and  $\nabla_p$  denotes the partial derivatives with respect to the parameters. In this work the minibatch stochastic gradient descent [27, 24, Chapter 8.1.3] is used together with the Adam algorithm [35] to adapt the step width automatically.

### 3.1 Generating the data set

To enable the MLP to approximate a given function training data is needed. As stated above, decision maker functions are approximated that take as input a feature vector of the solution on a single cell and return either 1 (mark the cell) or 0 (do not mark a cell). To generate training data the following problem is fixed.

**Example 1** (HMM example). Let  $\Omega = (0, 1)^2$  be the unit square and  $\boldsymbol{b} = (\cos(-\pi/3), \sin(-\pi/3))^T$ , c = f = 0. On the whole boundary Dirichlet boundary conditions are prescribed, i.e.,  $\Gamma_D = \partial \Omega$ , by choosing

$$g = \begin{cases} 1 & (y = 1 \land x > 0) \text{ or } (x = 0 \land y > 0.75), \\ 0 & \text{else.} \end{cases}$$

This example is a modification of a classical benchmark problem stated in [29] in which the discontinuity point of the Dirichlet boundary conditions is chosen slightly different to match the requirements of applying a DG method on a uniform grid.

For small diffusion coefficients  $\varepsilon$ , the solution possesses two boundary layers at the outflow boundary and an interior layer in the direction of the convection, see Figure 1 for a sketch of the solution.

To generate training data, the discrete problem (3) can be solved on a series of uniformly refined meshes starting with the initial meshes depicted in Figure 2. On each level, the discrete solution is calculated and afterwards on each cell the features of *LinTriaReco*, *ConstTriaReco*, *ConstJumpMod*, *ConstJumpNorm* and the corresponding labels are computed and stored, see Equations (7)–(13). Since a data point for each *cell* is created, a lot of data can be generated easily since the number of cells scales quadratically when the grid is refined. Here it comes in handy that the decision maker functions act locally.

The data is generated using discontinuous piecewise linear finite elements  $P_1$  for the above mentioned problem with a diffusion coefficient  $\varepsilon = 10^{-8}$ . The SIPG discretization is chosen with upwind stabilization, i.e.,  $\kappa = 1$  and  $\eta = 1$  in Equations (4) to (6). Let  $n_0$  be the number of vertices in each cell,



Figure 1: Sketch of the solution to Example 1 for  $\varepsilon = 10^{-8}$  obtained with a nonlinear algebraic flux-corrected (AFC) finite element method with Kuzmin limiter, see [5].



Figure 2: Initial meshes for Example 1. The grid on the left-hand side is referred to as regular and the grid on the right-hand side as irregular grid.

i.e.,  $n_0 = 3$  on the triangular grids depicted in Figure 2. Guided by [44, Chapter 2.7.1], the penalty parameter  $\sigma = 2\varepsilon n_0(p+1)(p+2)/2 = 18\varepsilon$  is used. All simulations are performed with ParMooN [22, 49] and the direct solver UMFPACK [12] is used to solve the linear systems of equations. In *LinTriaReco* and *ConstTriaReco* the tolerance *tol* is set to  $10^{-11}$ . For *ConstJumpMod* the parameters  $C_0 = 1$ , L = 1.5 and  $u_0 = 1$ , and  $\alpha_{ref} = 4$  are chosen. Lastly, the reference value  $\beta_{ref}$  is set to be the arithmetic mean of all  $\beta_E$  in *ConstJumpNorm*.

#### 3.1.1 Rotation invariance of the data

Unfortunately, the features defined in Equations (7), (9), (11) and (13) depend on the numbering of the edges and hence, so does the data. To overcome this problem, or in other words to introduce some sort of rotation invariance, each data point point in the data set is stored three times, one for each particular counter clockwise numbering of the edges.

### 3.1.2 Magnitude invariance

In [2] the authors have decided to scale the features to introduce some form of magnitude invariance. In contrast to this, here the features are *not* scaled. The reason for this is that all decision maker functions essentially compare the magnitude of a feature with other features. If features are scaled feature-wise as in [2], the ratio of the magnitude of features can be changed. In this way inconsistent data can be created, i.e., the label does not fit to the data anymore. To prevent this situation, a scaling of the features is therefore not applied.

### 3.2 Restricting the data set

Following the above describe procedure a lot of data can be generated, e.g., refining the regular grid nine times and the irregular grid eight times leads to 4.456.437 data points. Unfortunately, a lot of duplicates exist in the data, e.g., due to the fact that the solution of the problem is piecewise constant in huge parts of the domain and hence, the features can be equal. This can be the case for an individual limiter but also for any combination of limiters, e.g., also for all limiters at the same time. Our approach consists in removing the duplicates to prevent the MLP from learning a pattern specific to the duplicates and to prevent overfitting to the duplicates and hence, ending up in a MLP that does not generalize well to unseen data. That is, either the duplicates in the data of a single limiter are removed if a single limiter is learned, or duplicates of the data of all limiters if all limiters are learned at the same time, see also the numerical examples below.

After having removed the duplicates it can further be noted that there are for each limiter individually significantly less marked cells than not-marked cells, e.g., for *LinTriaReco* after removing the duplicates there are around 77% cells that are not marked and 23% marked cells. When inspecting the whole data set it can be seen that cells that are not marked by any limiter are more common (ca. 93.6%) than cells that are marked at least by one limiter (ca. 6.4%). It is well known that care has to be taken when it comes to such unbalanced data sets [36, Chapter 11.2]. To have a better balance between marked cells and not-marked cells, resp., in the distribution of the label combinations, the data set is further restricted to have either as many marked cells as cells that are not marked in the case that a single limiter is approximated or the amount of the combination where no limiter marks a cell is reduced to equal the amount of the second most occurring label combination in the case all limiters are

approximated at a single time. The rows that are removed are chosen randomly using a fixed random seed to guarantee reproducibility.

### 3.3 Splitting the data set

Even after deleting duplicates and decreasing the amount of cells that are not marked, resp., the amount of the most occurring label combination, a lot of training data remains, e.g., 379.539 when all limiters are learned at the same time, and 260.436 if only *ConstJumpNorm* is considered. On the one hand, the more data exists the more likely it is that the network approximates the function that lies behind the data, but on the other hand, more training data increases the optimization time when the network is trained. Hence, it is recommended to have less data of higher quality, i.e., showing relevant features of the function, than more data with lower quality. In this data set it might be difficult to choose "good" data points a priori but it might be still useful to choose only a subset of the data points for performance reasons. Therefore, a subset of only 7500 data points is randomly chosen to be the training data for the networks.

Furthermore, to prevent overfitting of the data, another 1875 are chosen to be the validation data set, see also [24, Chapter 5.3] for an introduction to overfitting and validation sets. During the training the validation set is evaluated as well to see how well the network generalizes to unseen data. At some point the networks might only fit better the training data but they become worse on the validation data set, which is why the optimization can be stopped at this point to prevent overfitting.

Last but not least, a third set is introduced with which the trained networks are rated how well they work. The so-called test set consists of the validation set and all remaining data. After the training has finished the networks are applied on the test set to measure the overall performance of the networks.

### 3.4 Measuring the performance of the networks

To measure the performance of the trained networks the measures

accuracy 
$$\operatorname{acc} := \frac{t_p + t_n}{t_p + f_p + t_n + f_n},$$
  
precision  $\operatorname{prc} := \frac{t_p}{t_p + f_p},$   
recall  $\operatorname{rec} := \frac{t_p}{t_p + f_n},$ 

are used, where  $t_p$  denotes the true positive,  $t_n$  the true negative,  $f_p$  the false positive and  $f_n$  the false negative classifications. The measure accuracy is the ratio of correct classified data to all data, i.e., it measures how good the networks performs overall. While the second measure gives information about the proportion of positive classifications that was actually correct, recall answers what proportion of actual positives was identified as such. Since for reducing spurious oscillations it is worse to not detect a cell that should be marked than to mark a cell that should not, recall might be considered more important than precision. Therefore, the total rating  $r_{\rm tot}$  of the limiters is set to be a weighted combination of the measures, namely

$$r_{\text{tot}} \coloneqq \frac{2}{5} \operatorname{acc} + \frac{1}{5} \operatorname{prc} + \frac{2}{5} \operatorname{rec},$$

where  ${\rm acc},\,{\rm prc}$  and  ${\rm rec}$  are computed based on the test set.

/I** I**	<u> </u>
hidden layers	[256, 128, 64], [128, 128, 128], [256, 128, 64, 32],
	[100, 100, 100, 100], [256, 128, 64, 32, 16], [90, 90, 90, 90, 90]
learning rate	$0.005, 0.001, 0.0005, 0.0001, 5 \cdot 10^{-5}, 1 \cdot 10^{-5}$
batch size	32, 64, 128
activation	ReLU, tanh
initialization seed	40, 41, 42

Table 1: Hyperparameters that are tested resulting in 648 different combinations.

# 4 Numerical studies

For the implementation of the MLP networks the open source library TensorFlow is used [1, 13]. As stated above, the finite element computations are performed with ParMooN [22, 49] and CppFlow [30] is used to open and deploy stored TensorFlow models in ParMooN. Note that the data and most parts of the code that are used in this section are publicly available at www.doi.org/10.20347/40vd-f944 [21].

### 4.1 Architecture of the MLPs

Given a specific mapping that should be approximated by a MLP, it is in general almost impossible to come up a priori with the optimal architecture of the MLP, i.e., the number of layers, activation functions, number of nodes per layer. To find a suitable architecture, in this work, different architectures are tested. Six different combinations of number of hidden layers and number of nodes which corresponds to the number of columns in the weight matrices  $W_i$ , two different activation functions, three different batch sizes, six different learning rates and three different initializations of the parameters are used which are coded by different seeds, resulting in 648 different architectures that are investigated, see Table 1. Each combination therefore can also be identified by a number between one and 648 which is done below. The size of the input and the output layer are determined by the task to solve. While for all hidden layers the same activation function is used, i.e., one of the functions given in Table 1, the activation function for the output layer depends on the experiment and is therefore stated in the experiments below. The parameters of the layers are initialized using the Glorot normalized initialization [23] with different seeds for each layer based on the seeds given in Table 1. Also the loss function  $\mathcal{L}$  depends on the experiment and hence is given below.

The networks are trained for at most 10000 epochs and the training is stopped earlier, if the loss of the validation set does not decrease for 100 epochs. The model with the best accuracy is then saved as trained model.

### 4.2 Learning single limiters

The first experiment figures out whether the individual feature mappings (8), (10), (12), and (14) can by approximated by a MLP. Since for all functions  $\mathcal{Y} = \{0, 1\}$ , the output layer consists of a single node and uses the sigmoid activation function. The size of the input layer is defined by the input of the decision maker functions, i.e., eight for *LinTriaReco* and *ConstTriaReco*, and four for *ConstJumpMod* and *ConstJumpNorm*. As loss the binary cross entropy loss  $\mathcal{L}(D)$  given in (15) is applied. The data is prepared as described in Sections 3.1 to 3.3 and the measures defined in Section 3.4 are used to measure the performance of the networks.

The total rating  $r_{tot}$  for the networks for each limiter is plotted in Figure 3, where *MLP* (*lim*) denotes the *MLP* networks that approximate the decision maker function of lim. In general it can be seen that the results of *MLP* (*LinTriaReco*) and *MLP* (*ConstTriaReco*) look similar as well as the results of *MLP* (*ConstJumpMod*) and *MLP* (*ConstJumpNorm*). On the one hand, all architectures are able to approximate *ConstJumpMod* very well and *ConstJumpNorm* can be approximated by almost all architectures. On the other hand, *LinTriaReco* and *ConstTriaReco* cannot be approximated that well with the investigated architectures. The best total rating for *MLP* (*LinTriaReco*) and *MLP* (*ConstJumpNorm*), see also Table 2. As indicated by the standard deviation, the quality of the approximation of *MLP* (*LinTriaReco*) and *MLP* (*ConstJumpNorm*), which in turn is more dependent than the approximation of *MLP* (*ConstJumpMod*). It can further be noted that in Figure 3 there is a pattern indicating which architectures work worse for *MLP* (*LinTriaReco*), *MLP* (*ConstTriaReco*) and *MLP* (*ConstJumpMod*). But since this experiment is intended to check if the decision maker functions can be approximated at all, it is not further investigated which parameters are responsible for the loss in quality.

Table 2: Statistics about the total rating of the trained networks of Section 4.2. Standard deviation is abbreviated by std.

	$r_{ m tot}$ LinTriaReco	$r_{ m tot}$ ConstTriaReco	$r_{ m tot}$ ConstJumpMod	$r_{\rm tot}$ ConstJumpNorm
max	0.744	0.737	0.999	0.997
mean	0.669	0.670	0.997	0.994
std	0.034	0.035	0.001	0.019



Figure 3: Rating for all architectures for all limiters for Section 4.2.

### 4.3 Overcoming the difficulties when learning LinTriaReco and ConstTriaReco

As seen in the previous experiment, the decision maker functions of *LinTriaReco* and *ConstTriaReco* could not be approximated well and *ConstJumpMod* and *ConstJumpNorm* could be approximated by the chosen architectures if only the features of the respective limiter are used. This experiment investigates if enriching the feature set enables the architectures to fit the decision maker function of *LinTriaReco*. To this end, the output layer consists again of a single node and the sigmoid activation function is used. The idea in this experiment is to use all features of *LinTriaReco*, *ConstTriaReco*,

*ConstJumpMod*, and *ConstJumpNorm*. Hence, the input layer, in contrast to the previous experiment, is now larger and consists of  $n_{\rm LTR} + n_{\rm CTR} + n_{\rm CJM} + n_{\rm CJN} = 24$  nodes. Again the binary cross entropy loss (15) is applied to train the networks. The data is loaded, restricted and split as before and the measure  $r_{\rm tot}$  is used to rate the trained MLPs.

Figure 4 shows the result for the tested architectures. All tested architectures have a similar good rating, i.e., it seems that there is a mapping from all features to the label of *LinTriaReco* that can be approximated with the used architectures. The best rating (0.979) is slightly worse than the best results for *ConstJumpMod* (0.999) and *ConstJumpNorm* (0.997) of the previous experiment but could increase the rating of *LinTriaReco* by around 0.235. Also all architectures are stable in the sense of producing similar good results as shown by the mean that is close to the best rating and the small standard deviation, see Table 3.

Table 3: Statistics about the total rating of the trained networks of Section 4.3. Standard deviation is abbreviated by std.

_	$r_{ m tot}$ LinTriaReco	
max	0.979	
mean	0.977	
std	0.002	



Figure 4: Results for all architectures from Section 4.3.

### 4.4 Learning all limiters simultaneously based on vectors

The previous experiment raises the questions whether it is possible to learn all decision maker functions at once. The idea is to train a network that approximates the map from all features to all labels simultaneously, i.e., the size of the input layer is  $n_{\rm LTR} + n_{\rm CTR} + n_{\rm CJM} + n_{\rm CJN} = 24$  and the output size is four, since we want to approximate four decision maker functions at once. In this sense the problem is a multi-label classification task. As in Section 4.2, the activation function of the output layer is set to be the sigmoid function such that the output of the network is in  $[0, 1]^4$ . This means that by construction the networks return a vector of four predicted labels at once. Furthermore, the loss

$$\mathcal{L}(D) \coloneqq \frac{1}{4} \sum_{j=1}^{4} \left( -\frac{1}{N} \sum_{i=1}^{N} y_{i,j} \ln(\hat{y}_{i,j}) + (1 - y_{i,j}) \ln(1 - \hat{y}_{i,j}) \right)$$

is used, where N is again the number of training data in  $\mathcal{D}$ ,  $y_{i,j}$  is the *j*-th component of the *i*-th training data point and  $\hat{y}_{i,j}$  is the *j*-th component of the prediction  $\hat{y}_i$  of the MLP. Comparing with Equation (15), this loss is the average of the binary cross entropy loss of the four decision maker functions that are learned at once. The data is prepared following the procedure described in Section 3.1 to Section 3.3 and the total measure  $r_{\text{tot}}$  from Section 3.4 is used to rate the trained MLPs. The measures are evaluated element-wise and not vector-wise for the output of the MLPs.

Table 4: Statistics about the total rating of the trained networks of Section 4.4. Standard deviation is abbreviated by std.

	$r_{\rm tot}$
max	0.950
mean	0.938
std	0.007

The results for all configurations are depicted in Figure 5. It can be seen that almost all configurations are able to approximate all decision maker functions at once. As indicated also by Table 4 the best network achieves a total rating of around 0.95 which is slightly lower than the best rating in Section 4.2 but can still considered to be a good result. Both the mean of 0.938 that is close to the maximal total rating and the small standard deviation (0.007) indicate that there are only few configurations that work worse than the best one. There seems to be again a pattern of those configurations but it is not further investigated since only the best configuration is of interest in this study.



Figure 5: Results for all architectures from Section 4.4.

### 4.5 Learning all limiters simultaneously based on classes

The multi-label problem of the previous section can be transformed to a multi-class problem. After preparing the data as before, every unique label combination is assigned to a unique number, e.g.,

$$[0, 0, 0, 0] \mapsto 0, \qquad [0, 0, 0, 1] \mapsto 1, \qquad [0, 0, 1, 0] \mapsto 2$$

and so on. This number j is then assigned to a probability vector, i.e., the components are non-negative and sum to 1, by mapping j to the vector  $[0, \ldots, 0, 1, 0, \ldots, 0]$  that has the 1 at the j-th component. Every entry in this vector gives the probability that the input belongs to the respective class. The input layer size is again 24 and the output layer size is the number of classes which are in this experiment 12 since not all possible label combinations occur in the dataset. After splitting the data into training, validation, and test set it can be observed that not all label combinations occur in the training set since some combinations are very rare. The activation function of the output layer is chosen to be the softmax function  $\sigma(x)_i = \exp(x_i) / \sum_{j=1}^{12} \exp(x_j)$ ,  $i = 1, \ldots, 12$ , such that the output is a probability vector. Hence, in contrast to the previous section, the output gives probabilities that the input belongs to the possible label combinations and does not return a specific combination. The usual loss for multi-class problems is used, namely the categorical cross entropy

$$\mathcal{L}(D) \coloneqq -\sum_{i=1}^{N} \sum_{j=1}^{12} y_{i,j} \ln(\hat{y}_{i,j}),$$

where N is the number of training data points,  $y_{i,j}$  is the *j*-th component of the *i*-th training data vector and  $\hat{y}_{i,j}$  the prediction of the network. Note that  $y_{i,j}$  is 0 for all except one entry where it is 1. To measure the performance of the networks the outputs of the networks are mapped back to label vectors by multiplying the probability of the classes with the corresponding label vectors of the classes and summing up the results. In other words, a weighted sum of all label vectors is computed where the weights correspond to the predicted propabilities. Afterwards the measures given in Section 3.4 can be computed.

In Figure 6 the total rating of the trained networks is plotted. The results are similarly to the results of Section 4.4 as also indicated by the values in Table 5. The best and the mean are negligibly smaller than the values obtained in the previous experiment. Hence, it does not matter whether to deal with the problem as a multi-label or a multi-class problem, at least in this particular setting with the used training set and the measures.



Figure 6: Results for all architectures from Section 4.5.

Table 5: Statistics about the total rating of the trained networks of Section 4.5. Standard deviation is abbreviated by std.

	$r_{\rm tot}$
max	0.949
mean	0.937
std	0.007

## 4.6 Applying a MLP limiter to higher polynomial degrees

Until here the experiments have investigated how good the MLPs can approximate the data but they have not been applied to the DG solution of (other) convection-dominated convection-diffusion

equations. To this end, the MLP from Section 4.4 with the best total rating is used, which has four hidden layers of 100 nodes and the hyperbolic tangent as activation function, and is trained with a learning rate of 0.0005, a batch size of 128, and is initialized with seed 42. After the DG solution of a convection-diffusion problem is solved, all features of the conventional limiter are calculated and the MLP is asked to predict the label combination given these features. A cell is finally marked if at least  $n \in [1, 2, 3, 4]$  of the four predicted labels are true, i.e., larger than 0.5. If a cell is marked then the solution is locally replaced by its integral mean, i.e.,  $\Pi_{MLP,K}(u_h) := \overline{u}_{h,K}$  since this choice has produced the best results in [19] and [20]. This limiter is below called *MLP* limiter.

#### 4.6.1 Determining the minimum number of predicted marks *n*

Since it is not a priori clear which value of n to choose, this is determined in a first step. The smaller n, the more cells are marked, which on the one hand hopefully leads to less spurious oscillations but on the other hand, marking too many cells might reduce the order of accuracy and leads to unnecessary computational overhead. Therefore, n should be chosen in such a way that enough but not too many cells are marked. To find the optimal n, Example 1 is used with exactly the same setting as in Section 3.1, i.e., the setting with which the data is created. Since the limiter is trained with this data it can be expected that it predicts the labels of the traditional limiter correctly in most of the cases. Since for Example 1 an analytical solution is not known, the discrete solution  $u_h$  cannot be compared against the exact solution. As in [19, 20], to assess the quality of the limited discrete solution therefore the measures

$$\operatorname{osc}_{\max}(u_h) = \max_{(x,y)\in\overline{\Omega}} u_h(x,y) - u_{\max} + u_{\min} - \min_{(x,y)\in\overline{\Omega}} u_h(x,y),$$
  
$$\operatorname{osc}_{\max}(u_h) = \frac{1}{|\mathcal{T}_h|} \sum_{K\in\mathcal{T}_h} \left[ \max\{0, \max_{(x,y)\in K} u_h(x,y) - u_{\max}\} + \max\{0, u_{\min} - \min_{(x,y)\in K} u_h(x,y)\} \right],$$

are used to measure the maximal size and a mean value of spurious oscillations, where  $u_{\text{max}}$  and  $u_{\text{min}}$  are the largest and smallest value of the weak solution, resp., and  $|\mathcal{T}_h|$  denotes the number of cells in the triangulation. In Example 1 it is  $u_{\text{min}} = 0$  and  $u_{\text{max}} = 1$ . To compute the desired quantities,  $u_h$  is evaluated at certain points, which are the points of the nodal functionals defining continuous  $P_p$  finite elements of the same order.

The results of the *MLP* limiter with n = 1, 2, 3, 4 on both the regular and the irregular grid are shown in Figure 7 and compared with *ConstJumpMod* and *ConstJumpNorm*, which are the classical limiters that works best for this problem [20].

It can be seen that the *MLP* limiter with n = 1 and n = 2 behaves similarly as well as those with n = 3 and n = 4. While the former ones are as good as the classical *ConstJumpMod* and *ConstJumpNorm* limiters, the latter ones behave much worse, meaning they lead to larger mean and maximal oscillations. As a consequence, in what follows, the *MLP* limiter with n = 2 is used since it produces better results than the ones with n = 3, 4 and should by definition mark less or the same amount of cells than the one with n = 1.



Figure 7: Results of  $osc_{mean}$  and  $osc_{max}$  for Example 1 on the regular and irregular grid depicted in Figure 2 for  $P_1$  finite elements with two classical limiters and various versions of the *MLP* limiter.



Figure 8: Results for measures for various limiters and various polynomial degrees obtained for Example 1 on the regular grid from Figure 2.



Figure 9: Results for measures for various limiters and various polynomial degrees obtained for Example 1 on the irregular grid from Figure 2.

#### 4.6.2 Higher polynomial degrees

Since the *MLP* limiter is fixed it can be applied to other problems. To start, again Example 1 is used but the limiter is applied to the discrete solution obtained with finite elements with higher polynomial degrees, namely  $P_2$ ,  $P_3$ , and  $P_4$  finite elements. The rest of the problem is not varied, i.e.,  $\varepsilon = 10^{-8}$ ,  $\kappa = 1$ ,  $\eta = 1$ . As in Section 3.1 the penalty parameter is chosen to be  $\sigma = 2\varepsilon n_0(p+1)(p+2)/2$ , where again  $n_0$  denotes the number of vertices a cell has. Also the parameters used in the classical limiters are kept the same. The problems are solved on the series of uniformly refined grids starting as above with the initial meshes depicted in Figure 2. In what follows *Galerkin* denotes the DG solution from Equation (3) without being limited.

The results for the measures  $osc_{mean}$  and  $osc_{max}$  for the best conventional limiter as well as the *MLP* limiter on both types of meshes are shown in Figures 8 and 9. It can be seen that the *MLP* limiter reduces both the mean and the maximal oscillations significantly compared to *Galerkin*. While on coarser grids it acts worse than *ConstJumpNorm* but better than *ConstJumpMod*, on finer grids all limiters almost coincide. A reason for this could be the fact that way more training data obtained on finer grinds is available compared to data from coarser grids, since the number of available data scales exponentially with the number of the refinement.

### 4.7 Applying a MLP limiter to the Hemker problem

Finally, in this section the *MLP* limiter is applied to a different example, namely the Hemker benchmark problem. It was proposed in [26] and it is a very popular benchmark problem for convection-dominated convection-diffusion equations. It models the transport of energy from a body through a channel and shows many features of problems that are also relevant in applications. The structure of the solution is similar to the solution of the HMM example, e.g., it is constant in most regions, and hence there is hope that the *MLP* limiter is able to limit the solution in a reasonable way.

**Example 2** (Hemker example). The problem is stated in  $\Omega = \{(-3,9) \times (-3,3)\} \setminus \{(x,y) :$ 

 $x^2 + y^2 \le 1$ }, and has the coefficients  $\boldsymbol{b} = (1, 0)^T$ , c = f = 0. If x = -3 and at the circular boundary, Dirichlet boundary conditions are prescribed by setting

$$g = 0$$
 if  $x = -3$ ,  $g = 1$  at the circle.

Everywhere else homogeneous Neumann conditions are applied. The solution is sketched in Figure 10 and takes values in [0, 1].



Figure 10: Sketch of the solution to Example 2 for  $\varepsilon = 10^{-8}$  obtained with a nonlinear algebraic flux-corrected (AFC) finite element method with Kuzmin limiter, see [5].



Figure 11: Initial mesh for Example 2 used in Section 4.7.

As before, the diffusion coefficient is set to  $\varepsilon = 10^{-8}$  and  $\kappa = 1$ ,  $\eta = 1$ , and  $\sigma = \varepsilon n_0 (p+1)(p+2)$  are used as parameters in the DG method. The problem is solved on a series of grids starting from the one depicted in Figure 11. The characteristic length scale for this problem is L = 13.5 and the remaining parameters of the limiter stay the same.

In Figure 12 the results for both measures for the limited solution and the original solution for various polynomial degrees are shown. As before *ConstJumpMod*, *ConstJumpNorm*, and the *MLP* limiter are able to reduce the oscillations significantly compared to *Galerkin*. For  $P_1$  and  $P_2$  the *MLP* limiter is slightly worse than the traditional ones on coarser grids but on finer grids it behaves equally well. For  $P_3$  and  $P_4$  it is worse than the classical companions, except for the finest grid and except for  $P_4$  for  $osc_{max}$  on one coarser level. We observed that the *MLP* limiter is always better than *LinTriaReco* in both measures, better than *ConstJumpNorm* for  $osc_{max}$  for all polynomial degree, and for  $osc_{mean}$  for  $P_1$  and for  $P_2$  to  $P_4$  on the two finer levels, but these results are not presented for the sake of brevity. A visualization of the limited  $P_4$  solution with *ConstJumpNorm* and *MLP* on the second finest grid is

shown in Figure 13. It can be observed that the *MLP* limiter limits most of the cells correctly but forgets to mark some cells with undershoots. This is also is the reason why it has worse  $osc_{mean}$  and  $osc_{max}$  values compared to *ConstJumpNorm*.



Figure 12: Results for measures for various limiters and various polynomial degrees obtained for Example 2.



Figure 13: Discrete  $P_4$  solution to Example 2 limited by the *ConstJumpNorm* limiter (left) and the *MLP* limiter (right).

## 5 Summary and outlook

This paper is a contribution to deeper understanding how neural network based slope limiters can be created and applied. In contrast to previous papers, it was focussed on constructing a multilayer perceptron model for limiting the discrete solution of an elliptical problem, namely convection-diffusion equations in the convection-dominated regime. It was shown how data from a lowest order discretization can be used to train a limiter that then can be applied to the discrete solution of higher order methods. The results have indicated that the limiter works almost equally well as classical methods for higher order methods for the same problem but somewhat worse than these methods when applied to the solution of a different problem.

In the future the impact of the features will be considered. Open questions include: Can the feature space be made smaller? Can the local degrees of freedom of the discrete solution be the features, i.e., can a MLP find a mapping between the local degrees of freedom and the labels from the decision maker functions? Can this MLP limiter also be used to mark cells for either adaptive refinement strategies or for adjusting the stabilization parameters in stabilized discretizations? These and other topics will be investigated in future work.

# Acknowledgement

The authors express their gratitude to Dr. Ulrich Wilbrandt for many time consuming but always valuable discussions.

# References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. Software available from https://www.tensorflow.org/.
- [2] Rémi Abgrall and Maria Han Veiga. Neural Network-Based Limiter with Transfer Learning. *Communications on Applied Mathematics and Computation*, Sep 2020.
- [3] Matthias Augustin, Alfonso Caiazzo, André Fiebach, Jürgen Fuhrmann, Volker John, Alexander Linke, and Rudolf Umla. An assessment of discretizations for convection-dominated convection–diffusion equations. *Computer Methods in Applied Mechanics and Engineering*, 200(47):3395– 3409, 2011.
- [4] Blanca Ayuso and L. Donatella Marini. Discontinuous Galerkin methods for advection-diffusion-reaction problems. *SIAM J. Numer. Anal.*, 47(2):1391–1420, 2009.
- [5] Gabriel R. Barrenechea, Volker John, Petr Knobloch, and Richard Rankin. A unified analysis of algebraic flux correction schemes for convection-diffusion equations. *SeMA J.*, 75(4):655–685, 2018.
- [6] Andrea Beck, David Flad, and Claus-Dieter Munz. Deep neural networks for data-driven les closure models. *Journal of Computational Physics*, 398:108910, 2019.
- [7] Andrea D. Beck, Jonas Zeifang, Anna Schwarz, and David G. Flad. A neural network based shock detection and localization approach for discontinuous Galerkin methods. *Journal of Computational Physics*, 423:109824, 2020.

- [8] Andrea Cangiani, Zhaonan Dong, Emmanuil H. Georgoulis, and Paul Houston. hp-version discontinuous Galerkin methods on polygonal and polyhedral meshes. SpringerBriefs in Mathematics. Springer, Cham, 2017.
- [9] Philippe G. Ciarlet. *The Finite Element Method for Elliptic Problems*. Classics in applied mathematics. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2002.
- [10] Bernardo Cockburn and Chi-Wang Shu. The Runge–Kutta Discontinuous Galerkin Method for Conservation Laws V: Multidimensional Systems. *Journal of Computational Physics*, 141(2):199–224, Apr 1998.
- [11] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.
- [12] Timothy A. Davis. Algorithm 832: UMFPACK V4.3–an unsymmetric-pattern multifrontal method. *ACM Trans. Math. Software*, 30(2):196–199, 2004.
- [13] TensorFlow Developers. TensorFlow, May 2022. v2.9.1, https://doi.org/10.5281/ zenodo.6574233.
- [14] Daniele Antonio Di Pietro and Alexandre Ern. Mathematical Aspects of Discontinuous Galerkin Methods, volume 69 of Mathématiques et Applications. Springer Verlag, Berlin, Heidelberg, 1st edition, 2012.
- [15] Vít Dolejší and Miloslav Feistauer. Discontinuous Galerkin Method: Analysis and Applications to Compressible Flow, volume 48 of Springer Series in Computational Mathematics. Springer International Publishing, Cham, 1 edition, 2015.
- [16] Vít Dolejší, Miloslav Feistauer, and Christoph Schwab. On some aspects of the discontinuous Galerkin finite element method for conservation laws. *Mathematics and Computers in Simulation*, 61(3–6):333–346, Jan 2003.
- [17] Vít Dolejší, Miroslav Feistauer, and Christoph Schwab. On discontinuous Galerkin methods for nonlinear convection-diffusion problems and compressible flow. In *Proceedings of EQUADIFF 10*, volume 127, page 163–179, Prague, 2002.
- [18] Vít Dolejší and Pavel Solin. *hp*-discontinuous Galerkin method based on local higher order reconstruction. *Appl. Math. Comput.*, 279:219–235, 2016.
- [19] Derk Frerichs and Volker John. On reducing spurious oscillations in discontinuous Galerkin (DG) methods for steady-state convection–diffusion equations. *Journal of Computational and Applied Mathematics*, 393:113487, Sep 2021.
- [20] Derk Frerichs-Mihov and Volker John. On a technique for reducing spurious oscillations in DG solutions of convection–diffusion equations. *Applied Mathematics Letters*, 129:107969, Jul 2022.
- [21] Derk Frerichs-Mihov, Ulrich Wilbrandt, Linus Henning, and Volker John. Data and code for using deep neural networks for detecting spurious oscillations in discontinuous Galerkin solutions of convection-dominated convection-diffusion equations, 2022. This work is licensed under CC BY 4.0, https://doi.org/10.20347/40vd-f944.

- [22] S. Ganesan, V. John, G. Matthies, R. Meesala, S. Abdus, and U. Wilbrandt. An object oriented parallel finite element scheme for computing PDEs: Design and implementation. In *IEEE 23rd International Conference on High Performance Computing Workshops (HiPCW) Hyderabad*, pages 106–115. IEEE, 2016.
- [23] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, volume 9 of Proceedings of Machine Learning Research, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 2010. PMLR.
- [24] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.
- [25] J. Gopalakrishnan and G. Kanschat. A multilevel discontinuous Galerkin method. *Numer. Math.*, 95(3):527–550, 2003.
- [26] P. W. Hemker. A singularly perturbed model problem for numerical computation. J. Comput. Appl. Math., 76(1-2):277–285, 1996.
- [27] Catherine F. Higham and Desmond J. Higham. Deep Learning: An Introduction for Applied Mathematicians. *SIAM Review*, 61(4):860–891, Jan 2019.
- [28] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- [29] Thomas J. R. Hughes, Michel Mallet, and Akira Mizukami. A new finite element formulation for computational fluid dynamics. II. Beyond SUPG. *Comput. Methods Appl. Mech. Engrg.*, 54(3):341–355, 1986.
- [30] Sergio Izquierdo. CppFlow, Sep 2022. v2.0.0, https://github.com/serizba/ cppflow, https://serizba.github.io/cppflow/.
- [31] Volker John and Petr Knobloch. A computational comparison of methods diminishing spurious oscillations in finite element solutions of convection-diffusion equations. In T. Vejchodsky J. Chleboun, K. Segeth, editor, *Proceedings of the International Conference on Programs and Algorithms of Numerical Mathematics 13*, pages 122–136, Prague, 2006. Academy of Science of the Czech Republic.
- [32] Volker John and Petr Knobloch. On Discontinuity-Capturing Methods for Convection-Diffusion Equations. In Alfredo Bermúdez de Castro, Dolores Gómez, Peregrina Quintela, and Pilar Salgado, editors, *Numerical Mathematics and Advanced Applications*, pages 336–344, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [33] Subodh M. Joshi, Thivin Anandh, Bhanu Teja, and Sashikumaar Ganesan. On the choice of hyper-parameters of artificial neural networks for stabilized finite element schemes. *International Journal of Advances in Engineering Sciences and Applied Mathematics*, 13:278–297, 2020.
- [34] Guido Kanschat. *Discontinuous Galerkin Methods for Viscous Incompressible Flow*. Advances in numerical mathematics. Teubner Research, Dt. Univ.-Verl, Wiesbaden, 1st edition, 2007.
- [35] Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method For Stochastic Optimization. In *ICLR 2015*, page 13. arXiv, 2014.

- [36] Miroslav Kubat. *An introduction to machine learning*. Springer Nature Switzerland, Cham, 3rd edition, 2021.
- [37] Yang Liu, Yutong Lu, Yueqing Wang, Dong Sun, Liang Deng, Fang Wang, and Yan Lei. A CNN-based shock detection method in flow visualization. *Computers & Fluids*, 184:1–9, 2019.
- [38] Nils Margenberg, Christian Lessig, and Thomas Richter. Structure preservation for the deep neural network multigrid solver. *ETNA Electronic Transactions on Numerical Analysis*, 56:86–101, 2021.
- [39] Rômulo Montalvão Silva and Alvaro Coutinho. PINNs for Parametric Incompressible Newtonian Flows. In Proceedings of the XLII Ibero-Latin-American Congress on Computational Methods in Engineering and III Pan-American Congress on Computational Mechanics, ABMEC-IACM, 11 2021.
- [40] Nathaniel R. Morgan, Svetlana Tokareva, Xiaodong Liu, and Andrew Morgan. A machine learning approach for detecting shocks with high-order hydrodynamic methods. In *AIAA Scitech 2020 Forum*, 2020.
- [41] Deep Ray and Jan S. Hesthaven. An artificial neural network as a troubled-cell indicator. *Journal of Computational Physics*, 367:166–191, Aug 2018.
- [42] Deep Ray and Jan S. Hesthaven. Detecting troubled-cells on two-dimensional unstructured grids using a neural network. *Journal of Computational Physics*, 397:108845, 2019.
- [43] W.H. Reed and T.R. Hill. Triangular mesh methods for the neutron Transport Equation. Technical Report LA-UR-73-479, Los Alamos Scientific Laboratory, Los Alamos, NM, 1973.
- [44] Béatrice Rivière. Discontinuous Galerkin Methods for Solving Elliptic and Parabolic Equations: Theory and Implementation, volume 35 of Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, Jan 2008.
- [45] Raul Rojas. Networks of width one are universal classifiers. In *Proceedings of the International Joint Conference on Neural Networks*, volume 4, pages 3124–3127, 2003.
- [46] Hans-Görg Roos, Martin Stynes, and Lutz Tobiska. Robust Numerical Methods for Singularly Perturbed Differential Equations: Convection-Diffusion-Reaction and Flow Problems, volume 24 of Springer Series in Computational Mathematics. Springer-Verlag, Berlin, Heidelberg, 2nd edition, 2008.
- [47] Maria Han Veiga and Rémi Abgrall. Towards a general stabilisation method for conservation laws using a multilayer perceptron neural network: 1d scalar and system of equations. In European Conference on Computational Mechanics and VII European Conference on Computational Fluid Dynamics, number 1, pages 2525–2550. ECCM, Juni 2018.
- [48] Henry von Wahl and Thomas Richter. Using a deep neural network to predict the motion of underresolved triangular rigid bodies in an incompressible flow. *International Journal for Numerical Methods in Fluids*, 93(12):3364–3383, 2021.
- [49] Ulrich Wilbrandt, Clemens Bartsch, Naveed Ahmed, Najib Alia, Felix Anker, Laura Blank, Alfonso Caiazzo, Sashikumaar Ganesan, Swetlana Giere, Gunar Matthies, Raviteja Meesala, Abdus Shamim, Jagannath Venkatesan, and Volker John. ParMooN–A modernized program package based on mapped finite elements. *Comput. Math. Appl.*, 74(1):74–88, 2017.