

Weighted Maximum Likelihood for Controller Tuning

Angel Romero*, Shreedhar Govil*, Gonca Yilmaz*, Yunlong Song, Davide Scaramuzza

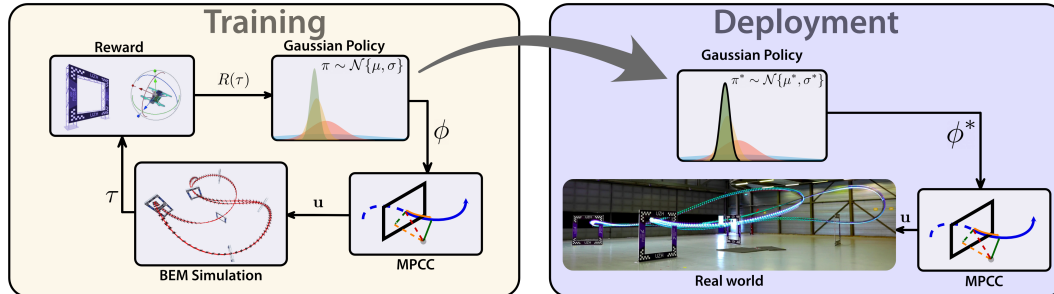


Fig. 1: During training, we sample a set of tuning parameters ϕ , run a series of simulations with a Model Predictive Contouring Controller in the loop, and obtain a set of trajectories τ . The reward signal computed from τ is then used to update the policy. The best policy is directly deployed in the real world and validated against both the hand-tuned controller and the state-of-the-art auto-tuning baseline. Our approach outperforms these methods and achieves an unprecedented lap time, reaching speeds of 75 km h^{-1} .

Abstract—Recently, Model Predictive Contouring Control (MPCC) has arisen as the state-of-the-art approach for model-based agile flight. MPCC benefits from great flexibility in trading-off between progress maximization and path following at runtime without relying on globally optimized trajectories. However, finding the optimal set of tuning parameters for MPCC is challenging because (i) the full quadrotor dynamics are non-linear, (ii) the cost function is highly non-convex, and (iii) of the high dimensionality of the hyperparameter space. This paper leverages a probabilistic Policy Search method—Weighted Maximum Likelihood (WML)—to automatically learn the optimal objective for MPCC. WML is sample-efficient due to its closed-form solution for updating the learning parameters. Additionally, the data efficiency provided by the use of a model-based approach allows us to directly train in a high-fidelity simulator, which in turn makes our approach able to transfer zero-shot to the real world. We validate our approach in the real world, where we show that our method outperforms both the previous manually tuned controller and the state-of-the-art auto-tuning baseline reaching speeds of 75 km/h.

SUPPLEMENTARY MATERIAL

Video of the experiments: <https://youtu.be/21junNZ-M38>

I. INTRODUCTION

Many fields in modern robotics, such as autonomous robot navigation or legged locomotion, become impeded by the complexity of simultaneously handling multiple task objectives while also dealing with environment changes over time. For instance, when flying a vision-based quadrotor, it is essential to include a perception-aware objective in the

cost function due to the limited field of view of the camera and the underactuated nature of the quadrotor dynamics. Another example comes from quadrupedal robots, where the robot needs to keep walking forward despite facing highly irregular profiles, deformable terrain, slippery surfaces, and overground obstructions [1].

Recent controller designs relying on optimization have shown the potential to cope with multiple task objectives and environment changes. Falanga et al. [2] proposed a perception-aware model predictive controller that can simultaneously track a trajectory and maximize the visibility of a point of interest in the camera image. Similarly, Model Predictive Contouring Control [3], [4] applied to quadrotor flight [5], [6] addresses the time-optimal flight problem by solving the time-allocation and the contouring control problem simultaneously, in real time, manifesting great adaptation against model mismatches and unknown disturbances.

However, despite the flexibility provided by numerical optimization, the design of a well-formulated task objective is challenging. To achieve real-time control performance, several approximations are introduced in the cost formulation, resulting in time-consuming hand tuning which leads to sub-optimal performance. Machine learning can be leveraged to select optimal parameters for optimization-based controllers automatically.

There are different existing approaches for automatic controller tuning. Specifically, one approach that has shown promising results in the domain of car racing is [7], which uses Bayesian Optimization to jointly optimize the parameters of the dynamic model and tuning parameters of an MPCC architecture. Similarly, AutoTune [8] tackles the receding horizon controller tuning for agile flight by leveraging statistical learning. It has the benefit that it can directly cope with sparse metrics such as lap times or gate collisions by sampling the metric function using Metropolis-Hasting, an efficient version of Monte Carlo sampling that is non-

*Equal contribution. The authors are with the Robotics and Perception Group, Department of Informatics, University of Zurich, and Department of Neuroinformatics, University of Zurich and ETH Zurich, Switzerland (<http://rpg.ifi.uzh.ch>). This work was supported by the Swiss National Science Foundation (SNSF) through the National Centre of Competence in Research (NCCR) Robotics, the European Union’s Horizon 2020 Research and Innovation Programme under grant agreement No. 871479 (AERIAL-CORE), and the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).

gradient based. While AutoTune has shown to outperform existing methods, we show that it has difficulties when dealing with controllers where the dimensionality of the parameter space increases. Both these related works have only dealt with a relatively small dimensionality in the parameter space (less than 10 parameters).

Recently, Policy-Search methods have successfully been used for designing a Model Predictive Controller (MPC) to perform agile flight within very complex environments [9]. The authors exploited the advantages of Policy Search to learn complex policies using data from past experience and applied them to achieve optimal-control performance in the task of agilely flying through fast-moving gates.

In this paper, we propose the use of Policy Search—in particular, Weighted Maximum Likelihood (WML)—for searching the space of hyperparameters of a Model Predictive Contouring Controller (MPCC). In contrast to end-to-end learning, where the agent needs to learn everything from data, having an optimal controller in the loop leverages the use of a model by the optimizer, which allows a sample efficiency that permits training in a high-fidelity simulator [10]. We show that the controller has the ability to transfer zero-shot to the real world by flying extremely agile maneuvers.

Contributions

The contributions of this paper are: (i) we show that we are able to find the optimal hyperparameters for a state-of-the-art MPCC controller in the task of time-optimal flight, pushing the performance of the controller to the best lap times to date. We show this by outperforming the best tuning that a human expert was able to achieve after trying for several weeks, in both simulation and the real world; and (ii) we demonstrate the ability to transfer the results obtained in simulation to the real world by flying through two different tracks at speeds up to 75 km h^{-1} .

II. RELATED WORK

A. Automatic controller tuning

There exist different methods for automatically tuning controllers [11]–[13]. In line with adaptive control [14], the classic approach (known as the MIT rule [15]) for controller tuning analytically finds the relationship between a performance metric, e.g. tracking error or trajectory completion, and optimizes the parameters with gradient-based optimization [16]–[18]. However, expressing the long term measure (in our case the lap time and gate passing metrics) as a function of the tuning parameters is impractical and generally intractable. Instead of analytically computing it, another line of work proposes to iteratively estimate the optimization function, and use the estimate to find optimal parameters [19]–[21]. However, these methods make over-simplifying assumptions on the objective function, e.g. convexity or relative Gaussianity between observations. Such assumptions are generally not suited for controller tuning to high-speed flight, where the function is highly non-convex.

Other approaches are mainly data-driven and directly focus on using a receding horizon optimal control approach in the

center of a learning algorithm. Learning-based MPC [22]–[32] can leverage real-world data to improve dynamic modeling or learn a cost function for MPC. It allows for a more robust and flexible MPC design. In particular, sampling-based MPC [27] algorithms are developed for handling complex cost criteria and general nonlinear dynamics. This is achieved by combining neural networks for the system dynamics approximation with the model predictive path integral (MPPI) control framework [27] for real-time control optimization. A crucial requirement for the sampling-based MPC is to generate a large number of samples in real-time, where the sampling procedure is generally performed in parallel by using graphics processing units (GPUs). Hence, it is computationally and memory expensive to run sampling-based MPC on embedded systems. These methods generally focus on learning dynamics for tasks where a dynamical model of the robots or its environment is challenging to derive analytically, such as aggressive autonomous driving around a dirt track [28]. More recent approaches [13], [33], [34] pose the optimization problem as an additional differentiable layer through which one can apply gradient-based optimization methods, e.g. back-propagation. In [33], they are able to directly find the gradient of the MPC controller with respect to its high level hyperparameters. They use the KKT (Karush-Kuhn-Tucker) conditions of the convex approximation to find this gradient. This allows treating MPC controllers as an extra layer in an end-to-end learning structure. Another work in this regard applies reverse mode auto-differentiation and proves to be robust to uncertainties in the system dynamics and environment in simulation [13].

B. High-speed quadrotor flight

The literature on high-speed, high agile quadrotor starts with polynomial planning and classical control. The main focus was directed to exploiting the differential flatness property of the quadrotor and leverage the use of polynomial for planning [35]–[38].

More recently, optimization-based methods have achieved true time-optimal trajectory planning using the quadrotor dynamics and numerical optimization [39]. However, due to high computational complexity, it normally requires several minutes or hours to solve, and the trajectory can only be computed offline. Additionally, a small deviation from the planned trajectory leads to catastrophic crashes since the vehicle is already operating on the edge of its actuator limits. The solution to this problem is the current state-of-the-art minimum-time flight: Model Predictive Contouring Control (MPCC) [5], which simultaneously does online time-optimal trajectory planning and vehicle control. In contrast to standard trajectory tracking control (e.g., MPC), MPCC balances the maximization of the progress along a given path and minimizes the deviation from it. MPCC has the freedom to optimally select at runtime the states at which the reference trajectory is sampled such that the progress is maximized while the platform stays within actuator bounds.

Latest developments in quadrotor simulations have allowed for evaluation environments where we can not only

train but also evaluate and zero-shot transfer control policies to the real-world. The state of the art on realistic simulations for quadrotors is [10], which introduces a hybrid aerodynamic quadrotor model which combines blade element momentum theory with learned aerodynamic representations from highly aggressive maneuvers.

III. PRELIMINARIES

In this section, we introduce the basics of Policy Search followed by Model Predictive Contouring Control. Then, we briefly introduce quadrotor dynamics.

A. Policy Search

We summarize policy search by following the derivation from [9]. In particular, we focus on episode-based policy search (or episodic reinforcement learning). Episode-based policy search adds perturbations in the policy parameter space instead of the policy output. A reward function $R(\tau)$ is used to evaluate the quality of trajectories τ that are generated by sampled parameters ϕ . In our case, we search for a constant Gaussian policy where the set of MPCC parameters ϕ that maximizes the reward is found. Therefore, the policy parameters ϕ are updated by maximizing the expected return of sampled trajectories

$$J_\phi = \mathbb{E}[R(\tau)|\phi] \approx \int R(\tau)p_\phi(\tau)d\tau. \quad (1)$$

We focus on a probabilistic model in which the search for high-level decision variables ϕ in the MPCC optimization is treated as a probabilistic inference problem, and use the Expectation-Maximization algorithm to update the policy parameters. We use a Gaussian distribution $\mathbf{z} \sim \pi_\phi = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ to represent the policy, where $\boldsymbol{\mu}$ is a mean vector and $\boldsymbol{\Sigma}$ is a diagonal covariance matrix. The covariance matrix is needed in order to incorporate exploration. Therefore, the MPCC parameters are $\phi = [\boldsymbol{\mu}, \boldsymbol{\Sigma}]$.

B. Model Predictive Contouring Control

Consider the discrete-time dynamic system with continuous state and input spaces, $\mathbf{x}_k \in \mathcal{X}$ and $\mathbf{u}_k \in \mathcal{U}$ respectively. Let us denote the time discretized evolution of the system $f: \mathcal{X} \times \mathcal{U} \mapsto \mathcal{X}$ such that $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k)$.

In contrast to traditional MPC approaches, where the objective is to track a dynamically feasible trajectory, the main objective of MPCC is to maximize the predicted traveled distance along a path. This idea is encoded in the controller by introducing the concept of *progress*, which is the travelled arc-length measured along the nominal path. In this paper, we aim to tune this MPCC controller for the task of drone racing. To that end, we consider the MPCC problem in its full form:

$$\begin{aligned} \pi(\mathbf{x}) = \underset{\mathbf{u}}{\operatorname{argmin}} \quad & \sum_{k=0}^N \|e^l(\theta_k)\|_{q_l}^2 + \|e^c(\theta_k)\|_{q_c}^2 + \|\boldsymbol{\omega}_k\|_{Q_\omega}^2 \\ & + \|\Delta v_{\theta_k}\|_{\tau_{\Delta v}}^2 + \|\Delta \mathbf{f}_k\|_{\mathbf{R}_{\Delta f}}^2 - \mu v_{\theta, k} \\ \text{subject to} \quad & \mathbf{x}_0 = \mathbf{x} \\ & \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) \\ & \mathbf{g}_{lb} \leq \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k) \leq \mathbf{g}_{ub} \end{aligned} \quad (2)$$

where \mathbf{g}_{lb} and \mathbf{g}_{ub} encode the lower and upper bounds on states and inputs, θ_N is the progress at the last horizon step, and e_k^c is the contour error at time k , which is the position error between the current position and the nominal path. For more details about the problem definition we refer the reader to [5].

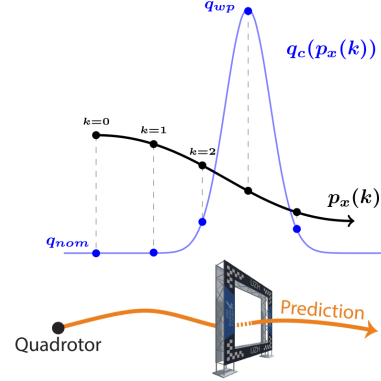


Fig. 3: Dynamic allocation of the weight in the x axis. When closer to the gate, the contouring weight $q_c(p_x(k))$ (shown in blue) grows to q_{wp} , decreasing the tracking error around the gate. Everywhere else q_c is kept constant and equal to q_{nom} such that the MPCC has more freedom to optimize for progress.

This MPCC algorithm has been adapted for the task of drone racing by adding a varying state-dependent cost to the q_c parameters, in the form of a scaled Gaussian distribution located at every waypoint (or gate) that needs to be passed. For every gate, the height h and the width w of the Gaussian are defined, which are additional tuning parameters. This way, a region of attraction at every gate is created. Therefore, the set of tunable parameters ϕ are:

$$\phi = \left[h_0, w_0, \dots, h_{n_g-1}, w_{n_g-1}, q_{nom}, r_{\Delta v}, r_{\Delta f}, \mu \right]$$

where n_g is the total number of gates. The dimensionality of the hyperparameter space is therefore $2n_g + 4$. The cost function of optimization problem (2) can then be written as $J_\phi(\mathbf{x}, \mathbf{u})$, where we need to find the set of optimal hyperparameters ϕ^* that solve a given higher level task, which will be encoded by a reward function $R(\tau)$.

C. Quadrotor Dynamics

In this section, we describe the specific dynamics used, referred in (2) as $f(\mathbf{x}_k, \mathbf{u}_k)$. The quadrotor's state space is described from the inertial frame I to the body frame B , as $\mathbf{x} = [\mathbf{p}_{IB}, \mathbf{q}_{IB}, \mathbf{v}_{IB}, \boldsymbol{\omega}_B]^T$ where $\mathbf{p}_{IB} \in \mathbb{R}^3$ is the position, $\mathbf{q}_{IB} \in \mathbb{S}\mathbb{O}(3)$ is the unit quaternion that describes the rotation of the platform, $\mathbf{v}_{IB} \in \mathbb{R}^3$ is the linear velocity vector, and $\boldsymbol{\omega}_B \in \mathbb{R}^3$ are the bodyrates in the body frame. The input of the system is given as the collective thrust $\mathbf{f}_B = [0 \ 0 \ f_{Bz}]^T$ and body torques $\boldsymbol{\tau}_B$. For readability, we drop the frame indices as they are consistent throughout the description. The dynamic equations are

$$\begin{aligned} \dot{\mathbf{p}} &= \mathbf{v} & \dot{\mathbf{v}} &= \mathbf{g} + \frac{\mathbf{R}(\mathbf{q})\mathbf{f}_T}{m} - \mathbf{R}(\mathbf{q})\mathbf{D}\mathbf{R}^T(\mathbf{q}) \cdot \mathbf{v} \\ \dot{\mathbf{q}} &= \frac{\mathbf{q}}{2} \odot [0 \ \boldsymbol{\omega}]^T & \dot{\boldsymbol{\omega}} &= \mathbf{J}^{-1}(\boldsymbol{\tau} - \boldsymbol{\omega} \times \mathbf{J}\boldsymbol{\omega}) \end{aligned} \quad (3)$$

where \odot represents the Hamilton quaternion multiplication, $\mathbf{R}(\mathbf{q})$ the quaternion rotation, m the quadrotor's mass, and \mathbf{J} the quadrotor's inertia.

Additionally, the input space given by \mathbf{f} and $\boldsymbol{\tau}$ is decomposed into single rotor thrusts $\mathbf{f} = [f_1, f_2, f_3, f_4]$ where f_i is the thrust at rotor $i \in \{1, 2, 3, 4\}$.

$$\mathbf{f}_T = \begin{bmatrix} 0 \\ 0 \\ \sum f_i \end{bmatrix} \quad \text{and} \quad \boldsymbol{\tau} = \begin{bmatrix} l/\sqrt{2}(f_1 + f_2 - f_3 - f_4) \\ l/\sqrt{2}(-f_1 + f_2 + f_3 - f_4) \\ c_\tau(f_1 - f_2 + f_3 - f_4) \end{bmatrix} \quad (4)$$

with the quadrotor's arm length l and the rotor's torque constant c_τ . In order to approximate the most prominent aerodynamic effects, the quadrotor's dynamics include a linear drag model [40]. In (3), \mathbf{D} is the diagonal matrix with drag coefficients d such that $\mathbf{D} = \text{diag}(d_x, d_y, d_z)$.

IV. PROBLEM FORMULATION

In this section, we combine the approaches introduced in Sections III-A and III-B. We treat MPCC as a controller $\boldsymbol{\tau} = \text{MPCC}(\mathbf{z})$ that is parameterized by the high-level decision variables \mathbf{z} . Here, $\boldsymbol{\tau} = [\mathbf{u}_h, \mathbf{x}_h]_{h \in 1, \dots, H}$ is a trajectory generated by MPCC given \mathbf{z} , where \mathbf{u}_h are control commands and \mathbf{x}_h are corresponding states of the robot. By perturbing \mathbf{z} , MPCC can result in completely different trajectories $\boldsymbol{\tau}$. For the MPCC to execute the given task in minimum time, the optimal \mathbf{z} has to be defined in advance.

To formulate the policy search as a latent variable inference problem, similar to [9], we introduce a binary *reward event* as an observed variable, denoted as $E = 1$. Maximizing the reward signal implies maximizing the probability of this *reward event*. The probability of this reward event is given by $p(E|\boldsymbol{\tau}) \propto \exp\{R(\boldsymbol{\tau})\}$, where $R(\boldsymbol{\tau})$ is a reward function for evaluating the goodness of the MPCC solution $\boldsymbol{\tau}$ with respect to a given evaluation metric of the task. This leads to the following maximum likelihood problem [41]:

$$\max_{\phi} \log p_{\phi}(E = 1) = \log \int_{\boldsymbol{\tau}} p(E|\boldsymbol{\tau}) p_{\phi}(\boldsymbol{\tau}) d\boldsymbol{\tau}, \quad (5)$$

which is intractable to solve directly and can be approximated efficiently using Monte-Carlo Expectation-Maximization (MC-EM) [42], [43]. MC-EM algorithms find the maximum likelihood solution for the log marginal-likelihood (5) by introducing a variational distribution $q(\boldsymbol{\tau})$, and then, decompose the marginal log-likelihood into two terms:

$$\log p_{\theta}(E = 1) = \mathcal{L}_{\theta}(q(\boldsymbol{\tau})) + D_{\text{KL}}(q(\boldsymbol{\tau}) || p_{\theta}(\boldsymbol{\tau}|E)) \quad (6)$$

where the D_{KL} is the Kullback–Leibler (KL) divergence between $q(\boldsymbol{\tau})$ and the reward-weighted trajectory distribution $p_{\theta}(\boldsymbol{\tau}|E)$.

The MC-EM algorithm is an iterative method that alternates between performing an Expectation (E) step and a Maximization (M) step. In the expectation step, we minimize the KL-divergence D_{KL} , which is equivalent to setting $q(\boldsymbol{\tau}) = p_{\theta}(\boldsymbol{\tau}|E) \propto p(E|\boldsymbol{\tau})p_{\theta}(\boldsymbol{\tau})$. In the maximization step,

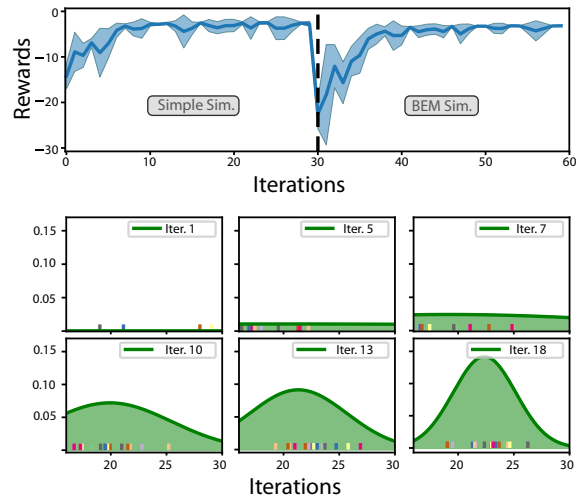


Fig. 4: Top: reward evolution as training progresses. The rewards have been computed for 5 training runs averaged. In the middle of the training we switch from simple simulator to BEM simulator. This way, we leverage the faster running time in simple simulation and transfer the policy from simple to BEM. Bottom: example of how a the policy evolves and converges for a single parameter through the iterations.

we update the policy parameters by maximizing the expected complete data log-likelihood

$$\phi^* = \arg \max_{\phi} \sum_i p(E|\boldsymbol{\tau}^{[i]}) \log p_{\phi}(\boldsymbol{\tau}^{[i]}) \quad (7)$$

where each sample $\boldsymbol{\tau}^{[i]}$ is weighted by the probability of the *reward event*, denoted as $p(E|\boldsymbol{\tau})$. The trajectory distribution $p_{\theta}(\boldsymbol{\tau}^{[i]})$ can be replayed by the high-level policy π_{θ} . To transform the reward signal $R(\boldsymbol{\tau}^{[i]})$ of a sampled trajectory $\boldsymbol{\tau}^{[i]}$ into a probability distribution of the *reward event*, we use the exponential transformation [41], [44], [45]:

$$d^{[i]} = p(E|\boldsymbol{\tau}) = \exp\{\beta R(\boldsymbol{\tau}^{[i]})\} \quad (8)$$

where the parameter $\beta \in \mathbb{R}_+$ denotes the inverse temperature of the soft-max distribution, higher value of β implies a more greedy policy update.

A. Learning-Based Hyperparameter Tuning for MPCC

As explained in Section IV, to find the optimal hyperparameters ϕ^* , we need to sample a set of trajectories $\boldsymbol{\tau}^{[i]}$ for each episode. We start with a random guess for π_{ϕ} , then we sample N different $\phi^{[i]}$. From every $\phi^{[i]}$, we run a simulation experiment that consists of completing L laps with the MPCC in the loop and a reward $R(\boldsymbol{\tau}^{[i]})$ is calculated. After all trajectories $\boldsymbol{\tau}^{[i]}$ are completed an update step can be made.

We designed a reward function specifically for the drone racing task such that it results in robust learning and can be directly used to find the optimal policy in the large action space for different tracks. The reward function choice is driven by three objectives: minimizing the lap time, avoiding gate collisions, and maximizing the number of gates passed.

TABLE I: We show the results obtained in simulation in both the Split-S track and in the Marv track. Our approach is able to find optimal solutions in both tracks, whereas the baselines struggle to find solutions in high dimensional hyperparameter spaces.

	Split-S track (18 hyperparameters)				Marv track (30 hyperparameters)			
	Lap Time [s]		Success Rate [%]		Lap Time [s]		Success Rate [%]	
	Simple	BEM	Simple	BEM	Simple	BEM	Simple	BEM
Random	6.024 ± 0.519	6.031 ± 0.517	34	27	10.45	10.48	1	1
Hand-Tuned	5.24	5.42	-	-	9.17	Collision	-	-
AutoTune	6.111 ± 0.342	6.471 ± 0.473	63	50	Collision	9.08	0	1
WML (Ours)	5.02 ± 0.008	5.274 ± 0.012	99	90	8.339 ± 0.106	8.381 ± 0.022	100	100

We define the reward function as follows.

$$\pi(\phi) = \operatorname{argmax}_{\phi} \mathbb{E} \left[\sum_{i=0}^N R(\tau^{[i]}) \right] \quad (9)$$

where $R(\tau^{[i]})$ is defined as:

$$R(\tau^{[i]}) = -t_1 - 2t_2 - 2^{10r_{miss}} + r_{pass} + \frac{n_{gp}}{n_g} \quad (10)$$

where t_1 and t_2 are the lap times of the first and second lap respectively, n_{gp} is the number of gates passed, n_g is the total number of gates, and r_{miss} is the sum of gate passing errors and

$$r_{pass} = \begin{cases} 1 & \text{if } r_{miss} \leq 0.01 \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

The miss reward, r_{miss} , is added exponentially in order to heavily penalize if the quadrotor passes the gate outside of a given threshold, which is in our case 0.5 m. We do not add r_{miss} as a linear component to ensure that the optimal solution does not miss the gates and can optimize for lower lap time, since this error can be very small compared to other terms in the reward (10). During the training r_{miss} can never be greater than 1.5 times the gate passing threshold hence the exponential never explodes. If the quadrotor passes the gate with a distance larger than 1.5 times the gate passing threshold, then we count this as a gate crash and give a large negative reward. If it passes all the gates with zero error we give a positive reward in r_{pass} (11). Additionally, we give higher weight to the t_2 than t_1 because the first lap is slower than the second one and the training prefers to decrease t_1 more. By giving a higher weight to the t_2 we decrease them both simultaneously.

V. EXPERIMENTS

A. Baselines

For the baselines, we choose three different approaches. First, we compare against the complete random approach, where we randomly sample the parameters. Second, we compare against the best results achieved by the previous literature where the controller was tuned by hand by an expert [5]. And third, we implement AutoTune, the current state-of-the-art in controller tuning [8] for our MPCC controller. All approaches have the same initialization strategy and allowed search space. Both AutoTune and our approach have exactly the same reward function, which is described in Eq. (10).

B. Training

The training experiments are conducted using *Flightmare* [46] and two different simulation environments for each track. One called *Simple* simulator, where the drone model is the one described in Section III-C and does not capture complex aerodynamic environments. The other one is a high-fidelity simulator called *BEM* [10], which includes a high-complexity aerodynamic simulation derived from blade element momentum theory.

During training, for each sample, we simulate $L = 2$ consecutive laps. We check whether the quadcopter is flying through all the gates and measure the lap times in order to compute the reward (10). We sample $N = 16$ set of parameters from the same Gaussian policy, and then do a policy update step for the episode. During training, our criterion of gate passing is that the drone passes within 0.5m of the gate center.

When we train directly in the BEM simulator from a random initialization, it usually takes longer to converge or it does not necessarily find the optimal solution. To address this, we first train in the simple simulator for 30 episodes and then initialize the training in BEM from the solution obtained in the Simple simulator. We reset the variance when we begin the training in the BEM environment so that the policy recovers its full exploration capabilities and perform another 30 episodes in BEM. Each episode takes around 35 seconds in Simple simulator and 200 seconds in BEM simulator around. The evolution of the rewards in Simple and BEM is shown in Figure 4.

C. Results in Simulation

In this section, we present the results achieved by our approach when compared against the baselines for two different tracks. The reference for the Split-S track has been generated with [39] while for the Marv track, we used [47]. For the comparison, we run 10 trials with different random initializations with the three approaches (Random, AutoTune and Ours). We draw 100 samples from our policy and the best 100 samples from Random and AutoTune. We present the results in Table I and in Fig. 5. We validate our results with gate passing threshold 0.6 m to the gate center, as the gate width is 1.2 m in real-world. As shown in Table I, our approach greatly outperforms both in laptimes, speed, success rate and consistency to all other baselines. Even for the Marv track, where hand-tuning was impossible due

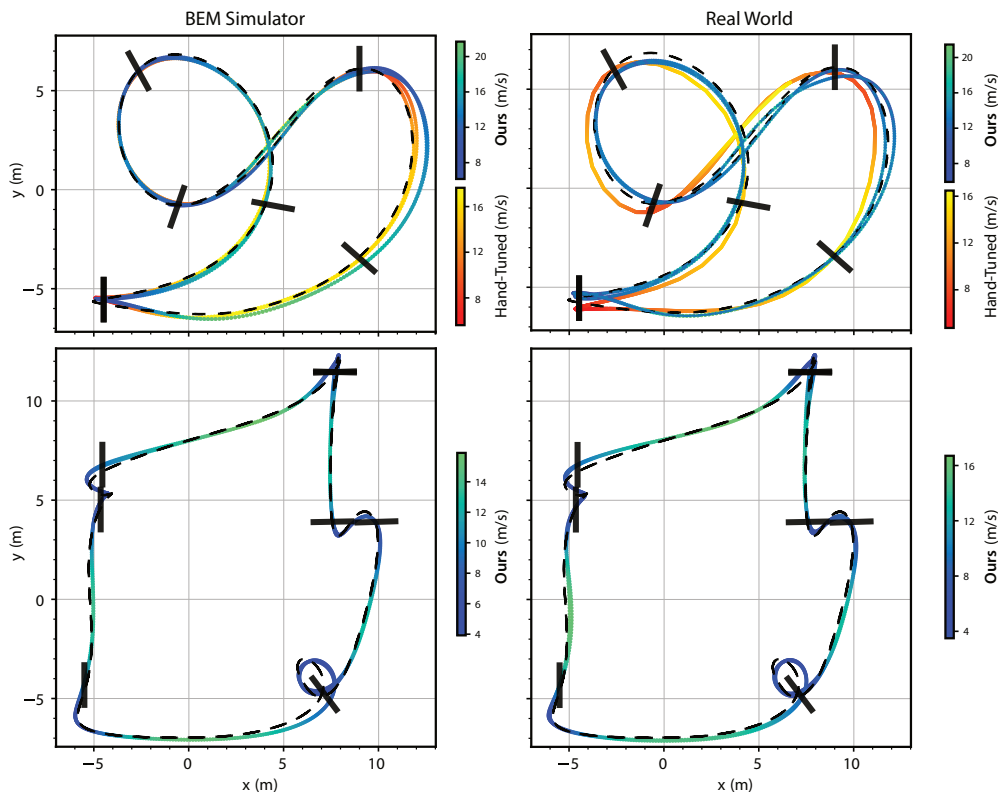


Fig. 5: Top: Split-S track. Bottom: Marv track. References are shown in dashed. Comparison of hand-tuned and our method, in simulation and in the real-world. The speeds achieved by our method are consistently higher than those achieved by the hand-tuned approaches. In fact, the human expert was not able to hand-tune a controller that was able to fly the Marv track. The tuning in simulation transfers to the real-world in both tracks.

to the complexity of the track and the higher amount of tuning parameters, i.e. 30, our approach finds a set of optimal hyperparameters that fly through the track with 100% success rate, whereas the baselines struggle to find a successful solution.

D. Results in Real World

We deploy the algorithm to the physical platform, which was built in-house from off-the-shelf components. A more detailed description of our platform can be found in [48].

The flying results in real-world for both tracks are shown in Fig. 5. One remark is that the learning-based policy transfers zero-shot to the real-world for both tracks without issue. We can see this by comparing them in XY position and in maximum velocity (blue colorbars). There is no hand-tuned data for the Marv track because it was not possible for a human expert to tune a controller that could complete the track without any collision.

Additionally, it is noteworthy to mention that with learning-based approach, we are able to get the fastest lap times ever achieved with the MPCC controller in these tracks. For the Split-S track, the fastest lap time of the hand-tuned controller in the real world, as reported in [5], is of 5.8s, with a maximum speed of 60 km/h, whereas the fastest lap time achieved by our learning-based tuning method is 5.29 s, with a maximum speed of 75 km/h. This improvement of around half a second shows how the proposed method can bring out the full potential of the MPCC controller.

VI. CONCLUSION

In this paper we show that by using the proposed learning-based algorithm to tune a model-based controller (in our case a MPCC), we can achieve performance limits that were impossible to achieve using any other approach. The data efficiency acquired by having a model-based controller in the loop allows us to train directly in a high fidelity simulator, which in turn has shown to allow zero-shot transfer to the real-world. A possible extension of this work would be to, instead of fixing the regions of attraction to Gaussians (as depicted in Fig. 3), one could let the data-driven algorithm to directly learn the best appropriate set of weights in the horizon at every time-step.

One limitation of our method is that when starting from a random initialization for ϕ , in some cases the algorithm leads to non-increasing rewards. A solution for this is to use a warm-starting strategy. We partially address this by initiating the training in the simple simulator, then continuing in the BEM simulator. Another limitation is that we need to re-train for every new track, since the size of the parameter space depends on the number of gates. This can be tackled by resorting to more complex network architectures and making the parameter set track independent. Lastly, the reward function is specifically tailored for drone racing objectives and to adapt it to other domains, one would need to re-design it. These results open the door to finding new real-world applications for controllers where the high-level objective can be expressed in terms of a reward function directly.

REFERENCES

- [1] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [2] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "PAMPC: Perception-aware model predictive control for quadrotors," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2018, pp. 1–8.
- [3] D. Lam, C. Manzie, and M. Good, "Model predictive contouring control," in *49th IEEE Conference on Decision and Control (CDC)*, IEEE, 2010, pp. 6137–6142.
- [4] D. Lam, C. Manzie, and M. C. Good, "Model predictive contouring control for biaxial systems," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 2, pp. 552–559, 2012.
- [5] A. Romero, S. Sun, P. Foehn, and D. Scaramuzza, "Model predictive contouring control for time-optimal quadrotor flight," *IEEE Transactions on Robotics*, pp. 1–17, 2022.
- [6] A. Romero, R. Penicka, and D. Scaramuzza, "Time-optimal online replanning for agile quadrotor flight," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7730–7737, 2022.
- [7] L. P. Fröhlich, C. Küttel, E. Arcari, L. Hewing, M. N. Zeilinger, and A. Carron, "Contextual tuning of model predictive control for autonomous racing," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 10555–10562.
- [8] A. Loquercio, A. Saviolo, and D. Scaramuzza, "Autotune: Controller tuning for high-speed flight," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4432–4439, 2022.
- [9] Y. Song and D. Scaramuzza, "Policy search for model predictive control with application to agile drone flight," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2114–2130, 2022.
- [10] L. Bauersfeld, E. Kaufmann, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," *RSS: Robotics, Science, and Systems*, 2021.
- [11] A. Schperberg, S. Di Cairano, and M. Menner, "Auto-tuning of controller and online trajectory planner for legged robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7802–7809, 2022.
- [12] M. Zanon and S. Gros, "Safe reinforcement learning using robust mpc," *IEEE Transactions on Automatic Control*, vol. 66, no. 8, pp. 3638–3652, 2020.
- [13] S. Cheng, M. Kim, L. Song, Z. Wu, S. Wang, and N. Hovakimyan, "DiffTune: Auto-tuning through auto-differentiation," *arXiv preprint arXiv:2209.10021*, 2022.
- [14] D. Hanover, P. Foehn, S. Sun, E. Kaufmann, and D. Scaramuzza, "Performance, precision, and payloads: Adaptive nonlinear mpc for quadrotors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 690–697, 2021.
- [15] K. J. Åström, "Theory and applications of adaptive control—a survey," *automatica*, vol. 19, no. 5, pp. 471–486, 1983.
- [16] M. J. Grimble, "Implicit and explicit lqg self-tuning controllers," *IFAC Proceedings Volumes*, vol. 17, no. 2, pp. 941–947, 1984.
- [17] K. J. Åström, T. Hägglund, C. C. Hang, and W. K. Ho, "Automatic tuning and adaptation for pid controllers—a survey," *Control Engineering Practice*, vol. 1, no. 4, pp. 699–714, 1993.
- [18] M. A. Mohd Basri, A. R. Husain, and K. A. Danapalasingam, "Intelligent adaptive backstepping control for mimo uncertain nonlinear quadrotor helicopter systems," *Transactions of the Institute of Measurement and Control*, vol. 37, no. 3, pp. 345–361, 2015.
- [19] M. Menner and M. N. Zeilinger, "Maximum likelihood methods for inverse learning of optimal controllers," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 5266–5272, 2020.
- [20] F. Berkenkamp, A. P. Schoellig, and A. Krause, "Safe controller optimization for quadrotors with gaussian processes," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2016, pp. 491–496.
- [21] A. Marco, P. Hennig, J. Böhm, S. Schaal, and S. Trimpe, "Automatic lqr tuning based on gaussian process global optimization," in *2016 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2016, pp. 270–277.
- [22] W. Edwards, G. Tang, G. Mamakoukas, T. Murphey, and K. Hauser, "Automatic tuning for data-driven model predictive control," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 7379–7385.
- [23] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.
- [24] N. A. Spielberg, M. Brown, and J. C. Gerdes, "Neural network model predictive motion control applied to automated driving with unknown friction," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 5, pp. 1934–1945, 2021.
- [25] K. Y. Chee, T. Z. Jiahao, and M. A. Hsieh, "Knode-mpc: A knowledge-based data-driven predictive control framework for aerial robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2819–2826, 2022.
- [26] A. Saviolo, G. Li, and G. Loianno, "Physics-inspired temporal learning of quadrotor dynamics for accurate model predictive trajectory tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10256–10263, 2022.
- [27] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [28] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic mpc for model-based reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2017, pp. 1714–1721.
- [29] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, "Robust constrained learning-based nmpc enabling reliable mobile robot path tracking," *The International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.
- [30] U. Rosolia and F. Borrelli, "Learning how to autonomously race a car: a predictive control approach," *IEEE Transactions on Control Systems Technology*, 2019.
- [31] G. Torrente, E. Kaufmann, P. Föhn, and D. Scaramuzza, "Data-driven mpc for quadrotors," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3769–3776, 2021.
- [32] M. Mehndiratta, E. Camci, and E. Kayacan, "Automated tuning of nonlinear model predictive controller by reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2018, pp. 3016–3021.
- [33] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable mpc for end-to-end planning and control," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 8289–8300.
- [34] L. Pineda, T. Fan, M. Monge, S. Venkataraman, P. Sodhi, R. T. Chen, J. Ortiz, D. DeTone, A. S. Wang, S. Anderson *et al.*, "TheSeus: A library for differentiable nonlinear optimization," in *Advances in Neural Information Processing Systems*.
- [35] M. Mueller, S. Lupashin, and R. D'Andrea, "Quadcopter ball juggling," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2011, pp. 4972–4978.
- [36] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robot. Autom. Mag.*, vol. 19, no. 3, pp. 20–32, 2012.
- [37] D. Mellinger, N. Michael, and V. Kumar, "Trajectory generation and control for precise aggressive maneuvers with quadrotors," *Int. J. Robot. Research*, vol. 31, no. 5, pp. 664–674, 2012.
- [38] M. W. Mueller, M. Hehn, and R. D'Andrea, "A computationally efficient algorithm for state-to-state quadcopter trajectory generation and feasibility verification," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2013.
- [39] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, 2021. [Online]. Available: <https://robotics.sciencemag.org/content/6/56/eabh1221>
- [40] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE Robot. Autom. Lett.*, vol. 3, no. 2, pp. 620–626, Apr. 2018.
- [41] M. P. 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.
- [42] J. Kober and J. R. Peters, "Policy search for motor primitives in robotics," in *Advances in neural information processing systems*, 2009, pp. 849–856.
- [43] N. Vlassis and M. Toussaint, "Model-free reinforcement learning as mixture learning," in *Proceedings of the 26th Annual International Conference on Machine Learning*, 2009, pp. 1081–1088.
- [44] G. Neumann *et al.*, "Variational inference for policy search in changing

- situations,” in *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, 2011, pp. 817–824.
- [45] M. Toussaint, “Robot trajectory optimization using approximate inference,” in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 1049–1056.
- [46] Y. Song, S. Naji, E. Kaufmann, A. Loquercio, and D. Scaramuzza, “Flightmare: A flexible quadrotor simulator,” in *Proceedings of the 2020 Conference on Robot Learning*, 2021, pp. 1147–1157.
- [47] Y. Song, M. Steinweg, E. Kaufmann, and D. Scaramuzza, “Autonomous drone racing with deep reinforcement learning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [48] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio, and D. Scaramuzza, “Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight,” *AAAS Science Robotics*, 2022.