# Non-intrusive reduced order modeling of nonlinear problems using neural networks 

J. S. Hesthaven ${ }^{\text {a }}$, S. Ubbiali ${ }^{\text {a,b }}$<br>${ }^{a}$ École Polytechnique Fédérale de Lausanne (EPFL), Route Cantonale, 1015 Lausanne, Switzerland<br>${ }^{b}$ Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milan, Italy


#### Abstract

We develop a non-intrusive reduced basis (RB) method for parametrized steady-state partial differential equations (PDEs). The method extracts a reduced basis from a collection of high-fidelity solutions via a proper orthogonal decomposition (POD) and employs artificial neural networks (ANNs), particularly multi-layer perceptrons (MLPs), to accurately approximate the coefficients of the reduced model. The search for the optimal number of neurons and the minimum amount of training samples to avoid overfitting is carried out in the offline phase through an automatic routine, relying upon a joint use of the latin hypercube sampling (LHS) and the Levenberg-Marquardt training algorithm. This guarantees a complete offline-online decoupling, leading to an efficient RB method - referred to as POD-NN - suitable also for general nonlinear problems with a non-affine parametric dependence. Numerical studies are presented for the nonlinear Poisson equation and for driven cavity viscous flows, modeled through the steady incompressible Navier-Stokes equations. Both physical and geometrical parametrizations are considered. Several results confirm the accuracy of the POD-NN method and show the substantial speed-up enabled at the online stage as compared to a traditional RB strategy.


Keywords: non-intrusive reduced basis method, proper orthogonal decomposition, multi-layer perceptron, Levenberg-Marquardt algorithm, Poisson equation, driven cavity flow

## 1. Introduction

Many applications in engineering and the applied sciences involve mathematical models expressed as parametrized partial differential equations (PDEs), in which boundary conditions, material properties, source terms, loads or geometric features of the underlying physical problem are expressed by a parameter $\boldsymbol{\mu}$ [18,23, 26]. A list of notable examples includes parameter estimation [6], topology optimization [5], optimal control [30] and uncertainty quantification [29]. In these examples, one is typically interested in a real-time evaluation of an output of interest (defined as a functional of the state variable [15]) for many parameter entries, i.e., for many configurations of the problem.

The increasing computational power and the simultaneous algorithmic improvements enable nowadays the highfidelity numerical resolution of complex problems via standard discretization procedures, such as finite difference (FD), finite volume (FV), finite element (FE), or spectral methods [41]. However, these schemes remain prohibitively expensive in many-query and real-time contexts, both in terms of CPU time and memory demand, due to the large amount of degrees of freedom (DOFs) they need to accurately solve the PDE [1]. In light of this, reduced order modeling (ROM) methods have received a significant attention in the last decades. The objective of these methods is to replace the full-order system by one of significant smaller dimension, to decrease the computational burden while leading to a controlled loss of accuracy [11].

Reduced basis (RB) methods constitute a well-known and widely-used example of reduced order modeling techniques. They are generally implemented pursuing an offline-online paradigm [31]. Based upon an ensemble of snapshots (i.e., highfidelity solutions to the parametrized differential problem), the goal of the offline step is to construct a solution-dependent basis, yielding a reduced space of globally approximating functions to represent the main dynamics of the full-order model [2, 11]. For this, two major approaches have been proposed in the literature: proper orthogonal decomposition (POD) [22, 46] and greedy algorithms [24]. The former relies on a deterministic or random sample in the parameter space to

[^0]generate snapshots and then employs a singular value decomposition (SVD) to recover the reduced basis. In the second approach, the basis vectors coincide with the snapshots themselves, carefully selected according to some optimality criterion. As a result, a greedy strategy is typically more effective and efficient than POD, as it enables the exploration of a wider region of the parameter space while entailing the computation of many fewer high-fidelity solutions [23]. However, there exist problems for which a greedy approach is not feasible, simply because a natural criterion for the choice of the snapshots is not available [2].

Once a reduced order framework has been properly set up, an approximation to the truth solution for a new parameter value is sought online as a linear combination of the RB functions, with the expansion coefficients determined via a projection of the full-order system onto the reduced space [7]. To this end, a Galerkin procedure is the most popular choice.

Despite their established effectiveness, projection-based RB methods do not provide any computational gain with respect to a direct (expensive) approach for complex nonlinear problems with a non-affine dependence on the parameters. This is a result of the cost to compute the projection coefficients, which depends on the dimension of the full-order model. In fact, a full decoupling between the online stage and the high-fidelity scheme is the ultimate secret for the success of any RB procedure [41]. For this purpose, one may recover an affine expansion of the differential operator through the empirical interpolation method (EIM) [3] or its discrete variants [10, 37]. However, for general nonlinear problems this is far from trivial.

A valuable alternative to address this concern is represented by non-intrusive RB methods, in which the high-fidelity model is used to generate the snapshots, but not in the projection process [11]. The projection coefficients are obtained via interpolation over the parameter domain of a database of reduced order information [9]. However, since reduced bases generally belong to nonlinear, matrix manifolds, standard interpolation techniques may fail, as they cannot enforce the constraints characterizing those manifolds, unless employing a large amount of samples [1, 4].

In this work, we develop a non-intrusive RB method employing POD for the generation of the reduced basis and resort to (artificial) neural networks, in particular multi-layer perceptrons, in the interpolation step. Hence, in the following we refer to the proposed RB procedure as the POD-NN method. Being of non-intrusive nature, POD-NN is suitable for a fast and reliable resolution of complex nonlinear PDEs featuring a non-affine parametric dependence. To test this assertion, the POD-NN method is applied to the one- and two-dimensional nonlinear Poisson equation and to the steady incompressible Navier-Stokes equations. Both physical and geometrical parametrizations are considered.

The paper is organized as follows. Section 2 defines the (parametrized) functional and variational framework which is required to develop a finite element solver, briefly outlined in Subsection 2.2. The standard projection-based POD-Galerkin (POD-G) RB method is derived in Section 3. Section 4 discusses components, topology and learning process for artificial neural networks. This is preparatory for the subsequent Section 5, which details the non-intrusive POD-NN RB procedure; both theoretical and practical aspects are addressed. Several numerical results, aiming to show the reliability and efficiency of the proposed RB technique, are offered in Section 6 for the Poisson equation (Subsection 6.1) and the lid-driven cavity problem for the steady Navier-Stokes equations (Subsection 6.2). Finally, Section 7 gathers some relevant conclusions and suggests future developments.

## 2. Parametrized partial differential equations

Assume $\mathscr{P}_{p h} \subset \mathbb{R}^{P_{p h}}$ and $\mathscr{P}_{g} \subset \mathbb{R}^{P_{g}}$ are compact sets, and let $\boldsymbol{\mu}_{p h} \in \mathscr{P}_{p h}$ and $\boldsymbol{\mu}_{g} \in \mathscr{P}_{g}$ be respectively the physical and geometrical parameters characterizing the differential problem, so that $\boldsymbol{\mu}=\left(\boldsymbol{\mu}_{p h}, \boldsymbol{\mu}_{g}\right) \in \mathscr{P}=\mathscr{P}_{p h} \times \mathscr{P}_{g} \subset \mathbb{R}^{P}$, with $P=P_{p h}+P_{g}$, represents the overall input vector parameter. While $\boldsymbol{\mu}_{p h}$ addresses material properties, source terms and boundary conditions, $\boldsymbol{\mu}_{g}$ defines the shape of the computational domain $\widetilde{\Omega}=\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right) \subset \mathbb{R}^{d}, d=1,2$. We denote by $\widetilde{\Gamma}\left(\boldsymbol{\mu}_{g}\right)=\partial \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)$ the (Lipschitz) boundary of $\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)$, and by $\widetilde{\Gamma}_{D}\left(\boldsymbol{\mu}_{g}\right)$ and $\widetilde{\Gamma}_{N}\left(\boldsymbol{\mu}_{g}\right)$ the portions of $\widetilde{\Gamma}\left(\boldsymbol{\mu}_{g}\right)$ where Dirichlet and Neumann boundary conditions are enforced, respectively, with $\widetilde{\Gamma}_{D} \cup \widetilde{\Gamma}_{N}=\widetilde{\Gamma}$ and $\stackrel{\circ}{\Gamma}_{D} \cap \stackrel{\circ}{\Gamma}_{N}=\varnothing$.

Consider a Hilbert space $\widetilde{V}=\widetilde{V}\left(\boldsymbol{\mu}_{g}\right)=\widetilde{V}\left(\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)\right)$ defined over the domain $\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)$, equipped with the scalar product $(\cdot, \cdot) \widetilde{V}_{\widetilde{V}}$ and the induced norm $\|\cdot\|_{\widetilde{V}}=\sqrt{(\cdot, \cdot)_{\widetilde{V}}}$. Furthermore, let $\widetilde{V}^{\prime}=\widetilde{V}^{\prime}\left(\boldsymbol{\mu}_{g}\right)$ be the dual space of $\widetilde{V}$. Denoting by $\widetilde{G}: \widetilde{V} \times \mathscr{P}_{p h} \rightarrow \widetilde{V}^{\prime}$ the map representing a parametrized nonlinear second-order PDE, the differential (strong) form of the problem of interest reads: given $\boldsymbol{\mu}=\left(\boldsymbol{\mu}_{p h}, \boldsymbol{\mu}_{g}\right) \in \mathscr{P}$, find $\widetilde{u}(\boldsymbol{\mu}) \in \widetilde{V}\left(\boldsymbol{\mu}_{g}\right)$ such that

$$
\begin{equation*}
\widetilde{G}\left(\widetilde{u}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right)=0 \quad \text { in } \widetilde{V}^{\prime}\left(\boldsymbol{\mu}_{g}\right), \tag{2.1}
\end{equation*}
$$

namely

$$
\left\langle\widetilde{G}\left(\widetilde{u}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right), v\right\rangle_{\tilde{V}^{\prime}, \tilde{V}}=0 \quad \forall v \in \widetilde{V}\left(\boldsymbol{\mu}_{g}\right),
$$

with $\langle\cdot, \cdot\rangle_{\tilde{V}^{\prime}} \tilde{V}: \widetilde{V}^{\prime} \times \widetilde{V} \rightarrow \mathbb{R}$ the duality pairing between $\widetilde{V}^{\prime}$ and $\widetilde{V}$.
The finite element method requires problem (2.1) to be stated in a weak (or variational) form [40]. To this end, let us introduce the form $\widetilde{g}: \widetilde{V} \times \widetilde{V} \times \mathscr{P} \rightarrow \mathbb{R}$, with $\widetilde{g}(\cdot, \cdot ; \boldsymbol{\mu})$ defined as:

$$
\widetilde{g}(w, v ; \boldsymbol{\mu})=\left\langle\widetilde{G}\left(w ; \boldsymbol{\mu}_{p h}\right), v\right\rangle_{\tilde{V}^{\prime}, \widetilde{V}} \quad \forall w, v \in \widetilde{V}
$$

The variational formulation of (2.1) then reads: given $\boldsymbol{\mu}=\left(\boldsymbol{\mu}_{p h}, \boldsymbol{\mu}_{g}\right) \in \mathscr{P}$, find $\widetilde{u}(\boldsymbol{\mu}) \in \widetilde{V}\left(\boldsymbol{\mu}_{g}\right)$ such that

$$
\widetilde{g}(\widetilde{u}(\boldsymbol{\mu}), v ; \boldsymbol{\mu})=0 \quad \forall v \in \widetilde{V}\left(\boldsymbol{\mu}_{g}\right)
$$

### 2.1. From physical to reference domain

As discussed in the Introduction, a reduced basis method seeks an approximated solution to a problem as a combination of (few) well-chosen basis vectors. These typically result from a suitable combination of a collection of high-fidelity approximations, called snapshots. Therefore, when addressing problems defined on variable shape domains, ensuring the compatibility among snapshots is crucial. To this end, it is practice to formulate and solve the differential problem over a fixed, parameter-independent domain $\Omega \subset \mathbb{R}^{d}$ [35]. This can be accomplished by introducing a parametrized map $\boldsymbol{\Phi}: \Omega \times \mathscr{P}_{g} \rightarrow \widetilde{\Omega}$ such that

$$
\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)=\boldsymbol{\Phi}\left(\Omega ; \boldsymbol{\mu}_{g}\right)
$$

The transformation $\boldsymbol{\Phi}\left(\cdot ; \boldsymbol{\mu}_{g}\right)$ allows one to restate the general problem (2.1). Let $V$ be a suitable Hilbert space over $\Omega$ and $V^{\prime}$ be its dual. Suppose $V$ is equipped with the scalar product $(\cdot, \cdot)_{V}$ and the induced norm $\|\cdot\|_{V}=\sqrt{(\cdot, \cdot)_{V}}$. Given the parametrized map $G: V \times \mathscr{P} \rightarrow V^{\prime}$ representing the (nonlinear) PDE over the reference domain $\Omega$, we focus on problems of the form: given $\boldsymbol{\mu} \in \mathscr{P}$, find $u(\boldsymbol{\mu}) \in V$ such that

$$
\begin{equation*}
G(u(\boldsymbol{\mu}) ; \boldsymbol{\mu})=0 \quad \text { in } V^{\prime} \tag{2.2}
\end{equation*}
$$

The weak formulation of problem (2.2) reads: given $\boldsymbol{\mu} \in \mathscr{P}$, seek $u(\boldsymbol{\mu}) \in V$ such that

$$
\begin{equation*}
g(u(\boldsymbol{\mu}), v ; \boldsymbol{\mu})=0 \quad \forall v \in V \tag{2.3}
\end{equation*}
$$

where $g: V \times V \times \mathscr{P} \rightarrow \mathbb{R}$ is defined as

$$
g(w, v ; \boldsymbol{\mu})=\langle G(w ; \boldsymbol{\mu}), v\rangle_{V^{\prime}, V} \quad \forall w, v \in V, \forall \boldsymbol{\mu} \in \mathscr{P}
$$

with $\langle\cdot, \cdot\rangle_{V^{\prime}, V}: V^{\prime} \times V \rightarrow \mathbb{R}$ the dual pairing between $V$ and $V^{\prime}$. Observe that the explicit expression of $g(\cdot, \cdot ; \boldsymbol{\mu})$ involves the map $\Phi\left(\cdot ; \boldsymbol{\mu}_{g}\right)$, thus keeping track of the original domain $\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)$. Then, the solution $\widetilde{u}(\boldsymbol{\mu})$ over the original domain $\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)$ can be recovered as

$$
\widetilde{u}(\boldsymbol{\mu})=u(\boldsymbol{\mu}) \circ \boldsymbol{\Phi}\left(\boldsymbol{\mu}_{g}\right)
$$

In our numerical tests, we employ the squared reference domain shown on the right in Fig. 6.3 and we resort to a particular choice for $\boldsymbol{\Phi}\left(\cdot ; \boldsymbol{\mu}_{g}\right)$ - the boundary displacement-dependent transfinite map (BDD TM) proposed by Jaggli et al. [26].

### 2.2. Discrete full-order model

Let $V_{h} \subset V$ be a suitable FE subspace of $V$ of (finite) dimension $N_{h}$, with $h \geq 0$ being the characteristic size of the underlying mesh $\Omega_{h}$ which discretizes $\Omega$. The FE approximation of the weak problem (2.3) can be cast in the form: given $\boldsymbol{\mu} \in \mathscr{P}$, find $u_{h}(\boldsymbol{\mu}) \in V_{h}$ such that

$$
\begin{equation*}
g\left(u_{h}(\boldsymbol{\mu}), v_{h} ; \boldsymbol{\mu}\right)=0 \quad \forall v_{h} \in V_{h} \tag{2.4}
\end{equation*}
$$

From an algebraic standpoint, letting $\left\{\phi_{1}, \ldots, \phi_{N_{h}}\right\}$ be a Lagrangian basis for $V_{h}$ and denoting by $\mathbf{u}_{h}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{h}}$ the vector collecting the nodal values $\left\{u_{h}^{(1)}(\boldsymbol{\mu}), \ldots, u_{h}^{\left(N_{h}\right)}(\boldsymbol{\mu})\right\}$ of $u_{h}(\boldsymbol{\mu})$, problem (2.4) is equivalent to: given $\boldsymbol{\mu} \in \mathscr{P}$, find $\mathbf{u}_{h}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{h}}$ such that

$$
\begin{equation*}
\mathbf{G}_{h}\left(\mathbf{u}_{h}(\boldsymbol{\mu}) ; \boldsymbol{\mu}\right)=\mathbf{0} \in \mathbb{R}^{N_{h}} \tag{2.5}
\end{equation*}
$$

where the $i$-th component of the residual vector $\mathbf{G}_{h}(\cdot ; \boldsymbol{\mu})$ is given by

$$
\begin{equation*}
\left(\mathbf{G}_{h}\left(\mathbf{u}_{h}(\boldsymbol{\mu}) ; \boldsymbol{\mu}\right)\right)_{i}=g\left(u_{h}(\boldsymbol{\mu}), \phi_{i} ; \boldsymbol{\mu}\right), \quad i=1, \ldots, N_{h} . \tag{2.6}
\end{equation*}
$$

Observe that, due to the nonlinearity of $g(\cdot, \cdot ; \boldsymbol{\mu})$ in its first argument, one has to resort to some iterative method, e.g., Newton's method, to solve the Galerkin problem (2.5).

## 3. Projection-based reduced basis method

As outlined above, the finite element discretization of the $\boldsymbol{\mu}$-dependent nonlinear differential problem (2.3), combined with Newton's method, entails the assembly and solution of (possibly) many linear systems, whose dimension $N_{h}$ is directly related to $(i)$ the size of the underlying grid and (ii) the order of the polynomial FE space adopted. Since the accuracy of the resulting discretization heavily relies on these two factors, a direct numerical approximation of the full-order model implies severe computational costs. Therefore, this approach is not affordable in many-query and real-time contexts. This motivates the use of reduced order models. Particularly, reduced basis methods seek an approximate solution to problem (2.3) as a linear combination of parameter-independent functions $\left\{\psi_{1}, \ldots, \psi_{L}\right\} \subset V_{h}$, called reduced basis functions, built from a collection of high-fidelity snapshots $\left\{u_{h}\left(\boldsymbol{\mu}^{(1)}\right), \ldots, u_{h}\left(\boldsymbol{\mu}^{(N)}\right)\right\}$, where the discrete and finite set $\Xi_{N}=\left\{\boldsymbol{\mu}^{(1)}, \ldots, \boldsymbol{\mu}^{(N)}\right\} \subset \mathscr{P}$ may consist of either a uniform lattice or randomly generated points over the parameter domain $\mathscr{P}$ [23]. The basis functions $\left\{\psi_{l}\right\}_{1 \leq l \leq L}$ generally follow from a principal component analysis (PCA) of the set of snapshots (in that case, $N>L$ ), or they might coincide with the snapshots themselves (in that case, $N=L$ ). In the latter approach, typical of any greedy method, the parameters $\left\{\boldsymbol{\mu}^{(n)}\right\}_{1 \leq n \leq N}$ must be carefully chosen following some optimality criterium (see, e.g., [11]). Here, we pursue the first approach, employing the well-known proper orthogonal decomposition (POD) method [22, 46], detailed in the following subsection. For now, assume that a reduced basis is available and let $V_{\mathrm{rb}} \subset V_{h}$ be the associated reduced basis space, i.e.,

$$
V_{\mathrm{rb}}=\operatorname{span}\left\{\psi_{1}, \ldots, \psi_{L}\right\} .
$$

A reduced basis solution $u_{L}(\boldsymbol{\mu})$ is sought in the form

$$
u_{L}(\boldsymbol{x} ; \boldsymbol{\mu})=\sum_{l=1}^{L} u_{\mathrm{rb}}^{(l)}(\boldsymbol{\mu}) \psi_{l}(\boldsymbol{x}) \in V_{\mathrm{rb}},
$$

with

$$
\mathbf{u}_{\mathrm{rb}}(\boldsymbol{\mu})=\left[u_{\mathrm{rb}}^{(1)}(\boldsymbol{\mu}), \ldots, u_{\mathrm{rb}}^{(L)}(\boldsymbol{\mu})\right]^{T} \in \mathbb{R}^{L}
$$

being the reduced coefficients (also called generalized coordinates) for the expansion of the RB solution in the RB basis functions. To recover $u_{L}(\boldsymbol{\mu})$, we proceed to project the variational problem (2.3) onto the RB space $V_{\mathrm{rb}}$ by pursuing a standard Galerkin approach, leading to the following reduced basis problem: given $\boldsymbol{\mu} \in \mathscr{P}$, find $u_{L}(\boldsymbol{\mu}) \in V_{\mathrm{rb}}$ so that

$$
\begin{equation*}
g\left(u_{L}(\boldsymbol{\mu}), v_{L} ; \boldsymbol{\mu}\right)=0 \quad \forall v_{L} \in V_{\mathrm{rb}} . \tag{3.1}
\end{equation*}
$$

Denoting by $\boldsymbol{\psi}_{l} \in \mathbb{R}^{N_{h}}$ the vector gathering the nodal values of $\psi_{l}$, for $l=1, \ldots, L$, let us introduce the matrix

$$
\mathbb{V}=\left[\boldsymbol{\psi}_{1}|\ldots| \boldsymbol{\psi}_{L}\right] \in \mathbb{R}^{N_{h} \times L} .
$$

For any $v_{L} \in V_{\mathrm{rb}}, \mathbb{V}$ encodes the change of variables from the RB basis to the Lagrangian FE basis, i.e.,

$$
\begin{equation*}
\mathbf{v}_{L}=\mathbb{V} \mathbf{v}_{\mathrm{rb}} . \tag{3.2}
\end{equation*}
$$

Then, due to (2.6) and (3.2), the algebraic formulation of the reduced basis problem (3.1) reads: given $\boldsymbol{\mu} \in \mathscr{P}$, seek $\mathbf{u}_{\mathrm{rb}}(\boldsymbol{\mu}) \in \mathbb{R}^{L}$ such that

$$
\begin{equation*}
\mathbf{G}_{\mathrm{rb}}\left(\mathbf{u}_{\mathrm{rb}}(\boldsymbol{\mu}) ; \boldsymbol{\mu}\right)=\mathbb{V}^{T} \mathbf{G}_{h}\left(\mathbf{u}_{L}(\boldsymbol{\mu}) ; \boldsymbol{\mu}\right)=\mathbb{V}^{T} \mathbf{G}_{h}\left(\mathbb{V} \mathbf{u}_{\mathrm{rb}}(\boldsymbol{\mu}) ; \boldsymbol{\mu}\right)=\mathbf{0} \in \mathbb{R}^{L} . \tag{3.3}
\end{equation*}
$$

### 3.1. Proper orthogonal decomposition

Consider a collection of $N$ snapshots $\left\{u_{h}\left(\boldsymbol{\mu}^{(1)}\right), \ldots, u_{h}\left(\boldsymbol{\mu}^{(N)}\right)\right\}$, corresponding to the finite and discrete parameter set $\Xi_{N}=\left\{\boldsymbol{\mu}^{(1)}, \ldots, \boldsymbol{\mu}^{(N)}\right\} \subset \mathscr{P}$, and let $\mathscr{M}_{\Xi_{N}}$ be the associated subspace, i.e.,

$$
\mathscr{M}_{\Xi_{N}}=\operatorname{span}\left\{u_{h}\left(\boldsymbol{\mu}^{(1)}\right), \ldots, u_{h}\left(\boldsymbol{\mu}^{(N)}\right)\right\} .
$$

We assume that $\mathscr{M}_{N}$ provides a good approximation of the discrete solution manifold $\mathscr{M}_{h}$,

$$
\mathscr{M}_{h}=\left\{u_{h}(\boldsymbol{\mu}): \boldsymbol{\mu} \in \mathscr{P}\right\},
$$

as long as the number of snapshots is sufficiently large (but typically much smaller than the dimension $N_{h}$ of the FE space). Then, we aim at finding a parameter-independent reduced basis for $\mathscr{M}_{\Xi_{N}}$, i.e., a collection of FE functions $\left\{\psi_{1}, \ldots, \psi_{L}\right\} \subset \mathscr{M}_{\Xi_{N}}$, with $L \ll N_{h}$ and $L$ independent of $N_{h}$, so that the associated linear space constitutes a low-rank approximation of $\mathscr{M}_{\Xi_{N}}$, optimal in some sense to be defined later. To this end, consider the snapshot matrix $\mathbb{S} \in \mathbb{R}^{N_{h} \times N}$ gathering the nodal values of the snapshots in a column-wise sense, i.e.,

$$
\mathbb{S}=\left[\mathbf{u}_{h}\left(\boldsymbol{\mu}^{(1)}\right)|\ldots| \mathbf{u}_{h}\left(\boldsymbol{\mu}^{(N)}\right)\right] .
$$

Denoting by $R$ the rank of $\mathbb{S}$, with $R \leq \min \left\{N_{h}, N\right\}$, the singular value decomposition (SVD) of $\mathbb{S}$ ensures the existence of two orthogonal matrices $\mathbb{W}=\left[\mathbf{w}_{1}|\ldots| \mathbf{w}_{N_{h}}\right] \in \mathbb{R}^{N_{h} \times N_{h}}$ and $\mathbb{Z}=\left[\mathbf{z}_{1}|\ldots| \mathbf{z}_{N}\right] \in \mathbb{R}^{N \times N}$, and a diagonal matrix $\mathbb{D}=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{R}\right) \in$ $\mathbb{R}^{R \times R}$, with $\sigma_{1} \geq \sigma_{2} \geq \ldots \geq \sigma_{r}>0$, such that

$$
\mathbb{S}=\mathbb{W}\left[\begin{array}{cc}
\mathbb{D} & 0 \\
0 & 0
\end{array}\right] \mathbb{Z}^{T}=\mathbb{W} \Sigma \mathbb{Z}^{T},
$$

where the zeros denote null matrices of appropriate dimensions. The real values $\left\{\sigma_{i}\right\}_{1 \leq i \leq R}$ are called singular values of $\mathbb{S}$, the columns $\left\{\mathbf{w}_{m}\right\}_{1 \leq m \leq N_{h}}$ of $\mathbb{W}$ are called left singular vectors of $\mathbb{S}$, and the columns $\left\{\mathbf{z}_{n}\right\}_{1 \leq n \leq N}$ of $\mathbb{Z}$ are called right singular vectors of $\mathbb{S}$, and they are related by the following relations:

$$
\begin{array}{ll}
\mathbb{S S}^{T} \mathbf{w}_{m}= \begin{cases}\sigma_{m}^{2} \mathbf{w}_{m} & \text { for } 1 \leq m \leq R, \\
\mathbf{0} & \text { for } R+1 \leq m \leq N_{h},\end{cases} & \mathbb{S}^{T} \mathbb{S} \mathbf{z}_{n}= \begin{cases}\sigma_{n}^{2} \mathbf{z}_{n} & \text { for } 1 \leq n \leq R, \\
\mathbf{0} & \text { for } R+1 \leq n \leq N,\end{cases}  \tag{3.4}\\
\mathbb{S} \mathbf{z}_{i}=\sigma_{i} \mathbf{w}_{i} \quad \text { for } 1 \leq i \leq R, & \mathbb{S}^{T} \mathbf{w}_{i}=\sigma_{i} \mathbf{z}_{i} \quad \text { for } 1 \leq i \leq R .
\end{array}
$$

At the algebraic level, our goal is to approximate the columns of $\mathbb{S}$ by means of $L$ orthonormal vectors $\left\{\widetilde{\mathbf{w}}_{1}, \ldots, \widetilde{\mathbf{w}}_{L}\right\}$, with $L<R$. It is an easy matter to show that for each $\mathbf{s}_{n}, n=1, \ldots, N$, the element of $\operatorname{span}\left\{\widetilde{\mathbf{w}}_{1}, \ldots, \widetilde{\mathbf{w}}_{L}\right\}$ closest to $\mathbf{s}_{n}$ in the Euclidean norm $\|\cdot\|_{\mathbb{R}^{N_{h}}}=\sqrt{(\cdot, \cdot)_{\mathbb{R}^{N_{h}}}}$ is given by

$$
\sum_{l=1}^{L}\left(\mathbf{s}_{n}, \widetilde{\mathbf{w}}_{l}\right)_{\mathbb{R}^{N_{h}}} \widetilde{\mathbf{w}}_{l} .
$$

Hence, we could measure the error committed by approximating the columns of $\mathbb{S}$ via the vectors $\left\{\widetilde{\mathbf{w}}_{l}\right\}_{1 \leq l \leq L}$ through the quantity

$$
\begin{equation*}
\varepsilon\left(\widetilde{\mathbf{w}}_{1}, \ldots, \widetilde{\mathbf{w}}_{L}\right)=\sum_{n=1}^{N}\left\|\mathbf{s}_{n}-\sum_{l=1}^{L}\left(\mathbf{s}_{n}, \widetilde{\mathbf{w}}_{l}\right)_{\mathbb{R}^{N_{h}}} \widetilde{\mathbf{w}}_{l}\right\|_{\mathbb{R}^{N_{h}}}^{2} . \tag{3.5}
\end{equation*}
$$

The Schmidt-Eckart-Young theorem $[17,44]$ states that the POD basis of $\operatorname{rank} L\left\{\mathbf{w}_{1}, \ldots, \mathbf{w}_{L}\right\}$, consisting of the first $L$ left singular values of $\mathbb{S}$, minimizes (3.5) among all the orthonormal bases of $\mathbb{R}^{L}$. Therefore, in the POD-Galerkin RB method, we set $\boldsymbol{\psi}_{l}=\mathbf{w}_{l}$, for all $l=1, \ldots, L$, so that

$$
\mathbb{V}=\left[\mathbf{w}_{1}|\ldots| \mathbf{w}_{L}\right] .
$$

From a computational viewpoint, the first $L$ left singular vectors $\left\{\mathbf{w}_{l}\right\}_{1 \leq l \leq L}$ of $\mathbb{S}$ can be efficiently computed through the so-called method of snapshots. We should distinguish two cases:
(a) if $N_{h} \leq N$ : directly solve the eigenvalue problems $\mathbb{S} \mathbb{S}^{T} \mathbf{w}_{l}=\lambda_{l} \mathbf{w}_{l}$, for $1 \leq l \leq L$;
(b) if $N_{h}>N$ : compute the correlation matrix $\mathbb{M}=\mathbb{S}^{T} \mathbb{S}$ and solve the eigenvalue problems $\mathbb{M} \mathbf{z}_{l}=\lambda_{l} \mathbf{z}_{l}$, for $1 \leq l \leq L$. Then, by (3.4), set $\mathbf{w}_{l}=\left(\lambda_{l}\right)^{-1 / 2} \mathbb{S}_{l}$, for $1 \leq l \leq L$.

### 3.2. Implementation: details and issues

The numerical procedure presented so far can be efficiently carried out within an offline-online framework [39]. The parameter-independent offline step consists of the generation of the snapshots through a high-fidelity, expensive scheme and the subsequent construction of the reduced basis via POD. To determine an appropriate dimension for the basis, which ensures a desired degree of accuracy, one can resort to empirical criteria, like, e.g., the relative information content [41]. Then, given a new parameter value $\boldsymbol{\mu} \in \mathscr{P}$, the nonlinear reduced system (3.3) is solved online.

However, to enjoy a significant reduction in the computational burden with respect to traditional (full-order) discretization techniques, the complexity of any online query should be independent of the original size of the problem. Yet, due to the nonlinearity of the underlying PDE and the non-affinity in the parameter dependence (partially induced by the transformation map $\boldsymbol{\Phi}\left(\cdot ; \boldsymbol{\mu}_{g}\right)$ ), the assembly of the reduced problems has to be embodied directly in the online stage, thus seriously compromising the efficiency of the overall procedure [3]. Without escaping the algebraic framework, this can be overcome by resorting to suitable techniques such as the discrete empirical interpolation method (DEIM) [10] or its matrix variant (MDEIM) [37], aiming at recovering an affine dependency on the parameter $\boldsymbol{\mu}$. However, the implementation of such techniques is problem-dependent and of an intrusive nature, as it requires one to modify the assembly routines of the corresponding computational code [9]. Moreover, any interpolation procedure unavoidably introduces a further level of approximation. As a matter of fact, typically one needs to generate a larger number of snapshots in the offline stage and then retain a larger number of POD modes to guarantee the same accuracy provided by the standard POD-Galerkin method [3].

## 4. Artifical neural networks

Inspired by the biological information processing system (see, e.g., [21, 28]), an artificial neural network (ANN), often referred to as neural network, is a computational model able to learn from observational data, i.e., by example, thus providing an alternative to the algorithmic programming paradigm [36]. As its original counterpart, it consists of a collection of processing units, called (artificial) neurons, and a set of directed weighted synaptic connections among the neurons. Data travel among neurons through the connections, following the direction imposed by the synapses. Hence, an artificial neural network is an oriented graph, with the neurons as nodes and the synapses as oriented edges, whose weights are adjusted by means of a training process to configure the network for a specific application [45].

In the following, we discuss the structure and training of a neural network, starting by detailing the working principles of an artificial neuron.

### 4.1. Neuronal model

An artificial neuron represents a simplified model of a biological neuron [28]. To introduce the components of the model, let us consider the neuron $j$ represented on the left in Fig. 4.1. Suppose that it is connected with $m$ sending neurons $\left\{s_{1}, \ldots, s_{m}\right\}$, and $n$ receiving (target) neurons $\left\{r_{1}, \ldots, r_{n}\right\}$. Denoting by $y_{\alpha}(t) \in \mathbb{R}$ the (scalar) output of a generic neuron $\alpha$ at time $t$ and by $w_{\alpha, \beta}$ the weight of the connection $\left(\alpha, \beta\right.$ ), neuron $j$ gets the weighted inputs $w_{s_{k}, j} y_{s_{k}}(t), k=1, \ldots, m$, at time $t$, and sends out the output $y_{j}(t+\Delta t)$ to the target neurons $\left\{r_{1}, \ldots, r_{n}\right\}$ at time $t+\Delta t$. In particular, neuron $r_{i}, i=1, \ldots, n$, receives as input $w_{j, r_{i}} y_{j}(t+\Delta t)$. Note that in the context of ANNs, time is discretized by introducing the timestep $\Delta t$. This is clearly not plausible from a biological viewpoint; however, it substantially simplifies the implementation. In the following, we will avoid to specify the dependence on time unless strictly necessary, to lighten the notation.

An artificial neuron $j$ is completely characterized by three functions: the propagation function, the activation function and the output function. The propagation function $f_{\text {prop }}$ converts the vectorial input $\mathbf{p}=\left[y_{s_{1}}, \ldots, y_{s_{m}}\right]^{T} \in \mathbb{R}^{m}$ into a scalar $u_{j}$ often called net input, i.e.,

$$
u_{j}=f_{\text {prop }}\left(w_{s_{1}, j}, \ldots, w_{s_{m}, j}, y_{s_{1}}, \ldots, y_{s_{m}}\right)
$$

A common choice for $f_{\text {prop }}$ (used also in this work) is the weighted sum, adding up the scalar inputs multiplied by their respective weights:

$$
f_{\text {prop }}\left(w_{s_{1}, j}, \ldots, w_{s_{m}, j}, y_{s_{1}}, \ldots, y_{s_{m}}\right)=\sum_{k=1}^{m} w_{s_{k}, j} y_{s_{k}} .
$$




Figure 4.1. Visualization of the generic $j$-th neuron of an artificial neural network, including (right) or not (left) a bias neuron. On the left, the neuron accumulates the weighted inputs $\left\{w_{s_{1}, j} y_{s_{1}}, \ldots, w_{s_{m}, j} y_{s_{m}}\right\}$ respectively coming from the sending neurons $\left\{s_{1}, \ldots, s_{m}\right\}$; on the right, the neuron accumulates the weighted inputs $\left\{w_{s_{1}, j} y_{s_{1}}, \ldots, w_{s_{m}, j} y_{s_{m}},-\theta_{j}\right\}$ respectively coming from the sending neurons $\left\{s_{1}, \ldots, s_{m}, b\right\}$, with $b$ the bias neuron. In both situations, the neuron then fires $y_{j}$, sent to the target neurons $\left\{r_{1}, \ldots, r_{n}\right\}$ through the synapsis $\left\{w_{j, r_{1}}, \ldots, w_{j, r_{n}}\right\}$. The neuron threshold is reported in brackets within its body.

At each timestep, the activation state $a_{j}$, often referred to as activation, quantifies to which extent neuron $j$ is currently active or excited. It results from the activation function $f_{\text {act }}$, which combines the net input $u_{j}$ with a threshold $\theta_{j} \in \mathbb{R}$ [28]:

$$
a_{j}=f_{a c t}\left(u_{j} ; \theta_{j}\right)=f_{a c t}\left(\sum_{k=1}^{m} w_{s_{k}, j} y_{s_{k}} ; \theta_{j}\right)
$$

Note that the threshold $\theta_{j}$ is a parameter of the network and as such one may choose to adapt it through a training process, exactly as it can be done for the synaptic weights. To ease the runtime access of $\theta_{j}$, it is common practice to introduce a bias neuron in the network. A bias neuron is a continuously firing neuron, with constant output $y_{b}=1$, which is directly connected with neuron $j$, assigning the bias weight $w_{b, j}=-\theta_{j}$ to the connection. As can be deduced by the representation on the right in Fig. 4.1, $\theta_{j}$ is now treated as a synaptic weight, while the neuron threshold is set to zero. Therefore, the net input and the activation state can respectively be expressed as

$$
u_{j}=\sum_{k=1}^{m} w_{s_{k}, j} y_{s_{k}}-\theta_{j} \quad \text { and } \quad a_{j}=f_{a c t}\left(\sum_{k=1}^{m} w_{s_{k}, j} y_{s_{k}}-\theta_{j}\right) .
$$

There exist various choices for the activation function. The sigmoid activation functions have been widely used for the realization of artificial neural networks due to their graceful combination of linear and nonlinear behaviour [21]. Sigmoid functions are s-shaped, monotically increasing, and assume values in a bounded interval; a well-known instance is given by the hyperbolic tangent,

$$
f_{a c t}(\nu)=\frac{e^{v}-e^{-v}}{e^{v}+e^{-v}} .
$$

Finally, the output function $f_{\text {out }}$ calculates the scalar output $y_{j} \in \mathbb{R}$ based on the activation state $a_{j}$ of the neuron:

$$
y_{j}=f_{\text {out }}\left(a_{j}\right) .
$$

Often, $f_{\text {out }}$ is the identity function, so that activation and output of a neuron coincides, i.e., $y_{j}=f_{\text {out }}\left(a_{j}\right)=a_{j}$. The output $y_{j}$ could then be sent either to other neurons or constitute a component of the overall output vector of the network, as for the neurons in the output layer of a feedforward neural network, illustrated in the following subsection.

It should be pointed out that the neural model presented so far refers to the so called computing neuron, i.e., a neuron which processes input information to provide a response. However, in a neural network one may also identify source neurons, supplying the network with the respective components of the activation pattern (input vector), without performing any computation [21].


Figure 4.2. A three-layer feedforward neural network, with three input neurons, two hidden layers each one consisting of six neurons, and four output neurons. Within each connection, information flows from left to right.

### 4.2. Network topology: the feedforward neural network

The interconnection of neurons within a network defines the topology of the network itself, i.e., its design. In the literature, many network architectures have been proposed, sometimes tailored to a specific application. Among all, feedforward neural networks, also called perceptrons [43], have been preferred in function regression tasks.

In a feedforward neural network, neurons are arranged into layers, with one input layer of $M_{I}$ source neurons, $K$ hidden layers, each one consisting of $H_{k}$ computing neurons, $k=1, \ldots, K$, and an output layer of $M_{O}$ computing neurons. As a characteristic property, neurons in a layer can only be connected with neurons in the next layer towards the output layer. Then, an activation pattern $\mathbf{p} \in \mathbb{R}^{M_{I}}$, supplied to the network through the source nodes in the first layer, provides the input signal for the neurons in the first hidden layer. For each hidden layer, its output signals give the input pattern for the following layer. In this way, information travels towards the last layer of the network, i.e., the output layer, whose outputs constitute the components of the overall output $\mathbf{q} \in \mathbb{R}^{M_{O}}$ of the network ${ }^{1}$. Hence, a feedforward network establishes a map between the input space $\mathbb{R}^{M_{I}}$ and the output space $\mathbb{R}^{M_{O}}$. This does make this network architecture particularly suitable for continuous function approximation.

Feedforward networks can be classified according to the number of hidden layers or, equivalently, the number of layers of trainable weights. Single-layer perceptrons (SLPs) consist of the input and output layer, without any hidden layer. Because of their simple structure, the range of application of SLPs is rather limited. Indeed, only linearly separable data can be properly represented using SLPs [28]. Conversely, multi-layer perceptrons (MLPs), with at least one hidden layer, are universal function approximators, as stated by Cybenko [12, 13]. In detail:
(i) MLPs with one hidden layer and differentiable activation functions can approximate any continuous function;
(ii) MLPs with two hidden layers and differentiable activation functions can approximate any function.

Therefore, in many practical applications there is no reason to employ MLPs with more than two hidden layers. However, (i) and (ii) do not give any practical advice neither on the number of hidden neurons nor the number of samples required to train the network: these should be found pursuing a trial-and-error (and likely time-consuming) approach [20].

An instance of a three-layer (i.e., two hidden layer plus the output layer) feedforward network is offered in Fig. 4.2. In this case, we have $M_{I}=3$ input neurons (denoted with the letter $i$ ), $H_{1}=H_{2}=6$ hidden neurons (letter $h$ for both hidden layers), and $M_{O}=4$ output neurons (letter $o$ ). In particular, it represents an instance of a completely linked perceptron, since each neuron is directly connected with all neurons in the following layer.

[^1]
### 4.3. Training a multi-layer feedforward neural network

As previously mentioned, the principal characteristic of a neural network is its capability of learning from the surrounding environment, storing the acquired knowledge via its internal parameters, i.e., the synaptic and bias weights. Learning is accomplished through a training process, during which the network is exposed to a collection of examples, called training patterns. According to some performance measure, the weights are then adjusted by means of a well-defined set of rules. Therefore, the learning procedure is an algorithm, typically iterative, such that after a successfull training, the neural network provides reasonable responses for unknown problems of the same class of the training set. This property is known as generalization [28].

Training algorithms can be classified based on the nature of the training set, i.e., the set of training patterns. We can then distinguish three learning paradigms: supervised learning, unsupervised learning and reinforcement learning [20]. The choice of the learning paradigm is clearly task-dependent. Particularly, for function approximation, the supervised learning paradigm is the natural choice.

Consider the nonlinear unknown function $\boldsymbol{f}: \mathbb{R}^{M_{I}} \rightarrow \mathbb{R}^{M_{O}}$ and a set of labeled examples $\left\{\mathbf{p}_{i}, \mathbf{t}_{i}=\boldsymbol{f}\left(\mathbf{p}_{i}\right)\right\}_{1 \leq i \leq N_{t r}}$, which form the training set. Note that to each input pattern $\mathbf{p}_{i} \in \mathbb{R}^{M_{I}}, i=1, \ldots, N_{t r}$, corresponds an associated desired output or teaching input $\mathbf{t}_{i} \in \mathbb{R}^{M_{O}}$. The goal is to approximate $\boldsymbol{f}$ over a domain $D \subset \mathbb{R}^{M_{I}}$ up to a user-defined tolerance $\epsilon$, i.e.,

$$
\|\boldsymbol{F}(\mathbf{x})-\boldsymbol{f}(\mathbf{x})\|<\epsilon \quad \forall \mathbf{x} \in D,
$$

where $\boldsymbol{F}: \mathbb{R}^{M_{I}} \rightarrow \mathbb{R}^{M_{O}}$ is the actual input-output map established by the neural network and $\|\cdot\|$ is some suitable norm on $\mathbb{R}^{M_{O}}$. For this purpose, consider the synapsis between a sending neuron $i$ and a target neuron $j$. At the $t$-th iteration (also called epoch) of the training procedure, the weight $w_{i, j}(t)$ of the connection $(i, j)$ is modified by the time-dependent quantity $\Delta w_{i, j}(t)$, whose form depends on the specific learning rule. Hence, at the subsequent iteration $t+1$ the synaptic weight is simply given by

$$
w_{i, j}(t+1)=w_{i, j}(t)+\Delta w_{i, j}(t)
$$

The whole training process is driven by an error or performance function $E$, which measures the discrepancy between the neural network knowledge of the surrounding environment and the actual state of the environment itself. Therefore, every learning rule aims to minimize the performance $E$, thought as a scalar function of the free parameters of the network, namely

$$
E=E(\mathbf{w}) \in \mathbb{R},
$$

with $\mathbf{w} \in \mathbb{R}^{W}$ being the vector collecting all the synaptic and bias weights. Thus, the point over the error surface reached at the end of a successful training process provides the optimal configuration $\mathbf{w}_{\text {opt }}$ for the network. A common choice for the performance function is the (accumulated) mean squared error (MSE)

$$
\begin{equation*}
E(\mathbf{w})=\sum_{\mathbf{p} \in P} E_{\mathbf{p}}(\mathbf{w})=\sum_{\mathbf{p} \in P} \frac{1}{M_{O}} \sum_{j=1}^{M_{O}}\left(t_{\mathbf{p}, j}-q_{\mathbf{p}, j}\right)^{2}, \tag{4.1}
\end{equation*}
$$

where for each input pattern $\mathbf{p}$ belonging to the training set $P, \mathbf{t}_{\mathbf{p}}$ and $\mathbf{q}_{\mathbf{p}}$ denote the corresponding teaching input and actual output, respectively. Observe that (4.1) accounts for the error committed on each input pattern in the training set, leading to an offline learning procedure.

In our numerical tests, the weights are iteratively adjusted via the Levenberg-Marquardt algorithm [19, 32], whose cost increases nonlinearly with the number of neurons in the network, making it poorly efficient for large networks. However, it turns out to be efficient and accurate for networks with few hundreads of connections.

## 5. A non-intrusive reduced basis method using artificial neural networks

The scenario portrayed so far motivates the research for an alternative approach to tackle any online query within the reduced basis framework. To this end, let us remark that there exists a one-to-one correspondence between the reduced space $V_{\mathrm{rb}}$ and the column space $\operatorname{Col}(\mathbb{V})$ of $\mathbb{V}$. Indeed, letting $\left\{\phi_{1}, \ldots, \phi_{N_{h}}\right\}$ be a basis for $V_{h}$ and $\left\{\psi_{1}, \ldots, \psi_{L}\right\}$ be the reduced
basis, from Eq. (3.2) follows:

$$
V_{\mathrm{rb}} \ni v_{L}=\sum_{j=1}^{L} v_{\mathrm{rb}}^{(j)} \psi_{j}=\sum_{j=1}^{L} v_{\mathrm{rb}}^{(j)} \sum_{i=1}^{N_{h}} \mathbb{V}_{i, j} \phi_{i}=\sum_{i=1}^{N_{h}}\left(\mathbb{V} \mathbf{v}_{\mathrm{rb}}\right)_{i} \phi_{i} \quad \leftrightarrow \quad \mathbf{v}_{L} \in \operatorname{Col}(\mathbb{V}) .
$$

In particular, this implies that the projection of any $\nu_{h} \in V_{h}$ onto $V_{\mathrm{rb}}$ in the discrete scalar product $(\cdot, \cdot)_{h}$,

$$
\begin{equation*}
\left(\chi_{h}, \xi_{h}\right)_{h}=\sum_{i=1}^{N_{h}} \chi_{h}^{(i)} \xi_{h}^{(i)}=\left(\chi_{h}, \xi_{h}\right)_{\mathbb{R}^{N_{h}}} \quad \forall \chi_{h}, \xi_{h} \in V_{h}, \tag{5.1}
\end{equation*}
$$

algebraically corresponds to the projection $\mathbf{v}_{h}^{\mathbb{V}}$ of $\mathbf{v}_{h}$ onto $\operatorname{Col}(\mathbb{V})$ in the Euclidean scalar product, given by

$$
\mathbf{v}_{h}^{\mathbb{V}}=\mathbb{P} \mathbf{v}_{h} \quad \text { with } \quad \mathbb{P}=\mathbb{V} \mathbb{V}^{T} \in \mathbb{R}^{N_{h} \times N_{h}} .
$$

Note that $\mathbf{v}_{h}^{\mathbb{V}}$ is the element of $\operatorname{Col}(\mathbb{V})$ closest to $\mathbf{v}_{h}$ in the Euclidean norm, i.e.,

$$
\left\|\mathbf{v}_{h}-\mathbf{v}_{h}^{\mathbb{V}}\right\|_{\mathbb{R}^{N_{h}}}=\inf _{\mathbf{w}_{h} \in \operatorname{Col}(\mathbb{V})}\left\|\mathbf{v}_{h}-\mathbf{w}_{h}\right\|_{\mathbb{R}^{N_{h}}} .
$$

Therefore, the element of $V_{\mathrm{rb}}$ closest to the high-fidelity solution $u_{h}$ in the discrete norm $\|\cdot\|_{h}=\sqrt{(\cdot, \cdot)_{h}}$ can be expressed as

$$
u_{h}^{\mathbb{V}}(\boldsymbol{x} ; \boldsymbol{\mu})=\sum_{j=1}^{N_{h}}\left(\mathbb{V}^{T} \mathbf{u}_{h}(\boldsymbol{\mu})\right)_{j} \phi_{j}(\boldsymbol{x})=\sum_{i=1}^{L}\left(\mathbb{V}^{T} \mathbf{u}_{h}(\boldsymbol{\mu})\right)_{i} \psi_{i}(\boldsymbol{x}) .
$$

Motivated by this last equality, once a reduced basis has been constructed (e.g., via POD of the snapshot matrix), we aim at approximating the function

$$
\begin{align*}
\boldsymbol{\pi}: \mathscr{P} \subset \mathbb{R}^{P} & \rightarrow \mathbb{R}^{L} \\
\boldsymbol{\mu} & \mapsto \mathbb{V}^{T} \mathbf{u}_{h}(\boldsymbol{\mu}), \tag{5.2}
\end{align*}
$$

which maps each input vector parameter $\boldsymbol{\mu} \in \mathscr{P}$ to the coefficients $\mathbb{V}^{T} \mathbf{u}_{h}(\boldsymbol{\mu})$ for the expansion of $u_{h}^{V}$ in the reduced basis $\left\{\psi_{i}\right\}_{1 \leq i \leq L}$. Then, given a new parameter instance $\boldsymbol{\mu}$, the associated RB solution is simply given by the evaluation at $\boldsymbol{\mu}$ of the so-built approximation $\hat{\boldsymbol{\pi}}$ of $\boldsymbol{\pi}$, i.e.

$$
\mathbf{u}_{\mathrm{rb}}(\boldsymbol{\mu})=\hat{\boldsymbol{\pi}}(\boldsymbol{\mu}),
$$

and, consequently,

$$
\mathbf{u}_{L}(\boldsymbol{\mu})=\mathbb{V} \hat{\boldsymbol{\pi}}(\boldsymbol{\mu}) .
$$

Note that, provided that the construction of $\hat{\boldsymbol{\pi}}$ is carried out within the offline stage, this approach leads to a non-intrusive RB method, enabling a complete decoupling between the online step and the underlying full-order model. Moreover, the accuracy of the resulting reduced solution uniquely relies on the quality of the reduced basis and the effectiveness of the approximation $\hat{\boldsymbol{\pi}}$ of the map $\boldsymbol{\pi}$, which we assume sufficiently smooth.

In the literature, different approaches for the interpolation of (5.2) have been developed, e.g., exploiting some geometrical considerations concerning the discrete solution manifold $\mathscr{\mu}_{h}$ [1], or employing radial basis functions [11]. In this work we resort to artificial neural networks for the nonlinear regression of the map $\boldsymbol{\pi}$, leading to the POD-NN RB method. As described in Subsection 4.3, any neural network is tailored to the particular application at hand by means of a preliminary training phase. Here, we are concerned with a function regression task, thus we straightforwardly adopt a supervised learning paradigm, training the perceptron via exposition to a collection of (known) input-output pairs

$$
P_{t r}=\left\{\left(\boldsymbol{\mu}^{(i)}, \mathbb{V}^{T} \mathbf{u}_{h}\left(\boldsymbol{\mu}^{(i)}\right)\right)\right\}_{1 \leq i \leq N_{t r}} .
$$

According to the notation and nomenclature previously introduced, for any $i=1, \ldots, N_{t r}, \mathbf{p}_{i}=\boldsymbol{\mu}^{(i)} \in \mathbb{R}^{P}$ represents the input pattern and $\mathbf{t}_{i}=\mathbb{V}^{T} \mathbf{u}_{h}\left(\boldsymbol{\mu}^{(i)}\right) \in \mathbb{R}^{L}$ the associated teaching input; together, they constitute a training pattern. In this respect, note that the teaching inputs $\mathbb{V}^{T} \mathbf{u}_{h}\left(\boldsymbol{\mu}^{(i)}\right), i=1, \ldots, N_{t r}$, are generated through the FE solver. On the one hand, this ensures the reliability of the teaching patterns, given the assumed high-fidelity of the FE scheme. On the other hand,

```
Algorithm 5.1 The offline and online stages for the POD-NN RB method.
    function \(\left[\mathbb{V}, \mathbf{w}_{\mathrm{rb}}\right]=\) PODNN_OFFLINE \(\left(\mathscr{P}, \Omega_{h}, N\right)\)
        generate the parameter set \(\Xi_{N}=\left\{\boldsymbol{\mu}^{(1)}, \ldots, \boldsymbol{\mu}^{(N)}\right\}\)
        compute the high-fidelity solutions \(\left\{\mathbf{u}_{h}\left(\boldsymbol{\mu}^{(1)}\right), \ldots, \mathbf{u}_{h}\left(\boldsymbol{\mu}^{(N)}\right)\right\}\) via FE-Newton's method
        generate the POD basis functions \(\left\{\mathbf{w}_{1}, \ldots, \mathbf{w}_{L}\right\}\) via method of snapshots
        assemble \(\mathbb{V}=\left[\mathbf{w}_{1}|\ldots| \mathbf{w}_{L}\right]\)
        find an optimal network configuration \(\mathbf{w}_{\mathrm{rb}}\) relying upon LHS and Levenberg-Marquardt algorithm
    end function
    function \(\mathbf{u}_{L}^{\mathbb{N N}}(\boldsymbol{\mu})=\) PODNN_ONLINE \(\left(\boldsymbol{\mu}, \mathbb{V}, \mathbf{w}_{\mathrm{rb}}\right)\)
        evaluate the output \(\mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})\) of the network for the input vector \(\boldsymbol{\mu}\)
        compute \(\mathbf{u}_{L}^{\mathbb{N N}}(\boldsymbol{\mu})=\mathbb{V} \mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})\)
    end function
```

this also suggests to incorporate the learning phase of the perceptron in the offline step of the POD-NN RB method, as described in Algorithm 5.1. In doing so, we exploit the natural decoupling between the training and the evaluation of neural networks, thus fulfilling the necessary requirement to ensure online efficiency.

It should be pointed out that the design of an effective learning procedure may require a larger amount of snapshots than the generation of the reduced space does. Moreover, we mentioned in Subsection 4.2 the time-consuming yet unavoidable trial-and-error approach which one should pursue in the search for an optimal network topology. To this end, we propose an automatic procedure.

While Cybenko's theorems (see Subsection 4.2) allow one to consider perceptrons with no more than two hidden layers, no similar a priori and general results are available for the number $H$ of neurons per layer ${ }^{2}$. Hence, given an initial amount $N_{t r}$ of training samples (say $N_{t r}=100$ ), we train the network for increasing values of $H$, stopping when overfitting of training data occurs, due to an excessive number of neurons with respect to the number of training patterns. In case the best configuration, i.e., the one yielding the smallest error on a test data set, does not meet a desired level of accuracy, we generate a new set of snapshots, which will enlarge the current training set, and then proceed to re-train the network in different configurations. At this point, it is worth pointing out two issues.
(i) Once the training set has been enriched, we can limit ourselves to network configurations including a number of neurons no smaller than the current optimal network. Indeed, the error on the test data set is decreasing in the number of patterns the neural network is exposed to during the learning phase, and the larger the number of neurons, the faster the decay.
(ii) In order to maximize the additional quantity of information available for the learning, we should ensure that the new training inputs, i.e., parameter values, do not overlap with the ones already present in the training parameter set. To this aim, we pursue an heuristhic approach, employing, at each iteration, the latin hypercube sampling [25], which provides a good compromise between randomness and even coverage of the parameter domain.

The procedure is then iterated until a satisfactory degree of accuracy and generalization is attained, or the available resources (i.e., computing power, running time and memory space) are exhausted. Therefore, the speed-up enabled at the online stage comes at the cost of an extended offline phase.

In our numerical tests, we rely on the mean squared error (4.1) as performance function driving the neural network training process. To motivate this choice, let

$$
\mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu}) \in \mathbb{R}^{L}
$$

be the (actual) output provided by the network for a given input $\boldsymbol{\mu}$, and

$$
\mathbf{u}_{L}^{\mathrm{NN}}(\boldsymbol{\mu})=\mathbb{V} \mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu}) \in \operatorname{Col}(\mathbb{V}) \subset \mathbb{R}^{N_{h}} .
$$

[^2]Then (omitting the dependence on the input vector for ease of notation):

$$
\operatorname{MSE}\left(\mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}, \mathbb{V}^{T} \mathbf{u}_{h}\right) \propto\left\|\mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}-\mathbb{V}^{T} \mathbf{u}_{h}\right\|_{\mathbb{R}^{L}}^{2}=\left\|\mathbf{u}_{L}^{\mathrm{NN}}-\mathbb{V} \mathbb{V}^{T} \mathbf{u}_{h}\right\|_{\mathbb{R}^{N_{h}}}^{2}=\left\|u_{L}^{\mathrm{NN}}-u_{h}^{\mathbb{V}}\right\|_{h}^{2},
$$

where

$$
u_{L}^{\mathrm{NN}}(\boldsymbol{x} ; \boldsymbol{\mu})=\sum_{i=1}^{L}\left(\mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})\right)_{i} \psi_{i}(\boldsymbol{x}) \in V_{\mathrm{rb}} .
$$

Therefore, for any training input $\boldsymbol{\mu}^{(i)}, i=1, \ldots, N_{t r}$, by minimizing the MSE we minimize the distance (in the discrete norm $\|\cdot\|_{h}$ ) between the approximation provided by the neural network and the projection of the FE solution onto the reduced space $V_{\mathrm{rb}}$. The proper generalization to other parameter instances, not included in the training set, is then ensured by the implementation of suitable techniques (e.g., early stopping [33], generalized cross validation [27]) aiming at preventing the network to overfit the training data.

## 6. Numerical results

In this section, we present the numerical results obtained via the FE, POD-G and POD-NN methods applied to parametrized full-Dirichlet boundary value problems (BVPs) for the one-dimensional and two-dimensional nonlinear Poisson equation and the two-dimensional incompressible steady Navier-Stokes equations. In the one-dimensional case we deal uniquely with physical parameters, whereas in two spatial dimensions we consider purely geometric parametrizations.

The two RB techniques considered in this work are compared both in terms of accuracy and performance during the online stage. For this, let $\|\cdot\|$ be the canonical (Euclidean) norm on $\mathbb{R}^{N_{h}}$. In the online phases of the POD-G and POD-NN methods, for a new parameter value $\boldsymbol{\mu} \in \mathscr{P}$ (not involved either in the generation of the POD basis nor in the identification of an optimal neural network), the following relative errors with respect to the high-fidelity solution $\mathbf{u}_{h}(\boldsymbol{\mu})$ are analyzed:
(a) the POD-G relative error,

$$
\begin{equation*}
\varepsilon_{\mathrm{PODG}}(L, \boldsymbol{\mu})=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbf{u}_{L}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|}=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbb{V} \mathbf{u}_{\mathrm{rb}}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|} ; \tag{6.1}
\end{equation*}
$$

(b) the POD-NN relative error,

$$
\begin{equation*}
\varepsilon_{\mathrm{PODNN}}(L, \boldsymbol{\mu})=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbf{u}_{L}^{\mathrm{NN}}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|}=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbb{V} \mathbf{u}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|} ; \tag{6.2}
\end{equation*}
$$

(c) the relative projection error, i.e., the error committed by approximating the high-fidelity solution with its projection (in the discrete scalar product $(\cdot, \cdot)_{h}$, see (5.1)) onto the reduced space $V_{\mathrm{rb}}$,

$$
\begin{equation*}
\varepsilon_{\mathbb{V}}(L, \boldsymbol{\mu})=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbf{u}_{h}^{\mathbb{V}}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|}=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbb{V V}^{T} \mathbf{u}_{h}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|} . \tag{6.3}
\end{equation*}
$$

Clearly, (6.3) provides a lower-bound for both (6.1) and (6.2).
The above errors are evaluated on a test parameter set $\Xi_{t e} \subset \mathscr{P}$ consisting of $N_{t e}$ randomly picked samples. Given that a sufficiently large number of test values is chosen, statistics for $\varepsilon_{\text {PODG }}(L, \cdot), \varepsilon_{\text {PODNN }}(L, \cdot)$ and $\varepsilon_{\mathbb{V}}(L, \cdot)$ on $\Xi_{t e}$ can be reasonably assumed independent of the particular choice made for $\Xi_{t e}$, thus making any subsequent error analysis reliable. In particular, in our numerical studies we let $N_{t e}$ range from 50 up to 100, and we consider the average of the errors over the test data set, respectively denoted by

$$
\bar{\varepsilon}_{\mathrm{PODG}}=\bar{\varepsilon}_{\mathrm{PODG}}(L)=\frac{\sum_{\boldsymbol{\mu} \in \Xi_{t e} \varepsilon_{\mathrm{PODG}}(L, \boldsymbol{\mu})}}{N_{t e}}, \quad \bar{\varepsilon}_{\mathrm{PODNN}}=\bar{\varepsilon}_{\mathrm{PODNN}}(L)=\frac{\sum_{\boldsymbol{\mu} \in \Xi_{t e}} \varepsilon_{\mathrm{PODNN}}(L, \boldsymbol{\mu})}{N_{t e}}, \quad \bar{\varepsilon}_{\mathbb{V}}=\bar{\varepsilon}_{\mathbb{V}}(L)=\frac{\sum_{\boldsymbol{\mu} \in \Xi_{t e}} \varepsilon_{\mathbb{V}}(L, \boldsymbol{\mu})}{N_{t e}} .
$$

In the training phase of neural networks, we rely upon the well-known cross-validation procedure [27]. While the training samples $\Xi_{t r} \subset \mathscr{P}$ are generated via successive latin hypercube samplings (as explained in Section 5), the validation inputs $\Xi_{v a} \subset \mathscr{P}$ are randomly picked through a Monte Carlo sampling, as done for $\Xi_{t e}$. The ratio between the size of $\Xi_{v a}$ and $\Xi_{t r}$
is set to 0.3 . The training is then performed using the multiple restarts approach to prevent the results from depending on the way the weights are (randomly) initialized; see, e.g., $[28,33]$. Typically, five restarts are performed for each network topology. To ease the research for an optimal network configuration and by exploiting Cybenko's results (see Subsection 4.2), we limit ourselves to three-layers neural networks, with the same number of neurons for both hidden layers. For each topology, the hyperbolic tangent is used as activation function for the hidden neurons.

All the results presented in the following subsections have been obtained via a MATLAB ${ }^{\circledR}$ implementation of the FE, POD-G and POD-NN methods, with all the simulations carried out on a laptop equipped with a CPU Intel ${ }^{\circledR}$ Core $^{\mathrm{TM}}$ i7 @ 2.50 GHz .

### 6.1. Nonlinear Poisson equation

Despite a rather simple form, the Poisson equation has proved itself to be effective to model steady phenomena occurring in, e.g., electromagnetism, heat transfer and underground flows [34]. We consider the following version of the parametrized Poisson equation for a state variable $\widetilde{u}=\widetilde{u}(\boldsymbol{\mu})$ :

$$
\left\{\begin{array}{rlrl}
-\widetilde{\nabla} \cdot\left(\widetilde{k}\left(\widetilde{\boldsymbol{x}}, \widetilde{u}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right) \widetilde{\nabla} \widetilde{u}(\boldsymbol{\mu})\right) & =\widetilde{s}\left(\widetilde{\boldsymbol{x}} ; \boldsymbol{\mu}_{p h}\right) & \text { in } \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)  \tag{6.4a}\\
\widetilde{u}(\boldsymbol{\mu}) & =\widetilde{h}\left(\widetilde{\boldsymbol{\sigma}} ; \boldsymbol{\mu}_{p h}\right) & & \text { on } \widetilde{\Gamma}_{D} \\
\widetilde{k}\left(\widetilde{\boldsymbol{\sigma}}, \widetilde{u}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right) \widetilde{\nabla} \widetilde{u}(\boldsymbol{\mu}) \cdot \widetilde{\boldsymbol{n}} & =0 & & \text { on } \widetilde{\Gamma}_{N}
\end{array}\right.
$$

Here, for any $\boldsymbol{\mu}_{g} \in \mathscr{P}_{g}, \widetilde{\boldsymbol{x}}$ and $\widetilde{\boldsymbol{\sigma}}$ denote a generic point in $\widetilde{\Omega}$ and on $\widetilde{\Gamma}$, respectively, $\widetilde{\nabla}$ is the nabla operator with respect to $\widetilde{\boldsymbol{x}}$, $\tilde{\boldsymbol{n}}=\widetilde{\boldsymbol{n}}(\widetilde{\boldsymbol{\sigma}})$ denotes the outward normal to $\widetilde{\Gamma}$ in $\widetilde{\boldsymbol{\sigma}}, \widetilde{k}: \widetilde{\Omega} \times \mathbb{R} \times \mathscr{P}_{p h} \rightarrow(0, \infty)$ is the diffusion coefficient, $\widetilde{s}: \widetilde{\Omega} \times \mathscr{P}_{p h} \rightarrow \mathbb{R}$ is the source term, and $\widetilde{h}: \widetilde{\Gamma}_{D} \times \mathscr{P}_{p h} \rightarrow \mathbb{R}$ encodes the Dirichlet boundary conditions. To ease the subsequent discussion, we limit the attention to homogeneous Neumann boundary constraints.

Let us fix $\boldsymbol{\mu} \in \mathscr{P}$ and set

$$
\widetilde{V}=\widetilde{V}\left(\boldsymbol{\mu}_{g}\right)=H_{\widetilde{\Gamma}_{D}}^{1}\left(\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)\right)=\left\{v \in H^{1}\left(\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)\right):\left.v\right|_{\widetilde{\Gamma}_{D}}=0\right\}
$$

Multiplying (6.4a) by a test function $v \in \widetilde{V}$, integrating over $\widetilde{\Omega}$, and exploiting integration by parts on the left-hand side, yields:

$$
\begin{equation*}
\int_{\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)} \widetilde{k}\left(\widetilde{u}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right) \widetilde{\nabla} \widetilde{u}(\boldsymbol{\mu}) \cdot \widetilde{\nabla} v d \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)=\int_{\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)} \widetilde{s}\left(\boldsymbol{\mu}_{p h}\right) v d \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right) \tag{6.5}
\end{equation*}
$$

where we have omitted the dependence on the space variable $\widetilde{\boldsymbol{x}}$ for ease of notation. For the integrals in Eq. (6.5) to be well-defined, we require, for any $\boldsymbol{\mu}_{g} \in \mathscr{P}_{g}$,

$$
\left|\widetilde{k}\left(\widetilde{\boldsymbol{x}}, r ; \boldsymbol{\mu}_{g}\right)\right|<\infty \text { for almost any (a.a.) } \widetilde{\boldsymbol{x}} \in \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right), r \in \mathbb{R}, \text { and } \widetilde{\boldsymbol{s}}\left(\boldsymbol{\mu}_{p h}\right) \in L^{2}\left(\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)\right)
$$

Let then $\widetilde{l}=\widetilde{l}(\boldsymbol{\mu}) \in H^{1}\left(\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)\right)$ be a lifting function such that $\left.\widetilde{l}(\boldsymbol{\mu})\right|_{\widetilde{\Gamma}_{D}}=\widetilde{h}\left(\boldsymbol{\mu}_{p h}\right)$, with $\widetilde{h}\left(\boldsymbol{\mu}_{p h}\right) \in H^{1 / 2}\left(\widetilde{\Gamma}_{D}\right)$. We assume that such a function can be constructed, e.g., by interpolation of the boundary condition. Hence, upon defining

$$
\begin{array}{rlr}
\widetilde{a}(w, v ; \boldsymbol{\mu}) & :=\int_{\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)} \widetilde{k}\left(w+\widetilde{l}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right) \widetilde{\nabla} w \cdot \widetilde{\nabla} v d \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)+\int_{\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)} \widetilde{k}\left(w+\widetilde{l}(\boldsymbol{\mu}) ; \boldsymbol{\mu}_{p h}\right) \widetilde{\nabla} \widetilde{l}(\boldsymbol{\mu}) \cdot \widetilde{\nabla} v d \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right) & \forall w, v \in \widetilde{V},  \tag{6.6}\\
\widetilde{f}(v ; \boldsymbol{\mu}) & :=\int_{\widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right)} \widetilde{s}\left(\boldsymbol{\mu}_{p h}\right) v d \widetilde{\Omega}\left(\boldsymbol{\mu}_{g}\right) & \forall v \in \widetilde{V},
\end{array}
$$

the weak formulation of problem (6.4) reads: given $\boldsymbol{\mu} \in \mathscr{P}$, find $\widetilde{u}(\boldsymbol{\mu}) \in \widetilde{V}\left(\boldsymbol{\mu}_{g}\right)$ such that

$$
\begin{equation*}
\widetilde{a}(\widetilde{u}(\boldsymbol{\mu}), v ; \boldsymbol{\mu})=\widetilde{f}(v ; \boldsymbol{\mu}) \quad \forall v \in \widetilde{V}\left(\boldsymbol{\mu}_{g}\right) \tag{6.7}
\end{equation*}
$$

Then, the weak solution of problem (6.4) is given by $\widetilde{u}(\boldsymbol{\mu})+\widetilde{l}(\boldsymbol{\mu})$.
Let us now re-state the variational problem (6.7) onto the reference domain $\Omega$. For this purpose, let $\Gamma_{D}$ and $\Gamma_{N}$ be the portions of the boundary $\Gamma=\partial \Omega$ on which we impose Dirichlet and Neumann boundary conditions, respectively. Moreover, we denote by $ل_{\boldsymbol{\Phi}}\left(\cdot ; \boldsymbol{\mu}_{g}\right)$ the Jacobian of the parametrized map $\boldsymbol{\Phi}\left(\cdot ; \boldsymbol{\mu}_{g}\right)$, with determinant $\left|J_{\boldsymbol{\Phi}}\left(\cdot ; \boldsymbol{\mu}_{g}\right)\right|$, and we set

$$
k(\boldsymbol{x}, \cdot ; \boldsymbol{\mu})=\widetilde{k}\left(\boldsymbol{\Phi}\left(\boldsymbol{x} ; \boldsymbol{\mu}_{g}\right), \cdot ; \boldsymbol{\mu}_{p h}\right), s(\boldsymbol{x} ; \boldsymbol{\mu})=\widetilde{s}\left(\boldsymbol{\Phi}\left(\boldsymbol{x} ; \boldsymbol{\mu}_{g}\right) ; \boldsymbol{\mu}_{p h}\right) \text { and } h(\boldsymbol{x} ; \boldsymbol{\mu})=\widetilde{h}\left(\boldsymbol{\Phi}\left(\boldsymbol{x} ; \boldsymbol{\mu}_{g}\right) ; \boldsymbol{\mu}_{p h}\right)
$$

Letting

$$
V=H_{\Gamma_{D}}^{1}(\Omega)
$$

and exploiting standard change of variables formulas, the variational formulation of the Poisson problem (6.4) over $\Omega$ reads: given $\boldsymbol{\mu} \in \mathscr{P}$, find $u(\boldsymbol{\mu}) \in V$ such that

$$
a(u(\boldsymbol{\mu}), v ; \boldsymbol{\mu})=f(\nu ; \boldsymbol{\mu}) \quad \forall v \in V,
$$

with

$$
\begin{aligned}
& a(w, v ; \boldsymbol{\mu})= \int_{\Omega} k(w+l(\boldsymbol{\mu}) ; \boldsymbol{\mu}) \rrbracket_{\boldsymbol{\Phi}}^{-T}\left(\boldsymbol{\mu}_{g}\right) \nabla w \cdot J_{\boldsymbol{\Phi}}^{-T}\left(\boldsymbol{\mu}_{g}\right) \nabla v\left|\rrbracket_{\boldsymbol{\Phi}}\left(\boldsymbol{\mu}_{g}\right)\right| d \Omega \\
&+\int_{\Omega} k(w+l(\boldsymbol{\mu}) ; \boldsymbol{\mu}) \rrbracket_{\boldsymbol{\Phi}}^{-T}\left(\boldsymbol{\mu}_{g}\right) \nabla l(\boldsymbol{\mu}) \cdot \rrbracket_{\boldsymbol{\Phi}}^{-T}\left(\boldsymbol{\mu}_{g}\right) \nabla v\left|\rrbracket_{\boldsymbol{\Phi}}\left(\boldsymbol{\mu}_{g}\right)\right| d \Omega \\
& f(v ; \boldsymbol{\mu})=\int_{\Omega} s(\boldsymbol{\mu}) v\left|\rrbracket_{\boldsymbol{\Phi}}\left(\boldsymbol{\mu}_{g}\right)\right| d \Omega
\end{aligned}
$$

for any $w, v \in V$ and $\boldsymbol{\mu} \in \mathscr{P}$. Note that, as in (6.6), we resort to a lifting function $l(\boldsymbol{\mu}) \in H^{1}(\Omega)$ with $\left.l(\boldsymbol{\mu})\right|_{\Gamma_{D}}=h(\boldsymbol{\mu})$, such that the weak solution to problem (6.4) re-stated over $\Omega$ is obtained as $u(\boldsymbol{\mu})+l(\boldsymbol{\mu})$.

### 6.1.1. One-dimensional test case

Consider the following BVP for the one-dimensional Poisson equation, featuring an exponential nonlinearity (in the diffusion coefficient) and involving three parameters:

$$
\left\{\begin{array}{l}
-\left(\exp (u(\boldsymbol{\mu})) u(\boldsymbol{\mu})^{\prime}\right)^{\prime}=s(x ; \boldsymbol{\mu}) \quad \text { in } \Omega=\left(-\frac{\pi}{2}, \frac{\pi}{2}\right)  \tag{6.8}\\
\left.u(\boldsymbol{\mu})\right|_{x= \pm \pi / 2}=\mu_{2}\left(2 \pm \sin \left(\frac{\mu_{1} \pi}{2}\right)\right) \exp \left( \pm \frac{\mu_{3} \pi}{2}\right)
\end{array}\right.
$$

Here, $\boldsymbol{\mu}=\left(\mu_{1}, \mu_{2}, \mu_{3}\right) \in \mathscr{P}=[1,3] \times[1,3] \times[-0.5,0.5]$ and $s(\cdot ; \boldsymbol{\mu})$ is defined such that

$$
u_{e x}(x ; \boldsymbol{\mu})=\mu_{2}\left(2+\sin \left(\mu_{1} x\right)\right) \exp \left(\mu_{3} x\right)
$$

is the exact solution to the problem for any $\boldsymbol{\mu} \in \mathscr{P}$. The left plot of Fig. 6.1 shows the convergence of the POD-G and POD-NN solutions with respect to the number of basis functions $L$ retained in the model, generated via POD of $N=100$ FE snapshots


Figure 6.1. Convergence analysis for the POD-G and POD-NN methods applied to problem (6.8) (left) and comparison between the FE and the POD-NN solutions for three parameter values (right). These latter results have been obtained via a neural network with $H_{1}=H_{2}=35$ units per hidden layer and employing $L=10$ POD modes.


Figure 6.2. Error analysis for the POD-NN RB method applied to problem (6.8) for several numbers of training samples. The solid sections represent the steps followed by the automatic routine described in Section 5 .
computed over a uniform mesh of 100 nodes. Observe how the performance of the POD-NN method depends highly on the number of training patterns and hidden neurons used. In fact, for $N_{t r}=100$ and $H_{1}=H_{2}=15$ (upward-pointing triangles), the method can approximate only the first five generalized coordinates, actually providing more satisfactory results than POD-G does. Yet, the error stagnates for $L>5$. However, as we expand the training set, we obtain more accurate predictions for a larger number of POD coefficients, provided that we allow the size of the network to increase. Particularly, employing $N_{t r}=400$ training samples and $H_{1}=H_{2}=35$ internal neurons (right-pointing triangles), the POD-NN error features an exponential decay for $L \leq 10$, resembling the projection error (squares). The reliability of the proposed RB method is also confirmed by the right plot of Fig. 6.1, offering a comparison between the FE and POD-NN solutions for some parameter values.

Diving deeper into the analysis of the POD-NN method, Fig. 6.2 offers the results concerning the search for an optimal three-layers network configuration using the methodology described in Section 5 . The steps followed by the routine are represented by solid tracts; however, for the sake of completeness, we also report the test error (i.e., the minimum over multiple restarts) for any considered number of hidden neurons per layer ( $H_{1}$ and $H_{2}$ ) and any dimension of the training set ( $N_{t r}$ ). We remark that the basic assumption which the proposed routine relies on is fulfilled: as the number of available samples increases, the amount of neurons which allows to attain the minimum error increases as well, while the minimum error itself decreases.

### 6.1.2. Two-dimensional test case

Let $\boldsymbol{\mu}=\left[\mu_{1}, \mu_{2}, \mu_{3}\right] \in \mathscr{P}=[-0.5,0.5] \times[-0.5,0.5] \times[1,5]$ and $\widetilde{\Omega}(\boldsymbol{\mu})$ be the stenosis geometry reported on the left in Fig. 6.3, parametrized in the depths $\mu_{1}$ and $\mu_{2}$ of the bottom and top restrictions (or inflations), respectively, and the length $\mu_{3}$ of the vessel. The Poisson problem we deal with reads:

$$
\begin{cases}-\widetilde{\nabla} \cdot(\exp (\widetilde{u}(\boldsymbol{\mu})) \widetilde{\nabla} \widetilde{u}(\boldsymbol{\mu}))=\widetilde{s}(\widetilde{x}, \widetilde{y}) & \text { in } \widetilde{\Omega}(\boldsymbol{\mu}),  \tag{6.9}\\ \widetilde{u}(\boldsymbol{\mu})=\widetilde{\sigma}_{x} \sin \left(\pi \widetilde{\sigma}_{x}\right) \cos \left(\pi \widetilde{\sigma}_{y}\right) & \text { on } \partial \widetilde{\Omega}(\boldsymbol{\mu}) .\end{cases}
$$

with the source term $\widetilde{s}=\widetilde{s}(\widetilde{x}, \widetilde{y})$ properly chosen so that the exact solution to the problem is given by

$$
\tilde{u}_{e x}(\tilde{x}, \tilde{y})=\tilde{y} \sin (\pi \tilde{x}) \cos (\pi \tilde{y})
$$

for all $\boldsymbol{\mu} \in \mathscr{P}$. Although the state equation has an exponential nonlinearity and the computational domain presents curvilinear boundaries, both of the RB methodologies studied in this work provide accurate solutions, close to the optimal one, i.e., $u_{h}^{\mathbb{V}}(\boldsymbol{\mu})$. This can be seen in Fig. 6.4, offering the finite element solution (top left) to (6.9) when $\boldsymbol{\mu}=(0.349,-0.413,4.257)$,


Figure 6.3. The physical (left) and reference (right) domains for the Poisson problem (6.9).


Figure 6.4. FE solution (top left) to the Poisson problem (6.9) with $\boldsymbol{\mu}=(0.349,-0.413,4.257)$, and pointwise errors yielded by either its projection onto $V_{\text {rb }}$ (top right), or the POD-G method (bottom left), or the POD-NN method (bottom right). The results have been obtained by employing $L=30$ POD modes.
and the pointwise errors committed by either the projection onto $V_{\mathrm{rb}}$ (top right), or the POD-G method (bottom left), or the POD-NN method (bottom right). The computational mesh consists of 2792 nodes, resulting in $N_{h}=2632$ degrees of freedom (as many as the inner nodes) for the FE scheme. The reduced space $V_{\mathrm{rb}}$ is generated by $L=30$ basis functions, given by POD of $N=100$ snapshots. Specifically, the results concerning the POD-NN method have been obtained by employing a neural network equipped with $H_{1}=H_{2}=35$ neurons per hidden layer and trained with $N_{t r}=200$ learning samples. In this way, the POD-NN procedure leads to an error which shows patterns similar to those featured by the projection error. This is not surprising, due to the way neural networks are trained within the POD-NN framework (see Section 5).

A quantitative evidence of the efficacy of the RB approximations is given in the plot on the left in Fig. 6.5, which reports the convergence analysis for both the POD-G and the POD-NN methods with respect to the number $L$ of retained modal basis functions. As usual, the error has to be interpreted as an average over a parameter data set $\Xi_{t e}$, here consisting of $N_{t e}=50$ samples. However, the second plot of Fig. 6.5 reveals how the former scheme produces a response for any online query approximately $10^{3}$ times slower than the proposed POD-NN procedure does.


Figure 6.5. Error analysis (left) and online CPU time (right) for the POD-G and the POD-NN methods applied to problem (6.9). The reduced basis have been generated via POD, relying on $N=100$ snapshots. The second plot refers to RB models including $L=30$ modal functions; within the POD-NN framework, an MLP emboding 35 neurons per inner layer has been used.


Figure 6.6. Convergence analysis with respect to the number of hidden neurons (left) and modal functions (right) used within the POD-NN framework applied to problem (6.9). The results provided in the first plot have been obtained using $L=30$ modes; the solid tracts refer to the steps performed by the automatic routine carried out to find an optimal network configuration.

Let us further analyze the first plot of Fig. 6.5. When the proper orthogonal decomposition of a set of $N=100$ snapshots is coupled with a three-layers perceptron with 35 neurons per hidden layer, trained to approximate the map (5.2) relying on $N_{t r}=200$ learning patterns, the relative error is smaller than the error committed by the POD-Galerkin procedure for each $L \leq 30$. Yet, halfing the number of learning patterns employed, the POD-NN error curve stops decreasing at $L=20$, then incurring a plateu. In this regard, Fig. 6.6 suggests that as the size of the available training set decreases, one should reduce the amount of computing units to be embodied in the network, to limit the risk of overfitting. Particularly, when only $N_{t r}=100$ training samples are used, the optimal network configuration comprises 20 neurons per internal layer (left). However, doubling the dimension of the training set, the same configuration does not allow to fully exploit the augmented amount of information. In that case, a network with, e.g., 35 neurons per hidden layer is preferable (right).

The distinguishing and novel feature of the POD-NN method is represented by the employment of multi-layer perceptrons to recover the coefficients of the reduced model. To motivate this choice, let us pursue a more traditional approach in


Figure 6.7. Average relative errors (left) and online run times (right) on $\Xi_{t e}$ for the POD-CS and the POD-NN methods applied to problem (6.9). For the latter, the results refer to an MLP equipped with $H_{1}=H_{2}=35$ neurons per hidden layer.
the interpolation step by resorting to cubic splines [14]. To this end, let $\Xi_{\delta} \subset \mathscr{P}$ be a tensor-product grid on the parameter domain, based on, e.g., Chebyshev nodes. In the offline phase, for each $\boldsymbol{\mu}_{\delta} \in \Xi_{\delta}$, we compute the truth solution $\mathbf{u}_{h}\left(\boldsymbol{\mu}_{\delta}\right) \in \mathbb{R}^{N_{h}}$ and we extract the expansion coefficients $\mathbb{V}^{T} \mathbf{u}_{h}\left(\boldsymbol{\mu}_{\delta}\right) \in \mathbb{R}^{L}$. At the online stage, given a new parameter value $\boldsymbol{\mu} \in \mathscr{P}$, the $i$-th expansion coefficient, $i=1, \ldots, L$, is sought by cubic spline interpolation of the samples

$$
\left\{\left(\boldsymbol{\mu}_{\delta},\left(\mathbb{V}^{T} \mathbf{u}_{h}\left(\boldsymbol{\mu}_{\delta}\right)\right)_{i}\right)\right\}_{\boldsymbol{\mu}_{\delta} \in \Xi_{\delta}}
$$

Hence, denoting by $\mathbf{u}_{\mathrm{rb}}^{\mathrm{CS}}(\boldsymbol{\mu}) \in \mathbb{R}^{L}$ the approximation of $\mathbb{V}^{T} \mathbf{u}_{h}(\boldsymbol{\mu})$, the reduced order approximation of $\mathbf{u}_{h}(\boldsymbol{\mu})$ is given by $\mathbf{u}_{L}^{\mathrm{CS}}(\boldsymbol{\mu})=\mathbb{V} \mathbf{u}_{\mathrm{rb}}^{\mathrm{CS}}(\boldsymbol{\mu})$. Similarly to the POD-G and the POD-NN methods, the accuracy of the resulting POD-CS procedure can be assessed by evaluating and averaging the relative error $\varepsilon_{\text {PODCS }}(L, \boldsymbol{\mu})$, defined as

$$
\varepsilon_{\mathrm{PODCS}}(L, \boldsymbol{\mu})=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbf{u}_{L}^{\mathrm{CS}}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|}=\frac{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})-\mathbb{V} \mathbf{u}_{\mathrm{rb}}^{\mathrm{CS}}(\boldsymbol{\mu})\right\|}{\left\|\mathbf{u}_{h}(\boldsymbol{\mu})\right\|}
$$

on a test parameter set $\Xi_{t e} \subset \mathscr{P}$, with $\Xi_{t e} \cap \Xi_{\delta}=\varnothing$.
For the Poisson problem (6.9), we report in the left plot of Fig. 6.7 the relative errors by both the POD-CS and the PODNN methods. For the former, we provide the results obtained on tensor-product grids consisting of $N_{\delta}$ interpolation points, with $N_{\delta} \in\left\{5^{3}, 7^{3}, 10^{3}\right\}$. For the latter, $H_{1}=H_{2}=35$ neurons per hidden layer and $N_{t r}=200$ training patterns have been used. Although the online performance of the two procedures are basically the same, as shown by the second plot in Fig. 6.7, we observe that the level of predictive accuracy enabled by the POD-NN method can be attained or at least approached by cubic spline interpolation only when this relies on $N_{\delta}=1000$ samples. In fact, as mentioned in the Introduction, a standard interpolation technique may require a large number of samples to be able to enforce the constraints characterizing the nonlinear manifolds which the reduced bases typically belong to [1]. Hence, although this approach provides a valuable alternative for paremetrized problems with a few parameters, it may be infeasible in real-life applications involving a large number of parameters. Conversely, the hope behind the choice of a neural network-based regression is that the amount of required snapshots properly scales with the dimension of the parameter space. This would justify the overheads due to the training phase.

### 6.2. Steady incompressible Navier-Stokes equations

The system of the Navier-Stokes equations model the conservation of mass and momentum for an incompressible Newtonian viscous fluid confined in a two- or three-dimensional region [42]. Letting $\boldsymbol{\mu} \equiv \boldsymbol{\mu}_{g}$ and $\widetilde{\boldsymbol{v}}=\widetilde{\boldsymbol{v}}(\widetilde{\boldsymbol{x}} ; \boldsymbol{\mu})$ and $\widetilde{p}=\widetilde{p}(\widetilde{\boldsymbol{x}} ; \boldsymbol{\mu})$ be, respectively, the velocity and pressure of the fluid, the parametrized steady version of the Navier-Stokes equations
considered here reads

$$
\left\{\begin{array}{rr}
\widetilde{\nabla} \cdot \widetilde{\boldsymbol{v}}(\boldsymbol{\mu})=0 & \text { in } \widetilde{\Omega}(\boldsymbol{\mu}),  \tag{6.10a}\\
-v(\boldsymbol{\mu}) \widetilde{\Delta} \widetilde{\boldsymbol{v}}(\boldsymbol{\mu})+(\widetilde{\boldsymbol{v}}(\boldsymbol{\mu}) \cdot \widetilde{\nabla}) \widetilde{\boldsymbol{v}}(\boldsymbol{\mu})+\frac{1}{\rho(\boldsymbol{\mu})} \widetilde{\nabla} \widetilde{p}(\boldsymbol{\mu})=\mathbf{0} & \text { in } \widetilde{\Omega}(\boldsymbol{\mu}), \\
\widetilde{\boldsymbol{v}}(\boldsymbol{\mu})=\widetilde{\boldsymbol{h}} & \text { on } \widetilde{\Gamma}_{D}(\boldsymbol{\mu}), \\
\widetilde{p}(\boldsymbol{\mu}) \widetilde{\boldsymbol{n}}-v(\boldsymbol{\mu}) \widetilde{\nabla} \widetilde{\boldsymbol{v}}(\boldsymbol{\mu}) \cdot \widetilde{\boldsymbol{n}}=\mathbf{0} & \text { on } \widetilde{\Gamma}_{N}(\boldsymbol{\mu}) .
\end{array}\right.
$$

Here, $\widetilde{\boldsymbol{h}}$ denotes the velocity field prescribed on $\widetilde{\Gamma}_{D}(\boldsymbol{\mu})$, whereas homogeneous Neumann conditions are applied on $\widetilde{\Gamma}_{N}(\boldsymbol{\mu})$. Furthermore, $\rho(\boldsymbol{\mu})$ and $v(\boldsymbol{\mu})$ represent the uniform density and kinematic viscosity of the fluid, respectively. Note that, despite that these quantities encode physical properties, we let them depend on the geometrical parameters as well. Indeed, fluid dynamics can be characterized (and controlled) by means of a wealth of dimensionless quantities, e.g., the Reynold's number, which combine physical properties of the fluid with geometrical features of the domain. Therefore, a numerical study of the sensitivity of the system (6.10) with respect to $\boldsymbol{\mu}$ may be carried out by adapting either $\rho(\boldsymbol{\mu})$ or $v(\boldsymbol{\mu})$ as $\boldsymbol{\mu}$ varies, so to preserve a dimensionless quantity of interest.

To write the differential system (6.10) in weak form over a $\boldsymbol{\mu}$-independent configuration $\Omega$, let us introduce the velocity and pressure spaces

$$
\widetilde{X}(\boldsymbol{\mu})=\left[H_{\widetilde{\Gamma}_{D}}^{1}(\widetilde{\Omega}(\boldsymbol{\mu}))\right]^{d} \text { and } \widetilde{Q}(\boldsymbol{\mu})=L^{2}(\widetilde{\Omega}(\boldsymbol{\mu})),
$$

respectively, with

$$
X=\left[H_{\Gamma_{D}}^{1}(\Omega)\right]^{d} \text { and } Q=L^{2}(\Omega)
$$

be their respective counterparts over $\Omega$. By multiplying (6.10) by test functions $(\widetilde{\chi}, \widetilde{\xi}) \in \widetilde{V}(\boldsymbol{\mu})=\widetilde{X}(\boldsymbol{\mu}) \times \widetilde{Q}(\boldsymbol{\mu})$, integrating by parts and then tracing everything back onto $\Omega$ by means of the parametrized map $\Phi(\cdot ; \boldsymbol{\mu})$, we end up with the following parametrized weak variational problem: given $\boldsymbol{\mu} \in \mathscr{P}$, find $u(\boldsymbol{\mu})=(\boldsymbol{\nu}(\boldsymbol{\mu}), p(\boldsymbol{\mu})) \in V=X \times Q$ so that

$$
a(\boldsymbol{v}(\boldsymbol{\mu}), \boldsymbol{\chi} ; \boldsymbol{\mu})+c(\boldsymbol{v}(\boldsymbol{\mu}), \boldsymbol{v}(\boldsymbol{\mu}), \boldsymbol{\chi} ; \boldsymbol{\mu})+d(\boldsymbol{v}(\boldsymbol{\mu}), \boldsymbol{\chi} ; \boldsymbol{\mu})+b(p(\boldsymbol{\mu}), \nabla \cdot \boldsymbol{\chi} ; \boldsymbol{\mu})=f_{1}(\boldsymbol{\chi} ; \boldsymbol{\mu}),
$$

for all $(\boldsymbol{\chi}, \xi) \in V$, with, for any $\boldsymbol{\mu} \in \mathscr{P}$,

$$
\begin{array}{ll}
c(\boldsymbol{\psi}, \boldsymbol{\chi}, \boldsymbol{\eta} ; \boldsymbol{\mu})=\int_{\Omega}\left(\boldsymbol{\psi} \cdot J_{\Phi}^{-T}(\boldsymbol{\mu}) \nabla\right) \boldsymbol{\chi} \cdot \boldsymbol{\eta}\left|J_{\Phi}(\boldsymbol{\mu})\right| d \Omega, & d(\boldsymbol{\psi}, \boldsymbol{\chi} ; \boldsymbol{\mu})=c(\boldsymbol{l}(\boldsymbol{\mu}), \boldsymbol{\psi}, \boldsymbol{\chi} ; \boldsymbol{\mu})+c(\boldsymbol{\psi}, \boldsymbol{l}(\boldsymbol{\mu}), \boldsymbol{\chi} ; \boldsymbol{\mu}), \\
a(\boldsymbol{\psi}, \boldsymbol{\chi} ; \boldsymbol{\mu})=v(\boldsymbol{\mu}) \int_{\Omega} J_{\Phi}^{-T}(\boldsymbol{\mu}) \nabla \boldsymbol{\psi}: \sqrt[J]{\Phi}_{-T}(\boldsymbol{\mu}) \nabla \boldsymbol{\chi}\left|ป_{\Phi}(\boldsymbol{\mu})\right| d \Omega, & f_{1}(\boldsymbol{\psi} ; \boldsymbol{\mu})=-a(\boldsymbol{l}(\boldsymbol{\mu}), \boldsymbol{\psi} ; \boldsymbol{\mu})-c(\boldsymbol{l}(\boldsymbol{\mu}), \boldsymbol{l}(\boldsymbol{\mu}), \boldsymbol{\psi} ; \boldsymbol{\mu}), \\
b(\boldsymbol{\psi}, \xi ; \boldsymbol{\mu})=-\frac{1}{\rho(\boldsymbol{\mu})} \int_{\Omega}\left(J_{\Phi}^{-T}(\boldsymbol{\mu}) \nabla \cdot \boldsymbol{\psi}\right) \xi\left|ป_{\Phi}(\boldsymbol{\mu})\right| d \Omega, & f_{2}(\xi ; \boldsymbol{\mu})=-b(\boldsymbol{l}(\boldsymbol{\mu}), \xi ; \boldsymbol{\mu}),
\end{array}
$$

for all $\boldsymbol{\psi}, \boldsymbol{\chi}, \boldsymbol{\eta} \in X$ and $\xi \in Q$. In the definitions of $d(\cdot, \cdot ; \boldsymbol{\mu}), f_{1}(\cdot ; \boldsymbol{\mu})$ and $f_{2}(\cdot ; \boldsymbol{\mu}), \boldsymbol{l}(\boldsymbol{\mu}) \in\left[H^{1}(\Omega)\right]^{d}$ denotes the lifting vector field, with $\left.\boldsymbol{l}(\boldsymbol{\mu})\right|_{\Gamma_{D}}=\boldsymbol{h}(\boldsymbol{\mu}), \boldsymbol{h}(\boldsymbol{x} ; \boldsymbol{\mu})=\widetilde{\boldsymbol{h}}(\boldsymbol{\Phi}(\boldsymbol{x} ; \boldsymbol{\mu}))$ being the velocity field prescribed on $\Gamma_{D}$. Hence, the weak solution to (6.10) defined over the fixed domain $\Omega$ is given by $(\boldsymbol{\nu}(\boldsymbol{\mu})+\boldsymbol{l}(\boldsymbol{\mu}), p(\boldsymbol{\mu}))$.

### 6.2.1. The lid-driven cavity problem

In this subsection, we are concerned with the numerical simulation of a viscous flow within a parallelogram-shaped cavity, illustrated in Fig. 6.8 (left). The geometry of the domain is affected by three parameters: $\mu_{1} \in[1,2]$ and $\mu_{2} \in[1,2]$ define the length of the horizontal and slanting (possibly vertical) edges, respectively, whereas $\mu_{3} \in[\pi / 6,5 \pi / 6]$ denotes the angle between the oblique sides and the positive $\widetilde{x}$-semiaxis. The Dirichlet boundary conditions enforced on the velocity field, graphically represented in Fig. 6.8 (center), are such that the flow is driven by a unit horizontal velocity at the top boundary [38]. Hence, this benchmark is typically referred to as the lid-driven cavity problem. In addition, we fix the pressure at the lower-left corner, in order to retain the uniqueness of the solution [16]. Therefore, the computational mesh $\Omega_{h}$ (Fig. 6.8, right) is refined in the upper part of the domain, so to properly resolve the steep velocity gradient near the upper boundary and the pressure singularities at the top corners.


Figure 6.8. Computational domain $\widetilde{\Omega}(\boldsymbol{\mu})$ (left), enforced velocity at the boundaries (center) and computational mesh $\Omega_{h}$ (right) for the lid-driven cavity problem for the incompressible Navier-Stokes equations.


Figure 6.9. Velocity streamlines (top) and pressure distribution (bottom) for the lid-driven cavity problem, computed through the FE method. Three different parameter values are considered; for all configurations, the Reynold's number is 400 .

For the Navier-Stokes equations, a suitable choice of the FE spaces is crucial to fulfill the well-known inf-sup stability condition [42]. A common and effective choice consists in using quadratic finite elements for the components of the velocity field and linear finite elements for the pressure, leading to the so-called $\mathbb{P}^{2}-\mathbb{P}^{1}$ (or Taylor-Hood) FE discretization [38]. Figure 6.9 shows the velocity streamlines (top) and the pressure distribution (bottom) obtained through the resulting FE method for, from the left to the right, $\boldsymbol{\mu}=(1,2 / \sqrt{3}, 2 \pi / 3), \boldsymbol{\mu}=(1,1, \pi / 2)$ or $\boldsymbol{\mu}=(1,2 / \sqrt{3}, \pi / 6)$. Observe how the moving lid induces a large velocity gradient close to the top boundary, while it only slightly affects the fluid in the lower part of the
cavity. In turn, this gives rise to two vortices, whose extent highly depends on the shape of the domain, or, equivalently, $\boldsymbol{\mu}$. Conversely, the pressure field does not undergo noticeable alterations across the parameter space. For all the configurations, a low-pressure region takes place at the upper-left corner of the domain and in the center of the major vortex (although it may be not clearly visible). Instead, the upper-right corner represents a stagnation point for the $\widetilde{x}$-velocity, and so therein the pressure assumes values larger than in the rest of the cavity [16]. All the results presented in Fig. 6.9 refer to a Reynold's number of 400 . For the specific problem at hand, the Reynold's number reads:

$$
R e=\frac{\max \left\{\mu_{1}, \mu_{2}\right\}}{v(\boldsymbol{\mu})}
$$

In our analyses, $v(\boldsymbol{\mu})$ is adapted as the geometry varies so that the Reynold's number is either 200 or 400.
Regarding the reduced order modeling of the parametrized problem of interest, we slightly distance ourselves from the general workflow depicted in Sections 3 and 5. Indeed, for the Navier-Stokes equations it is convenient to construct two separate reduced spaces $[2,11,41]: X_{\mathrm{rb}}$, of dimension $L_{\boldsymbol{v}}$, for the velocity field and $Q_{\mathrm{rb}}$, of dimension $L_{p}$, for the pressure distribution, respectively represented by the matrices

$$
\mathbb{V}_{\boldsymbol{v}}=\left[\boldsymbol{\psi}_{1}^{\boldsymbol{v}}|\ldots| \boldsymbol{\psi}_{L_{\boldsymbol{v}}}^{\boldsymbol{\nu}}\right] \in \mathbb{R}^{N_{h}^{\boldsymbol{v}} \times L_{\boldsymbol{v}}} \quad \text { and } \quad \mathbb{V}_{p}=\left[\boldsymbol{\psi}_{1}^{p}|\ldots| \boldsymbol{\psi}_{L_{p}}^{p}\right] \in \mathbb{R}^{N_{h}^{p} \times L_{p}}
$$

with $N_{h}^{\boldsymbol{v}}$ (respectively, $N_{h}^{p}$ ) the dimension of the FE velocity (resp., pressure) space. Although the underpinning finite element solver has been designed to fulfill the inf-sup condition, thus ensuring the stability of the scheme, when dealing with incompressible flows this property may not be automatically inherited by the POD-Galerkin reduced order solver. As a result, it may show severe instabilities [8]. Indeed, one has to carefully choose the reduced velocity space $X_{\mathrm{rb}}$ so to meet an equivalent of the inf-sup condition at the reduced level. To this end, in our simulations we resort to a supremizer enrichment of $X_{\mathrm{rb}}$, as proposed in [2]. At each pressure snapshot $\mathbf{p}_{h}\left(\boldsymbol{\mu}^{(n)}\right) \in \mathbb{R}^{N_{h}^{p}}, n=1, \ldots, N$, we associate the supremizer solution $\mathbf{s}_{h}\left(\boldsymbol{\mu}^{(n)}\right) \in \mathbb{R}^{N_{h}^{v}}$, defined as the solution to a linear system of the form

$$
\mathbb{A} \mathbf{s}_{h}\left(\boldsymbol{\mu}^{(n)}\right)=\mathbb{B}\left(\boldsymbol{\mu}^{(n)}\right) \mathbf{p}_{h}\left(\boldsymbol{\mu}^{(n)}\right),
$$

with $\mathbb{A} \in \mathbb{R}^{N_{h}^{v} \times N_{h}^{v}}$ and $\mathbb{B}\left(\boldsymbol{\mu}^{(n)}\right) \in \mathbb{R}^{N_{h}^{\boldsymbol{v}} \times N_{h}^{p}}$. Letting

$$
\mathbb{V}_{\boldsymbol{s}}=\left[\boldsymbol{\psi}_{1}^{\boldsymbol{s}}|\ldots| \boldsymbol{\psi}_{L_{s}}^{\boldsymbol{s}}\right] \in \mathbb{R}^{N_{h}^{\boldsymbol{v}} \times L_{s}}
$$

be the POD basis of rank $L_{\boldsymbol{s}}$ extracted from the ensemble of snapshots $\left\{\mathbf{s}_{h}\left(\boldsymbol{\mu}^{(1)}\right), \ldots, \mathbf{s}_{h}\left(\boldsymbol{\mu}^{(N)}\right)\right\}$, the POD-G approximation of the velocity field is then sought in the form

$$
\overline{\mathbf{v}}_{L}(\boldsymbol{\mu})=\overline{\mathbb{V}}_{\boldsymbol{\nu}} \overline{\mathbf{v}}_{\mathrm{rb}}(\boldsymbol{\mu})=\sum_{i=1}^{L_{\boldsymbol{v}}} v_{\mathrm{rb}}^{(i)} \boldsymbol{\psi}_{i}^{\boldsymbol{\nu}}+\sum_{i=1}^{L_{\boldsymbol{s}}} s_{\mathrm{rb}}^{(i)} \boldsymbol{\psi}_{i}^{\boldsymbol{s}}
$$

where

$$
\overline{\mathbb{V}}_{\boldsymbol{v}}=\left[\mathbb{V}_{\boldsymbol{v}} \mathbb{V}_{\boldsymbol{s}}\right] \in \mathbb{R}^{N_{h}^{\boldsymbol{v}} \times\left(L_{\boldsymbol{v}}+L_{s}\right)} \quad \text { and } \quad \overline{\mathbf{v}}_{\mathrm{rb}}(\boldsymbol{\mu})=\left[\begin{array}{c}
\mathbf{v}_{\mathrm{rb}}(\boldsymbol{\mu}) \\
\mathbf{s}_{\mathrm{rb}}(\boldsymbol{\mu})
\end{array}\right] \in \mathbb{R}^{L_{\boldsymbol{v}}+L_{\boldsymbol{s}}}
$$

In [2], several numerical evidences suggest that a proper value for $L_{s}$ should lay in between $L_{p} / 2$ and $L_{p}$. In our simulations, we set $L_{\boldsymbol{s}}=L_{p}=L_{\boldsymbol{v}}$, leading to a stable reduced (nonlinear) system in $L=3 L_{\boldsymbol{v}}$ unknowns.

Concerning the POD-NN procedure, for any $\boldsymbol{\mu} \in \mathscr{P}$, reduced order discretizations $\mathbf{v}_{L_{\boldsymbol{v}}}^{\mathrm{NN}}(\boldsymbol{\mu})$ and $\mathbf{p}_{L_{p}}^{\mathrm{NN}}(\boldsymbol{\mu})$ of $\boldsymbol{v}(\boldsymbol{\mu})$ and $p(\boldsymbol{\mu})$, respectively, are sought in the form

$$
\mathbf{v}_{L_{\boldsymbol{v}}}^{\mathrm{NN}}(\boldsymbol{\mu})=\mathbb{V}_{\boldsymbol{v}} \mathbf{v}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu}) \quad \text { and } \quad \mathbf{p}_{L_{p}}^{\mathrm{NN}}(\boldsymbol{\mu})=\mathbb{V}_{p} \mathbf{p}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})
$$

Here, $\mathbf{v}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})$ and $\mathbf{p}_{\mathrm{rb}}^{\mathrm{NN}}(\boldsymbol{\mu})$ denote the output vectors provided by two distinct multi-layer perceptrons, respectively trained with the ensemble of input-output patterns

$$
\left\{\boldsymbol{\mu}^{(i)}, \mathbb{V}_{\boldsymbol{\nu}}^{T} \mathbf{v}_{h}\left(\boldsymbol{\mu}^{(i)}\right)\right\}_{1 \leq i \leq N_{t r}} \quad \text { and } \quad\left\{\boldsymbol{\mu}^{(i)}, \mathbb{V}_{p}^{T} \mathbf{p}_{h}\left(\boldsymbol{\mu}^{(i)}\right)\right\}_{1 \leq i \leq N_{t r}}
$$



Figure 6.10. Velocity (left) and pressure (right) error analysis for the POD-G and POD-NN methods applied to the lid-driven cavity problem with $R e=200$ (top) and $R e=400$ (bottom).

At this regard, let us point out two important remarks.
(i) The supremizer solutions are not involved (either directly or indirectly) in the POD-NN framework. Indeed, they ensure the overall stability of the POD-G procedure, but they do not improve the accuracy of the method, and so they can be disregarded whenever stability is not an issue.
(ii) The routine outlined in Section 5, aiming at finding an optimal network configuration, is applied separately to the perceptrons designated to predict the velocity field and to the perceptrons required to approximate the pressure distribution. As a result, once the offline phase of the POD-NN method is completed, we end up with two different networks, equipped with $H_{1}^{v}=H_{2}^{v}$ and $H_{1}^{p}=H_{2}^{p}$ computing units per hidden layer, respectively.
Figure 6.10 reports the error committed by both RB procedures when approximating the velocity field (left) and the pressure distribution (right) by means of $L_{\boldsymbol{v}}$ velocity modes, $L_{p}$ pressure modes and, exclusively for the POD-Galerkin method, $L_{\boldsymbol{s}}$ supremizer modes, with $5 \leq L_{\boldsymbol{v}}=L_{s}=L_{p} \leq 35$. The plots on the top refer to a Reynold's number of 200, the ones on the bottom to a Reynold's number of 400 . Note that for clarity of illustration, the symbols denoting the projection, POD-G and POD-NN errors have been identified with a superscript (either $\boldsymbol{v}$ or $p$ ) recalling the state variable they refer to.

While the error yielded on the velocity field shows an almost perfect exponential decay with the number of POD modes included in the RB model, the POD-G method faces difficulties in providing a correct recovery of the pressure, already for $R e=200$. Indeed, the corresponding error curve is not monotone, and generally is at least one order of magnitude


Figure 6.11. Convergence analysis with respect to the number of hidden neurons and training samples used within the POD-NN procedure to approximate the velocity field (left) and the pressure distribution (right) by means of 35 modal functions. The Reynold's number is either 200 (top) or 400 (bottom).
above the projection error. Conversely, the POD-NN method attains a satisfactory predictive accuracy. The advantages of resorting to a neural network-based nonlinear regression coupled with POD are evident when the approximation is built upon a few basis functions, say $L_{p}<20$. Moreover, the proposed routine recommends the use of fewer neurons to predict the POD coefficients for the pressure than for the velocity, thus resulting in a lighter perceptron.

This is confirmed also by Fig. 6.11, which provides a sensitivity analysis of the predictive accuracy featured by the POD-NN method with respect to the amount of neurons and training samples used. Observe that for the pressure (right), employing more than 30 neurons within each hidden layer is counter-productive, both for $R e=200(t o p)$ and $R e=400$ (bottom). This assertion agrees with the peculiar role played by the pressure in the parametrized lid-driven cavity problem, with similar patterns featured across the entire parameter domain.

On the contrary, we have seen that the velocity field presents more complex dynamics, highly varying with the domain configuration. As a result, an optimal approximation of the velocity is obtained for the maximum values of $H_{i}^{v}, i=1,2$, and $N_{t r}$ tested, that is, $H_{1}^{v}=H_{2}^{v}=35$ and $N_{t r}=300$ for $R e=200, H_{1}^{v}=H_{2}^{v}=40$ and $N_{t r}=300$ for $R e=400$. Moreover, we may not be able to exactly attain the same precision enabled by the standard POD-Galerkin procedure, although the results are quite similar. However, this (slight) loss of accuracy is offset by a (great) reduction in the online run time: the POD-NN method takes around $2 / 100$ of seconds per query, against the average 40 seconds required by the POD-G method; see Fig. 6.12. As for the test cases for the Poisson problem, this comes at the cost of a longer offline phase.


Figure 6.12. Online run times for the POD-G and the POD-NN method applied to the lid-driven cavity problem with $R e=200$ (left) and $R e=400(r i g h t) . N_{t e}=75$ test configurations are considered. For the POD-NN method, the reported times include the (sequential) evaluation of both neural networks for the velocity and pressure field.


Figure 6.13. $\widetilde{x}$-velocity contour at three parameter values, as computed through the FE (top row) and POD-NN (bottom row) method. For each configuration, the Reynold's number is 400 .

Another numerical evidence of the predictive accuracy of the proposed POD-NN RB method is provided in Fig. 6.13, showing the contour plots for the $\tilde{x}$-velocity computed through the FE (top row) and the POD-NN (bottom row) scheme. Three different configurations, corresponding to as many input vectors, are considered; the Reynold's number is fixed to 400. We can appreciate good agreement between the solutions given by the full-order and the reduced order methods.

Lastly, Fig. 6.14 compares the streamlines obtained through the direct method (top) and the proposed reduced basis approach (bottom) over the three configurations previously considered. Streamlines provide an interesting test, as minor variations in velocity contours may lead to substantial differences in the streamlines [11]. Nevertheless, we observe a good


Figure 6.14. Streamlines at three parameter values, as computed through the FE (top row) and POD-NN (bottom row) method. For each configuration, the Reynold's number is 400 .
agreement between the two methods. In particular, in the second example, the POD-NN method is still able to detect the two micro recirculation zones at the lower corners of the domain. On the other hand, the method partially fails in properly describing the velocity field at the bottom-left corner in the first configuration and at the bottom-right corner in the last configuration. However, these are effectively dead zones, and one can safely disregard these little imprecisions.

## 7. Conclusion

In this work, we propose a non-intrusive RB method (referred to as POD-NN) for parametrized steady-state PDEs. The method extracts a reduced basis from a collection of snapshots through a POD procedure and employs multi-layer perceptrons to approximate the coefficients of the reduced model. By exploiting the fundamental results by Cybenko (see Subsection 4.2), we limit ourselves to perceptrons endowed with two hidden layers. The identification of the optimal number of inner neurons and the minimum amount of training samples to prevent overfitting is performed during the offline stage through an automatic routine, relying upon the latin hypercube sampling and the Levenberg-Marquardt training algorithm. On the one hand, this guarantees a complete decoupling between the offline and the online phase, with the latter having a computational cost independent of the dimension of the full-order model. On the other hand, this extends the offline stage with respect to a standard projection-based RB procedure, making the POD-NN method practically convenient only when the underlying PDE has to be solved for many parameter values (many-query context).

The POD-NN method has been successfully tested on the nonlinear Poisson equation in one and two spatial dimensions, and on two-dimensional cavity viscous flows, modeled through the steady incompressible Navier-Stokes equations. In particular, the proposed RB strategy enables the same predictive accuracy provided by the POD-Galerkin method while reducing the CPU time required to process an online query by two or even three order of magnitudes.

All test cases considered in our numerical studies involved three parameters, affecting either physical or geometrical factors of the differential problem. The extension of the POD-NN method to time-dependent problems depending on many parameters is left as future work.

Lastly, let us point out that although in this work we used POD to recover a reduced space, this choice is not binding, and a greedy approach may also be pursued.

## References

[1] Amsallem., D. (2010). Interpolation on manifolds of CFD-based fluid and finite element-based structural reduced-order models for on-line aeroelastic predictions. Doctoral dissertation, Department of Aeronautics and Astronautics, Stanford University.
[2] Ballarin, F., Manzoni, A., Quarteroni, A., Rozza, G. (2014). Supremizer stabilization of POD-Galerkin approximation of parametrized Navier-Stokes equations. MATHICSE Technical Report, École Polytechnique Fédérale de Lausanne.
[3] Barrault, M., Maday, Y., Nguyen, N. C., Patera, A. T. (2004). An 'empirical interpolation' method: Application to efficient reduced-basis discretization of partial differential equations. Comptes Rendus Mathematique, 339(9):667-672.
[4] Barthelmann, V., Novak, E., Ritter, K. (2000). High-dimensional polynomial interpolation on sparse grids. Advances in Computational Mathematics, 12(4):273-288.
[5] Bendsøe, M. P., Sigmund, O. (2004). Topology optimization: Theory, methods and applications. Heidelberg, DE: Springer Science \& Business Media.
[6] Brown, P. F., Pietra, V. J. D., Pietra, S. A. D., Mercer, R. L. (1993). The mathematics of statistical machine translation: Parameter estimation. Computational linguistics, 19(2):263-311.
[7] Buffa, A., Maday, Y., Patera, A. T., Prud'Homme, C., Turinici, G. (2012). A priori convergence of the greedy algorithm for the parametrized reduced basis method. ESAIM: Mathematical Modelling and Numerical Analysis, 46:595-603.
[8] Burkardt, J., Gunzburger, M., Lee, H. C. (2006). POD and CVT-based reduced-order modeling of Navier-Stokes flows. Computer Methods in Applied Mechanics and Egninnering, 196:337-355.
[9] Casenave, F., Ern, A., Lelièvre, T. (2015). A nonintrusive reduced basis method applied to aeroacoustic simulations. Advances in Computational Mathematics, 41:961-986.
[10] Chaturantabut, S., Sorensen, D. C. (2010). Nonlinear model reduction via discrete empirical interpolation. SIAM Journal on Scientific Computing, 32(5):2737-2764.
[11] Chen, W., Hesthaven, J. S., Junqiang, B., Yang, Z., Tihao, Y. (2017). A greedy non-intrusive reduced order model for fluid dynamics. Submitted to American Institute of Aeronautics and Astronautics.
[12] Cybenko, G. (1988). Continuous valued neural networks with two hidden layers are sufficient. Technical Report, Department of Computer Science, Tufts University.
[13] Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals, and Systems, 2(4):303-314.
[14] De Boor, C. (1978). A practical guide to splines. New York, NY: Springer-Verlag.
[15] Deparis, S. (2008). Reduced basis error bound computation of parameter-dependent Navier-Stokes equations by the natural norm approach. SIAM Journal of Numerical Analysis, 46(4):2039-2067.
[16] Dhondt, G. (2014). CalculiX CrunchiX user's manual. Available at http://web.mit.edu/calculix_v2.7/CalculiX/ ccx_2.7/doc/ccx/node1.html.
[17] Eckart, C., Young, G. (1936). The approximation of one matrix by another of lower rank. Psychometrika, 1:211-218.
[18] Eftang, J. L. (2008). Reduced basis methods for partial differential equations. Master thesis, Department of Mathematical Sciences, Norwegian University of Science and Technology.
[19] Hagan, M. T., Menhaj, M. B. (1994). Training feedforward networks with the Marquardt algorithm. IEEE Transactions on Neural Networks, 5(6):989-993.
[20] Hagan, M. T., Demuth, H. B., Beale, M. H., De Jesús, O. (2014). Neural Network Design, 2nd Edition. Retrieved from http://hagan.okstate.edu/NNDesign.pdf.
[21] Haykin, S. (2004). Neural Networks: A comprehensive foundation. Upper Saddle River, NJ: Prentice Hall.
[22] Liang, Y. C., Lee, H. P., Lim, S. P., Lin, W. Z., Lee, K. H., Wu, C. G. (2002) Proper orthogonal decomposition and its applications - Part I: Theory. Journal of Sound and Vibration, 252(3):527-544.
[23] Hesthaven, J. S., Stamn, B., Rozza, G. (2016). Certified reduced basis methods for parametrized partial differential equations. New York, NY: Springer.
[24] Hesthaven, J. S., Stamn, B., Zhang, S. (2014). Efficienty greedy algorithms for high-dimensional parameter spaces with applications to empirical interpolation and reduced basis methods. ESAIM: Mathematical Modelling and Numerical Analysis, 48(1):259-283.
[25] Imam, R. L. (2008). Latin hypercube sampling. Encyclopedia of Quantitative Risk Analysis and Assessment.
[26] Jaggli, C., Iapichino, L., Rozza, G. (2014). An improvement on geometrical parametrizations by transfinite maps. Comptes Rendus de l'Académie des Sciences Paris, Series I, 352:263-268.
[27] Kohavi, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. Proceedings of the $40^{\text {th }}$ International Joint Conference on Artificial Intelligence, 2(12):1137-1143.
[28] Kriesel, D. (2007). A Brief Introduction to Neural Networks. Retrieved from http://www.dkriesel.com/en/science/ neural_networks.
[29] Le Maître, O., Knio, O. M. (2010). Spectral methods for uncertainty quantification with applications to computational fluid dynamics. Berlin, DE: Springer Science \& Business Media.
[30] Lee, E. B., Markus, L. (1967). Foundations of optimal control theory. New York, NY: John Wiley \& Sons.
[31] Maday, Y. (2006) Reduced basis method for the rapid and reliable solution of partial differential equations. Proceedings of the International Congress of Mathematicians, Madrid, Spain, 1255-1269.
[32] Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. Journal of the Society for Industrial and Applied Mathematics, 11(2):431-441.
[33] The MathWorks, Inc. (2016). Machine learning challenges: Choosing the best model and avoiding overfitting. Retrieved from https://it.mathworks.com/campaigns/products/offer/common-machine-learningchallenges.html.
[34] Mitchell, W., McClain, M. A. (2010). A collection of $2 D$ elliptic problems for testing adaptive algorithms. NISTIR 7668.
[35] Manzoni, A., Negri, F. (2016). Automatic reduction of PDEs defined on domains with variable shape. MATHICSE technical report, École Polytechnique Fédérale de Lausanne.
[36] Nielsen, M. A. (2015). Neural Networks and Deep Learning. Determination Press.
[37] Negri, F., Manzoni, A., Amsallem, D. (2015). Efficient model reduction of parametrized systems by matrix discrete empirical interpolation. Journal of Computational Physics, 303:431-454.
[38] Persson, P. O. (2002). Implementation offinite-element based Navier-Stokes solver. Massachussets Institue of Technology.
[39] Prud'homme, C., Rovas, D. V., Veroy, K., Machiels, L., Maday, Y., Patera, A. T., Turinici, G. (2002). Reliable realtime solution of parametrized partial differential equations: Reduced-basis output bound methods. Journal of Fluids Engineering, 124(1):70-80.
[40] Quarteroni, A. (2010). Numerical models for differential problems (Vol. 2). New York, NY: Springer Science \& Business Media.
[41] Quarteroni, A., Manzoni, A., Negri, F. (2015). Reduced basis methods for partial differential equations: An introduction (Vol. 92). New York, NY: Springer, 2015.
[42] Rannacher, R. (1999). Finite element methods for the incompressible Navier-Stokes equations. Lecture notes, Institute of Applied Mathematics, University of Heidelberg.
[43] Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. Psychological Review, 65:386-408.
[44] Schmidt, E. (1907). Zur theorie der linearen und nichtlinearen integralgleichungen. I. Teil: Entwicklung willkürlicher funktionen nach systemen vorgeschriebener. Mathematische Annalen, 63:433-476.
[45] Stergiou, C., Siganos, D. (2013). Neural Networks. Retrieved from https://www.doc.ic.ac.uk/~nd/surprise_96/ journal/vol4/cs11/report.html\#Introductiontoneuralnetworks.
[46] Volkwein, S. (2008). Model reduction using proper orthogonal decomposition. Lecture notes, Department of Mathematics, University of Konstanz.


[^0]:    Email addresses: jan.hesthaven@epfl.ch (J.S. Hesthaven), subbiali@phys.ethz.ch (S. Ubbiali)

[^1]:    ${ }^{1}$ Please note that while the output of a single neuron is denoted with the letter $y$, we use the letter $\mathbf{q}$ (bolded) to indicate the overall output of the network. Clearly, for the $j$-th output neuron, its output $y_{j}$ coincides with the corresponding entry of $\mathbf{q}$, i.e., $q_{j}=y_{j}$ for any $j=1, \ldots, M_{O}$.

[^2]:    ${ }^{2}$ Here we uniquely consider networks with the same number of neurons in both the first and the second hidden layer, but this is not required.

