

Cooperative object transport with a swarm of e-puck robots: robustness and scalability of evolved collective strategies

Muhanad H. Mohammed Alkilabi^{1,3} · Aparajit Narayan¹ · Elio Tuci² 

Received: 20 October 2016 / Accepted: 22 March 2017 / Published online: 31 March 2017
© Springer Science+Business Media New York 2017

Abstract Cooperative object transport in distributed multi-robot systems requires the coordination and synchronisation of pushing/pulling forces by a group of autonomous robots in order to transport items that cannot be transported by a single agent. The results of this study show that fairly robust and scalable collective transport strategies can be generated by robots equipped with a relatively simple sensory apparatus (i.e. no force sensors and no devices for direct communication). In the experiments described in this paper, homogeneous groups of physical e-puck robots are required to coordinate and synchronise their actions in order to transport a heavy rectangular cuboid object as far as possible from its starting position to an arbitrary direction. The robots are controlled by dynamic neural networks synthesised using evolutionary computation techniques. The best evolved controller demonstrates an effective group transport strategy that is robust to variability in the physical characteristics of the object (i.e. object mass and size of the longest object's side) and scalable to different group sizes. To run these experiments, we designed, built, and mounted on the robots a new sensor that returns the agents' displacement on a 2D plane. The study shows that the feedback generated by the robots' sensors relative to the object's movement is sufficient to allow the robots to coordinate their efforts and to sustain the transports for an extended period of time. By extensively analysing successful behavioural strategies, we illustrate the nature of the operational mechanisms underpinning the coordination and synchronisation of actions during group transport.

✉ Elio Tuci
e.tuci@mdx.ac.uk

Muhanad H. Mohammed Alkilabi
mhml@aber.ac.uk

Aparajit Narayan
apn3@aber.ac.uk

¹ Aberystwyth University, Aberystwyth, UK

² Computer Science Department, Middlesex University, London, UK

³ Computer Science Department, Kerbala University, Kerbala, Iraq

Keywords Collective transport · Swarm robotics · Evolutionary computation · Artificial neural networks

1 Introduction

Cooperative object transport in multi-robot systems is a process requiring coordination and synchronisation of pushing/pulling forces by a group of autonomous robots in order to transport items that can not be transported by a single agent (Groß and Dorigo 2004a). In natural systems, ants have evolved extremely effective competencies to cooperatively retrieving items that can be hundreds or even thousands times the weight an individual can carry (Czaczkes et al. 2011). Cooperative transport is relatively ubiquitous in ants, being known in at least 40 genera of ants (Moffett 1992). Ants primarily engage in cooperative transport to retrieve objects (e.g. food items) that are too heavy or too large to be moved by a single individual. Owing to cooperative transport, ants can perform faster prey retrieval reducing both the exposition of foragers to predators, and the risk of food being caught and eaten by other aggressive species (Hölldobler et al. 1978; Yamamoto et al. 2009). The fast retrieving of preys also reduces the time workers are involved in transport tasks, freeing them for other colony relevant tasks (Feener et al. 1990; Tanner 2008). Czaczkes and Ratnieks (2013) show that cooperative transport also reduces the energetic cost of transport by allowing carriers to keep up with the dense flow of traffic and by reducing the possibility of traffic jams.

As shown by McCreery and Breed (2014), not all species of ants are efficient in collectively transporting large or heavy items. The alignment of agents' travel directions that triggers and sustains the transport is a complex process that not all ants species manage to execute efficiently. In some species, workers simultaneously push and pull the object in opposite directions, resulting in transport processes that are generally slow and interrupted by frequent “deadlock” in which the workers cancel each other's forces. In other species, the transport is extremely efficient, fast, and without deadlocks (Czaczkes and Ratnieks 2013). 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 active engagement of individuals into the transport (McCreery and Breed 2014). It has also been observed that, for those ants species in which the transport is very fast and efficient, some individuals seem to be more important than others in sustaining and directing the movement during the collective effort (Gelblum et al. 2015).

In spite of the numerous studies focusing on this process, various hypotheses concerning the mechanisms for alignment and coordination of forces during cooperative transport in ants remain to be empirically verified. For example, it is still not clear what mechanisms are used to assess consensus or quorum information about directional movements (McCreery and Breed 2014). 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 (Robson and Traniello 1998).

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 (Bonabeau et al. 1999). Like for ants, swarm robotics systems are required to operate and to coordinate their actions without using centralised control or global information, and without any form of global communication, since this is very likely to impose restrictions on the scalability of solutions for very large swarm size (i.e. hundreds or thousands of individuals) (Dorigo and Şahin

2004). Detailed and comprehensive reviews of the state of the art in swarm robotics can be found in [Brambilla et al. \(2013\)](#); [Bayindir \(2016\)](#). Cooperative transport is one of the most studied phenomenon in swarm robotics, since it requires collective competencies that can make swam robotic systems extremely effective in a variety of real-world applications, such as waste retrieval and disposal, de-mining, or operations requiring object manipulation in environments where a direct human intervention is impossible or impractical, such as in space or in deep sea ([Woern et al. 2006](#); [Huntsberger et al. 2000](#); [Parker and Zhang 2006](#)).

Generally speaking, the objective of the studies in swarm robotics on cooperative transport is to provide new engineering solutions to improve the effectiveness of the collective responses (e.g. see [Habibi et al. 2014, 2015](#)). Some of these research works are also relevant for biology since they represent a proof-of-concept demonstration on physical hardware of the kind of mechanisms natural swarms could potentially use to coordinate their actions during transport. In this study, we follow this multidisciplinary stance by engineering a swarm robotics system that transports objects of various sizes and masses, using strategies that proved to be robust and scalable to larger group sizes. At the same time, our work demonstrates how a group of robots equipped with a minimalist sensory apparatus, and with no means of direct communication, can effectively accomplish a collective transport task. Throughout the paper, the adjective “minimalist” referred to the robots sensory apparatus is used to indicate, in qualitative terms, robots that lack the capability to sense pushing/pulling forces and to directly communicate to any of their swarm mates. Our results point to a rather parsimonious explanation of the mechanisms required by real ants to transport an object. In particular, we show that the indirect perception of the movement of the object to be transported modulates the frequency with which a robot changes the point of application of its pushing forces. The perception of movements of the object reinforces pushing behaviour on the same robot–object or robot–robot contact point. The perception of no object movement induces the robot to change the point of application of its pushing forces. This mechanism 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.

In the next section, we briefly review some of those swarm robotics studies whose results, like in our case, suggest that complex forms of social behaviour can be accomplished with less than what originally thought to be necessary (see Sect. 2). Sections 3, 4 and 5 describe the task and the simulation model, the robots’ neuro-controller and the algorithm used to set its parameters, and the fitness function. Section 6 illustrates the experimental results and the analysis of the operational mechanisms underpinning the single robot’s behaviour. In Sect. 7, we discuss and comment on the results of this study and we point to interesting directions for future work.

2 Background

As acknowledged in [Groß and Dorigo \(2009\)](#), sensitivity to size and shape of the object, and undesired negative effects during tests on scalability with respect to group cardinality, seem to be major obstacles to the design of efficient group transport strategies in swarm robotics systems. Various hardware solutions and different types of control policies have been tested to improve the performance of swarms of robots designed to accomplish collective transport tasks. [Chen et al. \(2015\)](#) review this rather heterogeneous body of literature according to the type of strategy used by the robots to transport the object, effectively dividing the systems in three groups: those in which the robots push the object, those in which the robots pull the object, and those in which the robots use a caging strategy (i.e. the robots first surround the

object and then move it). After a careful analysis of advantages and disadvantages of each approach, the authors describe an alternative group transport method which, rather than trying to overcome the limitations imposed by occlusion of the goal, it exploits occlusion. The robots are designed to push the object across the portion of its surface, where it occludes the direct line of sight to the goal. The authors also provide an analytical proof of the effectiveness of the method and show the results of successful empirical tests in physical robots with objects of different shapes (see [Chen et al. 2015](#)).

In this section, our objective is to briefly review the literature on cooperative transport in swarm robotics by mainly focusing on studies whose results have also a relevance for biology. Our goal is to provide the reader with a brief overview of works whose aim is twofold: (i) to illustrate new structural/functional elements that prove to be effective engineering solutions for swarm robotics systems engaged in object transport scenarios and (ii) to generate new hypotheses about the mechanisms that may underpin the process of collective transport in natural swarms. Similarly to these studies, we describe a swarm robotics system that proves to be particularly effective in transporting, using pushing strategies, objects of various masses and sizes. We also generate a specific hypothesis on the characteristics of the sensory apparatus required by natural organisms to align pushing forces and to sustain the transport of heavy objects. The distinctive feature of this work and our contribution to the literature are to show that force sensors are not required to initiate and sustain the collective transport.

The starting point of the kind of research studies mentioned above can be identified in the pioneering work of [Kube and Zhang \(1997\)](#) on box pushing by a multi-robot system. This study is considered the first research work that formally represented in “hardware” the dynamics of collective object transport. In this study, the effectiveness of the individual mechanisms underpinning cooperative transport is tested with respect to their sensitivity to the group size. The authors demonstrate that coordinated efforts in the box pushing task are possible without use of direct communication or robot differentiation. The work described in [Kube and Bonabeau \(2000\)](#) further develops the model described in [Kube and Zhang \(1997\)](#) with the addition of a stagnation recovery strategy. Stagnation refers to a deadlock condition in which robots cancel each other pushing forces due to the way in which they are distributed around the object. The authors provide the robots with a realignment or repositioning strategy to allow the agents to overcome stagnation by redistributing the pushing forces around the object. In this study, the authors also evaluate the group transport strategies with objects of different shapes in scenarios in which the objects have to be transported toward a moving target.

The study described in [Berman et al. \(2011\)](#) tries to mimic the behaviour of natural swarms by looking for the individual rules that generate robust group-level responses. The authors observed a particular species of ants *Aphaenogaster cockerelli* in order to extract and reproduce in a simulated swarm robotics system those rules that govern the individual actions during group transport. The authors created a detailed model based on a qualitative analysis of the role and contribution of single ants during transport of food items to the nest. The collected data have 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. [Wang et al. \(2004\)](#) propose a decentralised control algorithm to control a swarm of robots in which the cooperative transport is coordinated by a leader robot that knows the direction of transport. The algorithm is directly implemented and tested on a robotic system in which the robot can only push the object. The authors show that the follower robots can interact with the leader robot by simply sensing the forces/movements on the object. This form of indirect communication through the object is sufficient to allow the robots to exert forces and to move the object along the transport trajectory known by the leader. Similarly, [Wang and](#)

[Schwager \(2016\)](#) develop a mechanical model for leader/follower multi-robot systems. The followers calculate the force and the velocity needed to exert forces on the object in order to reinforce the leader's efforts. Although the robots do not require any explicit communication mechanism, the model requires the robots to know the mass and the friction coefficients of the object beforehand in order to measure the velocity and acceleration at the centre of mass of the object. The swarm robotics model described in [Groß and Dorigo \(2008\)](#) demonstrates that communication between robots involved into the collective transport need not 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 robotics model targeting cooperative transport for simple robots that can only exert pushing forces. In this work, a group of physical e-puck robots are required to push an elongated cuboid object which is heavy enough to require the combined efforts of all the members of the group to be transported. The robots have to coordinate and align their movements in order to agree on a common direction of transport and to push the object for an extended period of time. As mentioned above, the distinctive feature of our model is the minimalist sensory apparatus provided to the robots. Contrary to similar previous studies (see for example [Aiyama et al. 1999](#); [Groß and Dorigo 2004b](#); [Groß et al. 2006a, b](#); [Nouyan et al. 2006](#); [Pettinaro et al. 2005](#); [Wang et al. 2006, 2004](#); [Wang and Schwager 2015, 2016](#)), our robots have no means to directly perceive and measure forces applied either directly or indirectly on the object or to the robots themselves. We use robots equipped with a sensory apparatus made of proximity sensors, a camera, and an “optic-flow” sensor, appositely designed, built, and mounted on the e-puck chassis to allow each robot to get a precise estimate of its movement on a 2D plane. This information, in combination with the readings of the distance sensors, generates a sensory stimulation that effectively informs the robot on the direction of movement of the object.

The optic-flow sensor is a relatively cheap hardware component that can be easily integrated in various robotic platforms. It is also extremely robust and capable of generating highly precise reading in operating conditions in which the floor is relatively flat. The results of our study unambiguously demonstrate that the feedback generated by the optic-flow sensor is sufficient to allow a swarm of robots to align their pushing forces, to agree on a common direction of transport, and to sustain the transport of heavy objects for an extended period of time. In view of these promising results, it is important to clarify that, from an engineering perspective, we do not claim that alternative solutions to those based on force sensors have to be always privileged in swarm robotics. The specific conditions in which a swarm is required to operate, the physical structure of the robots, as well as other contingent phenomena, are elements that have to be carefully taken into account to make important methodological choices whose significance has to be empirically tested. In this spirit, we believe that our results, generated with this minimalist set-up, can be used in the future as a term of comparison to verify whether the use of whatever more complex or simply different sensory apparatus returns any substantial benefit in term of group performance.

This paper is based on and extends a preliminary study described in [Alkilabi et al. \(2016b\)](#). We complement the initial evaluation of the system already illustrated in [Alkilabi et al. \(2016b\)](#) with an extensive series of further analyses focused on the description of the individual behavioural strategies used by the robots to coordinate and to synchronise their efforts. We also evaluate the scalability of the group transport strategies with respect to the cardinality of the group, and their robustness to a larger set of operating conditions, in which we vary mass and size of the object. In this study, we explicitly avoid to directly or indirectly impose to the robots a direction of transport, because we are interested in the emergence of alignment of pushing forces by means of local interactions. In a further study described in [Alkilabi](#)


et al. (2016a), we extended the e-puck competencies by designing mechanisms that allow e-puck robots first to move a heavy object in an arbitrary direction and subsequently to push it toward a specific target location.

3 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. The parameters of the neuro-controllers are set in a simulation environment which models kinematic and dynamic features of the experimental conditions in which simulated e-pucks are required to operate. The robot's sensory apparatus includes infrared sensors, a camera, and the optic-flow sensor appositely designed, built, and integrated into the e-puck structure for this task (see Fig. 1a). During evolution, the robots are initially positioned in a boundless arena with flat terrain, at 50 cm from the object. The robots starting positions correspond to randomly chosen points on a circle's circumference of 50 cm radius that has the object in its centre (see Fig. 1b). This circle is divided in four equal 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. 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.

Using a trial and error procedure, we found that for a group of three e-puck robots, all robots are required to initiate (i.e. to change the zero linear momentum of the cuboid object) and sustain (i.e. to continue moving the cuboid object when it has already a linear momentum) the transport when the mass to be transported is within the interval [450 g, 530 g]. The transport of a cuboid object of mass lower than 450 g requires less than three robots to be sustained. The transport of a cuboid object of mass higher than 530 g requires more than three robots to be initiated. For example, the transport of a cuboid object of mass within the interval [530 g, 600 g] cannot be initiated, although it can be sustained, by a three robots groups. The transport of a cuboid object of mass higher than 600 g cannot be initiated nor sustained by a three robots groups. For a group of four e-puck robots, all robots are required to initiate and sustain the transport when the mass to be transported is within the interval [600 g, 680 g]. The transport of a cuboid object of mass lower than 600 g requires less than four robots to be sustained. The transport of a cuboid object of mass higher than 680 g requires more than four robots to be initiated. Based on these measurements, we worked out that for groups of cardinality N , the lowest (L^m) and the highest (H^m) mass requiring the contribution of all N robots to both initiating and sustaining the transport can be expressed using the following formulas: $L^m = N \times 150$; $H^m = (N \times 150) + 80$. We have empirically verified with tests on physical robots the reliability of the above-mentioned formulas. This approach clearly generates only a rough estimation of the lowest and highest object's masses that can be transported by groups of cardinality N . The reason for this being the influence of multiple factors that cannot be systematically controlled. For example, small variability of the terrain surface, dust, wheels' slippage, the distribution of the points of application of pushing forces can considerably contribute to determine the mass of the object whose transport requires all N robots of the group to be both initiated and sustained.

The robots can perceive the object with their camera, and when sufficiently close to it, they can sense it with their infrared sensors. The task requires the robots to independently



(a)



(b)

Fig. 1 a E-puck robot with details of the optic-flow sensor. b The robots starting position during the evolutionary phase. The *black circles* indicate the robots and the *grey rectangle* refers to the object

search for the object and move toward it. Once in the proximity of the object, the robots have to coordinate their actions in order to push the object by exerting the forces required to transport it as far as possible from its initial position.

To take into account the dynamic aspects of this group transport scenario (e.g. forces, torque, friction), the agents and their environment have been simulated using the *Bullet* physics engine. The object has a cuboid shape (30 cm length, 6 cm width, 6 cm height) with a mass of 600 g. As mentioned above, our simulation models an e-puck robot (Mondada et al. 2009). The robot model consists of three rigid bodies, a cylindrical chassis (3.55 cm radius, 6.2 cm height, 200 g mass), and two motorised cylindrical wheels (2.05 cm radius, 0.2 cm height, 20 g mass) connected to the chassis with hinge joints. Both wheels can rotate forward and backward at a maximum speed of 8 cm/s (see also Alkilabi et al. 2015 for a detailed description of the simulator).

Every robot is provided with eight infrared sensors (IR^i with $i \in \{0, \dots, 7\}$), which give the robot a noisy and nonlinear indication of the proximity of an obstacle (e.g. the object or

another robot). The IR sensor values are computed using a nonlinear regression model of the sensor readings collected from the physical e-puck. Once the values are computed, they are subject to noise: IR sensors perceiving objects are subject to noise drawn from a random uniform distribution in the range $\pm 25\%$ of the sensor maximum readings. IR sensors not perceiving any objects are subject to noise drawn from a random uniform distribution in the range $\pm 7\%$ of the sensor maximum readings. Each robot is also equipped with a camera that can perceive coloured items (i.e. the object which is green, or robots which are all red). The camera has a receptive field of 30° , divided in three equal sectors C_i , with $i = \{1, 2, 3\}$, each of which can return one of four possible values: 0 if no item falls within the sector's field of view; 0.4 if one or more red items are perceived; 0.7 if a green item is perceived; 1.0 if red and green items are perceived. The camera can detect coloured objects up to a distance of 50 cm.

The new optic-flow sensor 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. 18×18 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 movement of the robot. This information is subsequently communicated to the robot controller in the form of four normalised 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. The optic-flow sensor returns readings at a maximum speed of 12 ips (inches per second) and can operate on different type of flat surfaces. To improve portability of solutions to physical hardware, in simulation, $+X$, $-X$, $+Y$, and $-Y$ are subject to uniformly distributed random noise in $[-0.025, 0.025]$. A detailed description of the optic-flow sensor's characteristics can be found at <http://www.aber.ac.uk/en/cs/research/ir/dss/#swarm-robotics>.

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 differentiate between the former and the latter condition and to respond accordingly. The results of this study show that the extended sensory apparatus of the e-puck robots generates readings that are sufficiently informative to allow the robots to coordinate their effort in order to collectively transport in an arbitrary direction an object that cannot be moved by single robots.

4 The controller and the evolutionary algorithm

The robot controller is composed of a continuous time recurrent neural network (CTRNN) of 19 sensor neurons, 6 internal neurons, and 4 motor neurons (see Beer and Gallagher 1992, and also Fig. 2 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 Eqs. 1, 2, and 3.

$$y_i = gI_i; \quad i \in \{1, \dots, N\}; \quad \text{with } N = 19; \quad (1)$$

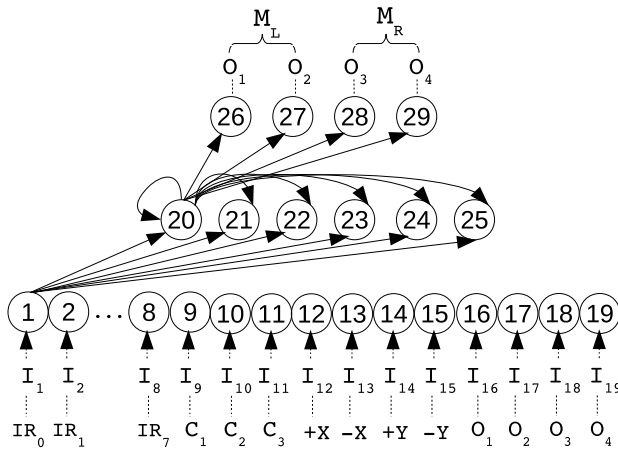


Fig. 2 The robot’s controller. The *continuous line arrows* indicate the efferent connections for only one neuron of each layer. Hidden neurons receive an afferent connection from each input neuron and from each hidden neuron, including a self-connection. Output neurons receive an afferent connection from each hidden neuron. Sensors to sensor neurons correspondence is indicated underneath the input layer, with IR_i referring to the infrared sensors, C_i to the camera sensors, $+X, -X, +Y, -Y$ to the readings of the optical flow sensor, and O_i referring to the output of the network at previous time step

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

$$y_i = \sum_{j=N+1}^{j=N+6} \omega_{ji} f_j; \quad i \in \{N + 7, \dots, N + 10\}; \tag{3}$$

with $f_j = \sigma(y_j + \beta_j)$; and $\sigma(x) = (1 + e^{-x})^{-1}$. 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 $i \in \{1, \dots, N\}$ is the activation of the i th sensor neuron (see Fig. 2 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, f_j the firing rate. 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 Eq. 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 (see Goldberg 1989). 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. Each genotype is a vector comprising real values coding for the network’s connection weights, decay constants, bias terms and gain factor. Initially, a random population of vectors is generated by initialising each component of each genotype to values chosen

uniformly random from the range $[0, 1]$. New genotypes, except the elite, are produced by applying recombination and mutation. Each new genotype has a 0.3 probability of being created by combining the genetic material of two parents. During recombination, one crossover point is selected. Genes from the beginning of the genotype to the crossover point are copied from one parent, and the other genes are copied from the second parent. Mutation entails that a random Gaussian offset is applied to each real-valued vector component encoded in the genotype, with a probability of 0.04. The mean of the Gaussian is 0, and its standard deviation is 0.1.

During evolution, all vector component values are constrained to remain within the range $[0, 1]$. Genetically encoded values were linearly mapped into CTRNN parameters with the following ranges: input neuron biases $\beta_I \in [-4, 4]$; hidden neuron biases $\beta_j \in [-5, 5]$; output neuron biases $\beta_O \in [-5, 5]$; connection weights from input to hidden neurons $\omega_{ji} \in [-8, 8]$; connection weights from hidden to hidden and from hidden to output neurons $\omega_{ji} \in [-10, 10]$; and gain factor $g \in [1, 13]$. Decay constants were first linearly coded in the range $\tau_i \in [0, 2.2]$ and then exponentially mapped into $\tau_i \in [10^0, 10^{2.2}]$. These parameter ranges were chosen on the basis of having proven useful in other CTRNN experiments. The large range of the exponentially mapped decay constants is meant to allow for evolution to select both neurons that tend to change their state (cell potential) radically every time step (i.e. neurons with small decay constants), and neurons that tend to change their state only minimally every time step (i.e. those with large decay constants).

5 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 m away from its initial position. At the beginning of each trial, the controllers are reset, and the robots are positioned in the arena as shown in Fig. 1b. 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 positions contribute 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.

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:

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

$$f_1 = \sum_{r=1}^R (1 - d_r); \quad f_2 = \frac{\Delta O^{pos}}{2}; \quad f_3 = t/T; \quad \text{with } T = 900; R = 4; \quad (5)$$

d_r is the normalised 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 at the end of the trial. 2 m is the maximum distance an object can be moved from its initial position. t is the trial duration in simulation steps. f_1 rewards groups in which the robots approach the object and remain close to it. f_2 rewards groups that transport the object as far as possible.

f_3 rewards groups for minimising the time required to move the object 2 metres away from its initial position. The fitness of a genotype (\bar{F}) is the average team evaluation score after it has been assessed $E = 12$ times: $\bar{F} = \frac{1}{E} \sum_{e=1}^E F_e$.

6 Results

The primary aim of this study is to design control systems for homogeneous groups of physical 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 robots' controller, we run 20 differently seeded evolutionary simulations, each simulation lasting 3000 generations. In order to choose the controller to be ported onto the physical robots, we re-evaluated, in simulation, the best genotypes from generation 1000 to generation 3000 for every run. During re-evaluations, we evaluated homogeneous groups of 3, 4, 5, and 6 simulated robots, for their capability to collectively transport rectangular cuboid objects of 30 and 40 cm lengths. Objects of each length are tested with two different masses. Object width and height are not changed with respect to evolutionary conditions. The total number of re-evaluation trials per group (i.e. 320) is given by all the possible combinations of the above-mentioned parameters (i.e. two object lengths, two object masses, four different values for group cardinality), with each combination repeated for 20 trials (i.e. five trials for each of the starting position shown in Fig. 3a). 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 the object is always heavy enough to require the combined effort of all robots of the group to be successfully transported. In each evaluation trial, the object is placed in the centre of a boundless flat 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 2 m away from its initial position. Only in this later case, the trial is considered successful.

The results of this re-evaluation test are shown in Fig. 3b, where the bars indicate the success rate (%) of the best evolved homogeneous group of each evolutionary run. In this graph, the runs are sorted from the best to the worst. The graph indicates that seven runs managed to generate a best evolved group with a success rate higher than 80%. Twelve out of 20 runs generated a best evolved group with a success rate higher than 70%. The very best group had a success rate higher than 90%. 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 physical e-pucks for further evaluations. In the next section, we describe the results of the first test with physical robots, and we compare these results with those of simulated robots controlled by the same controller and evaluated in similar operational conditions.

6.1 First evaluation test with physical e-pucks

The first evaluation test with physical 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 four robots and only one type of elongated cuboid object (30 cm length, 6 cm width and height, 600 g mass). During the test with physical robots, we evaluated homogeneous groups of 3, 4, 5, and 6 physical e-pucks, for their capability to collectively transport cuboid objects of 30 and 40 cm length.

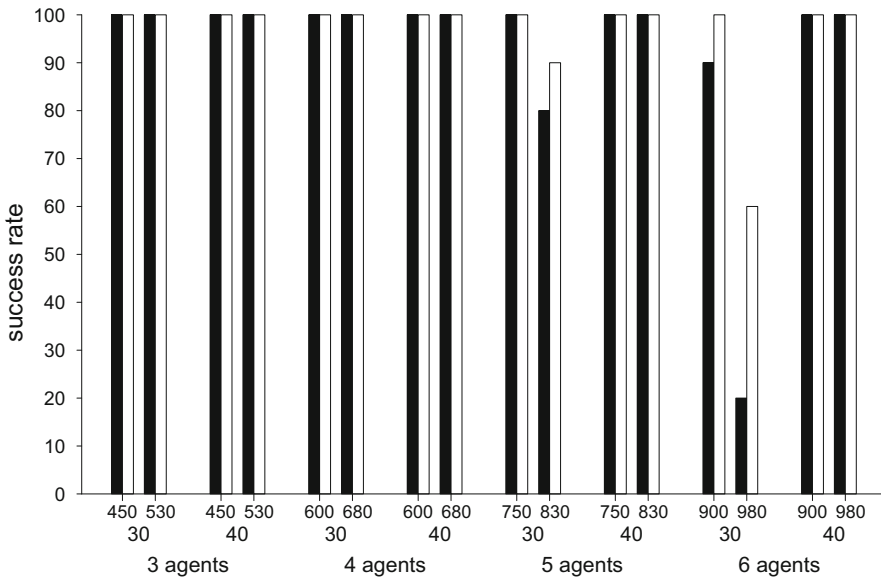


Fig. 4 Graph showing the success rate (%) of homogeneous groups controlled by the best evolved controller. *Black bars* refer to the performance of groups of physical e-pucks; *white bars* refer to the performances of groups of simulated e-pucks. The *x*-axis shows mass (in grams, first row) and length (in centimetres, second row) of the objects, and group cardinality. *Each bar* refers to the average performance on a set of ten trials

the starting positions 1 and 2 shown in Fig. 3a). In each evaluation trial, the object is placed in the centre of a bounded square arena (220 cm side length), and the robots are placed at about 50 cm from the object. Each evaluation trial can last up to 180 s. A trial is terminated earlier if the group is successful (i.e. if the object is transported 1 m away from its initial position). A *Vicon* tracking system is used to track the object movements in a 2D coordinates plane.

The results of the first evaluation test are shown in Fig. 4, where the bars indicate the success rate (%) of homogeneous groups controlled by the best evolved controller in 16 different evaluation conditions. Black bars refer to the performance of groups of physical 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 physical and simulated robots is meant to capture differences in performance when moving from simulation to reality. We notice that the performances of both physical and simulated robots are 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.¹

The results shown in Fig. 4 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 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

¹ Data and graphs not shown in the papers, as well as movies of physical and simulated robots can be found at <http://www.aber.ac.uk/en/cs/research/ir/dss/#swarm-robotics>.

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. In case of a detachment event, a robot 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. 4, the \bar{L} smaller than \bar{D} condition only minimally affects the performance of groups of 6 physical e-pucks transporting a slightly lighter object (see Fig. 4, six robots, 30 cm length, 900 g object, black bar). This is 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. 4 indicate that our simulation environment proved to be an effective methodological tool for the design of robust and reliable robot controllers, modifications to the characteristics of the robot–world model can be certainly made to improve the robustness of the evolved group strategies. For example, in the condition with 6 agents pushing a 30 cm length and 980 g we observe the largest difference between the performances of physical and simulated robots. In this condition, the longest side of the object is too short to allow all robots to directly push the object (i.e. small $(L:D)$ ratio). Thus, in order to initiate and sustain the transport, some robots have to push other robots that are in contact with the object. As discussed above, when physical robots push other robots we observe a relatively high number of detachment events, certainly higher than those that we observe in similar circumstances with simulated robots. Since detachment events tend to disrupt the group performance, physical robots perform worst than simulated one whenever the $(L:D)$ ratio is small. The results of this comparative test seem to suggest that our simulator tends to “simplify” the dynamics corresponding to the circumstance in which one robot pushes another robot that pushes the object, so that potentially disruptive phenomena that tend to occur with physical robots are not observed in simulation. We believe that if we could more accurately model the dynamics related to the configuration in which one robot pushes another robot that pushes the object we could design controllers that can more effectively cope with these dynamics, resulting in improved group performances and more robust and more scalable transport strategies.

In summary, the result of the first set of evaluations tell us that we succeeded in designing a controller to allow a swarm of physical 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 (see also Sect. 6.3). The cylindrical shape of the robots negatively impacts on the group performance when the length of the object is shorter than the sum of the robots’ diameter. This negative effect tends to disappear when the transport can be sustained by less robots than those required to initially move the object.




Fig. 5 Graph showing the sinuosity of the trajectories of the object when transported by physical robots during successful trials. The experimental conditions given by mass (in grams, first row) and length of the objects (in centimetres, second row), and group cardinality are indicated on the x-axis. Each point in the box refers to the group performance in a single trial. The number of successful trials per experimental condition is indicated above each box. Grey boxes refer to a condition in which there is a significant difference between the sinuosity recorded in trials with the light and heavy object (Wilcoxon rank sum test, $p < 0.05$). White boxes refer to conditions in which this difference is not significant. Boxes represent the inter-quartile range of the data, while dashed horizontal bars inside the boxes mark the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box

6.2 Behavioural analysis of physical robots’ groups using the sinuosity metric

In this section, we discuss the performance of physical robot groups with an analysis of the trajectories of the object in each of the 16 different evaluation conditions described above. This analysis is based on a metric referred to as sinuosity (S), which is defined as follows:

$$S = \frac{P}{\Delta O^{pos}} \tag{6}$$

where P corresponds to the object’s path length during transport, and ΔO^{pos} is the Euclidean distance between the position of the object’s centroid at the beginning and at the end of the trial. Both P and ΔO^{pos} are computed with reference to the object’s centre of mass. This metric has been originally proposed in [McCreery and Breed \(2014\)](#) to evaluate the effectiveness of group transport strategies in real ants. It is assumed that the higher the sinuosity the less efficient the transport strategy. This is because very sinuous trajectories are generated by frequent changes in directions of transport which generally result from a poor coordination between the agents while pushing the object. Low sinuosity instead tends to be associated to an effective coordination in which a consensus on the direction of transport is quickly found and maintained during the entire duration of the transport. The lowest value sinuosity can take is 1, which is achieved when the object is transported on the shortest possible path between the start and the final object position.

Figure 5 shows the sinuosity of the object trajectories of successful trials in 16 different evaluation conditions in which we vary the object mass, the object length, and the number of robots in the group. Recall that in each condition each group performs ten trials, and in

each trial, the group transport successfully ends when the object is transported to 1 m away from its initial position. The graph in Fig. 5 suggests that for each group size sinuosity tends to increase for heavy objects. Using the Wilcoxon rank sum test, we compared sinuosity in trials with the light and the heavy object for each object length of each group cardinality. The results of this analysis are shown in Fig. 5, where grey boxes refer to a condition in which there is a significant difference between the sinuosity recorded in trials with the light and the heavy object (Wilcoxon rank sum test, $p < 0.05$); white boxes refer to conditions in which this difference is not significant. The statistical test indicates that, for all group sizes, but only for objects of 40 cm, the sinuosity recorded with the heavy object is significantly higher than the sinuosity recorded with light object. The results of our analysis suggest that, for the 40 cm object, the heavier the object the more difficult for the group to coordinate their individual efforts to find a common direction of transport. In other words, the mass and the length of the longest object's side bear upon the quality of the transport trajectory. This can be due to multiple not mutually exclusive reasons. We observed that the longer and the heavier the object, the higher its tendency to generate rotational (i.e. object rotating on the axis perpendicular to the floor) rather than translational movements when subject to the robots pushing forces. Rotational movements, in turn, directly increases the sinuosity of the object transport trajectory, since they tend to produce a change in the direction of transport. However, future work is required to clarify this issue.

For the short object (i.e. 30 cm), only for the group with 5 robots, the sinuosity recorded with the heavy object is significantly higher than the sinuosity recorded with a light object (see Fig. 5, five agents, 30 cm). For the other three group size, there is no significant difference between the sinuosity recorded with the light and the heavy 30 cm object. How can we account for these results? For the condition with six agents, 30 cm object, the statistics are largely influenced by the fact that the heavy object turned out to be a too difficult transport task, with only two successful trials out of 10. The object turned out to be too small and too heavy with problems mainly caused by the fact that coordination was often achieved having more agents pushing other agents than those directly pushing the object (see also Sect. 6.1). For the other two conditions (i.e. 3 agents 30 cm object, and 4 agents 30 cm object), multiple factors can account for the not significant difference between the sinuosity recorded with the light and the heavy object including the uneven friction between the object and the ground. Uneven friction makes it harder to trigger the initial object movement, and it can interrupt a successful transport with a consequent lost of coordination. The latter phenomenon, that has been observed more often in trials with 30 cm than with 40 cm objects, has a clear impact on the sinuosity of the object trajectory, and it can account for the not significant difference between the sinuosity of light and heavy 30 cm objects.

The results of our tests with physical robots also suggest that both sinuosity and duration of transport tend to increase as the size of the group increases. We analysed the results using the Friedman test (Conover 1999). As the Friedman test is a rank-based nonparametric test, it does not require scaling the performance measure (i.e. sinuosity and duration) computed for each trial nor formulating any restrictive hypothesis on the underlying distribution of the different performance measures (see also Francesca et al. 2015). We represent the results of the Friedman test in a graphical way: two plots, one for duration and one for sinuosity of the trials, that show the expected rank obtained by each group size, together with a 95% confidence interval. If the confidence intervals of two groups do not overlap, the difference between the expected rank of the two is statistically significant. The graphs in Fig. 6 show that the six robots group performs significantly worse than all other groups both in terms of sinuosity and in terms of duration of trial. Unfortunately, with only six physical robots in our lab, we could not further explore this correlation between group size and performance drop




Fig. 6 Friedman test on aggregate data from the four group size concerning: **a** sinuosity; **b** duration of trial. Both in terms of sinuosity and in terms of time, the six robot group performs significantly worse than all other groups. The graphs show the expected rank obtained by each group size, together with a 95% confidence interval. When the confidence intervals of two groups do not overlap, the difference between the expected rank of the two is statistically significant. See also Sect. 6.2 for an explanation of how to read the plot

in terms of quality of transport. However, if what suggested by the graphs in Fig. 6 would be confirmed by further tests with larger groups, we would account for this trend calling upon the effect of robot–robot collisions. The larger the group, the more frequently these collisions occur. If a collision results in both robots occluding each other camera, as it happens during collisions with the object to be transported, the colliding robots may mistake each other for the object. Thus, instead of resolving the collision, they tend to reinforce it. Such an event inevitably increases the time required for the coordination of actions and directly affects the sinuosity and the time of transport. This is because, when two robots collide, the group tends to generate rotational instead of translational movements of the object, due to the lack of required forces to transport the item. We also need to remember that, in our set-up, the object mass increases as number of robots increases. Thus, in order to transport the object successfully, the coordination (i.e. the alignment of pushing forces on a common direction of transport) and the synchronisation (i.e. synchronous exertion of the pushing forces) of actions have to be progressively more accurate and effective. This means that, the object transportation tends to be interrupted more frequently and the direction of transportation also tends to change more frequently. These events are more likely to influence duration and trajectory of transportation in large groups than small groups.

In the next section, we complement our analysis of the behaviour of physical robot by showing the results of several post-evaluation tests aimed at unveiling the individual mechanisms used by the agents to coordinate and synchronise their actions.

6.3 An analysis of the individual mechanisms underpinning the group coordination of action

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 physical robot. In these tests, the robot undertakes multiple trials where it is required




Fig. 7 Graph showing the results of four tests conducted with a single physical robot. The tests are described in the text. Each box refers to the number of repositioning events on a set of ten trials. A repositioning event happens any time the robot stops pushing the object and immediately after starts pushing the object again in a slightly different position. Every trial last for 60 s. The experimental conditions *L*, *H*, *S*, and *R* refer to trials with different type of objects. In *L* (light), we use 30 cm length and 150 g mass object. In *H* (heavy), we use 30 cm length and 600 g mass object. In *S* (short), we use 400 g mass and 20 cm length object. In test *R* (long), we use 400 g mass and 40 cm length object

to push an object with varying characteristics (e.g. a light, a heavy, a short, 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. During each trial, we recorded the activation of the robot sensors' readings, and we use the readings from the front infrared sensors to count the number of repositioning events. A repositioning event corresponds to a variation of the front infrared sensors that deviate from maximum activation when the robot stop exerting pushing forces on the object and return to maximum activation when the robot regains physical contact with the object. 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. Persistence is low when the number of repositioning events is high, and vice versa. As discussed in [McCreery and Breed \(2014\)](#), 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, length, and its direction of movement with respect to the robot heading) and persistence.

In test A and test B, a physical robot is positioned in front of a cuboid object, facing the object at about 20 cm from it. In each trial, the robot is given 60 s to push the object. All tests are repeated for 10 trials. Condition *L* differs from condition *H* for the object mass. In *L*, the object length is 30 cm, and the object mass is 150 g. The robot can easily transport the object. In *H*, the object length is 30 cm, and the object mass is 600 g. The object is too heavy to be moved by the robot. Condition *S* differs from condition *R* for the object length. In *S*, the object mass is 400 g, and the object length is 20 cm. The robot has neither the capability to transport nor to rotate the object. In *H*, the object mass is 400 g, and the object length is 40 cm. Contrary to *S*, in *R* the robot can rotate the object by exerting pushing forces on either end of the longest side, but it cannot transport it. Figure 7 shows the number of

repositioning events counted during each trial of each test. The results clearly show that the number of repositioning events changes with respect to whether or not the object can be moved or simply rotated. When the object is so heavy that it cannot 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. 7, box H and box S, test A and test B). When the object is light enough to be moved or rotated by the robot, we observe a very low number of repositioning events. This indicates that the agent persistence is very high (see Fig. 7, box L and box R, test A and test B, respectively). 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 pushed.

In test C and test D, we look at whether the persistence can vary during a single trial in response to variations on how the object responds to the robots' actions. During group transport, changes in the way in which the object responds to the robot's actions are relatively frequent, since these responses are determined by the degree of alignment of the pushing forces. We simulated these variations by manually altering the object mass adding/removing an iron bar of 450 g to the object. When the object is light, it can be transported by the robot. When the object is heavy, it cannot be moved/rotated by the robot. In this series of tests, the object's length is fixed to 30 cm. For each test, we conducted 10 trials, each trial lasting 180 s. The object mass is altered every 60 s interval. At the beginning of each trial, the robot is position in front of the cuboid object, facing the object at about 20 cm from it. In test C, the object mass is heavy during the first 60 s ("H", 600 g), then light for the following 60 s ("L", 150 g), and then heavy again in the last 60 s interval ("H", 600 g). On the other hand, in test D, the object mass is light during the first 60 s ("L", 150 g), then heavy during the following 60 s ("H", 600 g), and then light again in the last 60 s interval ("L", 150 g).

The graph in Fig. 7 test C and test D shows that the number of repositioning events is altered by the way in which the object responds to the robot actions. The number of repositioning events is quite low when the object is light (see Fig. 7 condition *L* in test C, and test D), and it is higher when the object is heavy (see Fig. 7 condition *H* in test C, and test D). The graph also indicates that the behaviour of the robot is only determined by the current object mass, with no effect of previous experiences. When the object returns heavy after being light, the number of repositioning events increases (see Fig. 7 second condition *H* in test C), while when the object returns light after being heavy, the number of repositioning events drops (see Fig. 7 second condition *L* in test D). It seems that the controller maintains no memory of previous experience with respect to how the object responds to its actions. For more details see movies of the tests available in the supplementary material that can be found at <http://www.aber.ac.uk/en/cs/research/ir/dss/#swarm-robotics>.

We also ran 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

Table 1 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 |

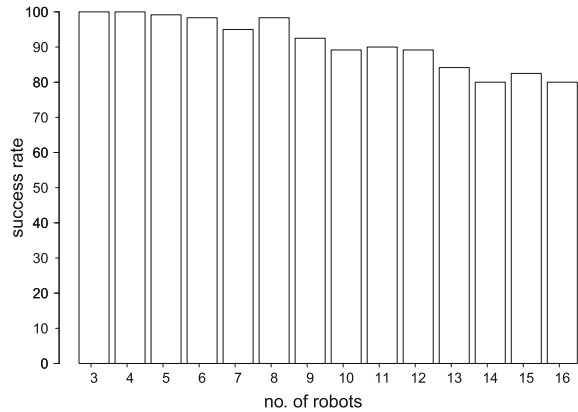
object) and when the object moves against its heading (non-static object). The results of this test, shown in Table 1, clearly indicate that no repositioning events are observed when the robot perceives the object moving against its heading. In other words, the robot quickly changes direction of pushing 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 object¹. 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 by slightly individual differences in persistence that emerge during the course of a trial. However, further investigations are required to better understand this process.

6.4 Scalability analysis

In this section, we illustrate the results of tests on scalability with respect to group size of the best evolved controller. In these tests, each homogeneous group undergoes a set of 120 simulated trials given by all the possible combinations of the parameters object length and mass—we used three object lengths and two object masses—with each combination repeated for 20 trials (i.e. 5 trials for each of the starting positions shown in Fig. 3a). In order to adjust the length of the longest object's side with respect to the group cardinality, we proceeded in the following: we first compute three reference ($L:D$) ratios for a group of four robots and objects of 20, 30, and 40 cm length. Then, for each group size we adjusted L in order to keep the ($L:D$) ratio equal to the three reference ratios. For what concerns the object mass, we set the mass in a way that the object is always heavy enough to require the combined effort of all robots of the group to be transported. The mass is set to 450 and 530 g for the smallest three robots group. For any extra robot added to the group, the object mass is increased by 150 g. For example, the four robots group is tested with an object of 600 g, and with an object of 680 g; the five robots group is tested with an object of 750 g, and with an object of 830 g, and so on, until the 16 robots group. A trial lasts 180 s (i.e. 1384 simulation steps). A trial is considered

Fig. 8 Graph showing the results of the scalability test in simulation. Each bar shows the success rate (%) of homogeneous groups controlled by the best evolved controller during 120 trials. A trial is considered successful if the group manages to transport the object 2 m away from its initial position



successful if the group manages to transport the object 2 m away from its initial position. The graph in Fig. 8 shows that for each group the average success rate is higher than 80%. This indicates that the group transport strategies of the best evolved controller scale relatively well with respect to the group cardinality. The average group performance tends to slightly degrade when the group size increases. This slight performance drop can be attributed, on the one hand, to an increase in the number of robot–robot collisions which hinders the process of alignment of the pushing forces, and on the other hand to the increment of the time required to align the pushing forces. Indirect evidence of the difficulties encountered by large groups comes from the observation of failed trials, where the groups proved unable to transport the objects for the required distance (2 m) within the trial time limits. However, the groups transported the object more than 50% of the distance required to judge the trial successful. We intentionally decided to keep the duration of each trial fixed for all group size. However, this has a clear negative impact on the performance of larger groups, for which more time is required to align the pushing forces.

7 Conclusions

We have described a study in which homogeneous groups of autonomous robots are required to perform a collective object transport task. The robots are controlled by dynamic neural networks synthesised using evolutionary computation techniques. The best evolved controller has been extensively tested on physical e-puck robots in various operating conditions in which we varied the size of the object, its mass, and the number of robots in the group. The results show that the collective transport strategies are fairly robust with respect to variations of the object characteristics and relatively scalable with respect to the group size.

The main contribution of this study is in providing a proof-of-concept demonstration of the effectiveness of group transport strategies generated by robots that cannot feel forces applied to the object to be transported. From an engineering perspective, this result demonstrates that robots equipped with a relatively simple sensory apparatus can develop complex group dynamics related to the alignment of pushing forces and the synchronisation of actions. These group dynamics proved to be effective in initiating and in sustaining the transport of heavy objects. Through an extensive analysis of the operational mechanisms underpinning individual responses, we have identified the main behavioural rules that guide the robots during the initial phase of the task, in which the agents have to agree on a common direction

of transport. In particular, we have shown that the perception of the movement of the object, mainly through the feedback generated by the optic-flow sensor, modulates the frequency with which a robot changes the point of application of its pushing forces. The perception of rotational and translational movements of the object reinforces the pushing behaviour on the same robot–object or robot–robot contact point. The perception of no object movement induces the robot to change the point of application of its pushing forces. These simple rules are 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.

From a biological perspective, the results of our study point to a rather parsimonious explanation of the nature of the mechanisms underpinning object transport in natural swarms. We suggest that the perception of the object movement could be an important cue that guides the behaviour of single ants during the coordination of actions in collective object transport tasks. Evidence in support of this hypothesis could be gathered replicating on real ants the tests described in Sect. 6.3, where the persistence metrics are used to describe the behaviour of individuals while facing either objects that are too heavy to be moved or simply rotated, or facing objects that can be only just rotated, or light enough to be moved.

Another lesson we learn from this study is on the effectiveness of the evolutionary robotics methodology used to design the robots controllers (i.e. artificial evolution in combination with artificial neural network controllers). The advantages and disadvantages of using the evolutionary robotics approach in the context of swarm robotics are extensively discussed in [Trianni et al. \(2014\)](#). We wish to emphasise that with the evolutionary approach applied to swarm robotics we managed to generate effective group transport strategies for physical robots in spite of the fact that controllers have been developed in a simulated environment. The substantial match between the performances of simulated and physical robots illustrated in Sect. 6.1 shows that the evolutionary approach can cope with the inevitable discrepancies between the reality and the simulation environment by generating robust and effective controllers even for complex tasks in which the dynamics of the system cannot be left out from the robot–world model (see [Birattari et al. 2016](#) for an analysis of the issues related to the transfer from simulation to reality of control software). Visual observations of the robot performances, in particular in unsuccessful trials, allowed us to identify and discuss phenomena, such as the case of robots pushing each other illustrated in Sect. 6.1, that in future work can be better modelled to improve the portability of the solutions on to the physical hardware. We also wish to clarify that the specific contribution of various implementation details (e.g. structural and functional properties of the robots’ controllers, the roulette wheel proportional fitness selection algorithm to search the problem space) remain to be empirically evaluated in specifically designed future comparative studies. Within reasonable limits, we have explored various different solutions and tested different combinations of values for the large set of parameters of the system. The implementation details illustrated in this paper are those that returned the best results in term of quality of the group transport strategies and effectiveness in generating these strategies with the evolutionary algorithm. However, we cannot exclude that alternative methodological solutions would not result in the emergence of equally effective or even more robust and more scalable group transport strategies. Finally, our experimental design (i.e. number of evolutionary runs, number of genotypes, number of evaluations per genotype) and the evaluation criteria for physical and simulated robots (i.e. number of trials) are a reasonable compromise between exploration of the problem and exploitation of the computational resources at our disposal. We considered that any action that would result in a longer evolutionary simulation time or longer evaluation time had to be avoided not only for the limits imposed by our computational resources, but also in view of an acceptable scientific practice in which hypotheses can be tested and parameters can be set in a reasonable amount of time.

Future work will concentrate on extending the behavioural capabilities of the robots to allow the group to operate in more complex environments. We have already run further experiments in which the robots capabilities have been extended to deal with tasks requiring the transport of object in an arbitrary direction and toward a specific target area (see Alkilabi et al. 2016a). We are currently looking at collective object transport scenarios in which the presence of obstacles along the transport trajectory requires the group to adjust the direction of motion. These types of scenarios can be extremely challenging because they require the group to maintain the consensus while making a number of serial decision on where to go next before reaching the target location. A recent research work with real ants (*P. longicornis*) has shown that these insects are highly effective at navigating environments in which various types of obstacles obstruct the transport of the object toward a nest area (see McCreery et al. 2016, for details). This study shows that ants implement a flexible collective transport strategy that proved to be effective even in difficult circumstances in which the group has to transport the object in the opposite direction of the target location. One of the aims of our future work is to verify whether the perception of the movement of the object is sufficient for the agents to maintain consensus and to repeatedly change directions of transport while avoiding eventual obstacles and searching for the way to the target location.

Acknowledgements Muhanad H. Mohammed Alkilabi thanks Iraqi Ministry of Higher Education and Scientific Research for funding his PhD. The authors would like to thank G. Francesca and M. Birattari for their support with the statistical analysis of the physical robots performances.

References

- Aiyama, Y., Hara, M., Yabuki, T., Ota, J., & Arai, T. (1999). Cooperative transportation by two four-legged robots with implicit communication. *Robotics and Autonomous Systems*, 29(1), 13–19.
- Alkilabi, M. H. M., Lu, C., & Tuci, E. (2015). Cooperative object transport using evolutionary swarm robotics methods. In P. Andrews, L. Caves, R. Doursat, S. Hickinbotham, F. Polack, S. Stepney, T. Taylor, & J. Timmis (Eds.), *Proceedings of the European conference on artificial life* (Vol. 1, pp. 464–471). Cambridge: MIT.
- Alkilabi, M. H. M., Narayan, A., Lu, C., & Tuci, E. (2016a). Evolving group transport strategies for e-puck robots: moving objects towards a target area. In R. Groß, et al. (Eds.), *Proceedings of the international symposium on distributed autonomous robotic systems (DARS)*. Berlin: Springer STAR.
- Alkilabi, M. H. M., Narayan, A., & Tuci, E. (2016b). Design and analysis of proximate mechanisms for cooperative transport in real robots. In M. Dorigo, M. Birattari, X. Li, M. Lpez-Ibez, K. Ohkura, C. Pinciroli, et al. (Eds.), *Proceedings of the 10th international conference on swarm intelligence (ANTS)* (pp. 101–112). Berlin: Springer.
- Bayindir, L. (2016). A review of swarm robotics tasks. *Neurocomputing*, 172, 292–321.
- Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1(1), 91–122.
- Berman, S., Lindsey, Q., Sakar, M. S., Kumar, V., & Pratt, S. C. (2011). Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems. *Proceedings of the IEEE*, 99(9), 1470–1481.
- Birattari, M., Delhaisse, B., Francesca, G., & Kerdoncuff, Y. (2016). Observing the effects of overdesign in the automatic design of control software for robot swarms. In M. Dorigo, M. Birattari, X. Li, M. Lpez-Ibez, K. Ohkura, C. Pinciroli, et al. (Eds.), *Proceedings of the 10th international conference on swarm intelligence (ANTS)* (pp. 149–160). Berlin: Springer.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford: Oxford University Press.
- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(2), 1–41.
- Chen, J., Gauci, M., Li, W., Kolling, A., & Gross, R. (2015). Occlusion-based cooperative transport with a swarm of miniature mobile robots. *IEEE Transactions on Robotics*, 31(2), 307–321.
- Conover, W. J. (1999). *Practical nonparametric statistics*. New York: Wiley.

- Czaczkes, T. J., Nouvellet, P., & Ratnieks, F. (2011). Cooperative food transport in the neotropical ant *Pheidole oxyops*. *Insectes Sociaux*, 58(2), 153–161.
- Czaczkes, T. J., & Ratnieks, F. (2013). Cooperative transport in ants (hymenoptera: Formicidae) and elsewhere. *Myrmecological News*, 18, 1–11.
- Dorigo, M., & Şahin, E. (2004). Guest editorial. *Special Issue: Swarm Robotics Automatic Robotics*, 17(2–3), 111–113.
- Feener, J. R., Donald, H., & Moss, K. A. G. (1990). Defense against parasites by hitchhikers in leaf-cutting ants: A quantitative assessment. *Behavioral Ecology and Sociobiology*, 26(1), 17–29.
- Francesca, G., Brambilla, M., Brutschy, A., Garattoni, L., Miletitch, R., Podevijn, G., et al. (2015). AutoMoDe-Chocolate: Automatic design of control software for robot swarms. *Swarm Intelligence*, 9(2–3), 125–152.
- Gelblum, A., Pinkoviezky, I., Fonio, E., Ghosh, A., Gov, N., & Feinerman, O. (2015). Ant groups optimally amplify the effect of transiently informed individuals. *Nature Communications*, 6, 1–9. doi:10.1038/ncomms8729.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley.
- Groß, R., & Dorigo, M. (2004a). Cooperative transport of objects of different shapes and sizes. In M. Dorigo, M. Birattari, C. Blum, L. M. Gambardella, F. Mondada, & T. Stützle, (Eds.), *Proceedings of the 4th international workshop on ant colony optimization and swarm intelligence*, volume 3172 of LNCS (pp. 106–117). Berlin: Springer.
- Groß, R., & Dorigo, M. (2004b). Group transport of an object to a target that only some group members may sense. In X. Yao, E. Burke, J. A. Lozano, J. Smith, J. J. Merelo-Guervs, J. A. Bullinaria, et al. (Eds.), *Proceedings of the 8th international conference on parallel problem solving from nature (PPSN)* (pp. 852–861). Berlin: Springer.
- Groß, R., & Dorigo, M. (2008). Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling. *Adaptive Behavior*, 16(5), 285–305.
- Groß, R., & Dorigo, M. (2009). Towards group transport by swarms of robots. *International Journal of Bio-Inspired Computation*, 1(1), 1–13.
- Groß, R., Mondada, F., & Dorigo, M. (2006a). Transport of an object by six pre-attached robots interacting via physical links. In *Proceedings of the IEEE international conference on robotics and automation (ICRA)* (pp. 1317–1323). IEEE.
- Groß, R., Tuci, E., Dorigo, M., Bonani, M., & Mondada, F. (2006b). Object transport by modular robots that self-assemble. In *Proceedings of the IEEE international conference on robotics and automation (ICRA)* (pp. 2558–2564). IEEE.
- Habibi, G., Kingston, Z., Xie, W., Jellins, M., & McLurkin, J. (2015). Distributed centroid estimation and motion controllers for collective transport by multi-robot systems. In A. Okamura et al. (Eds.), *Proceedings of the IEEE international conference on robotics and automation (ICRA)* (pp. 1282–1288). IEEE.
- Habibi, G., Xie, W., Jellins, M., & McLurkin, J. (2014). Distributed path planning for collective transport using homogeneous multi-robot systems. In M. Ani Hsieh, & G. Chirikjian, (Eds.), *Proceedings of the 12th international symposium on distributed autonomous robotic systems (DARS)* (pp. 151–164). Berlin: Springer.
- Hölldobler, B., Stanton, R., & Markl, H. (1978). Recruitment and food-retrieving behavior in novomessor (formicidae, hymenoptera). *Behavioral Ecology and Sociobiology*, 4(2), 163–181.
- Huntsberger, T., Rodriguez, G., & Schenker, P. (2000). Robotics challenges for robotic and human mars exploration. In W. Stone (Ed.), *Proceedings of the robotics 2000*. ASCE. doi:10.1061/9780784404768.
- Kube, C. R., & Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30, 85–101.
- Kube, C. R., & Zhang, H. (1997). Task modelling in collective robotics. In R. C. Arkin & G. A. Bekey (Eds.), *Robot colonies* (pp. 53–72). Berlin: Springer.
- McCreery, H. F., & Breed, M. D. (2014). Cooperative transport in ants: A review of proximate mechanisms. *Insectes Sociaux*, 61, 99–110.
- McCreery, H. F., Dix, Z. A., Breed, M. D., & Nagpal, R. (2016). *Collective strategy for obstacle navigation during cooperative transport by ants*. bioRxiv. <http://biorxiv.org/content/biorxiv/early/2016/06/29/061036.full>.
- Moffett, M. W. (1992). Ant foraging. *Research and Exploration*, 8, 220–231.
- Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., et al. (2009). The e-puck, a robot designed for education in engineering. In *Proceeding of the 9th international conference on autonomous robot systems and competitions* (vol. 1, pp. 59–65). IEEE.

- Nouyan, S., Groß, R., Dorigo, M., Bonani, M., & Mondada F. (2006). Group transport along a robot chain in a self-organised robot colony. In T. Arai et al., (Eds.), *Proceedings of the 9th international conference on intelligent autonomous systems* (pp. 433–442). IOS Press.
- Parker, C., & Zhang, H. (2006). Collective robotic site preparation. *Adaptive Behavior*, 14(1), 5–19.
- Pettinaro, G., Gambardella, L.-M., & Ramirez-Serrano, A. (2005). Adaptive distributed fetching and retrieval of goods by a swarm-bot. In B. Hannaford, G. Bekey, P. Fiorini, Y. Nakamura, L. Bushnell, D. Fox, et al., (Eds.), *Proceedings of the 12th international conference on advanced robotics (ICAR)* (pp. 825–832). IEEE.
- Robson, S. K., & Traniello, J. (1998). Resource assessment, recruitment behavior, and organization of cooperative prey retrieval in the ant *Formica schaufussi* (hymenoptera: Formicidae). *Insect Behavior*, 11, 1–22.
- Tanner, C. J. (2008). Resource characteristics and competition affect colony and individual foraging strategies of the wood ant *formica integroides*. *Ecological Entomology*, 33(1), 127–136.
- Trianni, V., Tuci, E., Ampatzis, C., & Dorigo, M. (2014). Evolutionary swarm robotics: A theoretical and methodological itinerary from individual neuro-controllers to collective behaviours. In P. A. Vargas, E. Di Paolo, I. Harvey, & P. Husbands (Eds.), *The horizons of evolutionary robotics* (pp. 153–178). Cambridge, MA: MIT Press.
- Wang, Y., & de Silva, C. W. (2006). Cooperative transportation by multiple robots with machine learning. In G. Yen, et al. (Eds.), *Proceedings of the IEEE congress on evolutionary computation (CEC)* (pp. 3050–3056). IEEE.
- Wang, Z., & Schwager, M. (2015). Multi-robot manipulation with no communication using only local measurements. In Y. Ohta et al. (Eds.), *Proceeding of the 54th IEEE conference on decision and control (CDC)* (pp. 380–385). IEEE.
- Wang, Z., & Schwager, M. (2016). Kinematic multi-robot manipulation with no communication using force feedback. In D. Kragic et al., (Eds.), *Proceedings of the IEEE international conference on robotics and automation (ICRA)* (pp. 427–432). IEEE.
- Wang, Z., Takano, Y., Hirata, Y., & Kosuge, K. (2004). A pushing leader based decentralized control method for cooperative object transportation. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems* (vol. 1, pp. 1035–1040). IEEE.
- Woern, H., Szymanski, M., & Seyfried, J. (2006). The i-swarm project. In K. Dautenhahn et al. (Eds.), *Proceeding of the 15th IEEE international symposium on robot and human interactive communication* (pp. 492–496). IEEE.
- Yamamoto, A., Ishihara, S., & Fuminori, I. (2009). Fragmentation or transportation: Mode of large-prey retrieval in arboreal and ground nesting ants. *Insect Behavior*, 22, 1–11.