

# On Embodied Memetic Evolution and the Emergence of Behavioural Traditions in Robots

Alan FT Winfield · Mehmet Dincer Erbas

Received: date / Accepted: date

**Abstract** This paper describes ideas and initial experiments in embodied imitation using e-puck robots, developed as part of a project whose aim is to demonstrate the emergence of artificial culture in collective robot systems. Imitated behaviours (memes) will undergo variation because of the noise and heterogeneities of the robots and their sensors. Robots can select which memes to enact, and – because we have a multi-robot collective – memes are able to undergo multiple cycles of imitation, with inherited characteristics. We thus have the three evolutionary operators: variation, selection and inheritance, and – as we describe in this paper – experimental trials show that we are able to demonstrate embodied movement-meme evolution.

**Keywords** Robot imitation · Artificial Culture · Memetic evolution · collective robotics

## 1 Introduction

This paper presents both ideas and initial results from an ongoing multi-disciplinary research project called “the emergence of artificial culture in robot societies” whose overall aim is to investigate the processes and mechanisms by which proto-cultural behaviours, better described as traditions, might emerge in a free running collective robot system. We accept, as a working hypothesis, the idea that mimesis and embodiment are

essential pre-requisites for cultural evolution [5]. It follows that since our aim is to demonstrate artificial culture we need a system of embodied artificial agents, i.e. robots, in which robots are able to learn socially from each other, by imitation. This group of robots, which we call ‘Copybots’<sup>1</sup>, require an environment in which behaviours can be copied, by imitation, from one robot to another and we refer to this environment as the ‘artificial culture lab’.

Importantly the Copybots implement a form of embodied imitation in which one robot can copy another’s behaviours only by observing with its physical senses and then transforming that sequence of sense-data (perceptions) into motor actions. We do not allow the robots to transmit behaviours (i.e. sequences of motor actions) directly from one to another. This means that the Copybots have to overcome essentially the same problems of inferring each others’ behaviours from possibly unreliable first-person perceptions as any embodied agents (robots, animals or humans), yet at same time the Copybots implement a rather minimal model of social learning by imitation. We argue that this embodied yet abstract model of social learning, by imitation, provides both a degree of biological plausibility and opportunities for unexpected emergence that would not be present in an agent-based simulation.

Dawkins coined the term ‘meme’ to describe a unit of cultural transmission [7], and we use this terminology here. We propose a definition of a robot meme as follows: *a contiguous sequence or package of behaviours copied from one robot to another, by imitation*. In the artificial culture lab we ‘seed’ each Copybot with initial behaviours which, in this paper, are self-contained movement sequences. We then free-run the Copybots so

---

A Winfield (✉) · M Erbas  
Bristol Robotics Laboratory  
University of the West of England, Bristol, UK  
Tel.: +44 117 328 2211  
E-mail: Alan.Winfield@uwe.ac.uk

M Erbas  
E-mail: Mehmet.Erbas@brl.ac.uk

---

<sup>1</sup>After [5], pages 106-107

that robots alternate between enacting movement sequences and watching (and learning) those sequences. The first-generation copy (child) of the initial seed behaviour may then itself be copied giving a second generation copy in a free-running iterative process that means we have behavioural heredity. The errors that inevitably occur in embodied robot–robot imitation, due to noisy sensors, imperfect observation and the estimation process of inferring motor actions, will give rise to variation in imitated behaviours. If, furthermore, we allow robots to select which socially learned behaviour to enact, then we have the three evolutionary operators necessary for behavioural (i.e. memetic) evolution. The artificial culture lab provides us with the infrastructure and tools needed to track and record these possibly complex sequences of interactions for later analysis.

We propose a definition for ‘artificial traditions’ as follows: *measurable and sustainable differences in the memes across different groups of robots, where those memes can be traced back to common ancestral memes.* We hypothesise that the system outlined in this paper, might provide the conditions in which we could observe both memetic evolution and the emergence of new memes which, perhaps because they are suited to the robot’s physical or sensory morphology, or for entirely contingent reasons, become dominant in the group of Copybots, i.e. new artificial traditions.

Before proceeding it is important to clarify the distinction between the terms memetic evolution, used in this paper, and memetic algorithm. A memetic algorithm is generally taken to be a population-based genetic algorithm in which some form of individual learning also takes place; for a review of memetic algorithms see [15]. In the work of this paper, however, there is no genetic evolution. Instead discrete behaviours, which are transmitted by imitation (social learning) are themselves subject to variation, selection and inheritance. Thus we use the term memetic evolution to mean: a Darwinian evolutionary process in which memes, not genes, evolve.

This paper is organised as follows. In Section 2 we define embodied imitation and place this work in the context of the research literature on robot imitation. Then in Section 3 we describe the robots and infrastructure of the artificial culture lab. Section 4 describes the robot–robot movement imitation algorithm at the core of the Copybots, detailed in Section 5. Section 6 outlines the results from initial trials with robots. Finally Section 7 concludes the paper with discussion and further work.

## 2 Embodied Imitation

We define embodied imitation as: *the imitation of a teacher-robot’s behaviour, or sequence of behaviours, by a learner-robot, where the learner-robot uses only its embodied sensors to perceive the teacher-robot.* This definition precludes robot–robot ‘telepathy’, i.e. one robot transferring part of its controller directly to another, and our insistence on embodied perception of one robot by another means that an implementation of embodied imitation needs to solve the so-called ‘correspondence problem’; a term which refers to the learner’s problem of translating a set of perceptual inputs to motor actions that correspond with the perceived actions of the teacher [18].

The study of imitation and social learning in robots, humans and animals has received cross-disciplinary attention in recent years [19,17]. Not surprisingly much attention has been given to the problem of humanoid robots imitating humans, since this presents a way of programming a robot by demonstration rather than coding, see for instance [20,14].

There has been less work describing experiments in embodied robot-robot imitation. The earliest is perhaps the work of Hayes and Demiris which describes an approach with one (pre-programmed) robot finding its way through a maze and another following it and observing its actions (turns). The following (learner) robot then associates each observed action with its own (time delayed) perception of the environment and hence learns how to navigate the maze, by imitation; this kind of imitation is called ‘matched dependent behaviour’ [11]. Dautenhahn’s seminal 1995 paper for the first time proposed a larger context of social interaction between embodied robots, in which a similar approach to imitation, i.e. one robot learning by following another, played a part [6]. Following a connectionist approach Billard and Hayes proposed the DRAMA architecture (Dynamical Recurrent Associative Memory Architecture); they provide one case study that interestingly involves the active participation of the teacher robot in the process of imitative learning [4].

Following their 1994 work, [11] Demiris et al, went on to propose the HAMMER architecture (Hierarchical, Attentive, Multiple Models for Execution and Recognition). In an important series of papers Demiris et al, developed an imitation architecture based upon the building block of paired inverse and forward models; the inverse model outputs motor commands but, instead of actually driving the motors, those commands are fed to the forward model and the output of the forward model compared with the input of the inverse model. Thus the robot is able to internally ‘rehearse’ possible

actions and compare these with its perception of the actions it is trying to imitate [8, 12, 9].

Alissandrakis et al, developed the ALICE architecture (Action Learning via Imitation between Corresponding Embodiments) in order to address the problem of robot-robot imitation across dissimilar embodiments; although not tested with real robots, ALICE contributes a powerful generalised solution to the correspondence problem for agents (or robots) with different morphologies [1, 3, 2]. ALICE works by creating a correspondence library relating the actions (and importantly effects) of the teacher to actions (or action sequences) that the learner is capable of.

### 3 The Artificial Culture laboratory

The artificial culture lab comprises a physical space (arena) designed for and populated by miniature wheeled mobile robots. The arena is closed in the sense that its physical boundaries define the edges of the robots' world, out of which they cannot physically stray. The arena is open to the environment, thus robots (since they have both light and sound sensors) are affected by ambient lighting or noise levels. Providing that these external environmental influences do not overwhelm the robots' sensors, they are not a problem. Indeed, a low-level of background 'noise' in the environment, for instance local variations in light levels, may usefully contribute to imperfect robot-robot imitation.

We make use of the well known *e-puck* robots [16]. Importantly, the *e-puck* robots can sense and track the movements of other robots nearby (albeit imperfectly because of their limited sensors); thus robots have the physical means for imitation. They have multi-coloured programmable lights (LEDs), simple cameras; microphones and speakers giving a wide range of options for robot-robot interaction. Robots can signal to each other with movement, light, or sound, one-to-one or one-to-many. Despite its clear strengths, the basic *e-puck* lacks the computational power for the image processing needed in vision based imitation and, to overcome this limitation, we have designed an open-hardware Linux extension board for the *e-puck*, based on the 32-bit ARM9 microcontroller [13]. The extension board also provides WiFi communication, essential for monitoring and data logging, and Player [10] integration to facilitate code development.

The artificial culture lab is fully instrumented. A vision tracking system from Vicon<sup>TM2</sup> provides high precision position tracking. A dedicated swarm-server combines the tracking data from the Vicon system with sta-

tus information, such as internal state data, from robots (via the WiFi network) and saves the experimental logs for offline analysis and interpretation. Fig. 1(a) shows the artificial culture lab arena in the Bristol Robotics Lab (BRL); Fig. 1(b) shows one of the *e-puck* robots fitted with Linux extension board, red skirt and tracking 'hat'.

### 4 Robot-robot imitation of movement

We are concerned here with the embodied imitation of behaviour, but 'behaviour' is too broad a term. Within embodied imitation we can identify at least three types of imitation:

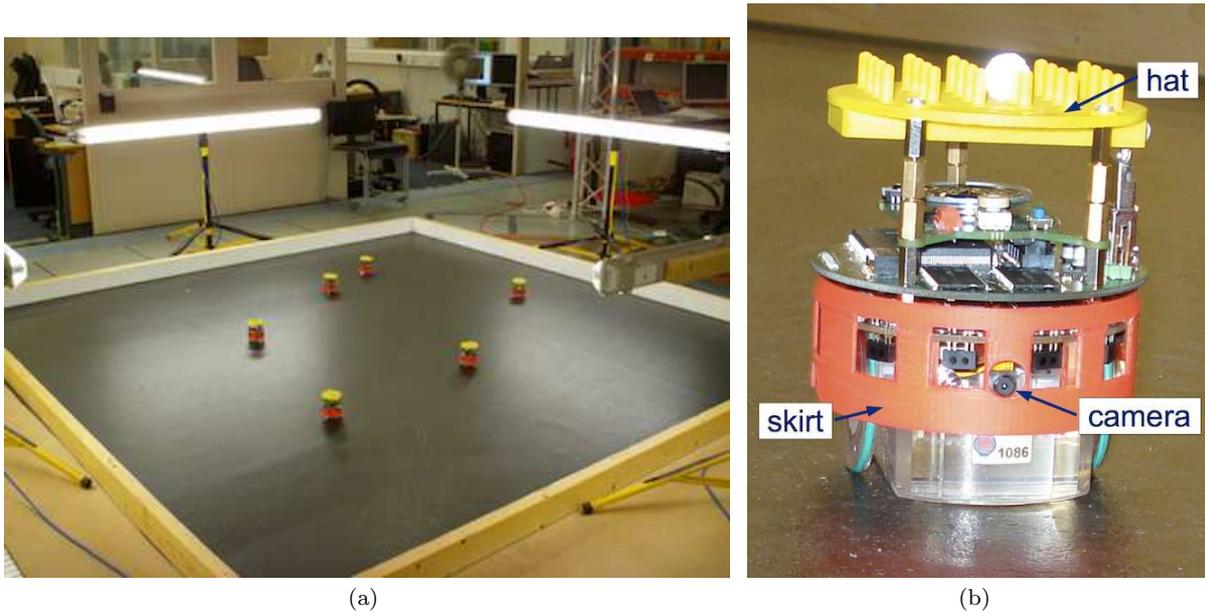
- imitation of actions only, i.e. one robot copying another's sequence of movements, sounds or lights;
- imitation of action and perception, i.e. one robot copying another's interactions with objects or other robots: we label this the 'imitation of interaction' and it is distinct from the imitation of actions because it requires that the learner-robot infers how changes to sensory inputs trigger actions from observation of the teacher-robot;
- imitation of goals, i.e. one robot copying the goals or intentions of another using, perhaps, a completely different set of actions.

Because it is the simplest we have first implemented the imitation of actions, and specifically the imitation of movement. The second, imitation of interaction, and third, imitation of goals, are outside the scope of this paper but will be considered as further work. We now describe the imitation-of-movement algorithm and experimental results obtained.

Before outlining the imitation algorithm we first need to describe the basic setup and some simplifying assumptions. In this approach the teacher-robot performs a sequence of movements while the learner-robot watches it and attempts to learn the observed sequence. The roles of teacher and learner are not fixed but interchangeable and – since we are interested in propagation of imitated behaviours – robots alternate between teacher and learner modes. In teacher mode when a robot is ready to perform a movement sequence it will first signal this by flashing its red LEDs. After completing the movement sequence the teacher-robot will signal again with its red LEDs. The learner-robot remains stationary while observing, only turning on-the-spot if necessary to keep the teacher-robot in its field of view.

The learner-robot watches the teacher robot with its onboard camera and, in order to facilitate the recognition of the teacher-robot and its movements, robots are

<sup>2</sup>(<http://www.vicon.com>)



**Fig. 1** (a) Artificial culture lab showing 6 robots in the arena. (b) An e-puck with Linux board fitted in between the e-puck motherboard (lower) and the e-puck speaker board (upper). Note both the red skirt, and the yellow ‘hat’ (which provides a matrix of pins for the reflective spheres which allow the tracking system to identify and track each robot).

fitted with coloured skirts that contrast with the background (i.e. arena boundaries), as shown in Fig.1(b). Since robots have only one camera and hence monoscopic vision the learner robot must judge the relative direction of movement and distance of the teacher robot by tracking the position and size of the teacher’s coloured skirt in its field of view. Although estimating relative size and position of the teacher robot is straightforward, it is error prone because of the relatively low resolution camera (640x480) and the presence of other robots and, furthermore, the learner robot cannot see the teacher robot’s turns – only infer them from changes in direction, thus we simplify the correspondence problem by limiting movement sequences to be composed of on-the-spot turns (rotations) and straight line segments at a constant speed. In the experiments of this paper the straight line speed was fixed at 5.2 cm/second and turns take, on average, 1 second.

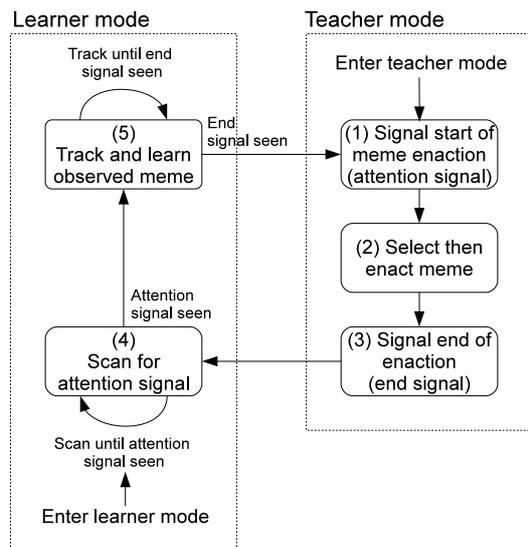
The imitation-of-movement algorithm thus has three stages:

1. while observing captured visual frames, for each frame estimate the relative position of the teacher robot, storing these positions in a linked list;
2. after the teacher’s sequence is complete, process the linked list using a regression line-fitting approach to convert the estimated positions into straight line segments (vectors);
3. transform the straight line segments, and their intersections, into a sequence of motor commands (turns and moves).

The imitation-of-movement algorithm outlined here does not have the complexity of the architectures outlined above in section 2, although it does clearly share a number of common elements. There are a number of reasons for the relative simplicity of our approach. Firstly, we are here imitating movement only and not interaction with objects, or other robots: thus the learner needs only to deduce action sequences and not perception-action sequences. Secondly our robots are homogeneous (similarly embodied), thus when the learner robot transforms its estimate of the teacher robot’s movement trajectory into ego-centric motor commands it can assume it has the same motion dynamics as the teacher robot; all robots move, while executing straight line segments, at a constant velocity. Thirdly, we are interested primarily in meme propagation across the robot collective, so our approach to imitation is only as complex as needed to create the conditions for movement-meme evolution.

## 5 Embodied memetic evolution: the Copybots

Fig. 2 shows the Finite State Machine (FSM) for our initial trials in embodied memetic evolution. Each robot runs the same FSM and we call these robots Copybots, since they have no behaviours other than imitation [5]. The Copybots alternate between learner and teacher modes. In teacher mode, on the right hand side of Fig. 2, a robot progresses through three states:



**Fig. 2** Finite State Machine for Copybots. Unlabelled transitions take place when the previous state completes (timings are explained in the text). Different entry points for robots starting in teacher or learner modes are shown. Once started the FSM loops indefinitely, alternating between teacher and learner modes.

1. Signal start of meme enaction. Here the robot lights its red body LEDs for a pre-set time period (35 seconds), to signal to other robots that it is about to enact an imeme<sup>3</sup> from its memory. We call this the ‘attention’ signal.
2. Select and Enact meme. Firstly, the robot selects an imeme from its memory for enaction (the choice of selection operator is described below). Secondly, the robot enacts the selected imeme by executing the sequence of turns and straight-line moves specified by the imeme. The time spent in this state depends on the number of turn-and-move pairs in the sequence and the length of the moves. (As an example a movement pattern which describes a square with 25 cm sides will take about 24 seconds to complete.)
3. Signal end of meme enaction. Here the robot again lights its red body LEDs for a fixed pre-set time period (1.6 seconds), this time to signal that meme enaction is complete (the ‘end’ signal). After completing this action the robot exits teacher mode and enters learner mode, state 4.

In learner mode, on the left hand side of Fig. 2, a robot progresses through two states:

4. Scan for attention signal. Here the robot scans for another robot’s ‘attention signal’ by rotating and scanning 360°. When it sees the attention signal (state 1 above) the robot transitions to the ‘track and learn observed meme state’ described next.
5. Track and learn observed meme. Here the robot tracks the movements of the other robot that has got its attention (rotating if necessary to keep it in its field of view), executing the algorithm described in Section 4. The robot will continue to track until it sees the ‘end’ signal (state 3 above), at which point it will complete the process and store the new learned meme in its memory. The robot will then exit learner mode and enter teacher mode, state 1.

In order to coordinate the behaviour of teacher- and learner-robots while at the same time allowing robots to alternate between these states we arrange that robots are organised at the start of an experiment in pairs. In each pair one robot starts in teacher-mode and the other in learner-mode. The entry points to the FSM for these two modes are shown in Fig. 2. This means that while the teacher-robot (of the pair) is in state 1 and making the attention signal, the other robot (the learner) is in state 4, scanning for an attention signal. Since each robot runs the same FSM, which loops indefinitely, then the two robots of the pair will alternate between teacher and learner modes, each in the other mode. Note also that in experiments with more than one pair of robots it is perfectly possible for a learner-robot to catch sight of the attention signal of the teacher-robot of a different pair. We regard this as a desirable side-effect of the attention-scanning mechanism of states 1 and 4 as it allows different teacher-learner pairs to emerge and hence behaviours to propagate across the entire robot collective.

Consider now the selection operator of teacher mode, state 2, above. There are clearly many criteria that could be used to select which stored imeme to enact but, in initial experimental trials, we have chosen to test operators that do not select for particular features or characteristics of stored memes. We thus avoid an explicit fitness function in order to allow for the possibility of undirected, or open-ended memetic evolution. Instead we apply simple probabilistic selection criteria, defined as follows:

- Always select the most recently learned meme. This is, in effect, memoryless imitation.
- Select from the meme list at random, but weighted in favour of the memes judged to be seen most often. Here not all memes are stored, only memes that are judged to be new, i.e. greater than a threshold value of similarity, are added to the memory.

<sup>3</sup>We use ‘imeme’ as shorthand for ‘internal representation of a meme’.

- Select from the meme list by choosing, at random, with equal probability. Here we assume that all observed memes are stored, regardless of their similarity.

## 6 Experimental Results

We have to date conducted trials with 2 and 4 robots, limiting ourselves initially to this small number of robots in order to focus on understanding the fidelity of imitation (variation) observed in embodied imitation, and the effect of different simple selection operators. We present here representative results of three trials, which – in turn – test the three selection operators outlined above. For each trial we show trajectory plots generated from the Vicon tracking data (robot x,y positions) logged, together with robot internal state data, by our swarm-server. We present brief qualitative results for the first two selection operators: select the most recently observed meme, and random selection with weighted probability. For the third selection operator, random selection with equal probability, we present a more detailed analysis in which we define a quality of imitation metric and then use the metric to quantitatively analyse the first two imitation events before then tracking the history of movement-meme evolution through the course of the experimental trial.

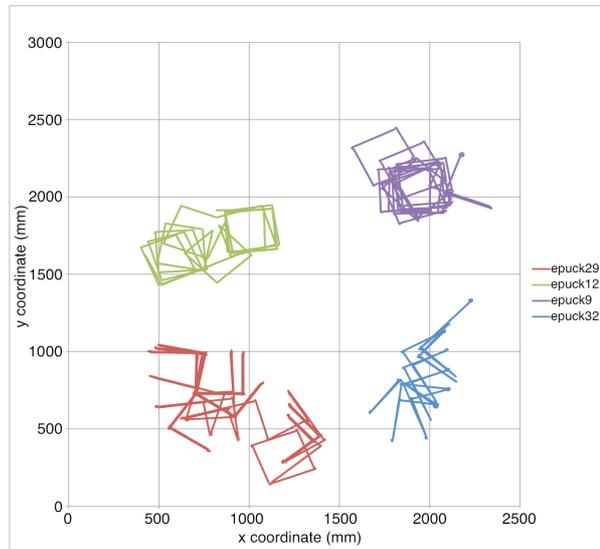
### 6.1 Two robots: Select most recently observed meme

In Fig. 3 we see a trajectory plot of a two robot experiment in which each robot alternates between teacher and learner, and enacts the most recently seen meme. Initially e-puck 9 is in teacher mode, and e-puck 12 in learner mode, and e-puck 9’s memory has been preloaded with a movement sequence that describes a square with sides of 25cm. The experimental trial illustrated in Fig. 3 takes about 20 minutes and consists of about twelve complete cycles of the FSM in Fig. 2, i.e. six teacher-learner imitation events per robot. It is important to note that the robot trajectories in Fig. 3 are not continuous, with each robot stopping after each meme enactment in order to scan for and then track the other robot. A screen recording can be viewed or downloaded from [22], which shows a ‘playback’ of the first 14 minutes of the trajectories in Fig. 3, in Stage [21], with status information. Here we change colour each cycle to facilitate analysis.

The annotation on Fig. 3 show the initial square (e-puck 9), and the 1st, 2nd, and 3rd generation copies. It is immediately clear that each successive copy introduces new imitation errors, in both turn angles and

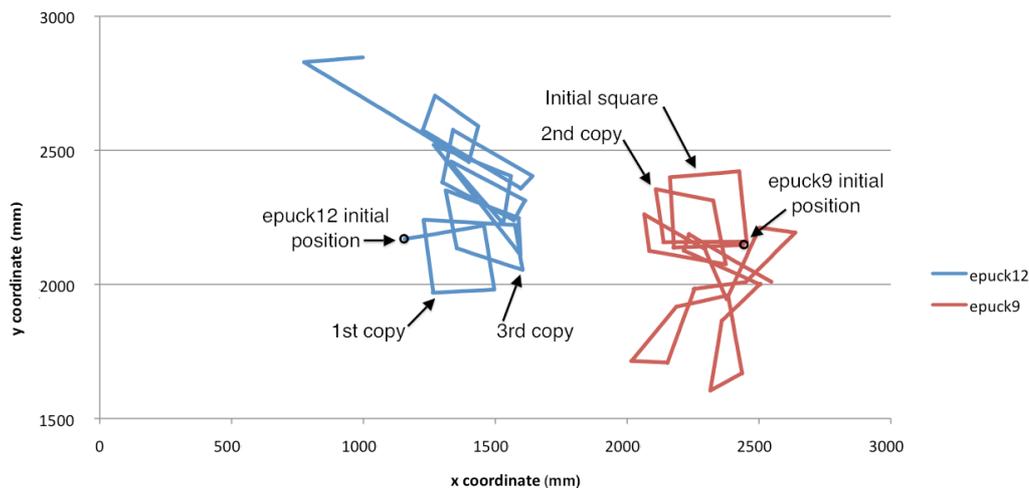
straight-line lengths, which accumulate and quickly result in significantly distorted variants of the original square. However, analysis shows that the number of segments remains at four until the robots become too far apart to see each other (and the experiment is halted), thus there is an inherited characteristic in this very simple example of movement-meme evolution.

### 6.2 Four robots: Random selection with weighted probability



**Fig. 4** Trajectory plot: four robot movement-meme evolution in which only dissimilar memes are stored and meme selection is random, weighted in favour of similarly observed memes. The four e-pucks are organised as two pairs: e-pucks 9 and 12, and e-pucks 29 and 32. At the start of the experiment e-pucks 9 and 29 are initialised in teacher mode and e-pucks 12 and 32 in learner mode.

In the trajectory plot of Fig. 4 we see four robots organised as two pairs: e-pucks 9 (purple) and 12 (green) are initialised with a movement pattern which describes a square trajectory and e-pucks 29 (red) and 32 (blue) initialised with a closed 4-segment L-shaped pattern. This experimental trail differs from those above in that here we have constrained the movement-imitation algorithm (both in estimating the turn angles, and in enacting movements) to multiples of  $90^\circ$ . The effect of this constraint is evident in the qualitatively high fidelity copies of both square and L-shaped movement patterns. What is, however, interesting is that robots in the 9/12 pair have, from time to time observed and stored the meme patterns from the 29/32 pair, and vice-versa. Because of the weighted selection criterion the



**Fig. 3** Trajectory plot: two robot imitation. The experiment starts with e-puck 9 in teacher mode, following a movement trajectory that describes a square with sides of 25cm; e-puck 12 is initialised in learner mode.

odds are against one of these being selected and enacted but, nevertheless, we see a square pattern enacted by e-puck29 (lower left) and an L-shaped pattern enacted by e-puck9 (upper right).

### 6.3 Two robots: Random selection with equal probability

#### 6.3.1 A quality of imitation metric

In order to quantitatively assess the fidelity of imitation (i.e. similarity of learned memes) we need to define a quality-of-imitation function,  $Q_i$ . Since movement-memes consist of a series of turn and straight line segments (vectors) we can compare the similarity of two memes by separately estimating three quality indicators: quality of angle (turn) imitation, quality of length imitation, and quality of segment imitation. In effect we decompose each movement-meme into its constituent turn angles, straight-line move lengths and number of straight-line segments. The quality of angle imitation between original meme ( $O$ ) and learned meme ( $L$ ) is calculated as follows:

$$Q_a = 1 - \frac{\sum_m |a_m^L - a_m^O|}{\sum_m a_m^O} \quad (1)$$

where  $a_m$  is the turn angle of move  $m$ , in degrees. Here we determine the ratio of the sum of angle differences between the moves of original and learned memes, to the total turn angle of the moves of the original meme. If original and learned memes have a different number of segments,  $N^O$  and  $N^L$  respectively, then we sum only over the number of segments in the smaller:  $\min(N^L, N^O)$ . A value of 1 indicates perfect fidelity

imitation. The quality of length imitation similarly calculates the length errors between original and learned memes:

$$Q_l = 1 - \frac{\sum_m |l_m^L - l_m^O|}{\sum_m l_m^O} \quad (2)$$

where  $l_m$  is the length of move  $m$ , in mm. Again a value of 1 indicates perfect fidelity imitation. The quality of segment imitation simply considers the difference between the number of segments (vectors) between original and learned memes:

$$Q_s = 1 - \frac{|N^L - N^O|}{N^O} \quad (3)$$

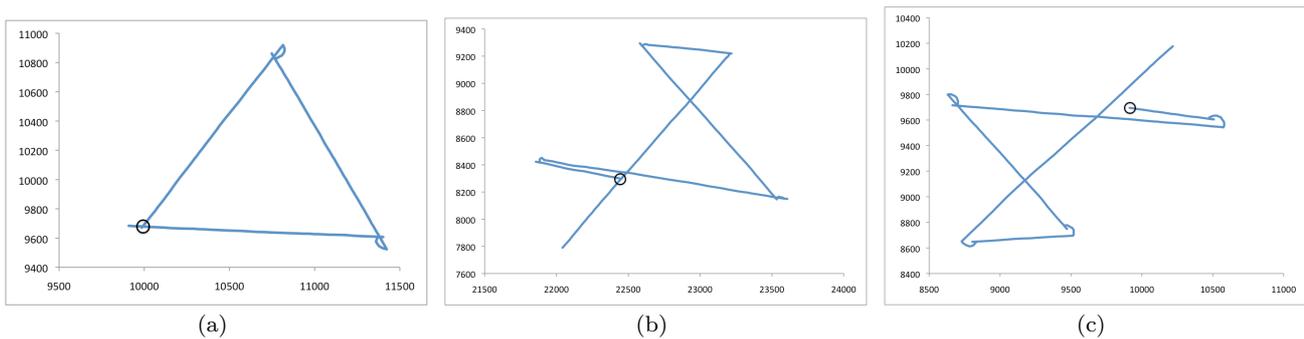
where  $N^L$  and  $N^O$  are the number of segments of learned and original memes, respectively. We now calculate the weighted sum of the three quality indicators, to arrive at a composite overall quality-of-imitation score:

$$Q_i = A Q_a + L Q_l + S Q_s \quad (4)$$

where  $A$ ,  $L$  and  $S$  are weighting coefficients, and  $A + L + S = 1$ . In the results given here we choose to give equal weighting to each of the three quality indicators, thus  $A = L = S = 0.33$ .

#### 6.3.2 Robot-robot imitation with variation

Fig. 5 shows two examples of embodied social learning, of movement, by imitation. Each of the three subfigures in Fig. 5 plots tracking data recorded, from the experimental infrastructure described in Section 3, when an e-puck robot enacts a single movement sequence. Here e-puck 9 has been initialised with a sequence of three turns and moves that describe an equilateral triangle, with 15 cm sides, and Fig. 5(a) shows e-puck 9 enacting the triangle. In this trial e-puck 12 watched e-puck



**Fig. 5** (a) Meme 1: initial movement meme enacted by e-puck 9. (b) Meme 2: imitation of meme 1 by e-puck 12,  $Q_i = 0.47$ . (c) Meme 3: imitation of meme 2 by e-puck 9,  $Q_i = 0.94$ . X and y axes are labelled in multiples of 0.1 mm and the circles mark the start positions of each movement-meme.

9 enact the triangle and, using the procedure outlined above, attempted to learn the movement sequence, by embodied imitation; the result is shown in Fig. 5(b) and it is immediately clear that this is a poor-fidelity copy. Although the copy clearly retains characteristics of the original triangle, i.e. the 3 longest line segments form a triangle and intersect at angles close to  $45^\circ$ , two additional short segments have been inserted, one at the start, followed by a u-turn, and another at the top apex of the triangle. Given these two additional segments it is not surprising that our quality of imitation score is poor:  $Q_i = 0.47$ . The quality of length imitation is much higher:  $Q_l = 0.75$ .

In this trial e-puck 9 then watched Meme 2, enacted by e-puck 12, and attempted to learn it, thus Fig. 5(c) is an imitation of Fig. 5(b). In contrast with the poor fidelity Meme 1  $\rightarrow$  Meme 2<sup>4</sup> imitation, we see that Meme 2  $\rightarrow$  Meme 3 imitation is much higher fidelity. Meme 3 is of course rotated with respect to Meme 2, but that is exactly what we would expect. Meme 3 retains the rather complex five segment structure of Meme 2, and gives a very high quality of imitation score of  $Q_i = 0.94$ ; e-puck 9 has certainly learned the complex ‘dance’ of e-puck 12. We have thus demonstrated both robot–robot social learning, by imitation, and shown that we obtain variation in socially learned behaviours ‘for free’ as a consequence of embodiment.

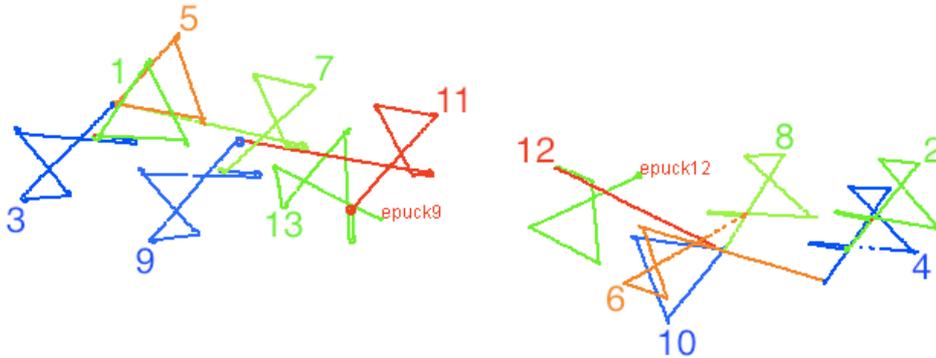
The question of precisely why the first imitation here, Meme 1  $\rightarrow$  Meme 2, was of much lower fidelity than the second, Meme 2  $\rightarrow$  Meme 3, is hard to answer. However, there are several factors at work. One is that the seeded movement pattern, the triangle, was chosen arbitrarily and without regard to which movement patterns may be easy or hard for the robots to copy. We must not assume that any pattern we choose can be imitated with high fidelity since the robots physical

and sensory morphology, and imitation algorithm, will almost certainly favour some patterns over others; this question is the subject of further work. Another factor is ambient environmental variation - it is possible that changes in light level across the arena made it harder for e-puck 12 to see e-puck 9 than vice-versa. Such contingencies are inevitable in embodied experiments of this nature. Furthermore, the fact that the imitation algorithm must infer turns by looking for the intersection of straight lines fitted to discrete estimated positions (as outlined in Section 4), combined with errors in estimating the position of the teacher-robot, may well be responsible for imitation artefacts such as the straight line segment inserted into the apex of the triangle at the top of Fig. 5(b).

### 6.3.3 Towards open-ended memetic evolution

Fig. 6 plots the position data captured during a two robot experiment in which each robot alternates between teacher-mode and learner-mode. Each robot learns and stores the meme enacted by the other, but then – when in teacher-mode – chooses which meme to enact using the equal-weighting random-selection operator. For clarity each movement sequence is shown here in a different colour, and labelled with the order in which the movement-memes were enacted by the two robots. In this run each robots memory is initialised with one imeme: a pattern of movements that describe an equilateral triangle with sides of 15 cm, and e-puck 9 is initially in teacher mode. The plot in Fig. 6 is a screen capture from a tool that we have developed for Stage [21], which ‘plays back’ the trajectories recorded from the real-robot experiments. In this tool we change colour each cycle to facilitate analysis. The screen recording from which Fig. 6 was captured (and then annotated) may be viewed or downloaded from [23].

<sup>4</sup>We use  $A \rightarrow B$  as shorthand for B is a learned and enacted copy of A.



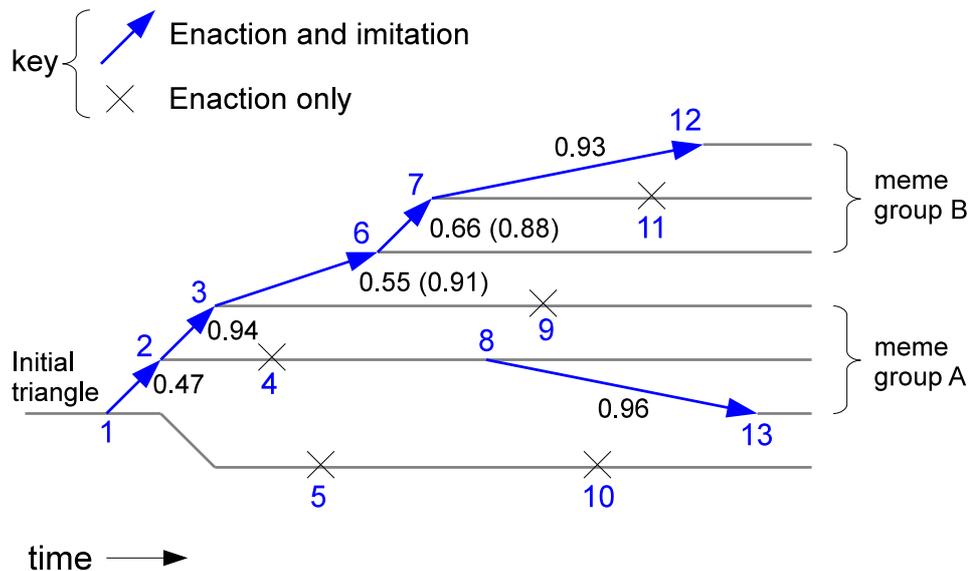
**Fig. 6** Trajectory plot: two robot movement-meme evolution in which all observed memes are stored and meme selection is random, with equal probability. The experiment starts with e-puck 9 (left) in teacher mode, following a movement trajectory that describes a triangle with sides of 15 cm.

We now apply the graphical meme-tracking approach proposed in [24], in order to trace the evolution of memes in the experiment of Fig. 6. Inspection of Fig. 6 shows that a ‘figure of eight’ meme appears to dominate (movement memes 2, 3, 4, 6, 8, 9 and 13), and the meme evolution diagram in Fig. 7 explains why.

Fig. 7 shows the evolution and heredity of memes in the two robot experimental trial of Fig. 6. It does not identify robots, but instead traces the evolution of memes – something which is not obvious from the trajectory plot of Fig. 6. To create Fig. 7 requires us to analyse the recorded experimental logs of each robot’s activity, step by step, in order to link the stored imemes of the 2 robots; in other words determine for each learned imeme, in the memory of each robot, which imeme of the other robot it is a copy of. Thus, each horizontal grey line in Fig. 7 represents a timeline for each single imeme. When that imeme is selected and enacted there are two possibilities: one is that the enaction was not, at any time during the experimental trial, imitated (i.e. learned *and* enacted) – these are shown as crosses (and labelled in the key ‘enaction only’). The other possibility is that the enaction was imitated during the experimental trail – these are shown as blue arrows (and labelled in the key ‘enaction and imitation’). The enaction only events (crosses) are labelled with the meme-enaction number in trajectory plot Fig. 6; the enaction and imitation events (blue arrows) are labelled with originating (parent) meme-enaction number, and learned meme-enaction number from Fig. 6. Each enaction and imitation event is also labelled with the quality-of-imitation score  $Q_i$ . Note that each enaction and imitation event results in a new imeme timeline

which continues throughout the trial. This reflects the fact that our robots, in this experiment, have unlimited imeme memories. If we instead had either a mechanism for robots ‘forgetting’ imemes (according to some criteria) or robots themselves ‘dying’, then some imeme timelines would terminate.

Thus we see, in Fig. 7, that Meme 2 is a poor-fidelity copy of Meme 1 (0.47) – the first ‘figure of eight’ movement pattern. Significantly, Meme 3 happens to be a high-fidelity copy of Meme 2 (0.94), and furthermore there are no further enaction and imitation events originating from Meme 1 – just two enaction only events: 5 and 10. Thus, all second and later generation memes have, as an ancestor, Meme 2. This fact, together with the high-fidelity copy of Meme 8 → Meme 13 (0.96) means that Memes 2, 3, 4, 8 and 13 are all either the same or very closely related and we label these Meme group A. Consider now imitation event Meme 3 → Meme 6, which appears to be relatively poor quality (0.55). However inspection shows that Meme 6, which has four segments, has lost the initial short segment of Meme 3; if we ignore the first segment of Meme 3 and re-calculate  $Q_i$  for Meme 3 → Meme 6, we obtain 0.91 – which more closely reflects the subjective similarity of Memes 3 and 6. By chance imitation event Meme 6 → Meme 7 has inserted a new short segment so that Meme 7 returns to five segments and, by similarly ignoring the new segment in Meme 7 and re-calculating  $Q_i$ , we obtain 0.88. Thus we see that Meme group B is both quantitatively and subjectively similar to Meme group A, with strongly inherited characteristics retained across five generations of meme: 2 to 12. We now understand



**Fig. 7** A visualisation of meme evolution within the (two) robot collective. At the start of the period just one movement-meme (triangle) is present; horizontal lines represent the ‘life course’ of each meme from left to right. Events (enaction only or enaction and imitation) are labelled with numbers, in blue, which correspond with memes in Fig. 6. Enaction and imitation events (blue arrows) are labelled with the quality-of-imitation score. Meme group A comprises memes (2, 3, 4, 8, 9 and 13) and meme group B is memes (6, 7, 11 and 12); memes are shown and labelled in Fig. 6.

why the emergent figure of eight movement pattern has become dominant.

Of course this particular meme evolution is highly contingent. The emergence of the same kind (‘species’) of dominant ‘figure of eight’ movement memes is most unlikely to happen again (and indeed, in repeat trials, has not). But this is exactly what one would expect of an embodied evolutionary process. Perhaps what is surprising is that in an open-ended evolutionary system one kind of meme becomes dominant (at least in this particular trial) – but this is simply explained by the fact that if there is a group of closely related memes in the robots’ memories (because of high-fidelity learning) then our equal probability random selection operator is more likely to select one of these. Note also just how important the initial few imitation events are to the later evolution of the system; the happenstance initial sequence of a poor-fidelity imitation event followed by a high-fidelity imitation event strongly (although not completely) determined the later evolutionary course of our trial system. Again this is strongly characteristic of an evolutionary system.

## 7 Discussion and further work

This paper has described a multi-robot laboratory for experiments in embodied imitation and memetic evolution. The project, of which this work forms a part is, at the time of writing, ongoing and it would be premature to draw any general conclusions with regard to the project aims of illuminating the processes and mechanisms for the emergence of artificial traditions across a robot collective. However, we can at this stage claim that embodied imitation does indeed give rise to meme variation ‘for free’, in the sense that those variations arise from both embodiment and the process of estimation inherent in solving the correspondence problem. The initial trials described in this paper demonstrate promising memetic evolution.

We can also conclude that the choice of selection operator is critical to our aim of demonstrating the emergence and evolution of stable new behavioural traditions. Our experimental trials clearly show that if a robot always copies the last thing it learned we have too much variation, as one might expect. On the other hand with a random selection operator weighted to favour the most frequently seen memes we have too little variation.

The simpler random selection with equal probability operator appears to give the best balance between too much and too little variation.

Further work will:

- run much longer trials to investigate the dynamics, over time, of dominant meme-groups and convergent evolution;
- explore the relationship between embodiment, including sensor characteristics and robot morphology, and the fidelity of imitation and the nature of variation, and address questions relating to the stability of meme transmission;
- introduce additional low-level robot behaviours and thus associate imitation with behaviours that have utility;
- extend the imitation algorithm to implement the imitation of interaction so that interactions between robots can be imitated and propagated across the collective, with richer ‘social learning’;
- further explore the mechanisms of meme selection together with environmental variation, in order to model the spatial and temporal dynamics of meme propagation across the robot collective and the possibility of the emergence of artificial traditions.

**Acknowledgements** This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC), grant reference EP/E062083/1 and the authors gratefully acknowledge project co-investigators and researchers. We additionally wish to acknowledge Wenguo Liu, designer of both the e-puck Linux extension board and the Stage tool which generated the attached movies, and Jean-Charles Antonioli, developer of the data logging system. We are grateful to Susan Blackmore for suggesting the random selection operator of section 6.3. Finally we are grateful to the anonymous reviewers for their detailed and helpful comments.

*Open science data:* all datasets from the experimental trials in this paper may be downloaded from the project web pages: <http://sites.google.com/site/artcultproject/>

## References

1. Alissandrakis, A., Nehaniv, C., Dautenhahn, K.: Imitation with ALICE: Learning to imitate corresponding actions across dissimilar embodiments. *IEEE Trans. Systems, Man and Cybernetics - Part A: Systems and Humans* **32**(4), 482–496 (2002)
2. Alissandrakis, A., Nehaniv, C., Dautenhahn, K.: Solving the correspondence problem in robotic imitation across embodiments: synchrony, perception and culture in artefacts. In: C. Nehaniv, K. Dautenhahn (eds.) *Imitation and Social Learning in Robots, Humans and Animals*, chap. 12, pp. 249–273. Cambridge University Press (2007)
3. Alissandrakis, A., Nehaniv, C.L., Dautenhahn, K.: Synchrony and perception in robotic imitation across embodiments. In: *IEEE Int. Symp. on Computational Intelligence in Robotics and Automation*, pp. 923–930 (2003)
4. Billard, A., Hayes, G.: DRAMA, a connectionist architecture for control and learning in autonomous robots. *Adaptive Behaviour* **7**(1), 35–63 (1999)
5. Blackmore, S.: *The Meme Machine*. Oxford University Press (1999)
6. Dautenhahn, K.: Getting to know each other - artificial social intelligence for autonomous robots. *Robotics and Autonomous Systems* **16**, 333–356 (1995)
7. Dawkins, R.: *The Selfish Gene*. Oxford University Press (1976)
8. Demiris, Y., Hayes, G.: Imitation as a dual-route process featuring predictive and learning components: a biologically-plausible computational model. In: K. Dautenhahn, C. Nehaniv (eds.) *Imitation in Animals and Artefacts*, chap. 13. MIT Press (2002)
9. Demiris, Y., Khadhoury, B.: Hierarchical, attentive, multiple models for execution and recognition (HAMMER). In: *Proc. ICRA 2005* (2005)
10. Gerkey, B.P., Vaughan, R.T., Stoy, K., Howard, A., Sukhatme, G.S., Mataric, M.J.: Most valuable Player: A robot device server for distributed control. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1226–1231 (2001)
11. Hayes, G., Demiris, J.: Robot controller using learning by imitation. In: *Proc. 2nd Int. Symp. on Intelligent Robotic Systems* (1994)
12. Johnson, M., Demiris, Y.: Abstraction in recognition to solve the correspondence problem for robot imitation. In: *Proc. Towards Autonomous Robotic Systems (TAROS)*, pp. 63–70 (2004)
13. Liu, W., Winfield, A.F.T.: Open-hardware e-puck Linux extension board for experimental swarm robotics research. *Microprocessors and Microsystems* **35**(1) (2011)
14. Mataric, M.J.: Getting humanoids to move and imitate. *IEEE Intelligent Systems and their Applications* **15**(4), 18–24 (2000)
15. Meuth, R., Lim, M.H., Ong, Y.S., Wunsch, D.: A proposition on memes and meta-memes in computing for higher-order learning. *Memetic Computing* **1**, 85–100 (2009)
16. Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klapotocz, A., Magnenat, S., Zufferey, J.C., Floreano, D., Martinoli, A.: The e-puck, a robot designed for education in engineering. In: *Proc. 9th Conference on Autonomous Robot Systems and Competitions*, pp. 59–65 (2009)
17. Nehaniv, C., Dautenhahn, K. (eds.): *Imitation and Social Learning in Robots, Humans and Animals*. Cambridge University Press (2007)
18. Nehaniv, C.L., Dautenhahn, K.: The correspondence problem. In: *Imitation in animals and artifacts*, pp. 41–61. MIT Press, Cambridge, MA, USA (2002)
19. Nehaniv, C.L., Dautenhahn, K. (eds.): *Imitation in Animals and Artefacts*. MIT Press (2002)
20. Scassellati, B.: Knowing what to imitate and knowing when you succeed. In: *Proc. AISB’99 Symposium on Imitation in Animals and Artefacts*, pp. 105–113 (1999)
21. Vaughan, R.: Massively multi-robot simulation in Stage. *Swarm Intelligence* **2**(2–4), 189–208 (2008)
22. Winfield, A.F.T., Erbas, M.D.: Screen recording showing animation of fig. 3 (2010). URL <http://sites.google.com/site/artcultproject/mc-fig3>
23. Winfield, A.F.T., Erbas, M.D.: Screen recording showing animation of fig. 6 (2010). URL <http://sites.google.com/site/artcultproject/mc-fig6>
24. Winfield, A.F.T., Griffiths, F.: Towards the emergence of artificial culture in collective robot systems. In: P. Levi, S. Kernbach (eds.) *Symbiotic Multi-robot Organisms: Reliability, Adaptability, Evolution*. Springer (2010)