

Optical Flow on a Flapping Wing Robot

Fernando Garcia Bermudez
Ronald S. Fearing, Ed.
Bruno Olshausen, Ed.



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2010-96

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-96.html>

June 14, 2010

Copyright © 2010, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

This material is posted here with permission of the IEEE. Such permission does not in any way imply IEEE endorsement of any UC Berkeley product or service. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising/promotional purposes or for creating new collective works for resale/redistribution must be obtained from the IEEE. By choosing to view this material, you agree to all provisions of the copyright laws protecting it.

Optical Flow on a Flapping Wing Robot

Fernando Garcia Bermudez

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval Committee

Ronald S. Fearing, Research Advisor
Bruno A. Olshausen, Second Reader

Optical Flow on a Flapping Wing Robot
Copyright © 2010 Fernando Garcia Bermudez

This report is an adaptation and an extension of work previously published in Garcia Bermudez and Fearing (2009) that is copyrighted © 2009 IEEE. Reprinted with permission.

The contents of Chapter 5 have not been previously published and are hereby licensed under the Creative Commons Attribution 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

Abstract

Optical flow sensing techniques are promising for obstacle avoidance, distance regulation, and moving target tracking, particularly for small mobile robots with limited power and payload constraints. Most optical flow sensing experimental work has been done on mobile platforms which are relatively steady in rotation, unlike the pitching motion expected on flapping wing flyers. In order to assess the feasibility of using optical flow to control an indoor flapping flyer, a 7 gram commercially available ornithopter airframe was equipped with on-board camera and CPU module with mass of 2.5 grams and 2.6 gram battery. An experiment was conducted capturing optical flow information during flapping and gliding flight on the same platform. As expected, flapping introduced substantial systematic bias to the direction estimates to the point of flipping the true direction periodically. Nonetheless, since the optical flow results oscillated at the same frequency as the flapping wings, one could disambiguate the jittering optical flow measurements by correlating these with real-time feedback from the motor current or voltage. Motor Back-EMF was introduced as a filtering signal after corroborating its correlation with the wingstroke position. A reconstructed wingstroke signal was then used to disambiguate the motion estimated in an experiment that constrained the robot to flap down a string. It is envisioned that this technique can be implemented on-board the flapping flyer due to its low computational complexity.

Contents

1	Introduction	1
2	Robotic Ornithopter Platform	3
3	Optical Flow Computation	7
4	Flapping Oscillations Overwhelm Optical Flow Estimates	11
4.1	Experimental Setup and Data Processing Details	11
4.2	Estimation Under Flapping and Gliding	13
4.3	Concluding Remarks	15
5	Using the Motor Back-EMF to Disambiguate Optical Flow Estimates	19
5.1	Experimental Setup and Data Capture Differences	19
5.2	Correlation between Back-EMF and Optical Flow Estimates .	21
5.3	Implementing a Back-EMF-based filter	23
5.4	Concluding Remarks	27
6	Conclusions	33
	Acknowledgments	35
	Bibliography	37

List of Figures

2.1	The flying robot	3
2.2	The image processing board	4
2.3	Block diagram of the image processing board	5
2.4	Dimensions of the robotic platform	5
3.1	Block diagram of an EMD	8
4.1	Block diagram of the image processing	12
4.2	Behavioral diagrams	14
4.3	Subsequent frames from flapping and gliding experiments	15
4.4	Experimental results	17
4.5	Control for the experiments in Figure 4.4	18
4.6	Full flapping cycle of the robot	18
5.1	Samples of raw data captured	20
5.2	Details of the Back-EMF correlation experiment	22
5.3	Results of the Back-EMF correlation experiment	23
5.4	A detail of Figure 5.3.	24
5.5	Back-EMF-based filtering of optical flow signal	25
5.6	Block diagram of the filtering algorithm	26
5.7	Details of the string-sliding experiment	28
5.8	Results of the string-sliding experiment	29
5.9	First control for the string-sliding experiment	30
5.10	Second control for the string-sliding experiment	31

Chapter 1

Introduction

Optical flow vision algorithms for use in robotic sensing have been implemented both in simulation (Cameron et al., 1998) and in robotic platforms that have generally presented steady motion for the camera, such as in the case of wheeled robots (Coombs and Roberts, 1992; McCarthy and Barnes, 2004; Mura and Franceschini, 1996; Srinivasan et al., 2004), fixed wing micro air vehicles (Barrows et al., 2003; Zufferey and Floreano, 2006; Zufferey et al., 2010), airships (Iida, 2003; Zufferey et al., 2006), quad-rotors (Herisse et al., 2010; Zingg et al., 2010), and tethered (Ruffier and Franceschini, 2005) and untethered (Srinivasan et al., 2004) helicopters. One of the most unsteady platforms is a robot that mimics the fly motion (Reiser and Dickinson, 2003), but its movement is constrained to stay within an artificially textured indoor arena. The group whose work is closest to the one presented in this report, (Wagter et al., 2007), proposes using optic flow for estimating the altitude of a flapping vehicle, but to date is mostly simulated, with the real video sequences used to test off-board algorithms having smooth motion.

The use of steady platforms for optical flow experimentation simplifies comparisons of several algorithms on sequences on which they all perform relatively well (Barron et al., 1994; Liu et al., 1998; Mammarella et al., 2008; McCane et al., 2001; McCarthy and Barnes, 2004). In addition, well structured environments also simplify extracting ground truth. There are some platforms that have been shown to work well in the outdoors (Barrows et al., 2003; Srinivasan et al., 2004; Zufferey et al., 2010), but most of the indoor environments used to test real robots use artificially textured walls and objects to improve contrast and thus the performance of optic flow algorithms. Probably the only exception is the work by Zingg et al. (2010), which presents

a quad-rotor visually navigating a heavily textured, yet unmodified, real corridor under normal lighting conditions.

In contrast to robotic optic flow, insects such as flies and bees, which use optical flow for motion detection and navigation (Reichardt, 1987; Srinivasan et al., 1999; Tammero and Dickinson, 2002), perform remarkably well in both outdoor and indoor environments. Their small size and power-to-weight ratio enable them to perform quick maneuvers (Tammero and Dickinson, 2002) deemed impossible for larger platforms such as planes. They are also robust to outside disturbances and the occasional error that sends them crashing onto transparent surfaces like windows. Flapping flight in insects probably increases unsteadiness in their visual input, something they might counteract by actuating their neck muscles to stabilize their gaze (Huston and Krapp, 2008). While a high speed camera or mirror mount could be used to compensate for body motion in flapping robot flight, as well as heavier mechanical solutions, we examine in this report the significance of flapping artifacts in optical flow sensing and then propose the motor’s Back-EMF¹ signal as a proxy for estimating the wingstroke position and disambiguating the oscillatory optical flow by filtering the out-of-phase estimates.

¹Back-electromotive force.

Chapter 2

Robotic Ornithopter Platform

Figure 2.1 shows the flying robot, a modified version of Interactive Toy's VAMP RC ornithopter including custom electronics. The image processing board used to acquire and pre-process the data is pictured in more detail on Figure 2.2. The board weighs 1.1grams, measures 15x35mm and is mainly comprised of a Microchip dsPIC33FJ128MC706 16bit microprocessor running at 40MHz, an OmniVision OV7660FSL VGA camera module, and an ATMEL AT45DB161D 2 megabyte (MB) DataFlash memory. The board was fabricated using a 25 μ m thick FR4 core printed circuit board (PCB).

For wireless communication with a PC, a 1.3g Roving Networks RN-41 Bluetooth 2.0 module of roughly the same dimensions as the image processing board was connected through the dsPIC's serial communication interface. Figure 2.3 shows the block diagram representing these two boards as well as the 70mg motor driving board.

Data was requested from the camera in grayscale at 160x120 (QQVGA) resolution and a rate of 25 frames per second (fps), pre-processed depending on the experimental requirements (detailed in §4.1 and §5.1), and then saved




Figure 2.1: The flying robot, a modified version of Interactive Toy's VAMP RC ornithopter, includes custom electronics used for data acquisition and part of the image processing.




Figure 2.2: The front and back of the image processing board. The OmniVision OV7660FSL camera module is visible at the top, a Microchip dsPIC33FJ128MC706 microprocessor on the right picture, and an ATMEL AT45DB161D flash memory on the left one. It weighs 1.1g and measures 15x35mm.

to the dsPIC’s 16KB random access memory (RAM). (The 2MB ATMEL memory was not utilized in this report.) These data-sets were offloaded to a computer at 230.4 Kbps over the Bluetooth RS-232 link at the end of the acquisition. Note that the custom electronics as well as the robot’s flapping motor are running on a 90mAh FULLRIVER lithium-polymer battery that weighs 2.6g.

The ornithopter uses a DC motor for flapping and steering. In the modified version used in this work, the Vamp’s RC electronics as well as the foam body are removed and the custom motor driving board used just actuates the flapping motor. The robot measures around 35x25cm, as can be seen in Figure 2.4, and weighs 12.6g when unmodified. Normally, as it flies forward at full throttle, it interleaves climbing periods with stall recovery periods at around 1Hz. This, coupled to the fact that it flaps at around 12-17Hz dependent on battery charge, are the main sources of unsteadiness of this platform (see Figure 4.2a).

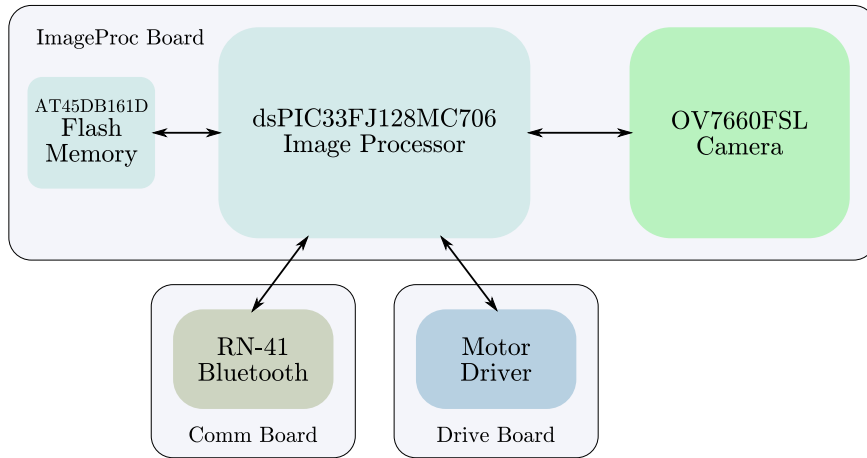


Figure 2.3: Block diagram of the image processing board alongside the Bluetooth communication module and the motor driver board.




Figure 2.4: Dimensions of the robotic platform. Note the camera module positioned to the side of the wing transmission mechanism. The optical axis is aligned with the direction of flight, which in the case of the figure would be the vertical axis.

Chapter 3

Optical Flow Computation

Once the data is acquired and pre-processed on-board, it is wirelessly sent to a PC to be further analyzed using Python¹. Even though this part of the processing is done off-board, the algorithms are still chosen according to the computational complexity that can be implemented on-board this type of hardware, since this is the end goal.

The optical flow algorithm chosen for this work is the standard elementary motion detector (EMD) correlation algorithm (Hassenstein and Reichardt, 1956; Reichardt, 1987), which is not only easy to implement on a fixed-point architecture such as the dsPIC's, but is also considered neurobiologically plausible in insects (Franceschini et al., 1989) and has been used in biological models of the fruit fly (Reiser and Dickinson, 2003; Srinivasan et al., 1999; Tammero and Dickinson, 2002). It's worth noting that EMDs have an inherent capacity for adapting to different velocity ranges, a feature not present in algorithms based on gradient techniques (Borst, 2007). This enables EMDs to work over a wide range of signal-to-noise ratios. Figure 3.1 shows the block diagram of an EMD, and is adapted from Reichardt (1987). Explicitly, (3.1) shows the formulas that this block diagram represents for a local pixel patch transitioning from frame k to $k + 1$:

$$\begin{aligned}u_{i,j}(k) &= I_{i,j}(k+1) \cdot I_{i+1,j}(k) - I_{i+1,j}(k+1) \cdot I_{i,j}(k), \\v_{i,j}(k) &= I_{i,j}(k+1) \cdot I_{i,j+1}(k) - I_{i,j+1}(k+1) \cdot I_{i,j}(k).\end{aligned}\tag{3.1}$$

u and v represent the horizontal and vertical optical flow component

¹Scientific Tools for Python: <http://www.scipy.org/>




Figure 3.1: Block diagram of an EMD (adapted from Reichardt, 1987). Note that, in our case, the delays are represented by consecutive frames in the video sequence captured on-board the platform.

matrices while I is the pixel intensity matrix. A further processing step, (3.2), integrates the fields spatially, summing the motion vectors over each motion field and normalizing by the corresponding Frobenius norm:

$$\begin{aligned}
 U(k) &= \frac{\sum_i \sum_j u_{i,j}(k)}{\sqrt{\sum_i \sum_j |u_{i,j}(k)|^2}}, \\
 V(k) &= \frac{\sum_i \sum_j v_{i,j}(k)}{\sqrt{\sum_i \sum_j |v_{i,j}(k)|^2}}.
 \end{aligned}
 \tag{3.2}$$

The integrated optical flow, U and V , has some information about the overall flow field and thus about the general motion. It is known, though, that the optical flow field is a nonlinear representation of the true 3D motion field and thus doing a linear combination of its vectors will rarely yield accurate results. Nonetheless, neurobiological observations of the fly’s nervous system support this computation (Single and Borst, 1998). This is most probably due to the fact that individual EMD outputs contain not only a direct current (DC) component corresponding to the stimulus motion direction, but also an alternating current (AC) component following the local intensity modulations, which carries no directional information and is phase-shifted with respect to neighboring EMDs (Haag et al., 2004). Thus, spatial integration of many adjacent EMDs is beneficial in that it preserves only the

directional DC component.

It is useful to keep in mind that a slightly more computationally complex pre- and post-processing of the data in and out of an EMD alongside a larger field-of-view, all of which is inspired by biology, result in one of the most efficient motion detection algorithms for use in high-dynamic range natural scenarios (Brinkworth and O'Carroll, 2009).

Chapter 4

Flapping Oscillations Overwhelm Optical Flow Estimates

This chapter introduces the problem that flapping oscillations bring to ego-motion estimation by comparing the EMD output under flapping and gliding. These results were previously published in Garcia Bermudez and Fearing (2009).

4.1 Experimental Setup and Data Processing Details

The modified robot used in the experiments of this chapter weighs 13.6g, which results in a dampening out of the slow climb/stall oscillations mentioned in Chapter 2 because the ornithopter, unable to climb, essentially performs a smooth landing on its body. Thus, the only significant source of unsteadiness in these experiments comes from the flapping. Note that the 13.6g weight includes the 7g airframe, 2.5g of boards, a 2.6g battery, and 1.5g of wiring and mounting hardware.

To reduce storage and transmission requirements, image data was reduced by subsampling and averaging, which resulted in 60 frames of heavily subsampled images at a final resolution of 18x13, comprising 2.4 sec of visual motion data. For applications such as wall following or terrain avoidance, a low resolution such as 18x13 is adequate. (For example, Barrows et al. (2003)

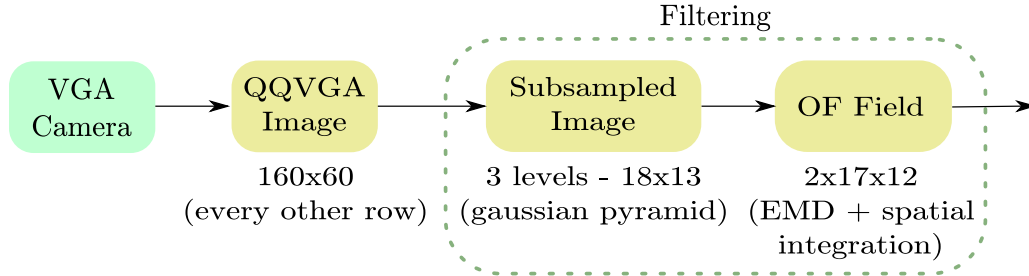


Figure 4.1: Block diagram representing the image processing performed.

used a 1x18 array for ground height regulation.) Averaging also improves the signal-to-noise ratio for the image data.

The first filtering block shown in Figure 4.1 outlines the image pre-processing happening at the camera board that yields the 18x13 frames. Basically, from the 160x120 image that the camera is sending to the dsPIC, the processor captures only every other line, yielding a 160x60 frame. At this point, the processor convolves the image with a 3x3 pixel discrete gaussian filter,

$$f_{x,y} = \begin{pmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{pmatrix},$$

which is equivalent to applying

$$f_x = \begin{pmatrix} 1 & 2 & 1 \end{pmatrix}$$

to the rows and

$$f_y = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}$$

to the columns. Thus, in order to perform the subsampling, f_x is applied three times to each row as they arrive, discarding every other pixel at each step, while f_y is applied only twice to each column as soon as the 60 rows are received and discarding pixels in the same manner, which results in the final image size of 18x13. If instead of the processor only capturing every other row of the input image, it would capture all 120 rows, one would apply f_y once more to the columns and the processing would be equivalent to applying a 3-level gaussian pyramid to the image received as suggested in Figure 4.1.

Note that in this experiment the camera was mounted vertically, thus yielding a 13x18 image. Since the camera field of view is approximately 37° and 50° in the x and y axes respectively, each reduced pixel subtends an angle of 4° and 4.5° . The EMD algorithm was applied to pairs of these pixels in the image both in the horizontal as well as the vertical direction. This yielded two 17x12 motion fields for each frame pair, coming to a total of 59 horizontal and vertical fields. Considering a maximum image shift of 1 pixel and 25 fps, the maximum sensed velocity by the EMDs would be 72 and 70 degrees per second in x and y image plane axes.

4.2 Estimation Under Flapping and Gliding

To assess the effect on the optical flow calculations of the inherent pitch and roll oscillations related to flapping, reduced image data sets were captured while flapping or gliding in an indoor environment. The video sequences were collected using the same hardware in both cases and the flapping motor was either powered on (flapping) or off (gliding). During the flapping experiments, the robot generally flew in a left circular trajectory with roughly a 5m radius (due to a slight weight imbalance) until it landed smoothly on its body. During the gliding flights, upon being launched manually forward, the robot usually climbed up quickly until stalling and then nose-dived into the ground (see Figure 4.2b).

Figure 4.3 shows three consecutive frames of a representative data-set for each experiment. These frames have the optical flow field overlaid on top of them as well as the integration result at the center of each frame. As one can visualize in the figure, the inferred direction that the optical flow integration outputs varies smoothly in the gliding experiment whereas it switches abruptly, frame to frame, in the flapping experiment. In the case of the gliding frame sequence, the inferred direction is that of motion of the robot with it's nose diving into the ground. For the flapping frame sequence, the inferred direction is only correct in the outer two frames, since the robot is circling around that direction. The middle frame indicates the opposite motion most probably due to the flapping induced pitch oscillations, which introduce substantial fluctuations to an otherwise smooth circular trajectory of the robot. Although the gliding and flapping trajectories were quite different, the lighting conditions were almost the same since all experiments were done at the same time of the day and in the same indoor environment.

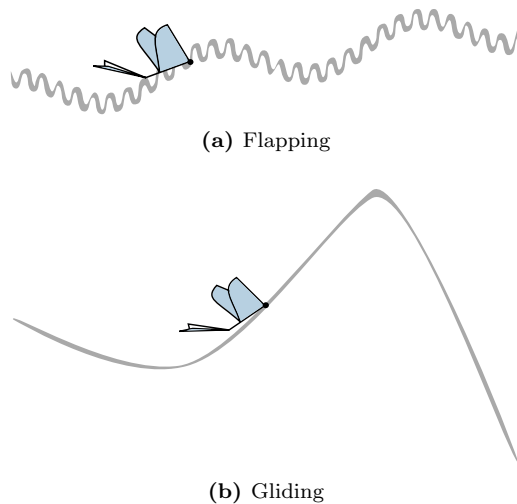



Figure 4.2: Behavioral diagrams: (a) robot’s behavior while flapping, with a slow climb/stall frequency and a faster flapping frequency; (b) robot’s behavior while gliding, with a fast climb until stalling and a sharp nose-dive into the ground. Note the black dot where the robot intersects the trajectory as it traverses it. This is where the camera is positioned during flight with its axis pointing in the flight direction.


We claim that the erroneous optical flow integration results for the flapping experiment are indeed due to the oscillations induced by flapping. To prove this point, the time-varying normalized vector signals $(U, V)^T$ were first passed through a Hann window of length 59 and then processed under a discrete Fourier transform, resulting in the plots of Figure 4.4.

As is evident from looking at the flapping results in Figure 4.4a, the optic flow vectors are oscillating at around 11-12Hz. There seems to be a small oscillation at around 2Hz, which could be explained as being related to the damped climb/stall cycle. It could also be related to the relatively short capture period of 2.4 sec, since a few coincidental events during this period can seem like a slow oscillation. This is in fact what can be seen in the gliding results in Figure 4.4b, since in this case it is known that the capture took place just as the robot was reaching the maximum altitude, stalling, and recovering from the stall. Thus, the bump around 1Hz most probably comes from that single event during the 2.4 sec of capture.

To estimate the noise present in the camera, image capture, and optical flow estimation process, a control experiment that consisted of capturing a still image sequence under the same lighting conditions was performed, and the results of it are included in Figure 4.5. According to this figure, the level of error present in the system is around 10% with no camera motion,



(a) Flapping



(b) Gliding

Figure 4.3: Subsequent frames for both the flapping (a) and gliding (b) experiments, which include the optical flow field overlaid as well as the integration result at the center. Note that this central arrow changes direction much more smoothly in the gliding experiment, while giving abrupt changes in direction for the flapping experiment. This behavior is consistent throughout the captured video sequences for each type of experiment.


and thus argues that almost everything other than the larger peaks in the resulting frequency spectrums might be noise. The source of the peak at 6 Hz is not known, but this component is small compared to wing flapping or slow turning peaks.

In order to verify that during the flapping experiment the optical flow algorithm result was indeed oscillating at the frequency that the robot was flying at, the flapping trajectory was captured on high speed video. Figure 4.6 shows a representative sequence of frames depicting a full flapping cycle of the robot during the same flapping experiment analyzed above. The total number of frames was 24, spanning 80ms if one takes into account that the video was captured at 300fps. This would indicate that at that point the robot was flapping at a low frequency of 12.5Hz, most probably due to low battery charge during the experiment. If one performs this same analysis at different positions throughout the trajectory, the same frequency is found. From the high speed video, the pitch range induced by flapping is estimated to be $\pm 5^\circ$.


4.3 Concluding Remarks

In order to separate the pitch oscillation from the optical flow direction estimates when flapping, we propose to concurrently capture the motor current

or voltage alongside the video sequence so as to later correlate optical flow integration errors to specific current or voltage profiles due to cyclic wing loading conditions. For example, the images could be captured in phase with the wing motion at top-dead-center or bottom-dead-center of the wing trajectory. This would enable almost exact nulling of the pitch rate disturbance, for example by calculating optic flow from pairwise frames $I_{i,j}(k)$ and $I_{i,j}(k + 2)$.



(a) Flapping



(b) Gliding

Figure 4.4: Experimental results. In each: (Top) components of the optical flow integration vector during the span of the captured data; (Bottom) single-sided amplitude spectrum of the above signal. Note: (a) the large peak at around 11-12Hz and the smaller peak around 2Hz; (b) the peak around 1Hz.




Figure 4.5: Control for the experimental results shown in Figure 4.4. Note that the error is around 10% of the previously plotted signals.




Figure 4.6: A sequence of frames representing a full flapping cycle of the robot.

Chapter 5

Using the Motor Back-EMF to Disambiguate Optical Flow Estimates

This chapter presents an investigation into how well the Back-EMF is correlated with the integrated optical flow signal and then proposes this signal as a way to eliminate, through filtering, the flapping-induced oscillations from the overall motion estimated.

5.1 Experimental Setup and Data Capture Differences

The robot differs from that of the experiment in Chapter 4 in that the camera is oriented downward 45° from the flight direction. This was modified in preparation for performing height regulation using optical flow on this platform and will also be important in §5.3, which presents an experiment that combines speed and height estimation.

To simplify analysis, just a 160-pixel column was kept from each image, making estimation one-dimensional. The column captured spanned the center of the vertically-oriented camera. To enable yet more flexibility in post-processing, these columns weren't subsampled on-board, as in Chapter 4, which resulted in 85 columns from subsequent frames being accumulated in the dsPIC's RAM before the experiments finished. This yielded 3.4 sec of video of these columns that were later sent off-board.




Figure 5.1: Two samples of raw data from the on-board camera. Both of these are composed of 85 frames, each of which is a single 160-pixel column. The sequence in (a) was captured while sliding down a string, while the one in (b) shows no motion other than the flapping-induced oscillations. Time is on the horizontal axis, spanning a total of 3.4 sec.

Figure 5.1 shows a couple of raw data samples. Figure 5.1a corresponds to an experiment where the camera accelerated up to a constant forward speed while getting closer to a texture. Some of the jaggedness shown here corresponds to flapping-induced oscillations (the rest is just pixelation). Figure 5.1b, on the other hand, represents the output when the robot is held stationary in translation and the only motion observed are the flapping-induced oscillations. These datasets will be studied more in depth in §5.3 and §5.2, respectively.

Alongside the image data, the motor’s Back-EMF was sampled at 250Hz. As is known, Back-EMF is an indirect measure of motor speed, which in turn should be correlated to the wingstroke and thus the robot’s slider-crank transmission position. Thus, the oscillations induced by flapping should be present in the Back-EMF signal. To synchronize the Back-EMF measurements with the images captured by the camera, the latest Back-EMF sample was noted both before and after the column was sampled. The fact that the Back-EMF measurements are spaced by 4 ms means that both signals can be synchronized only up to ± 2 ms.

5.2 Correlation between Back-EMF and Optical Flow Estimates

The experiment designed to explore the relationship between the two signals while flapping consisted in capturing the subsequent columns and the Back-EMF while holding the robot stationary. The on-board column sequence captured in this experiment is shown in Figure 5.1b. The only motion allowed was the oscillation induced by flapping. No translation or rotation of another kind was observed. The oscillatory motion was filmed at 1200fps using a Casio EX-F1 HS Exilim camera looking directly at the slider-crank mechanism that drives the wings, as can be seen in Figure 5.2a. The position of the crank in time was extracted visually and synchronized with the signals captured on-board by means of an LED that toggled at the start and end of each captured frame. The instants when two distinct extreme positions of the crank, denoted by the angles of 0° and 180° , as shown in Figure 5.2c and 5.2d, were recorded and plotted alongside the on-board signals. Figure 5.2b shows the on-board camera viewpoint.

The plot shown at the top of Figure 5.3 shows the integrated optical flow in the vertical direction, $V(t)$, as a stem plot. Given that each image is synchronized with the Back-EMF, each of the optical flow measurements was centered between the synchronization pulses of the two images involved in its calculation. The crank-angle signal plotted at the bottom of Figure 5.3, on the other hand, was synchronized to the start of the first image captured, using the corresponding Back-EMF synchronization pulse as the offset. Note that only the crank-angle extremes were extracted from the high-speed video because we want to highlight the highest velocity points. Lastly, the Back-EMF is presented both raw and low-pass filtered at the middle part of Figure 5.3. The low-pass filter used is based on a 3^{rd} order Butterworth filter with a cutoff frequency of 100Hz. To prevent phase delay after the application of this filter causally, the function `filtfilt`, present in the signal processing toolboxes of both Python's Scipy and Matlab, was used to apply the filter both causally and noncausally, eliminating phase delay, but doubling the filter's order in the process.

To bring attention to certain specifics, Figure 5.4 presents a detail of Figure 5.3. In particular, the optical flow estimate is generally negative during the upstroke (180° crank angle) and positive during the downstroke (0° crank angle). This level of correlation is expected because the only motion

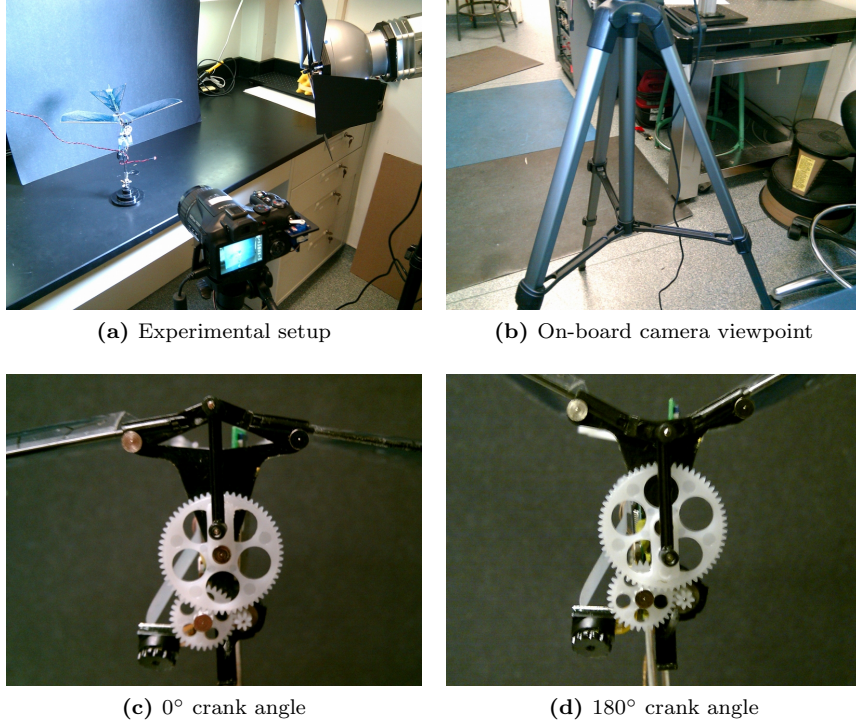


Figure 5.2: Details of the Back-EMF correlation experiment. The two extreme crank angles, pictured in (c) and (d), were recorded from the high-speed video sequence and then plotted against the Back-EMF and the integrated optical flow signals.

of the platform is due to the flapping-induced oscillations. On the same note, the filtered Back-EMF is minimal whenever the crank angle is at 90° and -90° and maximal whenever it is at one of the two extremes of 0° and 180° . The last detail to observe is the slight difference in the peak value of the Back-EMF cycles related to the upstroke and downstroke respectively. This is slightly more apparent in Figure 5.3.

To further point out that the crank-angle signal can be inferred from the Back-EMF, the flapping frequency (which was estimated to be 6.47Hz, half of the Back-EMF's fundamental frequency, in this experiment) and phase (estimated from the location of the higher peaks) was extracted from the Back-EMF signal and then used for a naïve sinusoidal reconstruction of the crank-angle signal shown at the bottom of Figure 5.3. A slightly modified version of this reconstructed signal could be used as a basis for filtering the integrated optical flow signal and getting rid of the flapping-induced




Figure 5.3: (Top) The integrated optical flow signal, $V(t)$, is plotted synchronized to the (Middle) Back-EMF and the (Bottom) crank-angle signals. Note that the Back-EMF signal is plotted both raw and low-pass filtered. The bottom plot has not only the visually extracted crank angles as a stem plot, but also a reconstructed sinusoid estimated from the fundamental frequency of the Back-EMF signal and the phase of its larger peaks.

oscillations by simulating phase-locked sampling of video data. This idea will be expanded on in the following section.

5.3 Implementing a Back-EMF-based filter

In the previous section, it was shown that a sinusoidal reconstruction of the flapping signal was possible just by analyzing the Back-EMF to extract the flapping frequency and phase. A more accurate representation of this signal could be inferred by computing a running estimate of the frequency and phase, thus accounting for speedups and slowdowns due to experimental




Figure 5.4: A detail of Figure 5.3.

dynamics. This would yield an asymmetric sinusoid which still wouldn't perfectly fit the actual flapping signal, due to dynamic effects that cannot be represented just by a single cosine term. Leaving these ideas for future work, this section will strive to show how an estimated flapping/crank-angle signal could serve as a filter to the integrated optical flow one, eliminating the flapping-induced oscillations from the estimated motion.

As a first-order approximation to this filtering signal, we scale the reconstructed flapping signal to span from 0 to 1 as the crank rotates from 0° to 180° and back. The idea is that when the integrated optical flow signal, $V(t)$, is multiplied by this filtering signal, the optical flow estimates closer to the crank angle of 180° will remain mostly untouched while the ones closer to a crank angle of 0° will tend to disappear. A further processing step, not taken in this work, could consist of interpolating the signals closer to 180° to fill-in the samples previously eliminated. This process would aim to reproduce in




Figure 5.5: Integrated optical flow signal alongside a filtered version of itself. This corresponds to the experiment first introduced in Figure 5.3.

software the equivalent of image capture phase-locked to a certain position in the robot’s flapping cycle at a higher frequency than that of flapping.

One of the main issues to overcome is the fact that the robot’s flapping frequency can be as high as 17Hz, dependent on the duty cycle the flapping motor is ran at. If the flapping frequency is higher than 12.5Hz, since our on-board image capture runs at 25fps, the resulting integrated optical flow signal will be aliased. To overcome aliasing, we pass the Back-EMF-filtered signal through a Kaiser window of $\beta = 3$ and then filter it again with a 2.5Hz cutoff frequency 3rd order Butterworth filter both causally and noncausally, using `filtfilt`. This results in the $V_{filt}(t)$ signals which are plotted next to the unfiltered integrated optical flow signals. Figure 5.5 shows the filtered signal corresponding to the experiment discussed in §5.2 and previously plotted in Figure 5.3. Since this experiment only presented oscillatory flapping-




Figure 5.6: Block diagram of the filtering algorithm based on the Back-EMF signal that was applied to the integrated optical flow estimates in order to get rid of the flapping-induced oscillations.

induced motion, because translation was constrained, one would expect that after cancelling the flapping-induced oscillations, the filtered signal should be almost null. A Block diagram of the filtering algorithm is included in Figure 5.6.

To further assess the usefulness of this filtering strategy, an experiment that consisted in allowing the robot to flap down a constraining string as it got nearer to a texture was performed. The experimental setup used can be seen in Figure 5.7 and consists in the robot sliding down a string, which was stretched at a -8° angle, against a striped background. The on-board camera took sequential frames of which only one column was kept, resulting in the translating striped pattern shown on Figure 5.1a.

A typical run of this experiment would result in the plots seen in Figure 5.8. The top plot shows the filtered integrated optical flow signal and can be divided into three sections. The first 1s of data represents transient oscillations of the robot on the string prior to the start of its slide down the incline. From then until $t = 2$ sec, the robot is sliding at a fairly constant speed down the string and getting closer to the texture, thus resulting in an increasing optical flow signal. The last 1.5 sec culminate the run with the robot slowing down as it gets closer to the texture. Ground truth data for both velocity and height over time has been included for comparison.

Thus, Figure 5.8 clearly shows the expected output for this experiment, but in order to make sure the signals make sense we'll resort to two control experiments. The first one consists on capturing video data while sliding the robot down the string with its motor off. The robot was launched manually down the string and the results can be seen in Figure 5.9. The first 0.5s correspond to the winding-up of the robot right before it's launched down the string. This is why for a half second the plot shows negative speed, as the robot goes slightly up the string, and then the speed quickly peaks as it starts sliding down and slows down for the rest of the run until it stops. The


ground truth velocity data supports these observations.

The second control experiment was even simpler and consisted in leaving the robot stationary on the string looking into a patch of texture. Processing these still images give the noise level present due to only the image capture and optical flow processing. The results can be seen in Figure 5.10. Note that even though none of these controls were filtered using the Back-EMF, since it wasn't present in these experiments, $V(t)$ was still passed through the same Kaiser window and low pass filter, so as to better compare them to previous results.


Lastly, note that even though the filtered integrated optical flow signal shows correlation to the ground truth velocity, it is debatable how much the height variation is affecting its output. In fact, if you observe the raw data in Figure 5.1a, it can be seen that the spatial frequency doesn't significantly increase as height decreases throughout the experiment. This means that the expansion signal is either small or even negligible and that for it to be significant, the height difference would need to be more drastic or the ground texture of higher spatial frequency.

5.4 Concluding Remarks


The set of experiments presented in this chapter not only show the correlation between the Back-EMF signal and the flapping-induced oscillations, which was used to construct an estimated flapping signal, but also show that this reconstructed signal can effectively be used to filter out the flapping-induced oscillations, leaving the rest of the motion content present in the filtered integrated optical flow signal. This approach is promising for on-board integration into the robotic platform due to its fairly low computational requirements. More experiments are required to guarantee that the procedure can also recover expansion information for use in height regulation, due to the shallow angle of the string and the low spatial frequency of the ground texture in these experiments.



(a) Experimental setup



(b) Diagram of the experimental setup (not to scale, dimensions are in mm)



(c) On-board camera viewpoint

Figure 5.7: Details of the string-sliding experiment.

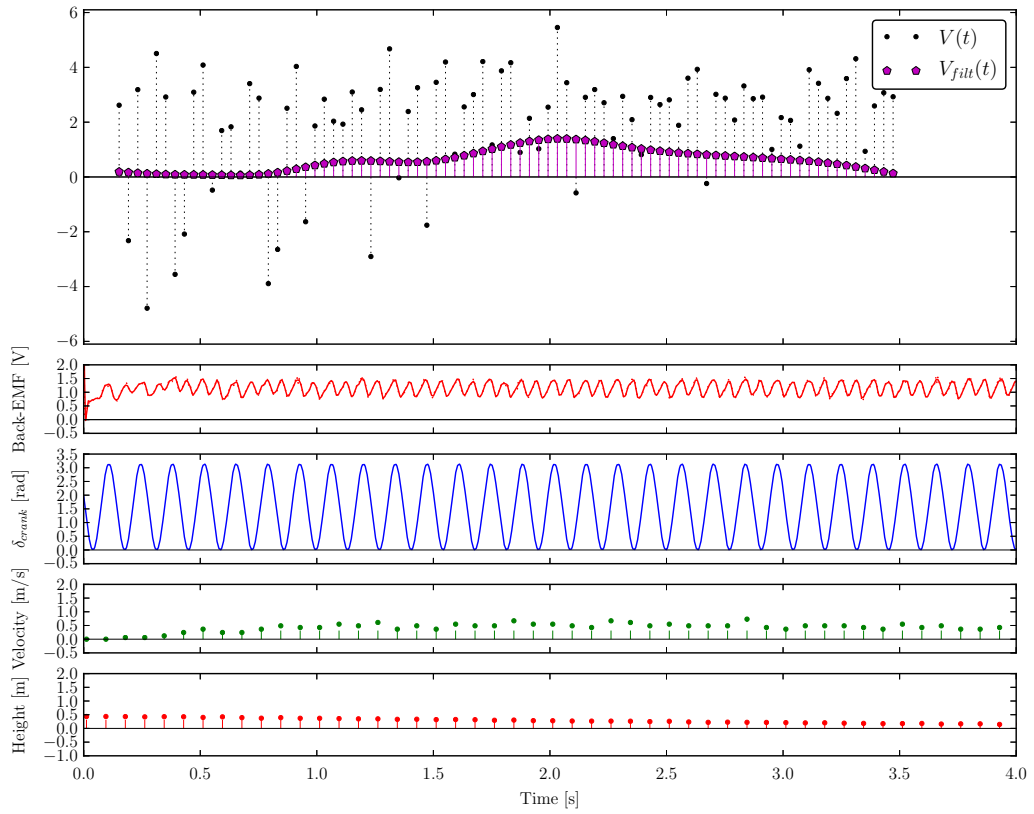


Figure 5.8: Results from a typical run of the string-sliding experiment. Note the increasing signal starting roughly a second into the experiment that represents both translation and expansion of the video sequence. Ground truth of both velocity and height over time was extracted from the high-speed video and plotted for comparison. For details about the signals plotted, please refer to the caption of Figure 5.3.

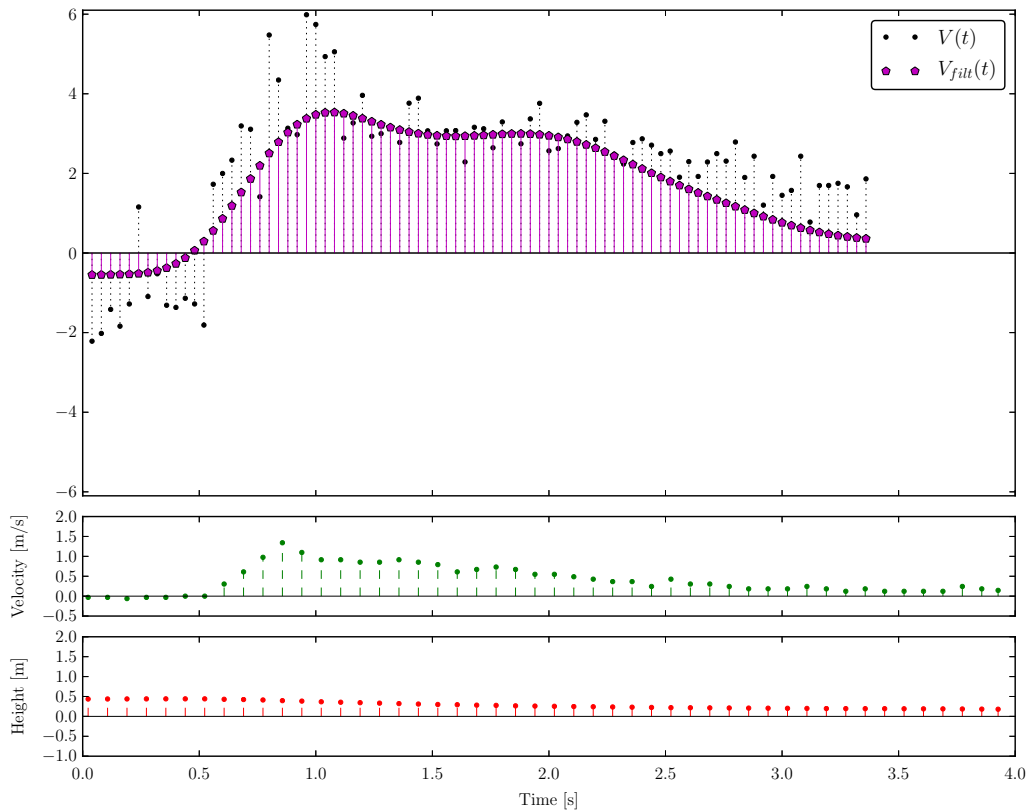


Figure 5.9: First control for the string-sliding experiment shown in Figure 5.8. It consisted in launching the robot manually down the string with its motor turned off. This should result in a signal not affected by flapping-induced oscillations. Note that the first negative section up to 0.5s represents motion up the string as the robot is winded-up right before the launch. After that, a peak can be clearly seen as the robot speeds up to its maximum velocity down the string and then slows towards a stop. The ground truth velocity clearly represents this same behavior.




Figure 5.10: Second control for the string-sliding experiment shown in Figure 5.8. This one consisted in leaving the robot stationary on the string while looking at a patch of texture. The output gives an idea of the error present in the image capture and optical flow processing.

Chapter 6

Conclusions

An order 10 gram robot ornithopter was constructed using a commercial platform combined with a lightweight cell phone camera and wireless interfaces. Subsampled, low resolution video data was captured during flapping or gliding flight and processed off-board. This experiment, using a simple yet efficient biomimetic optical flow algorithm which extracted net motion direction by averaging the flow field across the whole sensor, demonstrated the significance of pitch oscillations due to wing flapping on the optical flow direction estimates.

The small ornithopter used here demonstrates the coupling between body motion and optic flow sensing which can be expected without image stabilization mechanisms. The strong optical flow signal corresponding to the wing flapping frequency appears readily separable by a notch filter or synchronized sampling. Hence, active visual stabilization that might be used by insects such as flies does not appear critical.

We added motor Back-EMF measurements and corroborated that they were well correlated with the optical flow estimates by constraining the robot's motion to just oscillate due to flapping. The signal was then processed to infer the crank-angle/flapping signal. By directing attention to optical flow measurements that happen in phase with each other, one could discard out-of-phase optical flow estimates and fill-in interpolated estimates based on the ongoing rate of change.

The reconstructed crank-angle signal was used to effectively filter out flapping-induced oscillations on data from the robot flapping down a string, which resulted in the preservation of the rest of the motion content. This approach is promising for on-board integration into the robotic platform

due to its fairly low computational requirements. In future work, optical flow information will be used for robot steering in behaviors such as height regulation, wall following, and obstacle avoidance.

Acknowledgments

The author would like to thank Prof. Ronald Fearing for his support and guidance on the work leading to this report and for his expert advice on digital signal processing. He would also like to thank him and Prof. Bruno Olshausen for their insightful comments and editing that helped improve this report. The author is further indebted to Erik Steltz for his advice during the design, construction and debugging of the image processing board, Aaron Hoover and Stanley Baek for contributing to and debugging the code running on the image processing board, Paul Birkmeyer for his help with filming the experiments of Chapter 4, Kevin Peterson for a series of insightful discussions on signal analysis, Prof. Jitendra Malik for his advice on discrete image filtering, Fred Cheng from OmniVision for his prompt technical support for the camera module, Todd Templeton for his example and advice on obtaining and displaying permission from the IEEE for previously published material as well as the advice and support of members of the Biomimetic Millisystems Lab and the EECS community at UC Berkeley. The author wishes to especially thank his wife, Yanina Bogliotti, for her unconditional support, without which this report would not have been possible. This work supported by NSF Grants IIS-0705249 and CNS-0931463.

Bibliography

- 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.
- G. L. Barrows, J. S. Chahl, and M. V. Srinivasan. Biomimetic visual sensing and flight control. *The Aeronautical Journal*, 107(1069):159–168, March 2003.
- Alexander Borst. Correlation versus gradient type motion detectors: the pros and cons. *Phil. Trans. R. Soc. B*, 362(1479):369–374, March 2007. doi: 10.1098/rstb.2006.1964.
- Russell S. A. Brinkworth and David C. O’Carroll. Robust models for optic flow coding in natural scenes inspired by insect biology. *PLoS Comput Biol*, 5(11):e1000555, November 2009. doi: 10.1371/journal.pcbi.1000555.
- Seth Cameron, Stephen Grossberg, and Frank H. Guenther. A self-organizing neural network architecture for navigation using optic flow. *Neural Computation*, 10(2):313–352, February 1998.
- David Coombs and Karen Roberts. Bee-bot: using peripheral optical flow to avoid obstacles. In *Intelligent Robots and Computer Vision XI: Algorithms, Techniques, and Active Vision*, volume 1825, pages 714–721. SPIE, November 1992.
- N. Franceschini, A. Riehle, and A. Le Nestour. Directionally selective motion detection by insect neurons. In D. G. Stavenga and R. Hardie, editors, *Facets of vision*, pages 360–390. Springer, Berlin and New York, 1989.
- Fernando Garcia Bermudez and Ronald Fearing. Optical flow on a flapping wing robot. In *Intelligent Robots and Systems, 2009. IROS 2009*.

IEEE/RSJ International Conference on, pages 5027–5032, St. Louis, MO, USA, October 2009. doi: 10.1109/IROS.2009.5354337.

- J. Haag, W. Denk, and A. Borst. Fly motion vision is based on Reichardt detectors regardless of the signal-to-noise ratio. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16333–16338, November 2004. doi: 10.1073/pnas.0407368101.
- Von B. Hassenstein and W. Reichardt. Systemtheoretische analyse der zeit-, reihenfolgen- und vorseichenauswertung bei der bewegungsperzeption des rüsselkäfers *Chlorophanus*. *Zeitschrift für Naturforschung B*, 11(9):513–524, September 1956.
- Bruno Herisse, Sophie Oustrières, Tarek Hamel, Robert Mahony, and Francois-Xavier Russotto. A general optical flow based terrain-following strategy for a vtol uav using multiple views. In *Robotics and Automation, 2010. ICRA 2010. IEEE International Conference on*, Anchorage, Alaska, May 2010.
- Stephen J. Huston and Holger G. Krapp. Visuomotor transformation in the fly gaze stabilization system. *PLoS Biology*, 6(7):1468–1478, July 2008. doi: 10.1371/journal.pbio.0060173.
- Fumiya Iida. Biologically inspired visual odometer for navigation of a flying robot. *Robotics and Autonomous Systems*, 44(3-4):201–208, September 2003.
- Hongche Liu, Tsai-Hong Hong, Martin Herman, Ted Camus, and Rama Chelappa. Accuracy vs. efficiency trade-offs in optical flow algorithms. *Computer Vision and Image Understanding*, 72(3):271–286, December 1998.
- Marco Mammarella, Giampiero Campa, Mario L. Fravolini, Yu Gu, Brad Seanor, and Marcello R. Napolitano. A comparison of optical flow algorithms for real time aircraft guidance and navigation. In *Guidance, Navigation and Control Conference and Exhibit*, Honolulu, Hawaii, August 2008. AIAA.
- B. McCane, K. Novins, D. Crannitch, and B. Galvin. On benchmarking optical flow. *Computer Vision and Image Understanding*, 84(1):126–143, October 2001.

- Chris McCarthy and Nick Barnes. Performance of optical flow techniques for indoor navigation with a mobile robot. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, New Orleans, LA, April 2004. doi: 10.1109/ROBOT.2004.1302525.
- F. Mura and N. Franceschini. Obstacle avoidance in a terrestrial mobile robot provided with a scanning retina. In *Intelligent Vehicles Symposium, 1996., Proceedings of the 1996 IEEE*, pages 47–52, Tokyo, Japan, September 1996. doi: 10.1109/IVS.1996.566350.
- Werner Reichardt. Evaluation of optical motion information by movement detectors. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 161(4):533–547, July 1987.
- Michael B. Reiser and Michael H. Dickinson. A test bed for insect-inspired robotic control. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 361(1811):2267–2285, October 2003.
- Franck Ruffier and Nicolas Franceschini. Optic flow regulation: the key to aircraft automatic guidance. *Robotics and Autonomous Systems*, 50(4): 177–194, March 2005.
- Sandra Single and Alexander Borst. Dendritic integration and its role in computing image velocity. *Science*, 281(5384):1848–1850, September 1998.
- Mandyam V. Srinivasan, Michael Poteser, and Karl Kral. Motion detection in insect orientation and navigation. *Vision Research*, 39(16):2749–2766, August 1999.
- Mandyam V. Srinivasan, Shaowu Zhang, Javaan Chahl, G Stange, and M Garrat. An overview of insect-inspired guidance for application in ground and airborne platforms. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 218(6):375–388, January 2004. doi: 10.1243/0954410042794966.
- Lance F. Tammero and Michael H. Dickinson. The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. *The Journal of Experimental Biology*, 205:327–343, 2002.

- C. De Wagter, B. Bijnens, and J.A. Mulder. Vision-only control of a flapping mav on mars. In *AIAA Guidance, Navigation and Control Conference and Exhibit*, Hilton Head, SC, August 2007.
- Simon Zingg, Davide Scaramuzza, Stephan Weiss, and Roland Siegwart. Mav navigation through indoor corridors using optical flow. In *Robotics and Automation, 2010. ICRA 2010. IEEE International Conference on*, Anchorage, Alaska, May 2010.
- Jean-Christophe Zufferey and Dario Floreano. Fly-inspired visual steering of an ultralight indoor aircraft. *IEEE Transactions on Robotics*, 22(1): 137–146, February 2006.
- Jean-Christophe Zufferey, Alexis Guanella, Antoine Beyeler, and Dario Floreano. Flying over the reality gap: From simulated to real indoor airships. *Autonomous Robots*, 21(3):243–254, November 2006. doi: 10.1007/s10514-006-9718-8.
- Jean-Christophe Zufferey, Antoine Beyeler, and Dario Floreano. Autonomous flight at low altitude with vision-based collision avoidance and gps-based path following. In *Robotics and Automation, 2010. ICRA 2010. IEEE International Conference on*, Anchorage, Alaska, May 2010.