

Design and Analysis of Proximate Mechanisms for Cooperative Transport in Real Robots

Muhanad H. Mohammed Alkilabi^{1,2}, Aparajit Narayan¹, and Elio Tuci¹

¹Computer Science Department, Aberystwyth University, Aberystwyth, UK

²Computer Science Department, Kerbala University, Kerbala, Iraq
{mhm1, apn3, elt7}@aber.ac.uk

Abstract. This paper describes a set of experiments in which a homogeneous group of real e-puck robots is required to coordinate their actions in order to transport cuboid objects that are too heavy to be moved by single robots. The agents controllers are dynamic neural networks synthesised through evolutionary computation techniques. To run these experiments, we designed, built, and mounted on the robots a new sensor that returns the agent displacement on the x/y plane. In this object transport scenario, this sensor generates useful feedback on the consequences of the robot actions, helping the robots to perceive whether their pushing forces are aligned with the object movement. The results of our experiments indicated that the best evolved controller can effectively operate on real robots. The group transport strategies turned out to be robust and scalable to effectively operate in a variety of conditions in which we vary physical characteristics of the object and group cardinality. From a biological perspective, the results of this study indicate that the perception of the object movement could explain how natural organisms manage to coordinate their actions to transport heavy items.

Keywords: Swarm Robotics, Cooperative Transport, Evolutionary Robotics

1 Introduction

Collective object transport in a robot swarm is the ability of the robots to collect and transport objects that can not be transported by a single agent [5]. Cooperative transport is relatively ubiquitous in social insects, being known in at least 40 genera of ants [9]. It is primarily used to retrieve objects (e.g., food items) that are too heavy or too large to be moved by a single individual. For those ants species that live in environments in which the source of proteins are generally large carcass of insects, objects transport can be a solution to retrieve these precious food items reducing the time the food is exposed to competition [13]. It seems that a variety of parameters including the item's resistance to movement, the speed of transport, as well as the item size, shape, and mass play a significant role for the recruitment and for the active engagement of individuals into the transport [8]. However, cooperative transport in ants remains a poorly understood process, with various hypothesis concerning the mechanisms for alignment

and coordination of forces. There is not much empirical evidence to shed light on the proximate mechanisms underpinning this important cooperative process. In particular, it is still not clear what mechanisms are used to assess consensus or quorum information about directional movement [8]. Hypotheses vary from parsimonious explanations based on the perception of the object movement, to theories that require more complex structures for the perception of the forces exerted on the object, or for direct communication between the agents involved into the transport [11].

In recent years, the attempt of swarm roboticists to engineer groups of robots that generate interesting collective responses through self-organisation has provided biologists with an alternative method to investigate phenomena in social insects. The pioneering work of [7] on box-pushing by a multi-robot system has the merit of having formally represented in “hardware” the dynamics of collective object transport, pointing to issues of absolute relevance for a principled understanding of this form of cooperation. In the last 15 years, quite a few research works have tried to mimic the ants’ cooperative behaviour with the double aim of engineering robust and scalable multi-robot systems and understanding nature. In [3], the authors carefully observed under experimental conditions the behaviour of a colony of ants (*Aphaenogaster cockerelli*) engaged in a cooperative transport task. They created a detailed model based on qualitative analysis of the role and contribution of single ants during transport of food items to the nest. The collected data has been used to create a model of the ants’ behavioural rules during transport. The model has been validated by comparing the behaviour of simulated and real ants. In [12], the authors focus on the problem of alignment during transport showing that a robot leader that knows the direction of transport can induce the group to execute the desired cooperative manoeuvre by interacting with the group mates (i.e., followers) through forces exerted on the object. The swarm robotic model described in [6] demonstrated that communication between robots involved into the collective transport need not to be direct. Stigmergic forms of communication suffice to achieve coordination of forces and alignment in a group of robots retrieving heavy objects.

In this study, we describe a further swarm robotic model targeting cooperative transport. Real e-puck robots are required to push a cuboid object which, due to its mass/size, requires the cooperative effort of all the members of the group to be transported. The robots have to agree on a common direction of transport, to align their movements, and to push the object for an extended period of time. The distinctive feature of our model is the minimalist sensory apparatus provided to the robots. Contrary to the majority of previous similar studies, our robots have no means to feel forces. The objective of this study is to look at what the robots can collectively achieve with a sensory apparatus that allows them to indirectly perceive the movement of the object to transport. To run this study, we designed, built, and mounted a new sensor on the real e-pucks. This new sensor (hereafter, referred to as “optic-flow” sensor) is an optical camera positioned underneath the robot chassis, which returns the robot displacement on a x/y plane. In this collective transport scenario, the optic-flow

sensor, in combination with the distance sensors, generates a sensory stimulation that effectively informs the robot on the direction of movement of the object. The results of our study show that this simple feedback suffices to allow the robots to agree on a common direction of transport and to sustain the transport for an extended period of time. A significant contribution of our study is in showing the robustness of the transport strategies evaluated in different scenarios (i.e., with objects of different mass and length, and with groups of different cardinality), and in the analysis of the behavioural mechanisms used by the robots to coordinate their actions. Our results point to a rather parsimonious explanation of the mechanisms required by real ants to transport object. In particular, our results suggest that feedback on the movement of the object modulates the frequency with which a robot changes the point of application of its pushing forces. This modulation is sufficient for a robot to sense a quorum with respect to the direction of travel, and to break “deadlocks” in which the robots cancel each others’ forces. We also illustrate how the robots’ shape influences the group performances. Interesting future lines of investigation dictated by our results are also discussed in Section 6.

2 The Task and the Simulation Model

In this study, neuro-controllers are synthesised using artificial evolution to allow a homogeneous group of four autonomous robots to push an elongated cuboid object (30 cm length, 6 cm width and height, 600 g mass) as far as possible from its initial position. Our study is run with groups of e-puck robots [10]. The parameters of the neuro-controllers are set in a simulation environment which model kinematic and dynamic features of the experimental conditions in which the simulated e-pucks are required to operate. The robot sensory apparatus includes infra-red sensors, the camera mounted on the robot chassis, and the optic-flow sensor appositely designed, built, and integrated into the e-puck structure for this task (see Fig. 1a). In simulation, the robots are initially positioned in a boundless arena with flat terrain, at 50 cm from the object. The objective of the robots is to move the object 2 m away from its initial position. The object mass is set so that the coordinated effort of all four robots is required to move the object. The best evolved controller is ported onto real e-puck robots, and extensively tested in various experimental conditions. This paper mainly focuses on the results of the evaluations on real robots¹.

In the remaining of this Section, we illustrate the characteristics of the new optic-flow sensor, which is an optical camera mounted underneath the robot chassis and located inside the slot originally hosting the robot battery (see Fig. 1a). This sensor captures a sequence of low resolution images (i.e., 18x18 pixels) of the ground at 1500 frames per second. The images are sent to the on board DSP which, by comparing them, calculates the magnitude and the direction of

¹ A detailed description of the simulation environment, of the robot model, including noise applied to sensors and motors, as well as results of all re-evaluation tests, and movies can be found at <http://users.aber.ac.uk/e1t7/ANTS2016/>.

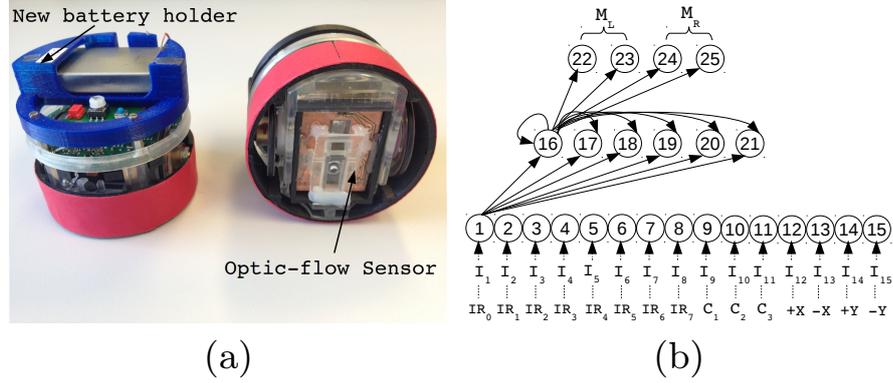


Fig. 1. (a) The e-puck robot with optic-flow sensor. (b) The robot's controller. The continuous line arrows indicate the efferent connections for only one neuron of each layer. Hidden neurons receive an afferent connection from each input neuron and from each hidden neuron, including a self-connection. Output neurons receive an afferent connection from each hidden neuron. Sensors to sensor neurons correspondence is indicated underneath the input layer.

movement of the robot. This information is subsequently communicated to the robot controller in the form of four normalized real values in $[0, 1]$: $+X$ and $-X$ representing the displacement on the positive and negative direction of the x axis, respectively; $+Y$ and $-Y$ representing the displacement on the positive and negative direction of the y axis, respectively. To improve portability of solutions to real hardware, in simulation, $+X$, $-X$, $+Y$, and $-Y$ are subjected to uniformly distributed random noise in $[-0.025, 0.025]$. The optic-flow sensor generates a sensory stimulus which is a direct feedback on the consequences of the signals sent to the motors. In a collective object transport scenario multiple contingencies can result in a robot failing to execute its desired action. For example, a forward movement command may not produce the desired action if the robot is pushing a stationary object, or an object that is moving in the opposite direction due to forces exerted by other robots. The optic-flow sensor generates readings that can be used by the agents to recognize these circumstances and to respond accordingly. The results of this study shows that this simple feedback, generated by the optic-flow sensor, is sufficient to allow a group of robots to coordinate their effort in order to collectively transport in an arbitrary direction an object that can not be moved by a single robot.

3 The Controller and the Evolutionary Algorithm

The robot controller is composed of a continuous time recurrent neural network (CTRNN) of 15 sensor neurons, 6 internal neurons, and 4 motor neurons (see [2] and also Fig. 1b which illustrates structure and connectivity of the network). The states of the motor neurons are used to control the speed of the left and

right wheels. The values of sensory, internal, and motor neurons are updated using equations 1, 2, and 3.

$$y_i = gI_i; i \in \{1, \dots, N\}; \quad \text{with } N = 15; \quad (1)$$

$$\tau_i \dot{y}_i = -y_i + \sum_{j=N+6}^{j=N+6} \omega_{ji} \sigma(y_i + \beta_j); i \in \{N+1, \dots, N+6\}; \quad (2)$$

$$y_i = \sum_{j=N+1}^{j=N+6} \omega_{ji} \sigma(y_j + \beta_j); i \in \{N+7, \dots, N+10\}; \quad (3)$$

with $\sigma(x) = (1 + e^{-x})^{-1}$. In these equations, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, g is a gain factor, I_i with $i = 1, \dots, N$ is the activation of the i^{th} sensor neuron (see Fig. 1b for the correspondence between robots sensors and sensor neurons), ω_{ij} the strength of the synaptic connection from neuron j to neuron i , β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate f_i . All sensory neurons share the same bias (β_I), and the same holds for all motor neurons (β_O). τ_i and β_i of the internal neurons, β_I , β_O , all the network connection weights ω_{ij} , and g are genetically specified networks' parameters. At each time step, the output of the left motor is $M_L = f_{N+7} - f_{N+8}$, and the right motor is $M_R = f_{N+9} - f_{N+10}$, with $M_L, M_R \in [-1, 1]$. Cell potentials are set to 0 when the network is initialised or reset, and equation 2 is integrated using the forward Euler method with an integration time step $T = 0.13$. A simple evolutionary algorithm using roulette wheel selection is employed to set the parameters of the networks [4]. The population contains 100 genotypes. Generations following the first one are produced by a combination of selection with elitism, recombination, and mutation. For each new generation, the eight highest scoring individuals (the elite) from the previous generation are retained unchanged. The remainder of the new population is generated by fitness proportional selection from the 60 best individuals of the old population. A detailed description of the evolutionary algorithm can be found in [1].

4 The Fitness Function

During evolution each group undergoes a set of $E = 12$ evaluations or trials. A trial lasts 900 simulation steps (i.e., 117 s, with 1 stimulation step corresponding to 0.13 s). A trial is terminated earlier if the group manages to displace the object 2 meters away from its initial position. At the beginning of each trial the controllers are reset, and the robots are positioned in the arena. Each trial differs from the others in the initialisation of the random number generator, which influences all the randomly defined features of the environment, such as the noise added to sensors and the robots initial position and orientation. The robots initial relative position with respect to the object is an important aspect which bears upon the complexity of this task. This is because the robots initial position contributes to determine the orientation with which they approach the object and consequently the nature of the manoeuvres required by the agents to

coordinate and synchronise their actions. During evolution, the robots starting positions correspond to randomly chosen points on a circle’s circumference of 50 cm radius that has the object in it’s centre. This circle is divided in four equals parts. Each robot is randomly placed in one part of this circle with random orientation in a way that the object can be within an angular distance of $\pm 60^\circ$ from its facing direction. These criteria should generate the required variability to develop solutions that are not sensitive to the robots initial positions.

In each trial (e), an evaluation function F_e rewards groups in which the robots remain close to the object, and transport the object as far as possible from its initial position. F_e is computed in the following:

$$F_e = f_1 + f_2 - f_3 \quad (4)$$

$$f_1 = \sum_{r=1}^R (1 - d_r); \quad f_2 = \Delta O^{pos}; \quad f_3 = t/T; \quad \text{with } T = 900; R = 4; \quad (5)$$

d_r is the Euclidean distance between the centroid of robot r and the centroid of the object. d_r is set to zero if the robot gets closer than 20 cm to the object. ΔO^{pos} is the Euclidean distance between the position of the object’s centroid at the beginning and the end of the trial. t is the trial duration in simulation steps. f_1 rewards groups in which the robots approach the object. f_2 rewards groups that transport the object as far as possible. f_3 rewards groups that performs the task faster (i.e., required less number of simulation steps). The fitness of a genotype (\bar{F}) is the average team evaluation score after it has been assessed $E = 12$ times: $\bar{F} = \frac{1}{E} \sum_{e=1}^E F_e$.

5 Results

The primary aim of this study is to design control systems for homogeneous groups of real e-pucks required to transport objects in a cooperative way. Our objective is to generate solutions that are robust with respect to the object mass and length, and scalable with respect to the group cardinality. To design the controllers, we run 20 differently seeded evolutionary simulations, each simulation lasting 3000 generations. In order to choose the controller to be ported onto the real robots, we re-evaluated, in simulation, the best genotypes from generation 1000 to generation 3000 for every run. During re-evaluations, groups of simulated robots are tested with objects of different length and mass. Moreover, the group cardinality and the robots initial positions and orientations are systematically varied (see description and results of re-evaluation tests at <http://users.aber.ac.uk/e1t7/ANTS2016/>).

The solution (i.e., the genotype coding for the controller) with the very best re-evaluation score has been selected to be ported onto the real e-pucks for further evaluations. In the next Section, we describe the results of the first test with real robots, and we compare these results with those of simulated robots controlled by the same controller, and evaluated in similar operational conditions.

5.1 First evaluation test with real e-pucks

The first evaluation test with real e-pucks has been designed to investigate the scalability of the controllers with respect to the number of robots in the group, as well as the robustness with respect to objects of different length and mass, and with respect to varying initial conditions. Recall that during evolution, we used only groups of 4 robots, and only one type of elongated cuboid object (30 cm length, 6 cm width and height, 600 g mass). During the test with real robots, we evaluated homogeneous groups of 3, 4, 5 and 6 real e-pucks, for their capability to collectively transport cuboid objects of 30 cm and 40 cm lengths. Objects of each length were tried with two different masses. Object width and height are not changed with respect to evolutionary conditions. The total number of re-evaluation trials (160) with real e-pucks is given by all the possible combinations of the above mentioned parameters (i.e., 2 lengths, 2 masses, and 4 different values for group cardinality), each combination repeated 10 times, 5 trials with all the robots positioned in front of the long sides of the objects, and 5 trials with all the robots positioned in front of the short sides of the cuboid objects. In each trial, half of the group faces one side and the other half faces the other side of the object. In order to enforce the requirement of collective transport, the masses of the object vary with respect to the cardinality of the group in a way that in each of the 160 evaluation trials, the object is heavy enough to require the combined effort of all robots of the group to be successfully transported. The object masses are indicated in Fig. 2a. In each evaluation trial, the object is placed in the centre of 220 cm bounded square arena, and the robots are placed at about 50 cm from the object. Each evaluation trial can last 180 s (i.e. 1384 simulation steps), and it is terminated earlier if the group manages to transport the object at least 1 m away from its initial position. Only in this later case, the trial is considered successful.

The results of the first evaluation test are shown in Fig. 2a, where the bars indicate the success rate (%) of homogeneous groups controlled by the best evolved neural network in 16 different evaluation conditions. Black bars refer to the performance of groups of real e-pucks; white bars refer to the performances of groups of simulated e-pucks evaluated in similar experimental conditions (i.e., same object length, same mass, same group cardinality, and approximately same robots initial positions). The comparison between real and simulated robots is meant to capture differences in performance when moving from simulation to reality. We notice that the performance of both real and simulated robots is close to or largely above 80% success rate in almost all evaluation conditions, demonstrating that the robots controller can successfully operate with larger groups than those used during the design phase, and with heavier and/or longer objects. At <http://users.aber.ac.uk/e1t7/ANTS2016/> the reader can find the results of further tests in simulations with groups of up to 16 robots. These tests could not be run on real e-pucks because in our Lab we only have 6 real e-pucks.

Results in Fig. 2a tell us that performances drop for the group of 6 e-pucks, transporting an object of 30 cm length, and of 980 g mass. This drop in performance can be explained with reference to two elements: the length of the

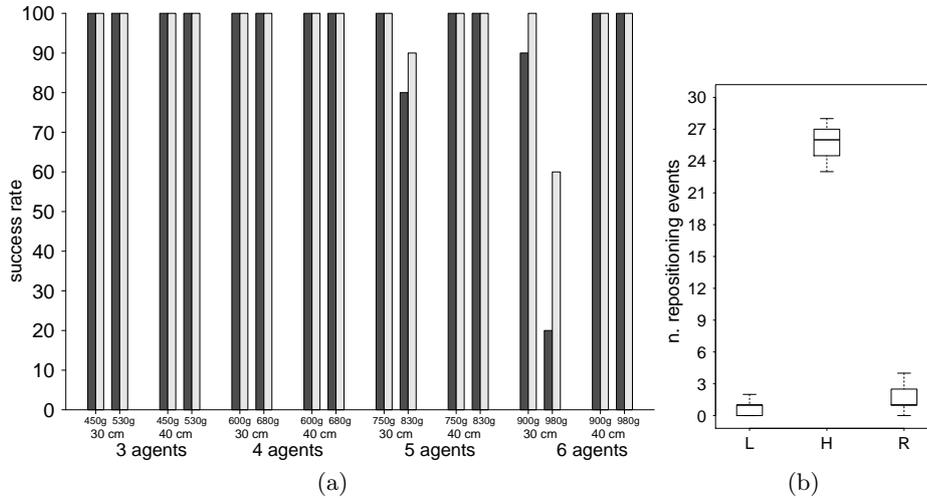


Fig. 2. (a) Graph showing the success rate (%) of homogeneous groups controlled by the best evolved neural network. Black bars refer to the performance of groups of real e-pucks; white bars refer to the performances of groups of simulated e-pucks. The x-axis shows mass and length of the object, and group cardinality. A trial is considered successful if the group manages to transport the object 1 m away from its initial position. Each bar refers to performances over 10 trials. (b) Graphs showing the results of evaluation tests with a single real robot. In test *L*, we use 30 cm length, and 150 g mass object. In test *H*, we use 30 cm length, and 600 g mass object. In test *R*, we use 400 g mass, and 40 cm length object.

longest size of the cuboid object (hereafter, referred to as \bar{L}), and the sum of the diameter of the robots in the group (hereafter, referred to as \bar{D}). The number of robots that are forced to indirectly push the object through physical contact with other robots progressively increases when \bar{L} becomes smaller than \bar{D} . The higher the number of robots pushing other robots, the higher the frequency of “detachment events” during transport. A detachment event refers to the case in which a robot loses physical contact with the element that it is currently pushing. Thus, it needs to relocate itself in a new position in order to keep on actively contributing to the collective transport. Detachment events have a negative impact on the group performance, since during such an event the group loses the contribution of one robot. Detachment events are more frequent when robots are required to push other robots than when robots directly push the object. This is because the e-pucks have a cylindrical shape which makes it relatively difficult for a robot to push another non-stationary robot. Generally speaking, we could say that the smaller the \bar{L} compared to \bar{D} , the higher the frequency of detachment events, the poorer the group performance. However, there are exceptions. As shown in Fig. 2a, the \bar{L} smaller than \bar{D} condition only minimally affects the performance of groups of 6 real e-pucks transporting a slightly lighter object (see Fig. 2a, 6 robots, 30 cm length, 900 g object, black bar). This is

Table 1. Table showing the number of repositioning events during each trial of the evaluation test in which a single robot pushes either a static object, or a non-static object intentionally moved in the opposite direction of the robot heading.

trial	1	2	3	4	5	6	7	8	9	10
static object	6	6	5	6	5	6	5	6	5	6
non-static object	0	0	0	0	0	0	0	0	0	0

because, as long as the group manages to exert a sufficient force to move the object, the smaller linear momentum due to the object’s lighter mass makes the detachment events less disruptive for the group performance. In other words, with a progressively lighter object, even if all the group members are required to initiate the transport, not all the robots are required to push a moving object to sustain the transport. Therefore, in this condition, detachment events are less disruptive with respect to the group performance.

Although, the results shown in Fig. 2a indicate that our simulation environment is a sufficiently accurate model of the “reality”, modifications can be certainly made to improve the robustness of group strategies. For example, the performances of groups of 6 simulated robots transporting 30 cm length, and 980 g object drop with respect to the performances of simulated groups in the other experimental conditions (see Fig. 2a, white bars). However, this performance drop is less evident for other groups of real e-pucks. Moreover, this is the evaluation condition in which we observe the largest difference between the performances of real and simulated robots. This suggests that our simulation does not accurately capture the effect of the \bar{L} smaller than \bar{D} condition. If we could model the effects of this phenomenon, we could perhaps improve the performances and robustness of the transport strategies.

The result of the first set of evaluations tell us that we succeeded in designing a controller to allow a swarm of real e-pucks to effectively transport heavy objects in a cooperative way. Performances are scalable and robust to deal with varying operating conditions. The results also demonstrate that group coordination of actions and alignment of pushing forces can be reached with a simple sensory apparatus made of distance sensors and the optic-flow sensor to indirectly perceive the object movement. The cylindrical shape of the robots negatively impacts the group performance when the length of the object is shorter than the sum of the robots radius. This negative effect tends to disappear when the transport can be sustained by less robots than those required to initially move the object.

5.2 Behavioural analysis

How do the robots manage to coordinate their actions to cooperatively transport the object? To answer this question, we describe the results of a further series of evaluation tests on a single real robot. In these tests, the robot undertakes multiple trials where it is required to push an object with varying characteristics (e.g., a light, a heavy, and a long object). During these tests, we record the number of

repositioning events. That is the number of times the robot changes the point of exerting forces on the object. A repositioning event happens anytime the robot stops pushing the object and immediately after starts pushing the object again in a slightly different position. In the biological literature, repositioning events are considered to be direct evidence of “persistence”: that is, the individual tendency to persevere with a given behavioural strategy. As discussed in [8], persistence is an individual-level parameter that modulates transport efficiency.

Our objective is to use the concept of persistence as a tool to move a step forward in the understanding of the operational mechanisms underlying the alignment of forces required for group transport. In particular, we are looking for relationships between characteristics of the object (i.e., its mass, and its direction of movement with respect to the robot heading) and persistence.

In the first series of tests, a real robot is positioned in front of a cuboid object, facing the object at 20 cm from it. In each trial, the robot is given 60 s to push the object. All tests are repeated for 10 trials. In test *L*, the object length is set to 30 cm, and the object mass is 150 g. The robot can easily transport the object. In test *H*, the object length is set to 30 cm, and the object mass is 600 g. The object is too heavy to be moved by the robot. In test *R*, the object length is set to 40 cm, and the object mass is 400 g. The robot can rotate the object by exerting pushing forces on either end of the longest side, but it can not transport it. Fig. 2b shows the number of repositioning events counted during each trial of each test. The results clearly show that the number of repositioning events change with respect to whether or not the object can be moved or simply rotated. When the object is so heavy that it can not be moved or rotated by the robot, we observe a very high number of repositioning events. This indicates that the agent persistence is low (see Fig. 2b, box H). When the object is light enough to be moved or to be rotated by the robot, we observe a very low number of repositioning events. This indicates that the agent persistence is very high (see Fig. 2b, box L, and box R). We conclude that, the agent perception of the object linear or rotational movement, through the optic-flow sensor, increases the agent persistence. In other words, the robot keeps on looking for new points on which to exert pushing forces if the object does not move. The robot does not change the point of contact with the object if the object moves while it is pushing it.

We also run a further test in which we looked at relationship between persistence and object movement. In this test, the object length is set to 30 cm, its mass to 600 g. This object is too heavy to be moved by the robot. We run 10 trials without interfering with the robot actions, and 10 trials in which we intentionally moved the object in the opposite direction of the robot heading while the robot is pushing the object. We refer to the trials with no experimenter interference as static object trials, and the trials with the intervention of the experimenter as non-static object trials. In each of the static and non-static trials, we counted the repositioning events with pushing forces exerted on the first touched long side of the object. We stopped counting as soon as the robot touches the other long side of the cuboid object. The aim of this test is to estimate how long it takes (in terms of repositioning events) the robot to invert the

direction of its pushing forces when the object does not move (static), and when the object moves against its heading. The results of this test, shown in Table 1, clearly indicate that no repositioning events are observed when the robot perceives the object moving against its heading. In other words, the robot quickly changes direction of pushing forces if it perceives the object moving against its heading. The response of the robot is to move away from the object with a circular trajectory that rather quickly takes it to the opposite side of the object. As shown in Table 1 for the static object, when no object movement is perceived, the robot keeps on looking for new points on which to exert pushing forces on the same side of the object.

By visually inspecting the robots' strategies during group transport, keeping in mind the results of our single robot evaluation tests, we noticed that robots heavily rely on the perception of the object (rotational) movement as a mean to align their forces. Robots exerting forces on the direction of the object rotation tend to have high persistence, while the robots exerting forces on the opposite direction of the object rotational movement tend to swap sides. When all the robots are on a single side, the force exerted on the object causes the object to switch from rotational to translation movement, and the transport begins. In the absence of rotational movements (e.g., with very heavy objects), the alignment certainly becomes more difficult, and it definitely takes longer for the robots to coordinate their efforts. Nevertheless, the robots eventually manage to position themselves on the same side of the object and to exert the required forces to move it. In these circumstances, we think that alignment is favoured by correlation between robot-robot interactions and individual persistence. However, further investigations are required to better understand this process.

6 Conclusions

As shown in [8], cooperative transport in ants is a poorly understood process, with not much empirical evidence to shed light on the proximate mechanisms underpinning this important cooperative process. The results of this study suggest that a rather parsimonious explanation based on the perception of the object movement could account for the alignment and the coordination of forces observed in natural organisms. We showed that groups of robots capable of perceiving whether their pushing forces are aligned with the object movement can use this cue to arrange themselves on the same side of an elongated cuboid object to exert a sufficient force to transport objects that are too heavy to be moved by single robots. We also provided a basic description of how this sensory information influences individual strategies. In particular, we showed that the object movement correlates with the agent persistence. In other words, agents that perceive their pushing forces aligned with the object movement tend to keep on exerting forces on the same point on the object. Perception of no object movement induces the robot to change the point of contact with the object. Perception of object movement against the robot heading induces the robot to exert forces on the opposite side of the object. We also showed that the robot con-

trollers synthesised using evolutionary computation techniques generate group strategies that are robust to deal with some variability in object length and mass, as well as scalable to successfully operate with groups of different cardinality. Future work will focus on the transport of objects with irregular shape and no symmetries, as well as on the analysis of the effects of robot-robot interactions on the development of successful group strategies.

Acknowledgements M.H. Mohammed Alkilabi thanks Iraqi Ministry of Higher Education and Scientific Research for funding his PhD, P. Todd and D. Lewis for their help and support in modifying e-puck robot.

References

1. Alkilabi, M., Lu, C., Tuci, E.: Cooperative object transport using evolutionary swarm robotics methods. In: Proc. of the European Conf. on Artificial Life. vol. 1, pp. 464–471. MIT (2015)
2. Beer, R., Gallagher, J.: Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior* 1(1), 91–122 (1992)
3. Berman, S., Lindsey, Q., Sakar, M., Kumar, V., Pratt, S.: Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems. *Proc. of the IEEE* 99(9), 1470–1481 (2011)
4. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA (1989)
5. Groß, R., Dorigo, M.: Cooperative transport of objects of different shapes and sizes. In: Dorigo, M., Birattari, M., Blum, C., Gambardella, L.M., Mondada, F., Stützle, T. (eds.) *Proc. of the 4th Int. Workshop on Ant Colony Optimization and Swarm Intelligence*. LNCS, vol. 3172, pp. 106–117. Springer (2004)
6. Groß, R., Dorigo, M.: Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling. *Adaptive Behavior* 16(5), 285–305 (2008)
7. Kube, C., Bonabeau, E.: Cooperative transport by ants and robots. *Robotics and Autonomous Systems* 30, 85–101 (2000)
8. McCreery, H., Breed, M.: Cooperative transport in ants: a review of proximate mechanisms. *Insects Sociaux* 61, 99–110 (2014)
9. Moffett, M.: Ant foraging. *Research and Exploration* 8, 220–231 (1992)
10. Mondada, F., et al.: The e-puck, a robot designed for education in engineering. In: *Proc. of the 9th Int. Conf. on Autonomous Robot Systems and Competitions*. vol. 1, pp. 59–65 (2009)
11. Robson, S., Traniello, J.: Resource assessment, recruitment behavior, and organization of cooperative prey retrieval in the ant *Formica schaufussi* (hymenoptera: Formicidae). *Insect Behavior* 11, 1–22 (1998)
12. Wang, Z., Takano, Y., Hirata, Y., Kosuge, K.: A pushing leader based decentralized control method for cooperative object transportation. In: *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*. vol. 1, pp. 1035–1040. IEEE (2004)
13. Yamamoto, A., Ishihara, S., Fuminori, I.: Fragmentation or transportation: mode of large-prey retrieval in arboreal and ground nesting ants. *Insect Behavior* 22, 1–11 (2009)