

AlphaPilot: Autonomous Drone Racing

Philipp Foehn*, Dario Brescianini*, Elia Kaufmann*,
Titus Cieslewski, Mathias Gehrig, Manasi Muglikar and Davide Scaramuzza

Abstract—This paper presents a novel system for autonomous, vision-based drone racing combining learned data abstraction, nonlinear filtering, and time-optimal trajectory planning. The system has successfully been deployed at the first autonomous drone racing world championship: the *2019 AlphaPilot Challenge*. Contrary to traditional drone racing systems, which only detect the next gate, our approach makes use of any visible gate and takes advantage of multiple, simultaneous gate detections to compensate for drift in the state estimate and build a global map of the gates. The global map and drift-compensated state estimate allow the drone to navigate through the race course even when the gates are not immediately visible and further enable to plan a near time-optimal path through the race course in real time based on approximate drone dynamics. The proposed system has been demonstrated to successfully guide the drone through tight race courses reaching speeds up to 8 m/s and ranked second at the *2019 AlphaPilot Challenge*.

Video of the race performance: <https://youtu.be/DGjwm5PZQT8>

I. INTRODUCTION

A. Motivation

Autonomous drones have seen a massive gain in robustness in recent years and perform an increasingly large set of tasks across various commercial industries; however, they are still far from fully exploiting their physical capabilities. Indeed, most autonomous drones only fly at low speeds near hover conditions in order to be able to robustly sense their environment and to have sufficient time to avoid obstacles. Faster and more agile flight could not only increase the flight range of autonomous drones, but also improve their ability to avoid fast dynamic obstacles and enhance their maneuverability in confined spaces. Human pilots have shown that drones are capable of flying through complex environments, such as race courses, at breathtaking speeds. However, autonomous drones are still far from human performance in terms of speed, versatility, and robustness, so that a lot of research and innovation is needed in order to fill this gap.

In order to push the capabilities and performance of autonomous drones, in 2019, Lockheed Martin and the Drone Racing League have launched the *AlphaPilot Challenge*^{1,2}, an open innovation challenge with a grand prize of \$1 million.

Authors with * contributed equally. All authors are with the Robotics and Perception Group, Dep. of Informatics, University of Zurich, and Dep. of Neuroinformatics, University of Zurich and ETH Zurich, Switzerland, <http://rpg.ifi.uzh.ch>. This work was supported by the Intel Network on Intelligent Systems, the Swiss National Science Foundation (SNSF) through the National Center of Competence in Research (NCCR) Robotics, the SNSF-ERC Starting Grant, and SONY.

¹<https://thedroneracingleague.com/airr/>

²<https://www.nytimes.com/2019/03/26/technology/alphapilot-ai-drone-racing.html>



Fig. 1. Our *AlphaPilot* drone waiting on the start podium to autonomously race through the gates ahead.

The goal of the challenge is to develop a fully autonomous drone that navigates through a race course using machine vision, and which could one day beat the best human pilot. While other autonomous drone races [1, 2] focus on complex navigation, the *AlphaPilot Challenge* pushes the limits in terms of speed and course size to advance the state of the art and enter the domain of human performance. Due to the high speeds at which drones must fly in order to beat the best human pilots, the challenging visual environments (e.g., low light, motion blur), and the limited computational power of drones, autonomous drone racing raises fundamental challenges in real-time state estimation, perception, planning, and control.

B. Related Work

Autonomous navigation in indoor or GPS-denied environments typically relies on simultaneous localization and mapping (SLAM), often in the form of visual-inertial odometry (VIO) [3]. There exists a variety of VIO algorithms, e.g., [4–7], that are based on feature detection and tracking that achieve very good results in general navigation tasks [8]. However, the performance of these algorithms significantly degrades during agile and high-speed flight as encountered in drone racing. The drone’s large translational and rotational velocities cause large optic flow, making robust feature detection and tracking over sequential images difficult and thus often causing substantial drift in the VIO state estimate [9].

To overcome this difficulty, several approaches exploiting the structure of drone racing with gates as landmarks have been developed, e.g., [10] and [11], where the drone locates

itself relative to gates. In [10], a handcrafted process is used to extract gate information from images that is then fused with attitude estimates from an inertial measurement unit (IMU) to compute an attitude reference that guides the drone towards the visible gate. While the approach is computationally very light-weight, it struggles with scenarios where multiple gates are visible and does not allow to employ more sophisticated planning and control algorithms which, e.g., plan several gates ahead. In [11], a convolutional neural network (CNN) is used to retrieve a bounding box of the gate and a line-of-sight-based control law aided by optic flow is then used to steer the drone towards the detected gate. The approach presented in [12] also relies on relative gate data but has the advantage that it works even when no gate is visible. In particular, it uses a CNN to directly infer relative gate poses from images and fuse the results with a VIO state estimate. However, the CNN does not perform well when multiple gates are visible as it is frequently the case for drone racing.

C. Contribution

The approach contributed herein builds upon the work of [12] and fuses VIO with a robust CNN-based gate corner detection using an extended Kalman filter (EKF), achieving high accuracy at little computational cost. The gate corner detections are used as static features to compensate for the VIO drift and to align the drones' flight path precisely with the gates. Contrary to all previous works [10–12], which only detect the next gate, our approach makes use of any gate detection and even profits from multiple simultaneous detections to compensate for VIO drift and build a global gate map. The global map allows the drone to navigate through the race course even when the gates are not immediately visible and further enables the usage of sophisticated path planning and control algorithms. In particular, a computationally efficient, sampling-based path planner (see e.g., [13], and references therein) is employed that plans near time-optimal paths through multiple gates ahead and is capable of adjusting the path in real time if the global map is updated.

II. ALPHAPILOT RACE FORMAT AND DRONE

A. Race Format

From more than 400 teams that participated in a series of qualification tests including a simulated drone race [14], the top nine teams were selected to compete in the 2019 *AlphaPilot Challenge*. The challenge consists of three qualification races and a final championship race at which the six best teams from the qualification races compete for the grand prize of \$1 million. Each race is implemented as a time trial competition in which each team is given three attempts to fly through a race course as fast as possible without competing drones on the course. Taking off from a start podium, the drones have to autonomously navigate through a sequence of gates with distinct appearances in the correct order and terminate at a designated finish gate. The race course layout, gate sequence, and position are provided ahead of each race up to approximately ± 3 m horizontal uncertainty, enforcing

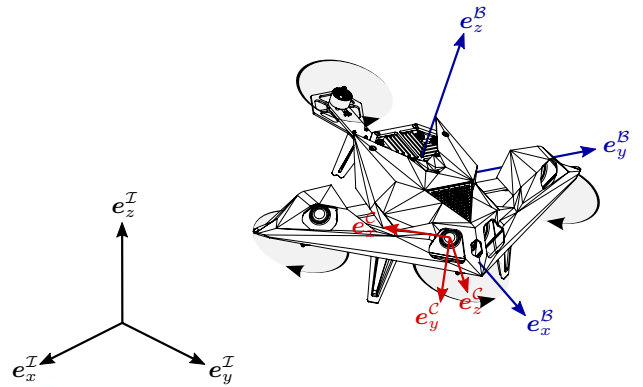


Fig. 2. Illustration of the race drone with its body-fixed coordinate frame \mathcal{B} in blue and a camera coordinate frame \mathcal{C} in red.

teams to come up with solutions that adapt to the real gate positions. Initially, the race courses were planned to have a lap length of approximately 300 m and required the completion up to three laps. However, due to technical difficulties, no race required to complete multiple laps and the track length at the final championship race was limited to about 74 m.

B. Drone Specifications

All teams were provided with an identical race drone (Fig. 1) that was approximately 0.7 m in diameter, weighed 3.4 kg, and had a thrust-to-weight ratio of 1.4. The drone was equipped with a NVIDIA Jetson Xavier embedded computer for interfacing all sensors and actuators and handling all computation for autonomous navigation onboard. The sensor suite included two $\pm 30^\circ$ forward-facing stereo camera pairs (Fig. 2), an IMU, and a downward-facing laser rangefinder (LRF). All sensor data were globally time stamped by software upon reception at the onboard computer. Detailed specifications of the available sensors are given in Table I. The drone was equipped with a flight controller that controlled the total thrust f along the drone's z -axis (see Fig. 2) and the angular velocity $\omega = (\omega_x, \omega_y, \omega_z)$ in the body-fixed coordinate frame \mathcal{B} .

C. Drone Model

Bold lower case and upper case letters will be used to denote vectors and matrices, respectively. The subscripts in $\mathcal{I}P_{CB} = \mathcal{I}P_B - \mathcal{I}P_C$ are used to express a vector from point C to point B expressed in frame \mathcal{I} . Without loss of generality, I is used to represent the origin of frame \mathcal{I} , and B represents the origin of coordinate frame \mathcal{B} . For the sake of readability,

TABLE I
SENSOR SPECIFICATIONS.

Sensor	Model	Rate	Details
Camera	Leopard Imaging IMX 264	60 Hz	global shutter, color, resolution: 1200×720
IMU	Bosch BMI088	430 Hz	range: $\pm 24g$, ± 34.5 rad/s resolution: $7e^{-4}g$, $1e^{-3}$ rad/s
LRF	Garmin LIDAR-Lite v3	120 Hz	range: 1-40 m resolution: 0.01 m

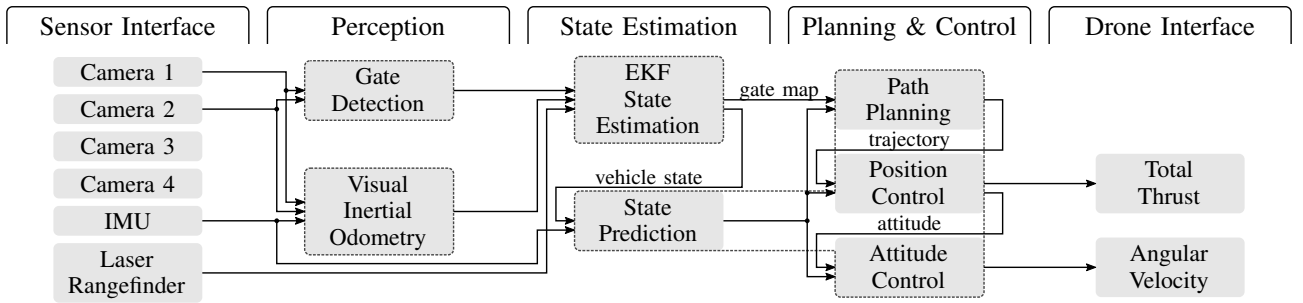


Fig. 3. Overview of the system architecture and its main components. All components within a dotted area run in a single thread.

the leading subscript may be omitted if the frame in which the vector is expressed is clear from context.

The drone is modelled as a rigid body of mass m with rotor drag proportional to its velocity acting on it [15]. The translational degrees-of-freedom are described by the position of its center of mass $\mathbf{p}_B = (p_{B,x}, p_{B,y}, p_{B,z})$ with respect to an inertial frame \mathcal{I} as illustrated in Fig. 2. The rotational degrees-of-freedom are parametrized using a unit quaternion \mathbf{q}_{IB} , where $\mathbf{R}_{IB} = \mathbf{R}(\mathbf{q}_{IB})$ denotes the rotation matrix mapping a vector from the body-fixed coordinate frame \mathcal{B} to the inertial frame \mathcal{I} [16]. A unit quaternion \mathbf{q} consists of a scalar q_w and a vector $\tilde{\mathbf{q}} = (q_x, q_y, q_z)$ and is defined as $\mathbf{q} = (q_w, \tilde{\mathbf{q}})$ [16]. The drone's equations of motion are

$$m\ddot{\mathbf{p}}_B = \mathbf{R}_{IB}f\mathbf{e}_z^{\mathcal{B}} - \mathbf{R}_{IB}\mathbf{D}\mathbf{R}_{IB}^T\mathbf{v}_B - m\mathbf{g}, \quad (1)$$

$$\dot{\mathbf{q}}_{IB} = \frac{1}{2} \begin{bmatrix} 0 \\ \boldsymbol{\omega} \end{bmatrix} \otimes \mathbf{q}_{IB}, \quad (2)$$

where f and $\boldsymbol{\omega}$ are the force and bodyrate inputs, $\mathbf{e}_z^{\mathcal{B}} = (0, 0, 1)$ is the drone's z -axis expressed in its body-fixed frame \mathcal{B} , $\mathbf{D} = \text{diag}(d_x, d_y, 0)$ is a constant diagonal matrix containing the rotor drag coefficients, $\mathbf{v}_B = \dot{\mathbf{p}}_B$ denotes the drone's velocity, \mathbf{g} is gravity and \otimes denotes the quaternion multiplication operator [16]. The drag coefficients were identified experimentally to be $d_x = 0.5 \text{ kg/s}$ and $d_y = 0.25 \text{ kg/s}$.

III. SYSTEM OVERVIEW

The system is composed of five functional groups: Sensor interface, perception, state estimation, planning and control, and drone interface (see Fig. 3). In the following, a brief introduction to the functionality of our proposed perception, state estimation, and planning and control system is given.

A. Perception

Of the two stereo camera pairs available on the drone, only the two central forward-facing cameras are used for gate detection (see Section IV) and, in combination with IMU measurements, to run VIO. The advantage is that the amount of image data to be processed is reduced while maintaining a very large field of view. Due to its robustness, multi-camera capability and computational efficiency, ROVIO [5] has been chosen as VIO pipeline. At low speeds, ROVIO is able to provide an accurate estimate of the quadrotor vehicle's pose and velocity relative to its starting position, however, at larger speeds the state estimate suffers from drift.

B. State Estimation

In order to compensate for a drifting VIO estimate, the output of the gate detection and VIO are fused together with the measurements from the downward-facing laser rangefinder (LRF) using an EKF (see Section V). The EKF estimates a global map of the gates and, since the gates are stationary, uses the gate detections to align the VIO estimate with the global gate map, i.e., compensates for the VIO drift.

Computing the state estimate, in particular interfacing the cameras and running VIO, introduces latency in the order of 130 ms to the system. In order to be able to achieve a high bandwidth of the control system despite large latencies, the vehicle's state estimate is predicted forward to the vehicle's current time using the IMU measurements.

C. Planning and Control

The global gate map and the latency-compensated state estimate of the vehicle are used to plan a near time-optimal path through the next N gates starting from the vehicle's current state (see Section VI). The path is re-planned every time (i) the vehicle passes through a gate, (ii) the estimate of the gate map or (iii) the VIO drift are updated significantly, i.e., large changes in the gate positions or VIO drift. The path is tracked using a cascaded control scheme (see Section VII) with an outer proportional-derivative (PD) position control loop and an inner proportional (P) attitude control loop. Finally, the outputs of the control loops, i.e., a total thrust and angular velocity command, are sent to the drone.

D. Software Architecture

The NVIDIA Jetson Xavier provides eight CPU cores, however, four cores are used to run the sensor and drone interface. The other four cores are used to run the gate detection, VIO, EKF state estimation, and planning and control, each in a separate thread on a separate core. All threads are implemented asynchronously to run at their own speed, i.e., whenever new data are available, in order to maximize data throughput and to reduce processing latencies. The gate detection thread is able to process all camera images in real time at 60 Hz, whereas the VIO thread only achieves approximately 35 Hz. In order to deal with the asynchronous nature of the gate detection and VIO thread and their output, all data are globally time stamped and integrated in the EKF accordingly. The EKF thread runs every time a new gate detection or LRF measurement is




Fig. 4. The gate detection module returns sets of corner points for each gate in the input image (fourth column) using a two-stage process. In the first stage, a neural network transforms an input image $\mathbf{I}_{w \times h \times 3}$ (first column) into a set of confidence maps for corners $\mathbf{C}_{w \times h \times 4}$ (second column) and Part Affinity Fields (PAFs) [17] $\mathbf{E}_{w \times h \times (4 \cdot 2)}$ (third column). In the second stage, the PAFs are used to associate sets of corner points that belong to the same gate. For visualization, both corner maps \mathbf{C} (second column) and PAFs \mathbf{E} (third column) are displayed in a single image each. While color encodes the corner class for \mathbf{C} , it encodes the direction of the 2D vector fields for \mathbf{E} . The yellow lines in the bottom of the second column show the six edge candidates of the edge class (TL, TR) (the TL corner of the middle gate is below the detection threshold), see Section IV-B. Best viewed in color.

available. The planning and control thread is executed at a fixed rate of 50 Hz. To achieve this, the planning and control thread includes the state prediction which compensates for latencies introduced by the VIO.

IV. GATE DETECTION

To correct for drift accumulated by the VIO pipeline, the gates are used as distinct landmarks for relative localization. In contrast to previous CNN-based approaches to gate detection, we do not infer the relative pose to a gate directly, but instead segment the four corners of the observed gate in the input image. This allows the detection of an arbitrary amount of gates, and allows for a more principled inclusion of gate measurements in the EKF through the use of reprojection error. Furthermore, it exhibits more predictable behavior for partial gate observations and overlapping gates. Since the exact shape of the gates is known, detecting a set of characteristic points per gate allows to constrain the relative pose. For the quadratic gates of the *AlphaPilot Challenge*, these characteristic points are chosen to be the inner corner of the gate border (see Fig. 4, 4th column). However, just detecting the four corners of all gates is not enough. If just four corners of several gates are extracted, the association of corners to gates is undefined (see Fig. 4, 3rd row, 2nd column). To solve this problem, we additionally train our network to extract so-called Part Affinity Fields (PAFs), as proposed by [17]. These are vector fields, which, in our case, are defined along the edges of the gates, and point from one corner to the next corner of the same gate, see column three in Figure 4. In Section IV-B, we describe how the PAFs are then used to solve the aforementioned gate

association problem.

A. Stage 1: Predicting Corner Maps and Part Affinity Fields

In the first detection stage, each input image \mathbf{I} is mapped by a neural network into a set of $N_C = 4$ corner maps and $N_E = 4$ PAFs. The network is trained in a supervised manner by minimizing the Mean-Squared-Error loss between the network prediction and ground-truth maps. In the following, ground-truth maps for both map types are explained in detail.

1) *Corner Maps*: For each corner class $j \in \mathcal{C}_j$, $\mathcal{C}_j := \{TL, TR, BL, BR\}$, a ground-truth corner map $\mathbf{C}_j^*(\mathbf{s})$ is represented by a single-channel map of the same size as the input image and indicates the existence of a corner of class j at pixel location \mathbf{s} in the image. The value at location $\mathbf{s} \in \mathbf{I}$ in \mathbf{C}_j^* is defined by a Gaussian as

$$\mathbf{C}_j^*(\mathbf{s}) = \exp\left(-\frac{\|\mathbf{s} - \mathbf{s}_j^*\|_2^2}{\sigma^2}\right), \quad (3)$$

where \mathbf{s}_j^* denotes the ground truth image position of the nearest corner with class j . The choice of the parameter σ controls the width of the Gaussian. We use $\sigma = 7$ pixel in our implementation. Gaussians are used to account for small errors in the ground truth corner positions that are provided by hand.

2) *Part Affinity Fields* : We define a PAF for each of the four possible classes of edges, defined by its two connecting corners as $(k, l) \in \mathcal{E}_{KL} := \{(TL, TR), (TR, BR), (BR, BL), (BL, TL)\}$. For each edge class (k, l) the ground-truth PAF $\mathbf{E}_{(k, l)}^*(\mathbf{s})$ is represented by a two-channel map of the same size as the input image and

points from corner k to corner l of the same gate, provided that the given image point \mathbf{s} lies within distance d of such an edge. We use $d = 10$ pixel in our implementation. Let G be the set of gates g and $\mathcal{S}_{(k,l),g}$ be the set of image points that are within distance d of the line connecting the corner points \mathbf{s}_k^* and \mathbf{s}_l^* belonging to gate g . Furthermore, let $\mathbf{v}_{k,l,g}$ be the unit vector pointing from \mathbf{s}_k^* to \mathbf{s}_l^* of the same gate. Then, the part affinity field $\mathbf{E}_{(k,l)}^*(\mathbf{s})$ is defined as:

$$\mathbf{E}_{(k,l)}^*(\mathbf{s}) = \begin{cases} \mathbf{v}_{k,l,g} & \text{if } \mathbf{s} \in \mathcal{S}_{(k,l),g}, \quad g \in G \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (4)$$

Note that a special case might exist in which the same point \mathbf{s} lies in $\mathcal{S}_{(k,l),g}$ of several gates. In that case, the $\mathbf{v}_{k,l,g}$ of all corresponding gates are averaged.

B. Stage 2: Corner Association

Discrete corner locations \mathbf{s}_j for each class $j \in \mathcal{C}_j$ are extracted from the corner map using non-maximum suppression and thresholding. This allows the formation of an exhaustive set of edge candidates $\{(\mathbf{s}_k, \mathbf{s}_l)\}$, see the yellow lines in Fig. 4. Given the corresponding PAF $\mathbf{E}_{(k,l)}(\mathbf{s})$, each edge candidate is assigned a score which expresses the agreement of that candidate with the PAF. This score is given by the line integral

$$\mathcal{S}((\mathbf{s}_k, \mathbf{s}_l)) = \int_{u=0}^{u=1} \mathbf{E}_{(k,l)}(\mathbf{s}(u)) \cdot \frac{\mathbf{s}_l - \mathbf{s}_k}{\|\mathbf{s}_l - \mathbf{s}_k\|} du, \quad (5)$$

where $\mathbf{s}(u)$ linearly interpolates between the two corner candidate locations \mathbf{s}_k and \mathbf{s}_l . In practice, the line integral \mathcal{S} is approximated by uniformly sampling the integrand.

As described in [17], extracting the “best” set of edges for class (k, l) according to this score is an instance of the maximum weight matching problem in a bipartite graph, as each corner $j \in \{k, l\}$ can only be assigned one edge. Once this problem is solved for each of the four edge classes, the pairwise associations can be extended to sets of associated points for each gate. We refer the reader to [17] for the detailed solutions of these problems.

C. Network Architecture, Training Data and Deployment

The network architecture deployed consists of a 5-level U-Net [18] with [12, 18, 24, 32, 32] convolutional filters of size [3, 3, 3, 5, 7] per level. At each layer, the input feature map is zero-padded to preserve a constant height and width throughout the network. As activation function, LeakyReLU with $\alpha = 0.01$ is used. The network is trained on a dataset consisting of 28k images recorded in 5 different environments. Each sample is annotated using the open source image annotation software labelme³, which is extended with KLT-Tracking for semi-automatic labelling. For deployment on the Jetson Xavier, the network is ported to TensorRT 5.0.2.6. To optimize memory footprint and inference time, inference is performed in half-precision mode (FP16) and batches of two images are fed to the network.

³<https://github.com/wkentaro/labelme>

V. STATE ESTIMATION

The nonlinear measurement models of the VIO, gate detection, and laser rangefinder are fused using an EKF [19]. In order to obtain the best possible pose accuracy relative to the gates, the EKF estimates the translational and rotational misalignment of the VIO origin frame \mathcal{V} with respect to the inertial frame \mathcal{I} , represented by \mathbf{p}_V and $\mathbf{q}_{\mathcal{I}\mathcal{V}}$, jointly with the gate positions \mathbf{p}_{G_i} and gate heading $\varphi_{\mathcal{I}G_i}$. It can thus correct for an imprecise initial position estimate, VIO drift, and uncertainty in gate positions. The EKF’s state space at time t_k is $\mathbf{x}_k = \mathbf{x}(t_k)$ with covariance \mathbf{P}_k , described by

$$\mathbf{x}_k = (\mathbf{p}_V, \mathbf{q}_{\mathcal{I}\mathcal{V}}, \mathbf{p}_{G_0}, \varphi_{\mathcal{I}G_0}, \dots, \mathbf{p}_{G_{N-1}}, \varphi_{\mathcal{I}G_{N-1}}). \quad (6)$$

The drone’s corrected pose $(\mathbf{p}_B, \mathbf{q}_{\mathcal{I}B})$ can then be computed from the VIO estimate $(\mathbf{p}_{VB}, \mathbf{q}_{\mathcal{V}B})$ by transforming it from frame \mathcal{V} into the inertial frame \mathcal{I} using $(\mathbf{p}_V, \mathbf{q}_{\mathcal{I}\mathcal{V}})$.

All estimated parameters are expected to be time-invariant but subject to noise and drift. This is modelled by a Gaussian random walk, simplifying the EKF process update to:

$$\mathbf{x}_{k+1} = \mathbf{x}_k, \quad \mathbf{P}_{k+1} = \mathbf{P}_k + \Delta t_k \mathbf{Q}, \quad (7)$$

where \mathbf{Q} is the random walk process noise. For each measurement \mathbf{z}_k with noise \mathbf{R} the predicted *a priori* estimate \mathbf{x}_k^- is corrected with measurement function $\mathbf{h}(\mathbf{x}_k^-)$ and Kalman gain \mathbf{K}_k resulting in the *a posteriori* estimate \mathbf{x}_k^+ , as

$$\begin{aligned} \mathbf{K}_k &= \mathbf{P}_k^- \mathbf{H}_k^\top (\mathbf{H}_k \mathbf{P}_k^- \mathbf{H}_k^\top + \mathbf{R})^{-1}, \\ \mathbf{x}_k^+ &= \mathbf{x}_k^- + \mathbf{K}_k (\mathbf{z}_k - \mathbf{h}(\mathbf{x}_k^-)), \\ \mathbf{P}_k^+ &= (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k^-, \end{aligned} \quad (8)$$

where $\mathbf{h}(\mathbf{x}_k^-)$ is the measurement function with jacobian \mathbf{H}_k .

To apply the EKFs linear update step on the over-parametrized quaternion, it is lifted to its tangent space description, similar to [20]. The quaternion $\mathbf{q}_{\mathcal{I}\mathcal{V}}$ is described by a reference quaternion $\mathbf{q}_{\mathcal{I}\mathcal{V}_{\text{ref}}}$, which is adjusted after each update step, and an error quaternion $\mathbf{q}_{\mathcal{V}_{\text{ref}}\mathcal{V}}$, of which only its vector part $\tilde{\mathbf{q}}_{\mathcal{V}_{\text{ref}}\mathcal{V}}$ is in the EKF’s state space.

A. Measurement Modalities

All measurements at time t_k are passed to the EKF together with the VIO estimate $\mathbf{p}_{VB,k}$ and $\mathbf{q}_{\mathcal{V}B,k}$ with respect to the VIO frame \mathcal{V} .

1) *Gate Measurements*: Gate measurements consist of the image pixel coordinates $\mathbf{s}_{C_{O_{ij}}}$ of a specific gate corner. Corners are denoted with top left and right, and bottom left and right, as in $j \in \{TL, TR, BL, BR\}$ and the gates are enumerated $i \in [0, N-1]$. All gates are of equal width w and height h , so that the corner positions in the gate frame \mathcal{G}_i can be written as $\mathbf{p}_{G_i C_{O_{ij}}} = \frac{1}{2} (0, \pm w, \pm h)$. The measurement equation can be written as the pinhole camera projection [21] of the gate corner into the camera frame. A pinhole camera maps the gate corner point $\mathbf{p}_{C_{O_{ij}}}$ expressed in the camera frame \mathcal{C} into pixel coordinates as

$$\mathbf{h}_{\text{Gate}}(\mathbf{x}) = \mathbf{s}_{C_{O_{ij}}} = \frac{1}{[\mathbf{p}_{C_{O_{ij}}}]_z} \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \end{bmatrix} \mathbf{p}_{C_{O_{ij}}}, \quad (9)$$

where $[\cdot]_z$ indicates the scalar z -component of a vector, f_x and f_y are the camera's focal lengths and (c_x, c_y) is the camera's optical center. The gate corner point $\mathbf{p}_{C_{o_{ij}}}$ is given by

$$\mathbf{p}_{C_{o_{ij}}} = \mathbf{R}_{IC}^T (\mathbf{p}_{G_i} + \mathbf{R}_{IG_i} \mathbf{p}_{G_i, C_{o_j}} - \mathbf{p}_C), \quad (10)$$

with \mathbf{p}_C and \mathbf{R}_{IC} being the transformation between the inertial frame \mathcal{I} and camera frame \mathcal{C} ,

$$\mathbf{p}_C = \mathbf{p}_V + \mathbf{R}_{IV} (\mathbf{p}_{VB} + \mathbf{R}_{VB} \mathbf{p}_{BC}), \quad (11)$$

$$\mathbf{R}_{IC} = \mathbf{R}_{IV} \mathbf{R}_{VB} \mathbf{R}_{BC}, \quad (12)$$

where \mathbf{p}_{BC} and \mathbf{R}_{BC} describe a constant transformation between the drone's body frame \mathcal{B} and camera frame \mathcal{C} (see Fig. 2). The Jacobian with respect to the EKF's state space is derived using the chain rule,

$$\frac{\delta}{\delta \mathbf{x}} \mathbf{h}_{\text{Gate}}(\mathbf{x}) = \frac{\delta \mathbf{h}_{\text{Gate}}(\mathbf{x})}{\delta \mathbf{p}_{C_{o_{ij}}}(\mathbf{x})} \cdot \frac{\delta \mathbf{p}_{C_{o_{ij}}}(\mathbf{x})}{\delta \mathbf{x}}, \quad (13)$$

where the first term representing the derivative of the projection remains the same for all components of the state space.

2) *Laser Rangefinder Measurement*: The drone's laser rangefinder measures the distance along the drone's negative z -axis to the ground, which is assumed to be flat and at a height of 0 m. The measurement equation can be described by

$$h_{\text{LRF}}(\mathbf{x}) = \frac{[\mathbf{p}_B]_z}{[\mathbf{R}_{IB} \mathbf{e}_z^B]_z} = \frac{[\mathbf{p}_V + \mathbf{R}_{IV} \mathbf{p}_{VB}]_z}{[\mathbf{R}_{IV} \mathbf{R}_{VB} \mathbf{e}_z^B]_z}. \quad (14)$$

The Jacobian with respect to the state space is again derived by $\frac{\delta h_{\text{LRF}}}{\delta \mathbf{p}_V}$ and $\frac{\delta h_{\text{LRF}}}{\delta \mathbf{q}_{IV}}$.

VI. PATH PLANNING

For the purpose of path planning, the drone is assumed to be a point mass with bounded accelerations as inputs. This simplification allows for the computation of time-optimal motion primitives in closed-form and enables the planning of time-optimal paths through the race course in real time. Although the dynamics of the quadrotor vehicle's acceleration cannot be neglected in practice, it is assumed that this simplification still captures the most relevant dynamics for path planning and that the resulting paths approximate the true time-optimal paths well. In the following, time-optimal motion primitives based on the simplified dynamics are first introduced and then a path planning strategy based on these motion primitives is presented.

A. Time-Optimal Motion Primitive

The minimum times T_x^* , T_y^* and T_z^* required for the vehicle to fly from an initial state, consisting of position and velocity, to a final state while satisfying the simplified dynamics $\ddot{\mathbf{p}}_B(t) = \mathbf{u}(t)$ with the input acceleration $\mathbf{u}(t)$ being constrained to $\underline{\mathbf{u}} \leq \mathbf{u}(t) \leq \bar{\mathbf{u}}$ are computed for each axis individually. Without loss of generality, only the x -axis is considered in the following. Using Pontryagin's maximum principle [22], it can be shown that the optimal control input is bang-bang in acceleration, i.e., has the form

$$u_x^*(t) = \begin{cases} \underline{u}_x, & 0 \leq t \leq t^*, \\ \bar{u}_x, & t^* < t \leq T_x^*, \end{cases} \quad (15)$$




Fig. 5. Example time-optimal motion primitive starting from rest at the origin to a random final position with non-zero final velocity. The velocities are constrained to ± 7.5 m/s and the inputs to ± 12 m/s². The dotted lines denote the per-axis time-optimal maneuvers.

or vice versa with the control input first being \bar{u}_x followed by \underline{u}_x . In order to control the maximum velocity of the vehicle, e.g., to constrain the solutions to ranges where the simplified dynamics approximate the true dynamics well or to limit the motion blur of the camera images, a velocity constraint of the form $\underline{\mathbf{v}}_B \leq \mathbf{v}_B(t) \leq \bar{\mathbf{v}}_B$ can be added, in which case the optimal control input has bang-singular-bang solution [23]

$$u_x^*(t) = \begin{cases} \underline{u}_x, & 0 \leq t \leq t_1^*, \\ 0, & t_1^* < t \leq t_2^*, \\ \bar{u}_x, & t_2^* < t \leq T_x^*, \end{cases} \quad (16)$$

or vice versa. It is straightforward to verify that there exist closed-form solutions for the minimum time T_x^* as well as the switching times t^* in both cases (15) or (16).

Once the minimum time along each axis is computed, the maximum minimum time $T^* = \max(T_x^*, T_y^*, T_z^*)$ is computed and motion primitives of the same form as in (15) or (16) are computed among the two faster axes but with the final time constrained to T^* such that trajectories along each axis end at the same time. In order for such a motion primitive to exist, a new parameter $\alpha \in [0, 1]$ is introduced that scales the acceleration bounds, i.e., the applied control inputs are scaled to $\alpha \underline{u}_x$ and $\alpha \bar{u}_x$, respectively. Fig. 5 depicts the position and velocity of an example time-optimal motion primitive.

B. Sampling-Based Receding Horizon Path Planning

The objective of the path planner is to find the time-optimal path from the drone's current state to the final gate, passing through all the gates in the correct order. Since the previously introduced motion primitive allows the generation of time-optimal motions between any initial and any final state, the time-optimal path can be planned by concatenating a time-optimal motion primitive starting from the drone's current (simplified) state to the first gate with time-optimal motion primitives that connect the gates in the correct order until the final gate. This reduces the path planning problem to finding the drone's optimal state at each gate such that the total time is minimized. To find the optimal path, a sampling-based strategy

is employed where states at each gate are randomly sampled and the total time is evaluated subsequently. In particular, the position of each sampled state at a specific gate is fixed to the center of the gate and the velocity is sampled uniformly at random such the velocity lies within the constraints of the motion primitives and the angle between the velocity and the gate normal does not exceed a maximum angle φ_{\max} . It is trivial to show that as the number of sampled states approaches infinity, the computed path converges to the time-optimal path.

In order to solve the problem efficiently, the path planning problem is interpreted as a shortest path problem. At each gate, M different velocities are sampled and the arc length from each sampled state at the previous gate is set to be equal to the duration T^* of the time-optimal motion primitive that guides the drone from one state to the other. Due to the existence of a closed-form expression for the minimum time T^* , setting up and solving the shortest path problem can be done very efficiently using, e.g., Dijkstra's algorithm [22]. In order to further reduce the computational cost, the path is planned in a receding horizon fashion, i.e., the path is only planned through the next N gates.

VII. CONTROL

This section presents a control strategy to track the near time-optimal path from Section VI. The control strategy is based on a cascaded control scheme with an outer position control loop and an inner attitude control loop, where the position control loop is designed under the assumption that the attitude control loop can track set point changes perfectly, i.e., without any dynamics or delay.

A. Position Control

The position control loop along the inertial z -axis is designed such that it responds to position errors $p_{B_{\text{err}},z} = p_{B_{\text{ref}},z} - p_{B,z}$ in the fashion of a second-order system with time constant $\tau_{\text{pos},z}$ and damping ratio $\zeta_{\text{pos},z}$,

$$\ddot{p}_{B,z} = \frac{1}{\tau_{\text{pos},z}^2} p_{B_{\text{err}},z} + \frac{2\zeta_{\text{pos},z}}{\tau_{\text{pos},z}} \dot{p}_{B_{\text{err}},z} + \ddot{p}_{B_{\text{ref}},z}. \quad (17)$$

Similarly, two control loops along the inertial x - and y -axis are shaped to make the horizontal position errors behave like second-order systems with time constants $\tau_{\text{pos},xy}$ and damping ratio $\zeta_{\text{pos},xy}$. Inserting (17) into the translational dynamics (1), the total thrust f is computed to be

$$f = \frac{[m(\ddot{\mathbf{p}}_{B_{\text{ref}}} + \mathbf{g}) + \mathbf{R}_{IB} \mathbf{D} \mathbf{R}_{IB}^T \mathbf{v}_B]_z}{[\mathbf{R}_{IB} \mathbf{e}_z^B]_z}. \quad (18)$$

B. Attitude Control

The required acceleration from the position controller determines the orientation of the drone's z -axis and is used, in combination with a reference yaw angle φ_{ref} , to compute the drone's reference attitude. The reference yaw angle is chosen such that the drone's x -axis points towards the reference position 5 m ahead of the current position, i.e., that the drone looks in the direction it flies. A nonlinear attitude controller similar to [24] is applied that prioritizes the alignment of the

drone's z -axis, which is crucial for its translational dynamics, over the correction of the yaw orientation:

$$\boldsymbol{\omega} = \frac{2 \operatorname{sgn}(q_w)}{\sqrt{q_w^2 + q_z^2}} \mathbf{T}_{\text{att}}^{-1} \begin{bmatrix} q_w q_x - q_y q_z \\ q_w q_y + q_x q_z \\ q_z \end{bmatrix}, \quad (19)$$

where q_w , q_x , q_y and q_z are the components of the attitude error $\mathbf{q}_{\mathcal{IB}}^{-1} \otimes \mathbf{q}_{\mathcal{IB}_{\text{ref}}}$ and where \mathbf{T}_{att} is a diagonal matrix containing the per-axis first-order system time constants for small attitude errors.

VIII. RESULTS

The proposed system was used to race in the *2019 AlphaPilot* championship race. The course at the championship race consisted of five gates and had a total length of 74 m. A top view of the race course as well as the results of the path planning and the fastest actual flight are depicted in Fig. 6 (left and center). With the motion primitive's maximum velocity set to 8 m/s, the drone successfully completed the race course in a total time of 11.36 s, with only two other teams also completing the full race course. The drone flew at an average velocity of 6.5 m/s and reached the peak velocity of 8 m/s multiple times. Note that due to missing ground truth, Fig. 6 only shows the estimated and corrected drone position.

The system was further evaluated at a testing facility where there was sufficient space for the drone to fly multiple laps (see Fig. 6, right), albeit the course consisted of only two gates. The drone was commanded to pass four times through *gate 1* before finishing in the *final gate*. Although the gates were not visible to the drone for most of the time, the drone successfully managed to fly multiple laps. Thanks to the global gate map and the VIO state estimate, the system was able to plan and execute paths to gates that are not directly visible. By repeatedly seeing either one of the two gates, the drone was able to compensate for the drift of the VIO state estimate, allowing the drone to pass the gates every time exactly through their center. Note that although seeing *gate 1* in Fig. 6 (right) at least once was important in order to update the position of the gate in the global map, the VIO drift was also estimated by seeing the *final gate*.

The results of the system's main components are discussed in detail in the following subsections, and a video of the results is attached to the paper.

A. Gate Detection

Even in instances of strong illumination changes, the gate detector was able to accurately identify the gates in a range of 2 – 17 m. Fig. 4 illustrates the quality of detections during the championship race (1st row) as well as for cases with multiple gates, represented in the test set (2nd/3rd row).

Gate detection is evaluated quantitatively on a separate test set of 4k images with respect to intersection over union (IoU) and false positive/negative corner predictions. While the IoU score only takes full gate detections into account, the false positives/negatives are computed for each corner detection. On




Fig. 6. Top view of the planned (left) and executed (center) path at the championship race, and an executed multi-lap path at a testing facility (right). Left: Fastest planned path in color, sub-optimal sampled paths in gray. Center: VIO trajectory as p_{VB} and corrected estimate as p_B .

the test set, the network achieves an IoU score with the human-annotated ground truth of 96.4%, an average false negative rate of 0.18 corners per image and an average false positive rate of 0.015 corners per image.

With the network architecture explained in Section IV, one simultaneous inference for the left- and right-facing camera requires computing 3.86 GFLOPS (40 kFLOPS per pixel). By implementing the network in TensorRT and performing inference in half-precision mode (FP16), this computation takes 10.5ms on the Jetson Xavier and can therefore be performed at the camera update rate.

B. State Estimation

Compared to a pure VIO-based solution, the EKF has proven to significantly improve the accuracy of the state estimation relative to the gates. As opposed to the works by [10–12], the proposed EKF is not constrained to only use the next gate, but can work with any gate detection and even profits from multiple detections in one image. Fig. 6 (center) depicts the flown trajectory estimated by the VIO system as p_{VB} and the EKF-corrected trajectory as p_B (the estimated corrections are depicted in gray). Accumulated drift clearly leads to a large discrepancy between VIO estimate p_{VB} and the corrected estimate p_B . Towards the end of the track at the two last gates this discrepancy would be large enough to cause the drone to crash into the gate. However, the filter corrects this discrepancy accurately and provides a precise pose estimate relative to the gates. Additionally, the imperfect initial pose, in particular the yaw orientation, is corrected by the EKF while flying towards the first gate as visible in the zoomed section in Fig. 6 (center).

C. Planning and Control

Fig. 6 (left) shows the nominally planned path for the *AlphaPilot* championship race, where the coloured line depicts the fastest path along all the sampled paths depicted in gray. In particular, a total of $M = 150$ different states are sampled at each gate, with the velocity limited to 8 m/s and the angle between the velocity and the gate normal limited to $\varphi_{\max} = 30^\circ$. During flight, the path is re-planned in a receding horizon fashion through the next $N = 3$ gates (see Fig. 6, center). It was experimentally found that choosing $N \geq 3$

greatly reduces the computational cost w.r.t. planning over the full track, while having only minimal impact on the flight time. Re-planning the paths takes less than 2 ms on the Jetson Xavier and can be done in every control update step.

Fig. 6 (right) shows resulting path and velocity of the drone in a multi-lap scenario, where the drone’s velocity was limited to 6 m/s. It can be seen that drone’s velocity is decreased when it has to fly a tight turn due to its limited thrust.

IX. DISCUSSION AND CONCLUSION

The proposed system managed to complete the course at a velocity of 5 m/s with a success rate of 100% and at 8 m/s with a success rate of 60%. At higher speeds, the combination of VIO tracking failures and no visible gates caused the drone to crash after passing the first few gates. This failure could be caught by integrating the gate measurements directly in a VIO pipeline, tightly coupling all sensor data. Another solution could be a perception-aware path planner trading off time-optimality against motion blur and maximum gate visibility.

The advantages of the proposed system are (i) a drift-free state estimate at high speeds, (ii) a global and consistent gate map, and (iii) a real-time capable near time-optimal path planner. However, these advantages could only partially be exploited as the races neither included multiple laps, nor had complex segments where the next gates were not directly visible. Nevertheless, the system has proven that it can handle these situations and is able to navigate through complex race courses reaching speeds up to 8 m/s and completing the championship race track of 74 m in 11.36 s.

While the *2019 AlphaPilot Challenge* pushed the field of autonomous drone racing, in particularly in terms of speed, autonomous drones are still far away from beating human pilots. Moreover, the challenge also left open a number of problems, most importantly that the race environment was partially known and static without competing drones or moving gates. In order for autonomous drones to fly at high speeds outside of controlled or known environments and succeed in many more real-world applications, they must be able to handle unknown environments, perceive obstacles and react accordingly. These features are areas of active research and are intended to be included in future versions of the proposed drone racing system.

REFERENCES

- [1] Hyungpil Moon, Yu Sun, Jacky Baltes, and Si Jung Kim. The IROS 2016 Competitions. *IEEE Robotics Automation Magazine*, 24(1):20–29, March 2017.
- [2] Hyungpil Moon, Jose Martinez-Carranza, Titus Cieslewski, Matthias Faessler, Davide Falanga, Alessandro Simovic, Davide Scaramuzza, Shuo Li, Michael Ozo, Christophe De Wagter, Guido de Croon, Sunyou Hwang, Sunggoo Jung, Hyunchul Shim, Haeryang Kim, Minhyuk Park, Tsz-Chiu Au, and Si Jung Kim. Challenges and implemented technologies used in autonomous drone racing. *Intelligent Service Robotics*, 12:137 – 148, 2019.
- [3] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, Jose Neira, Ian Reid, and John J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, Dec 2016. doi: 10.1109/TRO.2016.2624754.
- [4] Anastasios I. Mourikis and Stergios I. Roumeliotis. A multi-state constraint Kalman filter for vision-aided inertial navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3565–3572, April 2007.
- [5] Michael Bloesch, Sammy Omari, Marco Hutter, and Roland Siegwart. Robust visual inertial odometry using a direct EKF-based approach. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.
- [6] Tong Qin, Peiliang Li, and Shaojie Shen. VINS-Mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transactions on Robotics*, 34(4):1004–1020, 2018. doi: 10.1109/TRO.2018.2853729.
- [7] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. SVO: Semidirect visual odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2017. doi: 10.1109/TRO.2016.2623335.
- [8] Jeffrey Delmerico and Davide Scaramuzza. A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [9] Jeffrey Delmerico, Titus Cieslewski, Henri Rebecq, Matthias Faessler, and Davide Scaramuzza. Are we ready for autonomous drone racing? the UZH-FPV drone racing dataset. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [10] Shuo Li, Erik van der Horst, Philipp Duernay, Christophe De Wagter, and Guido de Croon. Visual model-predictive localization for computationally efficient autonomous racing of a 72-gram drone. *ArXiv*, abs/1905.10110, 2019.
- [11] Sunggoo Jung, Sunyou Hwang, Heemin Shin, and David Hyunchul Shim. Perception, guidance, and navigation for indoor autonomous drone racing using deep learning. *IEEE Robotics and Automation Letters*, 3(3): 2539–2544, July 2018. doi: 10.1109/LRA.2018.2808368.
- [12] Elia Kaufmann, Mathias Gehrig, Philipp Foehn, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. Beauty and the beast: Optimal methods meet learning for drone racing. *2019 International Conference on Robotics and Automation (ICRA)*, pages 690–696, 2018.
- [13] Steven M LaValle. *Planning algorithms*. Cambridge university press, 2006.
- [14] Winter Guerra, Ezra Tal, Varun Murali, Gilhyun Ryou, and Sertac Karaman. Flightgoggles: Photorealistic sensor simulation for perception-driven robotics using photogrammetry and virtual reality. *CoRR*, abs/1905.11377, 2019. URL <http://arxiv.org/abs/1905.11377>.
- [15] Jean-Marie Kai, Guillaume Allibert, Minh-Duc Hua, and Tarek Hamel. Nonlinear feedback control of quadrotors exploiting first-order drag effects. *IFAC-PapersOnLine*, 50(1):8189–8195, 2017.
- [16] Malcolm D. Shuster. Survey of attitude representations. *Journal of the Astronautical Sciences*, 41(4):439–517, Oct 1993.
- [17] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7291–7299, 2017.
- [18] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [19] Rudolf E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- [20] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Transactions on Robotics*, 33(1):1–21, 2017. doi: 10.1109/TRO.2016.2597321.
- [21] Richard Szeliski. *Computer Vision: Algorithms and Applications*. Texts in Computer Science. Springer, 2010. ISBN 9781848829343.
- [22] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.
- [23] Helmut Maurer. On optimal control problems with bounded state variables and control appearing linearly. *SIAM Journal on Control and Optimization*, 15(3):345–362, 1977.
- [24] Dario Brescianini and Raffaello D’Andrea. Tilt-prioritized quadrocopter attitude control. *IEEE Transactions on Control Systems Technology*, 2018.