

Deep polytopic autoencoders for low-dimensional linear parameter-varying approximations and nonlinear feedback design

Jan Heiland* Yongho Kim[†] Steffen W. R. Werner[‡]

*Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstraße 1, 39106 Magdeburg, Germany.

Email: heiland@mpi-magdeburg.mpg.de, ORCID: 0000-0003-0228-8522

Otto von Guericke University Magdeburg, Faculty of Mathematics, Universitätsplatz 2, 39106 Magdeburg, Germany.

Email: jan.heiland@ovgu.de

[†]Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstraße 1, 39106 Magdeburg, Germany.

Email: ykim@mpi-magdeburg.mpg.de, ORCID: 0000-0003-4181-7968

[‡]Department of Mathematics and Division of Computational Modeling and Data Analytics, Academy of Data Science, Virginia Tech, Blacksburg, VA 24061, USA.

Email: steffen.werner@vt.edu, ORCID: 0000-0003-1667-4862

Abstract: Polytopic autoencoders provide low-dimensional parametrizations of states in a polytope. For nonlinear PDEs, this is readily applied to low-dimensional linear parameter-varying (LPV) approximations as they have been exploited for efficient nonlinear controller design via series expansions of the solution to the state-dependent Riccati equation. In this work, we develop a polytopic autoencoder for control applications and show how it outperforms standard linear approaches in view of LPV approximations of nonlinear systems and how the particular architecture enables higher order series expansions at little extra computational effort. We illustrate the properties and potentials of this approach to computational nonlinear controller design for large-scale systems with a thorough numerical study.

1 Introduction

The design of controllers for nonlinear high-dimensional systems is a challenging task and there exists no generally applicable and computationally feasible approach. Here, we consider input-affine systems of the form

$$\dot{v}(t) = f(v(t)) + Bu(t), \quad y(t) = Cv(t), \quad (1)$$

where for time $t > 0$, $v(t) \in \mathbb{R}^n$ denotes the state, $u(t) \in \mathbb{R}^m$ and $y(t) \in \mathbb{R}^\ell$ denote the system's input and output, $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a possibly nonlinear function, and B and C are linear input and output operators.

For the control of systems like (1), methods like backstepping [17], feedback linearization [19, Chap. 5.3] and sliding mode control [12] require structural assumptions whereas general approaches based on the Hamilton-Jacobi-Bellman (HJB) equations are only feasible for very moderate system sizes; see, e.g., [8]. The commonly employed state-

dependent Riccati equation (SDRE) approximation to the HJB (see, e.g., [10] for a survey) lifts a significant portion of the computational complexity. Nonetheless, this approach became available for large-scale systems only in recently due to further simplifications; see [6] for a linear update scheme and [1] for approximations through series expansions.

In this work, we expand on the series expansion of the SDRE combined with the approximation of the nonlinear system through a low-dimensional linear parameter-varying (LPV) system. This allows us to provide a generic algorithm to nonlinear controller design; see our previous works [11, 16] and related discussions in [1, 2]. In particular, the results in [16] indicate a measurable improvement of the nonlinear feedback design over a related, well-tuned and robust LQR controller but left space for further performance improvements, for example, through

1. higher order expansions of the nonlinear feedback or

2. nonlinear parametrizations of the velocity states.

Both topics are linked to the question if a nonlinear parametrization can provide better low-dimensional approximations with fewer degrees of freedom and thus mitigating the natural limit, which is commonly expressed through the Kolmogorov n -width of any linear model order reduction scheme – like *proper orthogonal decomposition* (POD) which is a linearly optimal projection scheme; [13]. There have been many efforts and positive results concerning the use of nonlinear model order reduction approaches using, e.g., polynomial manifold maps [9], learning techniques [13], or the selection of local bases [3].

To consider higher order expansions of a nonlinear feedback law, a very low-dimensional parametrization is crucial. Among others, as we will lay out, the numbers of coefficient that need to be precomputed through large-scale matrix equations equals $\sum_{k=0}^p \binom{r+k-1}{k}$, where r is the dimension of the parametrization and p is the order of the expansion. For example, for $r = 10$ and a first-order expansion, the solve of $1+10 = 11$ matrix equations was required, cp. [16], whereas a second or third-order expansion would require 66 or 286 solves. However, for a parametrization of size $r = 5$, only 6, 21, or 56 matrix equation solves are needed for expansions of first, second and third order, respectively.

In this work, we rely on our recent development of polytopic autoencoders [14] that combines learning approaches with the selection of local bases in terms of polytope vertices. The concept of reconstructing the states in a polytope provides certain advantages over the reconstruction in a linear space. Generally, the observed states of a dynamical system are described in a bounded set rather than in an unbounded space. With the polytopic approach, the reconstruction is confined to a region where the states have been observed. This confinement stands against the freedom to explore the full state space, for controller design, however, it can be beneficial if tailored towards target regimes and since the resulting series expressions will be bounded by design.

The underlying ideas and major findings of this paper can be summarized as follows: We recall in Section 2 that a general nonlinear control system can be approximated by a low-dimensional LPV system, which can be exploited for controller design in particular if the parameter dependency itself is linear. In view of providing a performant but maximally low-dimensional parametrization, we resort to clustered polytopic autoencoders that encode the states of a dynamical system in a polytope and decode them with locally adapted bases; see Section 3. An interesting result in view of controller design is that this nonlinear decoding has a Taylor expansion with vanishing second-order terms (Lemma 1). For the actual controller design in Section 4, we

consider series expansions of state-dependent Riccati equation as it has been proposed and exercised before. Here, we extend the formulas to consider nonlinear parameter dependencies and higher order terms in the expansion. Finally, we report numerical experiments that illustrate the performance improvement of the polytopic autoencoder over standard approximation schemes, that explore the different levels of approximation to the underlying nonlinear SDRE feedback design, and that investigate the gain in controller performance for a commonly considered flow control problem; see Section 6.

2 Low-dimensional LPV approximations

For nonlinear systems like (1), we consider quasi-LPV approximations of the form

$$\dot{v}(t) = A(\rho(v(t))) v(t) + Bu(t), \quad y(t) = Cv(t). \quad (2)$$

Such systems can be derived from a state-dependent coefficient representation such as $f(v) = \tilde{A}(v)v$ and a parametrization of the state $v(t) \approx \tilde{v}(t) = \nu(\mu(v(t)))$, where μ is an encoder and ν is a decoder. In this case, we set $\rho = \mu(v)$ and $A(\rho) := \tilde{A}(\nu(\rho))$. Note that the parameter depends on the state, hence the term *quasi-LPV*.

Remark 1. In the example application we consider later, namely flows, we will work with coefficients \tilde{A} that are linear in the variable so that linear structures in the decoding $\rho \rightarrow \tilde{v}$ will lead to linear parameter dependencies as they have been proven beneficial for controller design; see, e.g., [4].

3 Polytopic parametrizations in clusters for controller design

As opposed to linear model order reduction approaches, where the state is parametrized in a linear subspace, in polytopic autoencoders (PAEs), the state is represented through coordinates with respect to vertices of an R -dimensional polytope. As a polytope is bounded, without loss of generality, the vertices \tilde{v}_k can be scaled and shifted such that the coordinates $\theta(t) \in \mathbb{R}^R$ that represent the reconstruction $\tilde{v}(t)$ in the polytope will be positive and sum up to one, i.e.,

$$\tilde{v}(t) = \sum_{k=1}^R \theta_k(t) \tilde{v}_k, \quad \text{with } \theta(t) \geq 0, \quad \sum_{k=1}^R \theta_k(t) = 1.$$

This design is not capable of exactly parametrizing a target value $v^* = 0$ unless the support vectors \tilde{v}_k come with

a specific linear dependency. To explicitly enable the reconstruction of the zero target value and to ensure that the series expansion of the SDRE is about $\theta(\tilde{v}) = 0$, we include $\tilde{v}_0 = 0$ in the reconstruction basis of the training, where we impose the positivity and summation constraint on the θ variable extended by θ_0 , i.e.,

$$\theta_0(t) + \sum_{k=1}^R \theta_k(t) = 1. \quad (3)$$

This way, the tuple $(1, \theta(t)) = (1, 0) \in \mathbb{R}^{R+1}$ will be a feasible value in the extended parameter space that represents the target value.

Following previous results that showed how clustering, i.e., how the choice of local bases can improve the reconstruction, we combine the polytopic parametrization with smooth clustering networks as we have developed it in [14]. In short, we will use coordinates $\theta = c(\rho) \otimes \rho$, where $\rho \in \mathbb{R}^r$ denotes the actual latent variable that encodes v and where c is a smooth clustering network that blends between q sets of local bases for the reconstruction.

Taking into consideration a zero target value, we propose a polytopic autoencoder that consists of three components:

1. A nonlinear encoder $\mu: \mathbb{R}^n \rightarrow \mathbb{R}^{r+1}$ so that

$$(\rho_0, \rho) = \mu(v) =: \text{softmax}(\mathbb{F}_e(v) + e_1),$$

with a scalar function ρ_0 and $\rho(t) \in \mathbb{R}^r$, where the function $\mathbb{F}_e: \mathbb{R}^n \rightarrow \mathbb{R}^{r+1}$ is a trainable nonlinear convolutional neural network with $\mathbb{F}_e(0) = 0$ and where $e_1 \in \mathbb{R}^{r+1}$ is the first unit vector. Since $\mathbb{F}_e(0) = 0$ and with the *softmax function*

$$\text{softmax}: x \rightarrow \left(\frac{x_i \tanh(ax_i)}{\sum_{j=1}^{r+1} x_j \tanh(ax_j)} \right)_i^{r+1},$$

for a parameter $a > 0$ (that we will set to $a = 10$ in the experiments), we ensure that the output $(\rho_0(t), \rho(t))$ satisfies

$$(\rho_0(t), \rho_k(t)) \geq 0 \quad \text{and} \quad \rho_0(t) + \sum_{k=1}^r \rho_k(t) = 1;$$

in particular $\mu(0) = (1, 0)$.

2. A smooth clustering network $c: \mathbb{R}^r \rightarrow \mathbb{R}^q$ so that

$$c(\rho) = \text{softmax}(\mathbb{F}_c(\rho) + \hat{e}_1), \quad (4)$$

where \mathbb{F}_c is a trainable neural network with $\mathbb{F}_c(0) = 0$ and where \hat{e}_1 is the first unit vector in \mathbb{R}^q . By design, for the output $c(\rho)$, we have that

$$c(\rho(t)) \geq 0 \quad \text{and} \quad \sum_{i=1}^q c(\rho(t))_i = 1$$

hold, so that $(\theta_0, \theta) := (\rho_0, c(\rho) \otimes \rho)$ satisfies the positivity and summation conditions (3).

3. A trainable linear map $W \in \mathbb{R}^{n \times qr}$ that realizes the reconstruction as

$$\tilde{v} = W(c(\rho) \otimes \rho). \quad (5)$$

We lay out the design and motivation in the following remarks:

- The property that $\mathbb{F}_e(0) = 0$ and $\mathbb{F}_c(0) = 0$ is achieved by using (convolutional) nonlinear layers without bias terms and with activation functions that cross the origin.
- The use of the extra ρ_0 variable ensures that ρ is zero at the origin despite being valid coordinates in a polytope.
- The design of the clustering network ensures $c(0) = e_1$ so that values of ρ close to zero are associated with the first cluster and that $c(\mu(0)) \otimes \mu(0) = (1, 0) \in \mathbb{R}^{qr}$ represents the target vector $v^* = 0$ exactly.
- With using the Kronecker product $c(\rho) \otimes \rho$, the k -th r columns of W denote the supporting basis for reconstruction in the k -th cluster; see [14].

Lemma 1. *With softmax: $x \rightarrow \left(\frac{x_i \tanh(ax_i)}{\sum_{j=1}^q x_j \tanh(ax_j)} \right)_i^q$, it holds that the Jacobian c' of the clustering network c as defined in (4) is zero at $\rho = 0$. In particular, the second-order Taylor expansion of $\rho \rightarrow A(\rho)$ coincides with the first-order expansion and reads*

$$A(\rho) \approx A(0) + \sum_j^r \rho_j \tilde{A}(w_j), \quad (6)$$

where w_j are the first r columns of $W \in \mathbb{R}^{n \times qr}$ as in (5).

Proof. By direct computations we find $c'(\mathbb{F}_c(\rho) + \hat{e}_1)|_{\rho=0} = c'(\hat{e}_1) = 0 \in \mathbb{R}^{q \times r}$. Since the coefficient map $c(\rho) \otimes \rho \rightarrow \tilde{A}(\tilde{v}(\rho))$ is linear, we can reduce the argument to the non-linear part $\rho \rightarrow c(\rho) \otimes \rho$. With $\rho = 0$, from

$$c(\rho + h) \otimes (\rho + h) - c(\rho) \otimes \rho = c(0 + h) \otimes h,$$

we derive the first and second-order expansion terms of the map $\rho \rightarrow c(\rho) \otimes \rho$ in the increment h at $\rho = 0$ as

$$c(0 + h) \otimes h = (c(0) + c'(0)h + o(|h|^2)) \otimes h$$

and as

$$c(\rho) \otimes \rho \approx c(0) \otimes \rho + (c'(0)\rho) \otimes \rho = c(0) \otimes \rho + (c'(0) \otimes I)(\rho \otimes \rho),$$

where $c'(0) \in \mathbb{R}^{q \times r}$ is the Jacobian of c at $\rho = 0$. With $c'(0) = 0$, the second-order terms vanish so that with $c(0) = \hat{e}_1$, the first (and second-order) approximation to the polytopic reconstruction reads

$$\tilde{v}(\rho) \approx W(\hat{e}_1 \otimes \rho) = \sum_{j=1}^r \rho_j w_j. \quad (7)$$

Therefore, formula (6) follows from the linearity of \tilde{A} . \square

4 Higher order series expansions of LPV approximations in SDRE feedback design

For the quasi-LPV system (2), the nonlinear SDRE feedback design defines the input u as

$$u(t) = -B^\top P(\rho(v(t)))v(t), \quad (8)$$

where $P(\rho)$ solves the SDRE

$$A(\rho)^\top P(\rho) + P(\rho)A(\rho) - P(\rho)BB^\top P(\rho) = -C^\top C. \quad (9)$$

Since the repeated solve of (9) is costly, series expansions of the feedback law (8) have been proposed.

We briefly derive the formulas for higher order expansions with nonlinear parameter dependencies in A . For linear parameter dependencies, the first-order approximation has been developed in [5] for a single parameter dependency and has been extended in [1] to the multivariate case.

If the state dependency of the SDRE solution is parameterized through ρ , the multivariate Taylor expansion of P about $\rho_0 = 0$ up to order p reads

$$P(\rho) \approx P_0 + \sum_{1 \leq |\alpha| \leq p} \rho^{(\alpha)} P_\alpha,$$

where $\alpha = (\alpha_1, \dots, \alpha_r) \in \mathbb{N}^r$ is a multiindex with $|\alpha| := \sum_{i=1}^r \alpha_i$, where $\rho^{(\alpha)} := \rho_1^{\alpha_1} \rho_2^{\alpha_2} \dots \rho_r^{\alpha_r}$, and where, importantly, P_α are constant; see [16] for more details.

For nonlinear parameter dependencies, we will consider a truncated series expansion of A :

$$A(\rho) \approx A_0 + \sum_{1 \leq |\alpha| \leq p} \rho^{(\alpha)} A_\alpha, \quad (10)$$

where A_α are constant matrices; see (6) that describes the first and second order expansion in the considered case with the structure of the nonlinear decoding as $\rho \rightarrow W(c(\rho) \otimes \rho)$.

We insert the expansions of A and P into the SDRE (9) and solve for the coefficients P_α by matching the individual

powers of ρ . For $\alpha = 0$, we obtain the standard Riccati equation

$$A_0^\top P_0 + P_0 A_0 - P_0 B B^\top P_0 = -C^\top C, \quad (11)$$

whereas for $|\alpha| = 1$, the coefficients P_α are obtained from

$$P_0 \left(\sum_{|\alpha|=1} \rho^{(\alpha)} A_\alpha \right) + \left(\sum_{|\alpha|=1} \rho^{(\alpha)} A_\alpha \right)^\top P_0 - \sum_{|\alpha|=1} \rho^{(\alpha)} P_\alpha B B^\top P_0 - P_0 B B^\top \sum_{|\alpha|=1} \rho^{(\alpha)} P_\alpha = 0$$

for arbitrary values of ρ . This amounts for the P_α to be given via the Lyapunov equations

$$(A_0 - B B^\top P_0)^\top P_\alpha + P_\alpha (A_0 - B B^\top P_0) = -A_\alpha^\top P_0 - P_0 A_\alpha, \quad (12)$$

for the r multiindices in $\alpha \in \mathbb{N}^r$ for which $|\alpha| = 1$; see also [1, 16] where the same formulas have been derived. Similarly, we may collect the terms in which ρ appears quadratic. Let α^* be a multiindex with $|\alpha^*| = 2$ and let β^* and δ^* be those two multiindices with $|\beta^*| = |\delta^*| = 1$ that satisfy $\beta^* + \delta^* = \alpha^*$. Then, the coefficient P_{α^*} that cancels the corresponding contributions in the equation is defined through the Lyapunov equation

$$(A_0 - B B^\top P_0)^\top P_{\alpha^*} + P_{\alpha^*} (A_0 - B B^\top P_0) = -P_0 A_{\alpha^*} - A_{\alpha^*}^\top P_0 - P_{\beta^*} A_{\delta^*} - A_{\delta^*}^\top P_{\beta^*} - P_{\delta^*} A_{\beta^*} - A_{\beta^*}^\top P_{\delta^*} + P_{\beta^*} B B^\top P_{\delta^*} + P_{\delta^*} B B^\top P_{\beta^*}. \quad (13)$$

Note that for given r and p , there exist exactly $\binom{r+p-1}{p}$ unique such multiindices α with the absolute value of $|\alpha| = p$. For such a series expansion of P , the feedback law (8) reads

$$u(t) = -B^\top \left(\sum_{|\alpha| \leq p} \rho^{(\alpha)} P_\alpha \right) v(t) =: u_{P_{|\alpha| \leq p}}(t), \quad (14)$$

which with the precomputed $B^\top P_\alpha$ is efficiently evaluated by adapting the values of ρ . We note that for $p = 0$ this formula reduces to the standard *linear quadratic regulator* (LQR) feedback with respect to a linearization about $v^* = 0$.

5 Solving the high-dimensional matrix equations

To solve the occurring high-dimensional linear and nonlinear matrix equations (11)–(13), some reformulations and state-of-the-art approaches are needed. Assume that the dimensions of the input and output operators m, ℓ are much

smaller than the state-space dimension n and that A_0 in (10) is sparsely populated. The first equation (11) is the standard factorized Riccati equation, which under the assumptions above can be efficiently solved for a low-dimensional approximation of the solution $Z_0 Z_0^T \approx P_0$, with $Z_0 \in \mathbb{R}^{n \times k}$ and $k \ll n$, e.g., via the low-rank Newton-Kleinman iteration [7]. For the two types of Lyapunov equations (12) and (13), some reformulations are needed to efficiently compute the solutions in the large-scale sparse setting; namely, the right-hand sides need to be reformulated into symmetric low-rank factorizations. In the case of (12), given the approximation $Z_0 Z_0^T \approx P_0$, the right-hand side can be rewritten as

$$-A_\alpha^T P_0 - P_0 A_\alpha \approx -[A_\alpha^T Z_0 \quad Z_0] J \begin{bmatrix} Z_0^T A_\alpha \\ Z_0^T \end{bmatrix},$$

where $J = \begin{bmatrix} 0 & I_k \\ I_k & 0 \end{bmatrix}$. For Lyapunov equations (12) with such symmetric factorized right-hand sides, the ADI approach has been extended to compute appropriate solution approximations of the form $L_\alpha D_\alpha L_\alpha^T \approx P_\alpha$; see [18]. Similarly, we can derive for the right-hand side of (13) low-rank factorizations of the form LDL^T , where

$$L^T = \begin{bmatrix} Z_0^T A_{\alpha^*} \\ Z_0^T \\ L_{\beta^*}^T (A_{\delta^*} - BK_{\delta^*}) \\ L_{\delta^*}^T (A_{\beta^*} - BK_{\beta^*}) \\ L_{\beta^*}^T \\ L_{\delta^*}^T \end{bmatrix}, D = \begin{bmatrix} J & 0 & 0 \\ 0 & 0 & D_{\beta^*, \delta^*} \\ 0 & D_{\beta^*, \delta^*} & 0 \end{bmatrix},$$

with $D_{\beta^*, \delta^*} = \begin{bmatrix} D_{\beta^*} & 0 \\ 0 & D_{\delta^*} \end{bmatrix}$, the low-rank factorizations $L_{\beta^*} D_{\beta^*} L_{\beta^*}^T \approx P_{\beta^*}$ and $L_{\delta^*} D_{\delta^*} L_{\delta^*}^T \approx P_{\delta^*}$, and the feedback matrices $K_{\beta^*} = B^T P_{\beta^*}$ and $K_{\delta^*} = B^T P_{\delta^*}$.

6 Numerical experiments

The code, data and results of the numerical experiments presented here can be found at [15].

We consider the same problem, numerical setup, and the same data points for training the models as in [16]. In short, the flow control problem of stabilizing the unstable steady-state of the flow behind a cylinder is modeled through the difference v of the velocity to the steady-stated and observed through outputs $y = Cv$. It is determined by the semi-discrete Navier-Stokes equations

$$M\dot{v} = N(v)v + Bu, \quad y = Cv, \quad v(0) = 0 \in \mathbb{R}^n, \quad (15)$$

written in so-called *divergence-free* coordinates of the difference state v , with $N: v(t) \rightarrow N(v(t)) \in \mathbb{R}^{n \times n}$ modeling the




Figure 1: Grid search for the selection of r and q with lowest averaged reconstruction error.

Table 1: Model complexity of the different PAE schemes and POD: # stands for *number of*. Further, we distinguish between linear (lin.) and nonlinear (nonl.) mappings in the en-/decoding.

Hyperparameters \ Scheme	POD-5	PAE-1.5	PAE-3.5	PAE-3.5 (p=1)
reduced dimension r	5	5	5	5
#clusters q	1	1	3	3
#encoding parameters	255970	42810	42810	42810
#encoding layers	1(lin.)	16(nonl.)	16(nonl.)	16(nonl.)
#decoding parameters	255970	255970	15+767910	255970
#decoding layers	1	1	1(nonl.)+1(lin.)	1

discrete diffusion-convection operator and $C \in \mathbb{R}^{6 \times n}$ and $B \in \mathbb{R}^{n \times 2}$ the discrete observation and control operators; see [16] for details. In the considered setup, the characteristic *Reynolds number* is set to 60 and the discretization results in $n = 51\,194$; see Figure 2 for an illustration. For the time integration, a semi-explicit Euler scheme is employed, which treats the feedback and the nonlinearity explicitly.

We consider the steady-state solution as target state and aim for the stabilization of the difference system (15). In order to trigger instabilities and for the collection of training data, a test input

$$u(t) = [\sin(t) \quad 0]^T, \quad (16)$$

is applied. As described in [14], the chosen deep convolutional encoder leverages depthwise separable convolutions along with a sparse interpolation matrix to mitigate computational complexity and memory requirements in this rather high-dimensional setup. As a result the model complexity in terms of parameters that define the encoder and decoder can be lower than for, say, a comparable POD reduction; see




Figure 2: Snapshots of the flow simulation setup.

Table 2: Average of the reconstruction errors (scaled with 10^3) of POD-r and PAE-q.r on the time range $[0, 0.5]$ used for training.

Scheme \ r	2	3	4	5	6	7	8
POD-r	1.340	0.610	0.305	0.149	0.076	0.038	0.020
PAE-1.r	1.335	0.614	0.326	0.181	0.118	0.121	0.087
PAE-3.r	0.189	0.135	0.104	0.066	0.075	0.097	0.079
PAE-3.r(p=1)	17.670	7.501	5.217	5.628	2.745	6.794	2.307

Table 1. For training the PAEs, we adopt a semi-supervised learning methodology with a combined loss function

$$\mathcal{L}(v^{(i)}) := \lambda \|\tilde{v}^{(i)} - v^{(i)}\|_M - l(\rho^{(i)}) \cdot \log(c(\rho^{(i)})),$$

where $v^{(i)}$ is the i -th data point, $\rho^{(i)} := \mu(v^{(i)})$ and $\tilde{v}^{(i)} := \nu(\rho^{(i)})$ for the corresponding encoder and decoder μ and ν , where l assigns labels that were precomputed by k -means clustering in \mathbb{R}^r , and where $\lambda > 0$ is a weighting factor between the reconstruction error (measured in the M -weighted discrete L^2 -norm) and the clustering error. We note that the chosen *cross-entropy* loss is commonly used to smoothly assign labels and that the labeling l is defined such that the target $\rho^* = 0$ is assigned to the first cluster. As it is standard practice, the loss is computed as the average over a *batch* of data points. In our experiments, we consider 401 data points equally distributed in the time interval $[0, 0.5]$ as training data and we use the *Adam optimizer* with a learning rate of 0.005, with 3200 epochs, a batch size of 64 and the loss weight $\lambda = 100$.

6.1 Reconstruction errors

We evaluate the reconstruction performance of the PAE in the averaged M -norm error

$$\frac{1}{N} \sum_{i=1}^N \|v^{(i)} - \tilde{v}^{(i)}\|_M,$$

where $v^{(i)}$ is the i -th data point and $\tilde{v}^{(i)} = \nu(\mu(v^{(i)}))$ for the corresponding encoder and decoder μ and ν . As laid out in Table 2, the PAE-1.r is comparable to the standard linear approach of POD for low dimensions $r = 2, 3, 4$, whereas the nonlinear encoding through the cluster selection significantly improves the reconstruction in low dimensions. As for the series expansion of the control law, only the first-order approximation is relevant, we tabulate this reconstruction error as well. Note that the bad average performance is due to the strong deviation for larger t , i.e. larger values of ρ , whereas for small t , the first-order approximation closely aligns with the nonlinear approach; see Figure 3.




Figure 3: Trajectory with pointwise reconstruction errors $\|v(t^{(i)}) - \tilde{v}(t^{(i)})\|_M$. The improvement of using 3 clusters is clearly visible as is the initial alignment with the nonlinear approach PAE-3.5 (and the subsequent deviation from it) of the first-order expansion.

Before turning to the controller design, we use a grid search to identify a suitable parameter setup for the reduced-order LPV approximation. As shown in Figure 1, the pair $(q, r) = (4, 5)$ achieves the lowest reconstruction error in the considered domain. Nonetheless, we selected $(q, r) = (3, 5)$ as a compromise of both accuracy and complexity.

6.2 SDRE approximation through series expansions

We check the approximation of the SDRE feedback along the example trajectory v generated by the test input (16). For that we compute the true SDRE feedback $u_P(t) = -B^T P(v(t))v(t)$, where P solves the SDRE, and compare to $u_{\tilde{P}}$ and $u_{P_{|\alpha| \leq p}}$ that denote the feedback computed through the LPV approximation in (9) (via PAE-q.5, $q = 1, 3$) and the expansions (14) for $p = 0, 1, 2$, respectively. The data in Figure 4 shows that the LPV approximation works particularly well for $t \leq 0.5$, which is the training regime, and that the first and second-order expansions significantly improve the zero-order approximation, which would be the LQR gain, but with a slight effect of the additional second-order terms.

6.3 Feedback performance

To evaluate the performance of the approximations for the controller design, we simulate the closed loop system with varying the two parameters $\gamma \in \{100, 10, 1, 0.1, 0.001\}$ that stands for the penalization factor of u in the underlying quadratic cost functional and the startup time $t_s > 0$ that marks the time after which the test input is turned off and the feedback is applied. In other words, γ is a parameter that weighs on the magnitude of the control and t_s regulates how far the system is deferred from the target state when the controller is activated. In Figure 5, for different orders $p = 0, 1, 2$ of approximation, we plot the averaged measured

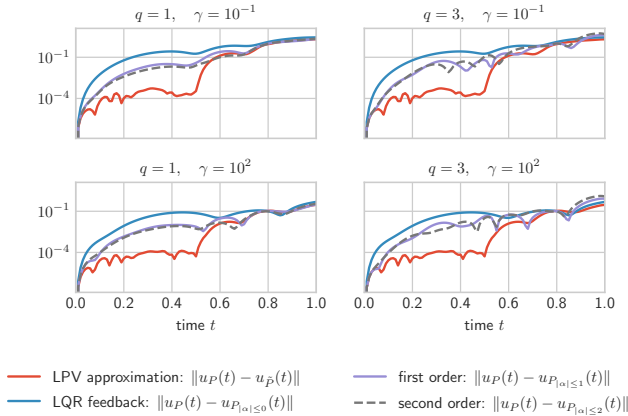


Figure 4: Differences to the exact SDRE feedback along a sample trajectory for feedback approximations through the LPV approximation of the nonlinearity by PAE- q .5, for $q = 1, 3$, and the corresponding zeroth, first and second-order series expansions of the SDRE feedback.

feedback magnitude

$$\frac{1}{t_e} \left(\int_{t_s}^{t_e} \|B^T \left(\sum_{|\alpha| \leq p} \rho(v(t))^{(\alpha)} P_\alpha \right) v(t)\|^2 dt \right)^{\frac{1}{2}}, \quad (17)$$

that tends to 0 for $t_e \rightarrow \infty$ if stabilization is achieved, and that is finite if and only if the simulation has not blown up before a finite end time t_e .

We plot the values of the performance index (17) for $t_e = 7.5$ and for various approximations to the nonlinear SDRE feedback in Figure 5 and the trajectories of selected cases in Figure 6. As we have observed before, the $p = 0$ expansion in (14), i.e., the LQR feedback, is remarkably performant for the optimal choice of the *Tikhonov* parameter γ . However, considering higher order expansions with $p = 1, 2$ together with the option to adapt the underlying approximation scheme, significantly widens the domain of attraction and can ensure reliable performance in particular for larger values of γ ; see, e.g., Figure 6 (top row) for the trajectory data for a setup with $\gamma = 10$, where only the second-order expansion achieves stabilization.

7 Conclusion

In the forward simulation model, the proposed nonlinear autoencoder with smooth clustering significantly improves the parametrization quality at very low dimensions. For the controller design, it was established that for second-order expansions of the SDRE feedback law, only the first-




Figure 5: Performance index (17) map of the linear and nonlinear feedbacks for PAE reduced-order models with $r = 5$, $q = 1, 3$, the SDRE feedback law expansion (14) for $p = 0, 1, 2$ and varying penalization parameter γ and startup times t_s . The darker the color, the smaller (i.e., the better) the performance index. At the blank spaces, the system blew up. The trajectories of the data marked with * are plotted in Figure 6.




Figure 6: Measured output and feedback input signals for the distinguished data points from Figure 5.

order expansions of the reduced-order LPV coefficient are relevant, which reduces the computational effort by large. Concerning the performance of the resulting feedback, the overall picture is diffuse. While the linear LQR approach, for suitable weight parameters, ensured stability also well away from the origin, the nonlinear feedback design based on different approximation schemes, clearly extended the basin of attraction. As the choice of the approximation was solely based on a single open-loop simulation in a time range that did not cover the range considered in the closed-loop simulation, in future work, we expect to gain additional performance by optimizing the approximation in view of feedback control.

Acknowledgments

Jan Heiland and Yongho Kim were supported by the German Research Foundation (DFG) Research Training Group 2297 “Mathematical Complexity Reduction (MathCoRe)”, Magdeburg. Steffen W. R. Werner acknowledges Advanced Research Computing (ARC) at Virginia Tech for providing computational resources and technical support that have contributed to the results reported within this paper.

References

- [1] A. Alla, D. Kalise, and V. Simoncini. State-dependent Riccati equation feedback stabilization for nonlinear PDEs. *Adv. Comput. Math.*, 49(1):9, 2023. doi:10.1007/s10444-022-09998-4.
- [2] A. Alla and A. Pacifico. A POD approach to identify and control PDEs online through state dependent Riccati equations. e-print 2402.08186, arXiv, 2024. Optimization and Control (math.OC). doi:10.48550/arXiv.2402.08186.
- [3] D. Amsallem, M. J. Zahr, and C. Farhat. Nonlinear model order reduction based on local reduced-order bases. *Int. J. Numer. Methods Eng.*, 92(10):891–916, 2012. doi:10.1002/nme.4371.
- [4] P. Apkarian, P. Gahinet, and G. Becker. Self-scheduled H_∞ control of linear parameter-varying systems: a design example. *Automatica*, 31(9):1251–1261, 1995. doi:10.1016/0005-1098(95)00038-X.
- [5] S. C. Beeler, H. T. Tran, and H. T. Banks. Feedback control methodologies for nonlinear systems. *J. Optim. Theory Appl.*, 107(1):1–33, 2000. doi:10.1023/A:1004607114958.
- [6] P. Benner and J. Heiland. Exponential stability and stabilization of extended linearizations via continuous updates of Riccati-based feedback. *Int. J. Robust Nonlinear Control*, 28(4):1218–1232, 2018. doi:10.1002/rnc.3949.
- [7] P. Benner, M. Heinkenschloss, J. Saak, and H. K. Weichelt. An inexact low-rank Newton-ADI method for large-scale algebraic Riccati equations. *Appl. Numer. Math.*, 108:125–142, 2016. doi:10.1016/j.apnum.2016.05.006.
- [8] T. Breiten, K. Kunisch, and L. Pfeiffer. Feedback stabilization of the two-dimensional Navier-Stokes equations by value function approximation. *Appl. Math. Optim.*, 80(3):599–641, 2019. doi:10.1007/s00245-019-09586-x.
- [9] P. Buchfink, S. Glas, and B. Haasdonk. Symplectic model reduction of Hamiltonian systems on nonlinear manifolds and approximation with weakly symplectic autoencoder. *SIAM J. Sci. Comput.*, 45(2):A289–A311, 2023. doi:10.1137/21M1466657.
- [10] T. Çimen. Survey of state-dependent Riccati equation in nonlinear optimal feedback control synthesis. *J. Guid. Control Dyn.*, 35(4):1025–1047, 2012. doi:10.2514/1.55821.
- [11] A. Das and J. Heiland. Low-order linear parameter varying approximations for nonlinear controller design for flows. e-print 2311.05305, arXiv, 2023. Optimization and Control (math.OC). doi:10.48550/arXiv.2311.05305.
- [12] S. J. Dodds. *Feedback Control: Linear, Nonlinear and Robust Techniques and Design with Industrial Applications*, chapter Sliding Mode Control and Its Relatives, pages 705–792. Advanced Textbooks in Control and Signal Processing. Springer, London, 2015. doi:10.1007/978-1-4471-6675-7_10.
- [13] S. Fresca and A. Manzoni. POD-DL-ROM: Enhancing deep learning-based reduced order models for nonlinear parametrized PDEs by proper orthogonal decomposition. *Comput. Methods Appl. Mech. Eng.*, 388:114181, 2022. doi:10.1016/j.cma.2021.114181.
- [14] J. Heiland and Y. Kim. Polytopic autoencoders with smooth clustering for reduced-order modelling of flows. e-print 2401.10620, arXiv, 2024. Machine Learning (cs.LG). doi:10.48550/arXiv.2401.10620.
- [15] J. Heiland, Y. Kim, and S. W. R. Werner. Code, data and results for numerical experiments in “Deep

- polytopic autoencoders for low-dimensional linear parameter-varying approximations and nonlinear feedback design” (version 1.0), March 2024. [doi:10.5281/zenodo.10783695](https://doi.org/10.5281/zenodo.10783695).
- [16] J. Heiland and S. W. R. Werner. Low-complexity linear parameter-varying approximations of incompressible Navier-Stokes equations for truncated state-dependent Riccati feedback. *IEEE Control Syst. Lett.*, 7:3012–3017, 2023. [doi:10.1109/LCSYS.2023.3291231](https://doi.org/10.1109/LCSYS.2023.3291231).
- [17] P. V. Kokotovic. The joy of feedback: nonlinear and adaptive. *IEEE Contr. Syst. Mag.*, 12(3):7–17, 1992. [doi:10.1109/37.165507](https://doi.org/10.1109/37.165507).
- [18] N. Lang, H. Mena, and J. Saak. An LDL^T factorization based ADI algorithm for solving large-scale differential matrix equations. *Proc. Appl. Math. Mech.*, 14(1):827–828, 2014. [doi:10.1002/pamm.201410394](https://doi.org/10.1002/pamm.201410394).
- [19] E. D. Sontag. *Mathematical Control Theory*, volume 6 of *Texts in Applied Mathematics*. Springer, New York, second edition, 1998. [doi:10.1007/978-1-4612-0577-7](https://doi.org/10.1007/978-1-4612-0577-7).