

Model Predictive Path Integral Control for Agile Unmanned Aerial Vehicles

Michal Minařík, Robert Pěnička, Vojtěch Vonásek, and Martin Saska

Abstract—This paper introduces a control architecture for real-time and onboard control of Unmanned Aerial Vehicles (UAVs) in environments with obstacles using the Model Predictive Path Integral (MPPI) methodology. MPPI allows the use of the full nonlinear model of UAV dynamics and a more general cost function at the cost of a high computational demand. To run the controller in real-time, the sampling-based optimization is performed in parallel on a graphics processing unit onboard the UAV. We propose an approach to the simulation of the nonlinear system which respects low-level constraints, while also able to dynamically handle obstacle avoidance, and prove that our methods are able to run in real-time without the need for external computers. The MPPI controller is compared to MPC and SE(3) controllers on the reference tracking task, showing a comparable performance. We demonstrate the viability of the proposed method in multiple simulation and real-world experiments, tracking a reference at up to 44 km h^{-1} and acceleration close to 20 m s^{-2} , while still being able to avoid obstacles. To the best of our knowledge, this is the first method to demonstrate an MPPI-based approach in real flight.

I. INTRODUCTION

Agile control of Unmanned Aerial Vehicles (UAVs), commonly referred to as drones, is a challenging problem, mainly when flying in environments cluttered with obstacles as they pose, in general, non-convex spatial constraints. A typical objective of the control task is to minimize the error between a reference state and the current UAV state [1]. Yet, if the reference is close to the obstacles, the controller should make sure not to collide with the obstacle. However, the objective highly depends on the application at hand. In drone racing [2], the objective is to minimize time or the progress along a reference path [3]. Nevertheless, even in minimum-time flight through cluttered environments, one might need to combine the time optimality objective with the safety of avoiding obstacles.

The ability to fly in unknown or partially known environments with obstacles is required in many real-world UAV applications, e.g., in bridge inspection [4], inspection of power lines [5], aerial coverage scanning [6], or even for digitalization of interiors of historical monuments [7]. Among others, the search-and-rescue scenario is the primary motivation for this work as the fast flight can enable the

The authors are with the Multi-robot Systems Group, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic (<http://mrs.felk.cvut.cz/>). This work has been supported by the Czech Science Foundation (GAČR) under research project No. 23-06162M, by the European Union under the project Robotics and Advanced Industrial Production (reg. no. CZ.02.01.01/00/22_008/0004590) and by CTU grant no SGS23/177/OHK3/3T/13. Computational resources were provided by the e-INFRA CZ project (ID:90254), supported by the Ministry of Education, Youth and Sports of the Czech Republic.

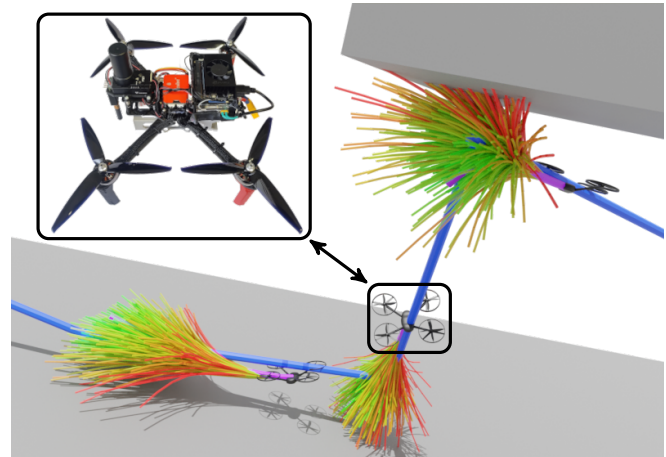


Fig. 1: The MPPI controller follows a target trajectory (blue) while avoiding collisions with obstacles (gray) by parallel prediction and weighing of hundreds of future trajectories (colored curves). Our proposed algorithm is able to run entirely onboard and online.

UAVs to reach possible victims faster, even in hard-to-fly environments such as forests or partially collapsed buildings.

The challenges of controlling a UAV in fast and agile flight include modeling obstacles, handling dynamics constraints, and accommodating objectives given by the application while computing the control commands online with high frequency. This is very demanding for the classical autonomous-flight pipeline, where planning is separated from control. While trajectory planning accounts for collisions and time allocation of UAV states, it usually does not account for physical constraints such as maximal thrust or angular speeds. On the other hand, the control usually does not account for obstacle avoidance, which poses a challenge during agile flight, where the deviation from the planned trajectory is typically larger. Moreover, the controller has to run online with high (e.g., 100 Hz) frequency, which is challenging with the increased control problem complexity, such as when considering obstacles or minimum time objectives.

Existing works that can handle various objectives and constraints are primarily based on optimization-based linear [8] or Nonlinear [9] Model Predictive Control (NMPC). The NMPC allows the use of a more precise nonlinear model of the drone dynamics, therefore allowing to constrain the UAV thrust and angular speed limits [2]. Nevertheless, the collision avoidance constraints are either considered as simplified spherical obstacles [10], or a limited number of obstacles is modeled [11], [12] at a control frequency of only 10–20 Hz. However, the other restrictions still hold,

such as universal avoidance of non-analytic¹ and non-convex obstacles, which are almost impossible to introduce into the NMPC framework. Moreover, the objective of the NMPC also requires to be an analytic function, usually in quadratic form.

In this paper, we use a different predictive control scheme, namely the Model Predictive Path Integral (MPPI) [13], [14] originally introduced for car racing. Using a sampling-based approach that leverages stochastic optimal control, the MPPI removes the restrictions on the cost function, modeled dynamics, and state constraints at the expense of a high computation demand. This demand is satisfied using the Compute Unified Device Architecture (CUDA) platform, allowing the parallelization of the calculations on a Graphics Processing Unit (GPU). We propose a method able to track a reference trajectory while using a metric approximation more suitable for UAV rotations. We show the ability to add an arbitrary collision detection module into the control cost function, thereby including collision avoidance into the control. We demonstrate all the mentioned abilities on a real-world platform and show that using our approach, it is possible to perform all computations online (at 100 Hz) and, more importantly, onboard a real drone while satisfying multiple constraints imposed on the controls and states. To the best of our knowledge, this is the first MPPI-based approach that can run online on a real drone (compared to existing related works done in simulation) flying fast and agile trajectories. We show tracking of a reference at speeds of up to 44 km h^{-1} and acceleration close to 20 m s^{-2} , even in situations where obstacles are present.

II. RELATED WORK

The proposed MPPI controller belongs to model-based control methods that use the UAV dynamics to compute the control commands for a given task. To this end, we overview the most relevant model-based control approaches [1].

One of the most popular controllers that allow agile flight is the geometric tracking controller [15] proposed on the Special Euclidean group $SE(3)$. The $SE(3)$ controller, also utilized within the MRS UAV Drone system [16], [17], and later compared to the proposed MPPI approach, is able to avoid singularities commonly associated with Euler angle formulations. The quadrotor dynamics are shown to be differentially flat [18], which enables the computation of both the full state and the control inputs to the system given high-order derivatives of the flat outputs (typically position and heading). This allows for precise control when the trajectory is given as a high-degree polynomial. The differential flatness can also be leveraged for control when modeling a multi-rotor UAV with linear rotor drag as shown in [19].

One of the most popular and robust control strategies for UAVs is the Model Predictive Control (MPC) [8]. Its nonlinear version (NMPC) has even been shown to be able to fly time-optimal agile trajectories designed for drone racing [2].

¹not described by well-known mathematical expressions, such as polynomials and trigonometric functions

A variant of NMPC specially designed for minimum-time flight, called Model Predictive Contouring Control (MPCC), is capable of computing time allocation of a path during flight [3], [20]. In [21], the authors propose L1-NMPC, which cascades NMPC to an L1 adaptive controller, showing robustness with respect to unmodeled aerodynamic effects, varying payloads, and parameter mismatch. A comparative study of nonlinear MPC-based and differential-flatness-based control for quadrotor agile flight is shown in [22]. While some MPC-based approaches can handle obstacles, they either use only simplified spherical obstacles [10] or a small number of more complex obstacles [11], [12]. At the same time, the control pipeline can run only at a frequency of only 10-20Hz. This prohibits the deployment of the obstacle-aware MPCs to fast flights.

The aforementioned methods are limited in integrating obstacle avoidance. To cope with obstacles, sampling-based approaches can be utilized. In [23], the authors plan a path for an aerial robot in an environment with obstacles by sampling a set of possible future motion primitives and selecting the best based on task-specific criteria. Such use of sampling allows general, non-convex and non-differentiable costs, but the result is a high-level path that needs to be followed using additional control algorithms, decoupling the planning and control tasks. To sample within low-level control sequences, Model Predictive Path Integral can be used, allowing the use of a wide range of cost functions [13], [24], depending on the desired task.

In [25], the authors tackle the task of autonomous drone racing, utilizing a cost function that aggregates 3D centerline following and gate passing, while maintaining the commanded speed and keeping the drone in a given air corridor. MPPI is used to compute desired collective thrust and body rates, which are then passed to a lower-level controllers. However, the algorithms are tested in a racing simulation environment with perfect state estimation and the authors do not discuss the hardware requirements needed to sample the numerous future trajectories (7200) at the desired rate (50 Hz). Using MPPI to navigate a drone through an environment with obstacles is discussed in [26]. The authors use MPPI to sample minimum jerk trajectories [27], whose collision with the environment is then evaluated. However, the computations are still performed off-board, and the experiments are done in an environment with Vicon.

In this work, we propose an approach allowing real-time and fully onboard agile control of a UAV. This is achieved by utilizing the MPPI algorithm and taking into account the low-level constraints imposed on the motors directly in the optimization task. Moreover, we leverage the ability of MPPI to handle more general cost functions, with an important example being universal obstacle avoidance. Most importantly, we design the rollout prediction and evaluation process with real-world deployment in mind and experimentally prove that the methods are able to run on board a real drone in real-time. To the best of our knowledge, this is the first MPPI-based used for flying outdoors with a real drone, with all computations running onboard.

III. METHODOLOGY

A. Model Predictive Path Integral Control

Model Predictive Path Integral (MPPI) is a predictive control algorithm designed to control nonlinear systems [13], [14]. The core idea is similar to Model Predictive Control (MPC) [8], however, instead of employing an optimization algorithm, a Monte Carlo sampling approach is used. This shift to sampling-based optimization allows for general, non-convex cost criteria [13], [24]. Moreover, since no gradient-based optimization is used to find and improve the solution, simple encodings of task descriptions with sparse (or non-existent) gradients can be utilized. An example would be testing for a set membership — yielding a constant in case the robot collides with an obstacle, zero otherwise.

The MPPI optimization algorithm relies on predicting K possible future trajectories (*rollouts*) over N discrete time samples (*prediction horizon*) in every iteration. The current state estimate is denoted \hat{x} and the nominal control sequence (sequence of N control actions obtained by initialization or previous optimization iterations) as \mathbf{u}^{nom} . K disturbance sequences of length N are sampled from a normal distribution with a zero mean and a covariance matrix Σ . We use lower index j to denote the discrete time index and upper index k to denote the rollout index:

$$\left. \begin{aligned} \mathbf{x}^k &= (\mathbf{x}_0^k, \dots, \mathbf{x}_j^k, \dots, \mathbf{x}_{N-1}^k, \mathbf{x}_N^k) \\ \mathbf{u}^k &= (\mathbf{u}_0^k, \dots, \mathbf{u}_j^k, \dots, \mathbf{u}_{N-1}^k) \\ \delta^k &= (\delta \mathbf{u}_0^k, \dots, \delta \mathbf{u}_j^k, \dots, \delta \mathbf{u}_{N-1}^k) \end{aligned} \right\} k = 1, \dots, K. \quad (1)$$

From the initial state \hat{x} , K rollouts are computed by forward simulation of the system dynamics, applying the disturbed nominal control

$$\left. \begin{aligned} \delta \mathbf{u}_j^k &\in \mathcal{N}(0, \Sigma) \\ \mathbf{u}_j^k &= \mathbf{u}_j^{\text{nom}} + \delta \mathbf{u}_j^k \\ \mathbf{x}_{j+1}^k &= \mathbf{x}_j^k + \mathbf{f}_{\text{RK4}}(\mathbf{x}_j^k, \mathbf{u}_j^k, \Delta t) \end{aligned} \right\} \begin{aligned} k &= 1, \dots, K \\ j &= 0, \dots, N-1. \end{aligned} \quad (2)$$

After the rollouts are computed, each rollout is evaluated by a task-specific cost function

$$\mathcal{S}_k = \text{ComputeCost}(\mathbf{x}^k, \mathbf{u}^k). \quad (3)$$

The lower the cost of the rollout, the better the trajectory is, considering the problem specification (visualized in Figure 2).

The costs are then transformed into weights $(\omega_1, \omega_2, \dots, \omega_K)$ by

$$\begin{aligned} \omega_k &= \frac{1}{\eta} \exp\left(-\frac{1}{\lambda} (\mathcal{S}_k - \rho)\right), \\ \eta &= \sum_{k=1}^K \exp\left(-\frac{1}{\lambda} (\mathcal{S}_k - \rho)\right), \\ \rho &= \min\{\mathcal{S}_1, \dots, \mathcal{S}_K\}. \end{aligned} \quad (4)$$

This procedure is the softmax transform used often in neural networks to normalize a vector of K values into a probability distribution, where the minimum ρ is subtracted

for numerical stability (without changing the result).

After computing the weights, the nominal control actions are updated by a weighted average of the disturbances

$$\mathbf{u}_j^{\text{nom}} := \mathbf{u}_j^{\text{nom}} + \sum_{k=1}^K \omega_k \cdot \delta \mathbf{u}_j^k. \quad (5)$$

The parameter λ scales the contribution of the trajectory rollouts to the result based on their evaluated cost, ranging from taking only the best rollout to averaging all disturbances. When λ is low, the control disturbances resulting in the best rollout have a significant impact on the updated control. As λ approaches zero, the weight vector approaches $(0, \dots, 0, 1, 0, \dots, 0)$ with one at the place of the best rollouts. Conversely, when λ is high, the weight vector approaches $(\frac{1}{K}, \dots, \frac{1}{K})$. The whole algorithm is presented in Alg. 1. To start the algorithm, we need to initialize the control over the control horizon N . In our case, we set the input to zero desired body rates and a constant collective thrust, resulting in a hover state.

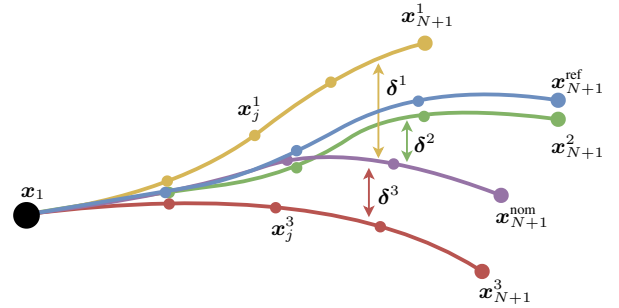


Fig. 2: To track a given trajectory (blue curve), the MPPI controller keeps nominal control (N consecutive control inputs) from previous optimization. When applied from the current state \mathbf{x}_1 , the nominal rollout is computed (purple). The K rollouts (other colored curves) are obtained by applying the actions perturbed by δ^k and evaluated with the cost function (green - best, red - worst).

B. Mathematical Model

To describe the state of the drone, we use position $\mathbf{p} \in \mathbb{R}^3$, unit quaternion rotation on the rotation group $\mathbf{q} \in \mathbb{S}\mathbb{O}(3)$ with $\|\mathbf{q}\| = 1$, velocity $\mathbf{v} \in \mathbb{R}^3$, and body rates $\boldsymbol{\omega} \in \mathbb{R}^3$. The evolution of the state $\mathbf{x} = [\mathbf{p}, \mathbf{q}, \mathbf{v}, \boldsymbol{\omega}]$ is given by dynamics

$$\begin{aligned} \dot{\mathbf{p}} &= \mathbf{v}, & \dot{\mathbf{q}} &= \frac{1}{2} \mathbf{q} \odot \begin{bmatrix} 0 \\ \boldsymbol{\omega} \end{bmatrix}, \\ \dot{\mathbf{v}} &= \frac{1}{m} \mathbf{R}(\mathbf{q}) \begin{bmatrix} 0 \\ 0 \\ F_t \end{bmatrix} + \mathbf{g}, & \dot{\boldsymbol{\omega}} &= \mathbf{J}^{-1} (\boldsymbol{\tau} - \boldsymbol{\omega} \times \mathbf{J} \boldsymbol{\omega}), \end{aligned} \quad (6)$$

where $\mathbf{R}(\mathbf{q})$ is the matrix representation of the quaternion \mathbf{q} , \odot represents quaternion multiplication, m is the drone's mass, \mathbf{g} is the acceleration due to gravity, and \mathbf{J} is the drone's inertia matrix. The control inputs are the torques $\boldsymbol{\tau} = (\tau_x, \tau_y, \tau_z)$ and the collective thrust F_t . The proposed MPPI controller will compute the body rates $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)$ and the collective thrust F_t , and send them to a lower-level flight controller (described in more detail in Section IV).

Algorithm 1: Model Predictive Path Integral Control

Input: Input initialization \mathbf{u}_{init} , Problem specific cost $\text{ComputeCost}(\mathbf{x}^k, \mathbf{u}^k)$, System dynamics \mathbf{f}_{RK4}
Params: Number of rollouts K and time steps N , noise covariance Σ , time step Δt

```
1 for  $j = 0, \dots, N-1$  do // Initialize the control
2    $\mathbf{u}_j^{\text{nom}} = \mathbf{u}_{\text{init}}$ 
3 while task not completed do
4    $\hat{\mathbf{x}} = \text{CurrentStateEstimate}()$ 
5   for  $k = 1, \dots, K$  do
6     // Simulate  $K$  rollouts
7      $\mathbf{x}_0^k = \hat{\mathbf{x}}$ 
8      $\delta^k = (\delta \mathbf{u}_0^k, \dots, \delta \mathbf{u}_{N-1}^k), \delta \mathbf{u}_j^k \in \mathcal{N}(0, \Sigma)$ 
9     //  $N$  steps into the future
10    for  $j = 0, \dots, N-1$  do
11       $\mathbf{u}_j^k = \mathbf{u}_j^{\text{nom}} + \delta \mathbf{u}_j^k$ 
12       $\mathbf{x}_{j+1}^k = \mathbf{x}_j^k + \mathbf{f}_{\text{RK4}}(\mathbf{x}_j^k, \mathbf{u}_j^k, \Delta t)$ 
13      // and evaluate their cost
14       $\mathcal{S}_k = \text{ComputeCost}(\mathbf{x}^k, \mathbf{u}^k)$ 
15    // Compute weights using Eq. 4
16     $(\omega_1, \omega_2, \dots, \omega_K) = \text{ComputeWeights}(\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K)$ 
17    for  $j = 0, \dots, N-1$  do
18       $\mathbf{u}_j^{\text{nom}} = \mathbf{u}_j^{\text{nom}} + \sum_{k=1}^K \omega_k \cdot \delta \mathbf{u}_j^k$  // Eq. 5
19    ApplyToSystem( $\mathbf{u}_0^{\text{nom}}$ )
20    for  $j = 0, \dots, N-2$  do
21       $\mathbf{u}_j^{\text{nom}} = \mathbf{u}_{j+1}^{\text{nom}}$  // Shift nominal control
22   $\mathbf{u}_{N-1}^{\text{nom}} = \text{Initialize}(\mathbf{u}_{N-1}^{\text{nom}})$ 
```

Therefore, we introduce *desired body rates* and *collective thrust*

$$\mathbf{u}_d = \begin{bmatrix} F_t \\ \omega_{xd} \\ \omega_{yd} \\ \omega_{zd} \end{bmatrix} = \begin{bmatrix} F_t \\ \boldsymbol{\omega}_d \end{bmatrix}, \quad (7)$$

used by the MPPI controller as the control inputs \mathbf{u}_j^k . However, we need to ensure that the commanded change in body rates is feasible (due to motor dynamics, we cannot expect the body rates to change precisely and arbitrarily as we command). We compute the needed change in body rates over the time step Δt as

$$\dot{\boldsymbol{\omega}}_d = \frac{1}{\Delta t} (\boldsymbol{\omega}_d - \boldsymbol{\omega}). \quad (8)$$

From $\dot{\boldsymbol{\omega}}_d$ we can compute the desired body torques

$$\boldsymbol{\tau}_d = \mathbf{J}\dot{\boldsymbol{\omega}}_d + \boldsymbol{\omega} \times \mathbf{J}\boldsymbol{\omega}, \quad (9)$$

and the desired single rotor thrusts \mathbf{T}_d generating the body torques $\boldsymbol{\tau}_d$ and collective thrust F_t

$$\mathbf{T}_d = \boldsymbol{\Gamma}^{-1} \begin{bmatrix} F_t \\ \boldsymbol{\tau}_d \end{bmatrix}. \quad (10)$$

The matrix $\boldsymbol{\Gamma}$ is the allocation matrix

$$\boldsymbol{\Gamma} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ -l/\sqrt{2} & l/\sqrt{2} & l/\sqrt{2} & -l/\sqrt{2} \\ -l/\sqrt{2} & l/\sqrt{2} & -l/\sqrt{2} & l/\sqrt{2} \\ -c_{tf} & -c_{tf} & c_{tf} & c_{tf} \end{bmatrix} \quad (11)$$

with l being the drone's arm length and c_{tf} being the rotor's torque constant. The single rotor thrusts \mathbf{T}_d can be clipped based on the motor constraints T_{\min} and T_{\max} of the drone

$$\mathbf{T}_{\text{clip}} = \text{clip}(T_{\min}, \mathbf{T}_d, T_{\max}). \quad (12)$$

The clipped rotor thrusts are then used to compute the clipped body torques and collective thrust

$$\begin{bmatrix} F_{t,\text{clip}} \\ \boldsymbol{\tau}_{\text{clip}} \end{bmatrix} = \boldsymbol{\Gamma} \mathbf{T}_{\text{clip}}. \quad (13)$$

Finally, we reconstruct the feasible inputs as

$$\dot{\boldsymbol{\omega}}_{\text{clip}} = \mathbf{J}^{-1} (\boldsymbol{\tau}_{\text{clip}} - \boldsymbol{\omega} \times \mathbf{J}\boldsymbol{\omega}), \quad (14)$$

$$\boldsymbol{\omega}_{\text{clip}} = \boldsymbol{\omega} + \dot{\boldsymbol{\omega}}_{\text{clip}} \cdot \Delta t. \quad (15)$$

The clipped values $[F_{t,\text{clip}}, \boldsymbol{\tau}_{\text{clip}}]$ from (13) are used to simulate the dynamics (6) and the inputs $[F_{t,\text{clip}}, \boldsymbol{\omega}_{\text{clip}}]$ from (13) and (15) are sent to the lower-level controller (line 15 in Alg. 1). This allows to consider the limits imposed on the motor thrusts by clipping them accordingly in (12).

IV. RESULTS

The proposed MPPI control algorithm was implemented and integrated into the multirotor Unmanned Aerial Vehicle control and estimation system developed by the Multi-robot Systems Group (MRS) at the Czech Technical University (CTU) [16], [17]. The experiments were conducted using a high-level computing unit Jetson Orin Nano 8 GB with 1.5 GHz 6-core Arm Cortex-A78AE CPU and 625 MHz 1024-core NVIDIA GPU. Thanks to the high number of GPU cores, we are able to roll out a large number of parallel simulations (as listed in Table II). In the simulated experiments, we used the multirotor dynamics simulation tools already available in the system [17]. For the real-world experiments, the computing unit was mounted on a quadrotor based on ReadytoSky Alien 7" frame with Emax Eco II 2807 1700 Kv motors equipped with 7" three-blade propellers. The UAV is equipped with Pixhawk Cube Orange+ with PX4 low-level flight controller and Holybro H-RTK F9P GNSS Series RTK GPS receiver for localization. We use the state estimation directly from the PX4 flight controller connected to the RTK GPS. The body rates and collective thrust control commands produced by the proposed method are sent directly to the PX4 flight controller that uses body-rate PID to track the commanded body rates. The reference trajectories are generated using an MPC-based linear trajectory tracker [17], which produces a smooth receding horizon trajectory that the proposed MPPI controller is tracking. For safety, additional parameters $\omega_{xy,\text{max}}$ and $\omega_{z,\text{max}}$ were introduced to limit the desired control applied to the system. The values used are listed in Table I. For an image of the quadrotor, we refer to Figure 1.

A. Speed of computation

The upper limit for one iteration is 10 ms per iteration, which is the update rate of body rate controllers in the MRS UAV system (100 Hz). Figure 3 shows how the number of rollouts K and the number of prediction steps N

TABLE I: Drone parameters and MPPI control limits.

UAV model		Control limits				
Parameter	Value	Parameter	Sim	Real		
m	[kg]	1.21	T_{min}	[N]	0.3	0.3
l	[m]	0.15	T_{max}	[N]	19.0	8.0
c_{tf}	[m]	0.012	$\omega_{xy,max}$	[rad s ⁻¹]	10.0	6.0
\mathbf{J}	[g m ²]	diag($\begin{bmatrix} 7.06 \\ 7.06 \\ 13.6 \end{bmatrix}$)	$\omega_{z,max}$	[rad s ⁻¹]	2.0	0.5

affect the iteration time. In the end, we found the values $K = 896$ and $N = 15$ to be a good trade-off between the number of trajectories and the prediction steps while making sure the computation finishes in the required time, even when sharing resources with other processes on the UAV.

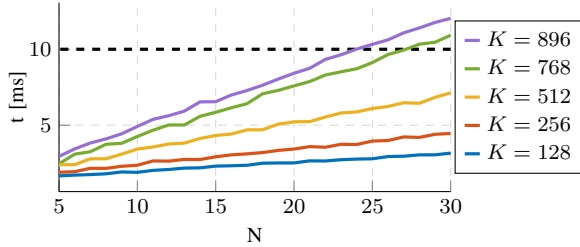


Fig. 3: Average time of one MPPI iteration (lines 3–18 in Alg. 1) depending on the number of rollouts K and number of prediction steps N . To run the controller at 100 Hz, we need to keep the iteration time below 10 ms (black dashed line). The computation time may still vary per iteration due to other processes running on the UAV.

B. Control input interpolation

To decouple the time step between two updates of the MPPI controller ($\Delta t = 10$ ms in our case) and the prediction time step, we employ an interpolation scheme. Decoupling the time steps allows the selection of the prediction horizon independently of the controller frequency. Otherwise, increasing the prediction horizon would be possible only by increasing N , however, that would increase the computational complexity, as shown in Figure 3. We increase the MPPI prediction time step to $n \cdot \Delta t$ and instead of shifting the nominal control by one each time (line 17 in Alg. 1), we interpolate linearly between them and run the optimization task every iteration. This allows the use of state feedback at 100 Hz, improving the performance of the controller. In our case, we use $n = 10$, resulting in MPPI having the prediction time step 0.1 s. With $N = 15$ prediction steps, the resulting prediction horizon becomes 1.5 s seconds.

C. Reference tracking

To design the cost function which assigns a real value to each rollout \mathbf{x}^k with inputs \mathbf{u}^k , we start with the standard weighing of input and input change

$$S_k = \sum_{j=0}^N \|\mathbf{u}_j^k\|_R^2 + \sum_{j=0}^{N-1} \|\Delta \mathbf{u}_j^k\|_{R_\Delta}^2, \quad (16)$$

where R and R_Δ are positive semidefinite cost matrices, $\Delta \mathbf{u}_j^k = \mathbf{u}_{j+1}^k - \mathbf{u}_j^k$ denotes the change of input, and $\|\cdot\|_P^2$ denotes the weighted Euclidean inner product $\|\mathbf{u}\|_P^2 = \mathbf{u}^T P \mathbf{u}$. The states \mathbf{x}_j^k are not penalized, since we

only expect to penalize high angular velocities, which can be done by weighing the desired body rates in \mathbf{u}_j^k .

To enforce tracking the reference, we introduce a term in the form

$$\sum_{j=0}^N \rho_{\text{ref}}(\mathbf{x}_j^k, \mathbf{x}_j^{\text{ref}}), \quad (17)$$

where ρ_{ref} is a metric (or at least an approximation thereof) on $\mathbb{R}^3 \times \mathbb{SO}(3) \times \mathbb{R}^3 \times \mathbb{R}^3$

$$\begin{aligned} \rho_{\text{ref}}(\mathbf{x}_j^k, \mathbf{x}_j^{\text{ref}}) = & \|\mathbf{p}_j^k - \mathbf{p}_j^{\text{ref}}\|_{c_p^{\text{ref}}}^2 + c_q^{\text{ref}} \cdot d_q(\mathbf{q}_j^k, \mathbf{q}_j^{\text{ref}})^2 \\ & + \|\mathbf{v}_j^k - \mathbf{v}_j^{\text{ref}}\|_{c_v^{\text{ref}}}^2 + \|\boldsymbol{\omega}_j^k - \boldsymbol{\omega}_j^{\text{ref}}\|_{c_\omega^{\text{ref}}}^2, \end{aligned} \quad (18)$$

with the weighted Euclidean metric used for \mathbb{R}^3 and weighing coefficients $c_p^{\text{ref}}, c_q^{\text{ref}}, c_v^{\text{ref}}$ and $c_\omega^{\text{ref}} \in \mathbb{R}$. The Euclidean norm is not a proper metric on $\mathbb{SO}(3)$ (quaternion \mathbf{q} representing the rotation) — this can be most prominently seen for quaternions \mathbf{q} and $-\mathbf{q}$, which represent the same rotation in \mathbb{R}^3 , but the Euclidean metric will yield a non-zero result. Therefore, we need to address this inconvenience when adding the reference tracking part to be able to handle distances between quaternions correctly. There exist two commonly used options for the function $d_q : \mathbb{SO}(3) \times \mathbb{SO}(3) \rightarrow \mathbb{R}$. The first one is computing the angle of rotation required to get from one orientation to the other, which is given by

$$\theta = \cos^{-1}(2\langle \mathbf{q}_1, \mathbf{q}_2 \rangle^2 - 1), \quad (19)$$

where $\langle \cdot, \cdot \rangle : \mathbb{SO}(3) \times \mathbb{SO}(3) \rightarrow \mathbb{R}$ denotes the quaternion inner product

$$\langle \mathbf{q}_1, \mathbf{q}_2 \rangle = w_1 w_2 + x_1 x_2 + y_1 y_2 + z_1 z_2. \quad (20)$$

However, the evaluation of this function is computationally demanding, and we use instead a common alternative, which roughly corresponds the exact angle (upto a multiplicative constant)

$$d_q(\mathbf{q}_1, \mathbf{q}_2) = 1 - \langle \mathbf{q}_1, \mathbf{q}_2 \rangle^2. \quad (21)$$

The advantage of both the exact angle and its approximation over the Euclidean metric is that it better respects the behavior of quaternions, mainly $d_q(\mathbf{q}, \mathbf{q}) = d_q(\mathbf{q}, -\mathbf{q}) = 0$. The comparison can be seen in Figure 4.

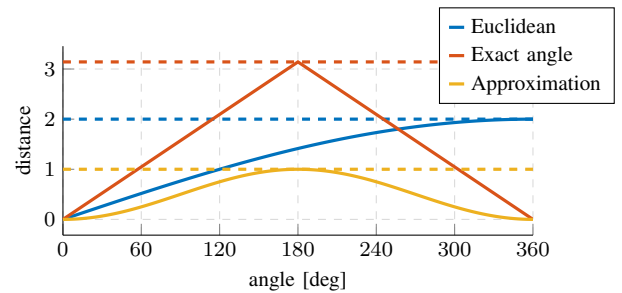


Fig. 4: Comparison of the Euclidean distance, exact angle (Eq. 19), and approximation thereof (Eq. 21). The object makes a full 360-degree rotation around a single axis, and the distance to the starting state is computed. The Euclidean distance yields a value of 2 when the object rotates fully, while the angle and its approximation correctly reach maximum at 180° and return to zero after the full rotation is completed.

We generate four reference trajectories (shown in Figure 5) and track them for 20 loops in simulation using the proposed MPPI, MPC [17], and SE(3) [15] controllers. The MPPI controller parameter values are listed in Table II, and the resulting tracking errors are presented in Table III (with visualizations in Figure 6). Our controller achieves lower tracking error (around 60% improvement) compared to the MPC controller in all cases and reaches results comparable to the SE(3) controller for the *line* and *eight* trajectories. The SE(3) controller being better than MPPI at the tracking task was expected, since it is specifically made for tracking a trajectory where high-order derivatives of desired states are available. However, if the derivatives were not available or if the task was more general than simple tracking (e.g., tracking while avoiding obstacles, as presented below), the SE(3) controller cannot be used, whereas MPPI can be used with only minor modifications.

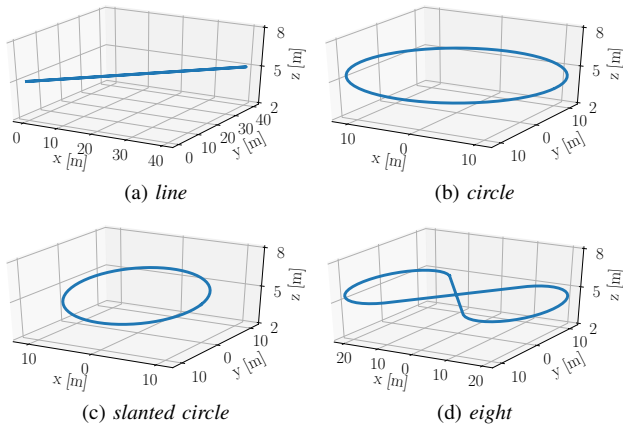


Fig. 5: Trajectories used for tracking error comparison.

TABLE II: Parameter values used in the MPPI controller

MPPI parameters					
Δt	100 ms	c_p^{ref}	400.0	Σ	diag(0.60, 0.15, 0.15, 0.05)
K	896	c_v^{ref}	40.0	R	diag(0.01, 0.05, 0.05, 0.10)
N	15	c_q^{ref}	20.0	R_Δ	diag(0.05, 0.10, 0.10, 0.30)
λ	10^{-4}	c_ω^{ref}	20.0		

After verifying our controller in the simulation, we used the same trajectories in real-world validation. We let our MPPI controller track the four trajectories, each for three laps. The controller is able to track the trajectories reliably, and the results are comparable to those in simulations, as shown in Table III and Figure 7. For the *eight* trajectory, the tracking error is increased by only 2.52% when compared to the simulations, which shows a great sim-to-real transferability. On the other hand, the tracking error on *line* is increased by 53.4%, mainly in the z-axis (Fig. 7a). We believe that this is caused by battery voltage drop at the turn points when under high load, changing the motors' thrust-to-throttle mapping, and could be in the future improved by using a method that accounts for the battery voltage, e.g., by using a voltage-based thrust map. In Figure 8, we visualize the drone position history in the video recording of tracking the *slanted circle* trajectory.

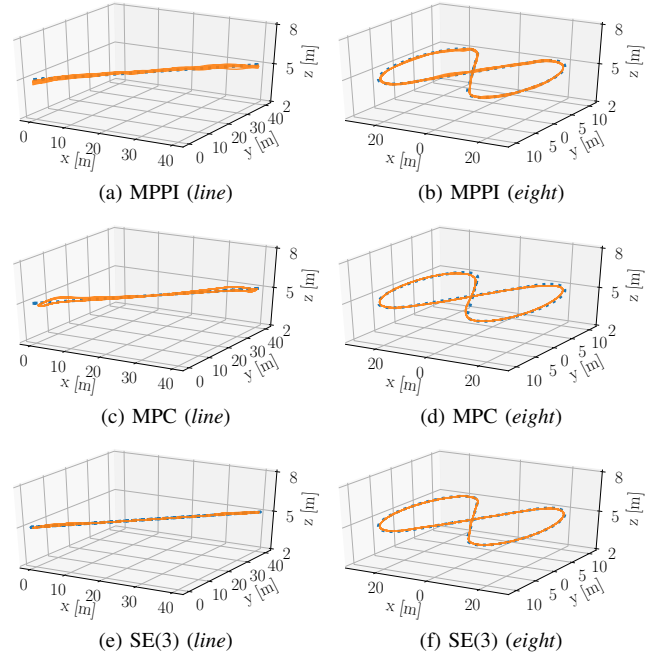


Fig. 6: Visualization of the drone trajectories tracked by the MPPI, MPC, and SE(3) controllers in the simulation.

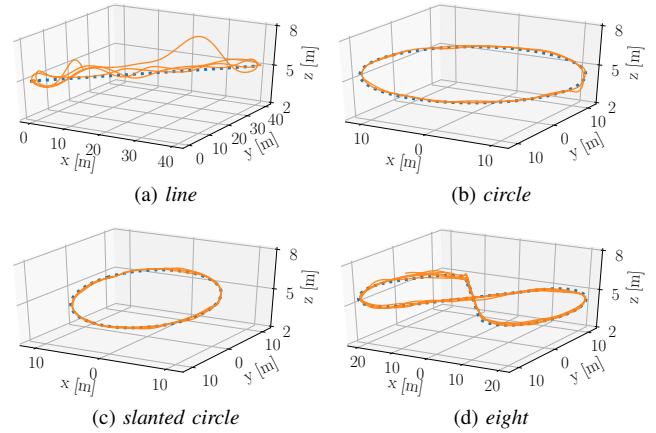


Fig. 7: Visualization of the drone trajectories tracked by the MPPI controller in real world.



Fig. 8: Drone tracked in the video recording of Fig. 7c.

TABLE III: Tracking error in simulation (left) and realworld (right) experiments

Trajectory	$\ \mathbf{v}\ _{\max}$ [m s ⁻¹]	$\ \mathbf{a}\ _{\max}$ [m s ⁻²]	Error [m] (<i>simulation</i>)			Error [m] (<i>realworld</i>)
			MPPI	MPC	SE(3)	MPPI
<i>eight</i>	8.853	7.571	0.633 ± 0.123	1.668 ± 0.396	0.430 ± 0.132	0.649 ± 0.356
<i>slanted circle</i>	5.652	5.289	0.412 ± 0.068	0.921 ± 0.163	0.061 ± 0.054	0.518 ± 0.130
<i>line</i>	12.271	19.782	0.691 ± 0.310	1.704 ± 0.617	0.500 ± 0.198	1.060 ± 1.148
<i>circle</i>	8.459	7.034	0.646 ± 0.127	1.591 ± 0.202	0.114 ± 0.083	0.804 ± 0.125

D. Obstacles

One of the key advantages of the MPPI method is the ability to include obstacles in the cost function and force the controller to avoid them, which is impossible for the other approaches presented in Section IV-C. This can be achieved by adding a term to the cost function that heavily penalizes states where the drone collides with the environment

$$S'_k = S_k + \sum_{j=0}^N c_{\text{obs}} \cdot \mathbb{1}_{\mathbf{x}_j^k \in \mathcal{C}_{\text{obs}}}, \quad (22)$$

where S_k is the original cost of the rollout, c_{obs} is the cost coefficient for the collision term ($c_{\text{obs}} = 10^6$ in our case), and $\mathbb{1}_{\mathbf{x}_j^k \in \mathcal{C}_{\text{obs}}}$ is an indicator function, which is 1 if the drone at state \mathbf{x}_j^k is in collision with the environment (i.e., \mathbf{x}_j^k lies in the obstacle region \mathcal{C}_{obs} determined by a collision detection module). Collision detection can be implemented in multiple ways (e.g., using mesh collision detection such as *Rapid* [28] or discrete grid-based methods, such as *OctoMap* [29]), and the exact implementation depends on the platform capabilities (e.g., Light Detection and Ranging (LiDAR)).

First, we validate the ability to avoid obstacles in an environment with a single cylindrical pillar, shown in Figure 9. Further, we add virtual obstacles to two of the tracked trajectories (*line* and *slanted circle*) and task the controller to track the trajectories 30 times in a row in the simulation, keeping all of the parameters as in Section IV-C. In Figure 10 we show that the MPPI controller is able to handle and avoid the obstacles, successfully finishing the tracking task without any collisions. It is even able to consistently cope with non-convex obstacles and situations where the drone needs to retreat after being led to a dead-end, as is the case near the leftmost obstacle in Fig. 10a. Even though this behavior could be avoided by increasing the prediction horizon, it would still appear for obstacles larger than the rollout distance. Therefore, leaving the dead-end detection task to a high-level planning algorithm providing the reference states $\mathbf{x}_j^{\text{ref}}$ might prove more suitable.

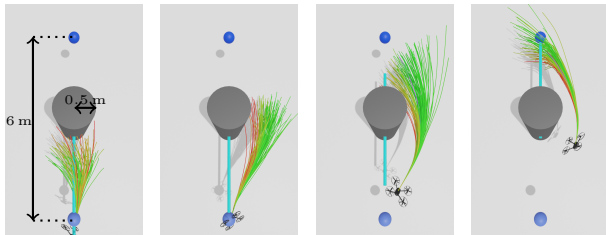


Fig. 9: Tracking a straight line moving between the blue spheres at a constant speed of 6 m s⁻¹ (cyan), while avoiding an obstacle (gray cylinder). The colliding rollouts are heavily penalized (red curves) in comparison to the non-colliding ones (green).

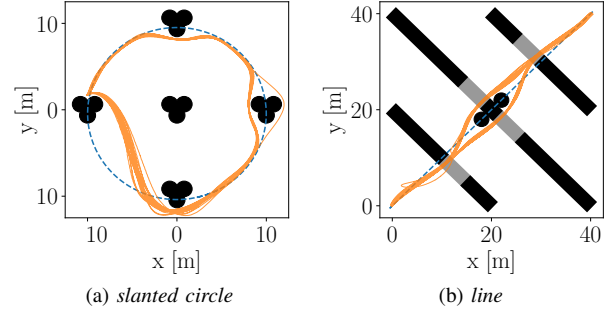


Fig. 10: We test our MPPI algorithm on *slanted circle* (Fig. 5c) with pillar-like obstacles (black) and on *line* (Fig. 5a) with walls (black, windows in gray) and let the drone loop the desired trajectory (blue curve) 30 times. The recorded UAV positions are shown in orange.

E. Discussion and future work

Future work extending the proposed method includes using the so-called *Delay compensation by prediction* [30]. The time of the computation of each iteration is consistent thanks to a fixed number of computations performed per iteration. Therefore, instead of addressing the problem starting at the current state, we can forward-simulate the system and start the optimization at a state at which the system will be when we would have computed the optimization iteration.

The selection of the noise parameters Σ could also be extended. For now, the values were hand-tuned for fast flying and would not be suitable for, e.g., hovering. Introducing a method to dynamically set the disturbance noise based on the reference variance could be more suitable, adjusting the noise variance during the flight.

Most notably, extensive real-world testing of avoiding (even dynamic) obstacles is planned. For that, known obstacles need to be placed in the testing arena and inserted into the collision checking module, or the platform needs to be extended to include a sensor able to detect obstacles.

V. CONCLUSION

In this work, we presented a model predictive planning and control methodology called Model Predictive Path Integral (MPPI) designed to control a drone along a given trajectory, while allowing the use of arbitrary cost functions and constraints. We proposed a method that respects low-level constraints imposed by the motor thrust limits of the UAV, while using a nonlinear model of the system dynamics to predict the future dynamic states. Moreover, we discussed the shortcomings of the traditional Euclidean metric used for tracking a reference trajectory considering the rotation part of the state and proposed a more suitable approach.

Validation of the complete approach was conducted using a UAV both in simulation and in outdoor flight using a

real drone. We showed that our proposed and implemented MPPI controller is able to control a drone at the level of commanded angular velocities and collective thrust while running all the needed computations at 100 Hz online and onboard. To the best of our knowledge, this is the first deployment of MPPI on a real drone without relying on external hardware, such as an offboard PC with a large GPU.

We compared our controller with two other controllers used for reference tracking, namely MPC and SE(3). The results show that the proposed method consistently outperforms the original MPC algorithm, achieving around 60% better tracking error in all test scenarios. Although unable to reach the performance of the SE(3) controller, it has the ability to include obstacles in the control task, which is impossible in the SE(3) controller and greatly limited in the existing MPC controllers. The ability of the MPPI controller to avoid non-convex obstacles by using a general collision detection module is demonstrated in two experimental setups, tracking a reference at up to 44 km h^{-1} and acceleration close to 20 m s^{-2} while consistently avoiding obstacles. This proves the biggest strength of the proposed method, which lies in the ability to handle non-convex and even non-differentiable cost functions, therefore enabling the general collision avoidance mentioned.

REFERENCES

- [1] T. P. Nascimento and M. Saska, "Position and attitude control of multi-rotor aerial vehicles: A survey," *Annual Reviews in Control*, pp. 129–146, 2019.
- [2] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, Jul. 2021.
- [3] 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.
- [4] E. Jeong, J. Seo, and J. Wacker, "Literature review and technical survey on bridge inspection using unmanned aerial vehicles," *Journal of Performance of Constructed Facilities*, vol. 34, no. 6, p. 04020113, 2020.
- [5] A. Ollero, A. Suarez, C. Papaioannidis, I. Pitas, J. M. Marredo, V. Duong, E. Ebeid, V. Kratky, M. Saska, C. Hanoune, A. Afifi, A. Franchi, C. Vourtsis, D. Floreano, G. Vasiljevic, S. Bogdan, A. Caballero, F. Ruggiero, V. Lippiello, C. Matilla, G. Cioffi, D. Scaramuzza, J. R. M. de Dios, B. C. Arrue, C. Martin, K. Zurad, C. Gaitan, J. Rodriguez, A. Munoz, and A. Viguria, "Aerial-core: Ai-powered aerial robots for inspection and maintenance of electrical power infrastructures," 2024.
- [6] D. Datsko, F. Nekovar, R. Penicka, and M. Saska, "Energy-aware multi-uav coverage mission planning with optimal speed of flight," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2893–2900, 2024.
- [7] P. Petracek, V. Kratky, T. Baca, M. Petrlík, and M. Saska, "New Era in Cultural Heritage Preservation: Cooperative Aerial Autonomy for Fast Digitalization of Difficult-to-Access Interiors of Historical Monuments," *IEEE Robotics and Automation Magazine*, pp. 2–19, February 2023.
- [8] H. Nguyen, M. Kamel, K. Alexis, and R. Siegwart, "Model predictive control for micro aerial vehicles: A survey," in *2021 European Control Conference (ECC)*, 2021, pp. 1556–1563.
- [9] L. F. Recalde, B. S. Guevara, C. P. Carvajal, V. H. Andaluz, J. Varela-Aldás, and D. C. Gandolfo, "System identification and nonlinear model predictive control with collision avoidance applied in hexacopters uavs," *Sensors*, vol. 22, no. 13, 2022.
- [10] B. Lindqvist, S. S. Mansouri, A.-a. Agha-mohammadi, and G. Nikolakopoulos, "Nonlinear mpc for collision avoidance and control of uavs with dynamic obstacles," *IEEE RA-L*, vol. 5, no. 4, pp. 6001–6008, 2020.
- [11] E. Small, P. Sotasakis, E. Fresk, P. Patrinos, and G. Nikolakopoulos, "Aerial navigation in obstructed environments with embedded nonlinear model predictive control," in *ECC*, 2019, pp. 3556–3563.
- [12] G. Garimella, M. Sheckells, and M. Kobilarov, "Robust obstacle avoidance for aerial platforms using adaptive model predictive control," in *IEEE ICRA*, 2017, pp. 5876–5882.
- [13] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 International Conference on Robotics and Automation (ICRA)*. Stockholm: IEEE, May 2016, pp. 1433–1440.
- [14] —, "Information-Theoretic Model Predictive Control: Theory and Applications to Autonomous Driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, Dec. 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8558663/>
- [15] T. Lee, M. Leok, and N. H. McClamroch, "Nonlinear robust tracking control of a quadrotor uav on se(3)," *Asian Journal of Control*, vol. 15, no. 2, pp. 391–408, 2013.
- [16] D. Hert, T. Baca, P. Petracek, V. Kratky, R. Penicka, V. Spurny, M. Petrlík, M. Vrba, D. Zaitlík, P. Stoudek, V. Walter, P. Stepan, J. Horyna, V. Pritzl, M. Sramek, A. Ahmad, G. Silano, D. B. Licea, P. Stibinger, T. P. Nascimento, and M. Saska, "Mrs drone: A modular platform for real-world deployment of aerial multi-robot systems," *Journal of Intelligent & Robotic Systems*, vol. 108, pp. 1–34, 2023.
- [17] T. Baca, M. Petrlík, M. Vrba, V. Spurny, R. Penicka, D. Hert, and M. Saska, "The MRS UAV System: Pushing the Frontiers of Reproducible Research, Real-world Deployment, and Education with Autonomous Unmanned Aerial Vehicles," *Journal of Intelligent & Robotic Systems*, vol. 102, no. 1, p. 26, May 2021.
- [18] D. Mellinger, N. Michael, and V. Kumar, "Trajectory generation and control for precise aggressive maneuvers with quadrotors," *The International Journal of Robotics Research*, vol. 31, no. 5, pp. 664–674, 2012.
- [19] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 620–626, 2017.
- [20] 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.
- [21] 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.
- [22] S. Sun, A. Romero, P. Foehn, E. Kaufmann, and D. Scaramuzza, "A comparative study of nonlinear mpc and differential-flatness-based control for quadrotor agile flight," *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3357–3373, 2022.
- [23] M. Dharmadhikari, T. Dang, L. Solanka, J. Loje, H. Nguyen, N. Khedekar, and K. Alexis, "Motion primitives-based path planning for fast and agile exploration using aerial robots," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 179–185.
- [24] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [25] J. Pravitra, K. A. Ackerman, C. Cao, N. Hovakimyan, and E. A. Theodorou, "L1-adaptive mppi architecture for robust and agile control of multirotors," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 7661–7666.
- [26] J. Higgins, N. Mohammad, and N. Bezzo, "A model predictive path integral method for fast, proactive, and uncertainty-aware uav planning in cluttered environments," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 830–837.
- [27] M. W. Mueller, M. Hehn, and R. D'Andrea, "A computationally efficient motion primitive for quadcopter trajectory generation," *IEEE Transactions on Robotics*, vol. 31, no. 6, pp. 1294–1310, 2015.
- [28] S. Gottschalk, M. Lin, and D. Manocha, "Obbtrees: A hierarchical structure for rapid interference detection," *Computer Graphics*, vol. 30, 10 1997.
- [29] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013.
- [30] S. Gros and M. Diehl, "Numerical optimal control," 04 2022.