EVOLUTIONARY ALGORITHM BASED NEURAL NETWORK CONTROLLERS: AN APPLICATION TO MAV SWARMS

FABIO RUINI
ANGELO CANGELOSI
Adaptive Behaviour and Cognition Research Group
School of Computing, Communications and Electronics
University of Plymouth
Drake Circus, PL4 8AA, Plymouth, Devon, United Kingdom

1. Introduction

During the last decade several studies have been carried out on wheeled and underwater autonomous robots driven by neural network controllers (e.g. [1] and [2]). The application of the same principles to flying machines has not yet been investigated thoroughly. With the only notable exception of the systems developed by Floreano [3] and Holland [4], it seems that the current approaches on the development of autonomous controllers for aircraft mainly rely on techniques other than neural networks, that is behaviour-based robotics and genetic programming [5].

The work presented herein focuses on the use of embodied neural network controllers for MAV (Micro-unmanned Aerial Vehicles) swarms. The goal of this research is to demonstrate how evolutionary autonomous controllers for flying robots can be successfully developed through computer simulations based on multi-agent systems methodologies.

2. Description of the model[1]

The experiments outlined in the following sections use a "search and destroy" scenario in the context of urban counter-terrorism. The environment where the simulations take place is a two-dimensional rectangular area representing a portion of London's Canary Wharf. The target, which corresponds to a person/vehicle to neutralize, is deployed in a random position inside this map. A MAV swarm, composed by four unmanned aircraft, has to navigate through the

[1] The specifications of the MAVs employed in these simulations (size, speed and autonomy) have been inspired by the WASP Block III, produced by the American manufacture Aerovironment (http://www.avinc.com/downloads/WASP-III_datasheet_6_5_07.pdf).

environment to reach the target and finally to neutralize it carrying out a certain operation. The conclusive action performed by the MAV, despite of its success, provokes the loss of the aircraft.

The fundamental assumption on which this model relies is that the MAVs are always aware of the target's position.

The robots' behavior is governed by a three-layered feed-forward neural network that receives the sensorial inputs from the environment and in turn triggers the appropriate motor answer. Even if each individual is endowed with its own neural controller, the MAVs belonging to the same swarm share the same connection weights: they are, in fact, clones of each other.

The MAVs' controllers evolve through a genetic algorithm in which elitism and random mutations are the operators used (a more detailed description of the model can be found in [6]).

3. Experimental setups

Simulations have been carried out on four different experimental setups, as summarized in the following paragraphs.

3.1. *Plain environment*

The first experimental setup aims to identify the minimal set of inputs required by the network in order to evolve the desired behavior.

The simulated environment is free from any obstructions. Each swarm is tested four times, with the target placed in different positions.

The results obtained from these simulations suggest that the optimal setup consists of four input neurons. One encoding the distance between the MAV and the target (ranging from 0 to 1 with discrete intervals) and the others three representing the relative angle between the two agents (approximated through a three-bit Gray Code representation). On average, the population evolved with this set of sensors is able to successfully carry out the task 93.46% of the times.

3.2. *Environment with obstacles*

In the second set of simulations we have inserted some obstacles into the map, corresponding to the location and extension of the tallest buildings present in the urban area we are using as a model. For the sake of simplicity, the simulated environment is still two-dimensional. The buildings represent for the MAVs a "no-fly zone": if they try to enter these areas, they will be immediately destroyed.

We contrast different setups of ultrasonic sensors capable of detecting the presence of an object situated in front of each MAV, along a straight line.

The results show how the best configuration consists of three sensors. One sensor is in front of the MAV and the other two are respectively oriented at -20° and +20° with respect to the aircraft's facing direction. Simulations results indicate that on average 0.22 aircrafts crashed against a building during a test, with 87.18% of tests successfully concluded.

3.3. *Coordinated action*

In the third set of simulations two aircrafts have to reach the target and attack it in quick succession in order for the test to be considered succeeded. To allow the evolution of this coordinated behavior, the MAVs have been endowed with two new inputs. They are two Booleans neurons that get activated when the MAV perceives the presence of a teammate within a certain range and when the target is damaged (i.e., it had recently received a hit) respectively.

The emergent behavior is straightforward. When the first MAV reaches the target it starts to turn around it, waiting for the arrival of a teammate. Then, when another aircraft finally arrives, one of the two MAVs attacks the target, quickly imitated by the other.

The results obtained are encouraging, given the more complicated scenario. 73% of tests succeeded, with 88.14% of times when the target is hit at least once.

3.4. *Movable target*

The last experimental setup introduces a target able to move, trying to escape from the approaching MAVs. When the target detects the presence of a MAV within a certain range, it moves (at one sixth of the aircraft's speed) following the direction that will maximize the distance between the two agents.

The results show how the presence of a movable target does not meaningfully affect the swarms' performance until the task does not require to be performed in a cooperative way (83.38% of tests succeeded). Instead, when a movable target has to be neutralized via a coordinated action, the MAVs' performance dramatically decreases (41.28% of tests succeeded, 73.45% of times the target is hit at least once).

4. Conclusions and future developments

The results summarized in this abstract indicate how evolutionary neural network controllers could be successfully employed in the domain of flying robots.

The results are particularly interesting when we consider the simplicity of the neural network architectures used, which operate in real-time without requiring any kind of memory and relying just on a very basic set of input sensors. Nevertheless the networks are able to achieve sophisticated behaviors such as navigate through unknown environments and perform tasks requiring coordination.

Future work will proceed following two main research directions. First, we will introduce explicit forms of communication between the MAVs, since the last experimental setup analyzed suggests that communication might positively affect the swarm's performance. Second, a three-dimensional/real-physics simulator will be developed in order to evolve controllers that might be more easily transferred from computer simulations to real robots.

Acknowledgments[2]

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-07-1-3075. The U.S. Government is authorized to reproduce and distribute reprints for Government purpose notwithstanding any copyright notation thereon.

Disclaimer

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

References

1. G. Baldassarre, D. Parisi and S. Nolfi, *Artificial Life*, 12-3 (2006).
2. V. S. Kodogiannis, *Int. Journal of Systems Science*, 37-3 (2006).
3. D. Floreano, S. Hauert, S. Leven and J. C. Zufferey, *Int. Symposium on Flying Insects and Robots* (2007).
4. R. De Nardi, O. Holland, J. Woods and A. Clark, *Proc. of 21st Bristol UAV Systems Conference* (2006).
5. M. D. Richards, D. Whitley and J. R. Beveridge, *Proc. of Genetic and Evolutionary Computation Conference (GECCO 2005)*, (2005).
6. F. Ruini and A. Cangelosi, *Proc. of the 11th Int. Conference on Information Fusion (FUSION 2008)* (2008).

[2] The authors would also thank euCognition (Network Action NA097-3) for the support provided during the preliminary phases of this research.