

optiPilot: control of take-off and landing using optic flow

Antoine Beyeler* Jean-Christophe Zufferey Dario Floreano

Laboratory of Intelligent Systems
École Polytechnique Fédérale de Lausanne, Switzerland (EPFL)
<http://lis.epfl.ch/>

senseFly, Ecublens, Switzerland
<http://www.sensefly.com>

Note: the definitive version of this paper was published in the proceedings of the European Micro Air Vehicle conference and competition 2009 (EMAV 2009).

ABSTRACT

Take-off and landing manoeuvres are critical for MAVs because GPS-based autopilots usually do not perceive distance to the ground or other potential obstacles. In addition, attitude estimates based on inertial sensors are often perturbed by the strong accelerations occurring during launch. This paper shows how our previously developed control strategy, called *optiPilot*, can cope with take-off and landing using a small set of inexpensive optic flow sensors.

1 INTRODUCTION

As shown by several recent studies [1, 2, 3], Miniature Aerial Vehicles (MAVs) are on the verge of being transferred from research labs to industry, which will turn them into products for commercial applications. Inertial Measurement Units (IMU) and GPS are commonly used on these aerial robots to achieve flight stabilisation and waypoint navigation [1, 4]. While these systems are able to autonomously control MAVs high in the sky, some issues remain that prevent safe operation in the vicinity of the ground.

We recently proposed a control strategy, called *optiPilot*, that exploits optic flow sampled all around the flight direction to provide MAVs with flight stabilisation and collision avoidance capabilities [5]. We implemented *optiPilot* using a series of low-cost computer mouse sensors with customised lenses for enhanced depth of field. Even in its early prototype stage, the 18-gram compound vision system made of seven optic flow sensors allowed us to demonstrate flight stabilisation, altitude control, and collision avoidance with a 400-gram flying wing¹. In principle, *optiPilot* is generic enough to fit any platform that has its translation vector mostly aligned with its main body axis, which is the case of fixed-wing MAVs as well as most rotorcraft in translation flight.

*Email address: antoine.beyeler@a3.epfl.ch

¹Videos of these experiments are available on <http://lis.epfl.ch/microflyers>

In this paper, we go one step further and describe how the same control strategy enables automatic take-off and landing. These features are not only useful on their own, but ideally complement IMU- and GPS-based autopilots. Take-off manoeuvres indeed represent a difficult phase for two reasons. First IMUs are often momentarily disturbed as a result of the strong accelerations that tend to saturate the inertial sensors. At the same time, and in particular in case of hand-launched take-off, the attitude of the MAV can be dangerously banked just after release, which calls for fast and efficient counter actions. Regarding the landing phase, IMU/GPS systems do not provide a precise way of measuring height above ground especially when the terrain is not perfectly flat or its altitude is unknown. This paper shows how *optiPilot* enables a safe control of both the initial and final phases of the flight while taking into account the proximity and the shape of the underlying terrain.

Optic flow contains information about the distance to the surroundings that can be used to control take-off and landing. Barber et al used optic flow sensors to assist the control of altitude and landing [6]. However, they still relied on GPS and IMU for altitude and attitude control. Other approaches have included optic flow for altitude control of flying platforms [7, 8], but they did not specifically demonstrate take-off and landing and required manual control of the attitude. Chahl et al used optic flow to control altitude during the approach phase of a fixed-wing MAV, but human assistance was required and complete landings were not demonstrated [9]. Ruffier et al studied optic-flow-based altitude control and landing using a tethered rotorcraft [10], but external assistance was required to maintain the optic flow detector upright and achieve autonomous flight.

In the next section, we provide a description of the *optiPilot* control strategy. Section 3 presents the materials and methods used to assess its performance. The results are described in section 4.

2 CONTROL STRATEGY

OptiPilot is a way of mapping a series of optic flow signals sampled in well-defined viewing directions into roll and




Figure 1: Coordinate system of an aircraft reference frame, with the names of the three rotation directions. On typical aircraft, the roll is controlled using the ailerons, and the pitch rate using the elevator. On flying wings such as the one displayed here, roll and pitch rotations are controlled by the differential and, respectively, common mode of actuation of the two control surfaces called elevons. These two modes of actuation are functionally identical to the ailerons and elevator. Yaw is usually either passively stabilised using fixed vertical surfaces or controlled using a rudder. In normal flight, passive or active yaw regulation is used to produce so-called coordinated turns [11]. Only the ailerons and elevator, or the elevons, are used to actually steer the aircraft.




Figure 2: Overview of the steps required to map the data provided by the vision system and rate gyroscopes into control signals. Yaw is assumed to be passively regulated (see Figure 1).

pitch rate commands (Figure 1). This control strategy follows a simple three-stage scheme illustrated in Figure 2. In this section, we first discuss how optic flow can directly be interpreted as proximity signals. We then describe the viewing directions at which the optic flow detectors should point and how optic flow measurements are combined into control signals for steering an aircraft.

The fundamental property of optic flow that enables proximity estimation is often referred to as *motion parallax* [12]. Essentially, it states that the component of optic flow that is induced by translatory motion is proportional to the magnitude of this motion and inversely proportional to the distance to static obstacles in the environment. It is also proportional to the sine of the angle between the translation vector and the

viewing direction. This can be written:

$$p_{\mathbf{T}}(\theta, \psi) = \frac{|\mathbf{T}|}{D(\theta, \psi)} \sin(\alpha), \quad (1)$$

where $p_{\mathbf{T}}(\theta, \psi)$ is the amplitude of translational optic flow measured in direction (θ, ψ) (see Figure 3 for the polar coordinate system convention), \mathbf{T} is the translation vector, $D(\theta, \psi)$ is the distance to the obstacle seen in direction (θ, ψ) and α is the angle between the translation vector \mathbf{T} and the viewing direction (θ, ψ) .

The optic flow perceived by a free-flying aircraft also contains the component induced by its rotations in addition to the translational optic flow described above. Consequently, it is necessary to exclude the optic flow component due to rotations to estimate the proximity of obstacles, a process known as *derotation* of optic flow [13]. In a MAV, this can be achieved by predicting the optic flow generated by rotation, as measured by rate gyroscopes, and then subtracting the predicted optic flow from the measured optic flow.

In normal translational flight, the translation vector is essentially aligned with the aircraft's main axis at all times. The angle α in (1) can be assumed equal to the polar angle θ of the coordinate system introduced in Figure 3 (also known as *eccentricity*). Equ. (1) can therefore be rewritten as:

$$p_{\mathbf{T}}(\theta, \psi) = \frac{|\mathbf{T}|}{D(\theta, \psi)} \sin(\theta) = \mu(\theta, \psi) \cdot |\mathbf{T}| \cdot \sin(\theta), \quad (2)$$

where μ is the proximity of obstacle (i.e. the inverse of distance, also referred to as *nearness*). This means that the magnitude of translational optic flow measurements can be directly interpreted as proximity signals, scaled with the sine of the eccentricity at which the measurements are taken. Also, for decreasing velocity $|\mathbf{T}|$, the same optic flow will be perceived with increasing proximity μ , a property that can be exploited to implement landing behaviour.

Let us now consider the directions where optic flow should be measured and how to combine these measurements to generate control signals for the aircraft. In order to reduce the sensory and computational requirements, it is desirable to keep the number of measurements as low as possible. It also turns out that not all the viewing directions in the visual field have the same relevance for flight control. Directions pointing at $\theta > 90^\circ$ correspond to obstacles that are behind the aircraft and thus do not require avoidance. For θ values close to 0° (i.e. close to the centre of the visual field), the magnitude of the optic flow measurements tends to zero because of the $\sin(\theta)$ factor in (2). Since the resolution of the vision system limits the possibility of measuring small amounts of optic flow, proximity estimation is not reliable for small eccentricities. These two limits ($\theta < 90^\circ$ and $\theta > 0^\circ$) suggest that the area of interest lies around $\theta = 45^\circ$. Experimental and theoretical work has shown that $\theta = 45^\circ$ is indeed optimal in a variety of situations [5, 14].




Figure 3: (left) The image-plane coordinate system used throughout this paper. $\psi \in [0; 2\pi]$ is the azimuth angle, with $\psi = 0$ corresponding to the dorsal part of the visual field and positive extending leftward. $\theta \in [0; \pi]$ is the polar angle. (right) Perspective sketch of the spherical vision system. Note that ψ and θ completely define a viewing direction with respect to the optical and the aircraft main axis.




Figure 4: Possible sampling of the visual field. N viewing directions are uniformly spaced on a circle of polar angle $\hat{\theta}$. Each viewing direction is separated by an inter-azimuthal angle $\hat{\psi}$. On this illustration, $N = 7$, $\hat{\theta} = 45^\circ$ and $\hat{\psi} = 30^\circ$.

During take-off and landing, a MAV equipped with optic flow sensors can typically perceive the ground and lateral obstacles such as trees. In most case, however, the space above the aircraft is free of obstacle. We therefore propose to measure optic flow in N viewing directions along the specific polar angle $\theta = \hat{\theta}$ and with an inter-azimuthal angle $\hat{\psi}$, considering only directions that are on either side and below the aircraft, as illustrated in Figure 4.

In order to map optic flow estimations into control signals, we propose the use of a simple weighted sum, which can be

written as:

$$c^j = \frac{\xi^j}{N \cdot \sin(\hat{\theta})} \cdot \sum_{k=0}^{N-1} p_{\mathbf{T}}(\hat{\theta}, k \cdot \hat{\psi} + \frac{\pi}{2}) \cdot w_k^j, \quad (3)$$

where c^j is the j^{th} control signal (either roll or pitch control), w_k^j the associated set of weights and ξ^j a gain to adjust the amplitude of the control signal. This summation process is qualitatively similar to what is believed to occur in the tangential cells of flying insects [15, 16, 17, 18]; namely, a wide-field integration of a relatively large number of optic flow estimates into a reduced number of control-relevant signals.

In order to use this approach to steer an aircraft, two sets of weights $\{w_k^R\}$ and $\{w_k^P\}$, $k = 0, 1, \dots, N-1$ must be devised, for the roll and the pitch control, respectively. Let us first consider the pitch control signal c^P (Figure 5 left). Proximity signals from the ventral region (i.e. ψ near 180°) correspond to obstacles beneath the aircraft. The corresponding weights should thus be positive to generate a positive control signal that results in a pitch-up manoeuvre. Conversely, proximity signals from either side of the aircraft (i.e. ψ near 90° and 270°) should not influence the pitching behaviour and the corresponding weights should thus be set to zero. An example of such a weight distribution could be:

$$w_k^P = \sin(k \cdot \hat{\psi}). \quad (4)$$

Using a similar reasoning, one can derive the qualitative distribution needed for the weights related to the roll signal (Figure 5 right). Weights corresponding to the left of the aircraft should be positive, in order to initiate a rightward turn in reaction to the detection of an obstacle on the left. Inversely, weights on the right should be negative. Since obstacles in the ventral region ($\psi = 180^\circ$) are avoided by pitching only,




Figure 5: (left) Possible distribution of w_k^P for the generation of the pitch control signal. The arrow in the centre indicates pitch direction for a positive pitch signal. (right) Possible distribution of w_k^R for the generation of the roll control signal. The arrow in the centre indicates roll direction for a positive roll signal.

the weights in this region should be set to zero. The following equation is one way to implement such a weight distribution:

$$w_k^R = \cos(k \cdot \hat{\psi}). \quad (5)$$

3 MATERIALS AND METHODS

To validate the autonomous take-off and landing capabilities of *optiPilot*, we ran a series of experiments with a flying wing platform. In this section, we describe the platform, its sensors, and the experimental method we used.

3.1 Flying platform

The test flying platform is based on a *swinglet* from sensefly² (Figure 6). This aircraft has a wingspan of 80 cm and a total weight of 407 g, including approximately 50 g for the sensor payload specifically required by our experiments. No particular efforts have been made at this stage to reduce the weight of the sensor. The MAV is equipped with an electronic board including a Microchip dsPIC33FJ256GP506 microcontroller, on which our control strategy was implemented. This controller is interfaced to three Analog Devices ADXRS610 rate gyroscopes that we used for optic flow derotation. In order to measure airspeed, it is also equipped with a Freescale MPXV5004DP differential pressure sensor and a custom-built pitot tube. In order to regulate airspeed, a PID controller was implemented to control the thrust of this platform and to maintain a cruise airspeed of 14 m/s, with a precision of ± 2 m/s. In the configuration used for the experiment presented in this paper, the aircraft has a stall speed of about 7 m/s.




Figure 6: Top view of the *swinglet* used for the experiments. It has a wing span of 80 cm and a total weight of 407 g including about 50 g of sensor payload.

In order to record the state of the aircraft during the experiments, we use a Xsens MTi-G unit which provides a 6-degree-of-freedom state estimation. At no time however this unit is used for the control of the aircraft. The states of the aircraft and its sensors are monitored and recorded in real time using a 2.4 GHz Digi XBee-PRO radio-link, the *Ishtar* communication protocol [19] and the *e-motion* monitoring software³.

During the experiments, a human pilot can take over the control of the aircraft using a regular RC controller. This capability is used to steer the platform into specific situations and to subsequently activate *optiPilot* to assess its behaviour

²<http://www.sensefly.com>

³<http://gna.org/projects/e-motion>



Figure 7: From left to right: the Avago ADNS5050 optic mouse sensor, the custom-designed optics based on the Philips CAX100 collimator lens ($f = 10$ mm) and the assembled optic flow detector (weighing 0.8 g).

in autonomous mode. The data presented in this paper includes only fully autonomous flight sequences.

3.2 Optic flow detection

There are a number of technologies available to estimate optic flow, including standard cameras and classical computer vision (for a review: [20, 21]), dedicated sensors such as aVLSI or mixed-mode custom vision chips [22, 23, 7] or optic computer mouse chips [24, 25, 26, 27]. The latter have the advantages of being lightweight, available off-the-shelf, easy to interface to the electronics and not requiring further processing, as the optic flow extraction is done on-chip. Each sensor provides a single 2D optic flow estimation, which requires the installation of as many chips as required viewing directions. This limitation, for a small number of viewing directions, is outweighed by the advantages listed above.

Figure 7 illustrates the vision system we developed. It is made of seven Avago ADNS5050 optic mouse sensors fitted with Philips CAX100 collimator lenses ($f = 10$ mm) using a custom-designed lens mount that clips directly onto the chip casing. The optics were calibrated such as to maximise the measure of image quality provided by the sensor (i.e. the SQUAL output signal). Each optic flow detector weighs 0.8 g. The complete compound vision system, including the mechanical fixture and the wiring, is illustrated in Figure 8 and weighs 18 g.

3.3 Experimental method

To assess the capability of autonomous take-off, we ran a series of experiments where the aircraft was hand-launched over reasonably flat terrain while *optiPilot* was activated. Since initially the speed is null, the speed regulator will drive the motor to full power. As soon as the aircraft is launched, it will gain speed and climb until it reaches its cruise altitude, where the visual ground avoidance balances the pitch-down tendency due to the pitch bias (see explanation below). To test autonomous landing, we ran experiments where the thrust is manually cut off while the aircraft is autonomously flying above flat terrain. Without propulsion, the MAV will reduce its speed leading to a decrease in the optic flow signal. This




Figure 8: Close-up view of the vision system made of 7 optic flow sensors (Figure 7). The viewing directions of these sensors are pointed at 45° with respect to the longitudinal axis of the aircraft, looking from left to right in the ventral part (eccentricity angle $\hat{\theta} = 45^\circ$ and inter-azimuthal angle $\hat{\psi} = 30^\circ$). The complete compound vision system weighs 18 g.

Table 1: Parameter values.

Parameter	Value
eccentricity $\hat{\theta}$	45°
inter-azimuthal angle $\hat{\psi}$	30°
pitch gain ξ^P	8.1
roll gain ξ^R	8.1
pitch weights w_k^P	according to (4)
roll weights w_k^R	according to (5)
pitch bias	-20%

will result in a decreasing pitch control signal and a descent of the aircraft down to the ground. Movie recordings of both sets of experiments are available online⁴.

We previously showed that values of eccentricity and inter-azimuthal angle of $\hat{\theta} = 45^\circ$ and $\hat{\psi} = 30^\circ$ (corresponding to $N = 7$) are optimal for flight control [5]. The other parameters required by the control strategy are shown in Table 1 and were maintained constant throughout the experiments. In order to keep the aircraft near to the ground, we added a bias of -20% on the elevator deflection. This means that, for a null signal generated by the control strategy, the aircraft has a slight tendency to pitch downward. This value, as well as those of the pitch and roll control gains ξ^P and ξ^R , were manually tuned in flight to produce a response profile that matches the flight dynamics of our MAV.

4 RESULTS

Figure 9 shows the results obtained from five hand-launched take-offs. Time $t = 0$ corresponds to the instant at which the platform is launched. After that, the speed quickly increases to reach the thrust control set-point of 14 m/s while

⁴<http://lis.epfl.ch/microflyers>




Figure 9: Automatic take-off experiments. Data recorded from five flights are shown, during which the aircraft was hand-launched (at time $t = 0$) over reasonably flat terrain. Attitude angles, elevator control signal, airspeed and altitude are plotted along with the average translation-induced optic flow perceived by the aircraft in each viewing directions.

the aircraft adopts a nose-up attitude (as shown by the positive value of the pitch angle). The course of the elevator control signal starts at rest with the constant negative value corresponding to the pitch bias. As soon as the speed increases, the elevator control signal increases in response to the perceived optic flow in the ventral part, which steers the aircraft nose-up. After about 1.5 s of flight, only very small corrections are applied while the aircraft reaches its cruise altitude and speed, resulting in a smooth ascent. Meanwhile, the roll angle is efficiently stabilised after the initial perturbations generated by the manual launch and kept at a level attitude. The average optic flow perceived during the experiment shows strong signal just after take-off, corresponding to the portion of flight where the aircraft is close to the ground. Just after take-off, the optic flow perceived is not symmetrical, indicating that the aircraft was not level around the roll axis. After about 1 s, the distribution of optic flow becomes symmetrical and its amplitude decreases, as *optiPilot* steers the aircraft to a level attitude away from the ground.

Figure 10 illustrates five autonomous landing experiments. At $t = 0$, the landing procedure is triggered by cutting off the motor of the MAV. From this point on, the speed gradually decreased, resulting in a reduction of the perceived optic flow and, therefore, the elevator control signal. As the ground avoidance no longer compensate for the nose-down elevator bias, the aircraft starts a descent. As the ground gets closer, the elevator control signal progressively increased in reaction to the perceived optic flow, which further slows down the air-

craft. By the time the aircraft touches the ground, the elevator is fully deflected, which ensures a low speed (of about 9 m/s) at touch-down and a smooth landing. The average optic flow perceived shows that the distribution of optic flow was kept symmetrical during the whole landing procedure by the roll control. The amplitude gradually increased as the aircraft approached the ground to reach a maximum at touch-down.

It is important to notice that the regulation of both attitude and altitude implicitly derives from a control strategy originally designed for obstacle avoidance [5]. Neither attitude nor altitude are explicitly estimated nor measured by *optiPilot*. Rather, flight stabilisation emerges from the interaction between the ground and the avoidance behaviour that strives to keep obstacles in the ventral region of the aircraft. This contrasts with the typical regulation strategies used by classical autopilots that require explicit estimation of the 6 degree-of-freedom state of the aircraft, at the cost of expensive sensing and processing systems.

5 CONCLUSION

Take-off and landing with manned and unmanned aircraft are, in general, considered to be critical manoeuvres due to the relatively low speed and the proximity of both ground and obstacles. In this paper, we showed that a visual control strategy that relies only on few lightweight and inexpensive sensors (i.e. optic mouse chips, rate gyroscopes and an airspeed sensor) is able to autonomously handle both take-off and landing. This was achieved with a simple sensor-to-




Figure 10: Automatic landing experiments. Data recorded from five flights are shown, during which landings were triggered at time $t = 0$ by cutting off the thrust over reasonably flat terrain. Pitch angle, elevator control signal, airspeed and altitude are plotted along with the average translation-induced optic flow perceived by the aircraft in each viewing direction.

actuator mapping that exploits the properties of translation-induced optic flow and the dynamics of flying platforms, without requiring the estimation of state information such as altitude or attitude angles. Moreover, this exact same control strategy also provides obstacle avoidance and flight stabilisation in cruise conditions [5]. These results thus pave the way towards practical applications where autonomous take-off and landing on irregular terrain or cluttered environments is needed. Future work will include experimentation in situations where man-made and natural obstacles are present on the path of the aircraft during take-off and landing.

ACKNOWLEDGEMENTS

We would like to thank Severin Leven for the development of the embedded processor board [28, 29] and André Guignard for the design and construction of the lens mounts. The authors are also indebted to Severin Leven and Yannick Weibel for improving the manuscript. This work was supported by the Swiss National Science Foundation (grant 200020-116149). Part of this work has been submitted for patenting (application PCT/IB2008/051497).

REFERENCES

- [1] K. P. Valavanis. *Advances in Unmanned Aerial Vehicles*. Springer, 2007.
- [2] C. Fargeon and F. Lefaudeux. UAV systems for security and environmental-related purposes (USEP). special report in UVSI yearbook., 2007.
- [3] Study analyzing the current activities in the field of UAV, 2008. Frost & Sullivan.
- [4] R.W. Beard, D. Kingston, M. Quigley, D. Snyder, R.S. Christiansen, and W. Johnson. Autonomous vehicle technologies for small fixed-wing UAVs. *Journal of Aerospace Computing, Information, and Communication*, 2(1):92–108, 2005.
- [5] A. Beyeler, J.C. Zufferey, and D. Floreano. Vision-based control of near-obstacle flight. *Autonomous Robots*, 2009.
- [6] D. B. Barber, S. R. Griffiths, T. W. McLain, and R. W. Beard. Autonomous landing of miniature aerial vehicles. *Journal of Aerospace Computing, Information, and Communication*, 4(5):770–784, 2007.
- [7] G.L. Barrows, C. Neely, and K.T. Miller. Optic flow sensors for MAV navigation. In Thomas J. Mueller, editor, *Fixed and Flapping Wing Aerodynamics for Micro Air Vehicle Applications*, volume 195 of *Progress in Astronautics and Aeronautics*, pages 557–574. AIAA, 2001.
- [8] W.E. Green, P.Y. Oh, K. Sevcik, and G.L. Barrows. Autonomous landing for indoor flying robots using optic flow. In *ASME International Mechanical Engineering Congress and Exposition, Washington, D.C.*, volume 2, pages 1347–1352, 2003.

- [9] J.S. Chahl, M.V. Srinivasan, and H. Zhang. Landing strategies in honeybees and applications to uninhabited airborne vehicles. *The International Journal of Robotics Research*, 23(2):101–110, 2004.
- [10] F. Ruffier and N. Franceschini. Optic flow regulation: the key to aircraft automatic guidance. *Robotics and Autonomous Systems*, 50(4):177–194, 2005.
- [11] B. L. Stevens and F. L. Lewis. *Aircraft Control and Simulation, 2nd edn.* Wiley, 2003.
- [12] T.C. Whiteside and G.D. Samuel. Blur zone. *Nature*, 225:94–95, 1970.
- [13] A. A. Argyros, D. P. Tsakiris, and C. Groyer. Biomimetic centering behavior for mobile robots with panoramic sensors. *IEEE Robotics and Automation Magazine*, 11(4):21–68, 2004.
- [14] S. Hrabar and G. S. Sukhatme. Optimum camera angle for optic flow-based centering response. In *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, Beijing, China*, pages 3922–3927, 2006.
- [15] R. Wehner. Matched filters - neural models of the external world. *Journal of Comparative Physiology A*, 161:511–531, 1987.
- [16] H.G. Krapp, B. Hengstenberg, and R. Hengstenberg. Dendritic structure and receptive-field organization of optic flow processing interneurons in the fly. *Journal of Neurophysiology*, 79:1902–1917, 1998.
- [17] M. Egelhaaf and R. Kern. Vision in flying insects. *Current Opinion in Neurobiology*, 12(6):699–706, 2002.
- [18] K. Karmeier, J. H. van Hateren, R. Kern, and M. Egelhaaf. Encoding of naturalistic optic flow by a population of blowfly motion-sensitive neurons. *Journal of Neurophysiology*, 96:1602–1614, 2006.
- [19] A. Beyeler, S. Magnenat, and A. Habersaat. Ishtar: a flexible and lightweight software for remote data access. In *Proceedings of the 2008 European Micro Air Vehicle Conference EMAV08*, 2008.
- [20] M.V. Srinivasan. An image-interpolation technique for the computation of optic flow and egomotion. *Biological Cybernetics*, 71:401–416, 1994.
- [21] J.L. Barron, D.J. Fleet, and S.S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, 1994.
- [22] S. Mehta and R. Etienne-Cummings. Normal optical flow chip. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS 2003)*, pages 784–787, 2003.
- [23] R. Moeckel and S.-C. Liu. Motion detection circuits for a time-to-travel algorithm. In *IEEE International Symposium on Circuits and Systems (ISCAS), New Orleans*, pages 3079–3082, 2007.
- [24] D.B. Barber, S. Griffiths, T.W. McLain, and R.W. Beard. Autonomous landing of miniature aerial vehicles. In *AIAA Infotech@Aerospace*, 2005.
- [25] S. Griffiths, J. Saunders, A. Curtis, T. McLain, and R. Beard. Obstacle and terrain avoidance for miniature aerial vehicles. In K.P. Valavanis, editor, *Advances in Unmanned Aerial Vehicles: State of the Art and the Road to Autonomy*, volume 33, chapter I.7, pages 213–244. Springer, 2007.
- [26] A.F. Rodriguez, E. Andersen, J.M. Bradley, and C.N. Taylor. Wind estimation using an optical flow sensor on a miniature air vehicle. In *AIAA Conference on Guidance, Navigation, and Control*, 2007.
- [27] H. Dahmen, A. Millers, and H.A. Mallot. Insect inspired odometry by optic flow recorded with optical mouse chips. In D. Floreano, J.C. Zufferey, M.V. Srinivasan, and C.P. Ellington, editors, *Flying Insects and Robots*. Springer, 2009. in press.
- [28] S. Leven, J.-C. Zufferey, and D. Floreano. A simple and robust fixed-wing platform for outdoor flying robot experiments. In *International Symposium on Flying Insects and Robots*, pages 69–70, 2007.
- [29] S. Leven, J.-C. Zufferey, and D. Floreano. A minimalist control strategy for small UAVs. In *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, St. Louis*, 2009.