

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

QRnet: Optimal Regulator Design With LQR-Augmented Neural Networks

Permalink

<https://escholarship.org/uc/item/9sk4h56r>

Journal

IEEE Control Systems Letters, 5(4)

ISSN

2475-1456

Authors

Nakamura-Zimmerer, Tenavi

Gong, Qi

Kang, Wei

Publication Date

2021-10-01

DOI

10.1109/lcsys.2020.3034415

Peer reviewed

QRnet: optimal regulator design with LQR-augmented neural networks

Tenavi Nakamura-Zimmerer¹, Qi Gong¹, and Wei Kang²

Abstract—In this paper we propose a new computational method for designing optimal regulators for high-dimensional nonlinear systems. The proposed approach leverages physics-informed machine learning to solve high-dimensional Hamilton-Jacobi-Bellman equations arising in optimal feedback control. Concretely, we augment linear quadratic regulators with neural networks to handle nonlinearities. We train the augmented models on data generated without discretizing the state space, enabling application to high-dimensional problems. We use the proposed method to design a candidate optimal regulator for an unstable Burgers’ equation, and through this example, demonstrate improved robustness and accuracy compared to existing neural network formulations.

Index Terms—Optimal control, Machine learning, Neural networks, Distributed parameter systems

I. INTRODUCTION

WHILE the linear quadratic regulator (LQR) is firmly established as one of the most powerful tools in linear control, the design of optimal regulators for nonlinear systems continues to challenge the control community. The bottleneck in optimal feedback design is the need to solve a Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE). Due to the well-known “curse-of-dimensionality,” this can be extremely difficult for high-dimensional nonlinear systems.

For this reason, there is an extensive literature on methods for approximating solutions of HJB equations. Some key examples include series expansions [1]–[5], level set methods [6], patchy dynamic programming [7], [8], semi-Lagrangian methods [9], [10], method of characteristics and Hopf formula-based algorithms [11], [12], tensor-based methods [13], and polynomial approximation [14]. Unfortunately, many of these methods are limited to moderate dimensions, local solutions, or dynamics with certain algebraic structure.

In recent years, neural networks (NNs) have gained considerable attention as a promising tool for high-dimensional problems since they can avoid the use of spatial grids. Many NN-based methods represent the solution of the HJB equation – called the *value function* – with a NN and minimize the residual of the HJB PDE and boundary conditions at randomly sampled collocation points [15]–[17]. That is, they solve the

HJB PDE in the least-squares sense. [18] propose a method for learning a suboptimal policy and approximate value function locally around some nominal trajectories.

Finally, a number of recent works have demonstrated the potential of data-driven methods for HJB. The core idea is to generate data by solving a number of two-point boundary value problems (BVPs) which describe the characteristics of the value function. These BVPs can be solved independently, without a spatial mesh, and in parallel, thus making the algorithm *causality-free*. This property allows the method to be applied to high-dimensional problems. Given BVP data, one then constructs a model of the value function based on this data. In [19], [20] the value function is calculated with sparse grid interpolation, in [21]–[23] supervised learning is used to train a NN model, and in [24] the value function is approximated by sparse polynomial regression. Lastly, [25] consider a related approach which connects forward-backward stochastic differential equations with the HJB equation in stochastic optimal control.

In this paper we propose a physics-informed machine learning method to solve high-dimensional infinite horizon quadratic regularization problems. This important class of control problems arise in the design controllers for regularization or set-point tracking, and has numerous engineering applications in e.g. aerospace, robotics, chemical process control, and distributed parameter systems. The contributions of the present work are to 1) extend the framework introduced in [21]–[23] to infinite horizon problems; and 2) introduce a NN value function model which includes a quadratic term based on the LQR approximation for the linearized dynamics. The proposed model structure mirrors the form of a series expansion with the NN accounting for higher order terms.

We contend that the proposed model structure, which we call *QRnet*, has the following advantages:

- In common practice, linear and nonlinear parts of the control are often treated separately. On the other hand, the QRnet feedback controller smoothly integrates these components to achieve good performance on large domains while retaining the local robustness of LQR.
- Model training is LQR-initialized: rather than learning from scratch, the NN builds on the scaffolding of the LQR value function. As we show in Section IV-B, this can reduce sensitivity to variations in the data set and weight initialization.

¹Tenavi Nakamura-Zimmerer and Qi Gong are with the Department of Applied Mathematics, Baskin School of Engineering, University of California, Santa Cruz, CA. tenakamu@ucsc.edu, qigong@soe.ucsc.edu

²Wei Kang is with the Department of Applied Mathematics, Naval Postgraduate School, Monterey, CA. wkang@nps.edu

II. PROBLEM SETTING

We consider infinite-horizon nonlinear optimal control problems (OCPs) of the form

$$\begin{cases} \text{minimize}_{\mathbf{u}(\cdot)} & J[\mathbf{u}(\cdot)] = \int_0^\infty \mathcal{L}(\mathbf{x}, \mathbf{u}) dt, \\ \text{subject to} & \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}, \mathbf{u}), \\ & \mathbf{x}(0) = \mathbf{x}_0. \end{cases} \quad (1)$$

Here $\mathbf{x} : [0, \infty) \rightarrow \mathbb{R}^n$ is the state, $\mathbf{u} : [0, \infty) \rightarrow \mathbb{R}^m$ is the control, and $\mathbf{f}(\mathbf{x}, \mathbf{u}) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a Lipschitz continuous vector field. We consider problems with quadratic running cost,

$$\mathcal{L}(\mathbf{x}, \mathbf{u}) = (\mathbf{x} - \bar{\mathbf{x}})^T \mathbf{Q}(\mathbf{x} - \bar{\mathbf{x}}) + (\mathbf{u} - \bar{\mathbf{u}})^T \mathbf{R}(\mathbf{u} - \bar{\mathbf{u}}), \quad (2)$$

where $\mathbf{Q} \in \mathbb{R}^{n \times n}$ is positive semi-definite, $\mathbf{R} \in \mathbb{R}^{m \times m}$ is positive definite, and $\bar{\mathbf{x}} \in \mathbb{R}^n$, $\bar{\mathbf{u}} \in \mathbb{R}^m$ are a (possibly unstable) fixed point of the dynamics such that $\mathbf{f}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) = \mathbf{0}$. This standard cost function is a natural choice for regularization or set-point tracking problems where we want to stabilize the objective state $\bar{\mathbf{x}}$.

Throughout this paper, we assume that the OCP (1) is well-posed, i.e. an optimal feedback control $\mathbf{u}^*(t)$ exists. This ensures that the optimal controlled trajectory, $\mathbf{x}^*(t)$, satisfies $\lim_{t \rightarrow \infty} \mathcal{L}(\mathbf{x}^*(t), \mathbf{u}^*(t)) = 0$. Due to real-time application requirements, we would like to design a closed-loop feedback controller, $\mathbf{u}^*(t) = \mathbf{u}^*(\mathbf{x}(t))$, which can be evaluated online given any measurement of \mathbf{x} .

A. The Hamilton-Jacobi-Bellman equation

Following the standard procedure in optimal control, to compute the optimal feedback control we begin by defining the value function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ as the optimal cost-to-go of (1) starting at the point $\mathbf{x}(0) = \mathbf{x}$. That is,

$$V(\mathbf{x}) := J[\mathbf{u}^*(\cdot)]. \quad (3)$$

It can be shown that the value function is the unique viscosity solution [26] of the steady state HJB PDE,

$$\begin{cases} \min_{\mathbf{u}} \{ \mathcal{L}(\mathbf{x}, \mathbf{u}) + V_{\mathbf{x}}^T(\mathbf{x}) \mathbf{f}(\mathbf{x}, \mathbf{u}) \} = 0, \\ V(\bar{\mathbf{x}}) = 0, \end{cases} \quad (4)$$

where we denote $V_{\mathbf{x}} := \partial V / \partial \mathbf{x}$. If (4) can be solved (in the viscosity sense), then it provides both necessary and sufficient conditions for optimality.

Given the value function $V(\cdot)$, we define the Hamiltonian

$$\mathcal{H}(\mathbf{x}, V_{\mathbf{x}}, \mathbf{u}) := \mathcal{L}(\mathbf{x}, \mathbf{u}) + V_{\mathbf{x}}^T \mathbf{f}(\mathbf{x}, \mathbf{u}). \quad (5)$$

The optimal control satisfies the Hamiltonian minimization condition,

$$\mathbf{u}^*(\mathbf{x}) = \mathbf{u}^*(\mathbf{x}; V_{\mathbf{x}}(\mathbf{x})) = \arg \min_{\mathbf{u}} \mathcal{H}(\mathbf{x}, V_{\mathbf{x}}, \mathbf{u}). \quad (6)$$

Then if we can solve (4), the optimal feedback control is obtained online as the solution of (6).

But as discussed in Section I, solving (4) is extremely challenging. Thus a common approach is to linearize the dynamics about $(\mathbf{x} = \bar{\mathbf{x}}, \mathbf{u} = \bar{\mathbf{u}})$ to obtain the linear system

$$\begin{cases} \frac{d}{dt} (\mathbf{x} - \bar{\mathbf{x}}) \approx \mathbf{A} (\mathbf{x} - \bar{\mathbf{x}}) + \mathbf{B} (\mathbf{u} - \bar{\mathbf{u}}), \\ \mathbf{A} := \frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\bar{\mathbf{x}}, \bar{\mathbf{u}}), \quad \mathbf{B} := \frac{\partial \mathbf{f}}{\partial \mathbf{u}}(\bar{\mathbf{x}}, \bar{\mathbf{u}}). \end{cases} \quad (7)$$

Under the mild conditions that (\mathbf{A}, \mathbf{B}) is controllable and $(\mathbf{A}, \mathbf{Q}^{1/2})$ is observable, the value function of the OCP with the linear dynamics (7) and quadratic cost (2) is

$$V^{\text{LQR}}(\mathbf{x}) = (\mathbf{x} - \bar{\mathbf{x}})^T \mathbf{P}(\mathbf{x} - \bar{\mathbf{x}}), \quad (8)$$

where $\mathbf{P} \in \mathbb{R}^{n \times n}$ is a positive semi-definite matrix satisfying the Riccati equation,

$$\mathbf{Q} + \mathbf{A}^T \mathbf{P} + \mathbf{P} \mathbf{A} - \mathbf{P} \mathbf{B} \mathbf{R}^{-1} \mathbf{B}^T \mathbf{P} = \mathbf{0}. \quad (9)$$

Furthermore, the resulting state feedback controller is linear with constant gain:

$$\mathbf{u}^{\text{LQR}}(\mathbf{x}) = -\mathbf{K} \mathbf{x}, \quad \mathbf{K} = \mathbf{R}^{-1} \mathbf{B}^T \mathbf{P}. \quad (10)$$

This approach has yielded many successful engineering applications, but it is suboptimal and in some cases even fails to stabilize the nonlinear dynamics. For this reason we are interested in modeling the value function $V(\cdot)$ for the full nonlinear dynamics.

B. Pontryagin's Minimum Principle

To make use of (6) for general nonlinear systems, we need an efficient way to approximate the value function and its gradient. Like [19]–[24], rather than solve the full HJB equation (4) directly, we exploit the fact its characteristics evolve according to a two-point BVP, well-known in optimal control as Pontryagin's Minimum Principle (PMP, [27]):

$$\lim_{t_f \rightarrow \infty} \begin{cases} \dot{\mathbf{x}}(t) = \mathcal{H}_{\lambda} = \mathbf{f}(\mathbf{x}, \mathbf{u}^*(\mathbf{x}; \boldsymbol{\lambda})), & \mathbf{x}(0) = \mathbf{x}_0, \\ \dot{\boldsymbol{\lambda}}(t) = -\mathcal{H}_{\mathbf{x}}(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{u}^*(\mathbf{x}; \boldsymbol{\lambda})), & \boldsymbol{\lambda}(t_f) = \mathbf{0}, \\ \dot{v}(t) = -\mathcal{L}(\mathbf{x}, \mathbf{u}^*(\mathbf{x}; \boldsymbol{\lambda})), & v(t_f) = 0. \end{cases} \quad (11)$$

Here $\boldsymbol{\lambda} : [0, \infty) \rightarrow \mathbb{R}^n$ is called the *costate*. The two-point BVP (11) provides a necessary condition for optimality, and if we further assume that the solution is optimal, then along the characteristic $\mathbf{x} = \mathbf{x}^*(t; \mathbf{x}_0)$ we have

$$V(\mathbf{x}) = v(t), \quad V_{\mathbf{x}}(\mathbf{x}) = \boldsymbol{\lambda}(t), \quad \mathbf{u}^*(\mathbf{x}) = \mathbf{u}^*(t). \quad (12)$$

In general, the BVP (11) admits multiple solutions. So while the characteristics of the value function satisfy (11), there may be other solutions to these equations which are sub-optimal and thus not characteristics. In certain problems the characteristics can also intersect, giving rise to non-smooth value functions and difficulties in applying (12).

Optimality of solutions to (11) can be guaranteed under some convexity conditions (see e.g. [28]). For most dynamical systems it is difficult to verify such conditions globally, but we can guarantee optimality locally around an equilibrium point [2]. Addressing the challenge of global optimality is beyond the scope of the present work, so in this paper we assume that solutions of (11) are optimal. Under this assumption, the relationship between PMP and the value function given in (12) holds everywhere.

Note the proposed method can still be applied even when this assumption cannot be verified. In such cases PMP is the prevailing tool for finding *candidate optimal* solutions, and from these the proposed method yields a stabilizing feedback controller which satisfies necessary conditions for optimality.

III. NEURAL NETWORK VALUE FUNCTION MODELING

A. Data generation

Like [22]–[24], we generate data by solving the two-point BVP (11) for a set of randomly sampled initial conditions. Critically, these BVPs can be solved independently without knowledge of nearby solutions. Methods based on this idea are referred to as *causality-free* [19], [20], [22]–[24]. Note that the related method proposed in [21] differs slightly because it does not allow one to choose initial conditions freely.

1) *Generating infinite-horizon data*: Different from these prior works which consider fixed finite time OCPs, in this paper we are interested in infinite-horizon problems. Notice that the infinite-horizon PMP (11) is obtained with the limit $t_f \rightarrow \infty$ of a finite-horizon problem [27]. To reflect this, we solve (11) up to some fixed final time t_f . Next we check if the running cost $\mathcal{L}(\mathbf{x}(t_f), \mathbf{u}(t_f))$ is smaller than a desired tolerance. If not, we extend the time horizon and – using the previous solution as an initial guess – re-solve the BVP until the running cost is sufficiently small.

Once the running cost is small enough, it follows that the finite-horizon solution approximates the solution of the infinite-horizon problem as further integration should not change the cost significantly. This conclusion is reasonable as closed-loop stability is a necessary condition for finiteness of the value function. Then by applying (12) at each point along the trajectory $\mathbf{x} = \mathbf{x}(t; \mathbf{x}_0)$ and aggregating data from all infinite-horizon BVP solutions, we obtain a data set

$$\mathcal{D} = \left\{ \mathbf{x}^{(i)}, V^{(i)}, \boldsymbol{\lambda}^{(i)} \right\}_{i=1}^{|\mathcal{D}|},$$

where $V^{(i)} := V(\mathbf{x}^{(i)})$, $\boldsymbol{\lambda}^{(i)} := V_{\mathbf{x}}(\mathbf{x}^{(i)})$, and $|\mathcal{D}|$ denotes the number of data points in \mathcal{D} . Note that there is no need to distinguish data from different trajectories as the value function and its gradient are time-independent.

While generating data in this way is efficient because we extract a lot of data from each successful BVP solution, it has the side effect of concentrating a large amount of data near the equilibrium. On the other hand, we are interested in designing controllers which are effective over large regions of the state space and consequently we need data sets which support learning far from the equilibrium. To this end, instead of including the whole trajectory in the data set, we only take points with $t \leq t_f/T$ for a parameter $T \geq 1$. T is chosen to balance efficiency in data generation with the competing objective of adequately representing the entire state space; in this paper we set $T = 3$.

2) *LQR warm start for reliable BVP solution*: In this paper, we solve the two-point BVP using the SciPy [29] implementation of the BVP solver introduced in [30]. This algorithm is highly accurate but convergence is highly sensitive to the initial guess for $\mathbf{x}(t)$ and $\boldsymbol{\lambda}(t)$. Furthermore, convergence is increasingly dependent on good initializations as we increase the length of the time interval to approximate the infinite-horizon problem.

To mitigate this difficulty we simulate the dynamics up to some large final time t_f with an LQR controller (10) to close the loop. This provides a guess for the optimal state trajectory,

and a guess for the costate can be obtained with the LQR approximation $\boldsymbol{\lambda}(t) \approx 2\mathbf{P}\mathbf{x}(t)$. While the costate guess is often far from perfect, we find that it is usually close enough to facilitate reliable convergence over large time horizons. We refer to this strategy as *LQR warm-start*.

B. Neural network architecture

We employ a simple and intuitive architecture to model the value function. The main idea is to augment the LQR value function approximation with a NN which accounts for nonlinearities. The LQR value function is computed with respect to the dynamics linearized around $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$, and provides a good local approximation. The NN corrects and extends the approximation throughout the training domain.

As for the NN, we use a standard fully-connected feed-forward architecture. We denote the output of the network as $W^{\text{NN}}(\cdot)$. Feedforward NNs approximate complicated nonlinear functions by a composition of simpler functions, namely

$$W^{\text{NN}}(\mathbf{x}) = g_L \circ g_{L-1} \circ \dots \circ g_1(\mathbf{x}), \quad (13)$$

where the ℓ th layer, $\ell = 1, \dots, L$, is defined as $g_\ell(\mathbf{z}) = \sigma_\ell(\mathbf{W}_\ell \mathbf{z} + \mathbf{b}_\ell)$. Within each layer, \mathbf{W}_ℓ and \mathbf{b}_ℓ are the weight matrices and bias vectors, respectively, and $\sigma_\ell(\cdot)$ denotes a nonlinear *activation function* applied component-wise to its argument. We keep the scalar output layer $g_L(\cdot)$ linear, so $\sigma_L(\cdot)$ is the identity function.

We combine the raw NN prediction (13) with the LQR value function (8) for the linearized dynamics (7) as

$$V^{\text{NN}}(\mathbf{x}) = \frac{1}{c} \log [1 + cV^{\text{LQR}}(\mathbf{x})] + W^{\text{NN}}(\mathbf{x}), \quad (14)$$

with a trainable parameter $c > 0$. Intuitively, LQR provides a good approximation near $\bar{\mathbf{x}}$. There $V^{\text{LQR}}(\mathbf{x})$ is small and hence $c^{-1} \log [1 + cV^{\text{LQR}}(\mathbf{x})] \approx V^{\text{LQR}}(\mathbf{x})$ for all $c \in (0, \infty)$. Further away from $\bar{\mathbf{x}}$, we have $c^{-1} \log [1 + cV^{\text{LQR}}(\mathbf{x})] \ll V^{\text{LQR}}(\mathbf{x})$, thereby increasing the relative importance of the corrective NN. The parameter c governs the radius in which this term approximates $V^{\text{LQR}}(\mathbf{x})$; in particular $\lim_{c \rightarrow 0} c^{-1} \log [1 + cV^{\text{LQR}}(\mathbf{x})] = V^{\text{LQR}}(\mathbf{x})$. Notice that the model structure (14) is similar to a series expansion, except that we explicitly reduce the impact of lower order terms away from the linearization point.

Finally, the NN-based feedback control is evaluated by substituting $V_{\mathbf{x}}^{\text{NN}}(\cdot)$ into (6) in place of the gradient of the true value function:

$$\mathbf{u}^{\text{NN}}(\mathbf{x}) := \mathbf{u}^*(\mathbf{x}; V_{\mathbf{x}}^{\text{NN}}(\mathbf{x})). \quad (15)$$

It should be emphasized that the gradient $V_{\mathbf{x}}^{\text{NN}}(\cdot)$ is calculated using automatic differentiation, and is therefore *exact* and computationally efficient.

C. Physics-informed learning

Suppose we have generated a data set \mathcal{D} as discussed in Section III-A. This data takes the form of input-output pairs: $\mathbf{x}^{(i)}$ are the inputs and $(V^{(i)}, \boldsymbol{\lambda}^{(i)})$ are the outputs to be modeled. Let $\boldsymbol{\theta}$ denote the collection of model parameters:

$$\boldsymbol{\theta} := \{c\} \cup \{\mathbf{W}_\ell, \mathbf{b}_\ell\}_{\ell=1}^L.$$

In [21]–[23], the NN is trained by solving a supervised learning problem in which one minimizes the mean square regression loss,

$$\text{loss}_V(\boldsymbol{\theta}) := \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \left[V^{(i)} - V^{\text{NN}}(\mathbf{x}^{(i)}; \boldsymbol{\theta}) \right]^2, \quad (16)$$

plus a gradient regularization term

$$\text{loss}_\lambda(\boldsymbol{\theta}) := \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \left\| \boldsymbol{\lambda}^{(i)} - V_{\mathbf{x}}^{\text{NN}}(\mathbf{x}^{(i)}; \boldsymbol{\theta}) \right\|^2. \quad (17)$$

This serves as a form of *physics-informed* regularization. The term “physics-informed” is borrowed from [31] which partially inspired the authors’ previous work [22], [23]. By incorporating the prior knowledge that $\boldsymbol{\lambda}(t) = V_{\mathbf{x}}(\mathbf{x}(t))$, we maximize the information extracted from the available data and obtain more optimal feedback laws [21]–[23].

In this work we impose an additional penalty on deviating from the optimal control:

$$\text{loss}_u(\boldsymbol{\theta}) := \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \left\| \mathbf{u}^*(\mathbf{x}^{(i)}) - \mathbf{u}^{\text{NN}}(\mathbf{x}^{(i)}; \boldsymbol{\theta}) \right\|^2. \quad (18)$$

Minimizing this control penalty term contributes directly toward the ultimate goal of using the NN for optimal feedback by effectively enforcing the Hamiltonian minimization condition (6) on the learned feedback policy. We now arrive at the following optimization problem which we solve to train the NN:

$$\underset{\boldsymbol{\theta}}{\text{minimize}} \quad \text{loss}(\boldsymbol{\theta}) := \text{loss}_V(\boldsymbol{\theta}) + \mu_\lambda \text{loss}_\lambda(\boldsymbol{\theta}) + \mu_u \text{loss}_u(\boldsymbol{\theta}), \quad (19)$$

where $\mu_\lambda, \mu_u \geq 0$ are scalar weights.

To quantify the accuracy of the model, we generate two data sets from *independently drawn* initial conditions. During training, the network observes only data points from the training set $\mathcal{D}_{\text{train}}$. The other data set, which we call the validation set \mathcal{D}_{val} , is reserved for evaluating the NN accuracy *after* training. Good validation performance indicates that the NN generalizes well, i.e. it did not overfit the training data. The ability to empirically measure model accuracy in this way is a key feature of causality-free methods.

IV. APPLICATION TO DISTRIBUTED PARAMETER SYSTEM

In this section we explore the effectiveness of the proposed algorithm by solving a 64-dimensional OCP arising from a Chebyshev pseudospectral (PS) discretization of a modified Burgers’ equation with a destabilizing reaction term. Stabilization of Burgers’ equation is a common benchmark problem in distributed parameter systems and similar problems have recently been considered in e.g. [5], [14], [22], [23].

Let $X(t, \xi) : [0, \infty) \times [-1, 1] \rightarrow \mathbb{R}$ satisfy the following one-dimensional controlled PDE with Dirichlet boundary conditions:

$$\begin{cases} X_t = -\frac{1}{2}(X^2)_\xi + \nu X_{\xi\xi} + \alpha(\xi) X e^{-\beta X} + \mathbf{b}^T(\xi) \mathbf{u}(t) & \text{for } t > 0, \xi \in (-1, 1), \\ X(t, -1) = X(t, 1) & \text{for } t > 0, \\ X(0, \xi) = X_0(\xi) & \text{for } \xi \in (-1, 1). \end{cases} \quad (20)$$

Here $\nu, \beta > 0$ are scalar parameters, $\alpha : (-1, 1) \rightarrow \mathbb{R}$, and $\mathbf{b} : (-1, 1) \rightarrow \mathbb{R}^m$. The control $\mathbf{u} : [0, \infty) \rightarrow \mathbb{R}^m$ is designed to stabilize the *open-loop unstable* origin by solving the PDE-constrained OCP

$$\begin{cases} \min_{\mathbf{u}(\cdot)} & J[\mathbf{u}(\cdot)] = \int_0^\infty \left(\|X\|_{L^2_{(-1,1)}}^2 + R \mathbf{u}^T \mathbf{u} \right) dt, \\ \text{s.t.} & \text{Eq. (20)}, \end{cases} \quad (21)$$

where

$$\|X\|_{L^2_{(-1,1)}}^2 := \int_{-1}^1 X^2(t, \xi) d\xi.$$

We define

$$\alpha(\xi) = \begin{cases} -\kappa \left(\xi + \frac{1}{5} \right) \left(\xi - \frac{1}{5} \right), & \xi \in \left[-\frac{1}{5}, \frac{1}{5} \right], \\ 0, & \xi \notin \left[-\frac{1}{5}, \frac{1}{5} \right]. \end{cases}$$

and consider the case with $m = 2$ actuators active on compact supports defined by

$$\mathbf{b}(\xi) = \begin{pmatrix} \begin{cases} -\kappa \left(\xi + \frac{4}{5} \right) \left(\xi + \frac{2}{5} \right), & \xi \in \left[-\frac{4}{5}, -\frac{2}{5} \right], \\ 0, & \xi \notin \left[-\frac{4}{5}, -\frac{2}{5} \right], \\ -\kappa \left(\xi - \frac{2}{5} \right) \left(\xi - \frac{4}{5} \right), & \xi \in \left[\frac{2}{5}, \frac{4}{5} \right], \\ 0, & \xi \notin \left[\frac{2}{5}, \frac{4}{5} \right]. \end{cases} \end{pmatrix}$$

We set $\nu = 0.02$, $\beta = 0.1$, $\kappa = 25$, and $R = 0.5$.

A. Pseudospectral discretization

To solve (21) using our framework, we perform Chebyshev PS collocation to transform (20) into a system of n ordinary differential equations (ODEs). Following [32], let

$$\xi_j = \cos(j\pi/n), \quad j = 0, 1, \dots, n, n+1.$$

We collocate $X(t, \xi)$ at the non-boundary nodes ξ_j , $j = 1, 2, \dots, n$, and set $X(t, \xi_0) = X(t, \xi_{n+1}) = 0$ to account for the boundary conditions. We then define

$$\mathbf{x}(t) := (X(t, \xi_1), X(t, \xi_2), \dots, X(t, \xi_n))^T$$

and construct Chebyshev differentiation matrices $\mathbf{D}, \mathbf{D}^2 \in \mathbb{R}^{n \times n}$. Hence the PDE (20) becomes

$$\dot{\mathbf{x}} = -\frac{1}{2} \mathbf{D} \mathbf{x}^2 + \nu \mathbf{D}^2 \mathbf{x} + \boldsymbol{\alpha} \odot \mathbf{x} \odot e^{-\beta \mathbf{x}} + \mathbf{B} \mathbf{u}, \quad (22)$$

where $\mathbf{x}^2 := \mathbf{x} \odot \mathbf{x}$, “ \odot ” denotes elementwise multiplication, and $\boldsymbol{\alpha} \in \mathbb{R}^n$ and $\mathbf{B} \in \mathbb{R}^{n \times m}$ are the collocated versions of $\alpha(\xi)$ and $\mathbf{b}^T(\xi)$.

The integral appearing in the cost function is conveniently approximated by Clenshaw-Curtis quadrature [32]. Let w_j , $j = 1, 2, \dots, n$ be the non-boundary Clenshaw-Curtis quadrature weights. Then

$$\begin{cases} \|X\|_{L^2_{(-1,1)}}^2 = \int_{-1}^1 X^2(t, \xi) d\xi \approx \mathbf{x}^T \mathbf{Q} \mathbf{x}, \\ \mathbf{Q} = \text{diag}(w_1, w_2, \dots, w_n). \end{cases}$$

Now the original OCP (21) can be reformulated as a quadratic cost ODE-constrained problem,

$$\begin{cases} \min_{\mathbf{u}(\cdot)} & J_n[\mathbf{u}(\cdot)] = \int_0^\infty (\mathbf{x}^T \mathbf{Q} \mathbf{x} + R \mathbf{u}^T \mathbf{u}) dt, \\ \text{s.t.} & \text{Eq. (22)}. \end{cases} \quad (23)$$

B. Numerical results

We now present results of our method for solving the OCP (23) with $n = 64$ collocation points. The algorithm can handle higher dimensions, but we find that $n = 64$ points are enough to resolve the stiff dynamics for all initial conditions tested. We compare the accuracy of QRnet with that of a standard NN trained using supervised learning (comparable to [21]–[23]), as well as with a straight LQR approximation.

We consider initial conditions which are sums of sine functions with uniform random coefficients:

$$X_0(\xi) = \sum_{k=1}^{10} a_k \sin(k\pi\xi), \quad a_k \sim \mathcal{U}(-1/k, 1/k), \quad (24)$$

and use the following hyperparameters for training:

- in the NN component $W^{\text{NN}}(\cdot)$ we use $L = 5$ layers each with 32 neurons, and apply the $\tanh(\cdot)$ activation function to all hidden layers $\ell = 1, \dots, 4$;
- we set the weights in the loss function (19) to $\mu_\lambda = 0$ and $\mu_u = 5$, and optimize using L-BFGS-B [33];
- we implement the model in TensorFlow 1.11 [34] and train it on an NVIDIA RTX 2080Ti GPU.

First we study the sensitivity of the method with respect to variations in the data set and parameter initialization. This is important as collecting data can be time-consuming and NN training is a highly non-convex optimization problem. To this end we vary the number of trajectories in the training set $\mathcal{D}_{\text{train}}$, and for each different data set size, conduct ten trials with different randomly generated training trajectories and NN weight initializations. For validation we build a data set of 400 trajectories totaling $|\mathcal{D}_{\text{val}}| = 66060$ data points.

From the results shown in Fig. 1, it is clear that both the plain NN and QRnet vastly outperform LQR for value function reconstruction and optimal control prediction. Furthermore, the distribution of validation errors skews lower for QRnet than the plain NN. This suggests that QRnet is less sensitive to variations in the data set and parameter initialization, and thus enables the use of smaller data sets. The training time for smaller data sets is also shorter. Together, these properties facilitate rapid prototyping over an iterative design process, and combine naturally with the progressive model refinement strategy proposed in [22].

Next we select NN and QRnet models trained on the same set of 64 trajectories with nearly identical validation accuracy. We compare the performance of the LQR, NN, and QRnet feedback controllers to the open-loop optimal control, across 1200 simulations for initial conditions of different size, $\|X_0\|_{L^2_{(-1,1)}} = 0.1, 0.2, \dots, 1.2$. Results are shown in Fig. 2. As expected, the performance of LQR steadily degrades away from the origin, while the NN and QRnet controllers do well throughout the training domain. The latter two appear to be largely equivalent, except for a few initial conditions where the NN accumulates significant additional cost while LQR and QRnet do not. This supports the idea that QRnet inherits some local robustness from LQR, yielding a more reliable controller than a plain NN.

We conclude with a simulation of the dynamics (22) with QRnet feedback, for a typical random initial condition. Results

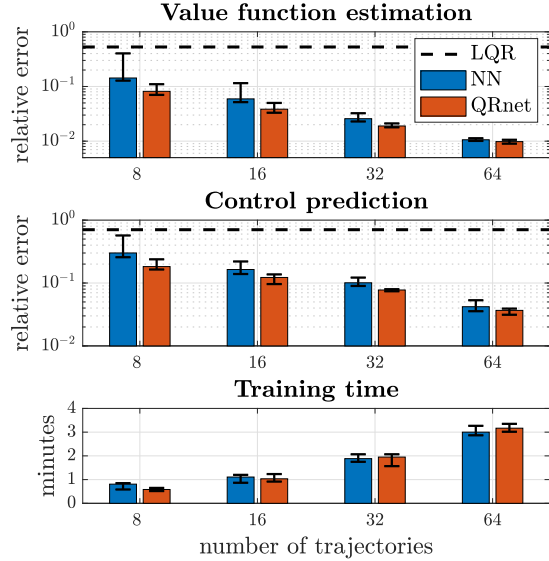


Fig. 1: Relative mean absolute error in estimating the value, relative mean L^2 error in predicting the optimal control, and training time, depending on the number of trajectories seen during training. The bar graph height shows the median over ten trials and error bars indicate the 25th and 75th percentiles.




Fig. 2: Difference between controlled and optimal costs, depending on the norm of the initial condition. Shaded areas cover the 15th to 85th percentiles, lines show medians, and symbols pick out the top three outliers for each group of initial conditions.

are plotted in Fig. 3. There we can see that QRnet stabilizes the system and closely approximates the optimal control, even where LQR deviates from it.

V. CONCLUDING REMARKS

In this paper we have presented *QRnet*, an extension of the causality-free physics-informed learning framework [22], [23] to infinite-horizon OCPs. This extension is comprised of two main features: efficient infinite-horizon data generation and a structurally-motivated NN architecture. By way of the Burgers’ benchmark problem, we have illustrated the potential for use in high-dimensional nonlinear systems and the improved robustness and performance the LQR-augmented model architecture.

Much remains to be explored in the development of deep learning approaches for feedback design. For instance, we are interested in applying the proposed method to problems




Fig. 3: Closed-loop dynamics and controls of the Burgers' system (22) for a random initial condition. The full simulation is over $t \in [0, 30]$ but we show only $t \in [0, 6]$.

with control constraints and non-quadratic costs with locally quadratic approximations. It will also be useful to study how beneficial the model architecture is, depending on how well the original problem can be approximated by a linear quadratic one. Finally, since the computational framework depends crucially on data generation, in future work we plan to study different strategies for improving the robustness and efficiency of this step.

REFERENCES

- [1] E. Al'brekht, "On the optimal stabilization of nonlinear systems," *J. Appl. Math. Mech.*, vol. 25(5), pp. 1254–1266, 1961.
- [2] D. Lukes, "Optimal regulation of nonlinear dynamical systems," *SIAM J. Control*, vol. 7, no. 1, pp. 75–100, 1969.
- [3] W. Kang, P. De, and A. Isidori, "Flight control in a windshear via nonlinear h_∞ methods," in *Proceedings of the 31st IEEE Conference on Decision and Control*, vol. 1, 1992, pp. 1135–1142.
- [4] C. Navasca and A. J. Krener, "Solution of Hamilton Jacobi Bellman equations," in *Proceedings of the 39th IEEE Conference on Decision and Control (CDC)*, vol. 1, 2000, pp. 570–574.
- [5] J. Borggaard and L. Zietsman, "The quadratic-quadratic regulator problem: Approximating feedback controls for quadratic-in-state nonlinear systems," in *American Control Conference (ACC)*, 2020, pp. 818–823.
- [6] S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *J. Comput. Phys.*, vol. 79, no. 1, pp. 12–49, 1988.
- [7] S. Cacace, E. Cristiani, M. Falcone, and A. Picarelli, "A patchy dynamic programming scheme for a class of Hamilton-Jacobi-Bellman equations," *SIAM J. Sci. Comput.*, vol. 34, no. 5, pp. A2625–A2649, 2012.
- [8] C. Navasca and A. J. Krener, "Patchy solutions of Hamilton-Jacobi-Bellman partial differential equations," in *Modeling, Estimation and Control*, ser. Lecture Notes in Control and Information Sciences, A. Chiuso, S. Pinzoni, and A. Ferrante, Eds. Springer-Verlag Berlin Heidelberg, 2007, vol. 364, pp. 251–270.
- [9] O. Bokanowski, J. Garcke, M. Griebel, and I. Klompaker, "An adaptive sparse grid semi-Lagrangian scheme for first order Hamilton-Jacobi Bellman equations," *J. Sci. Comput.*, vol. 55, no. 3, pp. 575–605, 2013.
- [10] M. Falcone and R. Ferretti, *Semi-Lagrangian Approximation Schemes for Linear and Hamilton-Jacobi Equations*. Society for Industrial and Applied Mathematics, 2013.
- [11] Y. T. Chow, J. Darbon, S. Osher, and W. Yin, "Algorithm for overcoming the curse of dimensionality for state-dependent Hamilton-Jacobi equations," *J. Comput. Phys.*, vol. 387, pp. 376–409, 2019.
- [12] I. Yegorov and P. M. Dower, "Perspectives on characteristics based curse-of-dimensionality-free numerical approaches for solving Hamilton-Jacobi equations," *Appl. Math. Optim.*, 2018.
- [13] S. Dolgov, D. Kalise, and K. Kunisch, "Tensor decompositions for high-dimensional Hamilton-Jacobi-Bellman equations," *arXiv:1908.01533 [math.OC]*, 2019.
- [14] D. Kalise and K. Kunisch, "Polynomial approximation of high-dimensional Hamilton-Jacobi-Bellman equations and applications to feedback control of semilinear parabolic PDEs," *SIAM J. Sci. Comput.*, vol. 40, no. 2, pp. A629–A652, 2018.
- [15] M. Abu-Khalaf and F. L. Lewis, "Nearly optimal control laws for nonlinear systems with saturating actuators using a neural network HJB approach," *Automatica*, vol. 41, no. 5, pp. 779–791, 2005.
- [16] Y. Tassa and T. Erez, "Least squares solutions of the HJB equation with neural network value-function approximators," *IEEE Trans. Neural Netw.*, vol. 18, no. 4, pp. 1031–1041, 2007.
- [17] J. Sirignano and K. Spiliopoulos, "DGM: A deep learning algorithm for solving partial differential equations," *J. Comput. Phys.*, vol. 375, pp. 1339–1364, 2018.
- [18] F. Jiang, G. Chou, M. Chen, and C. J. Tomlin, "Using neural networks to compute approximate and guaranteed feasible Hamilton-Jacobi-Bellman PDE solutions," *arXiv:1611.03158 [cs.LG]*, 2016.
- [19] W. Kang and L. C. Wilcox, "A causality free computational method for HJB equations with application to rigid body satellites," in *AIAA Guidance, Navigations, and Control Conference*, 2015.
- [20] —, "Mitigating the curse of dimensionality: Sparse grid characteristics method for optimal feedback control and HJB equations," *Comput. Optim. Appl.*, vol. 68, no. 2, pp. 289–315, 2017.
- [21] D. Izzo, E. Öztürk, and M. Märten, "Interplanetary transfers via deep representations of the optimal policy and/or of the value function," in *Genetic and Evolutionary Computation Conference*, 2019, pp. 1971–1979.
- [22] T. Nakamura-Zimmerer, Q. Gong, and W. Kang, "Adaptive deep learning for high-dimensional Hamilton-Jacobi-Bellman equations," *arXiv:1907.05317 [math.OC]*, 2019.
- [23] —, "A causality-free neural network method for high-dimensional Hamilton-Jacobi-Bellman equations," in *American Control Conference (ACC)*, 2020, pp. 787–793.
- [24] B. Azmi, D. Kalise, and K. Kunisch, "Data-driven recovery of optimal feedback laws through optimality conditions and sparse polynomial regression," *arXiv:2007.09753 [math.OC]*, 2020.
- [25] J. Han, A. Jentzen, and W. E., "Solving high-dimensional partial differential equations using deep learning," *Proc. Natl. Acad. Sci. USA*, vol. 115, no. 34, pp. 8505–8510, 2018.
- [26] M. G. Crandall and P.-L. Lions, "Viscosity solutions of Hamilton-Jacobi equations," *Trans. Amer. Math. Soc.*, vol. 277, no. 1, pp. 1–42, 1983.
- [27] L. S. Pontryagin, *Mathematical Theory of Optimal Processes*, ser. L.S. Pontryagin selected works. Taylor and Francis, 1987, vol. 4.
- [28] O. L. Mangasarian, "Sufficient conditions for the optimal control of nonlinear systems," *SIAM J. Control*, vol. 4, no. 1, pp. 139–152, 1966.
- [29] P. Virtanen, R. Gommers, T. E. Oliphant, and et. al., "SciPy 1.0: Fundamental algorithms for scientific computing in Python," *Nat. Methods*, vol. 17, pp. 261–272, 2020.
- [30] J. Kierzenka and L. F. Shampine, "A BVP solver based on residual control and the MATLAB PSE," *ACM Trans. Math. Softw.*, vol. 27, no. 3, pp. 299–316, 2001.
- [31] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019.
- [32] L. N. Trefethen, *Spectral Methods in MATLAB*. Society for Industrial and Applied Mathematics, 2000.
- [33] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM J. Sci. Comput.*, vol. 16, no. 5, pp. 1190–1208, 1995.
- [34] M. Abadi, A. Agarwal, P. Barham, et al., "TensorFlow: Large-scale machine learning on heterogeneous systems," *arXiv:1603.04467 [cs.DC]*, 2016.