2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

A Reinforcement Learning Approach to Lift Generation in Flapping MAVs: Experimental Results

Mehran Motamed and Joseph Yan

***Abstract*— In [17] we proposed an RL framework for control of flapping-wing MAVs. The algorithm has been discussed and simulation results using a quasi-steady model showed initial promise. In this paper, the results from an experiment on a *Drosophila*-based dynamically scaled model are presented and are used to verify the control framework. Moreover, a comparison between a biological *Drosophila melanogaster* and the experimental results shows the actual possibility of employing the proposed approach to MAV control problem.**

I. INTRODUCTION

Insect-scale flapping wing micro aerial vehicles (MAVs) have been proposed for a number of applications, especially ones where maneuverability at low speed is necessary often in confined spaces. These applications include internal inspection of pipes, search and rescue in hazardous or inaccessible areas, and surveillance of indoor environments. Some notable work in this area include those by Fearing et al. [9], Michelson and Reece [14], and Pornsin-Sirirak et al. [18].

Although one might consider designing low speed MAVs by scaling down dimensions of conventional aircraft (Reynolds number[1], $Re \approx 10^7$), doing so is not feasible because it would significantly magnify the relative effect of the viscous forces. In view of the remarkable capabilities of insect flight for low Reynolds numbers ($Re \approx 10^2$), numerous researchers have turned to the use of flapping wings to engineer a solution. The best understanding, both qualitative and quantitative, of insect unsteady aerodynamics has come from three general techniques: (a) direct force measurements on tethered biological insects (e.g., [1], [22]); (b) computational fluid dynamics (CFD) models which numerically solve the Navier-Stoke equations (e.g., [19] and [23]); and (c) instrumented dynamically-scaled models (e.g., [8], [2] and [11]). Sane provides a comprehensive survey on these methods, including their advantages and drawbacks in [20].

The authors previously discussed the problem of unsteady aerodynamics at low Reynolds numbers and proposed reinforcement learning (RL) as a suitable candidate for the flight control of flapping-wing MAVs [17]. The algorithm has been discussed and simulation results using a quasi-steady model showed initial promise by converging to a smooth flapping motion. The computer simulation demonstrated that the algorithms are suitable, and helped developing needed conditions to carry out the physical experiments but the adequacy of the quasi-steady model, itself, is questionable. Therefore, the RL algorithm was implemented on a dynamically scaled model to determine its use for true unsteady aerodynamics. This paper focuses on the findings from this experiment.

The remainder of this paper is organized as follows. Section 2 describes the dynamically scaled model in detail. Section 3 describes the experiment procedure and implementation while section 4 summarises the evaluation and results.

II. DYNAMICALLY SCALED MODEL

In the case of insect flight, the building of a dynamically scaled model will provide conditions that would be comparable to the real insect by conserving the underlying aerodynamics. The theory behind conserving the aerodynamics while scaling the geometric dimensions of a problem is commonly used in fluid mechanics studies (the basic concept can be found in a standard textbook on fluid dynamics such as [10]). The theory and design details of building the apparatus is previously discussed in [11] and therefore is not repeated in this paper. In this section we summarize the prototype and model specifications as well as the experimental setup.

A. Prototype and Model Specifications

The basis for specification of the scaled model is a typical *Drosophila melanogaster*, and the parameters are adjusted to geometrically and kinematically scale the *Drosophila* flight. Specifications for prototype and model are shown in Table I. "Prototype" refers to the biological *Drosophila melanogaster* and "model" refers to the dynamically scaled apparatus. Note that some parameters are not known before the end of experiment as the flapping motion is yet to be achieved.

B. Experimental Setup

The experiment setup is shown in Fig. 1. The base-joint is immersed in a $40'' \times 20'' \times 20''$ container filled with a mixture of glycerine and water to give the desired kinematic viscosity. The three stepper motors are able to provide pitch, roll and yaw motions[2] and are controlled by a PC through a GALIL DMC-2183 motion controller board. Force measurements are read from a sealed ATI NANO17 6-axis force/torque sensor installed on the wing base-joint.

This work was supported by NSERC Discovery Grant 262312-03.

M. Motamed is with Motion Metrics International Corp., Vancouver, BC, Canada, mehran@motionmetrics.com.

J. Yan is with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada, josephy@ece.ubc.ca.

[1]In aerodynamics, the Reynolds number (*Re*) is a measure of the ratio of inertial to viscous forces and it determines the general qualitative behaviour of the fluid interaction with an object.

[2]Only roll and yaw motions were used in the this experiment.

TABLE I

MORPHOLOGICAL AND FUNCTIONAL PARAMETERS OF THE PROTOTYPE AND THE MODEL.

Symbol	Parameter	Prototype (p)	Model (m)	
b	Wing span (mm)	4.94	Fixed:	400
A	Aspect ratio	5.487	Fixed:	5.487
$\hat{r}_2^2(S)$	Non-dimensional second moment of wing area	0.328	Fixed:	0.328
Φ	Stroke amplitude (deg)	$148-169$	Range: 0–160	160†
n	Stroke frequency (Hz)	$190-212$	Range: 0–0.420	0.325†
$\overline{\lvert d\hat{\phi}/d\hat{t}\rvert_{max}}$	maximum non-dimensional angular velocity	5.50	Bounded:	6.25‡
$\overline{(\lvert d\hat{\phi}/d\hat{t}\rvert)^2}$	Mean square of non-dimensional angular velocity	–	–	21.0†
$\overline{\lvert d\hat{\phi}/d\hat{t}\rvert^{\frac{5}{2}}}$	Mean $\frac{5}{2}$-power of the absolute value of the non-dimensional angular velocity	–	–	48.9†
$\hat{t}_{rot}$	Non-dimensional rotation time	–	–	0.25†
$\overline{F_L}$	Mean lift force (N)	–	–	1.82†
ρ	Fluid density ($kg\ m^{-3}$)	1.2	Range: 999–1261	1235†
ν	Fluid kinematic viscosity at $20°C$ (cSt)	15	Range: 1–1180	160†
Re	Reynolds number	165	Constraint:	165

Re_p calculated based on values of this table.
A_p and $\hat{r}_2^2(S)_p$ calculated from wing shape depicted in [27]; ν_p from [24]; all other prototype data from [12], [13].
Ranges are for the model and are based on model limitations.
†Parameters are found for the model after algorithm convergence.
‡The bound on maximum non-dimensional angular velocity is discussed in §III-A.
$\frac{5}{2}$-power of the absolute value of the non-dimensional angular velocity is calculated for future reference.

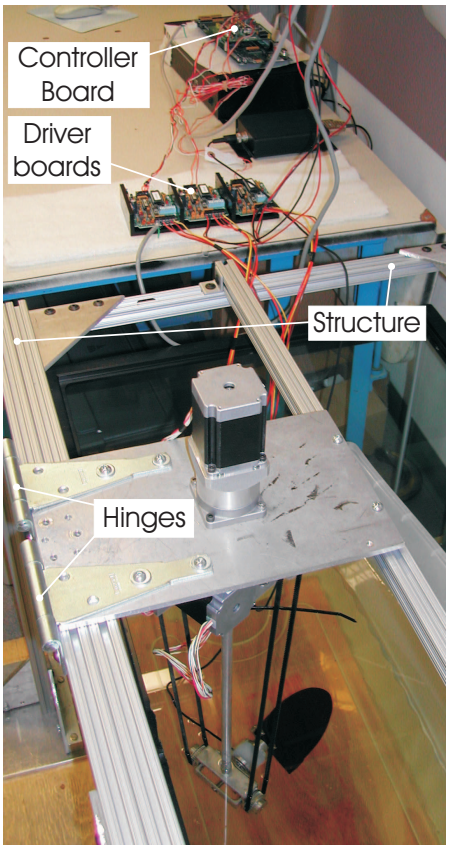

Fig. 1. Photo of experimental setup showing the aluminum frame reinforced for vibrations, single-wing driving mechanism and controller boards.

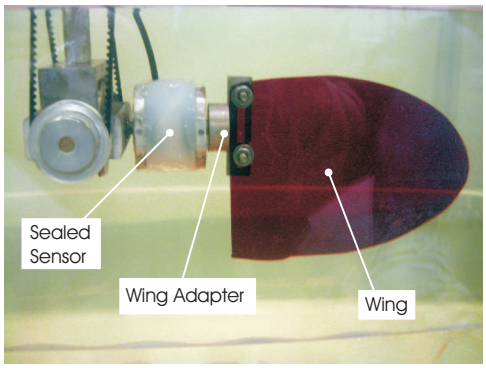

Fig. 2. The actual wing used in experiment, fabricated from a flat $\frac{1}{8}''$ plate of acrylic, shaped to have a planform geometrically similar to the wing of *Drosophila melanogaster*.

Forces are filtered using a zero-phase-delay, digital low-pass Butterworth filter with cutoff frequency of 2.5 Hz. The wing is fabricated from a flat $\frac{1}{8}''$ plate of acrylic, shaped to have a planform geometrically similar to the wing of a *Drosophila melanogaster* (wing shape from [27]). The wing, attached to the wing adapter and the sealed sensor, is shown in Fig. 2.

As can be seen from the Fig. 2, part of the wing proximal to the base rotation joint is missing due to the space occupied by the force sensor and other attachments. The aerodynamic force contribution of a differential wing element is proportional to the fourth power of its distance from the rotation axis. Thus, although the missing portion of the wing may appear to be large, its actual contribution has been estimated to be $+3.5\%$ so it is neglected (for more detail, refer to [15]). Another possible source of error is from the tank side walls, surface and ground. The total effect of the side walls, surface and ground on the mean lift-force coefficient calculated to be

$+13.5\%$ and the forces were corrected accordingly (for more detail, refer to [15]).

III. EXPERIMENT

The Q-learning algorithm used to carry out the experiment is previously explained by authors in [17]. All software were written in MathWorks' MATLAB®.

A. Power Requirement

The scaled-model apparatus, described above, ensures the underlying aerodynamics of *Drosophila*'s flight is conserved by matching the Reynolds numbers between the prototype and the model. Since the learning algorithm explores all the possible actuations it can make, one more attribute has to be matched before forces from the scaled model can be compared to *Drosophila*'s flight forces. *Drosophila*'s flight muscles have limited abilities to accelerate and decelerate the wing and therefore the scaled model, in our experimental case, should be made limited in software to enable future comparisons with *Drosophila*. In order to make the scaled model comparable to *Drosophila* in terms of actuation abilities, a maximum angular velocity has been found for *Drosophila* and is used as an upper bound on joint velocities in scaled model. From the mean specific inertial power, $\overline{P^*_{acc}}$ ($= \overline{P_{acc}}/mg$) from Ellington's formulation of power in [7], we have[3]:

$$\overline{P^*_{acc}} = \frac{\rho n^3 \Phi^2 R^3 (d\hat{\phi}/d\hat{t})^2_{max}}{2p_w}\Big[\frac{\rho_w}{\rho}\hat{h}\hat{r}_2^2(m)+\frac{\pi\hat{v}\hat{r}_2^2(v)}{2A}\Big], \quad (1)$$

where p_w is the wing loading. The two terms in the brackets represent the moments of inertia for the wing mass and virtual mass. These are determined by ρ_w, the density of cuticle, $\hat{h}$, the non-dimensional wing thickness, $\hat{r}_2^2(m)$, the non-dimensional second moment of wing mass, $\hat{v}$, the non-dimensional virtual mass, and $\hat{r}_2^2(v)$, the non-dimensional second moment of virtual mass. Virtual mass takes into account the inertia of the air around the wings that must be moved during the flapping motion. The parameter of interest in this equation is the non-dimensional angular velocity, $(d\hat{\phi}/d\hat{t})$, defined in [4] from the dimensional angular velocity, $(d\phi/dt)$, as:

$$(d\hat{\phi}/d\hat{t}) = \frac{2}{n\Phi}(d\phi/dt). \quad (2)$$

$|d\hat{\phi}/d\hat{t}|_{max}$ is the observed maximum non-dimensional angular velocity and can be used to find a theoretical upper bound on maximum non-dimensional angular velocity, i.e., $|d\hat{\phi}/d\hat{t}|_{bound}$. The actual value for $|d\hat{\phi}/d\hat{t}|_{bound}$ can be derived from $|d\hat{\phi}/d\hat{t}|_{max}$ of *Drosophila* in its *maximum* flight performance. Lehmann and Dickinson [12] investigated power requirements of *Drosophila* in its minimum, hovering and maximum performances. However, they assumed a constant maximum non-dimensional angular velocity of $|d\hat{\phi}/d\hat{t}|_{max} = 5.50$ in all three cases. Hence, for our purposes, there is no advantage in using their results on *Drosophila*'s maximum flight performance over other works on hovering *Drosophila*. Another way to estimate the non-dimensional angular velocity is from two simple models: in the *sawtooth* model, the wing-tip velocity follows a constant profile, and in the *harmonic* model, it follows a half-cosine profile [26] [4]. The values for non-dimensional angular velocity for each of these models has been calculated to be 4.00 and 6.28 respectively. Lehmann and Dickinson discussed that the actual value falls somewhere between these two extremes [12]. Consequently, the value $|d\hat{\phi}/d\hat{t}|_{bound} = 6.25$ has been selected to give a reasonable, yet practical estimate of the bound on non-dimensional angular velocity. This value is non-dimensional and therefore stays constant between the model and the prototype. We assume the same bound on the angular velocity for the muscles/motor of the other DOF, i.e., $|d\alpha/dt|_{bound} = |d\phi/dt|_{bound} = 6.25$.

[3]For the definitions and derivations of all the parameters in this equation, see [3] and [7].

B. Learning on Scaled Model

The state, action and reward functions are defined in the same way as previously stated in [17]:

$$\boldsymbol{s} = \begin{bmatrix} \phi \\ \alpha \\ \dot{\phi} \\ \dot{\alpha} \end{bmatrix}, \qquad \boldsymbol{a} = \begin{bmatrix} \Delta\dot{\phi} \\ \Delta\dot{\alpha} \end{bmatrix}, \quad (3)$$

$$r = R(\boldsymbol{s},\boldsymbol{a}) = \begin{cases} -R_{max} & \text{if } |\ddot{\phi}| > \ddot{\Phi}_{threshold}, \\ F_L(\boldsymbol{s},\boldsymbol{a}) & \text{otherwise.} \end{cases} \quad (4)$$

Due to algorithm limitations (refer to the algorithm in [15] for more detail), each of the two angular velocities has been quantized in 9 steps, that is:

$$\{\dot{\phi}, \dot{\alpha}\} \in \{-4,\ -3,\ -2,\ -1,\ 0,\ 1,\ 2,\ 3,\ 4\}. \quad (5)$$

Reynolds number can be estimated from a modified equation derived by Ellington in [7] that uses the time-averaged velocity of the wingtip instead of the flight velocity:

$$\overline{Re} = \frac{\Phi b^2 n}{\nu A}. \quad (6)$$

The actual flapping frequency and stroke amplitude is not known until the agent reaches the end of its learning episode and converges to a flapping trajectory. This is because the flapping motion is being learned by the agent and the parameters of this motion are not assumed beforehand. This introduces some difficulties into how well the Reynolds numbers can be matched. At the end of each run of the experiment, a new run has been started with a new kinematic viscosity to achieve matching Reynolds numbers. Therefore, the process of parameter matching (between prototype and model), in our case, is inevitably iterative. However, the simulation data presented in [17] provide insight into the possible ranges of these values, and has been used to estimate Φ and n and reduce the number of iterations.

The experiment was carried out for the last third part of the learning episode (i.e., 100K cycles) because of three reasons. First, the quasi-steady model is believed to be able

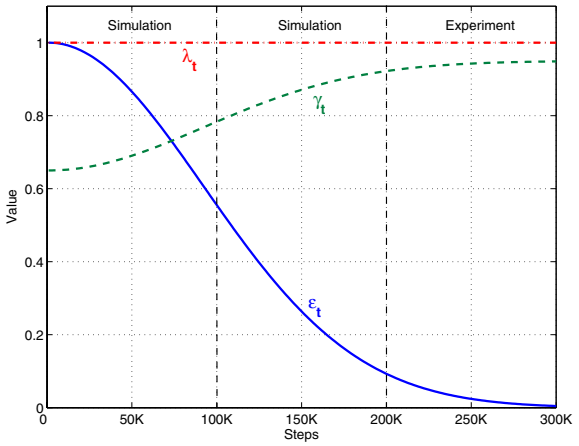

Fig. 3. Graph of ε_t, γ_t and λ_t used in experiment. The experiment was carried out using $\lambda = 1$ and ε and γ varying by equations described in [17]. As can be seen from the figure, the exploration rate ε is decreased and the discount factor γ is increased by the time step. The former implies decreased exploration near the end of the simulation, and the latter indicates increased reliability on expected future reinforcements as the agent becomes more "experienced".

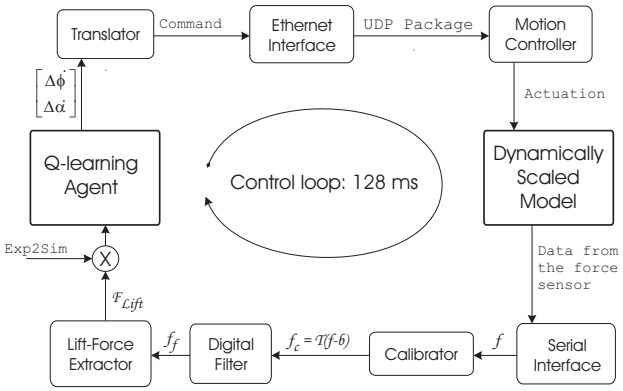

Fig. 4. Control loop of the experiment showing different components involved in actuation and sensing.

to sufficiently approximate the environment for the first two-thirds of the learning episode. This way the agent can learn the essential properties of the environment and then perform a more detailed search on an actual physical environment to fine-tune its behaviour. Second, the total learning episode could be very long considering the speed of the scaled model (e.g., an experimental episode would have taken more than 10 hours). And third, as the scaled model was not designed primarily to carry out such long experiments, the wear and tear on the mechanical parts was a concern. Fig. 3 shows how ε, γ and λ change over time. The three parameters are fed into the simulation for the first two-thirds of the learning episode, and then are used in the experiment for the last third of the episode.

Fig. 4 shows the data flow in the control loop that runs at $f_{control} = 7.81\ Hz$. The Q-learning agent is connected to the scaled model via actuation and sensing lines. A control

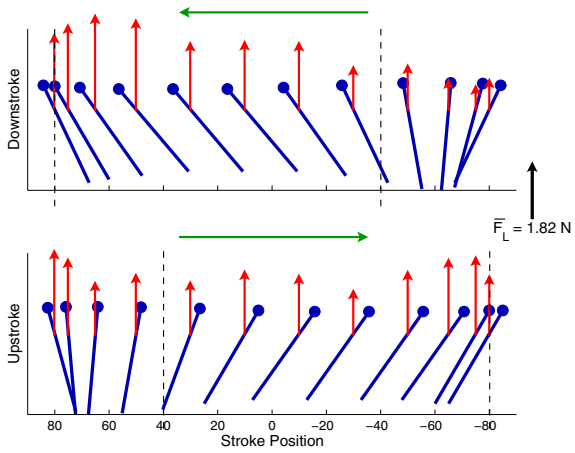

Fig. 5. Wing chord representation of the optimal policy found by the agent maximizing lift force in the experiment. The agent makes use of its full stroke amplitude. The rotations have been delayed until the start of the next half stroke, with a duration of $\hat{t}_{rot} \approx 0.25$ of the stroke cycle. The dashed lines indicate the approximate start and end of rotations and show how the rotations are delayed. The arrows on the wing chord show only the instantaneous lift forces.

cycle is initiated by the Q-learning agent issuing an action which in turn actuates the motors on the scaled model.

The force measurements at the wing base consist of gravitational, inertial and aerodynamic components. The lift force is the vertical aerodynamic component and should be extracted from this measured force. The inertial components represent the acceleration forces due to the mass of the sensor and the wing as well as the *added mass*[4] of the fluid around the wing, and are harder to isolate, since they are trajectory-dependent. However, as the stroke plane is horizontal, only the horizontal component of the total force is assumed to be affected by inertial forces. Moreover, since the base is fixed, inertial forces do not affect the *mean* aerodynamic forces, and the *mean* total force over a cycle can be obtained without calculating the inertial forces. Therefore, The vertical component of the net force can be derived by a coordinate transformation, followed by subtracting the gravitational contribution of the masses of the wing holder and the wing, calculated during the calibration process, from the vertical component. The lift force is then calculated by accounting for the total effect of the side-walls, surface and bottom of the tank.

IV. RESULTS

A. Experimental Results

Without any *a priori* assumptions on the desired trajectory, the agent converged to a smooth flapping motion as shown in Fig. 5 and 6, similar but slightly different than the one found from the simulation [17]. The agent preferred to accelerate at the beginning of each half stroke and to decelerate at the end to avoid sudden stall and at the same time to generate as much lift as possible to maximize the cumulative reward.

[4]When a wing accelerates, it sets the surrounding fluid in motion, causing in inertial forces by the fluid. This phenomenon is called the added mass effect.

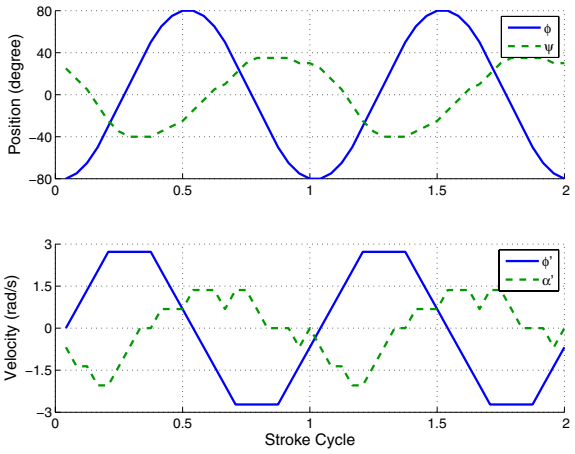

Fig. 6. The position and velocity profiles of the optimal policy of the experiment. Parameter $\psi = \alpha - 90^\circ$ is the rotation angle. The figure shows nearly sinusoidal motion on ϕ. The discontinuities in acceleration seen in velocity profile are due to quantization of velocities in 9 steps.

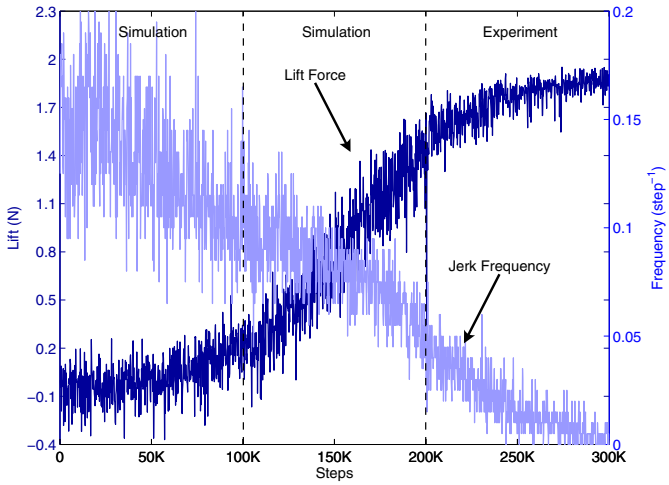

Fig. 7. Improvement on lift generation and jerk avoidance from the experiment. The graph shows the mean lift increasing by steps until it reaches the maximum mean lift force of $\overline{F_L} = 1.82\ N$. The graph also shows the agent minimized the frequency of jerk associated with a motion-range violation.

Fig. 7 shows that the algorithm converges to a mean lift of $\overline{F_L} = 1.82\ N$ in about $300K$ steps. From the figure it is also clear that the agent minimized the frequency of out-of-range motions due to associated punishment, as described in the algorithm [17].

As can be seen in Fig. 5, the rotations initiated before the end of each half stroke and the duration of each rotation occupy about one-quarter of the total stroke cycle. More technically, by defining non-dimensional time as $\hat{t} = nt$, we have $\hat{t}_{rot} \approx 0.25$. The rotations are almost delayed until the start of the next half stroke, as can be seen from the dashed lines of Fig. 5.

The experiment took approximately 3.5 hours[5] to learn the flapping flight. The speed of convergence was not an issue in this work, and the program has not been optimized for time.

The optimal policy has converged to the length of $N_{exp} = 24$ steps. That is, the final flapping motion cycle is completed in 24 time steps. The flapping frequency, then, calculated to be $n_m = 325\ mHz$. The agent made use of the whole available stroke motion range, as expected, and therefore $\Phi_m = 160^\circ$.

Given n_m and Φ_m, we can confirm that the actual Reynolds number is matched making the scaled model *comparable*, in this sense, to a typical *Drosophila melanogaster*. We have:

$$\overline{Re}_{m,actual} = \frac{\Phi_m b_m^2 n_m}{\nu_m A_m} = 165, \tag{7}$$

As can be seen, the $Re_{m,actual}$ and Re_p are matched. The final parameters for the the dynamically scaled model are summarized in Table I. The model parameters were set after several iterations to give a complete match with the prototype.

As qualitatively illustrated in Fig. 7, the experiment on the scaled model have successfully demonstrated the applicability of the control framework (presented in [17] and [15]) for the case study of lift generation.

B. Comparison with Drosophila melanogaster

In order to provide some assessment of the quantitative results, in this section, we compare some of the experimental findings from the scaled model to those of the fruitfly, *Drosophila melanogaster*. These comparisons must be prefaced with the cautionary warning that the uncertainty in the extrapolations are not insignificant, making some comparisons questionable. The scaled model is not intended to scale every attribute of biological insects, since the final goal of building an MAV capable of sustained flight will not be achieved by replicating biological insects (as also acknowledged in [16] and [15]). There are fundamental technological differences in mechanism and design (e.g., actuation) between the prototype (i.e., nature) and model (i.e., engineering). Moreover, there is a significant difference between the goal functions of the RL agent and an insect: The agent was not concerned with minimizing energy (and therefore had a greater freedom in generating aerodynamic forces), while it is believed that flapping power consumption certainly played a role in the evolution of biological insects. One can argue that based on this discussion, the promise of the findings from this experiment will not be realized until an actual MAV is built. In the meantime, however, nature can still be regarded as an intuitive guideline for confirming the final achievements and to some extent identifying the promise of the results.

The purpose of building any dynamically scaled model is to scale the prototype, gather information difficult to obtain in the prototype scale, and scale the results to get the corresponding results for the actual use on a smaller scale MAV. Therefore the comparison is made between a

[5] $T_{exp} = 100,000\ (steps) \times 0.128\ (s/step)$

TABLE II

PARAMETERS FOR THE BIOLOGICAL *Drosophila* IN ITS MAXIMUM PERFORMANCE AND FOR THE EXTRAPOLATED MAV USED IN COMPARISON.

Symbol	Parameter	*Drosophila m.* (p)	Extrapolated MAV (MAV)
b	Wing span (mm)	4.94	40
A	Aspect ratio	6.18	5.487
$\hat{r}_3^3(S)$	Non-dimensional second moment of wing area	0.242	0.230
Φ	Stroke amplitude (deg)	169	160
n	Stroke frequency (Hz)	212	3.04
$\overline{\lvert d\hat{\phi}/d\hat{t}\rvert^3}$	Mean cube of the absolute value of the non-dimensional angular velocity	104.6	114.9
$\overline{F_t}$	Mean total force (μN)	13.8	18.53
ρ	Air density ($kg\ m^{-3}$)	1.2	1.2
ν	Air kinematic viscosity (cSt)	15	15
κ	Rankine-Froude correction factor	1.28	1.28
Re	Reynolds number	165	165

biological *Drosophila* (denoted by subscript p, as before) and an extrapolated *Drosophila*-based MAV[6] of $4\ cm$ in wingspan[7] (denoted by subscript MAV to distinguish it from the scaled model denoted by subscript m).

To compare the generated force, it is more appropriate to compare the mean lift-force coefficients. This is because $\overline{C_L}$ is not affected by different scaling factors in different works, and can be directly compared to that of a biological insect. Since the prototype and the model have different policies for expending power, it is logical to compare the mean aerodynamic efficiencies of the two flapping trajectories in parallel[8]. Therefore, pairs of $(\overline{C_L},\ \overline{\eta_a})$ give a better metric for comparison.

The mean lift-force coefficient, $\overline{C_L}$, can be found from the following equation (modified from [5]):

$$\overline{C_L} = \frac{64\overline{F_L}}{\rho\Phi^2 n^2 b^4 A^{-1}\overline{\left(d\hat{\phi}/d\hat{t}\right)^2}\hat{r}_2^2(S)}. \tag{8}$$

As mentioned earlier, the force coefficient stays unchanged with scaling; therefore, in order to find $(\overline{C_L})_{MAV}$, model parameters can be plugged into (8).

The mean aerodynamic efficiency, $\overline{\eta_a}$, can be defined as the minimum power required to hover divided by the aerodynamic power expended (from Weis-Fogh [25] and Ellington [7]). That is:

$$\overline{\eta_a} = \frac{P^*_{RF}}{\overline{P^*_a}}, \tag{9}$$

where P^*_{RF} is the minimum power calculated for a steady downward Rankine-Froude momentum jet. P^*_{RF} for a horizontal stroke plane can be written as:

$$P^*_{RF} = \left(\frac{2p_w}{\rho\Phi A}\right)^{\frac{1}{2}}, \tag{10}$$

where p_w ($= F_L/S$ where $S = b^2A^{-1}$) is the wing loading. The mean specific aerodynamic power, $\overline{P^*_a}$, is the sum of the mean specific induced power, $\overline{P^*_{ind}}$, the mean specific profile power, $\overline{P^*_{pro}}$ and the mean specific parasite power, $\overline{P^*_{par}}$ (For more detail, refer to [7]). $\overline{P^*_{par}}$, the parasite power required to move the body through the air is commonly neglected when flight speed is assumed to be zero [7], and for tethered flight [12]. The power equations are therefore (modified from [12] and [7]):

$$\overline{P^*_{ind}} = \kappa\left(\frac{2\overline{F_t}}{\rho\Phi b^2}\right)^{\frac{1}{2}}, \tag{11}$$

$$\overline{P^*_{pro}} = \frac{\rho\Phi^3 n^3 b^5 A^{-1}\overline{\lvert d\hat{\phi}/d\hat{t}\rvert^3}\hat{r}_3^3(S)\overline{C_{D,pro}}}{128}, \tag{12}$$

$$\overline{P^*_a} = \overline{P^*_{ind}} + \overline{P^*_{pro}}, \tag{13}$$

where $\overline{F_t}$ is the mean total force, and κ is the Rankine-Froude correction factor required because of the periodic nature of vortex shedding in the wake [6]. The value for κ has been taken from [12]. $\hat{r}_3^3(S)$ is the third moment of wing area, $\overline{\lvert d\hat{\phi}/d\hat{t}\rvert^3}$ is the mean cube of the absolute value of the non-dimensional angular velocity and $\overline{C_{D,pro}}$ is the mean profile drag coefficient which according to the approximation of Ellington in [7] can be estimated as $7/(\sqrt{Re})$.

The parameters needed to calculate the lift and power for the extrapolated MAV of wingspan $b_{MAV} = 4\ cm$ are summarized in Table II. Plugging into the lift and power equations yields $(\overline{C_L},\ \overline{\eta_a})_{MAV} = (3.56,\ 0.40)$. The maximum observed lift coefficient for *Drosophila* can be found from [13] (also mentioned by Sane and Dickinson [21]) to be $(\overline{C_L})_p = 1.9$. The parameters needed to calculate the power for a biological *Drosophila* are summarized in Table II (from [12]). Using the same equations of power, $(\overline{\eta_a})_p = 0.25$. Therefore, $(\overline{C_L},\ \overline{\eta_a})_p = (1.9,\ 0.25)$.

The comparison is summarized in Table III. From the table, the first row indicates that the extrapolated MAV has demonstrated promising capability in generating a mean lift coefficient in the range of the *Drosophila melanogaster* prototype and beyond. In this comparison case and using the same equations for power, the second row suggests that the MAV has achieved a higher

[6]That is, extrapolated from the *Drosophila*-based scaled model.

[7]The choice of wing span was arbitrary.

[8]It can be argued that the mechanical efficiencies of the two systems should be compared. However, the inertial power depends on the actuation mechanism and design, and therefore can be quite different for an MAV and an insect.

TABLE III

COMPARISON RESULT BETWEEN A FRUIT FLY AND THE EXTRAPOLATED MAV IN TERMS OF LIFT GENERATION ($\overline{C_L}$) AND AERODYNAMIC POWER EFFICIENCY ($\overline{\eta_a}$).

	Drosophila m.	**extrapolated MAV**
$\overline{C_L}$	1.9	3.56
$\overline{\eta_a}$	0.25	0.40

mean aerodynamic efficiency. This increase was unexpected, considering the unlimited power budget of the MAV *versus* the limited budget of the biological insect. The reader is advised not to conclude a subtle but important implication of this sentence, that is the MAV being *overall* more efficient than *Drosophila*. Several reasons can be mentioned. First, the flapping trajectories were not the same. Second, the flapping trajectories compared do not represent the overall performances of either the MAV or the *Drosophila*. Finally, it is not discussed how much the biological insect has sacrificed power, in its maximum performance, to temporarily generate high aerodynamic forces.

Nevertheless, the comparison with the biological insect shows promising results and great confidence in the actual values achieved by the extrapolated MAV.

REFERENCES

[1] M.H. Dickinson and K.G. Götz. The wake dynamics and flight forces of the fruit fly, Drosophila Melanogaster. *Journal of Experimental Biology*, 199:2085–2104, 1996.

[2] M.H. Dickinson, F.-O. Lehmann, and S.P. Sane. Wing rotation and the aerodynamic basis of insect flight. *Science*, 284:1954–1960, 1999.

[3] C.P. Ellington. The aerodynamics of hovering insect flight. II. morphological parameters. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 305(1122):17–40, 1984.

[4] C.P. Ellington. The aerodynamics of hovering insect flight. III. kinematics. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 305(1122):41–78, 1984.

[5] C.P. Ellington. The aerodynamics of hovering insect flight. IV. aerodynamic mechanisms. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 305(1122):79–113, 1984.

[6] C.P. Ellington. The aerodynamics of hovering insect flight. V. a vortex theory. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 305(1122):115–144, 1984.

[7] C.P. Ellington. The aerodynamics of hovering insect flight. VI. lift and power requirements. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 305(1122):145–181, 1984.

[8] C.P. Ellington, C. van den Berg, A.P. Willmot, and A.L.R. Thomas. Leading edge vortices in insect flight. *Nature*, 384:626–630, Dec 1996.

[9] R.S. Fearing, K.H. Chiang, M.H. Dickinson, D.L. Pick, M. Sitti, and J. Yan. Wing transmission for a micromechanical flying insect. In *The IEEE International Conference on Robotics and Automation (ICRA'00)*, pages 1509–1516, San Francisco, CA, Apr 2000.

[10] R.W. Fox and A.T. McDonald. *Introduction to Fluid Mechanics*. John Wiley & Sons, New York, 3rd edition, 1985.

[11] W. Lai, J. Yan, M. Motamed, and S. Green. Force measurements on a scaled mechanical model of dragonfly in forward flight. In *The IEEE International Conference on Advanced Robotics (ICAR'05)*, pages 595–600, Seattle, WA, Jul 2005.

[12] F.-O. Lehmann and M.H. Dickinson. The changes in power requirements and muscle efficiency during elevated force production in the fruit fly Drosophila Melanogaster. *Journal of Experimental Biology*, 200:1133–1143, 1997.

[13] F.-O. Lehmann and M.H. Dickinson. The control of wing kinematics and flight forces in fruit flies (Drosophila spp.). *Journal of Experimental Biology*, 201:385–401, 1998.

[14] R.C. Michelson and S. Reece. Update on flapping wing micro air vehicle research - ongoing work to develop a flapping wing, crawling entomopter. In *The 13th Bristol International Conference on RPV/UAV Systems*, pages 30.1–30.12, Bristol, England, Mar-Apr 1998.

[15] M. Motamed. Reinforcement-learning control framework and sensing pradigm for flapping-wing micro aerial vehicles. Master's thesis, University of British Columbia, April 2006.

[16] M. Motamed and J. Yan. A review of biological, biomimetic and miniature force sensing for microflight. In *The IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'05)*, pages 3939–3946, Edmonton, AB, Aug 2005.

[17] M. Motamed and J. Yan. A reinforcement learning approach to lift generation in flapping MAVs: Simulation results. In *The IEEE International Conference on Robotics and Automation (ICRA'06)*, pages 2150–2155, Orlando, FL, May 2006.

[18] T.N. Pornsin-sirirak, S.W. Lee, H. Nassef, J. Grasmeyer, Y.C. Tai, C.M. Ho, and M. Keennon. Mems wing technology for a battery-powered ornithopter. In *The 13th IEEE International Conference on Microelectromechanical Systems*, pages 799–804, Miyazaki, Japan, Jan 2000.

[19] R. Ramamurti and W.C. Sandberg. A three-dimensional computational study of the aerodynamic mechanisms of insect flight. *Journal of Experimental Biology*, 205:1507–1518, 2002.

[20] S.P. Sane. The aerodynamics of insect flight. *Journal of Experimental Biology*, 206:4191–208, 2003.

[21] S.P. Sane and M.H. Dickinson. The control of flight force by a flapping wing: lift and drag production. *Journal of Experimental Biology*, 204:2607–2626, 2001.

[22] A. Sherman and M.H. Dickinsons. summation of visual and mechanosensory feedback in Drosophila flight control. *Journal of Experimental Biology*, 207:133–142, 2004.

[23] M. Sun and S.L. Lan. A computational study of the aerodynamic forces and power requirements of dragonfly (Aeschna Juncea) hovering. *Journal of Experimental Biology*, 207:1887–1901, 2004.

[24] S. Vogel. *Life in Moving Fluids*. Princeton University Press, Princeton, NJ, 1981.

[25] T. Weis-Fogh. Energetics of hovering flight in hummingbirds and in Drosophila. *Journal of Experimental Biology*, 56:79–104, 1972.

[26] J.M. Zanker. The wing beat of Drosophila melanogaster. I. kinematics. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 327:1–18, 1990.

[27] E. Zimmerman, A. Palsson, and G. Gibson. Quantitative trait loci affecting components of wing shape in Drosophila melanogaster. *Genetics*, 155(2):671–683, Jun 2000.