Automated aerodynamic performance estimates for flapping flight

Diana Chin

Department of Mechanical Engineering
Stanford University
Stanford, CA
ddchin@stanford.edu

Abstract—**The aerodynamic performance of freely flying birds is inherently difficult to measure directly without interfering with their natural flight. Many current methods thus combine kinematic parameters derived from high-speed video with theoretical force and power models. However, these methods often require manual tracking of points in each frame of a video to obtain kinematic data. This project utilizes image processing techniques to automatically extract kinematic data (stroke angle and wingbeat timing) from high speed videos of Pacific parrotlets flying from perch to perch. To analyze aerodynamic performance, this data is then combined with actuator disc theory to estimate the induced power during each wingbeat.**

Keywords—bird flight; image processing; aerodynamic performance; wingbeat; actuator disc

I. Introduction

A. Motivation

Quantifying the aerodynamic performance of freely flying birds is an open challenge in the field, because it is inherently difficult to directly measure the forces generated and the power required by flapping wings without impeding their natural motion. Force and power estimates are thus frequently derived from kinematic analysis of high speed video. Acquiring this data has traditionally involved tracking markers on the bird, as in [1], using software tools, such as the DLT program described in [2]. Although this program has some automatic tracking capabilities, they only work well with very clearly visible markers and relatively uncluttered backgrounds. However, artificial markers may interfere with the flight of smaller birds. Natural landmarks, such as the eye of the bird, can be used as an alternative, but tracking them often involves manual clicking of points in each frame of a high speed video. This is not only a tedious process, but can also introduce more noise into the resulting kinematic data. This project thus utilizes image processing techniques to automate the extraction of relevant kinematic data from high speed videos.

B. Actuator Disc Theory

Animal flight performance can be quantified through power estimates derived from models such as the actuator disc theory. The aerodynamic power required for flight is the sum of the induced power for overcoming the drag due to lift generation, and profile power for overcoming form drag and skin friction [3]. Although the actuator disc theory was originally developed for helicopter blades, it has been used to approximate induced power for hovering insects [4] and forward flying birds [5]. As shown in Fig. 1, the wings are modeled as an actuator disc, similar to a helicopter rotor, and lift (L) is estimated as the change in momentum as air passes through this disc into the wake below. For a disc with area A and air with density ρ and downwards velocity w,

$$L = 2\rho A w^2 \tag{1}$$

The induced power (P) can then be expressed by

$$P = Lw = L\left(\frac{L}{2\rho A}\right)^{0.5} = \left(\frac{L^3}{2\rho A}\right)^{0.5} \tag{2}$$

For a bird that sweeps out an angle ϕ with a wing of length r, the area of the actuated disc can be estimated as

$$A = \pi r^2 (\phi/180^\circ) \tag{3}$$

Although average power requirements can be roughly estimated by assuming a constant A, the angle ϕ can vary significantly between different wingbeats of a given flight. Instead, ϕ can be calculated as the peak-to-peak change in the stroke angle, defined here as the angle formed between the chord of the wing and the longitudinal axis of the body. Determining the stroke angle directly from high speed video will therefore enable more accurate power estimates for individual wingbeats.
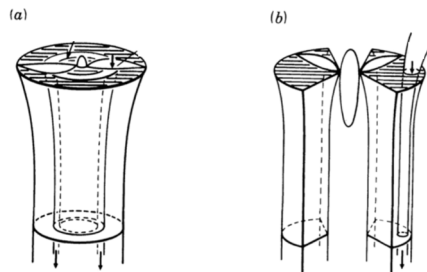

Fig. 1. The idealized wake from the actuator disc model for (a) a helicopter rotor and (b) a flapping-wing animal. Induced power can be estimated from the change in momentum through the shaded area, i.e. the "actuator disc". Figure adapted from [4].

II. RELATED WORK

A few previous EE 368 projects have also utilized image processing techniques to study flapping bird flight. Three studies are particularly relevant to this work [6-8]. The first study similarly automated the tracking of wing angles in a high-speed video. However, the algorithm was designed for one specific video of a hovering hummingbird. Unlike the videos of parrotlets processed in this project, the hummingbird stayed in roughly the same position throughout the video, and its wings did not deform signficantly during flapping [6]. The second project also identified wingbeat phases during the forward flight of a parrotlet. This was done using a video of the rear view of the flight and relied on image matching, instead of the point tracking used in this work (and in [6]). Image matching worked well for identifying when wingspan was the greatest, but did not work as well for identifying the beginning of the downstroke [7]. The third project also focused on wingbeat phase detection. However, this study classified wingbeat phases based on the eccentricity of the region containing the wing of the bird, and required two camera views for accurate results. Additionally, only the middle 2-3 wingbeats of flights were studied [8]. No previous studies have tracked stroke angles or successfully identified downstroke timing from takeoff to landing using only one camera view.

III. EXPERIMENTAL SETUP

Three Pacific parrotlets were trained to fly on cue from perch to perch inside of an aerodynamic force platform using positive reinforcement. The perches were set 75cm apart at inclination angles of 0, +20, and -20 degrees, and three to five flights at each inclination were recorded for each bird. Recordings were made at 3000 fps using a set of black-and-white Phantom Miro M310 high-speed cameras positioned outside of the setup, as shown in Fig. 2. To provide a bottom-up view of each flight, a panel of three mirrors was placed on the bottom plate of the force
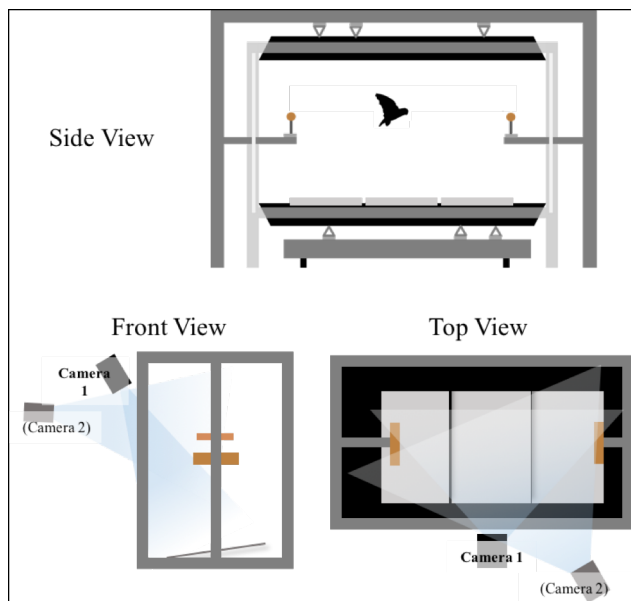

Fig. 2. Side, front, and top views of the experimental setup used to film the perch-to-perch flights anaylzed in this study.

platform at an angle of 11 degrees. The lens of the primary camera was angled 22^o from the horizontal (11^o from the mirror plane), as this was the smallest offset from the mirror angle that still enabled an unobstructed view of the mirrors. Actuator disc calculations were adjusted to account for these offsets. A second camera was also used to provide a side view of the flight. This secondary footage was not used for kinematic analysis, but was used to verify wingbeat timing results from the primary video, and to create time-lapse images of selected flights.

The cameras were synchronized with ATI Nano43 force sensors inside the platform that measure vertical forces (i.e. lift) generated during flight. Synchronization of all recordings enabled kinematic data derived from the high-speed videos to be combined with the time-resolved force data.

IV. METHODS

A. Image Processing

A series of image processing steps are first used to extract stroke angles from the high-speed video showing the bottom-up view of each flight:

1)   Importing video frames: The grayscale video frames are imported into Matlab. To save processing time, only every 6 frames are saved to convert the 3000 fps videos to 500 fps, and each frame is scaled down to 80% of its original size.

2)   Contrast adjustment and background suppression: For each saved frame, a user-specified value of gamma (typically around 1/3 for most videos) is applied to improve contrast, and the pixel values of the first and last frame are averaged and subtracted from each frame to suppress the background.

3)   Region selection: A smaller rectangular region that includes the complete flight is selected from each frame for analysis. If this region is to be manually selected, a time-lapse image is generated to assist the user in selecting a suitable region by overlaying ten frames spaced equally throughout the flight. The camera was not moved between flights recorded on the same day, so this region can also be generated automatically for flights recorded on a previously analyzed date.

4)   Binarizing regions: Each frame is then converted to a binary image based on a user-specified threshold (typically around 0.05). Morphological opening with a square structuring element of size 10 is used to remove unwanted white pixels in the background of each frame. Morphological closing with a square structuring element of size 7 is then used to remove unwanted black pixels in the thresholded bird. An example frame resulting from these first 4 steps is shown in Fig. 3.

5)   Centroid alignment: The centroid in each binarized region is identified, and its location is converted back into coordinates in the original frame. Frames where no centroid is found (when the majority of the bird was not in the analyzed region) are excluded. Frames before takeoff and after landing, identified based on when the centroid does not show significant movement between frames (more than 5 pixels), are also removed. The centroid path in the remaining frames is smoothed using the loess method with a span of 70% of the data. The original frames are then shifted to align the smoothed
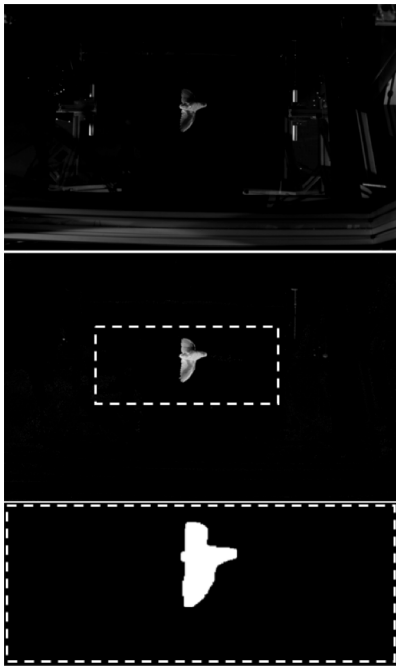
Fig. 3. Frames resulting from image processing steps 1-4. The top image shows the original frame. The middle image shows the result after step 2, with the selected region from step 3 represented by the dotted white line. The bottom image shows the binary result from step 4.

centroid coordinates so that the bird remains centered in every frame.

6) Edge detection: A new, smaller region is extracted from each shifted frame and binarized as in step 4, but with smaller square structuring elements of size 5 for opening and closing. A morphological edge detector, wherein the thresholded bird is subtracted from a dilated version (using a square structuring element of size 3), is then used to extract the outline of the bird in each frame.

7) Wingtip Identification. The edge pixels are converted to polar coordinates (R, θ), where R is the distance from the pixel to the right shoulder of the bird (assumed to be 3 pixels below the centroid) and θ is the clockwise angle formed between the horizontal and a line drawn from the centroid to the pixel. The pixels are sorted based on R, and the furthest pixel on the bird's right wing that falls within a pre-specified range for θ is identified as the wingtip. These θ ranges are specified for 4 parts of each video: the initial frames, usually near takeoff (from the first 1/8 of the flight); the first main section (1/8 to 1/3 of the flight); the middle section (1/3 to 2/3 of the flight); and the final section of the flight. The R and θ corresponding to the identified wingtip pixel are saved as the wing length and stroke angle of that frame, as shown in Fig. 4. The stroke angles extracted from the video are then used to determine the start and end times of each downstroke.

B. Wingbeat Extraction

The stroke angle and wing length traces for each flight are first smoothed using the loess method with a span of 4-7% of
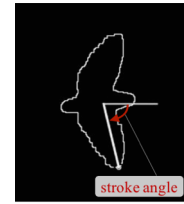

Fig. 4. Example frame showing the bird outline from step 6 with the identified wingtip location and stroke angle from step 7.

the data. The largest peaks are identified as potential downstroke start times. To avoid identifying the start of a bound, when the wings are folded against the body, as the start of a downstroke, potential start times are eliminated if the angle does not change significantly shortly after the peak time. The first downstroke is the only stroke that sometimes begins with an increase in the stroke angle, so the first downstroke start time is identified separately based on a user-specified threshold. The largest dips (local minima) are identified as potential downstroke end times. To account for wing rotation time, the end times are shifted 10ms before the minimum angle is reached. Any extra maxima or minima are then removed so that the starts and stops alternate in time. Downstrokes lasting less than 7ms (these sometimes resulted due to misidentified wingtips during the bound) are also recognized as falsely detected strokes and are removed.

C. Post-processing

The stroke angles and downstroke timing are then post-processed in three different ways:

1) Stroke angle video: A video showing how the stroke angle varies in time is generated for each flight. Each video is created by combining three components: a video showing the bird outlines with identified wingtip locations; a video of the actual bird flying in the original analysis region (from step 3 of image processing); and a graph of stroke angle vs. time, where the angle is plotted as each frame is shown in the two videos.

2) Time-lapse image: Frames corresponding to the downstroke start times are extracted from video recorded by the secondary camera (described in the experimental methods section). The first video frame is subtracted from each of these selected frames to suppress the background, and the resulting frames are overlayed. Because there is signficant overlap in the position of the bird between these frames, the frames are added in reverse order, starting with the last downstroke frame. To avoid saturation, a mask is created by binarizing the current image after each downstroke frame is added. This mask is then applied to the next downstroke frame before that frame is added to the time-lapse image.

3) Induced power estimates: Synchronized force data from the setup is combined with the kinematic data to estimate induced power for each downstroke using (2). Wing lengths are converted from pixels to meters using scaling factors calculated from a wand calibration performed in the setup. For ascending and descending flights, the scaling factor is assumed to vary linearly from takeoff to landing. Lift forces and wing lengths

are averaged over each downstroke, and the stroke amplitude ϕ from (3) is calculated as the difference between the maximum and minimum stroke angle during each downstroke

V. RESULTS

The processing steps described in the previous section were tested on 31 videos of perch-to-perch flights by the three parrotlets. Stroke angles and all wingbeats were successfully extracted from every video by specifying the same set of parameters for each flight (these are described in the next section). However in one video the bird did not move signficantly before making its first downstroke, so the Matlab code needed to be modified to avoid cutting off the first wingbeat. The results were checked visually using footage from both the primary and secondary camera. Examples of results from the three post-processing procedures described in the previous section are shown in Figs. 4-6.
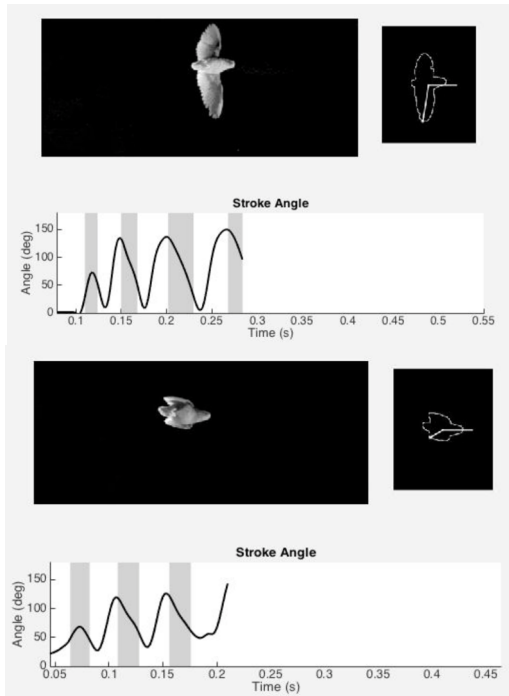

Fig. 4. Screenshots of videos created for flights by two different birds. The top video shows a parrotlet mid-downstroke, and the bottom video shows a parrotlet entering into a bound. Downstrokes are denoted by the gray shaded regions in the plots.
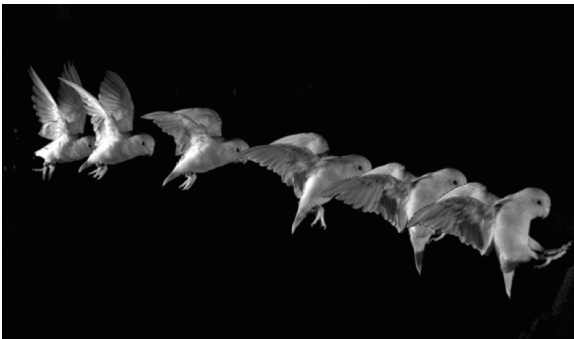

Fig. 5. Time-lapse image showing a parrotlet at the start of each downstroke during a perch-to-perch flight.
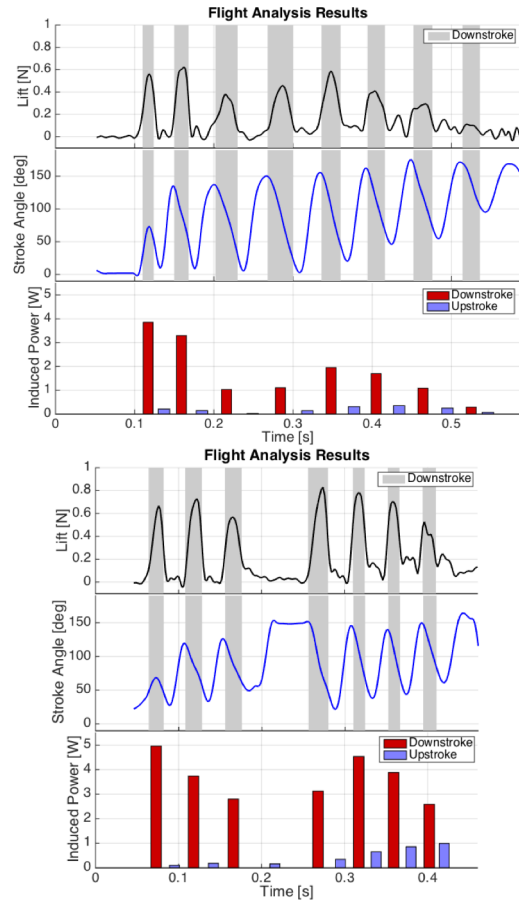

Fig. 6. Flight analysis results from the two flights shown in Fig. 4. Synchronized force data (lift), kinematic data (stroke angle), and induced power estimates per wingbeat are plotted.

VI. DISCUSSION

This project was successful in automating aerodynamic performance estimates through the extraction of stroke angles and wingbeat timing from high speed videos of flapping flight. Compared to traditional methods of identifying downstroke times by eye and manually clicking points frame-by-frame, the program developed in this study greatly expedites the extraction of kinematic data from high speed video. By automating the process, the program is also much more objective and repeatable than manual methods.

The post-processing results also enable greater understanding of bird flight aerodynamics than would be possible by only viewing the original videos. For example, Fig. 5 immediately shows how significantly the body pitch of the bird changes with each stroke from takeoff to landing. The flight analysis results in Fig. 6 demonstrate how different flight strategies impact aerodynamic power requirements; the first flight involves narrower lift force peaks with smaller stroke amplitudes compared to the second flight, so the induced power required for downstrokes in the first flight become signifcantly greater.

However, the program developed for extracting kinematic data is not without its limitations. In addition to the gamma for

contrast correction and the threshold for binarizing frames, several additional parameters are necessary for the program to work consistently across different flight videos. Seven parameters are needed to correctly identify the wingtip (6 angle ranges and 1 wing length threshold), and four other parameters are needed to accurately determine all downstrokes (the level of smoothing for the stroke angle, the range of frames to search for peaks in, the maximum number of peaks to search for, and the angle threshold for finding the initial downstroke start time). Wingtip identification proved to be particularly challenging because the wingtip was not always the furthest point from the centroid, and the path that it traversed varied over the course of the flight. This necessitated dividing the flight into four parts, with specific stroke angle limits for each section and a minimum wing length threshold for the middle two sections. Setting proper limits sometimes required several iterations to identify the wingtip with reasonable consistency throughout a flight.

In spite of these limitations, identifying the wingtip based on point tracking is expected to be more accurate than relying on other methods, such as image matching, due to the large variation in wingshapes over the course of the flight. As shown in Fig. 5, even the starts of downstrokes vary widely from takeoff to landing during these perch to perch flights.

VII. FUTURE WORK

The program developed in this project would be further improved by reducing the number of parameters needed for robust processing of different flight videos. A possible solution would be to predict the approximate location of the wingtip in each frame based on its previous trajectory so that all the angle limits would not be needed. By improving the wingtip identification step, fewer parameters would likely be required to accurately determine downstroke timing. Other venues for future work include automatically identifying the analysis region in a new video instead of requiring the user to select it,

and tracking the beak or body orientation of the bird. It was assumed that the bird's body stayed aligned with the horizontal axis of the camera frame throughout the flight, but birds sometimes showed some slight deviations in body angle during the recorded flights. Adjusting for these changes in the bird's body orientation would improve the accuracy of the stroke angles extracted from the videos.

REFERENCES

[1] Berg, A. M. and Biewener, A. A. (2010). Wing and body kinematics of takeoff and landing flight in the pigeon (*Columba livia*). *J. Exp. Biol.* **213**, 1651-1658.

[2] Hedrick, T. L. (2008). Software techniques for two- and three-dimensional kinematic measurements of biological and biomimetic systems. *Bioinsp. Biomim.* **3**, 034001.

[3] Pennycuick, C. J. (1968). Power requirements for horiztonal flight in the pigeon *Columba Livia*. *J. Exp. Biol.* 49, 527-555.

[4] Ellington, C. P. (1984). The Aerodynamics of Hovering Insect Flight. V. A Vortex Theory. *Phil. Trans. R. Soc. B. Biol. Sci.* **305**, 1122, 115-144.

[5] Muijres, F.T., Spedding, G. R., Winter, Y. and Hedenstrom, A. (2011). Actuator disk model and span efficiency of flapping flight in bats based on time-resolved PIV measurements. *Exp. Fluids*. **51**, 511-525.

[6] Deetjen, Marc. Hovering hummingbird automated segmentation. Spr 2014-15.

[7] Gutierrez, Eric. Quantitative measurements of forward flapping flight using image processing. Spr 2014-15.

[8] Ravnan, Eirik. Automatic phase detection of a bird wing in flight using stereo cameras. Aut 2013-14.