

Real-time Motion Planning for Aerial Videography with Dynamic Obstacle Avoidance and Viewpoint Optimization

Tobias Nageli¹, Javier Alonso-Mora^{2,3}, Alexander Domahidi⁴, Daniela Rus², Otmar Hilliges¹

Index Terms—Aerial Robotics, Motion and Path Planning, Collision Avoidance.

Abstract— We propose a method for real-time motion planning with applications in aerial videography. Taking framing objectives, such as position of targets in the image plane as input, our method solves for robot trajectories and gimbal controls automatically and adapts plans in real-time due to changes in the environment. We contribute a real-time receding horizon planner that autonomously records scenes with moving targets, while optimizing for visibility to targets and ensuring collision-free trajectories. A modular cost function, based on the re-projection error of targets is proposed that allows for flexibility and artistic freedom and is well behaved under numerical optimization. We formulate the minimization problem under constraints as a finite horizon optimal control problem that fulfills aesthetic objectives, adheres to non-linear model constraints of the filming robot and collision constraints with static and dynamic obstacles and can be solved in real-time. We demonstrate the robustness and efficiency of the method with a number of challenging shots filmed in dynamic environments including those with moving obstacles and shots with multiple targets to be filmed simultaneously.

I. INTRODUCTION

ROBOTICS and micro-aerial vehicles in particular are rapidly becoming an end-user facing technology. In particular, the application domain of filming with aerial vehicles currently receives great interest from industry and consumers. Professional camera teams leverage consumer-grade robots to create stunning visuals that previously required a helicopter and expensive camera gear. However, manually flying quadcopters remains a surprisingly hard task. Furthermore, automated flight modes in current commercial offerings are restricted to simple circling or target following.

Several algorithms have been proposed for the planning of quadcopter trajectories [1, 2, 3], taking both aesthetic objectives and the physical limits of the robot into consideration.

Manuscript received: September, 10, 2016; Revised: December, 5, 2016; Accepted: January, 21, 2017.

This paper was recommended for publication by Editor Jonathan Roberts upon evaluation of the Associate Editor and Reviewers' comments. This work was supported in part by Microsoft Research, SMARTS N00014-09-1051, pDOT ONR N00014-12-1-1000 and the Toyota Research Institute. We are grateful for their support. We also thank Parrot for hardware and tech support.

¹AIT Lab, Department of Computer Science, ETH Zurich, 8092 Zurich, Switzerland {naegelit|otmarh}@inf.ethz.ch

²CSAIL, Massachusetts Institute of Technology, 02139 Cambridge, MA, USA rus@csail.mit.edu

³DCSC, Delft Technical University, 2628CD Delft, Netherlands J.AlonsoMora@tudelft.nl

⁴embotech GmbH, 8092 Zurich, Switzerland domahidi@embotech.com

Digital Object Identifier (DOI): see top of this page.



Fig. 1: We propose a real-time method to film shots of dynamic scenes automatically. The user specifies target positions on screen. We solve for robot trajectories and gimbal controls, producing footage that adheres to aesthetic constraints, minimizes visual occlusion and produces collision-free trajectories in dynamic scenes.

These methods make the task of trajectory planning easier for non-experts. However, current approaches employ global optimization techniques, planning the entire trajectory a priori. While [2] allows to adjust the velocity along the planned trajectory at execution time, none of the existing methods are capable of re-planning a suitable trajectory in real-time, for example to avoid collisions with dynamic obstacles or to film targets that move in unpredictable ways (e.g., human actors).

In this paper we propose an approach to the automatic generation of quadcopter and gimbal controls in real-time while ensuring physical feasibility. In particular, we contribute a fast receding horizon planner based on numerical optimization for automatic aerial cinematography. The method takes high-level aesthetic objectives given by a cameraman as input and automatically records the scene while the targets move in an a priori unknown way. The system fulfils the high-level framing objective, in the least squares sense, while ensuring collision-free paths. The resulting appearance of targets is specified via set-points in screen space (extending ideas outlined in [4]) and the method minimizes the re-projection error alongside the viewing direction and scale of the target projections in real-time.

We formulate this cost minimization problem under constraints as a finite horizon model predictive control (MPC) non-linear program with the following properties: (i) a main

utility function to fulfill - as close as possible - the specifications from the user in a dynamic scene; (ii) input and state constraints to respect the dynamics and model constraints of the aerial vehicle; (iii) constraints for collision avoidance with obstacles and the targets being recorded. The resulting optimization problem can be solved numerically with state-of-the-art solvers in real-time. The inputs computed are applied for the first time step, and the optimization is repeated at the next sampling instance with updated information about the quadrotor state, obstacles and target positions. We believe that the method is general enough and could be adapted to other tasks for aerial vehicles which include collision avoidance with respect to moving obstacles.

II. RELATED WORK

Quadrotor trajectory generation is a well studied problem and various approaches have been proposed, including early work on MPC for collision avoidance applied to aerial vehicles [5, 6], global forward planning approaches to generate minimum snap trajectories [7], for generating collision-free trajectories via convex optimization [8] or for state interception with a quadrotor [9] based on a convex MPC formulation. Outside of the field of aerial vehicles, constrained optimization methods have also been proven to be powerful tools for trajectory generation [10]. We also formulate our problem in the MPC framework to attain real-time performance, using the full non-linear dynamics of the quadrotor and extending trajectory generation to include collision avoidance with dynamic targets and to respect cinematographic objectives.

Automatic camera control has been studied extensively in the context of virtual environments in computer graphics. We refer to the comprehensive review in [11], with the majority of methods using discrete optimization formulations. Notably in virtual environments it has become somewhat of a standard to define viewing constraints in screen-space [4, 12, 13] However, virtual environments are not limited by real-world physics and robot constraints and hence can produce arbitrary camera trajectories, velocities and viewpoints.

Extending the work on trajectory generation [7], several algorithms for planning aerial video shots of (mostly) static scenes [1, 2, 3] exist. Joubert et al.'s method [2] allows an operator to adjust the velocity of the MAV at execution time to keep moving targets in-frame. However, the method cannot re-plan the trajectory in cases where the target does deviate from an a-priori known path or in cases where dynamic obstacles temporarily obstruct the original trajectory. Lino et al. [14] propose a method to generate camera paths in virtual environments that ensure visibility of the faces of two actors in the image. This concept was later extended to quadrotors [15], albeit limited to obstacle free environments. To record a single moving target with multiple quadrotors, [16] introduced a particle swarm method for formation control. The method is again limited to obstacle-free environments and only tested in simulation. We propose a general constrained optimization approach for automatic viewpoint and trajectory computation allowing for multiple moving targets and obstacles in the environment. The approach can be solved in real-time.



Fig. 2: Illustration of cinematographic framing constraints: size, viewing angle and position on screen (from left to right).

Our method also accounts for occlusions and loosely relates to the fields of target tracking [17], visual servoing [18], active vision [19] and persistent monitoring [20], where our focus is on videographing a set of moving targets.

III. PRELIMINARIES

A. Cinematographic Objectives

While film making is a form of art and relies on human creativity and intuition, many aspects of it have been studied and categorized forming a ‘grammar’ of film [21]. Fig. 2 summarizes the most important aspects. We later formalize these mathematically for use in a cost minimization algorithm. In particular, we are interested in framing objectives – that is formal rules that specify how objects should appear on the screen. The first important notion is that of distance. Shots can be categorized into five types of shots (close-up, close-medium, full and long shots), see Fig. 2, left. A further important aspect is that of the relative viewing angle, which can also be categorized into ranges of pan- and tilt-angles), see Fig. 2, middle. Finally, the screen position of a filmed target is important in order to create aesthetically pleasing footage. In particular the rule of thirds (cf. [22]) prescribes the placement of significant vertical and horizontal elements along the horizontal and vertical thirds (Fig. 2, right).

1) *Shot Framing Requirements*: From the above film grammar we can extract objectives for an optimization method. In particular, we consider (i) the *screen position* of targets, where we seek to minimize the distance between the desired position on screen and the projection of the actual (3D) target position onto the screen. (ii) To relate shot distances to positions, we seek to optimize the *projected size* of a given target. (iii) Furthermore, we require the algorithm to consider the *relative viewing angle* between target and camera center, for example to keep the face of a person in view. (iv) Finally, we require the algorithm to account for an externally set *camera pose*, which could be specified directly by the user or a high-level global planning algorithm.

B. Notation and Coordinate Frames

We denote the set of positive definite matrices of size n by \mathbb{S}_{++}^n and the set of positive semi-definite matrices of size n by \mathbb{S}_{+}^n . Given a vector $\mathbf{x} \in \mathbb{R}^n$, we define the square of the norm as $\|\mathbf{x}\|_P^2 \triangleq \mathbf{x}^T P \mathbf{x}$ for $P \in \mathbb{S}_{++}^n$. Vectors are denoted in bold. If not indicated differently, vectors are expressed in the standard Earth North Up (ENU) coordinate system. We denote points in 3D as \mathbf{p} with a name as subscript (e.g., \mathbf{p}_h for \mathbf{p}_{head}). A relative vector between two points \mathbf{p}_a and \mathbf{p}_b is denoted as

\mathbf{r}_{ab} . A superscript \mathbf{r}_{ab}^c indicates the vector \mathbf{r}_{ab} is expressed in frame C (without subscript the vector is expressed in the inertial frame I). Pixel coordinates are denoted as μ . Desired setpoints are indicated with a subscript $(\cdot)_d$.

Rotation matrices performing rotations from frame A to frame B are denoted by $\mathbf{R}_{BA} = \mathbf{R}(\bar{q}_{BA}) \in SO(3)$, where \bar{q}_{BA} is the corresponding quaternion. We adhere to the JPL quaternion definition [23] and denote a quaternion by $\bar{q} = [q_x \mathbf{i} + q_y \mathbf{j} + q_z \mathbf{k} + q_w] = [\mathbf{q}, q_w]^T$. Quaternion multiplication is denoted by \otimes . Orientation errors are described in $so(3)$, the tangent space of $SO(3)$, and are written as $\delta\theta$.

C. Robot and Camera Model

Our method is agnostic to the quadrotor hardware. However, in our experimental setup we use a Parrot Bebop 2 and therefore we will use its model in our derivations. In general, the pose of a quadrotor can be described by its position $\mathbf{p}_q \in \mathbb{R}^3$ and orientation $\bar{q}_q \in SO(3)$. While the camera is not directly attached to the center of the quadrotor, experimentally we have found it to be sufficient to assume the camera position \mathbf{p}_c and the quadrotor's position \mathbf{p}_q to be identical. In the following consider $\mathbf{p}_c = \mathbf{p}_q$. The set of feasible states is denoted by \mathcal{X} . The state of the system is defined by

$$\mathbf{x} = [\mathbf{p}_c, v_x, v_y, \bar{q}_q]^T \in \mathbb{R}^5 \times SO(3), \quad (1)$$

where v_x, v_y are the velocities of the quadrotor along the x and y axis.

The inputs accepted by Parrot's SDK are roll angle ϕ_q , pitch angle θ_q , angular speed around the body- z axis ω_{q_z} and velocity in the body- z axis $v_z = \dot{p}_{q_z}$. The camera is attached to the robot via a pan-tilt gimbal (in case of the Bebop this is a software gimbal) where in this paper only the tilt part is used. Due to pre-stabilized gimbal dynamics, it can be modeled using a separate pitch angle θ_g . The set of feasible inputs is denoted by \mathcal{U} . The inputs to the system are

$$\mathbf{u} = [\phi_q, \theta_q, \omega_{q_z}, v_z, \theta_g] \in \mathbb{R}^5. \quad (2)$$

The sets \mathcal{X} and \mathcal{U} are given by the physical limits of the environment (e.g. $p_z > 0$) and by the internal constraints of the Parrot Bebop 2 (e.g. maximal vertical and horizontal velocities, maximal roll pitch angle). The limits are described in the documentation of the Parrot SDK¹.

The full non-linear model and its states are not exposed via the SDK, hence we use the following approximation for the dynamics. The system model of the flying camera is given by non-linear ordinary differential equations for the acceleration in the x and y axis, $\ddot{\mathbf{p}}_{c_{x,y}}$, and the velocity in body- z \dot{p}_{c_z} .

$$\ddot{\mathbf{p}}_{c_{x,y}} = \mathbf{R}_\psi(\bar{q}_q) \begin{bmatrix} -\tan(\theta_q) \\ \tan(\phi_q) \end{bmatrix} g \quad \text{and} \quad \dot{p}_{c_z} = v_z, \quad (3)$$

where $g = 9.81 \frac{m}{s^2}$ is the earth's gravity and $\mathbf{R}_\psi(\bar{q}_q) \in SO(2)$ is the rotation matrix of the quadrotor's yaw only. The camera's attitude is given by the quaternion

$$\dot{\bar{q}}_c = \bar{q}_g(\theta_g) \otimes (\bar{q}_q \otimes \bar{q}(\omega_{q_z})), \quad (4)$$




Fig. 3: Coordinate frames and physical quantities used in our method, here shown on the example of a human actor. The pose of the flying camera is given by its position \mathbf{p}_c , its orientation \bar{q}_c and the orientation of the gimbal \bar{q}_g . The target is modeled by an ellipsoid S_{Ω_s} centered at \mathbf{p}_t . The flying camera avoids collisions with this volume. In this case the objective is to record the head of the target, modeled as an oriented sphere with direction given by ψ , size h_h and position \mathbf{p}_h . The projection of the head onto the camera image space has size σ and is centered at μ_x, μ_y .

where $\bar{q}_g(\theta_g)$ is the quaternion defined by the gimbal pitch angle θ_g and $\bar{q}(\omega_{q_z})$ is defined as

$$\bar{q}(\omega_{q_z}) = \frac{1}{2\epsilon} \begin{bmatrix} 0 \\ 0 \\ \omega_{q_z} \sin(\epsilon) \\ 2\cos(\epsilon) \end{bmatrix} \quad \text{with} \quad \epsilon = \|\omega_{q_z}\|. \quad (5)$$

In the remainder of the paper the continuous non-linear quadrotor model $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ is discretized using a standard forward Euler approach: $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k)$ [24].

D. Target Model

We denote the position of each actor, or moving target, by $\mathbf{p}_t \in \mathbb{R}^3$, with K the number of targets. We then assume that the human motion is based on a constant velocity model $\dot{\mathbf{p}}_t = \mathbf{q}_\eta$ where $\mathbf{q}_\eta \sim \mathcal{N}(0, \sigma_t)$ is Gaussian noise. Given the current observations, we predict future positions of the actor with a standard linear Kalman filter [24].

IV. TRAJECTORY GENERATION FOR VIEWPOINT OPTIMIZATION

The goal of our algorithm is to find a feasible and locally optimal trajectory for the quadrotor and gimbal control inputs in real-time. The method produces imagery that minimizes the deviation with respect to the high-level image space specifications of Sec. III-A1 and generates trajectories that guarantee collision-free motion with respect to dynamic obstacles.

A scene is formed by K targets to be filmed, e.g. the faces of actors as seen in Fig. 3. Each target $i \in \{1, \dots, K\}$ is modeled as a sphere at position $\mathbf{p}_h \in \mathbb{R}^3$, with orientation ψ and of diameter h_h . Note that for clarity we typically omit

¹<http://developer.parrot.com/>

the additional subscript i unless explicitly dealing with several targets. Without loss of generality, in the following we only consider the orientation ψ around the z-axis, to model the gaze direction of the face (see Fig. 3). For collision avoidance and visibility, we consider that each target (not only the face) occupies a physical space defined by an ellipsoid S_{Ω_s} centered at $\mathbf{p}_t = \mathbf{p}_h + (\frac{h_h}{2} - \frac{h}{2})\mathbf{e}_z$, where h denotes the physical height of the target from the ground and \mathbf{e}_z the earth's up vector, illustrated in Fig. 3 for the case of a human actor.

A. Method Overview

We propose a receding horizon MPC formulation for trajectory generation given by the following variables, cost terms and constraints. Denote by Δ_t the time step and by N the time horizon of the MPC problem. At each time step a local trajectory of duration $N\Delta_t$ is computed. The first input is then applied. The optimization problem is re-solved at every sampling instance, leading to closed loop performance. This makes the approach robust against model uncertainties and unpredictable disturbances. At the next timestep $k+1$ we use the estimated states $\hat{\mathbf{x}}_0$ of the quadrotor and all targets as the initial state \mathbf{x}_0 and use the trajectory from timestep k as an initialization for the solver.

1) *Variables*: The optimization variables are the following.

(a) The states $\mathbf{x}_{0:N}$, which include the initial state \mathbf{x}_0 , the state trajectory $\{\mathbf{x}_1, \dots, \mathbf{x}_{N-1}\}$ and the final state \mathbf{x}_N . (b) The control inputs $\mathbf{u}_{0:N-1} = \{\mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$. And (c) the additional slack variables $\mathbf{s}_{0:N} = \{\mathbf{s}_0, \dots, \mathbf{s}_N\}$, one for each collision constraint and described in the forthcoming Sec. IV-D. The set of optimization variables is denoted by

$$\mathbf{z} = \{\mathbf{x}_{0:N}, \mathbf{u}_{0:N-1}, \mathbf{s}_{0:N}\}. \quad (6)$$

2) *Cost terms*: The final quantity we seek to minimize is the weighted sum of cost terms consisting of image framing, visibility and pose costs over all N stages of the planned trajectory.

3) *Constraints*: We include constraints to avoid collisions with moving obstacles and moving targets and to respect the dynamic model of the quadrotor introduced in Sec. III-C.

B. Viewpoint Optimization Problem

At the core of our algorithm lie a number of cost terms $c_{\text{image}}(\mathbf{x})$, $c_{\text{angle}}(\mathbf{x})$ and $c_{\text{size}}(\mathbf{x})$ which describe the deviation from the desired position of each target's projection on the image plane $\mu_{(x,y)_d}$, from the viewing angle \mathbf{a}_d and from the projection size σ_d . Each of these cost terms is computed for each state $k \leq N$ of the planning horizon and for each target $i \leq K$. For simplicity of exposition, the derivation of the costs is described for a general target. We also define the vector \mathbf{r}_{ch} which is the relative vector between the target and the camera as well as the rotation into the camera frame of \mathbf{r}_{ch} which is denoted as \mathbf{r}_{ch}^c :

$$\mathbf{r}_{ch} = \mathbf{p}_h - \mathbf{p}_c \quad \text{and} \quad \mathbf{r}_{ch}^c = R(\bar{\mathbf{q}}_c)\mathbf{r}_{ch},$$

1) *Image Space Positions*: Given the desired position of each target's projection on the image plane $\mu_{(x,y)_d}$, the pixel coordinates \mathbf{m}_d are computed via the camera intrinsics with focal point $C_{x,y}$ and focal length $\mathbf{f} = [f_x, f_y]$:

$$\mathbf{m}_d = \begin{bmatrix} (\mu_{xd} - C_x) f_x \\ (\mu_{yd} - C_y) f_y \end{bmatrix} \quad \text{with} \quad \mu_{(x,y)_d} \in \{0, 2C_{(x,y)}\}.$$

Consider the vector $\mathbf{r}_d^c = [\mathbf{m}_d, 1]^T \in \mathbb{R}^3$ pointing from the camera center through the desired pixel location \mathbf{m}_d in the image. We compute the quadratic image space location cost $c_{\text{image}}(\mathbf{x})$ using the residual given by the difference between the ray \mathbf{r}_{ch}^c from the camera to the target and the desired direction \mathbf{r}_d^c ,

$$c_{\text{image}}(\mathbf{x}) = \|\rho_m\|_{Q_m} \quad \text{with} \quad \rho_m = \frac{\mathbf{r}_{ch}^c}{\|\mathbf{r}_{ch}^c\|} - \frac{\mathbf{r}_d^c}{\|\mathbf{r}_d^c\|} \quad (7)$$

2) *Projection Size*: The size σ of the target in the image plane is computed by projecting the 3D sphere of diameter h_h to a circle in the 2D image plane. The target-size cost $c_{\text{size}}(\mathbf{x})$ is computed using the difference between the projected σ and the desired size σ_d .

$$c_{\text{size}}(\mathbf{x}) = \|\sigma - \sigma_d\|_{Q_\sigma} \quad \text{with} \quad \sigma = \frac{h_h \|\mathbf{f}\|}{\|\mathbf{r}_{ch}^c\|}. \quad (8)$$

3) *Relative Viewing Angle*: Given the desired viewing angles θ_d and ψ_d , and the current orientation ψ of the target, we define the desired viewing orientation \mathbf{a}_d between the center of the camera and the target:

$$\mathbf{a}_d = [\sin \theta_d \cos(\psi_d + \psi), \sin \theta_d \sin(\psi_d + \psi), \cos \theta_d]^T$$

Because the angle ψ_d is relative to the target we add the current rotation of the target ψ to obtain the setpoint. The view angle cost $c_{\text{angle}}(\mathbf{x})$ is computed using the error between the desired (relative to the target's normal) and the current viewing angle:

$$c_{\text{angle}}(\mathbf{x}) = \|\rho_a\|_{Q_a} \quad \text{with} \quad \rho_a = -\frac{\mathbf{r}_{ch}^c}{\|\mathbf{r}_{ch}^c\|} - \frac{\mathbf{a}_d}{\|\mathbf{a}_d\|}. \quad (9)$$

4) *Camera Pose*: In addition to local viewpoint optimization we can also consider the residual in the camera's position and orientation dynamics - although this is not a required cost. Here we assume that either the user or a global path planning algorithm can specify a desired pose \mathbf{p}_{cd} and $\bar{\mathbf{q}}_{qd}$ for the flying camera, which is compared to the actual pose of the camera given by \mathbf{p}_c and $\bar{\mathbf{q}}_c$. The position residual is given by the euclidean distance. The orientation residual ρ_θ is computed building the error quaternion and projecting it into the tangent space $so(3)$ of the $SO(3)$ manifold:

$$\rho_\theta = \delta\theta = [\bar{\mathbf{q}}_{e_x} \quad \bar{\mathbf{q}}_{e_y} \quad \bar{\mathbf{q}}_{e_z}] \in so(3) \quad \text{with} \quad \bar{\mathbf{q}}_e = \bar{\mathbf{q}}_c \otimes \bar{\mathbf{q}}_{cd}^{-1}.$$

This yields the camera pose cost $c_{\text{pose}}(\mathbf{x})$:

$$c_{\text{pose}}(\mathbf{x}) = \|\mathbf{p}_{cd} - \mathbf{p}_c\|_{Q_p} + \|\rho_\theta\|_{Q_\theta}. \quad (10)$$




Fig. 4: Schematic illustration of occlusion minimization based on viewpoint dependent horizon culling.

C. Occlusion Minimization

To allow for scenes containing objects that can move into the line-of-sight, we have to account for target visibility. Dynamic objects are modelled as ellipsoids S_{Ω_s} and we observe that point visibility can be decided by ‘horizon culling’, illustrated in In Fig. 4. The lines from the camera center p_c to the ellipsoid’s tangent define the horizon plane H_o . All points on the intersection of H_o and S_{Ω_s} are on the horizon, all points in the shaded region C_o are below the horizon and therefore not visible from the camera’s viewpoint. To this end we adopt a fast visibility test [25], which can be summarized as: (i) determine if a point is in front H_o (visible) and (ii) for those behind H_o whether they are within the infinite cone formed by the tangent lines (C_o).

To determine if a target t_i is in front of H_o , defined by t_j , we compute:

$$\begin{aligned} r_{chi} &= p_{hi} - p_c, \\ r_{ctj} &= p_{tj} - p_c, \\ p_{proj} &= r_{chi}^T r_{ctj}. \end{aligned}$$

Where r_{chi} is the vector to the target, r_{ctj} the vector to the center of ellipsoid S_{Ω_s} and p_{proj} is the component of r_{chi} in the direction of r_{ctj} . If $p_{proj} > r_{ctj}^T r_{ctj} - 1$ then the point is behind the plane.

For these cases we can determine if the point falls into the cone by comparing the opening angles α and β , see Fig. 4. Avoiding costly trigonometric functions we use directly the approach given in [25]. This gives us the visibility score:

$$d_v = \frac{p_{proj}}{r_{chi}^T r_{chi}} > r_{ctj}^T r_{ctj}.$$

And we define the visibility cost $c_{vis}(x)$:

$$c_{vis}(x) = \begin{cases} \|d_v\|_{Q_v} & \text{if } d_v > 0 \text{ and } p_{proj} > r_{ctj}^T r_{ctj} - 1, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

D. Collision Avoidance

A final concern of the planning algorithm is to generate collision-free trajectories. Here we assume known geometry of the scene and real-time position information of moving obstacles (targets are also considered as moving obstacles). To improve safety and at the same time guarantee performance and responsiveness of the system we adopt a two layered approach: (i) a potential field to repel the robot from obstacles and (ii) a hard constraint to stay outside a smaller enclosing ellipsoid S_{Ω_s} to enforce collision-free motion. We denote the relative vector between the center of the target and the camera by $r_{ct} = p_t - p_c$.

1) *Collision avoidance potential field*: First, we employ a potential field cost which becomes active as soon as the camera enters an ellipsoid S_{Ω_l} containing the target and with a buffer zone. This term helps to maintain a safe distance from moving obstacles and leaves some buffer for the un-modeled dynamics of the quadrotor and human motion. Formally, the distance to this ellipsoid is given by

$$d_c = r_{ct}^T \Omega_l r_{ct} - 1.$$

Which produces the cost term:

$$c_{coll}(x) = \begin{cases} \|d_c\|_{Q_c} & \text{if } d_c > 0 \\ 0 & \text{else} \end{cases}. \quad (12)$$

2) *Collision avoidance constraint*: Second, we introduce a non-linear constraint which becomes active within the collision ellipsoid S_{Ω_s} . This constraint guarantees collision-free motion, i.e. the position of the quadrotor must remain outside of the S_{Ω_s} at all times. Formally,

$$r_{ct}^T \Omega_s r_{ct} > 1 - s_c, \quad (13)$$

We introduce the slack variables s_c and, therefore, we adopt a soft-constrained approach for the collision-avoidance constraints to ensure that the optimizer always returns an answer in practice. It can be shown that under sufficiently high penalization of (a linear norm of) the slack variables s_c , the solution of the hard-constrained problem is recovered when it exists; otherwise, a plan with minimum deviation will be computed by the optimizer [26].

E. MPC formulation

The locally optimal trajectory and inputs for the quadrotor are then computed by solving a constrained optimization problem, which consists of the cost terms introduced in Eq. (7)-(12) and the constraints of Eq. (13), stacked together over all targets, obstacles and time-steps of the controller. Without loss of generality, we consider the obstacle set to be equal to the set of targets, since an obstacle can be treated as a target with only the collision avoidance constrained active - and no other cost term. The full constrained optimization problem solved as MPC non-linear program is then given by:

$$\begin{aligned}
\min_{\mathbf{x}, \mathbf{u}, \mathbf{s}} \quad & w_N^T c(\mathbf{x}_N, \mathbf{u}_N) + \sum_{k=0}^{N-1} w^T c(\mathbf{x}_k, \mathbf{u}_k) + \lambda \|\mathbf{s}_k\|_\infty \\
\text{s.t.} \quad & \mathbf{x}_0 = \hat{\mathbf{x}}_0 && \text{(Initial State)} \\
& \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), && \text{(Dynamics)} \\
& \mathbf{r}_{ct}^T \Omega_s \mathbf{r}_{ct} > 1 - s_k, && \text{(Collision Avoidance)} \\
& \mathbf{r}_{ct} = g(\mathbf{x}_k), && \\
& \mathbf{x}_k \in \mathcal{X}, && \text{(State Constraints)} \\
& \mathbf{u}_k \in \mathcal{U}, && \text{(Input Constraints)} \\
& s_k \geq 0, && \text{(Slack Constraints).}
\end{aligned} \tag{14}$$

where the stage cost function $c(\mathbf{x}_k, \mathbf{u}_k)$ is defined as:

$$c(\mathbf{x}_k, \mathbf{u}_k) = [c_{\text{image}} \quad c_{\text{size}} \quad c_{\text{angle}} \quad c_{\text{coll}} \quad c_{\text{vis}} \quad c_{\text{pose}}]^T_{(\mathbf{x}_k, \mathbf{u}_k)}$$

Where w are weights that can be set interactively by the user to control influence of the different framing constraints. This problem can readily be formulated in standard software, e.g. FORCES Pro [27, 28], and efficient code be generated to solve it in real-time.

V. EXPERIMENTAL RESULTS

A. Hardware Setup

We use a Parrot Bebop2 with an integrated electronic gimbal. Robot and targets are tracked with a Vicon motion capturing system. Experiments were conducted on a standard desktop PC (Quadcore Intel i7 CPU@3.5 GHz).

B. Experiments

We conducted five experiments to evaluate the performance of our proposed method under dynamic conditions. In our experiments we used a maximum number of three moving targets, although this is not an inherent limitation of the method. In the following results the boxes rendered on the images represent the setpoints. The size of a box determines the desired size of the projected target ellipse. The color coding is: Target 1 (red), Target 2 (purple) and Target 3 (blue).

Experiment 1 (Setpoint Change): First, we illustrate the effect of different framing objectives. A single target is filmed with different setpoints, including transitions from sitting to standing. Horizon length $N = 25$.

Experiment 2 (Single Target with Occlusion): Our first experiment demonstrates effectivity of the framing objectives (Eq. (7)-(9)) and occlusion handling (Eq. (11)). One face is set as target and two further actors move in and out of the line-of-sight. The goal is then to keep the main target always in view while minimizing occlusions from the other two targets. Horizon length $N = 25$.

Experiment 3 (Single target with collision avoidance): In the second experiment we modify above setup so that the two other targets walk randomly, including directly towards the quadrotor. Again the algorithm will try to keep the target in frame while avoiding collisions (Eq. (12)). Horizon length $N = 25$.

Experiment 4 (Multi target framing objective): In our final setup we declare all three actors to be targets and the algorithm tries to keep all three faces appearing along a single line in the image. Horizon length $N = 25$.

Experiment 5 (Execution time): To evaluate the real-time performance we conducted two experiments. We measure the execution time per timestep along a full trajectory (solving over the entire horizon per timestep). First, we use a constant number of targets $K = 2$ and vary horizon length ($N = 25, N = 40, N = 55$). Second, we keep the horizon fixed $N = 25$ and vary the number of targets.

C. Results

Experiment 1: Fig. 5 shows the algorithm working with a single target $K = 1$ and a horizon length of $N = 25$. With a single target we change the setpoints to: 3/4 Frontal right, 3/4 frontal left, right screen position, left screen position. The accompanying video also shows the target moving.

Experiment 2: The sequence in Fig. 6 illustrates how at time $t = 0s$ the initial condition is met and the face is visible according to the framing objective. At time $t = 1s$ target 3 moves into the line-of-sight and occludes target 2. The robot moves smoothly to the closest pose in terms of framing objectives but restores line-of-sight $t = 3s$.

Experiment 3: Fig. 7 shows representative views with the framing objective of creating a frontal shot of target 2. Although the targets 1 and 3 force the robot to perform multiple collision avoidance moves, the target remains in view and its screen space position remains relatively stable. Due to safety reasons collision avoidance takes the highest priority. Furthermore, collision avoidance (hard constraint) restricts the quadrotor motion stronger than the occlusion minimization (soft constraint).

Experiment 4: Fig. 8 shows representative frames from a shot with multi-target framing objective. Although there exists, except of some degenerated cases, a camera position with a zero re-projection error, the algorithm has to respect state constraints. In this case the quadrotor has to stay inside the physical room limits, reducing the range of motion drastically. Nevertheless all targets remain in view and the framing requirements are fulfilled as closely as possible.

Experiment 5: Fig. 9 plots the computation time for a trajectory with 2000 timesteps at each of which we solve Eq. (14) over the horizon length resulting in a total of $2000 \times N$ calls to the solver. We randomly vary the setpoints, ensuring that we don't initialize with unrealistic values too close to the solution. Furthermore, collision avoidance is turned on. The two experiments are parametrized: (i) First, a constant target number $K = 2$ and variable horizon length Fig. 9, top. The mean solve-times are: 0.01s, 0.015s and 0.025s respectively. (ii) Second, a constant horizon $N = 25$ and a variable number of targets K Fig. 9, bottom. The mean solve-times are: 0.009s, 0.011s and 0.011s. The computational time grows approximately linear with the length of the horizon. This result is expected according to the design of FORCES Pro. In general we see that the computation time for longer horizons and more targets also tend to have more variability in the solve-time.



Fig. 5: Exp. 1: The effect of viewpoint optimization under varying setpoints: 3/4 Frontal right, 3/4 frontal left, right screen position, left screen position (from left to right).




Fig. 6: Exp. 2: Effect of occlusion handling (onboard view). Entire duration is 3s. Inset shows robot position and non-visible areas.




Fig. 7: Exp. 3: Framing & collision avoidance. From left to right: Offboard view (robot in green). Onboard view (target in purple). Visualization of collision ellipsoids. Target position residual, gimbal reduces error in x but yawing the robot is slow (large error in y).




Fig. 8: Exp. 4: The Camera tries to keep the three targets in view. Distribution of the re-projection is shown (right).

We think this is due to the existence of a unique solution for a single target, while for multiple targets, depending on the setpoints, the solution may only exist in the least squares sense.

VI. CONCLUSION

In this paper we propose a novel trajectory generation method that solves for trajectory parameters from set-points defined in image space, and a formulation of this problem that lends itself to an implementation as a receding horizon optimal control program with non-linear constraints which can be solved numerically with state-of-the art solvers in real-time. We have demonstrated in experiments with an aerial

vehicle and multiple targets that the algorithm can satisfy framing constraints, derived from cinematographic rules, continuously minimize occlusions from dynamic objects in the environment and avoid collisions with these. Currently our algorithm accepts position and framing setpoints from a user. In future work this could be handled by a global planning algorithm that generates inputs for entire scenes rather than individual shots. We believe that the algorithm is general and could be adapted to other tasks for aerial vehicles which include collision avoidance with respect to moving obstacles. We release our work as open-source code².

²<http://ait.ethz.ch/mpc/>

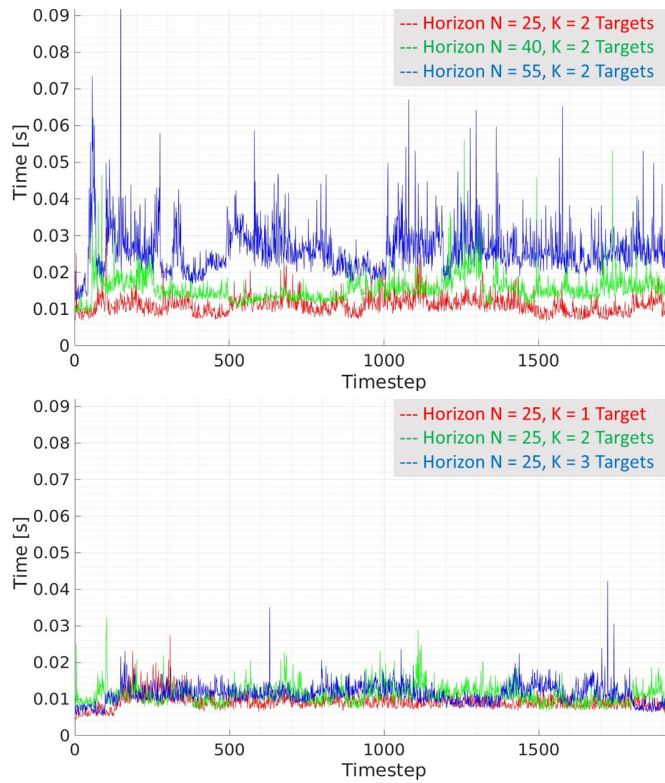


Fig. 9: Plots show solve time per evaluation of the MPC-problem (Eq. (14)) as function of horizon length (N) and number of targets (K). Note that we solve over N stages for each timestep (y-axis). *Top*: We change desired setpoints randomly throughout the trajectory but keep number of targets fixed. Colors indicate different horizon length. *Bottom*: Fixed horizon length, while changing number of targets.

REFERENCES

- [1] C. Gebhardt, B. Hepp, T. Naegeli, S. Stevsic, and O. Hilliges, "Airways: Optimization-based Planning of Quadrotor Trajectories according to High-Level User Goals," in *SIGCHI Conference on Human Factors in Computing Systems*, CHI '16, (New York, NY, USA), ACM, 2016.
- [2] N. Joubert, M. Roberts, A. Truong, F. Berthouzoz, and P. Hanrahan, "An interactive tool for designing quadrotor camera shots," *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2015)*, vol. 34, no. 6, p. 238, 2015.
- [3] M. Roberts and P. Hanrahan, "Generating dynamically feasible trajectories for quadrotor cameras," *ACM Transactions on Graphics (Proc. SIGGRAPH 2016)*, vol. 35, no. 4, 2016.
- [4] M. Gleicher and A. Witkin, "Through-the-lens camera control," vol. 26, (New York, NY, USA), ACM, July 1992.
- [5] D. S. US, H. Kim, and S. Sastry, *Decentralized reflective model predictive control of multiple flying robots in dynamic environment*. 2003.
- [6] A. Richards and J. How, "Decentralized model predictive control of cooperating UAVs," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pp. 4286–4291 Vol.4, IEEE, 2004.
- [7] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 2520–2525, IEEE, 2011.
- [8] J. Alonso-Mora, T. Naegeli, R. Siegwart, and P. Beardsley, "Collision Avoidance for Aerial Vehicles in Multi-agent Scenarios," *Auton. Robot.*, Jan. 2015.
- [9] M. W. Mueller and R. D'Andrea, "A model predictive controller for quadcopter state interception," in *Control Conference (ECC), 2013 European*, pp. 1383–1389, IEEE, 2013.
- [10] J. Schulman, Y. Duan, J. Ho, A. Lee, and I. Awwal, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, 2014.
- [11] M. Christie, P. Olivier, and J. M. Normand, "Camera control in computer graphics," *Computer Graphics Forum*, vol. 27, no. 8, pp. 2197–2218, 2008.
- [12] S. M. Drucker and D. Zeltzer, "Intelligent camera control in a virtual environment," in *In Proceedings of Graphics Interface 94*, pp. 190–199, 1994.
- [13] C. Lino, M. Christie, R. Ranon, and W. Bares, "The director's lens: An intelligent assistant for virtual cinematography," in *Proceedings of the 19th ACM International Conference on Multimedia*, MM '11, (New York, NY, USA), pp. 323–332, ACM, 2011.
- [14] C. Lino and M. Christie, "Intuitive and efficient camera control with the toric space," *ACM Transactions on Graphics (Proc. SIGGRAPH 2015)*, vol. 34, no. 4, p. 82, 2015.
- [15] Q. Galvane, J. Fleureau, F.-L. Tardieu, and P. Guillotel, "Automated Cinematography with Unmanned Aerial Vehicles," in *Eurographics Workshop on Intelligent Cinematography and Editing* (M. Christie, Q. Galvane, A. Jhala, and R. Ronfard, eds.), The Eurographics Association, 2016.
- [16] F. Poiesi and A. Cavallaro, "Distributed vision-based flying cameras to film a moving target," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pp. 2453–2459, IEEE, 2015.
- [17] K. Hausman, G. Kahn, S. Patil, J. Müller, and K. Goldberg, "Cooperative Occlusion-Aware Multi-Robot Target Tracking using Optimization," *rl.berkeley.edu*.
- [18] B. Espiau, F. Chaumette, and P. Rives, "A new approach to visual servoing in robotics," *Robotics and Automation, IEEE Transactions on*, vol. 8, no. 3, pp. 313–326, 1992.
- [19] J. Aloimonos, I. Weiss, and A. Bandyopadhyay, "Active vision," *Journal of computer vision*, 1988.
- [20] S. L. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 5448–5455, IEEE, 2011.
- [21] Daniel Arijon, *Grammar of the film language*. 1976.
- [22] University of Dallas, "Elements of Cinematography: Camera." <http://www.utdallas.edu/atec/midori/Handouts/camera.htm>. Accessed: 2016-09-02.
- [23] W. Breckenridge, "Quaternions proposed standard conventions," *Jet Propulsion Laboratory, Pasadena, CA, Interoffice Memorandum IOM*, pp. 343–79, 1999.
- [24] B. P. Gibbs, *Advanced Kalman filtering, least-squares and modeling: a practical handbook*. John Wiley & Sons, 2011.
- [25] P. Cozzi and F. Stoner, "Gpu ray casting of virtual globes," in *ACM SIGGRAPH 2010 Posters*, SIGGRAPH '10, (New York, NY, USA), ACM, 2010.
- [26] E. C. Kerrigan and J. M. Maciejowski, "Soft constraints and exact penalty functions in model predictive control," in *Control 2000 Conference*, Cambridge, 2000.
- [27] A. Domahidi, A. U. Zgraggen, M. N. Zeilinger, M. Morari, and C. N. Jones, "Efficient interior point methods for multistage problems arising in receding horizon control," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pp. 668–674, IEEE, 2012.
- [28] A. Domahidi and J. Jerez, "FORCES Pro: code generation for embedded optimization." <https://www.embotech.com/FORCES-Pro>, September 2016.