



## Brief paper

Large scale model predictive control with neural networks and primal active sets<sup>☆</sup>Steven W. Chen<sup>a,\*</sup>, Tianyu Wang<sup>b</sup>, Nikolay Atanasov<sup>b</sup>, Vijay Kumar<sup>a</sup>, Manfred Morari<sup>a</sup><sup>a</sup> GRASP Laboratory, University of Pennsylvania, Philadelphia, PA, United States of America<sup>b</sup> Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA, United States of America

## ARTICLE INFO

## Article history:

Received 28 October 2019  
 Received in revised form 5 July 2021  
 Accepted 24 August 2021  
 Available online xxx

## Keywords:

Fully connected neural network  
 Primal active set method  
 Model predictive control  
 Receding horizon control

## ABSTRACT

This work presents an explicit–implicit procedure to compute a model predictive control (MPC) law with guarantees on recursive feasibility and asymptotic stability. The approach combines an offline-trained fully-connected neural network with an online primal active set solver. The neural network provides a control input initialization while the primal active set method ensures recursive feasibility and asymptotic stability. The neural network is trained with a primal–dual loss function, aiming to generate control sequences that are primal feasible and meet a desired level of suboptimality. Since the neural network alone does not guarantee constraint satisfaction, its output is used to warm start the primal active set method online. We demonstrate that this approach scales to large problems with thousands of optimization variables, which are challenging for current approaches. Our method achieves a 2× reduction in online inference time compared to the best method in a benchmark suite of different solver and initialization strategies.

© 2021 Elsevier Ltd. All rights reserved.

## 1. Introduction

Model predictive control (MPC) is a dynamic optimization technique widely used in industrial process applications, such as oil refineries and chemical plants (Qin & Badgwell, 2003). Recently, MPC has found mainstream use in robotics for controlling ground (Richter, Vega-Brown, & Roy, 2018), aerial (Bouffard, Aswani, & Tomlin, 2012; Watterson & Kumar, 2015), humanoid (Erez et al., 2013), and other autonomous robots due to its versatility, robustness, and safety guarantees. The transition from the process industry to robotics brings additional challenges since the available computation time is reduced from hours to milliseconds.

MPC techniques can be categorized into *implicit* and *explicit*. Implicit MPC focuses on online computation of an open-loop control sequence, optimizing the system performance at its current

<sup>☆</sup> This work is supported in part by ARO grant W911NF-13-1-0350, DARPA, United States grant HR001151626/HR0011516850, NSF, United States CRII RI IIS-1755568, the Semiconductor Research Corporation (SRC), United States, the NVIDIA AI Labs program, ONR SAI, United States N00014-18-1-2828, and ARL DCIST CRA, United States W911NF-17-2-0181. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editors Gabriele Pannocchia, Martin Monnigmann under the direction of Editor Ian R. Petersen.

\* Corresponding author.

E-mail addresses: [chenste@seas.upenn.edu](mailto:chenste@seas.upenn.edu) (S.W. Chen), [tiw161@eng.ucsd.edu](mailto:tiw161@eng.ucsd.edu) (T. Wang), [natanasov@eng.ucsd.edu](mailto:natanasov@eng.ucsd.edu) (N. Atanasov), [kumar@seas.upenn.edu](mailto:kumar@seas.upenn.edu) (V. Kumar), [morari@seas.upenn.edu](mailto:morari@seas.upenn.edu) (M. Morari).

state. For example, Wang and Boyd (2010) exploit the structure of the quadratic program (QP) associated with the MPC problem to design efficient interior point methods. Vichik and Borrelli (2014) experiment with a different computation scheme and demonstrate that analog circuits can solve QPs in microseconds through the use of custom hardware. Rather than solving an optimization problem online, explicit methods manage the computational load by pre-computing a control law  $\mathbf{u} = \boldsymbol{\mu}(\mathbf{x})$  offline, as a function of all feasible states  $\mathbf{x}$ , where  $\boldsymbol{\mu}$  is known to be piecewise affine on polytopes determined by the system constraints (Borrelli, Bemporad, & Morari, 2017, Ch. 11). Computing an optimal explicit control law can become computationally intractable in large problems because the number of polytopic regions may grow exponentially (worst case) with the number of constraints. Even determining which region contains the current system state online can require nontrivial processing power or memory storage (Kvasnica & Fikar, 2012). One approach to address the computational challenges of explicit MPC control is to construct an approximate sub-optimal controller. Jones and Morari (2010) use a double-description method to build piecewise-affine approximations of the value function and use barycentric functions on the polytopic regions to obtain an associated control law. The authors prove recursive feasibility and asymptotic stability of the suboptimal controller. Despite various approximation techniques, explicit MPC has not been demonstrated to scale to the same problem sizes as those handled by implicit MPC methods.

This paper investigates the use of a Rectified Linear Unit (ReLU) fully connected neural network to approximate the piecewise affine explicit MPC control law. There have been several recent works using neural networks for MPC design. [Chen et al. \(2018\)](#) use a neural network with an orthogonal projection operation to approximate the optimal control law. [Hertneck, Köhler, Trimpe, and Allgöwer \(2018\)](#) use a neural network in a robust MPC framework to provide statistical guarantees of feasibility and stability. [Zhang, Bujarbaruah, and Borrelli \(2019\)](#) use a neural network to approximate the primal and dual variables, and provide statistical guarantees as well as certificates of suboptimality. Our work extends these approaches by providing deterministic guarantees on recursive feasibility and asymptotic stability through corrective steps generated by an online QP solver. Moreover, we demonstrate for the first time that neural-network-based approximations of MPC control laws scale to large problems with thousands of optimization variables.

A few closely related works combine the strengths of machine learning for scalability with the analytical tractability of QP problems in an explicit-implicit MPC approach. [Zeilinger, Jones, and Morari \(2011\)](#) compute a piecewise affine control law approximation, combine it with an active set method, and provide criteria to terminate the active set method early while still obtaining guarantees on recursive feasibility and asymptotic stability. [Klaučo, Kalúz, and Kvasnica \(2019\)](#) use classification trees and nearest neighbors to warm start an active set method and solve to the optimal solution. These approaches differ from our method because they do not utilize a neural network as the control policy approximator. In addition, they do not demonstrate nor address the challenges in scaling to large systems, specifically related to efficient generation of train data sets in high dimensions.

Generating large data sets of initial feasible states, optimal inputs, primal, and dual variables for high-dimensional systems to approximate an MPC controller has not been studied widely in the literature. Yet, a large data set is critical for accurate learning-based control law approximations, especially for high-dimensional systems. A naïve rejection sampling approach to generate feasible initial states has rapidly diminishing probability of success as the number of constraints and dimensions grow. Existing works use gridding ([Klaučo et al., 2019](#); [Summers, Jones, Lygeros, & Morari, 2011](#)) or random sampling ([Hertneck et al., 2018](#); [Zhang et al., 2019](#)) methods that do not scale to high dimensions, or rely on reinforcement and imitation learning techniques that perform closed-loop simulations leading to weaknesses discussed in [Ross, Gordon, and Bagnell \(2011\)](#). We propose an algorithm to efficiently generate large data sets in high dimensions using geometric random walks.

In summary, the goal of this work is to develop a neural-network approximation of an MPC control law that scales to large systems, performs faster than implicit methods online, and provides guarantees on recursive feasibility and asymptotic stability. Our contributions include:

- a primal-dual loss function that incorporates state and input constraints to train a deep ReLU network approximation of an explicit MPC control law (Section 3);
- an online approach that corrects the network output using a primal active set solver to guarantee recursive feasibility and asymptotic stability (Section 4);
- a geometric random walk algorithm for generating feasible sample data sets efficiently, necessary for training control law approximations for large systems (Section 5);
- a 2× reduction in online inference time against the best benchmark method on a system with thousands of optimization variables (Section 6).

## 2. Preliminaries

Consider a discrete-time linear time-invariant system,

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t), \quad (1)$$

subject to state and input constraints,

$$\begin{aligned} \mathbf{x}(t) \in \mathcal{X} &:= \{\mathbf{x} \in \mathbb{R}^n | \mathbf{A}_x \mathbf{x} \leq \mathbf{b}_x\} \quad \forall t \geq 0, \\ \mathbf{u}(t) \in \mathcal{U} &:= \{\mathbf{u} \in \mathbb{R}^m | \mathbf{A}_u \mathbf{u} \leq \mathbf{b}_u\} \quad \forall t \geq 0. \end{aligned} \quad (2)$$

Assuming that the pair  $(\mathbf{A}, \mathbf{B})$  is stabilizable ([Borrelli et al., 2017, Ch. 7](#)), we are interested in obtaining a receding horizon controller (RHC) ([Borrelli et al., 2017, Ch. 12](#)). At each time  $t$ , we solve a constrained finite-time optimal control problem ([Borrelli et al., 2017, Ch. 7](#)),

$$\begin{aligned} \min_{\mathbf{u}_{0:N-1}} J(\mathbf{u}_{0:N-1} | \mathbf{x}(t)) &= \mathbf{x}_N^\top \mathbf{P} \mathbf{x}_N \\ &+ \sum_{k=0}^{N-1} (\mathbf{x}_k^\top \mathbf{Q} \mathbf{x}_k + \mathbf{u}_k^\top \mathbf{R} \mathbf{u}_k) \\ \text{s.t. } \mathbf{x}_{k+1} &= \mathbf{A} \mathbf{x}_k + \mathbf{B} \mathbf{u}_k, \quad \mathbf{x}_0 = \mathbf{x}(t), \\ \mathbf{A}_x \mathbf{x}_k &\leq \mathbf{b}_x, \quad \mathbf{A}_u \mathbf{u}_k \leq \mathbf{b}_u, \quad \mathbf{A}_f \mathbf{x}_N \leq \mathbf{b}_f, \end{aligned} \quad (3)$$

where symmetric matrices  $\mathbf{Q}, \mathbf{P} \in \mathcal{S}_{>0}^n$ ,  $\mathbf{R} \in \mathcal{S}_{>0}^m$  define desired system behavior over planning horizon  $N$ . The terminal cost,  $\mathbf{x}_N^\top \mathbf{P} \mathbf{x}_N$ , and terminal constraints,  $\mathcal{X}_f := \{\mathbf{x} \in \mathbb{R}^n | \mathbf{A}_f \mathbf{x} \leq \mathbf{b}_f\}$ , ensure feasibility and asymptotic stability of the RHC as will be discussed in Section 2.3. The first input,  $\mathbf{u}_0$ , is applied to the system, leading to a new state  $\mathbf{x}(t+1)$ , and the process is repeated. Alternatively, instead of recomputing  $\mathbf{u}_{0:N-1}$  at every  $t$ , it may be desirable to obtain a function  $\mathbf{u}_0 = \mu(\mathbf{x})$  that specifies the input for an arbitrary state  $\mathbf{x}$ .

### 2.1. Batch formulation

The batch formulation of the problem in (3) is:

$$\begin{aligned} \min_{\mathbf{z}} J(\mathbf{z} | \mathbf{x}) &= \mathbf{z}^\top \mathbf{H} \mathbf{z} + \mathbf{x}^\top \mathbf{Q} \mathbf{x} \\ \text{s.t. } \mathbf{G}_{\text{eq}} \mathbf{z} &= \mathbf{E}_{\text{eq}} \mathbf{x}, \quad \mathbf{G}_{\text{in}} \mathbf{z} \leq \mathbf{w}_{\text{in}} + \mathbf{E}_{\text{in}} \mathbf{x} \end{aligned} \quad (4)$$

with the following definitions

$$\begin{aligned} \mathbf{z} &= [\mathbf{x}_1^\top \dots \mathbf{x}_N^\top \mathbf{u}_0^\top \dots \mathbf{u}_{N-1}^\top] \in \mathbb{R}^{N(n+m)} \\ \mathbf{H} &= \text{diag}(\mathbf{I}_{N-1} \otimes \mathbf{Q}, \mathbf{P}, \mathbf{I}_N \otimes \mathbf{R}) \in \mathbb{R}^{N(n+m) \times N(n+m)} \\ \mathbf{G}_{\text{eq}} &= [\mathbf{I}_{Nn} - \mathbf{L}_N \otimes \mathbf{A}; -\mathbf{I}_N \otimes \mathbf{B}] \in \mathbb{R}^{Nn \times N(n+m)} \\ \mathbf{E}_{\text{eq}} &= \mathbf{e}_1 \otimes \mathbf{A} \in \mathbb{R}^{Nn \times n} \\ \mathbf{G}_{\text{in}} &= \text{diag}(\mathbf{0}_{c_x}, \mathbf{I}_{N-1} \otimes \mathbf{A}_x, \mathbf{A}_f, \mathbf{I}_N \otimes \mathbf{A}_u) \\ \mathbf{E}_{\text{in}} &= -\mathbf{e}_1 \otimes \mathbf{A}_x \in \mathbb{R}^{(Nc_x + c_f + Nc_u) \times n} \\ \mathbf{w}_{\text{in}} &= [\mathbf{1}_N \otimes \mathbf{b}_x; \mathbf{b}_f; \mathbf{1}_N \otimes \mathbf{b}_u] \in \mathbb{R}^{(Nc_x + c_f + Nc_u)} \end{aligned} \quad (5)$$

where  $\mathbf{L}_N$  is the matrix of size  $N \times N$  with ones on the first subdiagonal and zeros elsewhere,  $\mathbf{e}_i$  is the  $i$ th standard basis vector,  $\mathbf{1}_N$  is the vector of all ones of size  $N$ ,  $c_x, c_f, c_u$  denote the number of constraints specified by the rows of  $\mathbf{A}_x, \mathbf{A}_f, \mathbf{A}_u$ , “;” denotes vertical concatenation, and  $\otimes$  denotes the Kronecker product. When the parameter  $\mathbf{x}$  is fixed, (4) is a quadratic program (QP) and the solution is a vector  $\mathbf{z}$ .

A common choice for the terminal cost matrix  $\mathbf{P}$  is the solution  $\mathbf{P}_\infty$  to the algebraic Riccati equation ([Borrelli et al., 2017, Ch. 8](#)):

$$\mathbf{P}_\infty = \mathbf{A}^\top \mathbf{P}_\infty \mathbf{A} + \mathbf{Q} - \mathbf{A}^\top \mathbf{P}_\infty \mathbf{B} (\mathbf{B}^\top \mathbf{P}_\infty \mathbf{B} + \mathbf{R})^{-1} \mathbf{B}^\top \mathbf{P}_\infty \mathbf{A}. \quad (6)$$

A common choice for  $\mathcal{X}_f$  is the maximal positively invariant set,  $\mathcal{O}_\infty^{\text{LQR}}$ , of the linear quadratic regulator (LQR), which is computable via reachability analysis ([Borrelli et al., 2017, Ch.11](#)) using standard toolboxes ([Herceg, Kvasnica, Jones, & Morari, 2013](#)). Note

that  $\mathcal{O}_{\infty}^{LQR}$  is easily computable even for large systems, and is different from the maximal control invariant set,  $\mathcal{C}^{\infty}$ , used in [Chen et al. \(2018\)](#), which is very challenging to compute for large systems.

## 2.2. Feasibility and duality

For a given state  $\mathbf{x}$ , (4) is a strictly convex QP because if  $\mathbf{Q}$ ,  $\mathbf{P}$  are positive definite then  $\mathbf{H}$  is positive definite too. While this assumption can be relaxed to allow positive semi-definite matrices, it simplifies the development of our approach for approximating the solution of (4), which is unique in this case. Let  $\mathcal{X}_0 \subseteq \mathcal{X}$  be the set of parameters  $\mathbf{x}$  for which (4) is feasible ([Borrelli et al., 2017](#), Defn. 6.3). Given  $\mathbf{x} \in \mathcal{X}_0$  and an associated primal feasible  $\mathbf{z}$ , the *suboptimality level* of  $\mathbf{z}$  is:

$$\sigma(\mathbf{z}|\mathbf{x}) := J(\mathbf{z}|\mathbf{x}) - J(\mathbf{z}^*|\mathbf{x}) \geq 0, \quad (7)$$

where  $\mathbf{z}^*$  is the unique minimizer of (4). We introduce *dual variables*  $(\mathbf{v}, \boldsymbol{\lambda})$  ([Boyd & Vandenberghe, 2004](#), Ch.5), and define the Lagrangian associated with (4) as:

$$\begin{aligned} \mathcal{L}(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) := & \mathbf{z}^{\top} \mathbf{H} \mathbf{z} + \mathbf{x}^{\top} \mathbf{Q} \mathbf{x} + \mathbf{v}^{\top} (\mathbf{G}_{\text{eq}} \mathbf{z} - \mathbf{E}_{\text{eq}} \mathbf{x}) \\ & + \boldsymbol{\lambda}^{\top} (\mathbf{G}_{\text{in}} \mathbf{z} - \mathbf{w}_{\text{in}} - \mathbf{E}_{\text{in}} \mathbf{x}). \end{aligned} \quad (8)$$

The Lagrangian dual of a minimization QP is a maximization QP over  $(\mathbf{v}, \boldsymbol{\lambda})$  with objective function  $d(\mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) := \inf_{\mathbf{z}} \mathcal{L}(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x})$ . The dual variables  $(\mathbf{v}, \boldsymbol{\lambda})$  are *dual feasible* if  $\boldsymbol{\lambda} \geq 0$ . For a given parameter  $\mathbf{x}$ , any primal feasible  $\mathbf{z}$ , any dual feasible  $(\mathbf{v}, \boldsymbol{\lambda})$ , and optimal primal-dual variables  $(\mathbf{z}^*, \mathbf{v}^*, \boldsymbol{\lambda}^*)$ , *strong duality* holds because Slater's conditions are satisfied ([Boyd & Vandenberghe, 2004](#), Ch.5):

$$\begin{aligned} d(\mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) \leq d(\mathbf{v}^*, \boldsymbol{\lambda}^*|\mathbf{x}) &= \mathcal{L}(\mathbf{z}^*, \mathbf{v}^*, \boldsymbol{\lambda}^*|\mathbf{x}) \\ &= J(\mathbf{z}^*|\mathbf{x}) \leq J(\mathbf{z}|\mathbf{x}). \end{aligned} \quad (9)$$

The *feasible duality gap* associated with  $\mathbf{x}$ ,  $\mathbf{z}$ ,  $\mathbf{v}$ ,  $\boldsymbol{\lambda}$  is

$$\eta(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) := J(\mathbf{z}|\mathbf{x}) - d(\mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) \geq \sigma(\mathbf{z}|\mathbf{x}) \geq 0, \quad (10)$$

which is an upper bound on the suboptimality level  $\sigma(\mathbf{z}|\mathbf{x})$  for any feasible  $(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda})$  due to (9).

## 2.3. Receding horizon control

Alternatively, (4) can be viewed as a *multiparametric quadratic program* (mp-QP) over functions that map  $\mathbf{x}$  to primal variables  $\mathbf{z}$  ([Borrelli et al., 2017](#), Ch. 6).

**Definition 1.** A planner is a function  $\pi : \mathcal{X}_0 \rightarrow \mathbb{R}^{N(n+m)}$  that maps a parameter  $\mathbf{x}$  to a decision variable  $\mathbf{z}$ . A planner  $\pi$  is primal feasible if  $\mathbf{z} = \pi(\mathbf{x})$  is a primal feasible variable  $\forall \mathbf{x} \in \mathcal{X}_0$  in (4). It is optimal, if  $\mathbf{z} = \pi(\mathbf{x})$  is the optimal solution  $\forall \mathbf{x} \in \mathcal{X}_0$ .

**Definition 2.** The receding horizon controller (RHC) corresponding to a planner  $\pi$  is a function  $\mu : \mathcal{X}_0 \rightarrow \mathbb{R}^m$  that returns the first control input in  $\mathbf{z} = \pi(\mathbf{x})$ , i.e.,  $\mathbf{u}_0 = \mu(\mathbf{x})$ .

Recursive feasibility and asymptotic stability ([Borrelli et al., 2017](#), Ch.12) are two properties that should be guaranteed for an RHC. Primal feasibility is a property of a planner  $\pi$  and the open-loop optimization problem, while recursive feasibility is a property of an RHC  $\mu$  and the corresponding closed loop system  $\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mu(\mathbf{x}(t))$  for  $t \geq 0$ . In general, primal feasibility of the open-loop optimization problem does not imply recursive feasibility of the corresponding closed-loop controller. Recursive feasibility of a control law is a necessary, but not sufficient, condition for asymptotic stability.

## 3. Explicit control policy approximation

The first step of our approach is offline training of a deep neural network  $\tilde{\pi}(\mathbf{x}|\boldsymbol{\theta})$  with parameters  $\boldsymbol{\theta}$  that provides a candidate solution  $\mathbf{z}$  for the QP in (4). We choose to approximate the entire primal prediction over the planning horizon instead of the control law  $\mu$ . Although this choice requires approximating  $N(n+m)$  variables instead of only  $m$ , the additional predicted variables may be used to obtain the desired guarantees through an online primal active set solver in Section 4.

**Definition 3.** A deep neural network (DNN)  $\tilde{\pi}(\mathbf{x}|\boldsymbol{\theta})$  with  $L$  layers is a composition of  $L$  affine functions:  $\lambda_l(\mathbf{x}) := \boldsymbol{\theta}_l^w \mathbf{x} + \boldsymbol{\theta}_l^b$ , each except the last one followed by a nonlinear activation function  $\mathbf{h}$ , so that:  $\tilde{\pi}(\mathbf{x}|\boldsymbol{\theta}) = \lambda_L \circ \mathbf{h} \circ \lambda_{L-1} \circ \dots \circ \mathbf{h} \circ \lambda_1(\mathbf{x})$ . The DNN parameters are  $\boldsymbol{\theta} := \{(\boldsymbol{\theta}_l^w, \boldsymbol{\theta}_l^b)\}_{l=1}^L$ . Each layer  $l$  has width defined by the number of rows of  $\boldsymbol{\theta}_l^w$  and  $\boldsymbol{\theta}_l^b$ .

The activation function  $\mathbf{h}$  is fixed (not optimized) and typically chosen as a sigmoid, hypertangent, or ReLU ([Goodfellow, Bengio, & Courville, 2016](#), Ch.6). As noted in [Chen et al. \(2018\)](#), a ReLU activation  $\mathbf{h}(\mathbf{x}) := \max\{\mathbf{0}, \mathbf{x}\}$ , where the max is applied elementwise, is well-suited for mp-QP problems because any piecewise-affine function on polyhedra, such as the solution for an mp-QP ([Borrelli et al., 2017](#), Thm. 6.7), can be represented exactly by a ReLU DNN ([Arora, Basu, Mianjy, & Mukherjee, 2018](#), Thm. 2.1). We restrict the approximation of a planner  $\pi$  to the class of functions represented by a ReLU DNN  $\tilde{\pi}$ .

A train data set  $\mathcal{D} := \{\mathbf{x}_i, \mathbf{z}_i^*, \mathbf{v}_i^*, \boldsymbol{\lambda}_i^*\}_i$  of optimal primal and dual variables is necessary for supervised learning of the neural network parameters  $\boldsymbol{\theta}$  that approximate an optimal planner  $\pi^*(\mathbf{x}) \approx \tilde{\pi}(\mathbf{x}|\boldsymbol{\theta})$ . Optimizing a least-squares loss function measuring the discrepancy between  $\mathbf{z}_i^*$  and  $\mathbf{z}_i := \tilde{\pi}(\mathbf{x}_i|\boldsymbol{\theta})$  is a common choice but it fails to explicitly account for the constraints in (4). We propose a loss function which incorporates information from the optimal primal and dual variables by measuring the discrepancy of the Lagrangian values:

$$\ell(\boldsymbol{\theta}) := \sum_{i=1}^{|\mathcal{D}|} (\mathcal{L}(\tilde{\pi}(\mathbf{x}_i|\boldsymbol{\theta}), \mathbf{v}_i^*, \boldsymbol{\lambda}_i^*|\mathbf{x}_i) - \mathcal{L}(\mathbf{z}_i^*, \mathbf{v}_i^*, \boldsymbol{\lambda}_i^*|\mathbf{x}_i))^2.$$

There are a variety of potential loss functions that utilize the Lagrangian. We chose this particular one because it allows supervision from both the optimal primal and dual variables and was effective in our experiments (Section 6). Since  $\pi^*(\mathbf{x})$  is piecewise-affine, it can be represented exactly by  $\tilde{\pi}(\mathbf{x}|\boldsymbol{\theta})$  with sufficient depth and width. By strict convexity, the primal variables  $\mathbf{z}_i^*$  are unique minimizers of each term in the loss  $\ell(\boldsymbol{\theta})$ . We train the network using stochastic gradient descent ([Kingma & Ba, 2015](#)), sampling a subset (mini-batch) of the train data  $\mathcal{D}$ , computing the loss,  $\ell(\boldsymbol{\theta})$ , over the mini-batch, and using backpropagation to compute the loss gradient and update the network parameters  $\boldsymbol{\theta}$ . The mini-batch sampling is continued until the data set is exhausted, which concludes an epoch, and is then repeated until convergence.

## 4. Primal active set method for guarantees

The trained neural network  $\tilde{\pi}(\mathbf{x}|\boldsymbol{\theta})$  is a planner that maps states  $\mathbf{x}$  to primal predictions  $\mathbf{z}$ . It can be implemented as an RHC  $\mu$  by applying the first control input  $\mathbf{u}_0$  from  $\mathbf{z} = \tilde{\pi}(\mathbf{x}(t)|\boldsymbol{\theta})$  at each encountered state  $\mathbf{x}(t)$  over time. [Theorem 4](#) summarizes the conditions that ensure recursive feasibility and asymptotic stability of  $\mu$ .

**Theorem 4.** Let  $\pi$  be a planner for the finite-horizon deterministic optimal control problem in (3) with  $\mathbf{b}_x, \mathbf{b}_u, \mathbf{b}_f > \mathbf{0}$ . Assume that the terminal constraint set  $\mathcal{X}_f$  is control invariant and the terminal cost  $p(\mathbf{x}) := \mathbf{x}^\top \mathbf{P} \mathbf{x}$  is a control Lyapunov function (Borrelli et al., 2017, Rmk 12.3) over  $\mathcal{X}_f$ . If for all  $\mathbf{x} \in \mathcal{X}_0$ ,  $\mathbf{z} = \pi(\mathbf{x})$  is primal feasible and there exists a function  $\gamma(\mathbf{x})$  such that:

$$0 \leq \sigma(\pi(\mathbf{x})|\mathbf{x}) \leq \gamma(\mathbf{x}) \leq \mathbf{x}^\top \mathbf{Q} \mathbf{x}, \quad (11)$$

then the RHC  $\mu$  corresponding to  $\pi$  is recursively feasible and asymptotically stable with domain of attraction  $\mathcal{X}_0$  for system (1) subject to constraints (2).

**Proof.** Recursive feasibility of  $\mu$  follows from Borrelli et al. (2017, Thm. 12.1), while asymptotic stability follows from Borrelli et al. (2017, Thm. 13.1).  $\square$

If we guarantee that the neural network output  $\mathbf{z}(t) = \tilde{\pi}(\mathbf{x}(t)|\theta)$  at each time step  $t$  is primal feasible and that there exists  $\gamma(\mathbf{x}(t))$  that satisfies threshold (11), then according to Theorem 4, the RHC  $\mathbf{u}(t) = \mu(\mathbf{x}(t))$  will be recursively feasible and asymptotically stable. In Section 4.1, we show how to construct  $\gamma(\mathbf{x})$  and check whether a given network output  $\mathbf{z}$  satisfies primal feasibility and threshold (11). Then, in Section 4.2, we use the network output  $\mathbf{z}$  as initialization of an online primal active set solver which optimizes  $\mathbf{z}$  until primal feasibility and the suboptimality certificate in (11) are both satisfied.

#### 4.1. Obtaining a suboptimality certificate

Given  $\mathbf{z} = \tilde{\pi}(\mathbf{x}|\theta)$ , primal feasibility can be checked using the constraints in (4). Section 4.2 details how to obtain primal feasibility given a primal infeasible initialization. Once  $\mathbf{z}$  is primal feasible, we obtain dual feasible variables  $(\mathbf{v}, \boldsymbol{\lambda})$  and define  $\gamma(\mathbf{x})$  as the feasible duality gap  $\eta(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x})$  in (10). This guarantees that  $\sigma(\mathbf{z}|\mathbf{x}) \leq \eta(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x})$  and we can directly check whether  $\eta(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) \leq \mathbf{x}^\top \mathbf{Q} \mathbf{x}$  to satisfy (11).

We use the KKT optimality conditions to derive  $(\mathbf{v}, \boldsymbol{\lambda})$ . First, we determine which constraints in (4) are active at  $\mathbf{z}$ . All  $N_n$  equality constraints  $\mathbf{G}_{\text{eq}} \mathbf{z} = \mathbf{E}_{\text{eq}} \mathbf{x}$  are active by definition. Let  $\mathcal{A}$  be the set of inequality constraints  $i$  such that  $\mathbf{e}_i^\top \mathbf{G}_{\text{in}} \mathbf{z} = \mathbf{e}_i^\top \mathbf{w}_{\text{in}} + \mathbf{e}_i^\top \mathbf{E}_{\text{in}} \mathbf{x}$ . Due to complementary slackness, dual variables corresponding to inactive constraints should be set to 0. The remaining dual variables  $(\mathbf{v}, \boldsymbol{\lambda}_{\mathcal{A}})$  must satisfy the KKT conditions for an equality constrained QP:

$$\begin{bmatrix} 2\mathbf{H} & \mathbf{G}_{\text{eq}}^\top & \mathbf{G}_{\mathcal{A}}^\top \\ \mathbf{G}_{\text{eq}} & \mathbf{0} & \mathbf{0} \\ \mathbf{G}_{\mathcal{A}} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \mathbf{v} \\ \boldsymbol{\lambda}_{\mathcal{A}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{E}_{\text{eq}} \mathbf{x} \\ \mathbf{w}_{\mathcal{A}} + \mathbf{E}_{\mathcal{A}} \mathbf{x} \end{bmatrix}, \quad (12)$$

where  $\mathbf{G}_{\mathcal{A}}$ ,  $\mathbf{w}_{\mathcal{A}}$ , and  $\mathbf{E}_{\mathcal{A}}$  contain the rows of  $\mathbf{G}_{\text{in}}$ ,  $\mathbf{w}_{\text{in}}$ , and  $\mathbf{E}_{\text{in}}$  corresponding to  $\mathcal{A}$ . The dual variables can be expressed directly in terms of  $\mathbf{z}$ :

$$\begin{bmatrix} \mathbf{G}_{\text{eq}} \mathbf{H}^{-1} \mathbf{G}_{\text{eq}}^\top & \mathbf{G}_{\text{eq}} \mathbf{H}^{-1} \mathbf{G}_{\mathcal{A}}^\top \\ \mathbf{G}_{\mathcal{A}} \mathbf{H}^{-1} \mathbf{G}_{\text{eq}}^\top & \mathbf{G}_{\mathcal{A}} \mathbf{H}^{-1} \mathbf{G}_{\mathcal{A}}^\top \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\lambda}_{\mathcal{A}} \end{bmatrix} = -2 \begin{bmatrix} \mathbf{G}_{\text{eq}} \\ \mathbf{G}_{\mathcal{A}} \end{bmatrix} \mathbf{z}. \quad (13)$$

Finally, we can ensure that  $\boldsymbol{\lambda}_{\mathcal{A}}$  are dual feasible by applying an element-wise  $\max\{\boldsymbol{\lambda}_{\mathcal{A}}, \mathbf{0}\}$ . Given a primal feasible  $\mathbf{z}$ , this always yields dual feasible  $\mathbf{v}, \boldsymbol{\lambda}$  and we can obtain  $\eta(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) = J(\mathbf{z}|\mathbf{x}) - d(\mathbf{v}, \boldsymbol{\lambda}|\mathbf{x})$ . The dual objective  $d(\mathbf{v}, \boldsymbol{\lambda}|\mathbf{x})$  can be computed efficiently by pre-computing  $\frac{1}{2} \mathbf{H}^{-1} [\mathbf{G}_{\text{eq}}; \mathbf{G}_{\text{in}}]^\top$  offline.

#### 4.2. Ensuring feasibility and bounded suboptimality with an online primal active set solver

The procedure in Section 4.1 only checks whether the neural network output  $\mathbf{z} = \tilde{\pi}(\mathbf{x}|\theta)$  is primal feasible and satisfies the suboptimality certificate in (10) but does not provide an approach

### Algorithm 1 Explicit–Implicit Planner

---

**Input:** initial state  $\mathbf{x}$ , NN parameters  $\theta$ , primal active-set QP solver  $\alpha$

- 1: **procedure** EXPLICITIMPLICITPLANNER( $\mathbf{x}, \theta, \alpha$ )
- 2:   Obtain neural network prediction  $\mathbf{z} = \tilde{\pi}(\mathbf{x}|\theta)$
- 3:   **if**  $\mathbf{z}$  is not primal feasible **then** ▷ Section 4.1
- 4:     Perform Phase I (feasibility) of  $\alpha$
- 5:   **while not** CERTIFYSUBOPTIMALITY( $\mathbf{x}, \mathbf{z}$ ) **do**
- 6:     Update  $\mathbf{z}$  via one Phase II iteration of  $\alpha$
- 7:   **return**  $\mathbf{z}$
- 8: **procedure** CERTIFYSUBOPTIMALITY( $\mathbf{x}, \mathbf{z}$ )
- 9:   Obtain dual feasible variables  $\mathbf{v}, \boldsymbol{\lambda}$  ▷ Section 4.1
- 10:   **if**  $\eta(\mathbf{z}, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{x}) \leq \mathbf{x}^\top \mathbf{Q} \mathbf{x}$  **then return** True ▷ Eq. (10)
- 11:   **else return** False

---

to modify  $\mathbf{z}$  in case any of the two requirements is violated. We use a primal active-set QP solver (Wright & Nocedal, 1999, Ch. 16) online to ensure satisfaction of both conditions. Primal active-set QP methods can be accelerated via warm start from a good initialization point, such as the neural network output  $\mathbf{z}$ .

A primal active-set method employs two phases. In Phase I, a linear feasibility program is solved to ensure that  $\mathbf{z}$  is primal feasible. Starting from this primal feasible point, Phase II updates the primal solution and the active constraints  $\mathcal{A}$  and solves a sequence of equality-constrained QPs. Phase II maintains primal feasibility throughout the iterations and, hence, can be terminated early, as soon as the suboptimality certificate in (11) is achieved, rather than continuing until optimality. We check the duality gap at intermediate iterates using the approach described in Section 4.1 and terminate Phase II as soon as the suboptimality certificate is obtained. Compared to other QP solvers, our approach accelerates online computation by reducing the initial iterations using good initializations from the neural network. It also reduces the final iterations using the early termination condition in (11). Interior-point methods are difficult to warm start (John & Yildirim, 2008), while non-primal active-set methods do not guarantee primal feasibility of intermediate iterates and cannot be terminated early (Ferreau, Bock, & Diehl, 2008). In general, however, the neural network output can be used to accelerate any QP solver which benefits from a good but potentially infeasible initialization point.

Alg. 1 presents our final explicit–implicit planner  $\pi$ . It is guaranteed to terminate since primal active-set QP methods reach an optimal solution in finite time for strictly convex QPs (Wright & Nocedal, 1999, Ch. 16) and the associated RHC  $\mu$  is guaranteed to be recursively feasible and asymptotically stable by Theorem 4.

### 5. Scaling to large systems

The efficiency of Alg. 1 online depends on the neural network  $\tilde{\pi}(\mathbf{x}|\theta)$  accurately approximating an optimal planner  $\pi^*(\mathbf{x})$  for the mp-QP in (4). If  $\tilde{\pi}(\mathbf{x}|\theta)$  approximates  $\pi^*(\mathbf{x})$  perfectly, Alg. 1 does not need to perform any computations online. However, if  $\tilde{\pi}(\mathbf{x}|\theta)$  is a poor approximation of  $\pi^*(\mathbf{x})$ , Alg. 1 may need to perform many Phase I and Phase II iterations.

Providing a large train set  $\mathcal{D} = \{\mathbf{x}_i, \mathbf{z}_i^*, \mathbf{v}_i^*, \boldsymbol{\lambda}_i^*\}_i$  is important for good neural network performance and generalization. This data set is generated offline using a QP solver to obtain optimal solutions  $(\mathbf{z}_i^*, \mathbf{v}_i^*, \boldsymbol{\lambda}_i^*)$  for initial states  $\mathbf{x}_i$  that render (4) feasible. The main challenge is that the set  $\mathcal{X}_0$  of initial states that make (4) feasible cannot be described explicitly, e.g., via halfspace or vertex representation. This makes sampling  $\mathbf{x}_i$  from  $\mathcal{X}_0$  challenging, especially for high-dimensional systems with long planning horizons.

The set  $\mathcal{X}_0$  can only be described by a membership oracle (Vempala, 2005), i.e., for a given  $\mathbf{x} \in \mathcal{X}$ , a QP solver can

**Algorithm 2** Data Set Generation

**Input:** number of goal points  $N_{tm}$ ,  $N_{bf}$ ,  $N_{tst}$  for the train set, buffer set, and test set; step size  $d > 0$ ; polyhedron  $\mathcal{X}$ ; primal active-set QP solver  $\alpha$

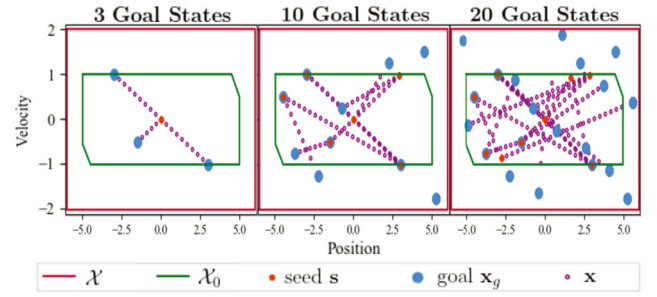
- 1: **procedure** GENERATEDATA( $N_{tm}$ ,  $N_{bf}$ ,  $N_{tst}$ ,  $d$ ,  $\mathcal{X}$ ,  $\alpha$ )
- 2: Let  $\mathbf{s}_0 = (\mathbf{x}_0, \mathbf{z}_0, \mathbf{v}_0, \boldsymbol{\lambda}_0, \mathbf{a}_0)$  be an initial seed with optimal primal, dual, and auxiliary solver variables
- 3: Let  $\mathcal{G}_{tm} \cup \mathcal{G}_{bf} \cup \mathcal{G}_{tst} \subset \mathcal{X}$  be a set of  $N_{tm} + N_{bf} + N_{tst}$  points of a Sobol sequence in  $\mathcal{X}$
- 4:  $\mathcal{S}_{tm}, \mathcal{D}_{tm} = \text{RANDOMWALK}(\mathcal{G}_{tm}, \{\mathbf{s}_0\}, d, \alpha)$
- 5:  $\mathcal{S}_{bf}, \mathcal{D}_{bf} = \text{RANDOMWALK}(\mathcal{G}_{bf}, \mathcal{S}_{tm}, d, \alpha)$
- 6:  $\mathcal{S}_{tst}, \mathcal{D}_{tst} = \text{RANDOMWALK}(\mathcal{G}_{tst}, \mathcal{S}_{bf} \setminus \mathcal{S}_{tm}, d, \alpha)$
- 7: **return**  $\mathcal{D}_{tm}, \mathcal{D}_{tst}$
- 8: **procedure** RANDOMWALK( $\mathcal{G}, \mathcal{S}, d, \alpha$ )
- 9:  $\mathcal{D} = \emptyset$
- 10: **for**  $\mathbf{x}_g \in \mathcal{G}$  **do**
- 11: Sample seed tuple  $\mathbf{s} \in \mathcal{S}$
- 12:  $\mathcal{D}' = \text{LINEQSOLVE}(\mathbf{s}, \mathbf{x}_g, d, \alpha)$
- 13:  $\mathcal{D} = \mathcal{D} \cup \mathcal{D}'$
- 14: Place last seed from  $\mathcal{D}'$  in  $\mathcal{S}$
- 15: **return**  $\mathcal{S}, \mathcal{D}$
- 16: **procedure** LINEQSOLVE( $\mathbf{s}_0, \mathbf{x}_g, d, \alpha$ )
- 17: Let  $\mathbf{x}$  be the first element of  $\mathbf{s}_0 = (\mathbf{x}_0, \mathbf{z}_0, \mathbf{v}_0, \boldsymbol{\lambda}_0, \mathbf{a}_0)$
- 18:  $n = \|\mathbf{x}_g - \mathbf{x}\|$ ,  $\mathbf{x}_n = (\mathbf{x}_g - \mathbf{x})/n$
- 19: **for**  $i = 1 \dots \lceil n/d \rceil$  **do**
- 20:  $\mathbf{x}_i = \mathbf{x} + i d \mathbf{x}_n$
- 21: Let QP( $\mathbf{x}_i$ ) be problem Eq. (4) with parameter  $\mathbf{x}_i$
- 22: (success,  $\mathbf{s}_i$ ) = Hot start  $\alpha$  with  $\mathbf{s}_{i-1}$  on QP( $\mathbf{x}_i$ )
- 23: **if** success **then**  $\mathcal{D} = \mathcal{D} \cup \mathbf{s}_i$
- 24: **else break**
- 25: **return**  $\mathcal{D}$

report whether (4) is feasible, i.e.,  $\mathbf{x} \in \mathcal{X}_0$  or not. Unfortunately, a simple approach, such as *rejection sampling*, which generates samples  $\mathbf{x}_i$  uniformly from  $\mathcal{X}$  and only keeps those that are feasible according to the QP solver, cannot be used because the probability of sampling feasible  $\mathbf{x}_i$  decreases quickly with the system dimension and the number of constraints in  $\mathcal{X}$ ,  $\mathcal{U}$ , and  $\mathcal{X}_f$ .

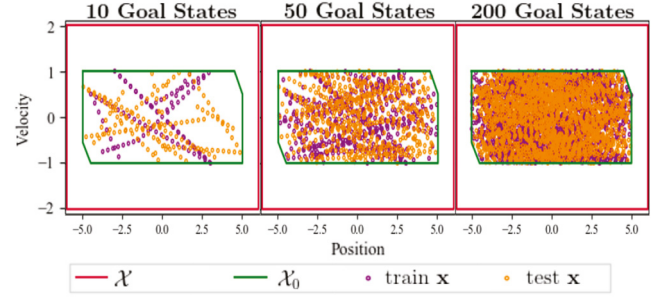
Instead of independent sampling from  $\mathcal{X}$ , we propose a *geometric random walk* technique which generates new samples based on previous successful samples from  $\mathcal{X}_0$ . The proposed approach starts with one feasible sample  $\mathbf{x}_0 \in \mathcal{X}_0$ , e.g., chosen as the system equilibrium state. We pick a random line  $l$  at the current feasible point and iteratively take small steps along the chord  $l \cap \mathcal{X}_0$ . At each potential point  $\mathbf{x}_i$  along  $l$ , we check whether the QP in (4) with parameter  $\mathbf{x}_i$  is feasible. If yes, the algorithm moves to  $\mathbf{x}_i$  and proceeds along  $l$ , otherwise, it stays at  $\mathbf{x}_{i-1}$  and picks a new random direction. To select random directions, we generate a set of goal states  $\mathbf{x}_g \in \mathcal{X}$  using a Sobol sequence (Sobol', 1967). Alg. 2 summarizes the proposed approach for data set generation. It generates one large data set and then splits it into a train set, buffer set, and test set. Due to the sequential nature of the data set generation, the buffer set is needed to ensure that the train and test sets do not contain overlapping seed and goal points. Fig. 1 illustrates the behavior of Alg. 2 for a double integrator system described in Section 6.1.

**6. Evaluation**

We demonstrate the proposed approach on four systems of increasing complexity. The experiments highlight the challenges that arise when scaling to high-dimensional systems. We use MPT3 (Herceg et al., 2013) to compute the terminal cost and constraints in (3), Tensorflow (Abadi et al., 2016) to train and evaluate the neural network planner  $\tilde{\pi}(\mathbf{x}|\theta)$ , and the primal active-set



(a) **Space Filling Generator:** Data generation procedure with varying number of goal states.



(b) **Train and Test Distribution:** Illustration of the train and test generation. Both fill  $\mathcal{X}_0$  with enough goal states.

**Fig. 1.** Illustration of the data set generation algorithm (Alg 2) for a double integrator system defined in Section 6.1.

System	Layer Widths	Num. Parameters
Sys. 1	<b>2</b> , 32, 32, <b>30</b>	2,142
Sys. 2	<b>12</b> , 32, 32, <b>300</b>	89,900
Sys. 3	<b>12</b> , 32, 64, 128, 256, <b>450</b>	159,522
Sys. 4	<b>36</b> , 128, 128, 256, 256, 512, 512, <b>2250</b>	1,668,554

**Fig. 2.** Neural network architectures. Each layer is fully connected and followed by the ReLU nonlinearity. Bolded are the input and output sizes of the network.

QP solver SQOPT (Gill, Murray, Saunders, & Wong, 2018) to implement the Phase I and Phase II iterations in Alg. 1. The neural network models were trained using dual Xeon E5-2683 v4 CPUs with 32 cores for data generation and two NVIDIA Titan X Pascal GPUs for training. The online inference speed evaluation was performed on an i7-7700K CPU with 4 cores and NVIDIA GeForce GTX 1080 Ti GPU.

The networks were trained with the Lagrangian loss function  $\ell(\theta)$  using the Adam optimizer (Kingma & Ba, 2015) for 100 epochs. The network architectures and training parameters are listed in Fig. 2. The network depths were chosen using techniques described in Chen et al. (2018) and the widths were chosen by trial and error. Other methods to choose network architectures are described in Lucia and Karg (2018). Given the large sizes of the train data sets (see Fig. 4), 100 epochs is sufficient to reach convergence. Depending on the system, we take 2.5–10% samples from the train set as a held-out validation set. The validation set is only used to evaluate the training loss to avoid over- or under-fitting and is not used to update the neural network weights.

System	$n$	$m$	$N$	$c_x$	$c_f$	$c_u$	$d_p$	$d_{in}$	$d_{eq}$
Sys. 1	2	1	10	4	6	2	30	66	20
Sys. 2	12	3	20	24	246	6	300	846	240
Sys. 3	12	3	30	24	104	6	450	1004	360
Sys. 4	36	9	50	72	408	18	2250	4908	1800

**Fig. 3.** Problem specifications: state dimension ( $n$ ), control dimension ( $m$ ), time horizon ( $N$ ), state constraints ( $c_x$ ), terminal state constraints ( $c_f$ ), control constraints ( $c_u$ ), primal variable dimension ( $d_p$ ), dual variable dimension for inequality constraints ( $d_{in}$ ) and equality constraints ( $d_{eq}$ ).

### 6.1. System descriptions

The dimensions of the four problem settings are summarized in Fig. 3. We choose the terminal region  $\mathcal{X}_f$  to be  $\mathcal{O}_{\infty}^{LQR}$ , and the terminal cost  $\mathbf{x}_N^T \mathbf{P}_{\infty} \mathbf{x}_N$ , defined in (6).

**System 1 (Double Integrator).** A mass under force input can be modeled as a double integrator system,

$$\mathbf{x} \in \mathbb{R}^2, \quad \mathbf{u} \in \mathbb{R}, \quad \mathbf{A} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 0.5 \\ 0.1 \end{bmatrix}.$$

We consider a problem with cost matrices  $\mathbf{Q} = \mathbf{I}_2$ ,  $\mathbf{R} = 1$ , state constraints  $\mathbf{A}_x = [\mathbf{I} \ -\mathbf{I}]^T$ ,  $\mathbf{b}_x = [5 \ 1 \ 5 \ 1]^T$ , and input constraints  $|u_k| \leq 2$ , for  $k = 1, \dots, N$ .

**System 2 (Quadrotor).** The differential flatness of quadrotor robot dynamics enables trajectory generation in the space of flat outputs (position, yaw and their derivatives) that are dynamically feasible for an underactuated quadrotor (Liu, Atanasov, Mohta, & Kumar, 2017). This leads to a time-invariant continuous-time system  $\dot{\mathbf{x}} = \mathbf{A}_c \mathbf{x} + \mathbf{B}_c \mathbf{u}$ , where  $\mathbf{A}_c = \mathbf{L}_4 \otimes \mathbf{I}_3$  and  $\mathbf{B}_c = \mathbf{e}_4 \otimes \mathbf{I}_3$ . The  $3 \times 3$  sub-matrices correspond to position, velocity, acceleration, and jerk. We discretize the system using Euler discretization with time step 0.1, and consider constraints  $\mathbf{A}_x = [\mathbf{I}_{12}, -\mathbf{I}_{12}]^T$ ,  $\mathbf{b}_x = [10 \cdot \mathbf{1}_3, 5 \cdot \mathbf{1}_3, 3 \cdot \mathbf{1}_3, 1 \cdot \mathbf{1}_3, 10 \cdot \mathbf{1}_3, 5 \cdot \mathbf{1}_3, 3 \cdot \mathbf{1}_3, 1 \cdot \mathbf{1}_3]^T$ ,  $\mathbf{A}_u = [\mathbf{I}_3, -\mathbf{I}_3]^T$ ,  $\mathbf{b}_u = [\mathbf{1}_3, \mathbf{1}_3]^T$  and cost matrices  $\mathbf{Q} = \mathbf{I}_{12}$ ,  $\mathbf{R} = \mathbf{I}_3$ .

**System 3 (Oscillating Masses).** Introduced in Wang and Boyd (2010), the oscillating masses is a linear system that can be scaled to large dimensions by increasing the number of masses and springs in the system. We use 6 masses with a mass of 1, and 3 springs with a spring constant  $c = 1$  and damping constant  $d = 0.1$ . Let  $a = -2c$  and  $b = -2$ . The system is defined as  $\dot{\mathbf{x}} = \mathbf{A}_c \mathbf{x} + \mathbf{B}_c \mathbf{u}$ , where

$$\mathbf{A}_c = \begin{bmatrix} \mathbf{0}_6 & \mathbf{I}_6 \\ a\mathbf{I}_6 + c\mathbf{L}_6 + c\mathbf{L}_6^T & b\mathbf{I}_6 + d\mathbf{L}_6 + d\mathbf{L}_6^T \end{bmatrix}, \quad \mathbf{B}_c = \begin{bmatrix} \mathbf{0} \\ \mathbf{F} \end{bmatrix}$$

$$\mathbf{F} = [\mathbf{e}_1 \quad -\mathbf{e}_1 \quad \mathbf{e}_2 \quad \mathbf{e}_3 \quad -\mathbf{e}_2 \quad \mathbf{e}_3]^T \in \mathbb{R}^{6 \times 3}. \quad (14)$$

We discretize this system using first-order hold and time step of 0.5. The cost terms are  $\mathbf{Q} = \mathbf{I}_{12}$ ,  $\mathbf{R} = \mathbf{I}_3$ , subject to state constraints  $\mathbf{A}_x = [\mathbf{I}, -\mathbf{I}]^T$ ,  $\mathbf{b}_x = 4 \cdot \mathbf{1}$ , and input constraints  $\mathbf{A}_u = [\mathbf{I}, -\mathbf{I}]^T$ ,  $\mathbf{b}_u = 0.5 \cdot \mathbf{1}$ .

**System 4 (High Dimensional Oscillating Masses).** The high dimensional oscillating masses system is a scaled-up version of Sys. 3. The matrix  $\mathbf{A}_c$  is the extension of (14) to 36 dimensions obtained by replacing the 6-dimensional submatrices with 18-dimensional counterparts. The input matrix is  $\mathbf{B}_c = [\mathbf{0} \ | \ \mathbf{I}_3 \ \otimes \ \mathbf{F}]^T$ . The cost and constraint terms are chosen similarly as for Sys. 3.

### 6.2. Data generation

As discussed in Section 5,  $\mathcal{X}_0$  cannot be sampled directly. Without an efficient data set generation strategy, obtaining large data sets is practically impossible. For example, a naive rejection-sampling approach yields feasible samples 98.6% of the time for Sys. 1, 16.7% for Sys. 2, 1.1% for Sys. 3, and 0.4% for Sys. 4. Generating a data set of 1 million feasible states for Sys. 4 using rejection sampling will take 2.5 years. In contrast, the geometric random walk approach in Alg. 2 generated a data set of 1.7 million feasible states in 5.2 h. Fig. 4 shows the data set sizes used for the different systems and the computation time for data generation using Alg. 2.

### 6.3. Open-loop metrics

On the test set, we compare our method (magenta) in Fig. 5 to various baseline methods. The baselines are the 6 combinations between 2 initialization methods – NN Warm start and Cold start – and 3 termination criteria – terminating after primal feasibility (*p.f.*), primal feasibility and suboptimality (*p.f. + sub.*), and optimality. *Sub.* is achieved at the time when the duality gap  $\eta(\mathbf{z}, \mathbf{v}, \lambda | \mathbf{x})$  defined in (10) is less than  $\mathbf{x}^T \mathbf{Q} \mathbf{x}$  according to optimality (olive). Fig. 5 displays the results on the four systems according to various initialization and termination criteria. We evaluate speedup by reporting the number of required iterations and evaluate performance by reporting the suboptimality levels. For Sys. 4, compared to the benchmark Cold *optimal* control law, our method (NN *p.f. + sub.*) reduces required iterations by 58 $\times$ . The results demonstrate that early termination is especially synergistic when paired with NN Warm starts, as it reduces the required iterations by 14.4 $\times$  versus only 1.3 $\times$  when paired with Cold starts.

Our method additionally performs well in terms of the suboptimality level. We define *open loop sub.* (%) as  $\sigma_{ol} := \frac{J(\mathbf{z} | \mathbf{x}) - J(\mathbf{z}^* | \mathbf{x})}{J(\mathbf{z}^* | \mathbf{x})}$ . On Sys. 4, it has an average  $\sigma_{ol}$  of 4.9% when only the *p.f.* criterion is met, and 3.5% when both *p.f. + sub.* criteria are met. This is in contrast to the Cold start method, where  $\sigma_{ol}$  is high (192.9%) when only the *p.f.* criterion is met, and obtaining both *p.f. + sub.* criteria is necessary to obtain a reasonable suboptimality (6.3%). This discrepancy demonstrates that the neural network  $\tilde{\pi}$  alone is obtaining a good guess of the optimal solution, and the online iterations serve mostly to satisfy the *p.f.* criterion. For the low dimensional Sys. 1,  $\sigma_{ol}$  of both the NN Warm start and Cold start methods for the *p.f.* criteria is high because the optimal cost is very small ( $\sim 0$ ).

### 6.4. Closed-loop metrics

We compare the closed-loop performance of our method to various benchmarks on Sys. 4 in Fig. 6. While a Cold start is unavoidable at the initial state  $\mathbf{x}(0)$ , we benchmark against Hot start initialization techniques (Ferreau, Kirches, Potschka, Bock, & Diehl, 2014) that exploit solutions of previously solved states for the subsequent states  $\mathbf{x}(t)$ . The first benchmark is the SQOPT solver (Gill et al., 2018) using Hot starts. The other benchmarks include state-of-the-art interior-point QP solvers, ECOS, Gurobi, and Mosek, from CVXPY (Diamond & Boyd, 2016). For the SQOPT methods, we evaluate 3 termination criteria: ours (*p.f. + sub.*), a fixed duality gap threshold  $\eta(\mathbf{z}, \mathbf{v}, \lambda | \mathbf{x}) < 0.1$ , and optimal. The CVXPY methods cannot be terminated early, and warm/hot starts are less effective.

We evaluate on trajectories generated from 128 initial states  $\mathbf{x}(0)$  randomly sampled from the test set. At each subsequent state  $\mathbf{x}(t)$ , we apply the first control  $\mathbf{u}_0$  obtained from each

Sys.	Data Set Sizes			Timing Statistics		
	Train	Buffer	Test	Data Gen.	Network Train	Test
	# Total (Seed) examples in thousands			total	total	per example
1	15 (2)	4 (0.4)	4 (0.4)	1.5 secs	100 secs	1.0 ms
2	863 (200)	173 (40)	173 (40)	3.0 mins	1.4 hrs	3.1 ms
3	174 (20)	35 (4)	35 (4)	1.0 mins	0.4 hrs	3.3 ms
4	1,754 (400)	175 (40)	175 (40)	5.2 hrs	20 hrs	38.8 ms

Fig. 4. Overview of the data set sizes and offline data generation, train, and test time statistics.

Sys.	Neural Network Warm Start			Cold Start		
	<i>p.f.</i>	<i>p.f. + sub.</i>	<i>optimal</i>	<i>p.f.</i>	<i>p.f. + sub.</i>	<i>optimal</i>
	Absolute Iterations Average (Worst)					
1	1.0 (5)	7.1 (15)	12.3 (21)	1.1 (2)	9.5 (12)	11.4 (22)
2	8.2 (44)	17.9 (70)	57.9 (110)	32.5 (97)	62.1 (124)	98.9 (144)
3	5.3 (123)	7.2 (133)	83.6 (145)	224.6 (408)	255.1 (448)	327.1 (498)
4	17.4 (291)	27.8 (391)	400.4 (1115)	1074.4 (2089)	1262.1 (2409)	1611.7 (2559)
	Open Loop Sub. (%) $\sigma_{ol}$ Average (Worst)					
1	1277.6 (41277.2)	9.8 (266.3)	—	201.3 (583.4)	13.0 (277.8)	—
2	1.2 (29.3)	0.3 (7.0)	—	19.7 (161.8)	0.8 (11.1)	—
3	3.1 (16.9)	2.6 (10.1)	—	98.3 (423.4)	6.4 (33.0)	—
4	4.9 (26.4)	3.5 (13.8)	—	192.9 (686.6)	6.3 (22.5)	—

Fig. 5. Open-loop metrics: average and worst case number of iterations and suboptimality. Our method is displayed in magenta, and the baseline method is displayed in olive. The evaluations are performed on independent states in the test set, and the baselines methods can only be initialized from a Cold start. On Sys. 4, our method achieves a 58 $\times$  reduction in required iterations compared to the baseline. Our methods achieve better suboptimality metric  $\sigma_{ol}$  compared to Cold starts due to the higher-quality NN initialization. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Sys. 4	Online Inference Time (ms)			Trajectory Sub. (%) $\sigma_{cl}$	
Method	Network	Initial $\mathbf{x}(0)$	Remaining $\mathbf{x}(t)$		Total
Average (Worst)					
NN <i>p.f. + sub.</i>	11 (65)	29 (168)	15 (92)	664 (1678)	8.8 (15.5)
NN ( $\eta < 0.1$ )	12 (63)	114 (168)	88 (216)	2185 (5380)	0.0 (0.0)
NN <i>optimal</i>	12 (63)	200 (221)	193 (259)	4420 (8355)	—
Hot <i>p.f. + sub.</i>	—	447 (576)	41 (186)	1311 (2988)	0.6 (1.5)
Hot ( $\eta < 0.1$ )	—	543 (644)	69 (288)	1962 (3873)	0.0 (0.0)
Hot <i>optimal</i>	—	593 (676)	89 (474)	2409 (4259)	—
CVXPY (ECOS)	—	276 (1541)	216 (386)	5114 (10141)	—
CVXPY (Gurobi)	—	147 (390)	151 (192)	3525 (6190)	—
CVXPY (Mosek)	—	182 (1437)	161 (225)	3789 (7032)	—

Fig. 6. Closed-loop metrics (Sys. 4): Average and worst case time required to implement the RHC for Warm starts, Hot starts, and Cold starts. The closed-loop trajectory suboptimality metric  $\sigma_{cl}$  (defined after (15)) is different from the open-loop suboptimality metric  $\sigma_{ol}$ . The former is a cumulative measure computed along the entire closed-loop trajectory, while the latter is an instantaneous measure at a given state  $\mathbf{x}$ . The CVXPY benchmarks are state-of-the-art QP solvers that provide a baseline for the online inference speed of the active set methods. The results show that relative to Hot starts, the NN Warm starts quickly reach the conditions specified in Theorem 4 (represented by *p.f. + sub.*). The NN Warm starts then switch from faster to slower relative to Hot starts captured at the fixed duality gap threshold  $\sigma(\mathbf{z}|\mathbf{x}) \leq \eta(\mathbf{z}, \mathbf{v}, \lambda|\mathbf{x}) < 0.1$ . The NN methods proceed slowly to the precise thresholds needed for optimality. Practically, these latter iterations are unnecessary, because when NN and Hot are terminated with  $\eta < 0.1$ , the resulting primal variables are essentially optimal, and the closed-loop trajectory matches the optimal one with  $\sigma_{cl} = 0.0$ . While NN *p.f. + sub.* does have a higher  $\sigma_{cl}$  compared to Hot *p.f. + sub.*, this difference indicates the effectiveness of pairing early termination criteria with NN warm starts, as Hot starts need to bear additional computation time to reach near optimal solutions before obtaining the required termination certificates.

method. We repeatedly execute this closed-loop controller until the system reaches a state  $\mathbf{x}(t) \in \mathcal{X}_f$ , at which point each method switches to an optimal LQR controller (Borrelli et al., 2017, Ch. 8). Inference speed is evaluated on states  $\mathbf{x}(t) \notin \mathcal{X}_f$ , and these states will vary across each method. Optimality performance is evaluated on the entire infinite-horizon trajectory:

$$J_{cl} = \sum_{t=0}^{\infty} \mathbf{x}^T(t) \mathbf{Q} \mathbf{x}(t) + \mathbf{u}^T(t) \mathbf{R} \mathbf{u}(t), \quad (15)$$

where the trajectory suboptimality is  $\sigma_{cl} = \frac{J_{cl} - J_{cl}^*}{J_{cl}^*}$ .

For methods that terminate early, rather than performing the suboptimality check at every iteration, it is more efficient to perform the check periodically. In our implementation with SQOPT, API restrictions require the suboptimality checks to be performed externally to the solver, and as a result, we periodically terminate the solver to check the intermediate iterates. We utilize the SQOPT API optimality tolerance parameter, which evaluates the size of the reduced gradients (Gill et al., 2018, pg. 39) to

prematurely terminate the solver and then perform the suboptimality check with the procedure described in Section 4.1 at each resulting iterate. If we are unable to obtain a *sub.* certificate, we reduce the optimality tolerance parameter by 2 and repeat from the current SQOPT solver state. For consistency, we apply the same termination strategy across all evaluated methods.

The closed-loop experiments with Hot starts (Fig. 6) exhibit similar patterns to the open-loop experiments with Cold starts (Fig. 5). On average, across the entire trajectories, our method achieves a  $2\times$  speedup compared to Hot *p.f. + sub.*, and a  $3.6\times$  speedup compared to Hot *optimal*. Similar to the open loop experiments, the closed loop experiments also demonstrate that early termination is especially synergistic with NN Warm starts compared to Hot starts. Versus optimal, early termination reduces inference time by  $12.9\times$  for NN compared to  $2.2\times$  for Hot. Further analysis suggests that early termination works well with NN because NN initialized methods reach approximate thresholds quickly, and then slowly reach the precise thresholds required for optimality. While, NN *p.f. + sub.* is  $2\times$  faster than Hot *p.f. + sub.*, it becomes  $1.8\times$  slower when solved to optimality. To demonstrate this switch, Fig. 6 evaluates each initialization method with a fixed duality gap threshold  $\eta < 0.1$  which shows that the NN Warm start inference time is only  $1.1\times$  slower than Hot. Importantly, when terminating with  $\eta < 0.1$ , the trajectories have effectively 0 suboptimality, demonstrating that improving from an approximate threshold to a precise optimality threshold has little benefit.

The  $\sigma_{cl}$  metrics in Fig. 6 further highlight synergies between NN Warm starts and early termination. NN *p.f. + sub.* has  $\sigma_{cl}$  of 8.8% compared to 0.6% for Hot *p.f. + sub.* While this may initially seem like a drawback for the NN methods, it actually demonstrates that early termination is not as effective in reducing online inference time for Hot starts, as satisfying the termination criteria requires near optimal solutions. The motivation of these methods is fast online inference speed with guarantees of feasibility and stability, not in achieving optimality.

## 7. Conclusion

We presented a hybrid explicit–implicit MPC procedure that combines an offline trained neural network with an online primal active set solver. We proposed a primal–dual loss function based on the Lagrangian to train the neural network. Using the primal variable predictions, we derived an algorithm to provide certificates of primal feasibility and suboptimality, the criteria necessary to guarantee recursive feasibility and asymptotic stability. Finally, we demonstrated how warm start and early termination can combine the primal active set solver with the neural network to accelerate inference times. The key challenge of function approximation in high dimensional mp-QPs is choosing good training points in a convex set defined by a membership oracle. We introduced a geometric random walk algorithm that can efficiently generate a data set for large problems. Our results indicate the importance of addressing this data generation problem to obtain a scalable solution. The combination of these ideas yields an RHC with guarantees on recursive feasibility and asymptotic stability, while achieving a  $2\times$  speedup versus the best benchmark method on a system with thousands of optimization variables.

## Acknowledgment

We are grateful to Elizabeth Wong for answering our many questions about SQOPT and providing guidance on initialization and accurate timing.

## References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). Tensorflow: A system for large-scale machine learning. In *Symposium on operating systems design and implementation*, vol. 16 (pp. 265–283).
- Arora, R., Basu, A., Mianjy, P., & Mukherjee, A. (2018). Understanding deep neural networks with rectified linear units. In *International conference on learning representations*. URL [https://openreview.net/forum?id=BIJ\\_rgWRW](https://openreview.net/forum?id=BIJ_rgWRW).
- Borrelli, F., Bemporad, A., & Morari, M. (2017). *Predictive control for linear and hybrid systems*. Cambridge University Press.
- Bouffard, P., Aswani, A., & Tomlin, C. (2012). Learning-based model predictive control on a quadrotor: Onboard implementation and experimental results. In *IEEE int. conf. on robotics and automation* (pp. 279–284).
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.
- Chen, S., Saulnier, K., Atanasov, N., Lee, D. D., Kumar, V., Pappas, G. J., et al. (2018). Approximating explicit model predictive control using constrained neural networks. In *American control conference* (pp. 1520–1527).
- Diamond, S., & Boyd, S. (2016). CVXPY: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 1–5.
- Erez, T., Lowrey, K., Tassa, Y., Kumar, V., Kolev, S., & Todorov, E. (2013). An integrated system for real-time model predictive control of humanoid robots. In *IEEE int. conf. on humanoid robots* (pp. 292–299).
- Ferreau, H., Bock, H., & Diehl, M. (2008). An online active set strategy to overcome the limitations of explicit MPC. *International Journal of Robust and Nonlinear Control*, 18(8), 816–830.
- Ferreau, H. J., Kirches, C., Potschka, A., Bock, H. G., & Diehl, M. (2014). QpOASES: A parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation*, 6(4), 327–363.
- Gill, P. E., Murray, W., Saunders, M. A., & Wong, E. (2018). *User's guide for SQOPT 7.7: Software for large-scale linear and quadratic programming: Technical report*, Department of Mathematics, UCSD.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Herceg, M., Kvasnica, M., Jones, C., & Morari, M. (2013). Multi-parametric toolbox 3.0. In *European control conference* (pp. 502–510).
- Hertneck, M., Köhler, J., Trimpe, S., & Allgöwer, F. (2018). Learning an approximate model predictive controller with guarantees. *IEEE Control Systems Letters*, 2(3), 543–548.
- John, E., & Yıldırım, E. A. (2008). Implementation of warm-start strategies in interior-point methods for linear programming in fixed dimension. *Computational Optimization and Applications*, 41(2), 151–183.
- Jones, C. N., & Morari, M. (2010). Polytopic approximation of explicit model predictive controllers. *IEEE Trans. on Auto. Control*, 55(11), 2542–2553.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations*.
- Klaučo, M., Kalúz, M., & Kvasnica, M. (2019). Machine learning-based warm starting of active set methods in embedded model predictive control. *Engineering Applications of Artificial Intelligence*, 77, 1–8.
- Kvasnica, M., & Fikar, M. (2012). Clipping-based complexity reduction in explicit MPC. *IEEE Trans. on Auto. Control*, 57(7), 1878–1883.
- Liu, S., Atanasov, N., Mohta, K., & Kumar, V. (2017). Search-based motion planning for quadrotors using linear quadratic minimum time control. In *IEEE/RSJ int. conf. on intelligent robots and systems* (pp. 2872–2879).
- Lucia, S., & Karg, B. (2018). A deep learning-based approach to robust nonlinear model predictive control. *IFAC-PapersOnLine*, 51(20), 511–516.
- Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7), 733–764.
- Richter, C., Vega-Brown, W., & Roy, N. (2018). Bayesian learning for safe high-speed navigation in unknown environments. In *Robotics research*, vol. 2 (pp. 325–341). Springer.
- Ross, S., Gordon, G., & Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In *Int. conf. on artificial intelligence and statistics* (pp. 627–635).
- Sobol', I. M. (1967). On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4).
- Summers, S., Jones, C. N., Lygeros, J., & Morari, M. (2011). A multiresolution approximation method for fast explicit model predictive control. *IEEE Transactions on Automatic Control*, 56(11), 2530–2541.
- Vempala, S. (2005). Geometric random walks: a survey. *Combinatorial & Computational Geometry*, 52, 577–616.
- Vichik, S., & Borrelli, F. (2014). Solving linear and quadratic programs with an analog circuit. *Computers & Chemical Engineering*, 70, 160–171.
- Wang, Y., & Boyd, S. (2010). Fast model predictive control using online optimization. *IEEE Transaction on Control System Technology*, 18(2), 267–278.
- Watterson, M., & Kumar, V. (2015). Safe receding horizon control for aggressive MAV flight with limited range sensing. In *IEEE/RSJ int. conf. on intelligent robots and systems* (pp. 3235–3240).




- Wright, S., & Nocedal, J. (1999). Numerical optimization. *Springer Science*, 35(67–68), 7.
- Zeilinger, M. N., Jones, C. N., & Morari, M. (2011). Real-time suboptimal model predictive control using a combination of explicit MPC and online optimization. *IEEE Transactions on Automatic Control*, 56(7), 1524–1534.
- Zhang, X., Bujarbaruah, M., & Borrelli, F. (2019). Safe and near-optimal policy learning for model predictive control using primal-dual neural networks. *American Control Conference*, 354–359.



**Steven W. Chen** received the M.S.E. degree in Robotics from the University of Pennsylvania in 2019, and the B.S. degree in economics, mathematics, and mathematical methods in the social sciences (MMSS) at Northwestern University in 2012. He is a Ph.D. student in Computer and Information Science at the GRASP Laboratory, University of Pennsylvania, and the co-founder and CEO of Treeswift, Philadelphia, PA, USA. His research focuses on robotics, control theory, and machine learning, applied to perception, action, and communication using ground and aerial robots. He

works on applications in precision agriculture and forestry, specifically in developing algorithms to deploy UAVs for autonomous under-canopy flights and semantic mapping.



**Tianyu Wang** received the B.S. (2016) in physics from Haverford College, PA, and the M.S. (2018) in electrical and computer engineering from University of California, San Diego, where he is currently pursuing a Ph.D. degree. He has held a summer position at Waymo LLC. His research interests include reinforcement learning, inverse reinforcement learning and autonomous driving.



**Nikolay Atanasov** is an Assistant Professor of Electrical and Computer Engineering at the University of California San Diego. He obtained a B.S. degree in Electrical Engineering from Trinity College, Hartford, CT, in 2008 and M.S. and Ph.D. degrees in Electrical and Systems Engineering from the University of Pennsylvania, Philadelphia, PA, in 2012 and 2015, respectively. His research focuses on robotics, control theory, and machine learning, applied to active sensing using ground and aerial robots. He works on probabilistic environment models that unify geometry and semantics and on optimal control and reinforcement learning approaches for minimizing uncertainty in these models. Dr. Atanasov's work has been recognized by the Joseph and Rosaline Wolf award for the best Ph.D. dissertation in Electrical and Systems Engineering at the University of Pennsylvania in 2015, the best conference paper award at the International Conference on Robotics and Automation in 2017, and the NSF CAREER award in 2021.

works on applications in precision agriculture and forestry, specifically in developing algorithms to deploy UAVs for autonomous under-canopy flights and semantic mapping.



**Vijay Kumar** received the B.Tech. degree from Indian Institute of Technology, Kanpur, India, in 1983, and the Ph.D. degree from The Ohio State University, Columbus, OH, USA, in 1987, both in mechanical engineering. He is currently the Nemirovsky Family Dean of Penn Engineering with appointments with the Department of Mechanical Engineering and Applied Mechanics, Department of Computer and Information Science, and Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. Since 1987, he has been on the Faculty with the Department

of Mechanical Engineering and Applied Mechanics with a secondary appointment in the Department of Computer and Information Science, University of Pennsylvania. He was the Deputy Dean for Research with the School of Engineering and Applied Science from 2000 to 2004. He was the Director of the General Robotics, Automation, Sensing and Perception Laboratory, a multidisciplinary robotics and perception laboratory, from 1998 to 2004. He was the Chairman of the Department of Mechanical Engineering and Applied Mechanics from 2005 to 2008. He was the Deputy Dean for Education with the School of Engineering and Applied Science from 2008 to 2012. He also was the Assistant Director of Robotics and Cyber Physical Systems, White House Office of Science and Technology Policy (2012–2013). His research interests include robotics, specifically multirobot systems, and microaerial vehicles. Dr. Kumar is a Fellow of the American Society of Mechanical Engineers (2003) a member of the National Academy of Engineering (2013). He has been on the editorial boards of IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION, IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING, ASME Journal of Mechanical Design, ASME Journal of Mechanisms and Robotics, and Springer Tract in Advanced Robotics. He is currently the Editor of ASME Journal of Mechanisms and Robotics and an Advisory Board Member of AAAS Science Robotics Journal. He was the recipient of the 1991 National Science Foundation Presidential Young Investigator award, the 1996 Lindback Award for Distinguished Teaching (University of Pennsylvania), the 1997 Freudenstein Award for significant accomplishments in mechanisms and robotics, the 2012 ASME Mechanisms and Robotics Award, the 2012 IEEE Robotics and Automation Society Distinguished Service Award, the 2012 World Technology Network Award, and the 2014 Engelberger Robotics Award. He was the recipient of the best paper awards at DARS 2002, ICRA 2004, ICRA 2011, RSS 2011, and RSS 2013, and has advised doctoral students who have won Best Student Paper Awards at ICRA 2008, RSS 2009, and DARS 2010.



**Manfred Morari** received the Diploma degree in chemical engineering from ETH Zürich, Zürich, Switzerland, and the Ph.D. degree in chemical engineering from the University of Minnesota, Minneapolis, MN, USA. He was a Professor and the Head of the Department of Information Technology and Electrical Engineering, ETH Zürich. He was the McCollum-Corcoran Professor of chemical engineering and the Executive Officer of control and dynamical systems with the California Institute of Technology (Caltech), Pasadena, CA, USA. He was a Professor with the University of Wisconsin, Madison,

WI, USA. He is currently with the University of Pennsylvania, Philadelphia, PA, USA. He supervised more than 80 Ph.D. students. Dr. Morari is a fellow of AIChE, IFAC, and the U.K. Royal Academy of Engineering. He is a member of the U.S. National Academy of Engineering. He was a recipient of numerous awards, including Eckman, Ragazzini, and Bellman Awards from the American Automatic Control Council (AACC); Colburn, Professional Progress, and CAST Division Awards from the American Institute of Chemical Engineers (AIChE); Control Systems Award and Bode Lecture Prize from IEEE; Nyquist Lectureship and Oldenburger Medal from the American Society of Mechanical Engineers (ASME); and the IFAC High Impact Paper Award. He was the President of the European Control Association. He served on the technical advisory boards of several major corporations.