

Autonomous MAV guidance with a lightweight omnidirectional vision sensor

Richard J. D. Moore, Karthik Dantu, Geoffrey L. Barrows & Radhika Nagpal

Abstract—This study describes the design and implementation of several bioinspired algorithms for providing guidance to an ultra-lightweight micro-aerial vehicle (MAV) using a 2.6 g omnidirectional vision sensor. Using this visual guidance system we demonstrate autonomous speed control, centring, and heading stabilisation on board a 30 g MAV flying in a corridor-like environment. In addition to the computation of wide-field optic flow, the comparatively high-resolution omnidirectional imagery provided by this sensor also offers the potential for image-based algorithms such as landmark recognition to be implemented in the future.

I. INTRODUCTION

Recent advances in airframe construction, flight dynamics and control, and sensor design are quickly driving insect-scale micro-aerial vehicles (MAVs) closer to mass production (e.g. the RoboBee [1], [2]). Ultra-lightweight MAVs could operate indoors and around humans more safely than heavier vehicles, and swarms of low-cost autonomous MAVs could perform tasks such as aerial micro-manipulation (e.g. crop pollination) or exploration of hazardous environments with greater efficiency than a more capable single agent acting alone [2]. However, due to the extreme resource constraints of the platform, novel guidance systems must be designed for such lightweight MAVs to navigate and perform their duties autonomously.

In prior work, we have described how wide-field optic flow, computed with a novel 2.6 g omnidirectional vision sensor, may be used to enable a lightweight upright-flying MAV (Fig. 1a) to hover in-place [3]. Here, we have re-designed the processing pathway to investigate visual guidance in an office-like environment. Using several simple bioinspired guidance algorithms we demonstrate autonomous speed control, centring, and heading stabilisation with our MAV in a corridor. Image processing, computation of optic flow, and low-level heading stabilisation are all performed on board at 10 Hz, although control commands are computed at a ground station to simplify development and logging.

Our research platform serves as a development proxy for the RoboBee, and our eventual goal is to transition these

This work was partially supported by the National Science Foundation (award number CCF-0926148).


Correspondence via rjdmoores@seas.harvard.edu

R. J. D. Moore is with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA.


K. Dantu is with the Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, NY 14260, USA.

G. L. Barrows is with Centeye Inc., Washington, DC 20008, USA.

R. Nagpal is with the School of Engineering and Applied Sciences and the Wyss Institute for Biologically Inspired Engineering, Harvard University, Cambridge, MA 02138, USA.



(a) Custom Blade mCX2 with sensor ring. (b) Sensor ring plan view.



(c) Sensor ring imagery during a traverse of an office corridor.

Fig. 1. (a) Our research platform, comprising a miniature coaxial helicopter, custom processing and control boards, and vision sensor ring; (b) a diagram of the sensor ring, showing the positions of the vision sensors, S_n , and their fields of view (blue), including overlapping regions (dark blue); and (c) omnidirectional imagery captured from the sensor ring at approximately 0.5 s intervals during a traverse of an office corridor.

guidance algorithms to a miniaturised omnidirectional vision system [4] on board the RoboBee.

II. RELATED WORK

The advantages of vision-based sensors over competing active sensing technologies has led to a plethora of studies on visual guidance for autonomous aircraft. Traditional computer vision-based approaches such as stereo vision or visual odometry can give the vehicle a detailed understanding of the surrounding environment and its own motion through it. However, such approaches are often burdened by significant computational overheads and thus require either aircraft capable of carrying substantial payloads [5]–[7], or that the raw sensor data be transmitted off-board for processing and computation of control signals [8].

Over the past several decades, much research has shown that biological vision systems can inspire novel strategies for aircraft guidance that can offer dramatically improved sensing and control efficiencies over more complex computer

vision-based approaches [9]. The discovery of optic flow-based guidance heuristics in honeybees [10] inspired a number of implementations of those algorithms for autonomous regulation of aircraft speed and height [11], avoidance of obstacles [12], and landing [13]. Specialised optic flow sensors have enabled the development of a new class of ultra-lightweight (< 50 g) autonomous aerial platforms that are capable of performing simple behaviours [14], [15] and reactive navigation within enclosed environments [16]. The guidance system for the 10 g fixed-wing aircraft described by [16] was inspired by the saccadic flight patterns of fruit flies and relied on two 1D lateral optic flow sensors, which were used to detect imbalances and frontal expansion in the optic flow field to generate appropriate steering commands.

Our system also falls into the class of ultra-lightweight autonomous MAVs, but we expect it to be more capable than other previously described platforms (*e.g.* [16]) due to the ability of our system to capture comparatively high-resolution ($128 \text{ px} \times 16 \text{ px}$) omnidirectional images at a high frame rate (~ 100 Hz). Thus, in addition to the computation of wide-field optic flow for reactive guidance, which is the extent of the present study, our vision system should also be capable of supporting image-based algorithms for proactive navigation, such as landmark recognition, relative heading determination, or visual servoing. At a total weight of approximately 30 g, we believe that this system represents one of the smallest autonomous aerial platforms described to date with this range of capabilities.

III. PLATFORM & HARDWARE

The research platform for this work was the Blade mCX¹ (Fig. 1a) – a low-cost ($< \$100$ USD), micro-scale (max. dimension 19 cm), radio-controlled, coaxial helicopter. Our customised prototype weighed approximately 30 g and had a flight time of the order of 5 – 10 minutes. Roll and pitch axes were stabilised mechanically via the flybar.

For this study we removed the proprietary control board and replaced it with two custom boards², the *heliboard* and *procboard*, which both incorporated a low-power 32-bit Atmel microcontroller (AT32UC3B1256). The *procboard* processed image data captured by the vision sensors to extract wide-field optic flow, which was then transferred to the *heliboard* via an I²C interface and fused with rotation rate data from a 3-axis gyroscope. De-rotated optic flow measurements were transmitted wirelessly to a ground station where control commands were computed and transmitted back to the helicopter (Fig. 2). We used the *Simbeotic* [17] test-bed to design flight tests and synchronise control commands and logging.

A. Vision sensor ring

The vision sensor ring (Fig. 1) comprised eight Centeye Faraya³ $64 \text{ px} \times 64 \text{ px}$ vision chips arranged around the

¹<http://www.e-fliterc.com/Products/Default.aspx?ProdID=EFLH2480>

²Designed by Centeye, Inc.

³<http://centeye.com/products/current-vision-chips>




Fig. 2. Our system architecture enables sensing or processing to be switched between on-board and ground-based components for modular development and testing. Details in text.

circumference of a ~ 10 cm diameter ring of flexible interconnect. The sensor ring and sensors weighed a total of 2.6 g. Each vision chip had a square field of view (FOV) of approximately $60^\circ \times 60^\circ$, which gave a combined FOV in excess of $360^\circ \times 60^\circ$ azimuth and elevation respectively. The spacings between vision chips on the sensor ring (Fig. 1b) were chosen to give a greater density of azimuthal optic flow measurements lateral to the helicopter, where the optic flow signals were expected to be strongest, whilst maintaining an omnidirectional FOV. Pixel binning was performed on the vision chips so that $16 \text{ px} \times 16 \text{ px}$ images at a bit depth of 10 bpp could be read out from all vision chips simultaneously (Fig. 1c).

B. Optic flow

Azimuthal optic flow was computed on the *procboard* using a simple block matching algorithm based on the sum of absolute disparity (SAD) [18] between blocks of pixels in the current and previous image frames from each vision sensor. We found this approach to be more robust to image noise than other methods (I2A [19] and LK [20]). Image motion was estimated for two horizontally adjacent subregions per sensor image, giving a total of 16 one-dimensional azimuthal optic flow vectors. We used a search range of ± 3 px and applied an equiangular fit [21] to the minimum and neighbouring SAD scores to estimate the true image offset in each subregion with sub-pixel accuracy. We determined empirically that a time lag between image frames of 100 ms maximised the sensitivity of our optic flow measurements while ensuring that the expected image velocities measured by the sensor ring did not result in displacements greater than our ± 3 px search range. Our measurement and control loop on board the helicopter was thus executed at 10 Hz.

IV. MODELLING ON-BOARD OPTIC FLOW

In the horizontal plane, optic flow, $\delta\theta$, measured at an azimuthal viewing direction, θ , can be stated simply as

$$\delta\theta = v \frac{\sin \theta}{d} + \omega, \quad (1)$$

where v and ω are respectively the observer's speed and yaw rate, and d is the range to the viewed object.

The rotational component of optic flow does not encode range and was discarded by subtracting the gyroscopically measured yaw rate of the helicopter from the azimuthal optic flow measured by the sensor ring. The residual optic flow therefore encodes both the magnitude and direction of the translation of the helicopter, as well as the range to objects within the environment, and thus provides useful guidance information for small autonomous agents such as insects or MAVs, where the inter-ocular separation is too small to provide useful stereoscopic range information during flight.

A. Characterising sensor performance

The performance of the sensor ring was analysed by manually translating it along the centre line of a representative office corridor at constant speed. A corridor represents a simplified indoor environment, which enabled the expected distribution of optic flow to be characterised easily. In the horizontal plane, range to a wall, d , at an azimuthal angle, θ , is given by $d = \frac{r}{\sin \theta}$, where r is the half-width of the corridor. Substituting into (1) gives the expected distribution of azimuthal optic flow in a corridor as a function of the viewing angle,

$$\delta\theta = \frac{v}{r} \sin^2 \theta. \quad (2)$$

The median image displacements recorded in each subregion during the pass of the corridor correlated well with the expected distribution (Fig. 3e). However, the instantaneous optic flow measurements were affected by the uncontrolled lighting and texture conditions in the office corridor. Therefore, several noise processes that could affect the computation of optic flow were explored with the aim of improving the robustness of the instantaneous measurements to varying environmental conditions.


1) *Image noise*: Each frame, four images were captured at the maximum transfer rate of the vision chips (~ 100 Hz), accumulated, and averaged to give a single image with reduced random pixel noise (Fig. 3b).

2) *Varying illumination*: A 3×3 horizontal Sobel filter kernel was convolved with the accumulated sensor images (Fig. 3c) to reduce spurious measurements of image motion caused by brightness differences between frames.

3) *Unmodelled motion*: Finally, the filtered images were compressed vertically by summing and combining multiple image rows at a ratio of 4:1 (Fig. 3d) to reduce the deleterious effects of unmodelled rolling or pitching motions on the computation of azimuthal optic flow.

V. OPTIC FLOW-BASED FLIGHT CONTROL

The optic flow that is produced by stable and centred flight in a corridor is predicted for all azimuthal viewing angles by (2). Similarly, deviations from this ideal scenario – either large-scale variations in the environment, or disturbances to the helicopter’s position or orientation – cause predictable variations in the measured distribution of wide-field optic flow. We used three simple and bio-inspired algorithms to map various deviations in the distribution of azimuthal optic flow to roll, pitch, and yaw control commands for the MAV.



(e) Median optic flow in corridor. (f) Coefficient of variation vs filtering.

Fig. 3. Three image processing steps were applied successively to the raw sensor ring image each frame to improve the robustness of the computed optic flow to varying environmental conditions. (a) A $32 \text{ px} \times 16 \text{ px}$ region cropped from the raw $128 \text{ px} \times 16 \text{ px}$ omnidirectional image; (b) four raw frames accumulated and averaged; (c) accumulated image convolved with a 3×3 horizontal Sobel kernel; and (d) Sobel filtered image compressed vertically by a ratio of 4 : 1. (e) The image processing steps preserved the expected distribution (best fit of (2) to filtered data recorded from office corridor) of median optic flow magnitudes at each viewing angle (circles), and (f) also reduced by $\sim 25\%$ the average coefficient of variation (CV), $\frac{\sigma}{\mu}$, of optic flow measured at each viewing angle.

A. Flow-based speed control

Translational optic flow encodes both the speed of the observer as well as the range to objects in the environment. Thus, by controlling pitch, or forwards speed, based on the total sum of observed translational optic flow, the helicopter is ensured to slow down in the presence of obstacles and to speed up in uncluttered spaces [10].

We computed a signed sum of optic flow to detect both forwards and backwards motion. Under pure translation, optic flow vectors in opposing visual hemispheres have opposite signs, due to the orientation of the sensors’ image planes with respect to the direction of motion. Thus, the total translational optic flow was computed by summing the difference between translational optic flow measured from diametrically opposite subregions,

$$\mathbf{F}_{\text{speed}} = \sum_{i=1}^8 (\delta\theta)_{T,i+8} - (\delta\theta)_{T,i}, \quad (3)$$

where $(\delta\theta)_{T,i}$ represents the translational component of optic flow measured in the i^{th} subregion (Fig. 4a). An integral-derivative controller ($k_i = 2.0, k_d = 0.5$) was used to derive a pitch control command that minimised the error between the total translational optic flow, $\mathbf{F}_{\text{speed}}$, and the set point

($\bar{\mathbf{F}}_{\text{speed}} = 0.2$ for closed-loop experiments). Integrated error gain dominated the pitch controller to permit a variable mapping between actual forwards speed and observed $\mathbf{F}_{\text{speed}}$, depending on the scale of the environment.




Fig. 4. Subdivision of azimuthal subregions for optic flow-based control, showing corresponding sensors, S_n , and subregions, R_n , as well as opposing subregions (red/blue) for (a) centring and speed control, and (b) heading control. Quadrants, Q_n , for heading control are labelled. Details in text.

B. Flow-based centring

Balancing the translational optic flow measured in each visual hemisphere enables the helicopter to maintain an even clearance to objects on both sides of its body simultaneously [10]. We computed a centring metric by summing the difference between the absolute translational optic flow in opposing hemispheres,

$$\mathbf{F}_{\text{centring}} = \sum_{i=1}^8 |(\delta\theta)_{T,i+8}| - |(\delta\theta)_{T,i}|, \quad (4)$$

where $|(\delta\theta)_{T,i}|$ is the absolute magnitude of the i^{th} translational optic flow vector (Fig. 4a). A proportional-integral controller ($k_p = 2.0, k_i = 0.05$) was used to derive a roll control command that minimised the error between the differential optic flow, $\mathbf{F}_{\text{centring}}$, and the set point ($\bar{\mathbf{F}}_{\text{centring}} = 0$ for closed-loop experiments). The small integrating error term was used to estimate the roll trim automatically during optic flow-controlled flight.

C. Flow-based heading control

For the centring algorithm to be most effective, the helicopter must maintain its heading such that the closest obstacle lies laterally (*e.g.* the MAV should maintain its heading parallel with a corridor's long axis). To realise this control strategy, the absolute translational optic flow was integrated according to

$$\mathbf{F}_{\text{heading}} = \sum_{i=1}^{16} \alpha \cdot \frac{|(\delta\theta)_{T,i}|}{\sin^2 \theta}, \quad \alpha \begin{cases} +1, & i \in \{Q_1, Q_3\} \\ -1, & i \in \{Q_2, Q_4\} \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

where Q_n represents the n^{th} view sphere quadrant (Fig. 4b). Absolute translational optic flow was negated in alternating quadrants so that a zero net yaw torque was generated when the distributions of measured optic flow magnitudes in the front and rear quadrants were symmetric for both the

left and right hemispheres, regardless of which hemisphere experienced the greater optic flow (*i.e.* proximity to one wall should not necessarily generate a yawing torque). Optic flow components were scaled by $\frac{1}{\sin^2 \theta}$ so that each component contributed equally to the sum under ideal conditions, *i.e.* (2). This heuristic was empirically determined to increase sensitivity to heading deviations. Optic flow vectors closest to the longitudinal axis of the helicopter were discarded ($\alpha = 0$) because their small viewing angle, θ , with respect to the helicopter's presumed direction of translation led to large variations in scaled flow magnitude.

A proportional controller ($k_p = 0.3$) was used to derive a yaw control command that minimised the error between the scaled translational optic flow, $\mathbf{F}_{\text{heading}}$, and the set point ($\bar{\mathbf{F}}_{\text{heading}} = 0$ for closed-loop experiments). Additionally, on the *heliboard*, the helicopter's gyroscopically measured yaw rate was used to dampen the vehicle's turn rate and also to estimate the vehicle's yaw rate trim.

VI. FLIGHT TEST RESULTS

A mock corridor was set up inside our indoor flight arena to analyse the performance of the sensor ring during flight (Fig. 5). A Vicon⁴ motion capture system was used to record the position and orientation of the helicopter during the flight tests. Data from both the Vicon system and on-board optic flow measurements were processed at the base station to generate control commands, which were communicated to the helicopter via the wireless link (Fig. 2). This experimental configuration permitted partial Vicon control and partial optic flow control during development of the system.



Fig. 5. A photograph of the corridor-like environment within the Vicon arena during an experiment. The dimensions of the corridor were approximately 5.0 m \times 2.6 m \times 1.5 m ($L \times W \times H$).

In order to analyse the optic flow signals generated during both rotational and translational motion, a simple flight plan was devised in which the helicopter traversed the length of the corridor and made two $\sim 180^\circ$ turns.

A. Closed-loop flight using a motion capture system

During initial flight testing, the helicopter was controlled automatically using position and heading information from the Vicon system and an independent control system; optic

⁴<http://www.vicon.com>

flow data was recorded. Rotation induced an approximately equal component of raw optic flow in each sensor in accordance with (1) and was discarded by subtracting the helicopter's gyroscopically measured yaw rate (Fig. 6). Residual optic flow was due only to the translational motion of the helicopter and closely matched the distribution predicted by (2) for flight in a corridor-like environment.





Fig. 6. (a) Raw optic flow recorded during the Vicon flight test, and (b) after de-rotation using gyroscopic yaw rate. Thick red bars correspond to the three translational flight segments, including a single long traverse of the length of the corridor (19 s \sim 25 s).

B. Closed-loop flight using on-board optic flow


The optic flow-based guidance algorithms were used to demonstrate autonomous and simultaneous control of the helicopter's forwards speed, lateral position, and heading direction while traversing the corridor. Optic flow-based control was engaged while the helicopter hovered at one end of the corridor. The evolution of the helicopter's flight path during each of the trials was then tracked using the Vicon motion capture system (Fig. 7a). The position and orientation of the helicopter at the start of the traverse were varied between trials to provide a range of initial conditions (± 0.5 m lateral offset and $\pm 20^\circ$ angular offset from the centre line). It was expected that the guidance algorithms would enable the helicopter to accelerate to an approximately constant cruise speed and to eventually obtain a stable trajectory down the centre line of the corridor. During the traverse, the helicopter flew autonomously using only sensory input from the vision sensors and gyroscope to control roll, pitch, and yaw, whilst the throttle was controlled using the Vicon system to maintain a height above ground of 0.75 m.

Once optic flow-based control was engaged, the helicopter accelerated forwards and took an average of 2.2 s to reach 85% of the total optic flow set point, $F_{\text{speed}} = 0.2$ (Fig. 7b centre). In some trials, the helicopter initially drifted backwards because the integral-dominant pitch controller took a finite time to overcome the negative pitch trim. Estimates of the steady-state performances of the guidance algorithms were achieved by arbitrarily considering just the final quarter of each closed-loop traverse (*i.e.* $y > 1$ m in Fig. 7). Using this metric, the average steady-state total optic flow was $F_{\text{speed}} = 0.19 \pm 0.08$ ($\pm 2\sigma$), which enabled the helicopter to hold an average forwards speed of 0.52 ms^{-1} .

Differential optic flow was driven towards zero by the centring algorithm (Fig. 7b upper), which resulted in an



(a) 2D position (solid coloured lines) and heading (coloured tabs) of the helicopter, measured via Vicon motion capture during six closed-loop trials under optic flow-based control. Also shown are the corridor walls (thick black lines) as well as the path of the helicopter under pure Vicon control (dashed black line) for comparison. Heading tabs are drawn at 30 ms intervals.



(b) Plots of the individual roll, F_{centring} (upper), and pitch, F_{speed} (centre), optic flow metrics used for centring and speed control respectively during the six closed-loop trials; as well as the heading direction (lower) measured via Vicon motion capture. Also shown are the respective optic flow set points (red dashed lines) and mean responses (solid black lines) from the five successful trials. Roll and pitch metrics are given in arbitrary units.

Fig. 7. Optic flow-based closed-loop control. Individual closed-loop trials are represented by corresponding solid coloured lines in (a) and (b). The beginning (circle) and end (cross) of the optic flow-based control segment is indicated for each trial. All data are plotted against the helicopter's position along the y axis during each trial to simplify comparison. Positive yaw rate corresponded to an anticlockwise rotation of the helicopter's heading angle, due to the orientation of the Vicon coordinate frame.

average unsigned steady-state centring (x position) error of $0.34 \text{ m} \pm 0.27 \text{ m}$ from the centreline of the 2.6 m-wide corridor. The relatively large variance in steady-state x position between trials compared to the centring performance under Vicon control (Fig. 7a) was attributed to the optic flow-based centring PI controller, which was tuned to accommodate the large and non-linear range of F_{centring} as a function of the helicopter's lateral position. However, the centring metric, F_{centring} (Fig. 7b upper), did correspond well with the actual x position of the helicopter throughout each trial (Fig. 7a),

thereby validating the guidance algorithm.

The helicopter maintained its heading direction approximately in line with the longitudinal axis of the corridor (Fig. 7b lower), as governed by (5). However, heading direction was also attracted towards the helicopter's instantaneous direction of translation, which was unexpected. The sharpest changes in heading direction corresponded to periods of highly transverse flight (Fig. 7a), when the helicopter's direction of translation was not aligned with its forwards axis. This behaviour likely contributed to the failure of one of the closed-loop optic flow trials by increasing the helicopter's angle of approach to the wall of the corridor.

VII. DISCUSSION

The inherent design constraints for small-scale MAVs generally restrict autonomy to specific behaviours, making direct comparisons between different platforms difficult. Of those that have demonstrated autonomous centring, the approach most similar to ours is that described by [22], who described a wide-field azimuthal optic flow-based guidance system and demonstrated centring and heading stabilisation on board a 680 g quadrotor, which incorporated a 500 MHz microprocessor and a wide-angle camera-mirror assembly. From published raw data, their platform was able to remain within $\pm \sim 0.4$ m of the centreline of their corridor over repeated trials. For comparison, we computed the maximum (2σ) steady-state centring error for our MAV to be 0.61 m.

A. Limitations

We did not evaluate our platform's ability to avoid front-on collisions. However, the failure of one of the optic flow-based closed-loop trials showed that our centring and speed control behaviours were only able to prevent a collision up to some maximum angle of approach to an obstacle. From our experimental data, the lower bound on this maximum angle of approach was $\sim 25^\circ$ (Fig. 7a). A simple collision avoidance strategy might be to detect patterns of expanding optic flow and steer accordingly (as in [23]).

A second limitation to the present system stems from our implicit assumption that the helicopter travels approximately forwards (pitch affects speed, roll affects centring). During highly transverse flight, coupling occurs between the wide-field azimuthal optic flow distributions caused by lateral translation and the vehicle's orientation with respect to the corridor [24]. This coupling would have reduced the efficacies of the described control algorithms and was likely responsible for the failure of the MAV's heading to converge to 90° (Fig. 7b lower). To resolve this issue, the MAV's direction of translation could be estimated from the wide-field distribution of azimuthal optic flow or from a downwards-facing 2D optic flow sensor.

VIII. CONCLUSION

Our 2.6 g omnidirectional vision sensor enables the computation of wide-field optic flow on board a MAV weighing approximately 30 g in total. We have implemented this system and described three simple guidance algorithms that

enable autonomous and simultaneous control of the MAV's flight speed, lateral position, and heading direction with respect to a corridor-like environment. All sensing and all significant processing was performed on board at 10 Hz. Future work includes studying goal-oriented navigation by a MAV using this vision system in office-like environments.

ACKNOWLEDGMENT

Travis Young (Centeye Inc.) for initial development of the embedded hardware and software, and Raphael Cherney (Harvard) for assisting with initial system characterisation.

REFERENCES

- [1] Ma, Chirattananon, Fuller, and Wood, "Controlled flight of a biologically inspired, insect-scale robot," *Science*, 2013.
- [2] Wood, Nagpal, and Wei, "Flight of the robobees," *Sci. Am.*, 2013.
- [3] Barrows, Young, Neely, Leonard, and Humbert, "Vision based hover in place," in *Proc. Aerospace Sciences Meeting*. AIAA, 2012.
- [4] Koppal, Gkioulekas, Young, Park, Crozier, Barrows, and Zickler, "Towards wide-angle micro vision sensors," *Trans. Pattern Anal. Mach. Intell.*, 2013.
- [5] Fraundorfer, Heng, Honegger, Lee, Meier, Tanskanen, and Pollefeys, "Vision-based autonomous mapping and exploration using a quadrotor MAV," in *Proc. IROS*. IEEE/RSJ, 2012.
- [6] Shen, Mulgaonkar, Michael, and Kumar, "Vision-based state estimation for autonomous rotorcraft MAVs in complex environments," in *Proc. ICRA*. IEEE/RAS, 2013.
- [7] Weiss, Scaramuzza, and Siegwart, "Monocular-SLAM-based navigation for autonomous micro helicopters in GPS-denied environments," *J. Field Rob.*, 2011.
- [8] Ahrens, Levine, Andrews, and How, "Vision-based guidance and control of a hovering vehicle in unknown, GPS-denied environments," in *Proc. ICRA*. IEEE/RAS, 2009.
- [9] Moore, "Vision systems for autonomous aircraft guidance," Ph.D. dissertation, 2012.
- [10] Srinivasan, Zhang, Lehrer, and Collett, "Honeybee navigation en route to the goal: visual flight control and odometry," *J. Exp. Biol.*, 1996.
- [11] Srinivasan, Moore, Thurrowgood, Soccol, and Bland, "From biology to engineering: Insect vision and applications to robotics," in *Frontiers in Sensing*. Springer, 2012.
- [12] Beyeler, Zufferey, and Floreano, "Vision-based control of near-obstacle flight," *Auton. Rob.*, 2009.
- [13] Green, Oh, Sevcik, and Barrows, "Autonomous landing for indoor flying robots using optic flow," in *Proc. IMECE*. ASME, 2003.
- [14] G. Bermudez and Fearing, "Optical flow on a flapping wing robot," in *Proc. IROS*. IEEE/RSJ, 2009.
- [15] Duhamel, Pérez-Arancibia, Barrows, and Wood, "Altitude feedback control of a flapping-wing microrobot using an on-board biologically inspired optical flow sensor," in *Proc. ICRA*. IEEE/RAS, 2012.
- [16] Zufferey, Klaptocz, Beyeler, Nicoud, and Floreano, "A 10-gram vision-based flying robot," *Adv. Robot.*, 2007.
- [17] Kate, Waterman, Dantu, and Welsh, "Simbeotic: A simulator and testbed for micro-aerial vehicle swarm experiments," in *Proc. IPSN*. ACM, 2012.
- [18] Hartley and Zisserman, *Multiple view geometry in computer vision*. Cambridge University Press, 2003.
- [19] Srinivasan, "An image-interpolation technique for the computation of optic flow and egomotion," *Biol. Cybern.*, 1994.
- [20] Lucas and Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. IJCAI*, 1981.
- [21] Shimizu and Okutomi, "Significance and attributes of subpixel estimation on area-based matching," *Sys. Comp. Japan*, 2003.
- [22] Conroy, Gremillion, Ranganathan, and Humbert, "Implementation of wide-field integration of optic flow for autonomous quadrotor navigation," *Auton. Rob.*, 2009.
- [23] Zufferey and Floreano, "Fly-inspired visual steering of an ultralight indoor aircraft," *Trans. Rob.*, 2006.
- [24] Humbert, Murray, and Dickinson, "Sensorimotor convergence in visual navigation and flight control systems," in *Proc. IFAC World Congress*, 2005.