Building an Artificial Bird: Goals and Accomplishments of the Robur Project

Stephane Doncieux Jean-Baptiste Mouret Adrien Angeli
Renaud Barate Emmanuel de Margerie Jean-Arcady Meyer

Université Pierre et Marie Curie - Paris 6,
UMR 7606 LIP6, 8, rue du Capitaine Scott,
Paris, 75015 France
http://animatlab.lip6.fr
{Stephane.Doncieux,Jean-Baptiste.Mouret,Adrien.Angeli,Emmanuel.de-Margerie,Jean-Arcady.Meyer}@lip6.fr
{Renaud.Barate}@ensta.fr

Abstract

The ROBUR project aims at developing a series of capacities that are inspired from those of birds, bats or insects, and that might contribute to the autonomy of UAV. However, although the ultimate goal is to integrate these capacities in a single flapping-wing platform, several preliminary studies described in this paper concern more classical platforms like planes or helicopters.

The capacities under study can be grouped in three different categories: flapping-wing flight, reflexes and high-level behaviours.

Research efforts in the first category concern the understanding of the aerodynamics of flapping-wing flight, and aim at designing appropriate morphologies and controllers that may serve to implement the corresponding behaviour on a robotic platform. The second category concerns the implementation of some reflexes, like those of obstacle-avoidance or speed-regulation, likely to contribute to an UAV's safety in its environment. As for high-level behaviours, they cover a wider range of capacities. Their role is to turn the UAV from a mere teleoperated engine to a fully autonomous robot. This entails capacities like being able to spare its energy expenditures, to know its current localization, and to decide what to do at every moment.

This article describes the major results already obtained within the framework of this project.

1 Introduction

Research on UAV experiences a very fast growth because the design and control of these engines raise interesting scientific issues, because several platforms are available off-the-shelf, and because they offer numerous applications, both military and civilian. For instance, a lot of work has been done to allow these platforms to take-off, to follow a route given by GPS waypoints, and to land. These capacities make various missions possible, mostly focussed on observation or surveillance tasks, but these missions usually require a priori knowledge of the terrain and detailed planning. Another approach is possible, that would release the need of human intervention as much as possible, letting the UAV make its own decisions, according to its objective and current state.

To reach such a decisional autonomy, the UAV should be endowed with several abilities. It first needs to be able to freely wander in its environment. Whereas today's UAV are usually flying in open spaces with as few obstacles as possible, an autonomous UAV should deal with any environment it can physically handle – a huge plane won't, of course, be able to fly in a urban canyon. Such engine should sense its environment to automatically adapt its behaviour, if only for detecting and avoiding obstacles. This latter capacity, in turn, may entail being able to adapt its speed.

Obstacle-avoidance and speed-regulation reflexes, associated with appropriate low-level controls, would allow an UAV to be teleoperated through higher-level orders like "go north", without the need to take additional care of the engine's safety. However, several other abilities are still required to reach full autonomy, as this approach would contribute reducing the range of admissible movements, but without specifying where to go. If the recourse to predefined GPS waypoints is excluded, several objectives may serve to drive an UAV, like its overall mission, or like secondary goals such as those involved in energy management. Indeed, to spare some energy, a specific trajectory, occasionally not directly leading to the target, may exploit winds or thermals in the surroundings. The UAV should also decide to divert its trajectory towards a refilling station, if its energy level becomes too low.

The Robur project of the AnimatLab aims at drawing inspiration from flying animals to develop and integrate abilities that would enhance the autonomy of a flapping-wing robot (figure 1). However, several preliminary studies described in this paper concern more classical platforms like planes or helicopters.

We will start by a presentation of results already obtained on the evolution of flapping-wing flight. We will then describe the implementation of several reflexes, before introducing the two high-level behaviours we are currently working on: simultaneous localization and mapping, on the one hand, and soaring behaviours, on the other hand.

2 Evolution of flapping-wing flight

In nature, birds and bats clearly demonstrate aptitudes for manoeuvrability and energy economy that are largely out of reach for current UAV of the same size. Additionally, flapping-flight control represents an interesting challenge for our learning and adapting

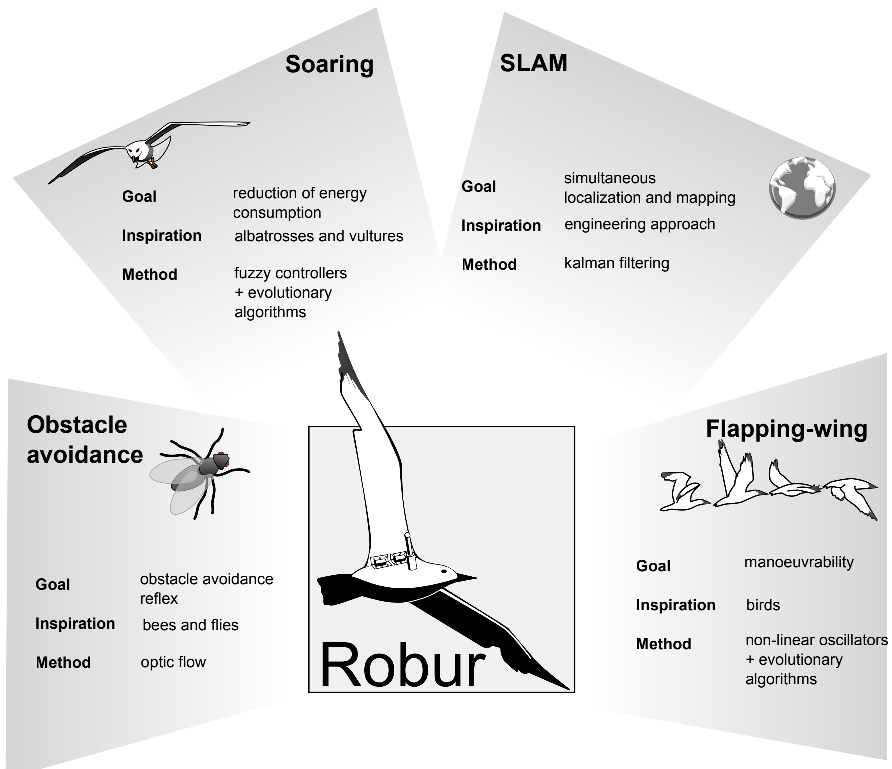

Figure 1: Overview of the adaptive capacities to be integrated on a flapping-wing platform within the framework of the Robur project.

algorithms. Indeed, most of the corresponding effector commands being oscillatory, some sensory feed-back is required to change their amplitude, frequency, or phase, while current engineering approaches to flapping-flight control call upon mere open-loop solutions [21, 2, 29, 19]. Moreover, biological observations clearly demonstrate that flying animals often change their wing shape during flight [28, 9], another consideration seldom taken into account in UAV's design and control.

Flapping-wing flight requires to precisely tune the wing orientation in order to maximise the generated lift and traction, while minimizing the energy expenditures. Birds efficiently exploit the numerous degrees-of-freedom (DOF) associated with their wings, while current artificial flapping-wing devices usually exhibit only one active DOF, the dihedral, while another one, the twist, is either passive or active. If several flapping-wing systems do fly nowadays [30, 22], they do not exploit the whole potential of this flying mode. In particular, biological observations [9] suggest that the control of four DOF per wing may be required to produce most of a bird's performance.

2.1 Closed-loop straight-forward flight

According to such considerations, we used a genetic algorithm to first evolve controllers for a two-winged engine, each wing being considered as made of two panels - a proximal one, close to the body, and a distal one, close to the wing tip. These controllers were networks of non-linear oscillators and classical neurons whose number, inner parameters and inter-connexions were settled by the evolutionary algorithm. They could be connected to a speed sensor which gave the difference between the aircraft's effective speed value and a target value the controller had to keep. Likewise, they could be connected to a wing's four DOF – the dihedral and twist of the proximal panel, and the twist and sweep of the distal panel – the two wings being supposed to beat in perfect symmetry (figure 2). Controllers securing a forward flight at constant speed and altitude, despite horizontal and vertical perturbations, were sought through a multi-objective optimization procedure taking four criteria into account: lift, traction, energy and stability.

These experiments were done using a custom-built simulator. This simulator, which has been partially validated through wind-tunnel measurements, was able to compute a good approximation of the aerodynamic forces resulting from the interaction of the airflow and a rigid panel, for any angle of attack.

After 1000 generations, efficient flying behaviours were generated. In particular, the artificial evolutionary process discovered kinematics generating traction and lift, whereas we didn't provide any information about the potential interest of such forces in the current context. This effect depended upon the following features:

- the twist reaches its maximum value at the middle of the down-stroke, and its minimum value at the middle of the up-stroke;
- the wing folding is maximal during the up-stroke.

Two different strategies may be observed, which possibly have some equivalence in the behaviour of real birds.

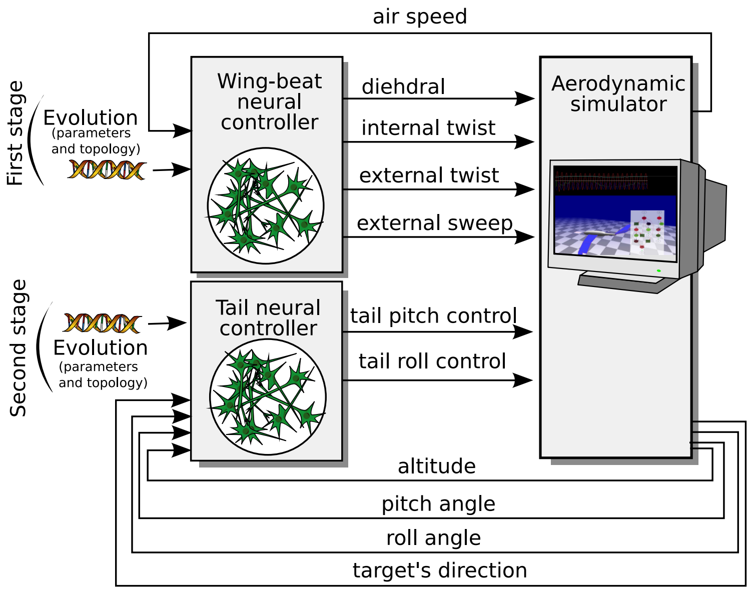

Figure 2: Overview of the control loops that were used in two stages: the first one served to evolve flapping-wing controllers, the second one to evolve tail controllers.

The first one consists in increasing the beating amplitude when accelerating, while decreasing it when slowing down. The internal and external twists are synchronized. Their amplitudes are accordingly adapted together with that of the dihedral: the greater they are, the faster the bird flies.

The other strategy consists in adapting the external twist only. Surprisingly, the external twist does not rely on any oscillator, and thus is not synchronized with wing beats. Actually, it is just controlled by a simple proportional controller with the relative air speed as input. As this part of the controller doesn't differentiate the up-stroke and the down-stroke, one may fear that a wrong twist during the up-stroke might dramatically increase the drag. However, this is not the case because the external sweep allows to fold the wing during this critical phase, thus minimizing the external twist effect. Thus, evolution discovered efficient solutions that use the internal panel to generate lift, and the external panel to generate traction, through two separate processes that were mixed in the former strategy.

Further details are to be found in [14].

2.2 Target-following flight

In a first attempt to obtain a fully-functional controller, we tried to evolve a tail controller to control the altitude, the pitch and the heading of a simulated bird.

We selected a flapping-wing controller previously evolved as described in the previous section, and we let the evolutionary algorithm search for a neural network able to efficiently control the tail. Although such an incremental approach - that divides the problem in two stages: evolving flapping-wing controllers for horizontal and straight flight first, then additionally evolving the capacity of using the tail for direction control - has proved to be inefficient in several contexts [6, 7] and is probably sub-optimal here, it has been used as a mean for providing reference results and for guiding further experimental work.

In these experiments, the inputs of the tail-controller were the altitude, the pitch and roll angles, as well as the target direction, whereas its outputs were the tail's pitch and roll effectors (figure 2).

While some minimally-efficient controllers were obtained [15], turns out that none of them was able to generate a sharp turn, probably because a symmetrical wing-beat controller is not adapted to this kind of manoeuvres. Additionally, these results suggest that, instead of artificially splitting the evolutionary process in several stages, it is probably wiser to simultaneously evolve two co-adapted wing-beat and tail controllers.

3 Low-level reflexes

Once a controller is able to stabilize the artificial bird and control its trajectory, some new skills are required to limit possible trajectories to those that will not damage the platform. This requires to perceive the UAV's environment and especially the surrounding obstacles that must be avoided.

As the available payload is limited, it is interesting to use sensors that may have several functionalities, like visual sensors that are currently used to provide visual feedback to human UAV operators. In particular, drawing inspiration from some flying animals, the optic flow detected by visual sensors may be exploited, together with simple controllers, to avoid lateral and frontal obstacles, and to adapt the flying speed to the dangerousness of the environment.

Lateral obstacle-avoidance relies on some properties of the optic flow created by forward translation, according to which perceived motions are inversely proportional to obstacles' distances. A simple strategy, which equalizes the optic flow perceived on both sides of the visual sensor, can thus make a UAV fly in the middle of corridors. Such a strategy, called the *balance strategy*, has been observed in flies and bees [24, 8, 25].

A second strategy may be used to avoid frontal obstacles. In pure longitudinal translation, a close object will generate more optic flow than a distant one, and a so-called time-to-collision can be evaluated on this basis for each pixel. When the mean time-to-collision becomes too low in front of the UAV, a dedicated reflex can be triggered to avoid collision. This strategy has been observed in gannets, that need to precisely evaluate when they will enter the water to decide to fold their wings [10, 11, 12].

This time-to-collision, when averaged over the whole visual field, provides information about the environmental clutter. A low value corresponds to an environment with lots of obstacles, while a high value denotes a more open space. Accordingly, a simple proportional controller may use this information to automatically adapt the UAV's maximal speed.

We implemented these three strategies in a realistic helicopter simulation calling upon a virtual 3D city. They allowed the helicopter to fly in three more or less cluttered environments without any collision, and they adapted the helicopter's speed to the local state of the environment, flying faster in open spaces, and slower near obstacles [17, 16].

4 Simultaneous Localization And Mapping

To be fully autonomous, our artificial bird will need to be able to build a map of its environment and to localize itself inside this map: this may help it to perform take-off and landing, to navigate, to memorize were to go to refill its batteries, or to localize which goal to reach to accomplish its mission. To this end, Simultaneous Localization And Mapping abilities (SLAM) are required.

The SLAM problem in mobile robotics has been addressed since 1987 [23], with some success on ground mobile robots (e.g. [5], [27]) through the use of Kalman or particle filters mixing sensory information with wheel-encoded odometry. However, this traditional scheme is generally relying on precise range-sensors such as lasers, radars or sonars that cannot be easily adapted on small UAV for which payload and energy are limited. Moreover, wheel-encoded robot odometry is inapplicable in the case of a UAV.

The use of small camera systems seems to be a good alternative as they are cheap, light, easy to manage, and may be used for other purposes, as previously demonstrated. However, while range-sensors directly provide the coordinates of landmarks relatively to the robot's position, when using vision, the corresponding information has to be extracted from the images. Some new image processing algorithms, like SIFT [13], allow accurate keypoint detection and large baseline matching even when the differences between images are large. Vision-based SLAM systems can therefore be designed with efficient substitutes to range-sensors and wheel-encoded odometry, as demonstrated, for example, in [4] and [20].

In the context of the Robur project, we implemented such a purely vision-based SLAM system for 2D MAV navigation [1]. We used a Kalman filter for simultaneously building a 2D metric map of visual ground landmarks and accurately computing an aircraft's localization in this map. 2D localization is a suitable preliminary step in the perspective of implementing cognitive behaviors like soaring, for example. In our approach, visual odometry is performed using SIFT feature-matching between consecutive images. In order to improve the map precision, we also estimate the radial distortion coefficient of the camera on-line, as an additional parameter in the Kalman filter. We obtained several conclusive results from images grabbed by the TwinStar UAV of the Paparazzi team[1] and by a home-made blimp.

5 Soaring

The platform we intend to build is supposed to be as generic as possible. Whatever its mission, it is supposed to behave in a way that will maximize its energetic autonomy. Birds

[1] www.nongnu.org/paparazzi

have the same constraints and found the solution long ago : slope winds, wind gradients, or thermals provide the necessary energy to remain aloft without efforts. Albatrosses, for instance, are able to fly for days without even flapping their wings [18], thanks to adapted trajectories above the waves.

We succeeded to reproduce such behaviours, and to let a simulated glider indefinitely fly in a wind gradient [3]. To this end, the glider must follow a very precise trajectory, according to which it starts diving, wind in the back. Near the water surface, it sharply turns to face the wind, and exploits the gained speed to reach the same altitude it started from (figure 3).

The corresponding controller was implemented with fuzzy rules [26], of which a preliminary set was empirically hand-designed. This set generated a globally correct trajectory, but poorly robust, as the glider ultimately crashed after a few cycles. Using an evolutionary algorithm, it has been possible to optimize this set of rules and to indefinitely avoid crashes. The robustness of this evolved controller to initial conditions, to the glider's morphology, and to sensory noise has been studied. The corresponding results indicate a poor sensibility to initial conditions, but a high sensibility to sensory noise. Likewise, the range of admissible flying directions relative to the wind has been evaluated, and revealed to be relatively small, as the glider wasn't able to remain aloft with an angle greater than 53^o relative to the wind. If further work remains then to be done to apply this approach to a real platform, or to exploit other energy-saving opportunities like slope winds for instance [18], it appears at this stage that a simple, but accurately tuned, controller can implement such a complex behaviour as soaring.

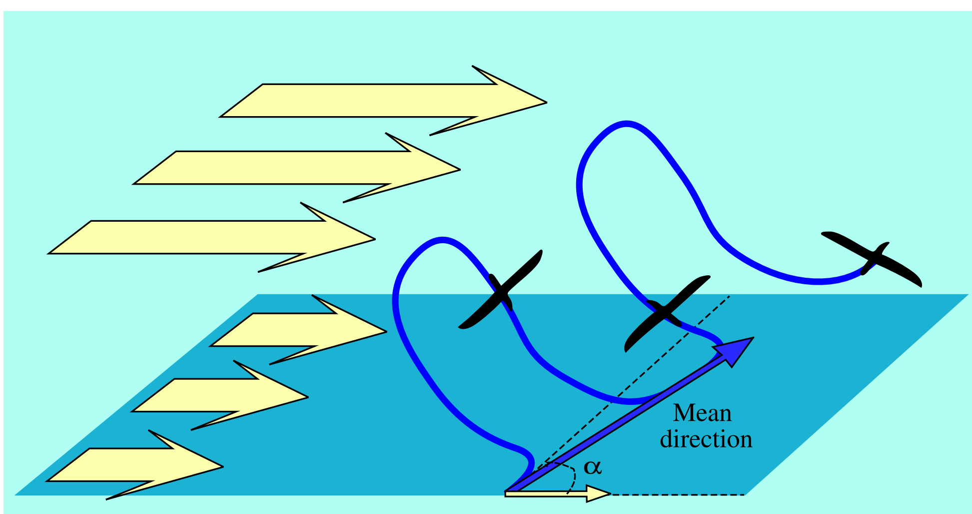

Figure 3: Trajectory used by an albatross during dynamic soaring.

Further work will be devoted to the evaluation of aerological conditions and to the autonomous planning of energy saving trajectories leading to a goal point. This will entail mandatory compromises between mission priority and energetic efficiency.

6 Discussion

The different results presented here have been obtained in simulation, except for the SLAM experiment. Several prototypes - notably a blimp, a motor-glider, a fixed-wing plane, and a quadri-rotor indoor helicopter - are currently developed that will make real-system applications possible. Likewise, a dedicated flapping-wing platform is expected to be available in the near future.

A critical point that has not been tackled yet is the integration of all the adaptive capacities that are, at the moment, necessarily developed more or less independently. However, the results already obtained provide some insights about the feasibility of such long-term objective. If one may easily foresee that the more stable the platform, the easier the integration will be, it also appears that the low-level reflexes responsible for obstacle-avoidance and speed-adaptation are relatively independent, and that it will be easy to connect them to the low-level flapping-flight controller. However, the perturbations created by the flapping movement will need to be as much damped as possible, and their effect on the optic flow computation system will need to be evaluated. As for the SLAM system and the soaring controllers, they are higher-level systems that will rely on both the flapping-flight controller and on the implemented reflex systems. The SLAM system is relatively independent and passive, at least for the moment. Running as a back-ground process to autonomously localize the future bird, it should not impose any particular constraint, except concerning the camera direction and, possibly, the exploration strategy for map-building. The soaring system will rely on two sub-systems. The first one will implement specific soaring behaviours, i.e., dynamic, slope or thermal soaring. This sub-system will choose the direction to follow and provide it as an input to the lower-level controllers previously described. The second sub-system will be connected to the SLAM system and to the mission planner to select a trajectory exploiting aerology to both save energy and reach the goal. The exact way such connection should be done remains to be specified, but is not expected to have a deep impact on other parts of the system.

7 Conclusion

We have presented the ROBUR project, whose goal is to build an autonomous flapping-wing system. This project focuses on decisional autonomy and aims at building the controllers required for this goal, from platform stabilization and control, to high-level systems providing decisional autonomy.

Current research efforts dedicated to the evolution of flapping-flight controllers based on neural networks and non-linear oscillator have been described. Efficient controllers able to adapt an artificial bird's speed and to turn in the direction of a given target have been generated with our methodology.

Obstacle-avoidance controllers have also been designed within the framework of this project. They rely on visual motion detection and make it possible to adapt the aircraft's speed to the degree of environmental clutter.

The Robur project also aims at studying higher-level capacities, like simultaneous localization and mapping. The system we described evaluates and memorizes landmarks' positions, as well as the relative position of the aircraft. The system is then able to determine its absolute trajectory and to return to its starting position without a GPS.

The last part of the system is responsible for energy-saving behaviours and aims at exploiting aerological conditions, like thermals, slope winds or wind gradients, just as birds do.

8 Acknowledgements

This work has been supported by a BQR grant from the Université Pierre et Marie Curie - Paris 6. Emmanuel de Margerie also benefited from a DGA/D4S post-doc grant.

References

[1] A. Angeli, D. Filliat, S. Doncieux, and J.-A. Meyer. 2d simultaneous localization and mapping for micro aerial vehicles. In *European Micro Aerial Vehicles (EMAV)*, 2006. to appear in EMAV 2006.

[2] P. Augustsson, K. Wolff, and P. Nordin. Creation of a learning, flying robot by means of evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO02)*, 2002.

[3] R. Barate, S. Doncieux, and J.-A. Meyer. Design of a bio-inspired controller for dynamic soaring in a simulated uav. 2006. submitted for publication.

[4] A.J. Davison. Real-time simultaneous localisation and mapping with a single camera. In *Ninth IEEE International Conference on Computer Vision (ICCV'03)*, 2003.

[5] M. W. M. Dissanayake, P. Newman, S. Clark, H. F. Durrant-White, and M. Csorba. A solution to the simultaneous localization and map building (slam) problem. *IEEE Transactions On Robotics and Automation*, 17(3):229–241, 2001.

[6] S. Doncieux and J.-A. Meyer. Evolving neural networks for the control of a lenticular blimp. In G. R. Raidl et al., editor, *Applications of Evolutionary Computing, EvoWorkshops2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, EvoSTIM*. Springer Verlag, 2003.

[7] S. Doncieux and J.A. Meyer. Evolution of neurocontrollers for complex systems: alternatives to the incremental approach. In *Proceedings of The International Conference on Artificial Intelligence and Applications (AIA 2004)*, 2004. (to appear).

[8] N. Franceschini, J.M. Pichon, and C. Blanes. From insect vision to robot vision. *Philosophical Transactions of the Royal Society of London B*, 4(4):283–294, 1992.

[9] T.L. Hedrick, B.W. Tobalske, and A.A. Biewener. Estimates of circulation and gait change based on a three-dimensional kinematic analysis of flight in cockatiels (nymphicus hollandicus) and ring ed turtle doves (streptopelia risoria). *Journal of Experimental Biology*, 205:1389–1409, 2002.

[10] D. N. Lee. A theory of visual control based on information about time-to-collision. *Perception*, 5:437–459, 1976.

[11] D. N. Lee. The optic flow field: The foundation of vision. *Philosophical Transactions of the Royal Society of London B*, 290:169–179, 1980.

[12] D. N. Lee and P. E. Reddish. Plummeting gannets: A paradigm of ecological optics. *Nature*, 293:293–294, 1981.

[13] D.G. Lowe. Distinctive image feature from scale-invariant keypoint. *International Journal of Computer Vision*, 60(2):91–110, 2004.

[14] J.-B. Mouret, S. Doncieux, T. Druot, and J.-A. Meyer. Evolution of closed-loop neuro-controllers for flapping-wing animats. 2006. submitted for publication.

[15] J.-B. Mouret, S. Doncieux, and J.-A. Meyer. Incremental evolution of target-following neuro-controllers for flapping-wing animats. In *Proceedings of the 9th International Conference on the Simulation of Adaptive Behavior (SAB)*, Rome, Italy, 2006. (to appear).

[16] L. Muratet, S. Doncieux, Y. Briere, and J.-A. Meyer. A contribution to vision-based autonomous helicopter flight in urban environments. *Robotics and Autonomous Systems*, 50(4):195–209, 2005.

[17] L. Muratet, S. Doncieux, J.-A. Meyer, P. Pirim, and T. Druot. Système d'évitement d'obstacles biomimétique basé sur le flux optique. Application à un drone à voilure fixe en environnement urbain simulé. In *Proceedings of Journées MicroDrones. CD-ROM ENSICA/SupAero*, Toulouse, 2003.

[18] C. J. Pennycuick. The flight of petrels and albatrosses (Procellariiformes), observed in South Georgia and its vicinity. *Phil. Trans. R. Soc. Lond. B*, 300:75–106, 1982.

[19] T. Rakotomamonjy, T. Le Moing, and M. Ouladsine. Kinematics optimization for a flapping-wing micro air vehicle. In *IFAC*, 2005.

[20] P. Saeedi, P.D. Lawrence, and D.G. Lowe. Vision-based 3-d trajectory tracking for unknown environments. *IEEE transactions on robotics*, 22(1):119–136, February 2006.

[21] L. Schenato, X. Deng, and S. Sastry. Flight control system for a micromechanical flying insect: architecture and implementation. *Robotics and Automation*, 2:1641–1646, 2001.

[22] L. Schenato, W. Chung Wu, and S. Sastry. Attitude control for a micromechanical flying insect via sensor output feedback. *IEEE Transactions on Robotics and Automation*, February 2004.

[23] R. Smith, M. Self, and P. Cheeseman. A stochastic map for uncertain spatial relationships. In *Workshop on Spatial Reasoning and Multisensor Fusion*, 1987.

[24] M.V. Srinivasan, M. Lehrer, W.M. Kirchner, and S.W. Zhang. Range perception through apparent image speed in freely flying honeybees. *Visual Neuroscience*, 6:519–535, 1991.

[25] M.V. Srinivasan, S.W. Zhang, M. Lehrer, and T.S. Collett. Honeybee nagigation en route to the goal: Visual flight control and odometry. *The Journal of Experimental Biology*, 199:237–244, 1996.

[26] T. Takagi and M. Sugeno. Fuzzy identification of systems and its application to modelling and control. In *IEEE Transactions on Systems, Man and Cybernetics*, volume 15. 1985.

[27] S. Thrun, M. Montemerlo, D. Koller, B. Wegbreit, J. Nieto, and E. Nebot. Fastslam: An efficient solution to the simultaneous localization and mapping problem with unknown data association. *Journal of Machine Learning Research*, -:–, 2004.

[28] B. Tobalske and K. Dial. Flight kinematics of black-billed magpies and pigeons over a wide range of speeds. *Journal of Experimental Biology*, 199(2):263–280.

[29] J. Wu and Z. Popovič. Realistic modeling of bird flight animations. *ACM Trans. Graph.*, 22(3):888–895, 2003.

[30] J. Yan, R. Wood, S. Avandhanula, M. Sitti, and R. S. Fearing. Towards flapping wing control for a micromechanical flying insect. In *IEEE Int. Conf. Robotics and Automation*, 2001.