

A 6000 Hz Computer Vision System for *Real-Time* Wing Beat Analysis of *Drosophila* *

Chauncey F. Graetzel^{†‡}, Steven N. Fry[‡], and Bradley J. Nelson[†]

[†]*Institute of Robotics and Intelligent Systems*
ETH Zurich
8092 Zurich, Switzerland
cgraetzel@ethz.ch and bnelson@ethz.ch

[‡]*Institute of Neuroinformatics*
ETH Zurich
8057 Zurich, Switzerland
steven@ini.phys.ethz.ch

Abstract—This paper reports the design and application of a novel high speed computer vision system for real-time analysis of the wing kinematics of tethered flying fruit flies (*D. melanogaster*). The system uses a camera with dynamic regions of interest (ROI) to increase temporal resolution from localized sampling, without loss of spatial resolution. An extended Kalman filter is employed to fit an *a priori* kinematic model to past wing position measurements, allowing the position of the next ROI to be predicted and providing a real-time readout of kinematic data. Using this approach, we sampled the wing position at 6250Hz with a precision of 1° , using a ROI of approximately 3600 pixels. This is more than four times faster than other computer vision based tracking system to date. Beyond the study of insect flight control, this paper demonstrates a novel approach to track complex and fast moving structures in real time applications, a challenge often faced in micro and nano technologies.

Index Terms—Biomechanics, tracking, machine vision, Kalman filtering, control systems.

I. INTRODUCTION

Interest in micro-scale biological phenomena has substantially increased in recent times. This interest arises from two complementary sides. On the one hand, biologists are aspiring to link findings obtained at the molecular level with larger scale phenomena. On the other hand, engineers are increasingly turning to small organisms to gain design principles for future microrobots.

As dimensions are scaled down, biomechanical analysis is generally challenged by the reduced interaction forces and time constants [1]. Recent advances in microelectromechanical systems (MEMS) and imaging technologies, as well as the increased computational power, offer new opportunities for conducting biomechanical research at the micro-scale.

The fruit fly (Fig. 1) offers a particularly interesting model for exploring biomimetic design principles in microrobots, due to its precise wing actuation at the millimeter scale [2] [3] and highly specialized neuro-motor control system [4]. Various techniques have been applied to study the flight biomechanics in this species, including upscaled robotic models [5]–[7], quasi steady-state models [7] [6] or computational fluid dynamic models [8]. These approaches rely on precise wing kinematic measurements of fruit flies.

*This work is supported by the Swiss National Science Foundation grant #205321-103910




Fig. 1. Fruit fly: *Drosophila melanogaster*

In view of the detailed understanding of the neural control of flight, behavioral experiments are of imminent importance. A powerful technique to explore sensory-motor control loops is the use of tethered flight “simulators”, in which measurements of the flight behavior are used to create a sensory (e.g. visual) feedback loop [9] [10]. This requires real-time kinematic measurements, for which we provide a novel approach.

II. STATE OF THE ART

A. Previous wing beat analyzers

In prior efforts, the measurement of wing motion has been performed using an infrared light that projected the shadow of each wing through a crescent mask onto a photodiode receptor [11] [12]. While this setup can measure the relative stroke amplitude of both wings, it is fundamentally limited by the reduction of a three degree of freedom movement, i.e. the wing angles, to a single dimensional analog value, i.e. the photodiode readout voltage. Furthermore, construction of the analogue hardware and the requirement for a precise alignment of the fly prior to measurement have hindered a more wide-spread use of this technique. Our new system based on computer vision has the advantage of being highly flexible, allowing additional functionality, such as automatic calibration features to be implemented. In addition, the two dimensional view combined with prior knowledge of an organism’s dimensions can be used to obtain new metrics, such as a full 3D reconstruction of the wings’ positions.

B. Offline wing beat analysis

Due to the limitations of prior wing beat analyzers, detailed kinematics have been measured using off-line procedures with previously acquired flight sequences. In the fruit fly, Zanker et al. [13] used stroboscopic techniques with a mirror setup to reconstruct wing positions from multiple views. Fry et al. [2] recorded flight sequences at 5000Hz and applied a computer-assisted manual procedure to fit 3D wing models to three views of the fly. These experiments showed that the typical wing kinematics are well characterized by at least 10 samples per stroke, which at a wing beat frequency of around 200Hz, requires a sampling rate of close to 2000Hz.

C. Real-time high speed computer vision

Real-time measurements are necessary in closed-loop behavioral experiments. Recently, there have been significant efforts to design tracking systems that successfully operate at high frequencies (see Fig. 2). The fundamental problem related to real time image processing is the large amount of data generated, which must be transferred from the camera and processed with minimal delays. The various approaches to reduce these bandwidth problems can be divided into three categories.

First, VLSI vision chips with a processing element embedded in each pixel have been used [14]–[16]. Such architectures integrate computing power within the acquisition system. Image pre-processing steps, such as linear convolutions [14], have been computed in parallel by this type of approach. Alternatively, the amount of data transferred can be reduced by only sending out the location of the pixels that have undergone a high variation in relative intensity [17]. The limited pixel array size and the difficulty in obtaining packaged products of such architectures has restricted their practical use.

Second, arrays of off-the-shelf camera systems have been used [18] [19]. By increasing the number of acquisition systems, the effective frame rate is increased. The images are synchronized both in time and space. For applications at the sub-millimeter scale, however, it is not always practical to arrange many acquisition systems in an environment that is usually highly constrained.

Third, camera manufacturers have recently marketed high speed acquisition systems with dynamic regions of interest, i.e. the ROI and acquisition parameters can be updated from frame to frame. Compared to other systems, the tradeoff between spatial and temporal resolution is minimized. However, such systems decrease the field of view (FOV). Consequently, the ROI must be repositioned to follow the object of interest, which in turn necessitates additional computational costs.

III. SYSTEM OVERVIEW

A. Experimental Setup

The experiments were performed on an optical table (see Fig. 3). Flies were immobilized on a cooling stage and glued to a tungsten probe using UV cured glue (Loctite,

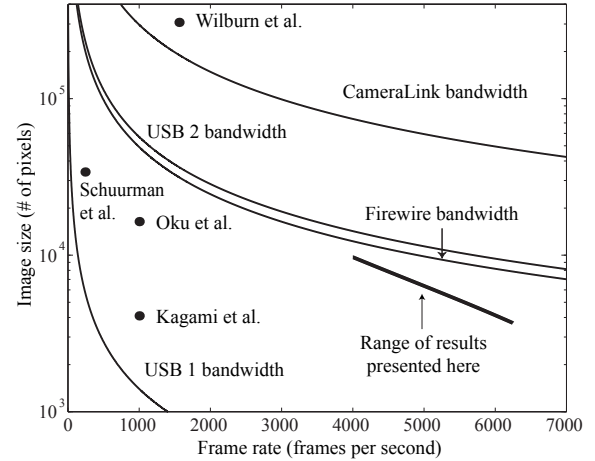


Fig. 2. Comparison between current high speed vision systems

Duro Clear Glass Adhesive) using standard techniques [20]. The camera used for the acquisition is manufactured by Photonfocus AG (Pfäffikon, Switzerland) and includes the dynamic ROI feature described in the previous section. A zoom lens (Edmund Optics VZM 300i) was mounted on the camera. The camera was connected to a Silicon Software (Mannheim, Germany) Micro Enable III Framegrabber via a Camera Link interface. The computer was a standard commercial PC running on an Intel Pentium IV 2.8Ghz processor with 512MB of RAM. The probe with the attached fly was positioned along 6 degrees of freedom with a micromanipulator (Sutter MP285). The fly was positioned such that the wing stroke plane coincided with the camera plane as best as possible while avoiding occlusions from the tether. A randomized bundle fiber guided the light from a 150W halogen light source (Schott ACE) under the fly. Diffusive tracing paper was placed 15mm in front of the light source, providing intense and homogenous back lighting. An additional 650nm high pass filter was used to eliminate the wavelengths visible to the fruit fly's visual system [21].

The software was programmed using Visual C++ 6.0. Standard image processing tools used OpenCV and Intel Performance Primitives (IPP). The program ran in the Windows environment with real-time thread priority to avoid significant delays due to system interrupts.

B. Acquisition Parameters

The zoom lens' aperture was fully closed to $\emptyset 1.5\text{mm}$ and the primary magnification was set to 0.75:1. The camera's exposure time was fixed to $50\mu\text{s}$ to avoid wing blurring. Tracking typically lasted two seconds, during which the images were transferred to the frame grabber's buffer and were accessed directly by the image processing tools. The camera was grabber controlled, i.e. the frames were acquired continuously, independently of the software status. Typically, no frames were dropped. To track both wings, the system sequentially exposed each side. Consequently, the effective temporal resolution for each wing was divided by two.

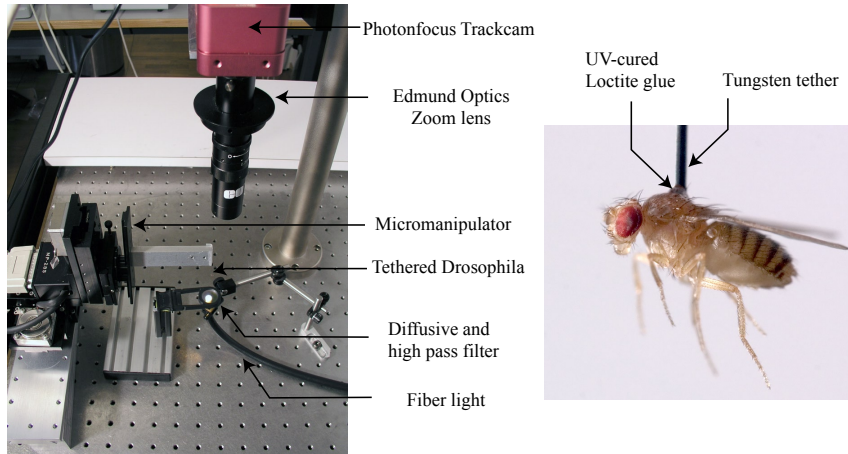


Fig. 3. Experimental setup of the high speed vision system

C. Software Architecture

The system is composed of two main sections with different requirements: calibration and tracking. The calibration is used to extract as much prior knowledge as possible from the experimental conditions. The tracking then relies on this information to calculate the positions of both wings in the smallest amount of time. Both of these tasks are discussed in detail in the subsequent sections.

IV. CALIBRATION

The three main goals of calibration are to extract: a) the position and orientation of the fruit fly's body, b) the path of each wing and c) the visual properties of the wings and the background. The output of the calibration procedure is a sequence of pixel locations describing the typical path of each wing (see Fig. 4). This sequence is then used during tracking in order to focus image processing only on the relevant pixels.

A. Sequence acquisition and statistical analysis

Initially a set of fifty full-frame images (1024*1024 pixels) is acquired at 80Hz during random phases within several wing strokes. Next, a statistical analysis of the values of each individual pixel location is made throughout the sequence. This results in two images: First, the median value of each pixel is used to build the background image. Second, the standard deviation image is binarized and employed as a mask to isolate the wings from the rest of the image.

B. Body property calculation

To segment the body, the median image is initially thresholded and morphological operators are applied to eliminate small blobs. Then, size and position criteria are applied to the remaining blobs. Once the blob containing the body is identified, a moment analysis is performed to determine its angular orientation. This orientation is used to differentiate the left and right side. The body blob is then enlarged to take into account small body movements,

and the result is used as an additional mask to isolate the wings.

C. Wing segmentation and path calculation

This set of masks is applied to each image in the sequence. The background image is subtracted to increase the signal-to-noise ratio by rendering the system robust to spatial variations in backlighting. The Max-Lloyd quantization algorithm is used to calculate an optimal binary threshold value (in the least square sense) to distinguish the wings from the background. The application of this threshold results in binary images representing the isolated wings. A morphological opening is performed to eliminate possible holes within the wing blobs caused by the semi-transparent sections of the wings. Next, a Canny edge detector is applied to isolate the wing edges. A Hough transform is employed to extract the strongest line on each side of the body. In 80% of the cases, this line corresponds to the leading edge of the wing. The intersection between each strongest line of the sequence is calculated. The median intersection is taken as an estimate of the wing hinge position.

V. TRACKING

The fundamental requirement for tracking is that it must operate as fast as possible without sacrificing robustness. In order to achieve high temporal resolution without having to give up spatial resolution, we chose to acquire only small regions of interest. *A priori* knowledge about the wing motion was used to predict the future position of the ROI. A wing beat kinematic model can be fitted to the measurements, thus increasing the accuracy of the estimate and, at the same time, providing a real-time analysis of the data. This analysis can then be used as a metric in the real-time experiments. Fig. 6 describes the system's state machine.

A Kalman filter was chosen as a means to fit the kinematic model to the measured data and to predict the future states and positions. We chose the Kalman filter because

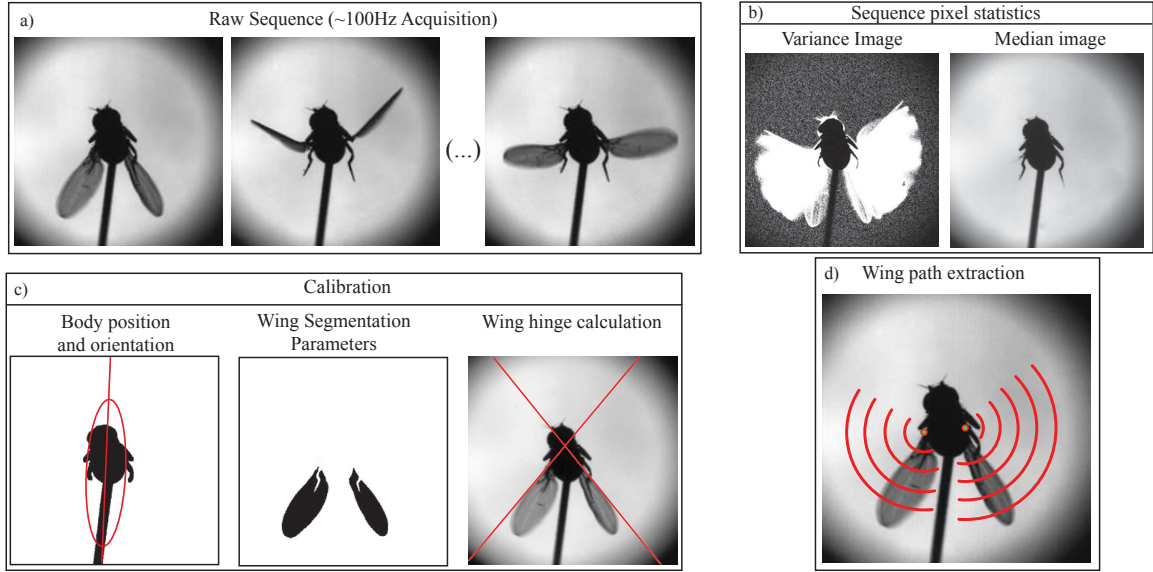


Fig. 4. Calibration procedure a) raw sequence images (1024x1024 pixels) b) statistical analysis c) body and wing position extraction d) wing path extraction: the red circular arcs represent the extracted wing paths




Fig. 5. Edge extraction algorithm applied during tracking: 1) The relevant pixels within the image, whose location have been computed during the calibration task, are copied into a 1D array. 2) The background image is subtracted from the sequence. 3) The resultant sequence is binarized by applying the threshold obtained from the calibration. 4) A morphological opening is applied to fill in gaps caused by the semi-transparent sections of the wing. 5) An edge detector extracts the leading and trailing edge of the wing.

it is computationally efficient (recursive) and optimal with respect to most statistical criteria [22].

A. Wing edge detection

To initialize the Kalman filter prior to wing edge following, the ROI is placed around one of the two bijective points of the wing path (maximum or minimum). Thus, the point corresponds to a unique phase of the wing beat kinematics, and the filter can adjust itself without losing tracking due to the small wing velocity at these points. Once the ROI is positioned, images are continuously acquired. The images are processed as described in Fig. 5. The system remains in the detection mode until both a leading and a trailing edge are found. The Kalman filter is then initialized and the system switches to the wing edge following mode.

B. Wing edge following

During wing edge following, the current ROI is grabbed and processed. The results of the processing, i.e. the leading and trailing edge locations, are used to update the parametric model of the wing kinematics, which is then used to predict the future position of the wing. The ROI is centered on the prediction, and the camera parameters are updated. The next frame is then grabbed using these parameters, and the process is repeated. If the process fails to detect the wing edges over an extended amount of time, the system switches back to the detection mode and sets the ROI back to its initial position, as described in Fig. 6.

C. Kalman filtering

A Kalman filter was used to fit an *a priori* model of the wing kinematics to the measured data and predict the subsequent wing position.

Since the wing beat is periodic, we use a Fourier series to model it:

$$\theta = a_0 + \sum_{i=1}^N a_i \cos(\omega t) + b_i \sin(\omega t) \quad (1)$$

where θ represents the 1D angular position of one of the wing edges (with respect to the wing hinge) and N is the number of Fourier terms (note: there are two separate estimates for each side, i.e. leading and trailing edge).

By constraining all parameters except amplitude, offset, frequency and phase, we maintained the original shape of the function, while scaling it in time and space. Eq.(1) is rewritten to incorporate these state variables:

$$\begin{cases} \theta = a_0 + \sum_{i=1}^N \alpha_i M \sin(\xi + \varphi_i) \\ \xi = \omega t \end{cases} \quad (2)$$

where α_i is the relative amplitude of the Fourier coefficient i , M is the amplitude of the function, ξ contains the phase

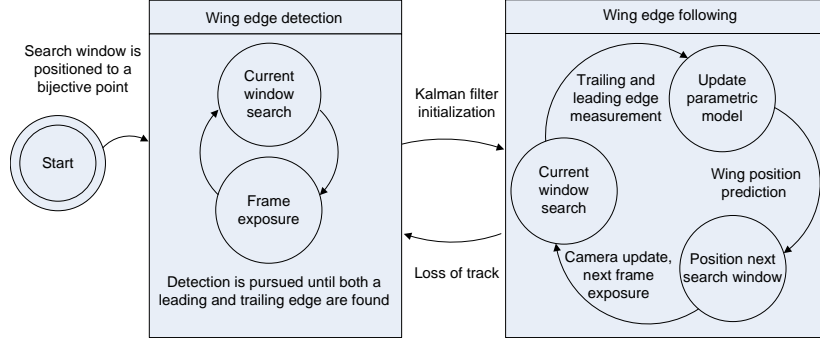


Fig. 6. Tracking state machine

information of the function with respect to the frequency $2\pi\omega$ (phase is dependent on frequency), and φ_i is the relative phase of the Fourier coefficient i . The α_i and φ_i coefficients are static values which describe the shape of the model. The state vector is $x_k = [a_0 \ M \ \xi \ \omega]^T$, and the derivation of the discrete state dynamics from (2) yields:

$$x_k = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T_s \\ 0 & 0 & 0 & 1 \end{bmatrix}}_A x_{k-1} + W \quad (3)$$

$W \sim N(0, Q)$

where T_s is the discrete time step, A relates the state at the previous time step to the state at the current step, W is the process noise and Q is the process noise covariance. This noise is introduced to take into account the changes in kinematics performed by the fly, as well as the imprecision inherent to the model.

The measurement equation is nonlinear. Therefore, an Extended Kalman Filter (EKF) is used, which linearizes the measurement equation at each time step:

$$\begin{aligned} \theta &= H x_k + v_k \\ H &= \begin{pmatrix} \frac{\delta\theta}{\delta a_0} & \frac{\delta\theta}{\delta M} & \frac{\delta\theta}{\delta \xi} & \frac{\delta\theta}{\delta \omega} \end{pmatrix} \\ v_k &\sim N(0, R) \end{aligned} \quad (4)$$

where v_k is the measurement noise and R is the measurement noise covariance. In the experiments, the static parameters α_i and φ_i with a total of 3 Fourier terms ($N = 3$) were chosen to fit the kinematic measurements of the wing beat presented in [3]. The two parameters of the Kalman filter, R and Q , were determined empirically from a first series of measurements. The initial estimate of the state vector is computed when both wing edges are detected. The EKF is then used at each time step to predict the position of the next window. Every time a new measurement is available the state vector is updated by solving the Ricatti equations.

VI. RESULTS

Fig. 7 shows the real-time wing position measurements for one of the two tracked wings.

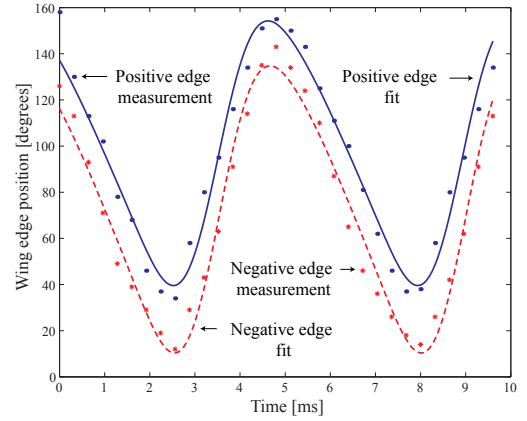


Fig. 7. Tracking results for one wing with the camera running at 6250Hz




Fig. 8. Process time line: Procedures related to frame i are in bold.

In the experiments, the resolution of the wing position was 1° . There are then two parameters that the user can choose: the field of view (FOV) or the frame rate. These two parameters are interdependent, since a larger field of view increases the amount of data to be transferred and processed. Therefore, a tradeoff must be found; a smaller field of view will increase the temporal resolution but at the risk of losing track of the wing. The results presented here vary between 4000Hz and 6250Hz.

Fig. 8 shows the timing of the processes at 6250 Hz. Fig. 9 shows a sequence of acquired images which were recorded by saving the remaining data on the frame buffer after tracking.

The system's frequency is presently limited by the acquisition and not by the processing. To increase the temporal resolution one could further narrow down the field of view.




Fig. 9. Sample from a tracking sequence, leading and trailing edges are indicated in white and black circles

However, the field of view must include the width of the wing. If the temporal resolution must be increased even further, it must be done at the cost of spatial resolution.

In the experiments the robustness of the tracking was high. Most errors were caused by the fly extending its legs into the wing path, however, this behavior was rare. In a few other cases the system lost track of the wing edges, typically around midstroke, where the wing moves the fastest. Usually, the system recovered tracking as the wing slowed down toward stroke reversal.

VII. CONCLUSION

This paper demonstrates the development and application of a novel high speed computer vision approach for the real-time analysis of wing kinematics in tethered *Drosophila melanogaster*.

Whereas previous approaches used analogue electronics to measure stroke amplitude and frequency in tethered flying insects, we apply digital real-time high speed videography to measure additional and more detailed parameters from tethered flying fruit flies. The new technique has several advantages. First, the implementation is based on software and commercially available hardware, allowing new systems to be implemented and modified easily. Second, the software allows additional functionality to be implemented, such as the calibration procedure to measure mean stroke paths. Third, our system provides a continuous measure of changes in wing motion due to the Kalman filter analysis. Fourth, additional important parameters such as an estimate of the angle of attack can be extracted from the image.

Besides ease of implementation and use, the relevance of our new technique lies in the specific adaptations that users can apply to adjust their system for different measurement conditions, such as various angles-of-view or insect species. Furthermore, the present system can be modified to perform at higher levels. For example, provided sufficient computing power, the use of additional cameras or mirrors to obtain multiple views could provide more information.

It is likely that real-time digital videography will find widespread use in experimental research and process control, whenever the status of a system must be rapidly evaluated as part of a control loop. For example, such situations are often found in product line inspection or in tracking at the micro and nano scale.

REFERENCES

- [1] M. Wautelet, "Scaling laws in the macro-, micro- and nanoworlds," *European Journal of Physics*, vol. 22, no. 6, pp. 601–611, 2001.
- [2] S. N. Fry, R. Sayaman, and M. H. Dickinson, "The aerodynamics of free-flight maneuvers in *Drosophila*," *Science*, vol. 300, no. 5618, pp. 495–498, 2003.
- [3] —, "The aerodynamics of hovering flight in *Drosophila*," *Journal of Experimental Biology*, vol. 208, no. 12, pp. 2303–2318, 2005.
- [4] M. A. Frye and M. H. Dickinson, "Fly flight: A model for the neural control of complex behavior," *Neuron*, vol. 32, no. 3, pp. 385–388, 2001.
- [5] M. H. Dickinson, F. O. Lehmann, and S. P. Sane, "Wing rotation and the aerodynamic basis of insect flight," *Science*, vol. 284, no. 5422, pp. 1954–1960, 1999.
- [6] W. B. Dickson and M. H. Dickinson, "The effect of advance ratio on the aerodynamics of revolving wings," *Journal of Experimental Biology*, vol. 207, no. 24, pp. 4269–4281, 2004.
- [7] S. P. Sane and M. H. Dickinson, "The aerodynamic effects of wing rotation and a revised quasi-steady model of flapping flight," *Journal of Experimental Biology*, vol. 205, no. 8, pp. 1087–1096, 2002.
- [8] S. Mao and H. Q. Jiang, "Aerodynamic force generation and power requirements in forward flight in a fruit fly with modeled wing motion," *Journal of Experimental Biology*, vol. 206, no. 17, pp. 3065–3083, 2003.
- [9] M. Heisenberg and R. Wolf, "On the fine-structure of yaw torque in visual flight orientation of *Drosophila-melanogaster*," *Journal of Comparative Physiology*, vol. 130, no. 2, pp. 113–130, 1979.
- [10] S. M. Tang, R. Wolf, S. P. Xu, and M. Heisenberg, "Visual pattern recognition in *Drosophila* is invariant for retinal position," *Science*, vol. 305, no. 5686, pp. 1020–1022, 2004.
- [11] M. H. Dickinson, F. O. Lehmann, and K. G. Götz, "The active control of wing rotation by *Drosophila*," *Journal of Experimental Biology*, vol. 182, pp. 173–189, 1993.
- [12] D. R. Reynolds and J. R. Riley, "Remote-sensing, telemetric and computer-based technologies for investigating insect movement: a survey of existing and potential techniques," *Computers and Electronics in Agriculture*, vol. 35, no. 2-3, pp. 271–307, 2002.
- [13] J. M. Zanker, "The wing beat of *Drosophila melanogaster* I. Kinematics," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 327, pp. 1–18, 1990.
- [14] G. L. Cembrano, A. Rodriguez-Vazquez, R. C. Galan, F. Jimenez-Garrido, S. Espejo, and R. Dominguez-Castro, "A 1000 FPS at 128 x 128 vision processor with 8-Bit digitized I/O," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 7, pp. 1044–1055, 2004.
- [15] H. Oku, N. Ogawa, M. Ishikawa, and K. Hashimoto, "Two-dimensional tracking of a motile micro-organism allowing high-resolution observation with various imaging techniques," *Review of Scientific Instruments*, vol. 76, no. 3, 2005.
- [16] S. Kagami, T. Komuro, and M. Ishikawa, "A high-speed vision system with in-pixel programmable ADCs and PEs for real-time visual sensing," in *Advanced Motion Control*, 2004, pp. 439–443.
- [17] P. Lichtsteiner and T. Delbruck, "A 64x64 event-driven logarithmic temporal derivative silicon retina," in *PRIME 2005*, EPFL, Switzerland, 2005.
- [18] B. Wilburn, N. Joshi, V. Vaish, E. V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy, "High performance imaging using large camera arrays," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 765–776, 2005.
- [19] D. C. Schuurman and D. W. Capson, "Robust direct visual servo using network-synchronized cameras," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 319–334, 2004.
- [20] L. F. Tammero and M. H. Dickinson, "Collision-avoidance and landing responses are mediated by separate pathways in the fruit fly, *Drosophila melanogaster*," *Journal of Experimental Biology*, vol. 205, no. 18, pp. 2785–2798, 2002.
- [21] M. Juusola and R. C. Hardie, "Light adaptation in *Drosophila* photoreceptors: I. Response dynamics and signaling efficiency at 25 degrees C," *Journal of General Physiology*, vol. 117, no. 1, pp. 3–25, 2001.
- [22] P. Maybeck, *Stochastic models, estimation and control*. Academic Press, 1979, vol. 1.