

Reduced Order Modeling for Nonlinear PDE-constrained Optimization using Neural Networks

Nikolaj Takata Mücke, Lasse Hjuler Christiansen, Allan Peter Engsig-Karup, John Bagterp Jørgensen

Abstract—Nonlinear model predictive control (NMPC) often requires real-time solution to optimization problems. However, in cases where the mathematical model is of high dimension in the solution space, e.g. for solution of partial differential equations (PDEs), black-box optimizers are rarely sufficient to get the required online computational speed. In such cases one must resort to customized solvers. This paper presents a new solver for nonlinear time-dependent PDE-constrained optimization problems. It is composed of a sequential quadratic programming (SQP) scheme to solve the PDE-constrained problem in an offline phase, a proper orthogonal decomposition (POD) approach to identify a lower dimensional solution space, and a neural network (NN) for fast online evaluations. The proposed method is showcased on a regularized least-square optimal control problem for the viscous Burgers' equation. It is concluded that significant online speed-up is achieved, compared to conventional methods using SQP and finite elements, at a cost of a prolonged offline phase and reduced accuracy.

I. INTRODUCTION

Nonlinear model predictive control is a well-established technique to solve optimal control problems within science, engineering and other areas [1], [2]. However, in cases where the underlying mathematical model is of high dimension, as is the case for PDEs, the optimization requires substantial computational power and time and cannot meet real-time constraints. In such mathematical models the curse of dimensionality is often a bottleneck [3]. As a consequence, black box optimizers are not sufficient and one must resort to customized solvers. The state of the art customized solvers often rely on preconditioned iterative methods [4], [5], high-performance computing [6] and/or reduced order modeling techniques [7], [8] to improve the online computation time.

The contribution of this paper towards real-time optimization of large-scale processes, is a scheme based on proper orthogonal decomposition (POD), sequential quadratic programming (SQP) and neural networks (NNs), denoted the POD-SQP-NN scheme. This approach is in the field of reduced order modeling. The work is inspired by [9], where a similar approach is presented for forward solution of nonlinear PDEs. The idea is to identify a low dimensional representation of the optimal control function, and then train a neural network to map inputs to this representation. The scheme utilizes a data driven approach to compute the low dimensional representation. One starts by exploring the solution space by computing optimal control functions, using the SQP algorithm [10], and then computes a low dimensional

representation using POD [11]. The idea is closely related to that of principal component analysis for dimensionality reduction of data representations [12]. This approach exploits the often existing low dimensional structures to decrease the required computational effort and increase the accuracy of an NN.

The computation of the low dimensional representation and the training of the NN is kept in the offline stage. Thus, it is only necessary to evaluate the NN for given input in the online stage, which is an inexpensive procedure. This way one circumvents the problem associated with the curse of dimensionality in the online stage by keeping the full problem in the offline stage.

The POD-SQP-NN scheme is tested on a regularized least-square optimal control problem for the viscous Burgers' equation to demonstrate its potential. The viscous Burgers' equation has a nonlinear advection term. When this term is dominating shock waves will appear in the uncontrolled case, which makes it challenging to control [13].

The paper is organized as follows. Section 2 introduces the optimal control problem and the optimality system. Section 3 presents the reduced order modeling framework. Section 4 outlines the use of NNs. Section 5 presents numerical results and conclusions, as well as future prospects, are made in section 6 and 7.

II. OPTIMAL CONTROL OF NONLINEAR PDES

Consider the class of nonlinear time-dependent PDEs of the form

$$\partial_t y - \epsilon \partial_{xx} y + N(y) = f, \quad \text{in } Q \quad (1a)$$

$$y(x, t) = 0, \quad \text{on } \Sigma \quad (1b)$$

$$y(x, 0) = y_0, \quad \text{in } \Omega \quad (1c)$$

where $Q = \Omega \times (0, T)$ is a space-time cylinder, $\Omega \in \mathbb{R}^d$, $d = 1, 2, 3$, $\Sigma = \Gamma \times (0, T)$, $\Gamma = \partial\Omega$, $\epsilon > 0$, $f \in L^\infty(Q)$ is an external forcing, and N is nonlinear operator. The initial value problem (IVP), (1), is a general form of advection-reaction-diffusion equations with linear diffusion. Processes of this type encapsulates a great amount of physical phenomena, such as chemical reactions [4], fluid flows [14], and predator-prey systems [15].

A. The Optimal Control Problem

To control (1) towards a desired state we formulate the PDE-constrained optimization problem:

$$\min_{y,u} J(y,u), \quad (2a)$$

$$\text{s.t. } \partial_t y - \epsilon \partial_{xx} y + N(y) = f + bu, \quad \text{in } Q \quad (2b)$$

$$y(x,t) = 0, \quad \text{on } \Sigma \quad (2c)$$

$$y(x,0) = y_0, \quad \text{in } \Omega \quad (2d)$$

The goal of problem (2) is to determine the optimal control function, $u \in L^2(Q)$, such that the state, y , is as close as possible to a pre-defined desired state, $y_d \in L^2(Q)$. For this purpose a typical choice of the cost functional, J , is given by the tracking-type functional:

$$J(y,u) = \frac{1}{2} \|y - y_d\|_2^2 + \frac{\beta}{2} \|u\|_2^2, \quad (3)$$

where $\|\cdot\|_2$ denotes the standard $L^2(Q)$ norm and $\beta > 0$ is the Tikhonov regularization parameter, ensuring that the optimal control has higher regularity [10]. The control can be restricted to a subdomain, $\Omega_c \subset \Omega$ by choosing $b = b(x,t)$ as a characteristic function, i.e. $b = b(x,t)\chi_{\Omega_c}(x,t)$. Note that in many real-world applications (2) is often accompanied by constraints on the control function, u . However, to illustrate the ideas and principles of this paper we leave out such constraints on u .

B. Optimality System

Solving (2) can be done in more than one way. This paper follows the *optimize-then-discretize* approach [16]. That is, we start by deriving the optimality conditions on the continuous level and then we discretize the resulting system. Utilizing the formal Lagrange method [10], we arrive at the following system:

$$\partial_t y - \epsilon \partial_{xx} y + N(y) - \frac{b^2}{\beta} p = f, \quad \text{in } Q, \quad (4a)$$

$$-\partial_t p - \epsilon \partial_{xx} p + N_y(y)p = y - y_d, \quad \text{in } Q, \quad (4b)$$

$$y(x,t) = p(x,t) = 0, \quad \text{on } \Sigma, \quad (4c)$$

$$y(x,0) = y_0, \quad \text{in } \Omega, \quad (4d)$$

$$p(x,T) = 0, \quad \text{in } \Omega. \quad (4e)$$

where p is the adjoint state. The control is recovered by $u = \frac{b}{\beta} p$. The IVP (4) is a system of nonlinear PDEs with the state evolving forward in time and the adjoint state backwards. We use a Newton/SQP approach to overcome the nonlinearities, which first linearizes (4) and then solve the resulting system until a given stopping criteria is fulfilled [10], [13]. Denoting by (y_i, p_i) , $1 \leq i \leq k$, the Newton iterates and $(y_{k+1}, p_{k+1}) := (y, p)$, the next iterate, we recover the following linear system to solve for (y, p) :

$$\partial_t y - \epsilon \partial_{xx} y + N_y(y_k)y - \frac{b^2}{\beta} p = F_1, \quad (5a)$$

$$-\partial_t p - \epsilon \partial_{xx} p + N_y(y_k)p - y + N_{yy}(y_k)p_k y = F_2, \quad (5b)$$

where

$$F_1 := f + N(y_k) - N_y(y_k)y_k, \quad (6a)$$

$$F_2 := -y_d + N_{yy}(y_k)p_k y_k, \quad (6b)$$

and with initial and boundary conditions as given in (4).

While (5) are linear, the problem of opposite time evolutions of the state and the adjoint state remains. To solve this problem we solve for all time steps at once. This is also referred to as an *all-at-once* approach [16].

Using the Galerkin finite element method (FEM), we aim to represent the solutions by

$$y = \sum_{i=1}^{N_\delta} \hat{y}_i(t) \phi_i(x), \quad p = \sum_{i=1}^{N_\delta} \hat{p}_i(t) \phi_i(x), \quad (7)$$

where $\phi_i \in H^1(\Omega)$ are the finite element basis functions and \hat{y}_i, \hat{p}_i are the generalized Fourier coefficients, i.e. the degrees of freedom, to be estimated. We denote the finite dimensional solution space $V_\delta = \text{span}\{\phi_i\}_{i=1}^{N_\delta}$, introduce a time discretization, $t_i = i\Delta t$, $i = 1 \dots N_t$, and define

$$\mathbf{y} = (\hat{y}_1(t_1), \dots, \hat{y}_1(t_{N_t}), \dots, \hat{y}_{N_\delta}(t_1), \dots, \hat{y}_{N_\delta}(t_{N_t})), \quad (8a)$$

$$\mathbf{p} = (\hat{p}_1(t_1), \dots, \hat{p}_1(t_{N_t}), \dots, \hat{p}_{N_\delta}(t_1), \dots, \hat{p}_{N_\delta}(t_{N_t})), \quad (8b)$$

as the vectors collecting the degrees of freedom for all time steps. Thus, we solve a system of the form

$$\underbrace{\begin{bmatrix} A_k & B_k^T \\ B_k^T & -C_k \end{bmatrix}}_{L_k} \begin{bmatrix} \mathbf{y} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{F}_2 \end{bmatrix}, \quad (9)$$

in each Newton iteration. Note that $L_k \in \mathbb{R}^{2N_\delta N_t}$ can potentially be a large matrix. Thus, solving (9) multiple times may be expensive, which motivates the use of specialized approaches to improve computation time.

III. REDUCED ORDER MODELING

Decreasing the size of (9) is one way to reduce the computational cost. To do so we need an efficient reduced basis, i.e. a basis comprising fewer functions without compromising the accuracy for given set of parameters. Hence, we seek solutions of the form

$$u_{\text{rb}} = \sum_{i=1}^{N_{\text{rb}}} \hat{u}_i(t) \psi_i(x), \quad (10)$$

where $N_{\text{rb}} \ll N_\delta$ and ψ_i is the reduced basis functions. We denote the reduced space $V_{\text{rb}} = \text{span}\{\psi_i\}_{i=1}^{N_{\text{rb}}} \subset V_\delta$. This also motivates naming V_δ the *full-order* or *high-fidelity* space. To derive such a basis, we utilize the POD [11], [17].

A. Parametrized PDEs

To describe the POD method, we begin by introducing the concept of parametrized PDEs. Here, we assume that the solution does not only depend on space and time but also on the parameters. Define the vector, $\boldsymbol{\mu} \in \mathbb{P} \subset \mathbb{R}^P$, to be the vector of all parameters. Such parameters could be

either physical, such as the diffusion rate, characteristics of the initial condition, or geometric, such the shape and size of the domain. The input parameters will often comprise an initial state in optimal control problems. Therefore, it makes sense to decompose the parameter vector into an initial state and other parameters, $\boldsymbol{\mu} = (\mathbf{y}_0, \boldsymbol{\lambda})$.

The set of all solutions for varying parameters makes up a manifold, denoted the *solution manifold*:

$$\mathcal{M} = \{u(\boldsymbol{\mu}) \mid \boldsymbol{\mu} \in \mathbb{P}\}. \quad (11)$$

Similarly, we define \mathcal{M}_δ and \mathcal{M}_{rb} to be the solutions manifold for the high-fidelity and the reduced basis approximations, respectively.

B. Proper Orthogonal Decomposition

The POD method to compute the reduced basis is a data driven approach. One starts by computing a series of *snapshots*. A snapshot is a solution of the full problem for a given set of parameters and time instance. Ideally the snapshots should make up a good representation of the solution manifold. Then we seek to minimize

$$\sqrt{\sum_{i=1}^{N_s} \inf_{v \in V_{\text{rb}}} \|u_\delta^i - v\|_{L^2(\Omega)}^2}, \quad (12)$$

where V_{rb} is the reduced space, $u_\delta^i \in V_\delta$ denotes a snapshot, i.e. the solution for a specific parameter choice and time instance, and $N_s = N_p N_t$ is the number of snapshots, where N_p and N_t are the number of parameter samples and time steps respectively. According to the orthogonal projection theorem the minimizing reduced basis representation of u_δ^i is the orthogonal projection [11]:

$$\sqrt{\sum_{i=1}^{N_s} \left\| u_\delta^i - \sum_{k=1}^{N_{\text{rb}}} \langle u_\delta^i, \psi_j \rangle \psi_j \right\|_{L^2(\Omega)}^2}. \quad (13)$$

In the $l^2(\Omega)$ case one can show that minimizing (13) is equivalent to the constrained maximization problem [18]:

$$\max_{\{\boldsymbol{\psi}_i\}_{i=1}^{N_s}} \sum_{i=1}^{N_s} \sum_{j=1}^{N_\delta} |\mathbf{u}_i^T \boldsymbol{\psi}_j|^2, \quad (14a)$$

$$\text{s.t. } \boldsymbol{\psi}_j^T \boldsymbol{\psi}_i = \delta_{ij}, \quad \forall i, j \quad (14b)$$

where $\boldsymbol{\psi}_j$ is the vector of coefficients of the reduced basis functions, expanded in terms of the FE functions, and \mathbf{y}_i is the degrees of freedom of a snapshot. By collecting all the snapshots in a snapshot matrix, S , and all the reduced basis functions as columns in a matrix, W , the solution to (14) is the solution to the eigenvalue problem [11]:

$$SS^T W = W \Lambda, \quad (15)$$

where Λ is a diagonal matrix with the eigenvalues, corresponding to the columns of W , on the diagonal. The matrix SS^T is denoted the correlation matrix [11]. Truncating the number of eigenvectors gives us the number of reduced basis functions in use. One transforms between the degrees of

freedom, in the high-fidelity basis and the POD basis, by [11]:

$$\mathbf{y}_\delta = W \mathbf{u}_{\text{rb}}, \quad \mathbf{u}_{\text{rb}} = W^T \mathbf{u}_\delta. \quad (16)$$

The projection error on the snapshots corresponds to the eigenvalues of the correlation matrix of the POD modes left out [11]:

$$\sum_{i=1}^{N_s} \left\| \mathbf{u}_i - \underbrace{W W^T \mathbf{u}_i}_{\text{Projection}} \right\|_2^2 = \sum_{i=N_{\text{rb}}+1}^r \lambda_i. \quad (17)$$

Hence, the eigenvectors corresponding to the largest eigenvalues contain the most information about the solution space. Thus, one should choose the reduced basis functions by looking at the size of their respective eigenvalues.

IV. ARTIFICIAL NEURAL NETWORKS

First, we will give a brief recap of the fundamental concepts regarding NNs. An NN can be considered a map, $F : \mathbb{R}^{M_I} \rightarrow \mathbb{R}^{M_O}$, that maps a vector of M_I features to a vector of M_O responses. In the case of a feedforward network, also denoted a perceptron, the map consists of a series of function compositions [19]:

$$F(x; \theta, H, L) = u_{L+1} \circ f_L \circ u_L \circ \dots \circ f_1 \circ u_1(x), \quad (18)$$

where L is the number of hidden layers, H is the number of neurons in each layer, θ is the set of weights and biases, f_i are the activation functions, and u_i are the propagation functions. The evaluation of a NN is called a forward propagation since it can be considered a propagation of the input features through the series of functions. The propagation function is often a linear combination with a bias term:

$$u_i(x) = w_{i-1}x + \beta_i, \quad (19)$$

where w_i 's are matrices containing the weights and β_i 's are vectors containing the bias terms. There are many choices for the activation functions. Recently the functions called rectified linear units (ReLU) have gained much success due to great results. In this paper we make use of the slightly modified version called the leaky ReLU [20]:

$$f(x) = \max(\alpha x, x), \quad (20)$$

where α is a small positive real number.

The motivation for using neural networks are their versatility. It can be shown that any continuous function $\varphi \in \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ can be approximated arbitrarily well by a deep neural network [19]. However, in practice it is difficult to get arbitrary precision due to various problems. These include i) a large parameter space, ii) nonlinearities, and iii) a non-convex optimization problem for determining the parameters.

Much time has gone into developing optimization algorithms for computing the weights and biases to get as close as possible to the theoretical potential. Second order methods can not, in general, be used since they only ensure local convergence and are too expensive in high dimensions. Therefore, stochastic variations of gradient descent algorithms, such as Adam, are often used [21].

V. THE POD-SQP-NN SCHEME

For PDE-constrained optimization problems one can consider a map, that takes input parameters and maps them to the high-fidelity optimal control solution manifold:

$$G : \mathbb{P} \subset \mathbb{R}^P \rightarrow \mathbb{R}^{N_\delta N_t}, \quad (y_0, \boldsymbol{\lambda}) \mapsto G(y_0, \boldsymbol{\lambda}) = \mathbf{u}_\delta. \quad (21)$$

$G(y_0, \boldsymbol{\lambda})$ is computed using the SQP method and finite elements. The image of the mapping, G , corresponds to the high-fidelity solution manifold for the control function. The general idea of the POD-SQP-NN scheme is to approximate this map in the POD space by using NNs. That is, by using the relation in (16) we aim to compute a map, G_{nn} , such that

$$G_{nn}(y_0, \boldsymbol{\lambda}) \approx W^T G(y_0, \boldsymbol{\lambda}), \quad \forall \boldsymbol{\mu} \in \mathbb{P}, \quad (22)$$

from which we recover the optimal control in the high-fidelity space by

$$\mathbf{u}_{nn}^\delta = W G_{nn}(y_0, \boldsymbol{\lambda}). \quad (23)$$

Consequently, the POD-SQP-NN scheme is decomposed into two stages:

- **The offline stage.**

- 1) Compute snapshots, $\{(y_0, \boldsymbol{\lambda})_i, G((y_0, \boldsymbol{\lambda})_i)\}_{i=1}^{N_s}$, using SQP and finite elements.
- 2) Compute POD basis functions and transform snapshots to get training data: $\{(y_0, \boldsymbol{\lambda})_i, W^T G((y_0, \boldsymbol{\lambda})_i)\}_{i=1}^{N_s}$.
- 3) Train the neural network, G_{nn} , on training data.

- **The online stage.**

- 1) Input parameters, $(y_0, \boldsymbol{\lambda})$, and compute $G_{nn}(y_0, \boldsymbol{\lambda})$ by forward propagation in NN.
- 2) Transform output to high-fidelity space, $\mathbf{u}_{nn}^\delta = W G_{nn}(y_0, \boldsymbol{\lambda})$.
- 3) (Optional) Compute the optimal state by a forward solving of the state equation.

This procedure is outlined in Fig. 1. It is apparent that the offline stage is quite extensive. It consists of several steps which are all, by themselves, computationally heavy tasks and are highly dependent on the dimension of the problem. However, the online stage is not computationally demanding, which is essential for real-time applications.

By approximating the solution in the reduced space we reduce the number of output neurons and, thereby, also the number of weights to be trained significantly. This reduces training time as well as online computation time. Furthermore, it increases the accuracy of the NN since it reduces the size of the parameter space.

A typical problem when training neural networks is lack of enough data. In the proposed method, as much data as needed can be generated, given no time constraints on the offline phase. This, however, does not ensure high precision due to other constraints in neural networks such as bias-variance trade-off, irreducible errors, etc. [12]. However, it is worth discussing potential problems regarding the dimension of the input. Due to the curse of dimensionality it might

be infeasible to compute the proposed map G_{nn} since it might require an insurmountable amount of simulations to explore the full solution manifold. However, in many cases it can be exploited that the initial state lies on a low dimensional surface parametrized by only a few parameters, which reduces this problem significantly.

We define Π_{rb} to be the projection operator onto the reduced basis space. By using the triangle inequality, one gets rough a priori error estimate:

$$\begin{aligned} \|u - u_{nn}^\delta\|_{L^2(Q)} &\leq \underbrace{\|u - u_\delta\|_{L^2(Q)}}_{=:\varepsilon_\delta} + \underbrace{\|u_\delta - \Pi_{rb} u_\delta\|_{L^2(Q)}}_{=:\varepsilon_{rb}} \\ &\quad + \underbrace{\|\Pi_{rb} u_\delta - u_{nn}^\delta\|_{L^2(Q)}}_{=:\varepsilon_{nn}}, \end{aligned} \quad (24)$$

ε_δ is the error between the high-fidelity and the exact solution. ε_{rb} is the discrepancy between the high-fidelity solution and its POD basis representation. It depends on the quality of the POD method which, in turn, depends on the how reducible the problem is and how well the solution manifold is sampled. In many cases this error decreases exponentially [17]. Lastly, ε_{nn} is the discrepancy between the POD representation and the NN output. It depends on the quality of the training of the network and is often the most difficult to drive down through training.

VI. NUMERICAL EXAMPLE

As a test example, we consider the viscous Burgers' equation as the state equation. Thus, the PDE-constrained optimization problem is:

$$\min_{u, y} \quad J(y, u), \quad (25a)$$

$$s.t. \quad \partial_t y - \epsilon \partial_{xx} y + y \partial_x y = f + u, \quad \text{in } Q, \quad (25b)$$

$$y(x, t) = 0, \quad \text{on } \Sigma, \quad (25c)$$

$$y(x, 0) = y_0, \quad \text{in } \Omega, \quad (25d)$$

with J defined as in (3) and $f \in L^\infty(Q)$ a known forcing on the system. The well-posedness of the (25) is discussed in [13]. The resulting linear problem to be solved in each Newton iteration is

$$\partial_t y - \epsilon \partial_{xx} y + \partial_x(y_k y) - \frac{1}{\beta} p = f + y_k \partial_x y_k, \quad (26a)$$

$$-\partial_t p - \epsilon \partial_{xx} p - y_k \partial_x p + (\partial_x p_k - 1)y = -y_d + y_k \partial_x p_k, \quad (26b)$$

For the spatial discretizations FEniCS is used [22] and for the NN the PyTorch framework is used [23]. The unconditionally stable implicit Euler is used for the time discretization.

A. Case Study

For the specific case study we choose $\Omega = [0, 1]$, $T = 0.5$ and the initial state:

$$y_0(x) = \begin{cases} h, & \text{for } \omega \leq x \leq 1 - \omega \\ 0 & \text{otherwise} \end{cases}. \quad (27)$$

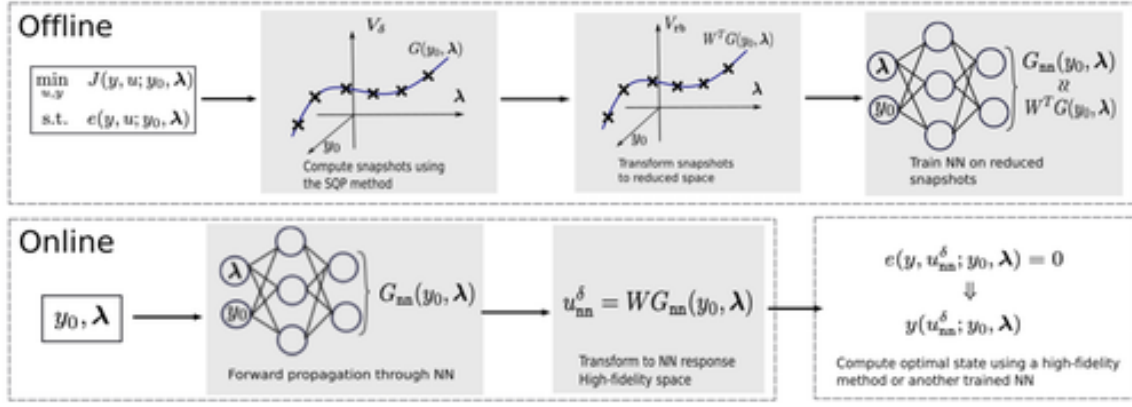


Fig. 1: Workflow for the POD-SQP-NN scheme divided into the offline and online phase. Note the significantly heavier computations in the offline phase in shape of repeated solving by the SQP method and the training of a neural network.

The desired state is to retain the initial condition over time. This problem is quite similar to the one considered in [13]. We define the vector of parameters to be varied as $\mu = (h, \omega, \epsilon) \in [0.5, 1] \times [0.1, 0.3] \times [0.1, 0.001]$. It is worth noting that h and ω parametrizes the initial condition while ϵ defines the diffusion rate. To sample the parameters the Latin Hypercube Sampling method is utilized to ensure a random distribution and that the full parameter space is explored [24].

In this paper, we use $N_p = 300$ and $N_t = 75$ which in total gives 22500 snapshots to compute the POD basis. Note that only 300 SQP all-at-once method solutions are needed to get all the snapshots. It is likely that that fewer snapshots would be sufficient since f is static and $y_d = y_0$.

To train the neural network in the POD-SQP-NN scheme a training set of 2500 optimal control solutions is used. For the architecture we train a feedforward network with three hidden layers with H neurons in each. Furthermore, the Adam optimizer was utilized using mini-batches of size 32 together with early stopping and the mean squared error with L^2 -regularization as the loss function. Furthermore, the weights were initialized according to the standard Gaussian distribution.

The error is measured to be the relative difference of the neural network response and the high-fidelity solution:

$$E = \frac{\|u_{nn}^\delta - u_\delta\|_{L^2}}{\|u_\delta\|_{L^2}}. \quad (28)$$

From Fig. 2 it is clear that the accuracy is increasing with the number of POD basis functions. Furthermore, more neurons in each layer gives a better approximation until a certain point at which the error stagnates. However, it is worth noting that the error is not decreasing monotonically, but is rather noisy. This points to the fact that the networks are not trained well enough in every configuration. In the best cases we reach an accuracy in the order of magnitude 10^{-2} seconds. In Fig. 3 we see the online computation time for the high-fidelity as well as for the POD-SQP-NN scheme. The high-fidelity scheme takes between 10 and 100 seconds while the POD-SQP-NN scheme is of the order of 10^{-3} seconds. This does not come as a surprise, since the online phase is just a

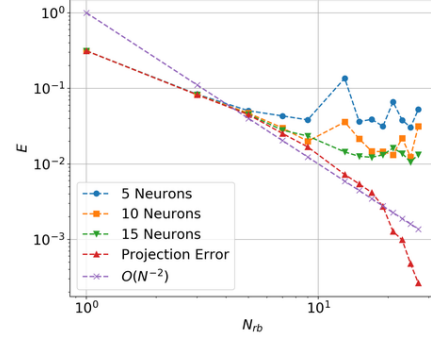


Fig. 2: Convergence of the POD-SQP-NN method for varying number of neurons in each layer.

single forward propagation for the POD-SQP-NN scheme.

Whether an accuracy of 10^{-2} is sufficient is highly dependent on the problem at hand. However, for many problems it is enough. In Figure 4 the relative difference in the objective functional,

$$\frac{\|J(y_{nn}^\delta, u_{nn}^\delta) - J(y_\delta, u_\delta)\|_{L^2}}{\|J(y_\delta, u_\delta)\|_{L^2}}, \quad (29)$$

is plotted. In all online test cases the discrepancy is no larger than 0.12, which suggests that an increase in accuracy will not yield significant changes in the objective functional. Whether such differences are acceptable depends on the problem at hand.

VII. CONCLUSIONS

A new and efficient scheme for online computations of nonlinear PDE-constrained optimization problems is proposed in the paper. We first showed that the problem at hand can be projected onto a lower dimensional manifold using the POD method, which lays the foundation for the POD-SQP-NN method. The scheme significantly decreases the online computation times for the case involving the viscous Burgers' equation. It did, however, come with a decrease in accuracy and a prolonged offline phase. Regarding the




Fig. 3: Execution time in seconds.




Fig. 4: Relative difference in objective function between the high-fidelity and the POD-SQP-NN solution.

decrease in accuracy it is possible that it does not pose any practical limitations due to the minor discrepancies in the objective functions. For these reasons, we see a potential in this method to be used to solve problems that are currently intractable in real-time.

VIII. FUTURE PROSPECTS

What was not discussed in sufficient detail is how the method performs on very high dimensional problems. In many NMPC applications it is not necessarily possible to parametrize the initial conditions by a small set of parameters, as was done here. In such cases the entire initial state must be the input, which could potentially cause problems. Furthermore, the reliability is not addressed as no rigorous a posteriori error bounds exists. This was, however, not a problem in the examples given.

In ongoing work, we will, furthermore, explore different NN architectures such as residual networks, and reinforcement learning which has shown to be well-suited for time dependent data. Alternatively, instead of using NN other new approaches can be utilized to address the curse of dimensionality, such as spectral tensor train decompositions [25].

REFERENCES

- [1] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control engineering practice*, vol. 11, no. 7, pp. 733–764, 2003.
- [2] E. F. Camacho and C. Bordons, "Nonlinear model predictive control: An introductory review," in *Assessment and future directions of nonlinear model predictive control*. Springer, 2007, pp. 1–16.
- [3] R. E. Bellman, *Adaptive control processes: a guided tour*. Princeton university press, 2015, vol. 2045.
- [4] L. H. Christiansen and J. B. Jørgensen, "A new lagrange-newton-krylov solver for pde-constrained nonlinear model predictive control," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 325–330, 2018.
- [5] J. W. Pearson and M. Stoll, "Fast iterative solution of reaction-diffusion control problems arising from chemical processes," *SIAM Journal on Scientific Computing*, vol. 35, no. 5, pp. B987–B1009, 2013.
- [6] G. Biros and O. Ghattas, "Parallel lagrange-newton-krylov-schur methods for pde-constrained optimization. part i: The krylov-schur solver," *SIAM Journal on Scientific Computing*, vol. 27, no. 2, pp. 687–713, 2005.
- [7] B. Haasdonk, "Reduced basis methods for parametrized pdes—a tutorial introduction for stationary and instationary problems," *Model reduction and approximation: theory and algorithms*, vol. 15, p. 65, 2017.
- [8] L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. van Bloemen Waanders, "Large-scale pde-constrained optimization: an introduction," in *Large-Scale PDE-Constrained Optimization*. Springer, 2003, pp. 3–13.
- [9] J. S. Hesthaven and S. Ubbiali, "Non-intrusive reduced order modeling of nonlinear problems using neural networks," *Journal of Computational Physics*, vol. 363, pp. 55–78, 2018.
- [10] F. Tröltzsch, *Optimal control of partial differential equations: theory, methods, and applications*. American Mathematical Soc., 2010, vol. 112.
- [11] A. Quarteroni, A. Manzoni, and F. Negri, *Reduced basis methods for partial differential equations: an introduction*. Springer, 2015, vol. 92.
- [12] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*. Springer series in statistics New York, NY, USA., 2001, vol. 1, no. 10.
- [13] F. Tröltzsch and S. Volkwein, "The sqp method for control constrained optimal control of the burgers equation," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 6, pp. 649–674, 2001.
- [14] D. Kuzmin and J. Hämäläinen, "Finite element methods for computational fluid dynamics: a practical guide," *SIAM Rev*, vol. 57, no. 4, p. 642, 2015.
- [15] B. Ainseba, M. Bendahmane, and A. Noussair, "A reaction-diffusion system modeling predator-prey with prey-taxis," *Nonlinear Analysis: Real World Applications*, vol. 9, no. 5, pp. 2086–2105, 2008.
- [16] F. Yilmaz and B. Karasözen, "An all-at-once approach for the optimal control of the unsteady burgers equation," *Journal of Computational and Applied Mathematics*, vol. 259, pp. 771–779, 2014.
- [17] J. S. Hesthaven, G. Rozza, B. Stamm, et al., *Certified reduced basis methods for parametrized partial differential equations*. Springer, 2016.
- [18] M. Gubisch and S. Volkwein, "Proper orthogonal decomposition for linear-quadratic optimal control," *Model reduction and approximation: theory and algorithms*, vol. 15, p. 1, 2017.
- [19] M. M. Wolf, "Mathematical foundations of supervised learning," July 2018.
- [20] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. icml*, vol. 30, no. 1, 2013, p. 3.
- [21] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [22] M. S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells, "The fenics project version 1.5," *Archive of Numerical Software*, vol. 3, no. 100, 2015.
- [23] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [24] W. G. Cochran, *Sampling techniques*. John Wiley & Sons, 2007.
- [25] D. Bigoni, A. P. Engsig-Karup, and Y. M. Marzouk, "Spectral tensor-train decomposition," *SIAM Journal on Scientific Computing*, vol. 38, no. 4, pp. A2405–A2439, 2016.