

GapFlyt: Active Vision Based Minimalist Structure-less Gap Detection For Quadrotor Flight

Nitin J. Sanket*, Chahat Deep Singh*, Kanishka Ganguly, Cornelia Fermüller, Yiannis Aloimonos

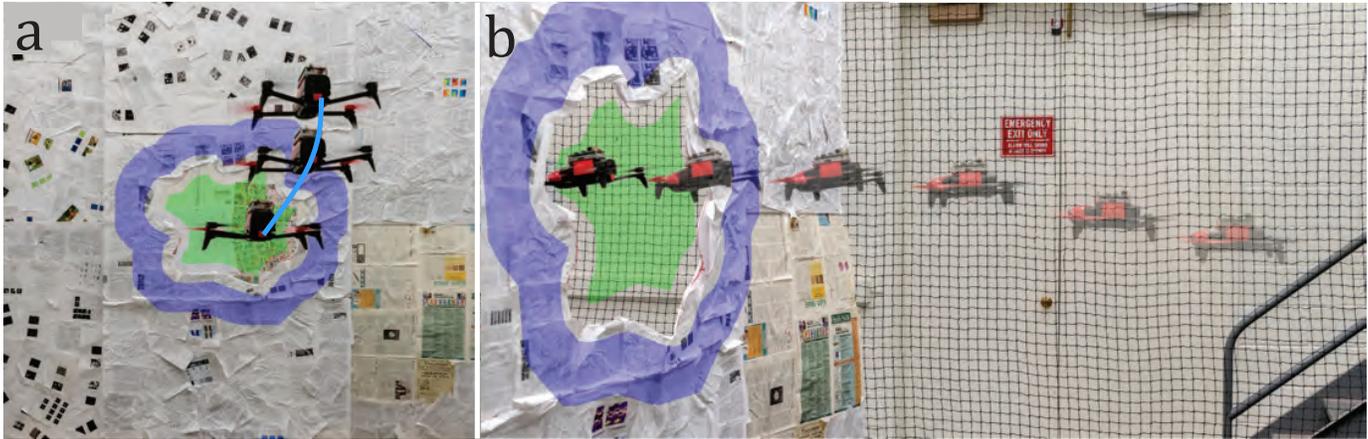


Figure 1. Different parts of the pipeline: (a) Detection of the unknown gap using active vision and TS²P algorithm (cyan highlight shows the path followed for obtaining multiple images for detection), (b) Sequence of quadrotor passing through the unknown gap using visual servoing based control. (The blue and green highlights represent the tracked foreground and background regions respectively. All the images in this paper are best viewed in color.)

Abstract—Although quadrotors, and aerial robots in general, are inherently active agents, their perceptual capabilities in literature so far have been mostly passive in nature. Researchers and practitioners today use traditional computer vision algorithms with the aim of building a representation of general applicability: a 3D reconstruction of the scene. Using this representation, planning tasks are constructed and accomplished to allow the quadrotor to demonstrate autonomous behavior. These methods are inefficient as they are not task driven and such methodologies are not utilized by flying insects and birds. Such agents have been solving the problem of navigation and complex control for ages without the need to build a 3D map and are highly task driven.

In this paper, we propose this framework of bio-inspired perceptual design for quadrotors. We use this philosophy to design a minimalist sensorimotor framework for a quadrotor to fly through unknown gaps without an explicit 3D reconstruction of the scene using only a monocular camera and onboard sensing. We successfully evaluate and demonstrate the proposed approach in many real-world experiments with different settings and window shapes, achieving a success rate of 85% at 2.5ms⁻¹ even with a minimum tolerance of just 5cm. To our knowledge, this is the first paper which addresses the problem of gap detection of an unknown shape and location with a monocular camera and onboard sensing.

Index Terms—Active Vision, Gap Detection, Quadrotor, Visual Servoing, Deep Learning in Robotics and Automation, Optical Flow, Tracking, Collision Avoidance, Computer Vision for

Automation, Aerial Systems: Perception and Autonomy, Visual Tracking and Visual-Based Navigation.

SUPPLEMENTARY MATERIAL

The supplementary hardware tutorial, appendix, code and video are available at prg.cs.umd.edu/GapFlyt.html.

I. INTRODUCTION AND PHILOSOPHY

THE quest to develop intelligent and autonomous quadrotors [1]–[3] has taken a center stage in the recent years due to their usefulness in aiding safety and intuitive control. To achieve any form of autonomy, quadrotors need to have perceptual capabilities in order to sense the world and react accordingly. A generic and fairly common solution to providing perceptual capabilities is to acquire a 3D model of its environment. Creating such a model is very useful because of its general applicability – one could accomplish many tasks using the same model. The process of obtaining a 3D model of the environment using onboard sensing and a myriad of different sensors has gained momentum in the last few years [4]. Sensors like the LIDAR, RGB-D and stereo camera cannot be used on a small quadrotor due to their size, weight, and power limitations. This constrains us to a monocular camera along with the already present onboard inertial sensor (IMU) and many algorithms have been developed based on the same principle [5] [6].

Instead of attempting to recover a 3D model of the scene, we want to recover a “minimal” amount of information that is sufficient to complete the task under consideration. We conceptualize an autonomous quadrotor as a collection of processes that allow the agent to perform a number of behaviors (or tasks) (Table I). At the bottom of the hierarchy is the task of kinetic stabilization (or egomotion). Next comes the ability for obstacle avoidance, where the obstacles could

Manuscript received: February, 15, 2018; Revised April, 26, 2018; Accepted May, 20, 2018. This paper was recommended for publication by Editor Jonathan Roberts upon evaluation of the the Associate Editor and Reviewers’ comments. This work was partly funded by the Brin Family Foundation, Northrop Grumman Corporation and National Science Foundation under grants SMA 1540917 and CNS 1544797 respectively. The support of ONR under award N00014-17-1-2622 is gratefully acknowledged. (Corresponding author: Nitin J. Sanket).

*Equal contribution. All authors are associated with University of Maryland Institute for Advanced Computer Studies, College Park. Emails: {nitin, chahat, kganguly, fer, yiannis}@umiacs.umd.edu

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

Table I
MINIMALIST DESIGN OF AUTONOMOUS QUADROTOR (DRONE) BEHAVIOURS.

Competence	Passive Approach	Active and Task-based Approach
Kinetic stabilization	Optimization of optical flow fields	Sensor fusion between optical flow and IMU measurements
Obstacle avoidance	Obtain 3D model and plan accordingly	Obtain flow fields and extract relevant information from them
Segmentation of independently moving objects	Optimization of flow fields	Fixation and tracking allows detection
Homing	Application of SLAM	Learn paths to home from many locations
Landing	Reconstruct 3D model and plan accordingly	Perform servoing of landing area and plan appropriate policy
Pursuit and Avoidance	Reconstruct 3D model and plan accordingly	Track while in motion
Integration: Switching between behaviors	Easy: The planner interacts with the 3D model	Hard: An attention mechanism on ideas switching between behaviors

be static or dynamic, followed by ability to perform homing, i.e., to find a specific location in an environment. Last in the hierarchy comes the ability to land (on a static or a dynamic surface) and the ability to pursue, or escape from, another agent. This hierarchy of competences, with each competence needing the one before it, constitutes the sensorimotor system of the quadrotor agent. These behaviors can be accomplished without an explicit 3D reconstruction because of the “active” nature of the quadrotor. The quadrotor can perform maneuvers and control the image acquisition process, thus introducing new constraints that were not there before - this is called “active” vision [7]–[9]. This methodology was inspired by the fruit fly [10]. Prior research has shown that fruit flies, and other insects [11] [12], do not perceive depth directly. It is achieved by a balance of active and exploratory procedures. In this paper, we focus on the second competence of obstacle avoidance. Specifically, the question this paper deals with is: “Can a quadrotor manage to go through an arbitrarily shaped gap without building an explicit 3D model of a scene, using only a monocular camera?” We develop the visual mathematics of the solution, evaluate and demonstrate the proposed approach in many real experiments with different settings and window shapes.

Traditional computer vision based approaches such as sparse or dense reconstruction [13]–[16] have been used to obtain a 3D structure of the scene over which sophisticated path planners have been used to plan collision free paths. Lately, deep-learning driven approaches have taken a center stage in solving the problem of fast collision avoidance and safe planning on quadrotors [17]. Most of these neural network approaches compute Fast Fourier Transforms (FFTs) for large filter sizes [18]. Such approaches can be processed on a Field-Programmable Gate Array (FPGA), rather than a Graphical Processing Unit (GPU) to drastically improve the computation performance and power efficiency [19] [20].

To our knowledge, this is the first paper which addresses the problem of gap detection with a monocular camera and onboard sensing. However, the problem of going through gaps has fascinated researchers from many years and some of the recent works which present algorithms for planning and control can be found in [21] [22]. Some of the works which paved way to the bio-inspired approach used in this paper can be found in [23]–[26].

The key contributions of this paper are given below:

- Active vision based structure-less minimalist gap detection algorithm – Temporally Stacked Spatial Parallax or TS²P (Fig. 1a).
- Visual servoing on a contour for the quadrotor to fly through unknown gaps (Fig. 1b).

A. Organization of the paper:

Sec. II presents the detection algorithm of a gap in an arbitrary shaped window using Deep Learning based optical

flow and the role of active vision in such algorithms. Sec. III describes the algorithm used for tracking the gap/safe point and the quadrotor control policy. Sec. IV illustrates the experimental setup and provides error analyses and comparisons with traditional approaches. We finally conclude the paper in Sec. V with parting thoughts on future work.

B. Problem Formulation and Proposed Solutions

A quadrotor is present in a static scene (Fig. 2), where the absolute depth of each point as ‘seen’ from the camera can be modelled as an univariate bimodal gaussian distribution. The location at which there is a maximum spatial depth discrepancy between pixels (projected point) is defined as the Contour (C) of the opening. In this paper, the words boundary, gap, window or opening refer to the same entity and will be used interchangeably. Any pixel close to the mode corresponding to the lower depth value is defined as the Foreground (\mathcal{F}) and similarly that corresponding to the higher depth value is defined as the Background (\mathcal{B}). A simple way of solving the problem of finding the gap for a near fronto-parallel pose is to find the depth discrepancy which is a trivial clustering problem when the depth is known. The depth can be known if multiple calibrated cameras are used or the entire scene is reconstructed in 3D. These methods are computationally expensive [4]. In this paper, we propose a ‘minimalist’ solution to find any arbitrary shaped gap for the quadrotor to go through using Temporally Stacked Spatial Parallax (TS²P) algorithm. A control policy based on contour visual servoing is used to fly through unknown gaps.

II. GAP DETECTION USING TS²P

Before we explain the procedure for gap detection, we need to formally define the notation used. (${}^aX_b, {}^aY_b, {}^aZ_b$) denotes the coordinate frame of b represented in the reference of a . The letters I, C, B and W are used as sub/superscript to denote quantities related to Inertial Measurement Unit (IMU), Camera, Body and World respectively. If a sub/superscript is omitted, the quantity can be assumed to be in W . C and I are assumed to be rigidly attached to each other with known

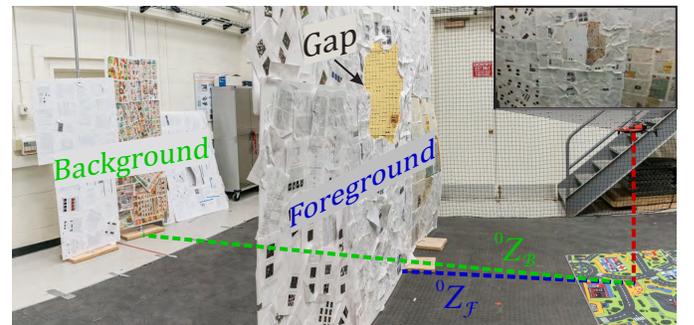


Figure 2. Components of the environment. On-set Image: Quadrotor view of the scene.

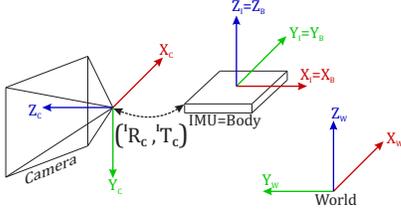


Figure 3. Representation of co-ordinate frames.

intrinsic and extrinsic. B is defined to be aligned with I (Fig. 3). A pinhole camera model is used for the formation of the image. The world point \mathbf{X} gets projected onto the image plane point \mathbf{x} .

The camera captures images/frames at a certain frequency. Let the frame at time t^i be denoted as \mathbb{F}^i and is called the i^{th} frame. Optical flow [27] between i^{th} and j^{th} frames is denoted by ${}^j_i\dot{p}_{\mathbf{x}}$ which is a sum of both translational (${}^j_i\dot{p}_{\mathbf{x},T}$) and rotational (${}^j_i\dot{p}_{\mathbf{x},R}$) components and are given by:

$${}^j_i\dot{p}_{\mathbf{x},T} = \frac{1}{Z_{\mathbf{x}}} \begin{bmatrix} xV_z - V_x \\ yV_z - V_y \end{bmatrix}; \quad {}^j_i\dot{p}_{\mathbf{x},R} = \begin{bmatrix} xy & -(1+x^2) & y \\ (1+y^2) & -xy & -x \end{bmatrix} \Omega$$

where $\mathbf{x} = [x \ y]^T$ denotes the pixel coordinates in the image plane, $Z_{\mathbf{x}}$ is the depth of \mathbf{X} (world point corresponding to pixel \mathbf{x}), $V = [V_x \ V_y \ V_z]^T$ and Ω are the linear and angular velocities of the C in W between t^i and t^j respectively.

The otherwise complex depth estimation problem can be trivialized by moving in a certain way [7]. These ‘‘active’’ vision principles dictate us to control the quadrotor’s movements so as to make the interpretation of optical flow easier. Since the quadrotor is a differentially flat system [28], the rotation about (X_B, Y_B, Z_B) or roll, pitch and yaw can be decoupled. As an implication, the movements can be controlled in a way to simplify the depth ($Z_{\mathbf{x}}$) estimation problem. The quadrotor is controlled in such a way that $\Omega \approx 0$, $V_z \ll V_x$ and $V_z \ll V_y$, then the optical flow can be modelled as:

$${}^j_i\dot{p}_{\mathbf{x}} = {}^j_iZ_{\mathbf{x}}^{-1} \begin{bmatrix} -{}^j_iV_x & -{}^j_iV_y \end{bmatrix}^T + \eta$$

where η is the approximation error. This shows that using the aforementioned ‘‘active’’ control strategy, we obtain an implicit 3D structure of the environment in the optical flow. The inverse depth in this ‘‘controlled’’ case manifests as a linear function of the optical flow.

The optical flow equation can be written for both foreground (\mathcal{F}) and background (\mathcal{B}) pixels independently. The magnitude of optical flow for \mathcal{F} is given by,

$$\left\| {}^j_i\dot{p}_{\mathbf{x},\mathcal{F}} \right\|_2 = {}^j_iZ_{\mathbf{x},\mathcal{F}}^{-1} \sqrt{{}^j_iV_x^2 + {}^j_iV_y^2} + \nu$$

where $\nu \sim \mathcal{N}(0, \sigma)$ is assumed to be an additive white Gaussian noise and is independent of the scene structure or the amount of camera movement between two frames (V, Ω). For such assumptions to be valid, the optical flow algorithm needs to work well for a wide range of camera movements in a variety of scene structures. Using fast traditional optical flow formulations based on methods like [29] or [30] voids such assumptions. This motivated us to use deep-learning based flow algorithms which excel at this task while maintaining a reasonable speed when running on a GPU. In this paper,

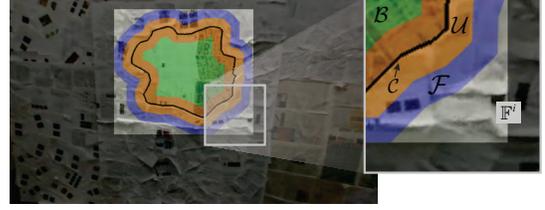


Figure 4. Label sets used in tracking. (blue: foreground region, green: background region, orange: uncertainty region, black line: contour, brighter part of frame: active region and darker part of frame: inactive region.)

FlowNet2 [31] is used to compute optical flow unless stated otherwise.

A simple method to reduce noise is to compute the mean of the flow magnitudes across a few frames. Let $\xi = \{\mathbb{F}^j, \dots, \mathbb{F}^k\}$ be a set of N frames from which the optical flow is computed with respect to some reference frame \mathbb{F}^i where the complete gap is assumed to be visible. Here, $N = \bar{\xi}$ is the cardinality of the set ξ . The mean flow magnitude at \mathbf{x} for \mathcal{F} ($\left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{F}} \right\|_2$) and \mathcal{B} ($\left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{B}} \right\|_2$) is given by,

$$\left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{F}/\mathcal{B}} \right\|_2 = (NZ_{\mathbf{x},\mathcal{F}/\mathcal{B}})^{-1} \sum_{r=j}^k {}^r_iV + \nu'$$

where $\nu' \sim \mathcal{N}(0, N^{-0.5}\sigma)$ and $V = \sqrt{V_x^2 + V_y^2}$. Clearly, the noise varies inversely with N .

Since $Z_{\mathbf{x},\mathcal{F}} < Z_{\mathbf{x},\mathcal{B}}$, we can say that $\left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{F}} \right\|_2 > \left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{B}} \right\|_2$. Now, $\left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{F}} \right\|_2 - \left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{B}} \right\|_2 \geq \tau$ can be used as a criterion for finding possible boundary regions. It was found experimentally that using inverse flow differences $\left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{B}} \right\|_2^{-1} - \left\| \xi_i\dot{p}_{\mathbf{x},\mathcal{F}} \right\|_2^{-1} \geq \tau'$ gave better numerical stability and better noise performance due to the scale compression by the inverse function. This is inherently the spatial derivative of inverse average (stacked) flow magnitudes and it can be written as $\Xi = \nabla \cdot \left\| \xi_i\dot{p}_{\mathbf{x}} \right\|_2^{-1}$, where, $\nabla = [\partial/\partial x \ \partial/\partial y]^T$. Note that this is the same as solving the edge detection problem in computer vision and any kernel or method like the Sobel operator or the Canny edge detection algorithm can be used.

III. HIGH SPEED GAP TRACKING FOR VISUAL SERVOING BASED CONTROL

This section presents a targeted solution for tracking a contour using label sets propagated using Focus of Expansion (FOE) constraints. A pixel at location \mathbf{x} is associated with a score $\chi(\mathbf{x}) \in [-1, 1]$ which denotes its score as foreground or background. The foreground and background pixel locations are defined by $\mathcal{F} = \{\mathbf{x} | \chi(\mathbf{x}) = +1\}$ and $\mathcal{B} = \{\mathbf{x} | \chi(\mathbf{x}) = -1\}$ respectively.

We define the opening \mathcal{O} on the image plane as $\mathcal{O} = \{\mathbf{x} | \chi(\mathbf{x}) < 0\}$. The pixel locations which cannot be perfectly classified as foreground or background belong to the uncertainty zone and are defined as $\mathcal{U} = \{\mathbf{x} | \chi(\mathbf{x}) \in (-1, +1)\}$. The contour location is defined as $\mathcal{C} = \{\mathbf{x} | \chi(\mathbf{x}) = 0\}$. Fig. 4 gives a visual representation of the different sets on a sample image.

The problem statement dictates the tracking of contour location \mathcal{C} across time. This problem is hard as the contour tracking relies on updating $\chi(\mathbf{x}) \forall \mathbf{x} \in \mathcal{U}$ over time which

Algorithm 1: Label Set Propagation using FOE constraints.

Data: $\mathbb{C}_{\mathcal{F}}^i, \mathbb{C}_{\mathcal{B}}^i, \mathbb{F}^i, \mathbb{F}^j, \mathcal{F}^i, \mathcal{B}^i$
Result: $\mathbb{C}_{\mathcal{F}}^j, \mathbb{C}_{\mathcal{B}}^j, \mathcal{F}^j, \mathcal{B}^j$

- 1 $\mathbb{C}_{\mathcal{F}}^j = \text{FeatureTracker}(\mathbb{C}_{\mathcal{F}}^i, \mathbb{F}^i, \mathbb{F}^j)$;
 - 2 $A = \begin{bmatrix} \mathbb{C}_{\mathcal{F},x}^i & -\mathbf{1} \end{bmatrix}$;
 - 3 $B = \begin{bmatrix} \mathbb{C}_{\mathcal{F},x}^j - \mathbb{C}_{\mathcal{F},x}^i \end{bmatrix}$;
 - 4 $\begin{bmatrix} \alpha \\ \beta \end{bmatrix} = A^\dagger B$;
 - 5 $x_{0,\mathcal{F}} = \beta/\alpha$; $y_{0,\mathcal{F}} = \left\langle \mathbb{C}_{\mathcal{F},y}^j, \frac{1}{\alpha} (\mathbb{C}_{\mathcal{F},y}^j - \mathbb{C}_{\mathcal{F},y}^i) \right\rangle$;
 - 6 $\mathcal{F}^j = \alpha \begin{bmatrix} \mathcal{F}_x^i - x_{0,\mathcal{F}} \\ \mathcal{F}_y^i - y_{0,\mathcal{F}} \end{bmatrix} + \begin{bmatrix} \mathcal{F}_x^j \\ \mathcal{F}_y^j \end{bmatrix}$;
 - 7 Repeat steps 1 through 6 for \mathcal{B} in parallel;
-

is non-trivial and computationally expensive. To simplify the problem, we use the dual formulation of the problem which is to track the pixel locations which belong to the set defined by $\{\mathbf{x} \notin \mathcal{U}\} = \mathcal{F} \cup \mathcal{B}$. This enables us to track the contour indirectly at high speeds as corner tracking is comparatively faster [32]. The trade-off in using the dual formulation is that we don't obtain the actual contour location across time - which might be needed for aggressive maneuvers, but this is not dealt in the scope of this paper. The label set propagation algorithm is described in Algorithm 1. Here, $\mathbb{C}_{\mathcal{F}}^i$ and $\mathbb{C}_{\mathcal{B}}^i$ represents a set of corner/feature points for the foreground and background pixels respectively in \mathbb{F}^i . A^\dagger denotes the pseudo inverse of the matrix A and $[x_{0,\mathcal{F}} \ y_{0,\mathcal{F}}]^T$ is the FOE for the foreground. Intuitively, we solve the linear equations of the horizontal flow field to obtain the divergence/time-to-contact α . The divergence in x-direction is used to also predict the y-coordinates.

A. Safe Point Computation and Tracking

We can assume that the real-world points corresponding to the background $\mathcal{B} \cup \mathbb{U}_B$ are far enough that we can approximate them to lie on a plane. The foreground points under consideration $\mathcal{F} \cup \mathbb{U}_F$ occupy a small area around the contour which can be assumed to be planar. Here, $\mathbb{U}_F \subset \mathcal{U}$ and $\mathbb{U}_B \subset \mathcal{U}$ are the sets which actually belong to the foreground and background respectively.

Now, the quadrotor can be represented as an ellipsoid with semi-axes of lengths a, b and c . As an implication of the above assumptions, the projection of the quadrotor on the window at any instant is an ellipse. Let us define the projection of the quadrotor on the image as \mathcal{Q} . The largest \mathcal{Q} can be written in terms of matrix equation of the ellipse centered at $[h, k]^T$ defined as $Q(h, k, R_\theta) = 0$.

Here, R_θ is a two-dimensional rotation matrix and θ is the angle the largest semi-axis of the ellipse makes with the X_C axis. The projection of the quadrotor on the image is given by $\mathcal{Q} = \{\mathbf{x} | Q(\mathbf{x}) \leq 0\}$. The safe region \mathcal{S} can be computed as

$$\mathcal{S} = \bigcup_{\forall \theta} \mathcal{O} \ominus \mathcal{Q}$$

where \ominus denotes the Minkowski difference of sets. Now, we define the 'safest point' (\mathbf{x}_s) as the barycentric mean of \mathcal{S} .

Remark. The above optimization problem can only be solved using convex optimization with a guaranteed global solution

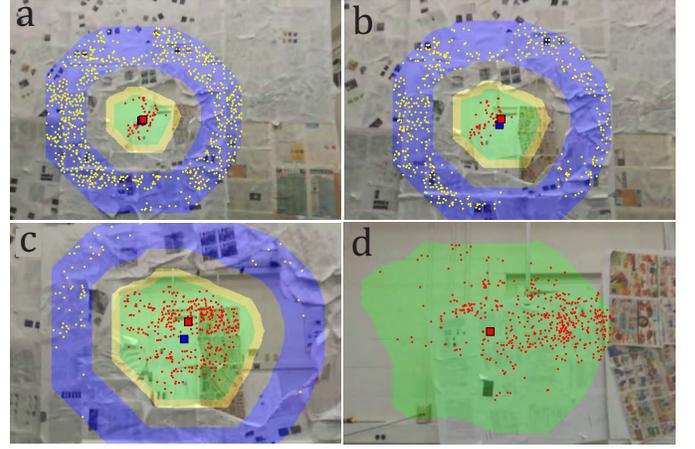


Figure 5. Tracking of \mathcal{F} and \mathcal{B} across frames. (a) shows tracking when $\mathbb{C}_{\mathcal{F}}^i > k_{\mathcal{F}}$ and $\mathbb{C}_{\mathcal{B}}^i > k_{\mathcal{B}}$. (b) When $\mathbb{C}_{\mathcal{B}}^i \leq k_{\mathcal{B}}$, the tracking for \mathcal{B} will be reset. (c) When $\mathbb{C}_{\mathcal{F}}^i \leq k_{\mathcal{F}}$, the tracking for \mathcal{F} will be reset. (d) shows tracking only with \mathcal{B} , when $\mathcal{F} = \emptyset$. (blue: \mathcal{F} , green: \mathcal{B} , yellow: \mathcal{O} , yellow dots: $\mathbb{C}_{\mathcal{F}}^i$, red dots: $\mathbb{C}_{\mathcal{B}}^i$, blue Square: $\mathbf{x}_{s,\mathcal{F}}$, red Square: $\mathbf{x}_{s,\mathcal{B}}$.)

when both \mathcal{O} and \mathcal{Q} are convex sets. A conservative solution to the above problem is fitting the largest scaled version of \mathcal{Q} inside \mathcal{O} when \mathcal{O} is a non-convex set and \mathcal{Q} is a convex set.

Note that as \mathcal{Q} is a chosen model, it can always be chosen to be a convex set, i.e., convex hull of the non-convex set. Also, from the above remark, the 'safest point' (\mathbf{x}_s) can be defined as the center of the largest ellipse which can be fit inside \mathcal{S} whilst maintaining the eccentricity equal to that defined by \mathcal{Q} . The optimization problem becomes,

$$\operatorname{argmax}_{a,\theta} \mathcal{S} \cap \mathcal{Q} \text{ s.t. } \mathcal{Q} \subseteq \mathcal{S} \text{ and } |\theta| \leq \theta_{max}$$

This problem can be solved using the procedure described in [33]. However, a desirable property of the safe region is that it should favor less-aggressive maneuvers. This can be modelled as a regularization penalty in the above formulation,

$$\operatorname{argmax}_{a,\theta} \mathcal{S} \cap \mathcal{Q} + \lambda \theta \text{ s.t. } \mathcal{Q} \subseteq \mathcal{S} \text{ and } |\theta| \leq \theta_{max}$$

Solving the above optimization problem is computationally intensive and not-possible without the prior knowledge of the scale/depth. For obtaining the minimalist solution, we assume that *the gap is large enough for the quadrotor to pass through* and replace the above optimization problem by an empirically chosen approximation. A simple and efficient safe point computation can be performed as the median of the convex set \mathcal{O} and is given by $\mathbf{x}_s \approx \operatorname{argmin}_{\mathbf{x}} \sum_{\forall o \in \mathcal{O}} \|o - \mathbf{x}\|_2$.

Remark. If the above approximation is used when \mathcal{O} is non-convex, the amount of deviation from the 'actual' safe point is a function of $\overline{\operatorname{Conv}(\mathcal{O})}/\overline{\mathcal{O}}$. Here $\operatorname{Conv}(\mathcal{O})$ is the convex hull of \mathcal{O} .

Keen readers would note that the formulation for the safe point is in-terms of \mathcal{O} which we wanted to avoid tracking in the first place. Indeed this is true and a simple solution takes care of this. Because we are propagating \mathcal{F} and \mathcal{B} with FOE constraints, the cross-ratios of $\{\|o, f\|_2 | o \in \mathcal{O}, f \in \mathcal{F}\}$ and $\{\|o, b\|_2 | o \in \mathcal{O}, b \in \mathcal{B}\}$ are preserved. Here, \mathcal{F} and \mathcal{B} are computed from the detected opening \mathcal{O} as follows: $\mathcal{F} = \{\mathcal{O} \oplus \epsilon_1 - \mathcal{O} \oplus \epsilon_2\}$ and $\mathcal{B} = \{\mathcal{O} \ominus \epsilon_3\}$. Here, ϵ_i is a user chosen kernel (circular in this paper).

The \mathcal{F} and \mathcal{B} are propagated from $\overline{\mathbb{F}}^i$ onwards where the detection was performed. The ‘safest point’ (\mathbf{x}_s) is computed as follows:

$$\mathbf{x}_s = \begin{cases} \mathbf{x}_{s,\mathcal{F}}, & \overline{\mathcal{F}} \geq \overline{\mathcal{B}} \\ \mathbf{x}_{s,\mathcal{B}}, & \text{otherwise} \end{cases}$$

where $\mathbf{x}_{s,\mathcal{F}}$ and $\mathbf{x}_{s,\mathcal{B}}$ are the safest points computed individually for \mathcal{F} and \mathcal{B} by taking their median respectively. Fig. 5 shows the tracking of \mathcal{F} and \mathcal{B} across time and how the tracking algorithm actively switches between $\overline{\mathbf{x}}_{s,\mathcal{F}}$ and $\overline{\mathbf{x}}_{s,\mathcal{B}}$ for the safest point \mathbf{x}_s . When $\overline{C}_{\mathcal{F}}^i \leq k_{\mathcal{F}}$ or $\overline{C}_{\mathcal{B}}^i \leq k_{\mathcal{B}}$, the tracker/feature points are reset in the current \mathcal{F} and \mathcal{B} sets respectively. Here, $k_{\mathcal{F}}$ and $k_{\mathcal{B}}$ are empirically chosen thresholds. For all experiments $k_{\mathcal{F}} = 40$ and $k_{\mathcal{B}} = 20$. Due to resetting and active switching, \mathbf{x}_s can jump around making the control hard, hence a simple Kalman filter with a forward motion model is used to smooth out the value of \mathbf{x}_s . From here on, safe point refers to the safest point \mathbf{x}_s .

B. Control Policy

We propose a control policy such that it follows the tracked \mathbf{x}_s . The quadrotor follows the dynamic model as given in [34]. The controller uses the traditional backstepping approach based on [34] and contains the following loops: Inner loop and outer loop controllers. Inner loop controls the attitude stability while the outer loop controller is responsible for the quadrotor position. It is important to note that frames are transformed from C to B .

Since the quadrotor is differentially flat, the altitude Z_B can be controlled independently from X_B and Y_B [28]. The control policy is to align the projection of the body center on the image plane with \mathbf{x}_s . The difference between the two centers is called the error e . The x and y component of the error e can be minimized by varying roll (ϕ) and net thrust (u_1) respectively. A simple Proportional-Integral-Derivative (PID) controller on e is used. This control policy only deals with the alignment of B to \mathbf{x}_s and does not deal with moving forward (Z_C). To move forward, the quadrotor pitch (ϕ) needs to be controlled. The rate of forward motion is controlled by the pitch angle θ_0 which is empirically chosen. The bigger the value of θ_0 the faster the quadrotor will fly towards the gap. It is important to note the implicit assumption made in this paper that the gap is large enough for the quadrotor to go through safely.

IV. EXPERIMENTS

A. Experimental Setup

The proposed framework was tested on a modified hobby quadrotor, Parrot[®] Bebop 2, for its cost effectiveness and ease of use. The Bebop 2 is equipped with a front facing camera, a 9-axis IMU and a downward facing optical flow sensor coupled with a sonar. The Parrot[®] Bebop 2 allows only high level controls in terms of body frame velocities using ROS. An NVIDIA Jetson TX2 GPU is mounted on the Bebop 2 as shown in Fig. 6 and is used to run all the perception and control algorithms onboard. The TX2 and Bebop 2 communicate via a WiFi connection, where the images are received at 30Hz. The overall weight of the flight setup is 680g with the dimensions being $32.8 \times 38.2 \times 12\text{cm}$.

All the experiments were prototyped on a PC running Ubuntu 16.04 with an Intel[®] Core i7 6850K 3.6GHz CPU,

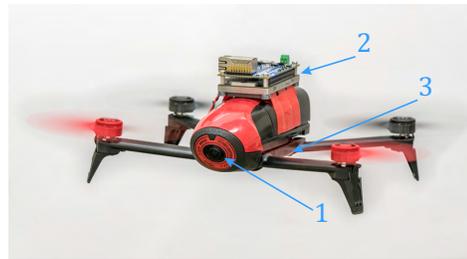


Figure 6. The platform used for experiments. (1) The front facing camera, (2) NVIDIA TX2 CPU+GPU, (3) Downward facing optical flow sensor (camera+sonar) which is only used for position hold.

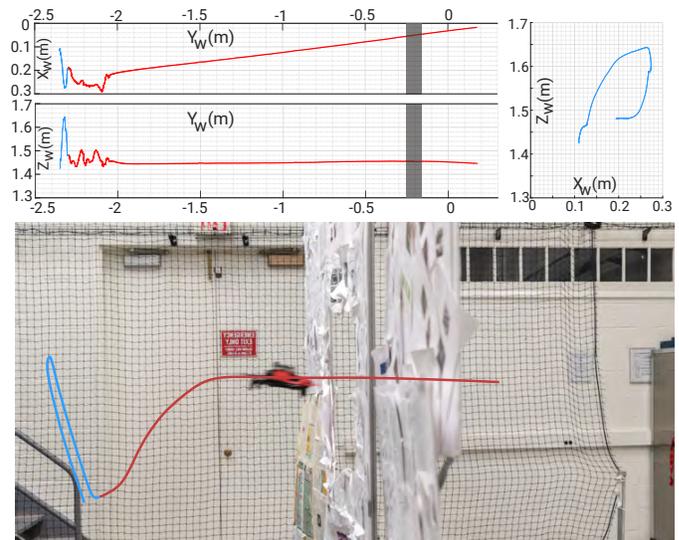


Figure 7. First two rows: (X_W, Y_W) , (Y_W, Z_W) and (X_W, Z_W) Vicon estimates of the trajectory executed by the quadrotor in different stages (gray bar indicates the gap). (X_W, Z_W) plot shows the diagonal scanning trajectory (the lines don't coincide due to drift). Last row: Photo of the quadrotor during gap traversal. (cyan: detection stage, red: traversal stage.)

an NVIDIA Titan-Xp GPU and 64GB of RAM in MATLAB using the Robotics Toolbox. The deep learning based optical flow runs on Python with TensorFlow back-end. All the final versions of the software were ported to Python to run on the NVIDIA Jetson TX2 running Linux for Tegra[®] (L4T) 28.2. A step-by-step tutorial on using Bebop 2 as a research platform is available at prg.cs.umd.edu/GapFlyt.html.

The environmental setup for the experiments consists of a rigid scene which has two near-planar structures, one for the foreground and the other for the background. As shown in Fig. 2, let us denote the initial perpendicular distance between the quadrotor body center and the foreground as ${}^0Z_{\mathcal{F}}$ and the background as ${}^0Z_{\mathcal{B}}$. The near-planar structures are made of foam-core with newspapers stuck on them to add texture to the scene. The gap is located near the center of the foreground and is of an arbitrary shape. For the detection of the window, the quadrotor executes a fixed diagonal straight line trajectory in the $X_W - Z_W$ plane as shown in Fig. 7 while taking a number of images along its path. The number of images used for the detection stage is a parameter denoted by N . Once the window is detected, \mathcal{F} and \mathcal{B} are tracked across time in order for quadrotor to go through the gap using visual servoing based control.

B. Experimental Results

The pipeline was evaluated on different variations of the environmental setup. In the first experiment, we test our

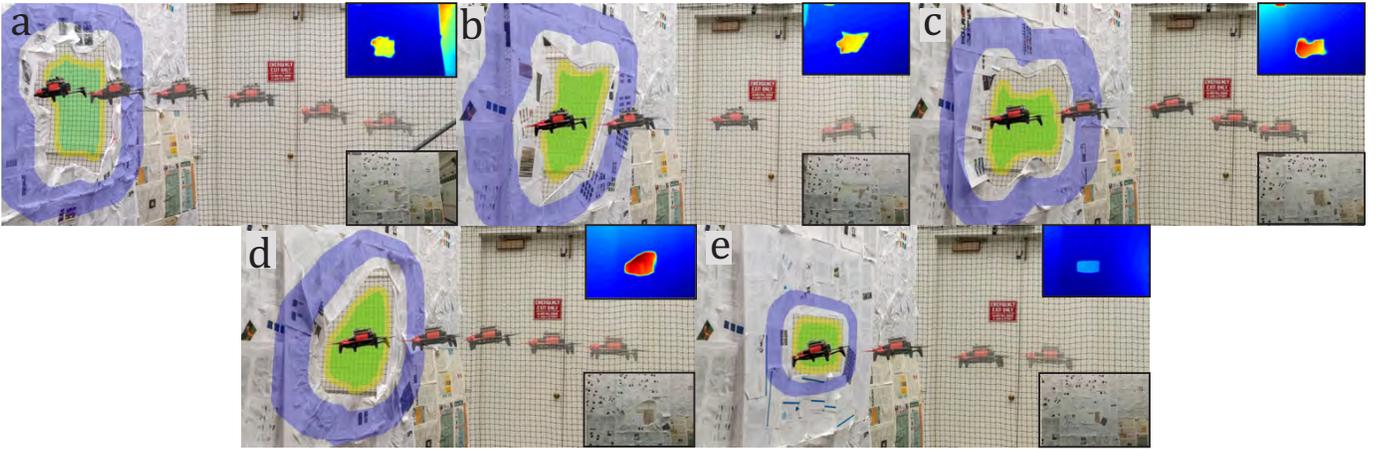


Figure 8. Sequence of images of quadrotor going through different shaped gaps. Top on-set: Ξ outputs, bottom on-set: quadrotor view.

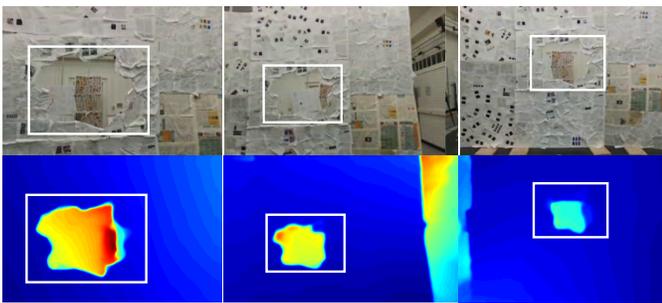


Figure 9. Top Row (left to right): Quadrotor view at ${}^0Z_{\mathcal{F}} = 1.5, 2.6, 3\text{m}$ respectively with ${}^0Z_{\mathcal{B}} = 5.7\text{m}$. Bottom Row: Respective Ξ outputs for $N = 4$. Observe how the fidelity of Ξ reduces as ${}^0Z_{\mathcal{F}} \rightarrow {}^0Z_{\mathcal{B}}$, making the detection more noisy. (white boxes show the location of the gap in Figs. 9 to 13.)

pipeline on five different arbitrarily shaped gaps as shown in Fig. 8. Unless otherwise stated ${}^0Z_{\mathcal{F}} \sim 2.6\text{m}$ and ${}^0Z_{\mathcal{B}} \sim 5.7\text{m}$. The aim here is to find biases in the detection pipeline. The windows were chosen to have a diversity in the geometrical sharpness of the corners, convexity of the shape and the size. As stated earlier, the only constraint imposed on the gaps is that they are large enough to go through with the quadrotor pitch angle close to zero and near-convex. The outputs of TS²P algorithm for different windows are shown in Fig. 8 with $N = 4$ along with their inverse average (stacked) flow magnitudes Ξ , illustrating that our detection algorithm is independent of shape and size of the opening. A canny edge detector is run on Ξ followed by morphological operations to obtain \mathcal{C} .

The second experiment is designed to test the noise sensitivity of TS²P. The intuition is that as ${}^0Z_{\mathcal{F}} \rightarrow {}^0Z_{\mathcal{B}}$, noisier the detection result. The outputs for different ${}^0Z_{\mathcal{F}}$ and ${}^0Z_{\mathcal{B}}$ are shown in Fig. 9 when $N = 4$. This is because the fidelity of Ξ becomes less and is more prone to noise. By increasing N the noise gets averaged out across frames improving the fidelity of Ξ .

In the third experiment, we present detection outputs for different values of N , image baselines and image sizes. The effect of N has been already discussed previously. Having a very small baseline results in effectively dropping the value of N and vice-versa. The results from different sized images as illustrated in Fig. 12 show that the detection algorithm can work even on a very small quadrotor which can only carry a very low-resolution camera (as low as 32×48 pixels). Our

algorithm can also handle dynamic noises very well though being modelled as a gaussian for the discussion. However, one can notice that the results improve significantly with increase in N (Fig. 13) demonstrating the advantage of TS²P.

Gap detection using TS²P almost always results in an eroded version of the true gap. This is good for safe maneuvers like the one considered in this paper. However, aggressive flight algorithms might suffer due to conservative results. This can be mitigated by tuning the values of image baselines, N and the trajectory of the quadrotor to obtain minimum erosion results. Tuning these parameters is easy when a prior about the scene is known or the flow algorithms are so fast that one can actively change the detection trajectory so as to maximize the coverage on the part of the contour least ‘seen’. The dynamic choice of these parameters comes into the scope of our future work.

In the last experiment we present alternative approaches including state-of-the-art methods which can be used to find the gap. The methods can be subdivided into structure based approaches and structureless approaches. The structure based approaches can be defined as the set of approaches where a full 3D reconstruction of the scene is computed, whereas, structureless approaches do not. The structure based approaches presented are DSO [16] – Direct Sparse Odometry, depth from hardware stereo cameras [15] and Stereo SLAM – Simultaneous Localization and mapping using Stereo Cameras and IMU [15]. The data for the structured approaches were collected using a Parrot[®] SLAMDunk [15]. The structureless approaches presented are MonoDepth [35] – deep learning based monocular depth estimation and the proposed TS²P on two different deep learning based dense optical flow algorithms, namely, FlowNet2 [31], SpyNet [36] and DIS [37]. Table II shows the comparison of the stated methods averaged over 150 trials.

Fig. 10 compares the results of DSO, stereo depth, MonoDepth and our method (TS²P) with the ground truth. It can be inferred that the MonoDepth results are extremely noisy (even with different models) making it impossible to detect the gap as the images in our paper were never ‘seen’ during training. Note that we don’t retrain or finetune any of the deep learning models in this paper. Retraining MonoDepth and other deep learning based methods used in this paper on our dataset might lead to better results. Whereas, DSO and stereo depth results can be used to detect the opening with some filtering. Stereo SLAM and DSO are slow in the map building

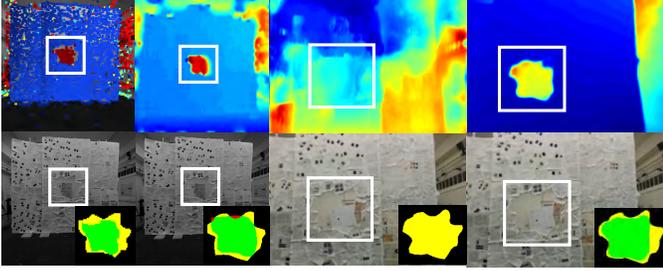


Figure 10. Comparison of different philosophies to gap detection. Top row (left to right): DSO, Stereo Depth, MonoDepth, TS²P. Bottom row shows the detected gap overlaid on the corresponding input image. (green: $\mathcal{G} \cap \mathcal{O}$, yellow: false negative $\mathcal{G}' \cap \mathcal{O}'$, red: false positive $\mathcal{G}' \cap \mathcal{O}$.)

stage (taking about 6s and 12s respectively), however, once the map is built and the algorithms are in the localization stage the depth (or scaled depth) are obtained at 20Hz. The Stereo SLAM and Stereo Depth were run on the SLAMDunk with an NVIDIA Jetson TK1 processor which is much slower than the NVIDIA Jetson TX2 processor used for running DSO and other methods.

Fig. 11 compares different optical flow methods used for TS²P. Though SpyNet and DIS optical flow are faster, FlowNet2 outputs significantly better results at the edges which is important for obtaining a good gap detection – this can be observed by looking at Ξ for each algorithm.

After the gap detection has been performed, \mathcal{F} and \mathcal{B} are computed from the detected gap \mathcal{C} . Fig. 5 shows \mathcal{F} and \mathcal{B} being propagated across time as the quadrotor is in pursuit of going through the gap with the update of \mathbf{x}_s . A comparison of tracking using different methods are given in Table III. Clearly, KLT outperforms all other methods with a theoretical maximum quadrotor speed of 8 ms^{-1} in the given scene. The theoretical maximum speed is calculated for a global shutter camera in such a way that the motion parallax is constrained within one pixel for the scene with ${}^0Z_{\mathcal{F}} \sim 2.6\text{m}$ and ${}^0Z_{\mathcal{B}} \sim 5.7\text{m}$. The calculation assumes that none of the tracking/matching methods work when the motion blur is more than one pixel. However, most of the methods can work well upto some pixels of motion blur, this will in-turn increase the theoretical maximum speed by this factor. If a rolling shutter camera is used without rolling shutter compensation, the theoretical maximum speed value has to be divided by the factor of blur caused by rolling shutter. We achieved a practical maximum speed of 2.5ms^{-1} in our experiments. We were limited to this speed due to the acceleration constraints on the Bebop 2 and the rolling shutter camera.

We achieved a remarkable success rate of 85% over 150 trials for different arbitrary shaped windows under a wide range of conditions which includes a window with a minimum tolerance of just 5cm (Fig. 14). Success is defined as window detection output \mathcal{O} having at least 75% overlap with the ground truth and traversal through the gap without collision. Failure cases also include the optimization failures and/or feature tracking failures for structure based approaches. For TS²P, we define Detection Rate (DR), Average False Negative (AFN) and Average False Positive (AFP) as follows (AFN and AFP are computed only for successful trails):

$$\text{DR} = \frac{\sum_{k=1}^{\text{Num. Trails}} (\lambda_D^k)}{\text{Num. Trails}}; \lambda_D^k = \left(\frac{\overline{\mathcal{G} \cap \mathcal{O}}}{\overline{\mathcal{G}}} \right)^k \geq 0.75$$

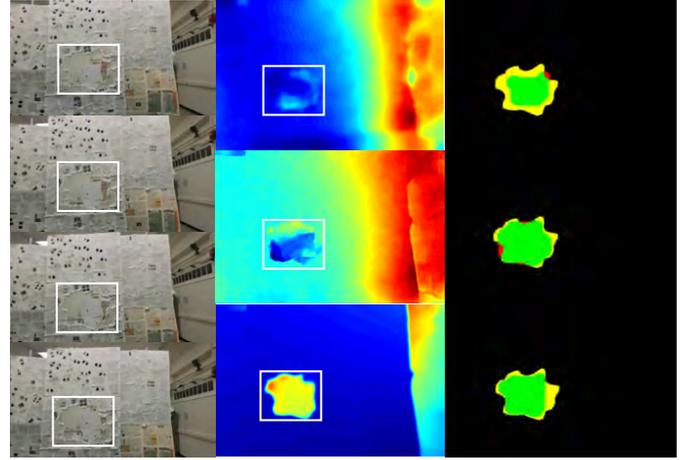


Figure 11. Left Column: Images used to compute Ξ . Middle Column (top to bottom): Ξ outputs for DIS Flow, SpyNet and FlowNet2. Right Column: Gap Detection outputs. (green: $\mathcal{G} \cap \mathcal{O}$, yellow: false negative $\mathcal{G}' \cap \mathcal{O}'$, red: false positive $\mathcal{G}' \cap \mathcal{O}$.)

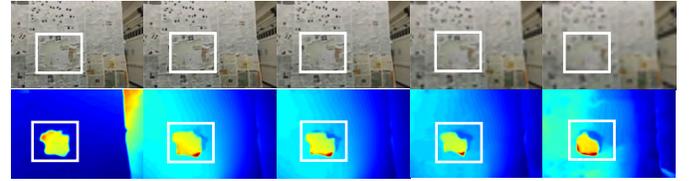


Figure 12. Top row (left to right): Quadrotor view at image sizes of 384×576 , 192×288 , 96×144 , 48×72 , 32×48 . Note all images are re-scaled to 384×576 for better viewing. Bottom row shows the respective Ξ outputs for $N = 4$.

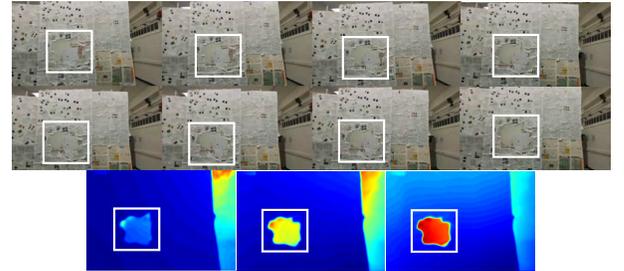


Figure 13. Top two rows show the input images. The third row shows the Ξ outputs when only the first 2, 4 and all 8 images are used.



Figure 14. Quadrotor traversing an unknown window with a minimum tolerance of just 5cm. (red dashed line denotes \mathcal{C} .)

$$\text{AFN} = \frac{\sum_{k=1}^{\text{Num. Succ. Trails}} (\lambda_N^k)}{\text{Num. Succ. Trails}}; \lambda_N^k = \left(\frac{\overline{\mathcal{G}' \cap \mathcal{O}'}}{\overline{\mathcal{G}'}} \right)^k$$

$$\text{AFP} = \frac{\sum_{k=1}^{\text{Num. Succ. Trails}} (\lambda_P^k)}{\text{Num. Succ. Trails}}; \lambda_P^k = \left(\frac{\overline{\mathcal{G}' \cap \mathcal{O}}}{\overline{\mathcal{G}'}} \right)^k$$

where \mathcal{A}' is the negation of the set \mathcal{A} .

Table II
COMPARISON OF DIFFERENT METHODS USED FOR GAP DETECTION.

Method	Sensor(s)	Run Time (Init. Time) in s	DR	AFN	AFP
DSO [16]	Monocular	0.05 (6)	0.88	0.20	0.00
Stereo Depth* [15]	Stereo	0.10	0.90	0.17	0.04
Stereo SLAM* [15]	Stereo + IMU	0.05 (12)	0.91	0.15	0.00
Mono Depth [35]	Monocular	0.97	0.00	–	–
TS ² P (FlowNet2 [31])	Monocular	1.00	0.93	0.14	0.02
TS ² P (SpyNet [36])	Monocular	0.12	0.74	0.16	0.05
TS ² P (DIS [37])	Monocular	0.45	0.62	0.20	0.04

* indicates algorithm tested on NVIDIA TK1 otherwise NVIDIA TX2.

Table III
COMPARISON OF DIFFERENT METHODS USED FOR TRACKING.

Method	Run Time (ms)	Theo. Max. Speed (ms ⁻¹)
GMS [38]	40	0.40
FAST [39] + RANSAC	8.3	1.92
Cuda-SIFT [40] + RANSAC	5	3.20
KLT [32]	2	8.00

V. CONCLUSIONS

We present a minimalist philosophy to mimic insect behaviour to solve complex problems with minimal sensing and active movement to simplify the problem in hand. This philosophy was used to develop a method to find an unknown gap and fly through it using only a monocular camera and onboard sensing. A comprehensive comparison and analysis is provided. To our knowledge, this is the first paper which addresses the problem of gap detection of an unknown shape and location with a monocular camera and onboard sensing. As a parting thought, IMU data can be coupled with the monocular camera to get a scale of the window and plan for aggressive maneuvers.

ACKNOWLEDGEMENT

The authors would like to thank Konstantinos Zampogiannis for helpful discussions and feedback.

REFERENCES

- Teodor Tomic et al. Toward a fully autonomous uav: Research platform for indoor and outdoor urban search and rescue. *IEEE robotics & automation magazine*, 19(3):46–56, 2012.
- Tolga Özaslan et al. Inspection of penstocks and featureless tunnel-like environments using micro UAVs. In *Field and Service Robotics*, pages 123–136. Springer, 2015.
- Nathan Michael et al. Collaborative mapping of an earthquake-damaged building via ground and aerial robots. *Journal of Field Robotics*, 29(5):832–841, 2012.
- Takafumi Taketomi et al. Visual slam algorithms: a survey from 2010 to 2016. *IPSI Transactions on Computer Vision and Applications*, 9(1):16, June 2017.
- Tong Qin and Shaojie Shen. Robust initialization of monocular visual-inertial estimation on aerial robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Vancouver, BC, Canada*, pages 24–28, 2017.
- Michael Bloesch et al. Iterated extended kalman filter based visual-inertial odometry using direct photometric feedback. *The International Journal of Robotics Research*, 36(10):1053–1072, 2017.
- John Aloimonos et al. Active vision. *International journal of computer vision*, 1(4):333–356, 1988.
- Jeannette Bohg et al. Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017.
- Ruzena Bajcsy et al. Revisiting active perception. *Autonomous Robots*, pages 1–20, 2017.
- Thomas S Collett. Insect vision: controlling actions through optic flow. *Current Biology*, 12(18):R615–R617, 2002.
- Bruno Mantel, Thomas A Stoffregen, Alain Campbell, and Benoît G Bardy. Exploratory movement generates higher-order information that is sufficient for accurate perception of scaled egocentric distance. *PLoS one*, 10(4):e0120025, 2015.
- Karl Kral and Michael Poteser. Motion parallax as a source of distance information in locusts and mantids. *Journal of insect behavior*, 10(1):145–163, 1997.
- Geoslam. <https://geoslam.com/>, 2018.
- Felix Endres et al. 3-D mapping with an RGB-D camera. *IEEE Transactions on Robotics*, 30(1):177–187, 2014.
- Parrot slamdunk. <http://developer.parrot.com/docs/slamdunk/>, 2018.
- Jakob Engel et al. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 2017.
- Nikolai Smolyanskiy et al. Toward low-flying autonomous mav trail navigation using deep neural networks for environmental awareness. *arXiv preprint arXiv:1705.02550*, 2017.
- Paolo D’Alberto et al. Generating FPGA-accelerated DFT libraries. In *Field-Programmable Custom Computing Machines, 2007. FCCM 2007. 15th Annual IEEE Symposium on*, pages 173–184. IEEE, 2007.
- Eriko Nurvitadhi et al. Can FPGAs beat GPUs in accelerating next-generation deep neural networks? In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 5–14. ACM, 2017.
- Song Han et al. ESE: Efficient speech recognition engine with sparse LSTM on FPGA. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 75–84. ACM, 2017.
- Giuseppe Loianno et al. Estimation, control, and planning for aggressive flight with a small quadrotor with a single camera and imu. *IEEE Robotics and Automation Letters*, 2(2):404–411, 2017.
- Davide Falanga et al. Aggressive quadrotor flight through narrow gaps with onboard sensing and computing using active vision. In *Robotics and Automation, 2017 IEEE International Conference on*, pages 5774–5781. IEEE, 2017.
- Nicolas Franceschini et al. From insect vision to robot vision. *Phil. Trans. R. Soc. Lond. B*, 337(1281):283–294, 1992.
- Mandyam V Srinivasan et al. Robot navigation inspired by principles of insect vision. *Robotics and Autonomous Systems*, 26(2-3):203–216, 1999.
- Julien R Serres and Franck Ruffier. Optic flow-based collision-free strategies: From insects to robots. *Arthropod structure & development*, 46(5):703–717, 2017.
- Kirk YW Scheper et al. Behavior trees for evolutionary robotics. *Artificial life*, 22(1):23–48, 2016.
- Longuet-Higgins et al. The interpretation of a moving retinal image. In *Proc. R. Soc. Lond. B*, volume 208, pages 385–397. The Royal Society, 1980.
- Koushil Sreenath et al. Trajectory generation and control of a quadrotor with a cable-suspended load—a differentially-flat hybrid system. In *Robotics and Automation, 2013 IEEE International Conference on*, pages 4888–4895. IEEE, 2013.
- Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. 1981.
- Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981.
- Eddy Ilg et al. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, 2017.
- Carlo Tomasi and Takeo Kanade. Detection and tracking of point features. 1991.
- Olaf Hall-Holt et al. Finding large sticks and potatoes in polygons. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 474–483. Society for Industrial and Applied Mathematics, 2006.
- Daniel Mellinger and Vijay Kumar. Minimum snap trajectory generation and control for quadrotors. In *Robotics and Automation, 2011 IEEE International Conference on*, pages 2520–2525. IEEE, 2011.
- Clément Godard et al. Unsupervised monocular depth estimation with left-right consistency. In *Computer Vision and Pattern Recognition, 2017 IEEE International Conference on*, 2017.
- Anurag Ranjan and Michael J. Black. Optical flow estimation using a spatial pyramid network. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Till Kroeger et al. Fast optical flow using dense inverse search. In *European Conference on Computer Vision*, pages 471–488. Springer, 2016.
- JiaWang Bian et al. Gms: Grid-based motion statistics for fast, ultra-robust feature correspondence. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Edward Rosten et al. Faster and better: A machine learning approach to corner detection. *IEEE transactions on pattern analysis and machine intelligence*, 32(1):105–119, 2010.
- Mårten Björkman et al. Detecting, segmenting and tracking unknown objects using multi-label mrf inference. *Computer Vision and Image Understanding*, 118:111–127, 2014.