

Learning Deep Control Policies for Autonomous Aerial Vehicles with MPC-Guided Policy Search

Tianhao Zhang, Gregory Kahn, Sergey Levine, Pieter Abbeel

Abstract—Model predictive control (MPC) is an effective method for controlling robotic systems, particularly autonomous aerial vehicles such as quadcopters. However, application of MPC can be computationally demanding, and typically requires estimating the state of the system, which can be challenging in complex, unstructured environments. Reinforcement learning can in principle forego the need for explicit state estimation and acquire a policy that directly maps sensor readings to actions, but is difficult to apply to unstable systems that are liable to fail catastrophically during training before an effective policy has been found. We propose to combine MPC with reinforcement learning in the framework of guided policy search, where MPC is used to generate data at training time, under full state observations provided by an instrumented training environment. This data is used to train a deep neural network policy, which is allowed to access only the raw observations from the vehicle’s onboard sensors. After training, the neural network policy can successfully control the robot without knowledge of the full state, and at a fraction of the computational cost of MPC. We evaluate our method by learning obstacle avoidance policies for a simulated quadrotor, using simulated onboard sensors and no explicit state estimation at test time.

I. INTRODUCTION

Model predictive control (MPC) is an effective and reliable method for controlling robotic systems, particularly autonomous aerial vehicles such as quadcopters, because of its robustness to moderate model errors [26], ability to use high-level objectives [41], and relative simplicity. However, applications of MPC can be computationally demanding, and typically require estimating the state of the system. The state estimation problem can be quite challenging in complex, unstructured environments. Reinforcement learning can in principle forego the need for explicit state estimation and acquire a policy that directly maps sensor readings to actions [5]. The power of reinforcement learning is derived from its ability to learn directly from the real-world behavior of the system. Unfortunately, this strength is also its major weakness when applied to unstable, fragile systems such as aerial vehicles, which can be damaged beyond repair by an unsuccessful, partially trained policy (e.g. by crashing into an obstacle). While alternative learning methods, such as learning from demonstration [36], [37], can address this issue, they typically require costly additional information, such as guidance from a human expert.

We propose to use an off-policy guided policy search algorithm in combination with a model predictive control




Fig. 1: Diagram of our method: the training phase alternates between running MPC to attempt the task and collect data under full state observation, and using this data to train a neural network policy that chooses actions based only on the vehicle’s onboard sensors. At test time, the neural network does not need the full state, and can control the vehicle in unstructured environments.

(MPC) scheme to train policies for autonomous aerial vehicles in a way that avoids catastrophic failure at training time. Guided policy search transforms RL into supervised learning, where an optimal control algorithm provides the supervision and the final control policy is trained with supervised learning. Typically, this optimal control algorithm is an offline trajectory optimization procedure, which either assumes a known model of the dynamics [21] or uses an iteratively learned model [19]. Both approaches are prone to failure during training because the known model may be inaccurate while the learned model is always inaccurate during the early stages of learning. By substituting MPC for offline trajectory optimization, we can obtain a variant of guided policy search that is robust to moderate model errors, and thus avoid catastrophic failures during training. Furthermore, since the final policy is trained with supervised learning, we can train complex, high-dimensional, and highly nonlinear policies, such as deep neural networks, which can represent a wide range of complex behaviors.

One might wonder why the guided policy search method is necessary if we already have access to an effective MPC procedure. In the case of autonomous aerial vehicles, training deep neural network policies with guided policy search affords us two main advantages. First, the neural network policy does not need to use the same inputs as MPC. In fact,

we can restrict its inputs to only those observations that are directly available from the vehicle’s onboard sensors, such as IMU readings and data from laser range finders, while the MPC training phase uses the full state of the system. Since the policy is represented by a deep neural network, it can even process complex, raw sensor information. For example, prior work has shown that guided policy search can learn policies that directly use camera images [23]. Since MPC is only used at training time, we can employ an instrumented training setup, where the full state is known at training time (e.g. using motion capture), but unavailable at test time. This instrumented training setup is one of the key benefits of our approach, since it allows for safe training with full state information, but still produces a policy that uses raw sensor readings and does not require explicit state estimation. The second benefit of this approach is that the final neural network policy is computationally much less expensive than MPC, and can be easily parallelized on specialized hardware. This advantage combines elegantly with the instrumented training setup, since the MPC solution can be computed offboard during training, while all policy computations may be performed onboard at test time.

Our main contribution is an MPC-guided policy search algorithm that can be used for learning control policies for autonomous aerial vehicles. This algorithm, illustrated in Figure 1, replaces the offline trajectory optimization that is typically used in guided policy search with online MPC, which continuously replans paths to the goal from the vehicle’s current state using an approximate model of the dynamics. Our modified MPC procedure also takes into account the actions that would be taken at each state by the current neural network in order to avoid actions that the network is unlikely to take. This ensures that, at convergence, the neural network achieves good long-horizon performance, despite being trained only with supervised learning. Our approach allows us to learn neural network policies that directly process raw observations from the vehicle’s onboard sensors, and are substantially faster to evaluate at test time than full MPC solutions. We demonstrate our method on a set of simulated quadrotor control tasks, including obstacle avoidance using simulated laser range sensors. We show that our approach can learn policies that are robust to a variety of perturbations and generalize successfully to large obstacle courses, without catastrophic failure during training.

II. RELATED WORK

Model predictive control (MPC) is an effective and popular technique for control of robotic systems, and is frequently used to control autonomous aerial vehicles such as quadrotors [38], [35], [34], [3], [2], [4], [29]. MPC is straightforward to apply when the state of the system is known (e.g. via a motion capture system), or when it can be measured accurately through sensors with well-understood observation models. However, vehicles navigating complex, unstructured environments must use more complex sensors, such as cameras and laser range finders. Incorporating such sensors into optimal control directly is challenging, since the

sensor reading depends on a complex and often unknown environment. This challenge is conventionally addressed by using localization and mapping algorithms to map out the environment [9] and then optimizing trajectories under the resulting map [11]. However, this kind of model-based approach is quite challenging when the vehicle is moving at high speed, or when onboard computation is limited.

On the other end of the spectrum from such model-based methods, reinforcement learning (RL) aims to directly learn control policies that map observations to controls [5]. This approach in principle removes the need for explicit state estimation and extensive computation at test time, by using a number of training episodes to iteratively improve the policy from real-world experience. RL has been used to train robotic controllers for games such as ball-in-cup and table tennis [16], manipulation [31], [6], and robotic locomotion [17], [40], [10], [8]. An overview of recent reinforcement learning methods in robotics can be found in a recent survey paper [15]. However, model-free RL is difficult to apply to unstable systems such as quadrotors, due to the possibility of catastrophic failure during training. Model-based RL can mitigate this problem by training a model from real-world experience, and then optimizing the policy under this model. While such methods have been successfully applied to aerial vehicles [1], [28], the requirement to be able to acquire an accurate model means that these methods share many of the challenges of MPC methods.

In this work, we use an off-policy reinforcement learning method called guided policy search, which incorporates the advantages of model-based methods at training time, while still training the policy to use only the onboard sensors of the robot, without explicit state estimation and using only real-world data. Guided policy search has been applied to locomotion [21], robotic manipulation [22], and vision-based robotic control [23], but all prior applications rely on an offline trajectory optimization phase to generate the controller that is then executed on the real system. While this offline optimization might use a learned model of the system dynamics [19], the resulting trajectory-centric controller is only adapted between episodes. This makes these methods liable to fail catastrophically when the model is inaccurate. We replace the offline trajectory optimization in guided policy search with MPC, which prevents catastrophic failures even during training, making the method suitable for learning policies for autonomous aerial vehicles.

One of the key advantages of guided policy search is its ability to train complex, high-dimensional, and highly nonlinear policies [20]. This allows us to use deep neural network representations for our policies, thus allowing them to handle complex, raw input from onboard sensors, without extensive engineering of the policy parameterization. While neural networks have been used for control for decades [12], [33], limitations on computation and algorithms have made large neural network policies very difficult to learn. More recently, deep neural network policies have been used for tasks ranging from robotic control [22], [23] to video game playing [27]. Deep neural networks have also been used to

learn models for MPC [18]. In this paper, we use neural networks to represent the policy, rather than the model, while MPC is used to help train this policy.

III. PRELIMINARIES

The goal of policy search is to minimize the expected cost $E_{\pi_\theta}[\sum_{t=1}^T \ell(\mathbf{x}_t, \mathbf{u}_t)]$ with respect to the parameters θ of a policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$. Here, \mathbf{x}_t denotes the full state at time t , \mathbf{o}_t the observation, \mathbf{u}_t the action, and $\ell(\mathbf{x}_t, \mathbf{u}_t)$ the cost function that defines the task. For example, a task that requires a quadrotor to fly to a position might have the cost be the distance between the vehicle and the goal. The expectation is taken with respect to $p(\tau) = p(\mathbf{x}_1) \prod_{t=1}^T p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t) \pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$, where $\tau = \{\mathbf{x}_1, \mathbf{u}_1, \dots, \mathbf{x}_T, \mathbf{u}_T\}$ denotes a trajectory. The observation \mathbf{o}_t is distributed according to some unknown observation distribution $p(\mathbf{o}_t|\mathbf{x}_t)$, which describes how the readings from the robot's sensors depend on the state.

Optimal control also seeks to solve problems of this type, under various assumptions about the dynamics $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$, observation function $p(\mathbf{o}_t|\mathbf{x}_t)$, and policy. For example, the differential dynamic programming (DDP) algorithm can be viewed as approximately optimizing the expected cost under a locally linear-Gaussian dynamics model and with time-varying linear [13], [24] or linear-Gaussian [20] policies that act directly on the state (i.e. $\mathbf{o}_t = \mathbf{x}_t$). While this method is not as general as policy gradient RL algorithms, which can optimize arbitrary parameterized policies with arbitrary observations under unknown dynamics [32], it is fast, simple, and often effective [39]. In order to combine the efficiency of DDP with the flexibility of general policy search methods, guided policy search uses DDP-like algorithms to solve the control problem from a variety of initial states and generate training data for arbitrary parameterized policies, which are then trained with supervised learning to mimic the behavior of the DDP solutions [20], [19].

Algorithm 1 presents a generic guided policy search method, where trajectory optimization is used to optimize a set of guiding trajectory distributions $p_i(\tau)$, defined by the corresponding linear-Gaussian controllers $p_i(\mathbf{u}_t|\mathbf{x}_t)$, and an arbitrary nonlinear policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ is trained using samples from all of these controllers. Since supervised learning does

not in general guarantee that the policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ will achieve good long-horizon performance [36], guided policy search alternates between optimizing the policy and optimizing each of the trajectory distributions, each time adjusting the trajectory cost and the policy optimization objective to ensure that the linear-Gaussian controllers $p_i(\mathbf{u}_t|\mathbf{x}_t)$ and policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ converge to the same behavior. The objective for trajectory optimization is modified by adding a penalty for the deviation from the policy, and the policy objective is modified by applying different weights to different samples [19] or using dual variables [23]. Convergence to a policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ that minimizes expected cost can be shown by casting this alternating optimization as a relaxation of a constrained optimization problem of the form

$$\begin{aligned} \min_{\theta, p} E_{\tau \sim p}[\ell(\tau)] \\ \text{s.t. } p(\mathbf{u}_t|\mathbf{x}_t) = \pi_\theta(\mathbf{u}_t|\mathbf{x}_t) \quad \forall t \end{aligned}$$

where $\ell(\tau)$ is shorthand for $\sum_{t=1}^T \ell(\mathbf{x}_t, \mathbf{u}_t)$, $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$ is shorthand for $\int \pi_\theta(\mathbf{u}_t|\mathbf{o}_t) p(\mathbf{o}_t|\mathbf{x}_t) d\mathbf{o}_t$, and $p(\tau)$ is a mixture of the guiding distributions $p_i(\tau)$. This constrained problem can be approximately transformed into the alternating optimization in Algorithm 1 by using sampling over the observations \mathbf{o}_t together with the framework of dual gradient descent [19] or BADMM [23]. In this paper, we use the BADMM version, which specifies the following objective for trajectory optimization:

$$\begin{aligned} \min_{p_i(\tau)} \sum_{t=1}^T E_{p_i(\mathbf{x}_t, \mathbf{u}_t)} \left[\ell(\mathbf{x}_t, \mathbf{u}_t) - \mathbf{u}_t^\top \lambda_{\mu t}^i \right. \\ \left. + \nu_t^i D_{\text{KL}}(p_i(\mathbf{u}_t|\mathbf{x}_t) \| \pi_\theta(\mathbf{u}_t|\mathbf{x}_t)) \right] \quad (1) \end{aligned}$$

where $\lambda_{\mu t}$ is a Lagrange multiplier on the mean action, and the third term is a KL-divergence penalty. Together, these terms serve to keep $p_i(\mathbf{u}_t|\mathbf{x}_t)$ close to $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$. The supervised objective for the policy is similarly given by

$$\begin{aligned} \min_{\theta} \sum_{i=1}^N \sum_{j=1}^M \sum_{t=1}^T \left[\nu_t^i D_{\text{KL}}(\pi_\theta(\mathbf{u}_t|\mathbf{o}_t^{i,j}) \| p_i(\mathbf{u}_t|\mathbf{x}_t^{i,j})) \right. \\ \left. + E_{\pi_\theta(\mathbf{u}_t|\mathbf{o}_t^{i,j})}[\mathbf{u}_t]^\top \lambda_{\mu t}^i \right] \quad (2) \end{aligned}$$

where N is the number of trajectory distributions p_i , M is the number of samples collected from each $p_i(\tau)$ and T is the length of each trajectory τ . This objective uses samples to estimate the integral $\int \pi_\theta(\mathbf{u}_t|\mathbf{o}_t) p(\mathbf{o}_t|\mathbf{x}_t) d\mathbf{o}_t$ and, in the case where the policy is given by the Gaussian $\mathcal{N}(\mu^\pi(\mathbf{o}_t), \Sigma^\pi)$, it corresponds to a weighted least squares objective on the mean $\mu^\pi(\mathbf{o}_t)$, while Σ^π can be solved for in closed form. For a detailed derivation of this method, as well as the update equations for ν_t^i and $\lambda_{\mu t}^i$, we refer the reader to previous work [23]. Note that the policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ only uses the observations \mathbf{o}_t as input, which means that, once it has been trained, it can be used in situations where the true state \mathbf{x}_t is unknown.

Prior guided policy search methods optimized the guiding trajectory distributions $p_i(\tau)$ using either offline trajectory optimization with known system dynamics [21], or trajectory-centric reinforcement learning [19]. The former

Algorithm 1 Generic guided policy search summary

- 1: **for** iteration $k = 1$ to K **do**
 - 2: Optimize trajectory distributions $p_i(\tau)$ to minimize $E_{p_i}[\ell(\tau)]$ and deviation from the policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$
 - 3: Generate samples $\{\tau_i^j\}$ from each $p_i(\tau)$
 - 4: Train nonlinear policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ to match the sampled trajectories $\{\tau_i^j\}$
 - 5: Update Lagrange multipliers to encourage agreement between $p_i(\mathbf{u}_t|\mathbf{x}_t)$ and $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$
 - 6: **end for**
 - 7: **return** optimized policy parameters θ
-

class of methods assumes that the true dynamics are known in advance, while the latter requires iteratively learning the dynamics by attempting to run potentially suboptimal controllers on the real physical system. In the case of unstable systems, such as autonomous aerial vehicles, neither approach is ideal, since the true dynamics are not known perfectly, and the suboptimal controller rollouts required for reinforcement learning might cause catastrophic failure, such as a crash. On the other hand, MPC methods that continuously recompute the vehicle's trajectory under an approximate model of the dynamics have been shown to exhibit good robustness to model errors [39]. In the next section, we discuss how MPC can be combined with guided policy search to learn effective control policies.

IV. MPC-GUIDED POLICY SEARCH

In this paper, we use MPC together with offline trajectory optimization to generate guiding samples for guided policy search. We assume that we have access to an approximate model of the system dynamics, which we use during training to choose actions that will accomplish the desired task, starting from a variety of initial states. These samples are then used as training data to train a nonlinear policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$, and this policy is included in the cost function for the next batch of samples. By repeatedly collecting new samples and training the policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ in this way, the method can acquire an effective nonlinear policy that generalizes to new states.

This method is a special case of the generic guided policy search framework presented in Algorithm 1, but a number of modifications are necessary to adapt the approach to use MPC to generate the guiding trajectory distributions. First, the MPC procedure must minimize the objective in Equation (1), which means that it must also minimize deviation from the neural network policy. Second, each MPC rollout produces a different locally linear controller, which necessitates a modification to the supervised policy learning phase. Lastly, since MPC uses a relatively short horizon, we generate target trajectories using an offline trajectory optimization phase, and then track these trajectories. We develop a formulation for this tracking objective that is compatible with guided policy search.

A. Model Predictive Control with DDP

The MPC method we use is based on differential dynamic programming (DDP) [13]. In particular, we use an efficient variant of this method called iterative LQG, which assumes access to an approximate model of the system dynamics and uses a local linear-quadratic expansion to solve the optimal control problem. We summarize the method in this section. However, our derivation largely follows prior work [41].

Iterative LQG assumes that the dynamics are given by a deterministic mean function $f(\mathbf{x}_t, \mathbf{u}_t) = E[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]$, with additive Gaussian noise. The algorithm computes a linear expansion of the dynamics around a nominal trajectory $\hat{\tau} = \{\hat{\mathbf{x}}_1, \hat{\mathbf{u}}_1, \dots, \hat{\mathbf{x}}_T, \hat{\mathbf{u}}_T\}$, as well as a quadratic

expansion of the cost. Without loss of generality, we assume that the nominal states and actions are zero for notational convenience. The linearized dynamics have the form $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t) = \mathcal{N}(f_{\mathbf{x}t}\mathbf{x}_t + f_{\mathbf{u}t}\mathbf{u}_t + f_{ct}, \mathbf{F}_t)$, and the quadratic cost approximation has the form

$$\ell(\mathbf{x}_t, \mathbf{u}_t) \approx \frac{1}{2}[\mathbf{x}_t; \mathbf{u}_t]^T \ell_{\mathbf{xu}, \mathbf{xu}t} [\mathbf{x}_t; \mathbf{u}_t] + [\mathbf{x}_t; \mathbf{u}_t]^T \ell_{\mathbf{xu}t} + \text{const},$$

where subscripts denote derivatives, e.g. $\tilde{\ell}_{\mathbf{xu}t}$ is the gradient of $\tilde{\ell}$ with respect to $[\mathbf{x}_t; \mathbf{u}_t]$, while $\tilde{\ell}_{\mathbf{xu}, \mathbf{xu}t}$ is the Hessian. Under this model of the dynamics and cost function, the optimal policy can be computed by recursively computing the quadratic Q -function and value function, starting with the last time step. These functions are given by

$$V(\mathbf{x}_t) = \frac{1}{2}\mathbf{x}_t^T V_{\mathbf{x}, \mathbf{x}t} \mathbf{x}_t + \mathbf{x}_t^T V_{\mathbf{x}t} + \text{const}$$

$$Q(\mathbf{x}_t, \mathbf{u}_t) = \frac{1}{2}[\mathbf{x}_t; \mathbf{u}_t]^T Q_{\mathbf{xu}, \mathbf{xu}t} [\mathbf{x}_t; \mathbf{u}_t] + [\mathbf{x}_t; \mathbf{u}_t]^T Q_{\mathbf{xu}t} + \text{const}$$

We can express them with the following recurrence:

$$Q_{\mathbf{xu}, \mathbf{xu}t} = \ell_{\mathbf{xu}, \mathbf{xu}t} + f_{\mathbf{xu}t}^T V_{\mathbf{x}, \mathbf{x}t+1} f_{\mathbf{xu}t}$$

$$Q_{\mathbf{xu}t} = \ell_{\mathbf{xu}t} + f_{\mathbf{xu}t}^T V_{\mathbf{x}t+1}$$

$$V_{\mathbf{x}, \mathbf{x}t} = Q_{\mathbf{x}, \mathbf{x}t} - Q_{\mathbf{u}, \mathbf{u}t}^T Q_{\mathbf{u}, \mathbf{u}t}^{-1} Q_{\mathbf{u}, \mathbf{x}t}$$

$$V_{\mathbf{x}t} = Q_{\mathbf{x}t} - Q_{\mathbf{u}, \mathbf{x}t}^T Q_{\mathbf{u}, \mathbf{u}t}^{-1} Q_{\mathbf{u}t},$$

which allows us to compute the optimal control law as $g(\mathbf{x}_t) = \hat{\mathbf{u}}_t + \mathbf{k}_t + \mathbf{K}_t(\mathbf{x}_t - \hat{\mathbf{x}}_t)$, where $\mathbf{K}_t = -Q_{\mathbf{u}, \mathbf{u}t}^{-1} Q_{\mathbf{u}, \mathbf{x}t}$ and $\mathbf{k}_t = -Q_{\mathbf{u}, \mathbf{u}t}^{-1} Q_{\mathbf{u}t}$. Performing a forward rollout using this control law allows us to find a new nominal trajectory, and the backward dynamic programming pass is repeated around this trajectory to take the next Gauss-Newton step.

To adapt DDP for performing MPC, we simply run the algorithm for a shorter horizon H at each time step t , so that the backward pass is performed from time step t to $t+H$.

B. Adapting MPC for Guided Policy Search

While we could simply adapt the DDP-based MPC algorithm in the previous section to optimize Equation 1, the short horizon typically used in MPC makes it difficult to accomplish complex tasks like obstacle avoidance, which require long-horizon lookahead, with only a high-level specification of the task, such as a desired flight direction and an obstacle collision penalty. Instead, we use an offline optimization based on iterative LQG [24] to first generate a reference trajectory, and then track this trajectory using MPC, with an additional term to account for differences from the neural network policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$.

Offline, we run iterative LQG with our known approximate model to optimize Equation 1. Since Equation 1 contains a KL-divergence term, the objective can be rewritten as

$$\min_{p_i(\tau)} \sum_{t=1}^T E_{p_i(\mathbf{x}_t, \mathbf{u}_t)} \left[\frac{1}{\nu_t^i} \ell(\mathbf{x}_t, \mathbf{u}_t) - \frac{1}{\nu_t^i} \mathbf{u}_t^T \lambda_{it}^i - \log \pi_\theta(\mathbf{u}_t|\mathbf{x}_t) - \mathcal{H}(p_i(\mathbf{u}_t|\mathbf{x}_t)) \right], \quad (3)$$

where \mathcal{H} represents the maximum entropy, and this maximum entropy objective can be optimized with iterative LQG.

The solution is a linear-Gaussian controller of the form $p_i(\mathbf{u}_t|\mathbf{x}_t) = \mathcal{N}(\mathbf{K}_t\mathbf{x}_t + \mathbf{k}_t, Q_{\mathbf{u},\mathbf{u}_t}^{-1})$, as shown in prior work [20]. Prior methods sample directly from the linear-Gaussian controller $p_i(\mathbf{u}_t|\mathbf{x}_t)$ [21], but since we would like to use MPC to robustly control the robot during the rollout, we instead construct a surrogate cost function $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$ for MPC that will allow us to robustly generate trajectories that have high probability under $p_i(\tau)$.

This surrogate cost should fulfill a number of criteria in order to be effective: first, it must encourage MPC to visit states that have high probability under $p_i(\tau)$; second, it must produce good long-horizon behavior even when optimized under a short horizon; and third, it must keep the generated behavior close to the neural network policy $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$. When we run MPC at time step t from the current state \mathbf{x}_t , we first compute $p_i(\mathbf{x}_{t'}|\mathbf{x}_t)$ for each $t' \in [t+1, t+H]$ under the known approximate dynamics and the time-varying linear-Gaussian controller obtained from the offline LQG optimization. Using $\mu_{t'}$ and $\Sigma_{t'}$ to denote the mean and covariance of $p_i(\mathbf{x}_{t'}|\mathbf{x}_t)$, we can compute these distributions according to the following recurrence:

$$\begin{aligned} \mu_{t'+1} &= \begin{bmatrix} f_{\mathbf{x}t'} & f_{\mathbf{u}t'} \end{bmatrix} \begin{bmatrix} \mu_{t'} \\ \hat{\mathbf{u}}_{t'} + \mathbf{k}_{t'} + \mathbf{K}_{t'}(\mu_{t'} - \hat{\mathbf{x}}_{t'}) \end{bmatrix} \\ \Sigma_{t'+1} &= \begin{bmatrix} f_{\mathbf{x}t'} & f_{\mathbf{u}t'} \end{bmatrix} \begin{bmatrix} \Sigma_{t'} & \Sigma_{t'}\mathbf{K}_{t'}^T \\ \mathbf{K}_{t'}^T\Sigma_{t'} & Q_{\mathbf{u},\mathbf{u}t'}^{-1} + \mathbf{K}_{t'}^T\Sigma_{t'}\mathbf{K}_{t'} \end{bmatrix} \begin{bmatrix} f_{\mathbf{x}t'}^T \\ f_{\mathbf{u}t'}^T \end{bmatrix} + \mathbf{F}_{t'} \end{aligned}$$

The intuition here is that we would like to figure out which states the offline LQG solution would prefer to visit, independently of the actions required to get to these states. This is important since, in the presence of model errors and perturbations, the nonlinear approximate model might indicate different actions when combined with MPC, but the overall distribution over states should remain similar. Once we obtain $p_i(\mathbf{x}_{t+1}, \dots, \mathbf{x}_{t+H}|\mathbf{x}_t)$, we can marginalize to obtain $p_i(\mathbf{x}_{t'})$ for each time step $t' \in [t+1, t+H]$, and we then construct the surrogate cost as

$$\begin{aligned} \tilde{\ell}(\mathbf{x}_{t'}, \mathbf{u}_{t'}) &= -\log p_i(\mathbf{x}_{t'}|\mathbf{x}_t) \\ &\quad - \nu_{t'}^i \log \pi_\theta(\mathbf{u}_{t'}|\mathbf{x}_{t'}) - \mathbf{u}_{t'}^T \lambda_{\mu_{t'}}^i. \end{aligned} \quad (4)$$

We then run MPC on this cost as described in the previous section to obtain a new linear-Gaussian controller for time step t of the form $\tilde{p}_{ij}(\mathbf{u}_t|\mathbf{x}_t) = \mathcal{N}(\tilde{\mathbf{K}}_{tij}\mathbf{x}_t + \tilde{\mathbf{k}}_{tij}, \tilde{Q}_{\mathbf{u},\mathbf{u}tij}^{-1})$, and choose the action by sampling from this linear-Gaussian. The subscript ij here denotes the j th sample (generated via MPC) from the i th trajectory distribution.

While samples generated in this way are not exactly samples from $p_i(\tau)$, but rather samples from a distribution formed by the product of independent marginals at each time step, we found that the resulting algorithm was still able to produce good training data for the neural network in guided policy search. Furthermore, the additional information provided by $p_i(\mathbf{x}_{t'}|\mathbf{x}_t)$ allowed MPC to succeed in the presence of model errors and disturbances, even on tasks such as obstacle avoidance that require long-horizon lookahead.

One final detail is that both the cost in Equation (4) and the offline optimization objective in Equation (3) require access to $\log \pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$, while we only have access to $\log \pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$, and \mathbf{o}_t is in general a complex and unknown function of \mathbf{x}_t , since the observation might include, e.g., laser

rangefinder readings, while the state might consist of the vehicle's position and orientation. To obtain $\log \pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$, we follow prior work and approximately linearize the policy by using the previous set of rollouts from the physical system. This can be done by fitting a time-varying linear-Gaussian model of $\log \pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$ to the samples, since each sample includes both \mathbf{x}_t and \mathbf{o}_t , allowing us to evaluate the policy at each sampled state \mathbf{x}_t . The fitting is done by using linear regression with a Gaussian mixture model prior, as in previous work [19].

C. Training the Nonlinear Policy

The final nonlinear policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ is trained using standard supervised learning, from samples collected via MPC. The objective for this supervised learning is given in Equation (2), though in the case of MPC-based samples, we substitute $\tilde{p}_{ij}(\mathbf{u}_t|\mathbf{x}_t)$ for $p_i(\mathbf{u}_t|\mathbf{x}_t)$. In the case of a conditionally Gaussian policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t) = \mathcal{N}(\mu^\pi(\mathbf{o}_t), \Sigma^\pi(\mathbf{o}_t))$, the KL-divergence $D_{\text{KL}}(\pi_\theta(\mathbf{u}_t|\mathbf{o}_t^{i,j})\|\tilde{p}_{ij}(\mathbf{u}_t|\mathbf{x}_t^{i,j}))$ in this objective can be written out as

$$\begin{aligned} D_{\text{KL}}(\pi_\theta(\mathbf{u}_t|\mathbf{o}_t^{i,j})\|\tilde{p}_{ij}(\mathbf{u}_t|\mathbf{x}_t^{i,j})) &= \\ &= \frac{1}{2}(\mu^\pi(\mathbf{o}_t) - \tilde{g}_{tij}(\mathbf{x}_t))\tilde{Q}_{\mathbf{u},\mathbf{u}tij}(\mu^\pi(\mathbf{o}_t) - \tilde{g}_{tij}(\mathbf{x}_t)) \\ &\quad - \frac{1}{2}\text{tr}[\tilde{Q}_{\mathbf{u},\mathbf{u}tij}\Sigma^\pi(\mathbf{o}_t)] + \frac{1}{2}\log|\Sigma^\pi(\mathbf{o}_t)| + \lambda_{\mu_t}^T\mu^\pi(\mathbf{o}_t), \end{aligned}$$

where $\tilde{g}_{tij}(\mathbf{x}_t) = \tilde{\mathbf{K}}_{tij}\mathbf{x}_t + \tilde{\mathbf{k}}_{tij}$. Note that this is simply a weighted least squares objective on the mean function $\mu^\pi(\mathbf{o}_t)$. In this work, we represent $\mu^\pi(\mathbf{o}_t)$ with a multilayer neural network, which allows us to train flexible and expressive policies. Since we prefer deterministic or nearly-deterministic policies, we choose $\Sigma^\pi(\mathbf{o}_t)$ to be constant, which means that we can solve for it in closed form according to

$$\Sigma^\pi(\mathbf{o}_t) = \left(\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M \tilde{Q}_{\mathbf{u},\mathbf{u}tij} \right)^{-1}.$$

The neural network mean function $\mu^\pi(\mathbf{o}_t)$ is optimized using stochastic gradient descent (SGD). As noted earlier, one of the key advantages of this type of training approach is that the input \mathbf{o}_t to the neural network policy need not match the state \mathbf{x}_t used during trajectory optimization and MPC, which allows us to train policies that operate directly on raw inputs from the onboard sensors.

Algorithm 2 MPC-guided policy search

- 1: **for** iteration $k = 1$ to K **do**
 - 2: Optimize $p_i(\tau)$ offline according to Equation (3)
 - 3: Run MPC M times from initial states $\mathbf{x}_1 \sim p_i(\mathbf{x}_1)$ to create $\{\tilde{p}_{ij}(\tau)\}$ and $\{\tau_{ij}\}$ using $\tilde{\ell}(\tau)$ in Equation (4)
 - 4: Train nonlinear policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ to match each $\tilde{p}_{ij}(\mathbf{u}_t|\mathbf{x}_t)$ along each τ_{ij} , using Equation (2)
 - 5: Fit time-varying linear-Gaussian model to estimate $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$ around each $p_i(\tau)$ using samples $\{\tau_{ij}\}$
 - 6: Update ν_t^i and $\lambda_{\mu_t}^i$ as in [23]
 - 7: **end for**
 - 8: **return** optimized policy parameters θ
-

D. Algorithm Summary

A summary of our method is presented in Algorithm 2. At each iteration, we first generate an offline solution by using iterative LQG to optimize the augmented objective in Equation (3). This offline solution allows us to initialize and construct the cost for MPC rollouts. We conduct M MPC rollouts for each trajectory distribution $p_i(\tau)$, constructing a new surrogate cost $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$ at each time step. These MPC rollouts provide us with sample trajectories $\{\tau_{ij}\}$ and MPC controllers $\{\tilde{p}_{ij}(\tau)\}$, which we can use to train the nonlinear neural network policy $\pi_\theta(\mathbf{u}_t|\mathbf{o}_t)$ as described in the previous section. After the policy is trained, we update our time-varying linear-Gaussian fit for $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$ by using the latest samples. Note that a separate linear-Gaussian estimate of $\pi_\theta(\mathbf{u}_t|\mathbf{x}_t)$ is constructed around each trajectory distribution $p_i(\tau)$. Finally, we adjust the dual variables as described in previous work [23].

V. EXPERIMENTAL EVALUATION

We evaluated our method on simulated quadrotor obstacle avoidance tasks.

A. Quadrotor System

The simulated quadrotor was modeled after 3DR’s IRIS+¹, which has width 0.47m, height 0.11m, and weight 1.5kg. The dynamics followed the formulation described by Martin and Salaun [25]. The true state of the vehicle $\mathbf{x}_t = (\mathbf{p}_t, \mathbf{v}_t, \mathbf{q}_t, \boldsymbol{\omega}_t) \in \mathbb{R}^{13}$ consisted of the position $\mathbf{p}_t = (x_t, y_t, z_t)$ and orientation \mathbf{q}_t , expressed as a quaternion, as well as their time derivatives, i.e. linear velocity \mathbf{v}_t and angular velocity $\boldsymbol{\omega}_t$. The controls $\mathbf{u}_t \in \mathbb{R}^4$ consisted of rotor velocities. The observation model $\mathbf{o}_t = (\mathbf{r}_t, \mathbf{v}_t, \mathbf{q}_t, \boldsymbol{\omega}_t) \in \mathbb{R}^{40}$ lacked the position \mathbf{p}_t and instead included readings \mathbf{r}_t from a set of 30 equally spaced laser rangefinders with max range 5m arranged in 180 degree fan in front of the vehicle. This type of observation model is quite challenging to integrate into simple control methods, such as time-varying linear controllers, but can be easily processed by a multilayer neural network policy.

B. Cost Function

The cost function for the iterative offline LQG optimization was

$$\begin{aligned} l(\mathbf{x}_t, \mathbf{u}_t) = & 10^3 \|\mathbf{v}_t - \mathbf{v}^*\|_2^2 + 500 \|z_t - z^*\|_2^2 + \\ & 10^4 \|\mathbf{q}_t - \mathbf{q}^*\|_2^2 + 250 \|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2 + \\ & 0.5 \|\mathbf{u}_t - \mathbf{u}_{\text{HOVER}}^*\|_2^2 + \\ & 10^4 \max(d_{\text{SAFE}} - \text{signed_distance}(\mathbf{x}_t), 0) \end{aligned}$$

where $\mathbf{x}_t = ((x_t, y_t, z_t), \mathbf{v}_t, \mathbf{q}_t, \boldsymbol{\omega}_t)$ is the full state as defined previously; $z^*, \mathbf{v}^*, \mathbf{q}^*, \boldsymbol{\omega}^*$ are the target height, linear velocity, orientation, and angular velocity, respectively; and $\mathbf{u}_{\text{HOVER}}^*$ is the desired rotor velocity for hovering. The final term is a hinge loss on the distance of the quadrotor to the nearest obstacle; there is no penalty if the nearest obstacle is further than d_{SAFE} .




Fig. 2: The quadrotor must learn to fly around a cylindrical obstacle and down a hallway using only onboard sensing. The blue semicircle is the range of the onboard laser range finders.

C. Neural Network Policy

The neural network policy consisted of two fully connected hidden layers, each with 40 rectified linear (ReLU) units [30]. For training the neural network in each iteration of guided policy search, we used the ADAM [14] algorithm, optimizing using 20,000 minibatches with a minibatch size of 50. We used default values for all learning parameters as presented in [14].

D. Experimental Domains

Figure 2 depicts the two simulated environments in which we trained our neural network policies: a single cylindrical obstacle of radius 0.5m and height 4m, and a straight hallway of width 5m and height 4m. For the cylinder avoidance task, we used $N = 18$ initial states with varying initial (x, y) positions in front of the cylinder; for the hallway task, we used $N = 6$ initial states with varying y -coordinate values. For both tasks, each initial state corresponds to a different trajectory distribution $p_i(\tau)$, with $M = 4$ samples from each distribution. Each trajectory τ has length $T = 150$, which equates to 7.5 seconds. Collision with an obstacle or the ground, or flying above an obstacle were considered as a training crash.

E. Baseline Methods and Model Errors

To evaluate the importance of using MPC, we trained neural network policies on each of the tasks in the presence of model errors, using three variants: the full MPC-guided policy search algorithm with the surrogate cost $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$ described in Section IV-B, a variant of MPC-guided policy search that uses the true cost $\ell(\mathbf{x}_t, \mathbf{u}_t)$ with the policy KL-divergence term and dual variables for MPC, and a variant that does not use MPC at all, and instead performs the rollouts by using the time-varying linear-Gaussian policy generated by the offline iterative LQG algorithm. All methods were trained for $K = 5$ guided policy search iterations.

We evaluated each of the above variants in the presence of four different types of model errors: no model error, the actual weight of the quadrotor was 0.05kg (3.3%) greater than the expected weight, the two rotors on one side of the quadrotor had an 8% multiplicative rotor velocity bias, and all model parameters (e.g. moments of inertia, drag coefficients) were perturbed by retaining only one significant digit.

¹<https://3dr.com/kb/iris/>

train: single cylinder; test: infinite forest												
	no model error			0.05kg mass error			8% rotor bias			perturbed model parameters		
method	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$
number of training crashes	1	0	0	1	0	0	46	0	0	N/A	4	0
average test flight duration (s)	56.9 ± 27.3	35.8 ± 22.5	53.4 ± 22.5	11.3 ± 5.2	6.9 ± 2.6	34.4 ± 8.4	60.0 ± 15.8	56.1 ± 30.8	95.7 ± 9.8	N/A	20.8 ± 9.8	60.5 ± 27.0

train: straight hallway; test: winding hallway												
	no model error			0.05kg mass error			8% rotor bias			perturbed model parameters		
method	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) offline $\ell(\mathbf{x}_t, \mathbf{u}_t)$	(baseline) MPC $\ell(\mathbf{x}_t, \mathbf{u}_t)$	full MPC $\tilde{\ell}(\mathbf{x}_t, \mathbf{u}_t)$
number of training crashes	0	0	0	0	0	0	76	0	0	N/A	0	0
average test flight duration (s)	46.2 ± 28.4	35.2 ± 13.3	35.2 ± 20.0	21.7 ± 5.8	15.7 ± 1.3	31.8 ± 15.7	26.0 ± 21.1	51.1 ± 28.6	28.5 ± 16.2	N/A	9.9 ± 4.2	55.2 ± 17.5

TABLE I: Training and test results comparing our full MPC-guided policy search method with two baselines. The test flight duration was averaged over 20 runs of the learned policies. In one scenario (top table), the policy search method variants were trained on a single cylinder and tested in an infinite forest. In the other scenario (bottom table), the policy search variants were trained in a straight hallway and tested in a winding hallway. Each policy search variant was tested with four different model errors. Experiments labelled N/A were unable to complete due to excessive crashing. In the majority of experiments, the full MPC-guided policy search method outperformed the two baselines, crashing less during training and with the final learned neural network policy flying for the longest duration in the test scenarios.

F. Results

Table I shows the number of crashes experienced by the quadrotor during training. These results indicate that MPC using the surrogate cost is able to train a successful neural network policy without experiencing catastrophic failure.

To evaluate the generalization of the learned policies, we ran the trained neural networks in two test scenarios: an infinite forest of cylinders of the same shape as in training, but at random positions an average distance of 5m apart, and a winding hallway with randomized turns of at most 30° every 5m. We ran the policies trained with a single cylinder in the infinite forest, and we ran the trained policies for the straight hallway in the winding hallway. The average flight duration of the final trained neural network policies are shown in Table I. With no model errors, our MPC-guided policy search algorithm was comparable to the other methods. When model errors were introduced, our method outperformed the two baselines in the majority of scenarios. Videos of the resulting flights are included as supplementary material, and may also be viewed on the project webpage².

Our evaluation shows that MPC-guided policy search is an effective algorithm for off-policy training of complex neural network policies for autonomous aerial vehicles. Our full method was able to learn each of the two behaviors without experiencing any catastrophic failures during training, and the trained policy was able to generalize effectively.

VI. DISCUSSION AND FUTURE WORK

We presented an algorithm for training deep neural network control policies for autonomous aerial vehicles, by using model predictive control to generate guiding samples for guided policy search. Our MPC-guided policy search uses

a modified MPC algorithm that trades off minimizing the cost against matching the current neural network policy, so as to generate good training data that can be used to train a better policy with standard supervised learning. Since the partially trained neural network policy is never used to choose actions at training time, the more robust and reliable MPC method provides a substantial improvement in safety over traditional reinforcement learning methods. Our results show that this algorithm is able to learn complex policies, such as high speed obstacle avoidance, using raw sensor inputs and low-level rotor command outputs.

One of the key ideas behind our method is the notion of an instrumented training setup, which allows MPC to be performed at training time with full state observations, which could be provided, for example, by using motion capture. At the same time, the vehicle gathers observations from its own onboard sensors, and trains the policy to mimic the action chosen by MPC using only the raw sensor readings, without relying on the full state. Acquiring the sensor readings is important, because accurately modeling complex sensors, such as laser range finders and cameras, is very difficult, while obtaining a model of the vehicle that is accurate enough to perform MPC is comparatively easier.

While our approach can train very complex, high-dimensional policies, it shares many of the limitations of prior guided policy search methods [23]. In particular, full state observations are required at training time, in order to perform MPC, even though the final neural network policy can perform the task using only onboard sensors. In the real world, this kind of state information could be obtained using an instrumented training environment (with, for example, motion capture). Since the instrumentation is only required during training, the final neural network is still able to act in the real world, so this approach is practical

²<http://rll.berkeley.edu/icra2016mpcpgs>

for a wide range of robotic tasks. However, not all aerial maneuvers can be learned in such an instrumented training setup, and combining explicit state estimation with guided policy search in future work could lead to a much more broadly applicable algorithm. Another direction that can be explored in future work is to combine guided policy search with more sophisticated MPC and planning algorithms. In principle, a wide variety of methods can be used to generate guiding samples, and more sophisticated methods might afford superior robustness and obstacle avoidance [7].

ACKNOWLEDGEMENTS

This research was funded in part by the Army Research Office through the MAST program, by DARPA under Award #N66001-15-2-4047, and by the Berkeley Vision and Learning Center (BVLC).

REFERENCES

- [1] P. Abbeel, A. Coates, M. Quigley, and A. Ng, "An application of reinforcement learning to aerobatic helicopter flight," in *Advances in Neural Information Processing Systems (NIPS)*, 2006.
- [2] K. Alexis, G. Nikolakopoulos, and A. Tzes, "Model predictive quadrotor control: attitude, altitude and position experimental studies," *Control Theory Applications, IET*, vol. 6, no. 12, pp. 1812–1827, Aug 2012.
- [3] F. Augugliaro, A. P. Schoellig, and R. D'Andrea, "Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach," in *International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [4] P. Bouffard, A. Aswani, and C. Tomlin, "Learning-based model predictive control on a quadrotor: Onboard implementation and experimental results," in *International Conference on Robotics and Automation (ICRA)*, 2012.
- [5] M. Deisenroth, G. Neumann, and J. Peters, "A survey on policy search for robotics," *Foundations and Trends in Robotics*, vol. 2, no. 1-2, pp. 1–142, 2013.
- [6] M. Deisenroth, C. Rasmussen, and D. Fox, "Learning to control a low-cost manipulator using data-efficient reinforcement learning," in *Robotics: Science and Systems (RSS)*, 2011.
- [7] R. Deits and R. Tedrake, "Efficient mixed-integer planning for UAVs in cluttered environments," in *International Conference on Robotics and Automation (ICRA)*, 2015.
- [8] G. Endo, J. Morimoto, T. Matsubara, J. Nakanishi, and G. Cheng, "Learning CPG-based biped locomotion with a policy gradient method: Application to a humanoid robot," *International Journal of Robotic Research*, vol. 27, no. 2, pp. 213–228, 2008.
- [9] F. Fraundorfer, L. Heng, D. Honegger, G. Lee, L. Meier, P. Tanskanen, and M. Pollefeys, "Vision-based autonomous mapping and exploration using a quadrotor mav," in *International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [10] T. Geng, B. Porr, and F. Wörgötter, "Fast biped walking with a reflexive controller and realtime policy searching," in *Advances in Neural Information Processing Systems (NIPS)*, 2006.
- [11] L. Heng, L. Meier, P. Tanskanen, F. Fraundorfer, and M. Pollefeys, "Autonomous obstacle avoidance and maneuvering on a vision-guided mav using on-board processing," in *International Conference on Robotics and Automation (ICRA)*, 2011.
- [12] K. J. Hunt, D. Sbarbaro, R. Zbikowski, and P. J. Gawthrop, "Neural networks for control systems: A survey," *Automatica*, vol. 28, no. 6, pp. 1083–1112, Nov. 1992.
- [13] D. Jacobson and D. Mayne, *Differential Dynamic Programming*. Elsevier, 1970.
- [14] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [15] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *International Journal of Robotic Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [16] J. Kober, E. Oztop, and J. Peters, "Reinforcement learning to adjust robot movements to new situations," in *Robotics: Science and Systems (RSS)*, 2010.
- [17] N. Kohl and P. Stone, "Policy gradient reinforcement learning for fast quadrupedal locomotion," in *International Conference on Robotics and Automation (IROS)*, 2004.
- [18] I. Lenz, R. Knepper, and A. Saxena, "Deepmpc: Learning deep latent features for model predictive control," in *Robotics: Science and Systems (RSS)*, 2015.
- [19] S. Levine and P. Abbeel, "Learning neural network policies with guided policy search under unknown dynamics," in *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [20] S. Levine and V. Koltun, "Guided policy search," in *International Conference on Machine Learning (ICML)*, 2013.
- [21] —, "Learning complex neural network policies with trajectory optimization," in *International Conference on Machine Learning (ICML)*, 2014.
- [22] S. Levine, N. Wagener, and P. Abbeel, "Learning contact-rich manipulation skills with guided policy search," in *International Conference on Robotics and Automation (ICRA)*, 2015.
- [23] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *arXiv preprint arXiv:1504.00702*, 2015.
- [24] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *ICINCO (1)*, 2004, pp. 222–229.
- [25] P. Martin and E. Salaun, "The true role of accelerometer feedback in quadrotor control," in *International Conference on Robotics and Automation (ICRA)*, 2010.
- [26] D. Q. Mayne, M. M. Seron, and S. V. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, Feb. 2005.
- [27] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing Atari with deep reinforcement learning," *NIPS '13 Workshop on Deep Learning*, 2013.
- [28] F. L. Mueller, A. P. Schoellig, and R. D'Andrea, "Iterative learning of feed-forward corrections for high-performance tracking," in *International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [29] M. Mueller and R. D'Andrea, "A model predictive controller for quadrotor state interception," in *European Control Conference (ECC)*, 2013.
- [30] V. Nair and G. Hinton, "Rectified linear units improve restricted boltzmann machines," in *International Conference on Machine Learning (ICML)*, 2010.
- [31] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in *International Conference on Robotics and Automation (ICRA)*, 2009.
- [32] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural Networks*, vol. 21, no. 4, pp. 682–697, 2008.
- [33] D. Pomerleau, "ALVINN: an autonomous land vehicle in a neural network," in *Advances in Neural Information Processing Systems (NIPS)*, 1989.
- [34] G. V. Raffo, M. G. Ortega, and F. R. Rubio, "An integral predictive/nonlinear control structure for a quadrotor helicopter," *Automatica*, vol. 46, no. 1, pp. 29 – 39, 2010.
- [35] A. G. Richards, "Robust constrained model predictive control," Ph.D. dissertation, Massachusetts Institute of Technology, 2004.
- [36] S. Ross, G. Gordon, and A. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," *Journal of Machine Learning Research*, vol. 15, pp. 627–635, 2011.
- [37] S. Ross, N. Melik-Barkhudarov, K. S. Shankar, A. Wendel, D. Dey, J. A. Bagnell, and M. Hebert, "Learning monocular reactive UAV control in cluttered natural environments," in *International Conference on Robotics and Automation (ICRA)*, 2013.
- [38] D. H. Shim, H. J. Kim, and S. Sastry, "Nonlinear model predictive tracking control for rotorcraft-based unmanned aerial vehicles," in *American Control Conference (ACC)*, 2002.
- [39] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- [40] R. Tedrake, T. Zhang, and H. Seung, "Stochastic policy gradient reinforcement learning on a simple 3d biped," in *International Conference on Intelligent Robots and Systems (IROS)*, 2004.
- [41] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.