Autonomous controllers based on neural networks for Micro-unmanned Aerial Vehicles: an Evolutionary Robotics approach

Fabio Ruini
Centre for Robotics and Neural Systems
University of Plymouth, UK

Overview

The research project described herein focuses on the design of autonomous controllers for MAVs (Micro-unmanned Aerial Vehicle). The design methodology adopted is Evolutionary Robotics (ER), which relies on the combined use of Neural Networks (NNs) and Evolutionary Algorithms (EAs). The main idea behind ER is that a robot is endowed with a set of sensors that allow it to 'sense' the surrounding environment. A neural network controls the robot's behaviour processing the information gathered by these sensors and producing a correspondent motor response. A proper set of connection weights and biases for the neural network is identified via an evolutionary algorithm which starts testing the performance obtained by 'random' controllers and progressively attempts to improve these until the desired behaviour is obtained. Given the nature of the design methodology adopted, computer simulations are required. Preliminary results have already been obtained using a simplified 2D model reproducing the layout of Canary Wharf, London. The ongoing work presented here focuses on the use of a more accurate 3D computer simulator instead.

Neural network controller

In the basic scenario, the controller used is a two-layer feed-forward NN. The network receives input information about the distance from the target, as well as the difference in horizontal and vertical angles compared to the 'ideal' trajectory. The output consists of two continuous neurons, generating real values used by the MAV to control yaw and pitch rotations, and one general purpose Boolean unit.

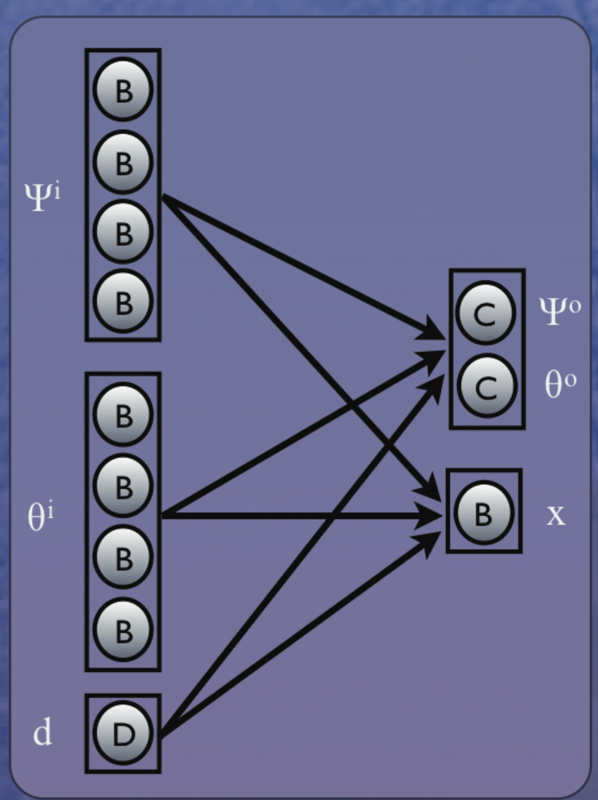

With the increasing complexity of the tasks analyzed, controllers with hidden layers, short-term memories and additional inputs are introduced. The encoding of the input information (normalised within a certain limited continuous range or discretised) varies as well for the different scenarios.

Basic navigation

In the first experimental setup, a single MAV is used. At the beginning of each test epoch the target is deployed at a random position within an obstacle-free environment - a parallelepiped box with the dimension 1,000 (X) x 1,500 (Z) x 600 (Y). The MAV has to navigate to a certain target area and, upon reaching it, perform a specific operation (represented as the activation of the Boolean output unit, which can only be operated once per epoch). The results have demonstrated that a controller fit for the execution of the above task can evolve in just a few hundreds generations. At the end of the evolution, the average MAV in the population is able to carry out the task 64.91% of times, while the best one scores 99.52%.

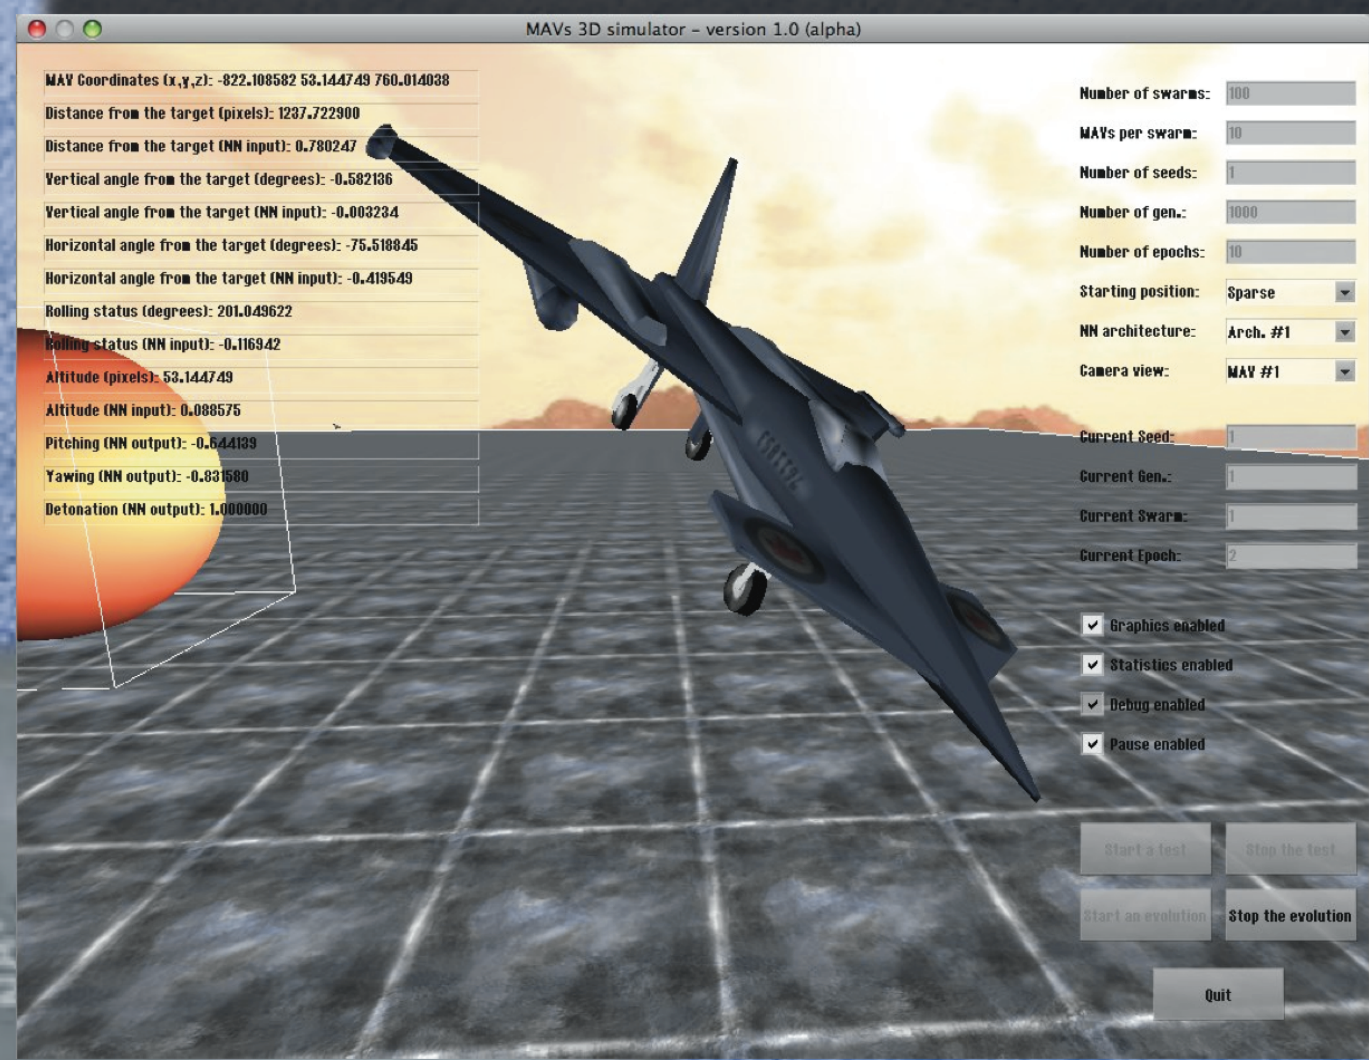

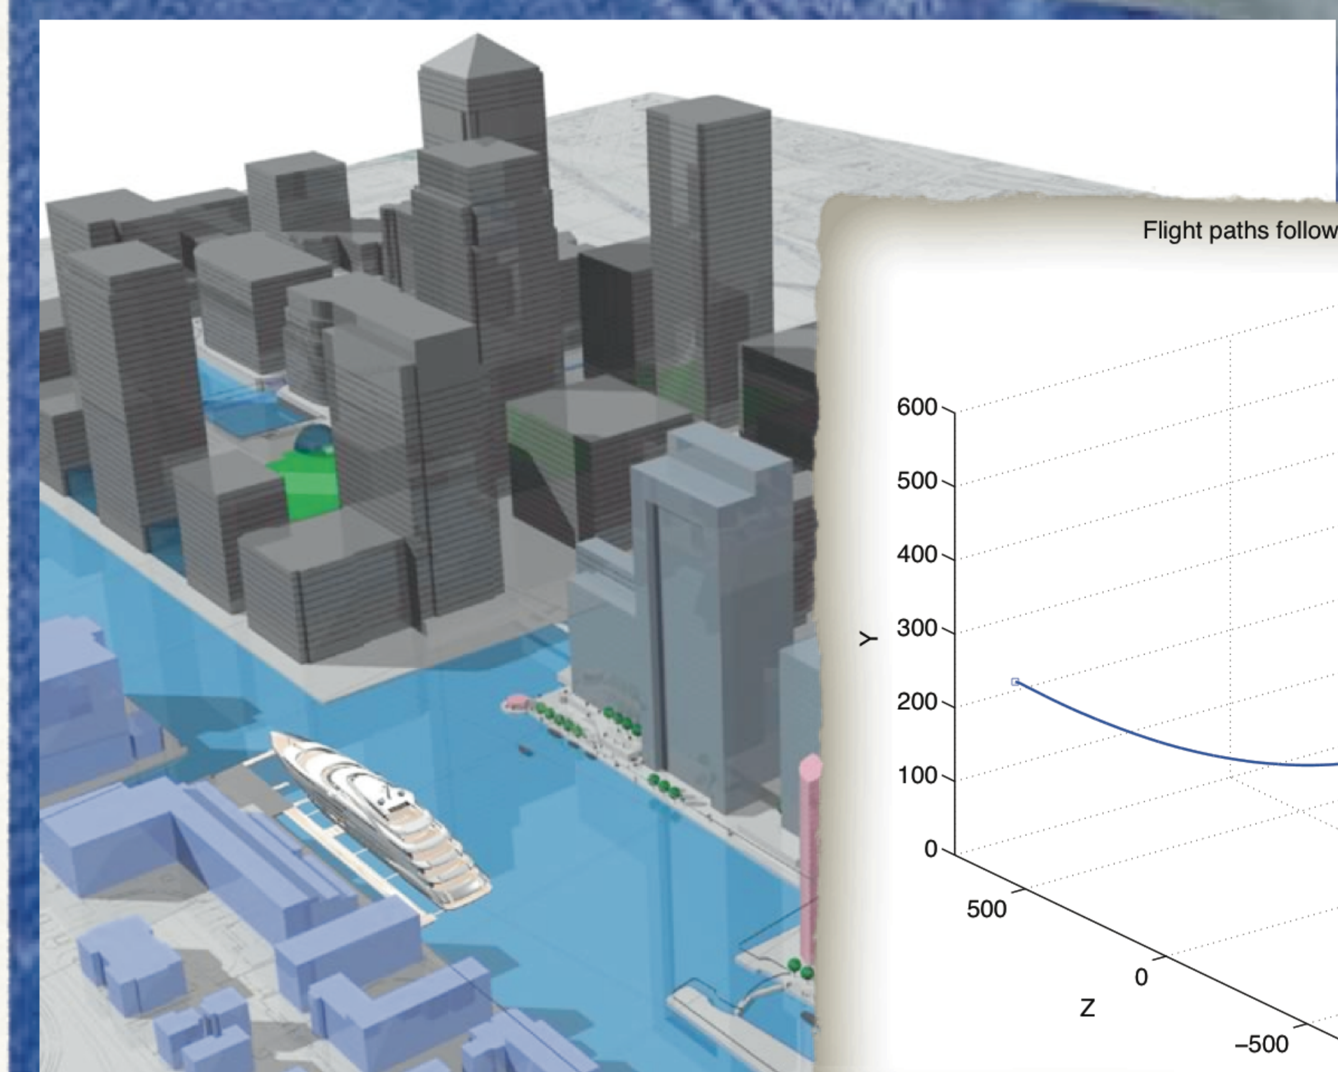

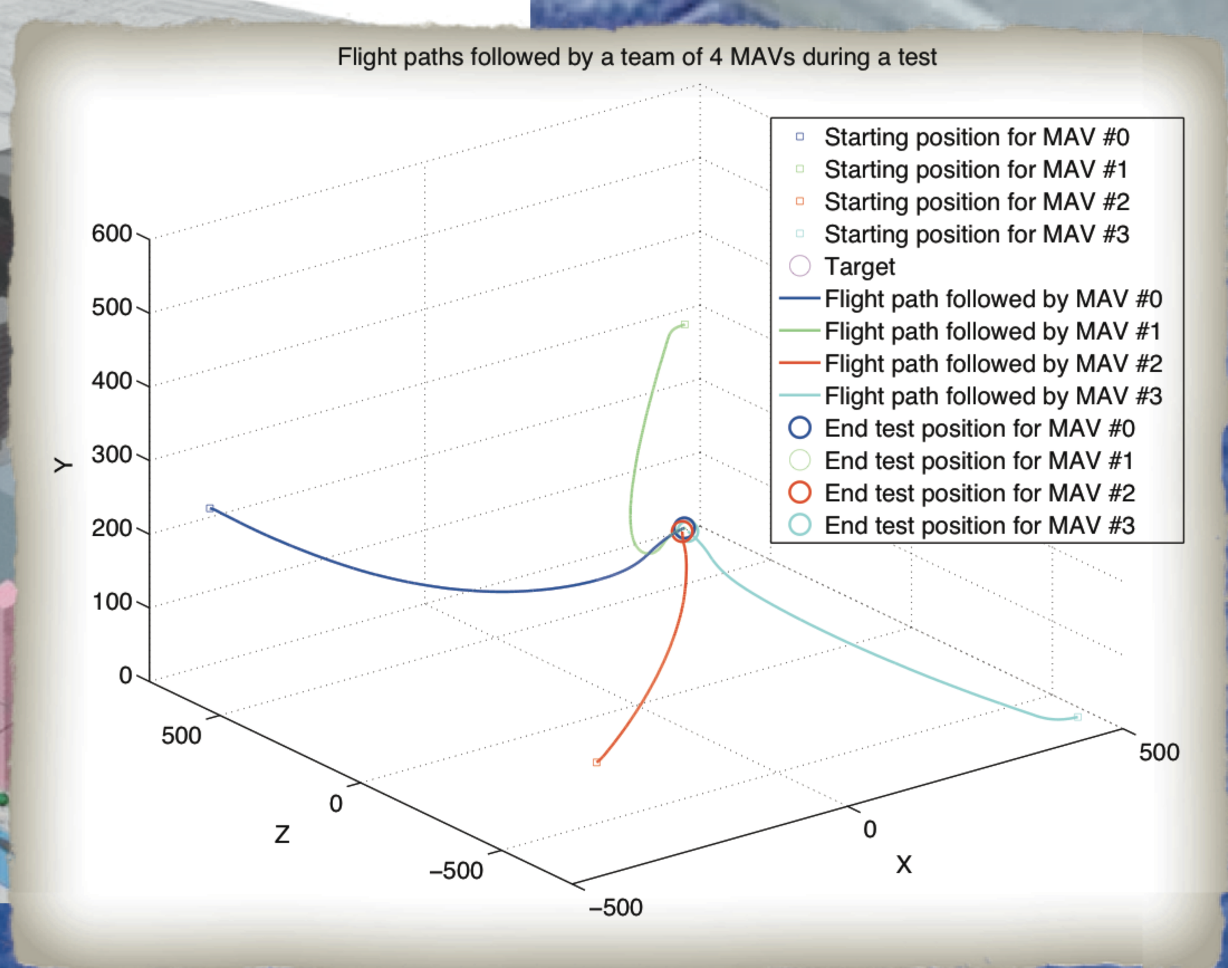

Top-right: a screenshot of the 3D simulator; top-left: the architecture of one of the NN controllers used; above this label: a 3D model of Canary Wharf (http://www.aquiva.co.uk/canarywharf) which will be implemented in future versions of the simulator; at right of this label: trajectories followed during a test by a team consisting of 4 MAVs sharing the same controller.

Movable target

The second experimental condition adds to the basic one a target able to detect an approaching MAV and then attempting to move away from it. The target detects the aircraft when within a certain distance range and 'escapes' at different speeds according to the setup analysed. The results have suggested that an MAV can keep a success rate comparable to the fixed-target scenario for a target moving up to half the speed of the aircraft (best success rate: 98.75% when target speed is 1/5th, 99.55% when 1/4th, 97.07% when 1/3rd, 90.5% when half). Over that threshold, the performance significantly decreases.

Implicit cooperation

In this experimental setup a team consisting of two MAVs is employed. In the test the team is subject to, the MAVs have to reach the target area independently and then, once there, activate the Boolean output neuron of their networks in quick succession. Compared to previous scenarios, the only additional information each MAV can now rely on is: (1) the presence of the teammate within a certain distance; (2) the fact that the teammate has recently activated its Boolean neuron inside the target area. With a target unable to move, the average success rate for the entire population reaches the 45.3%, while the best team scores a 81.92%. When the target can move, these values are 36.2% and 72.33% respectively.

Selected publications

Ruini, F., and Cangelosi, A. (2009), Extending the Evolutionary Robotics approach to Flying Machines: an Application to MAV Teams. *Neural Networks*, 22, 812-821.

Ruini, F., Cangelosi, A., and Zetule, F. (2009), Individual and Cooperative Tasks performed by MAV Teams driven by Embodied Neural Network Controllers. In *Proceedings of IJCNN 2009, International Joint Conference on Neural Networks* (pp. 2717-2724).

Ruini, F., and Cangelosi, A. (2008), Distributed Control in Multi-Agent Systems: A Preliminary Model of Autonomous MAV Swarms. In *Proceedings of FUSION 2008, Eleventh International Conference on Information Fusion* (pp. 1043-1050).

Efforts sponsored by the Air Force Office of Scientific Research, Air Office Material Command, USAF under grant number FA8655-07-1-3075.