# On the Emergence of Structure in Behaviours Evolved Through Embodied Imitation in a Group of Robots

Mehmet D Erbas[1] and Alan FT Winfield[1]

[1]Bristol Robotics Laboratory, University of the West of England, Bristol, UK
mehmet.erbas@brl.ac.uk

**Abstract**

This paper describes research in which we model social interactions between artificial agents using real robots. We show that variations that arise from embodiment allow certain behaviours, those that are more robust to the processes of embodied imitation, to emerge and evolve during multiple cycles of imitation. We test 3 memory strategies: no memory, limited memory and unlimited memory, and experimental results appear to show that with limited memory, those behaviours are more likely to become dominant within the robots' collective memory.

## Introduction

Social learning, which enables individuals to learn from each other, is a powerful mechanism in social animals, including humans. An important form of social learning is imitation, in which an individual observes and replicates another's actions. Imitation has been widely studied both by biologists and psychologists; biological research on imitation mostly focusses on its adaptive value for the organism, whereas psychologists are largely interested in the function of imitation and the mechanisms in which it plays a part (Zentall, 2001). There is continuing debate on the definition of imitation and whether it is unique to humans but what is not in doubt is that imitation clearly serves an important role in the development of social cognition in humans. For example, Dautenhahn et al reported that human babies are born with the ability to imitate a wide range of behaviours, including mouth opening and tongue protrusion (Dautenhahn et al., 2003). Meltzoff and Moore (Meltzoff and Moore, 1992) stated that human infants use imitation to enrich their understanding of people and their activities. Through imitation, humans are able to become part of a very complex social environment: human society. Imitation has also been seen as an important facet of cultural transmission; Dawkins argued (Dawkins, 1976) that imitation is a prerequisite for the evolution of culture, as it allows transmission of behaviours, with variation, between individuals.

The study of imitation in robotics has received cross-disciplinary attention in recent years. In the context of robotics research, Bakker and Kuniyoshi (Bakker and Kuniyoshi, 1996) defined imitation thus: “Imitation takes place when an agent learns a behaviour from observing the execution of that behaviour by a teacher”. This definition hints at how imitation is implemented and is used in most robotics research. Skill acquisition by human or robot demonstration has been widely investigated ((Scassellati, 1999); (Mataric, 2000)). This approach holds the promise that we may be able to overcome the necessity to program every behaviour a robot may need to perform, as the robot can learn new behaviours through observing demonstrations of those behaviours. However, as stated above, as well as supporting skill transmission between individuals, in human society, imitation has a social dimension, allowing individuals to become part of a social community. Alissandrakis et al. (Alissandrakis et al., 2004) stated that imitation may serve as a stepping stone towards the development of social cognition in artificial agents as it can form social integration with other artificial agents or with humans. Imitation research in robotics might also usefully address the question of how culture emerges and evolves as a novel property in groups of social animals. In (Winfield and Erbas, 2011) we introduce embodied imitation as a method for modelling the emergence of behavioural ‘traditions’ in social agents.

There has been some work examining the social dimension of imitation in robotics. Steels and Kaplan (Steels and Kaplan, 2001) stated that social learning can play a crucial role in initiating a humanoid robot into a linguistic culture. He used methods such as initiating open-ended dialogues among humans and robotic agents, in which social learning could be embedded. Billard (Billard, 1999) claimed that imitation can be used to enhance autonomous robots' learning of communication skills. The sharing of a similar perceptual context between the imitator and demonstrator can create the necessary social context in which language can develop. Billard devised some experiments in which robotic agents were able to learn a proto-language by using imitation to match their environmental perceptions with observed actions. In this paper, we aim to show that by sharing a similar perceptual context, agents involved in multiple cycles of imitation can – in a sense – agree on the structure of the information that can best be transferred by imitation (that is, what can be imitated). Multiple robots are programmed to observe and imitate each other's movement patterns and the imitated behaviours undergo multiple cycles of copying, in which they mutate because of noise and uncertainties in the real robots' sensors and actuators. We observe that some movement patterns, which can be imitated with high fidelity, emerge and evolve in the group of real robots.

Alisandrakis et al. (Alissandrakis et al., 2004) developed the ALICE architecture (Action Learning via Imitation between Corresponding Embodiments) to address the problem of imitation between dissimilar embodiments. They examined the rules of synchronisation, looseness of perceptual matching and proprioceptive matching in a series of experiments in which robotic arms with variably-sized and numbered joints imitate each other. They showed that patterns can be transmitted between simulated robotic arms and variations occur during these replications because of heterogeneities between the arms. They argue that these variations provide the evolutionary substrate for culture, as new behavioural patterns may emerge and be transferred between agents. In this paper, we describe a series of experiments in which real robots observe and imitate each other's movement patterns. We show that even in an homogeneous group of real robots, variations occur during the imitation process that allow certain behavioural patterns to emerge and evolve during multiple cycles of imitation. These evolved behaviours can be copied with higher fidelity, as they are more robust to uncertainties in the real robots' sensors and actuators.

## Imitation in Robots

As stated above, we have used real robots to model the social interactions between agents. The motivation for using real robots rather than simulated agents or biological social entities for modelling is:

- Real robots, with their less than perfect perception and actuators, provide natural variations in the imitation process which allow new behaviours to emerge and evolve. Using simulated agents in a simulated environment, we would have to control the degree and types of heterogeneities and noise, but this may preclude any emergent processes that are a part of imitation; the level of emergence in a simulated environment would be limited to the level of variance that is artificially introduced.
- Data about the imitative activity, including the internal data and calculations of the robots, can easily be extracted and examined. This would not be the case if biological social entities (for example, people or monkeys) were used.
- The implementation of imitation on real hardware makes clear how theoretical assumptions and hypotheses regarding imitation can be operationalised.

Figure 1: A Linux-extended e-puck robot. The robots are fitted with coloured skirts, to enable them to 'see' each other. The yellow hat on top of the robot provides a matrix of pins holding unique patterns of reflective markers that allow the tracking system to identify and track each robot.

### Hardware Setup

The artificial agents used to model social interactions are *e-puck* miniature robots (Mondada et al., 2009), 7 cm in diameter and 5 cm in height. They are equipped with 2 stepper motors, two wheels of 41 mm diameter, 8 proximity sensors, a CMOS image sensor, an accelerometer, a microphone, a speaker and a ring of coloured LEDs. Their on-board battery provides 3 hours of autonomy. The robots are enhanced with a Linux extension board (Liu and Winfield, 2011) based on the 32-bit ARM9 micro-controller with the Debian/Linux system installed. The board has a USB extension port, used to connect a wireless network card, and is equipped with a MicroSD card slot. These additions to the standard e-puck robot offer increased processing power and increased memory. The robots are also fitted with coloured 'skirts' to enable them to see each other using their built-in image sensors. The experiments are performed in an arena measuring 3 m x 3 m. A vision-tracking system provides high-precision position tracking and a dedicated swarm server combines the data from the tracking system and the internal data from robots for later analysis. Each robot is also fitted with a tracking 'hat' which provides a matrix of pins holding unique patterns of reflective markers that allow the tracking system to uniquely identify and track each robot (Fig. 1).

### Movement Imitation Algorithm

In this research, a robot-to-robot movement imitation algorithm is implemented on the Linux extended e-puck robots.

Each robot is able to track and copy the other robot's movement patterns. Since we are interested in embodied imitation, the algorithm completely depends on the visual data coming from the image sensor of the robots; no other type of communication is allowed between the robots.

There are 3 main stages in the imitation algorithm:

- Frame processing: While observing captured visual frames, the observing robot tracks the movement of the demonstrator robot. As stated above, the robots are fitted with coloured skirts; by determining the size and location of the skirt on the demonstrator robot, the observing robot estimates the relative position of the demonstrator and stores this information in a linked list of positions. In this way, up to 5 frames per second are processed.

- Data processing: After the demonstrator's movement pattern is completed, the observer robot processes the linked list of positions using a regression line-fitting method to convert the estimated positions into straight line segments.

- Pattern replication: The straight line segments and their intersections are converted into a sequence of motor commands (moves and turns).

In this way, the observing robot replicates the pattern demonstrated by the demonstrator robot.

### Quality of Imitation

To quantitatively assess the fidelity of imitation (that is, the similarity between the original movement pattern and its copy), a quality of imitation function needs to be defined. Since each movement pattern consists of straight moves and turns, there are 3 components to each pattern that can be copied: the number of segments (straight moves), the length of each move and the angle (turn) between each consecutive move. Therefore, the overall quality of a copy can be calculated by separately estimating 3 quality indicators. The quality of move length, $Q_l$, between the original path $O$ and its copy $C$ is calculated as follows:

$$Q_l = 1 - \frac{\sum_m |l_m^O - l_m^C|}{\sum_m l_m^C} \tag{1}$$

where $l_m$ is the length of move $m$ that is to be compared. Here, the ratio is calculated of move length differences between the original pattern and its copy and the total move length of the copy. If the original movement pattern and its copy have different numbers of segments, $N^O$ and $N^C$ respectively, the sum is calculated only over the number of segments in the smaller: $min(N^O, N^C)$. The quality of angle (turn) imitation similarly calculated as:

$$Q_a = 1 - \frac{\sum_m |a_m^O - a_m^C|}{\sum_m a_m^O} \tag{2}$$

where $a_m$ is the turn angle following the move $m$. The quality of segment imitation simply compares the difference between the number of segments of the original pattern and its copy. It is calculated as:

$$Q_s = 1 - \frac{|N^C - N^O|}{N^O} \tag{3}$$

where $N^O$ and $N^C$ are the number of segments of the original path and its copy. The overall quality of imitation, $Q_i$, is a combination of 3 quality indicators:

$$Q_i = \frac{LQ_l + AQ_a + SQ_s}{L + A + S} \tag{4}$$

where $L$, $A$ and $S$ are weighting coefficients.

To test the quality of imitation, a demonstrator robot is programmed to follow a sequence of straight line moves and turns that describes an equilateral triangle, while an imitator robot watches. Then, the imitator robot performs its copy of the demonstrator's pattern (Fig. 2). By comparing these two patterns, the quality of imitation is determined. The same scenario is repeated multiple times, with different distances between the robots. As shown in the figure, the best quality is achieved when the distance between robots is 1 m (Fig. 3). When the distance between robots is increased (to 1.5 m or more) the quality of imitation starts to degrade. This arises because the relative positional changes are estimated based on the size and location of the observed robot in the field of vision of the imitator robot. When the separation distance increases, the positional changes are harder to detect, as they cause smaller variations in the image of the observed robot. On the other hand, when the distance between robots is small (that is, 0.5 m or less), the demonstrator robot leaves the field of vision of the imitator robot many times, forcing the imitator robot to rotate itself each time and thus it may miss some turns of the demonstrator robot's trajectory while it is busy. Therefore, we have a separation range, between 0.5 m and 1 m, that is optimal for our vision based embodied imitation algorithm.

## Experiments

The notion of an imitation experiment is introduced to examine the effects of multiple cycles of imitation on the structure of the movement patterns that are being copied. During these experiments, 4 robots are placed in the arena, 1 m apart from each other (Fig. 4). Robots interact by copying each others' movement patterns using the imitation algorithm outlined in the previous section. Robots can be in one of two modes during the experiments: demonstrator or observer. When a robot enters demonstrator mode, it turns its LEDS on for 35 seconds to signal that it will start to demonstrate a movement pattern. During this period the demonstrator tries to grab the attention of one (or more) other robots. After that, the demonstrator robot turns its LEDs off and executes a movement pattern that consists

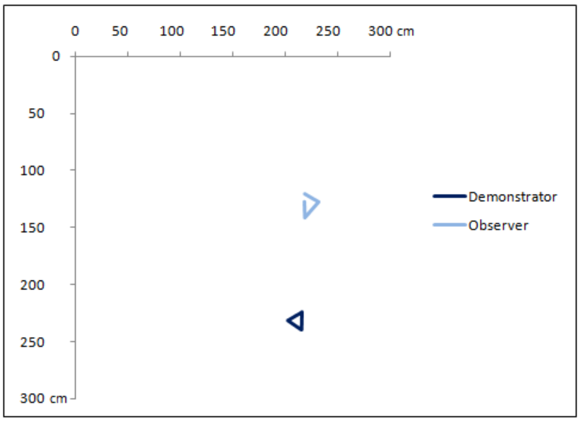


Figure 2: Plot of the trajectory of robots during an imitation run. The demonstrator robot moved in an equilateral triangular trajectory which was then copied by the observer robot. The robots were placed 1m apart.

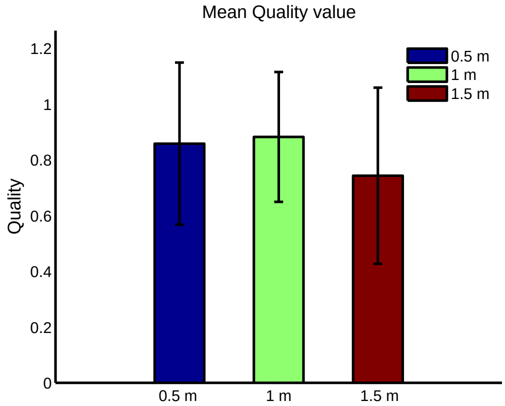


Figure 3: Mean quality of imitation ($Q_i$) value with 95% confidence intervals calculated at different distances between robots. Each bar shows mean quality value over 20 cycles of imitation in which an equilateral triangle (each side 15 cm) movement pattern described by the demonstrator robot is copied by an imitator robot. For quality of imitation calculation, each quality indicator was given equal weight: $L = A = S = 1$.

of straight-line moves and turns. When execution is complete, the demonstrator robot blinks its LEDs for one second to signal ‘finish’. Then the demonstrator robot returns to its original start position and enters the observer mode. When a robot enters observer mode, it searches for a start signal by scanning the arena while rotating itself. When it detects a start signal, it focuses its attention on the demonstrator robot and waits for the demonstration to start. After completion of the demonstration, the observer robot records what it has observed and enters demonstrator mode. The finite state machine of the controller of the robots is shown in Fig. 5. At the start of the experiment, two of the robots are in demonstrator mode (Robots A and B) while the other two are in observer mode (Robots C and D).The experiment is left free-running as the robots change roles while imitating each other. All internal calculations and movement patterns of the robots are recorded for later analysis.

## Imitation with no memory

In the first set of experiments, the robots are able to remember only the most recent pattern they have observed; any newly-observed pattern replaces the previous one. Robot A is initialised with a square trajectory and Robot B is initialised with a equilateral triangle trajectory. Fig. 6 shows the pattern evolution map of an experiment in which 39 successful imitations were completed in approximately 20 minutes. In the figure, each node represents a pattern. If an arrow exists at a node, this means one of the robots executed that pattern and it was imitated by another robot. The new copy is at the end of the arrow. If the copy is high-quality, ($Q_i >= 0.85$) the the node has a dark colour.

We first observe in this experiment that the original patterns deteriorate very quickly. At the beginning of the run, both robots (C and D), by chance, copied the square trajectory and the triangular trajectory vanished from the experiment. The square trajectory also deteriorates rapidly, as any low quality copy can easily replace it. Because some bad copies missed turns, eventually the robots ended up with a pattern consisting of a single forward move. These low quality copies do not occur often but just one is sufficient to disrupt the evolution of the movement patterns. In this experimental run, all patterns after pattern number 22 consist of one single move without any turns. These single move patterns can be copied with high quality but we still observe some poor copies. We conclude therefore that when robots have no memory, evolution of the movement patterns is acutely sensitive to imitation errors.

## Imitation with unlimited memory

In the second set of experiments robots have unlimited memory so they save all patterns that they have observed. When they enter demonstrator mode robots randomly select, with equal probability, one of the patterns in their memory and demonstrate it. Once again Robots A and D are initialised with a square trajectory and Robots B and C are initialised with an equilateral triangle trajectory. Fig. 7 shows the pattern evolution map for an experiment with this setting. In this run, 55 successful imitations were completed in 30 minutes. We first observe in these experiments that – as we would expect – the original movement patterns are more likely to be preserved (with variation), as each newly-observed pattern is stored in memory. Low quality copies occasionally occur, but as they do not replace previously ob-

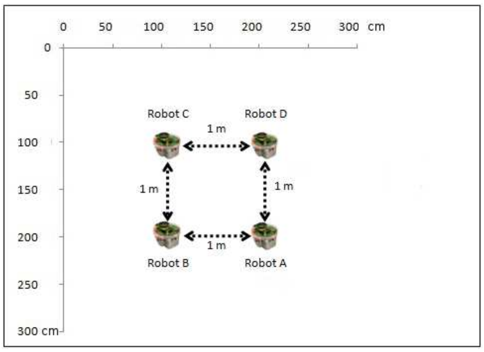


Figure 4: Each experiment presented in this section is performed in a 3 m by 3 m arena with 4 robots, placed 1 m apart and arranged as shown here. In all experiments, Robot A and Robot B are started in demonstrator mode while Robot C and Robot D are started in observer mode.

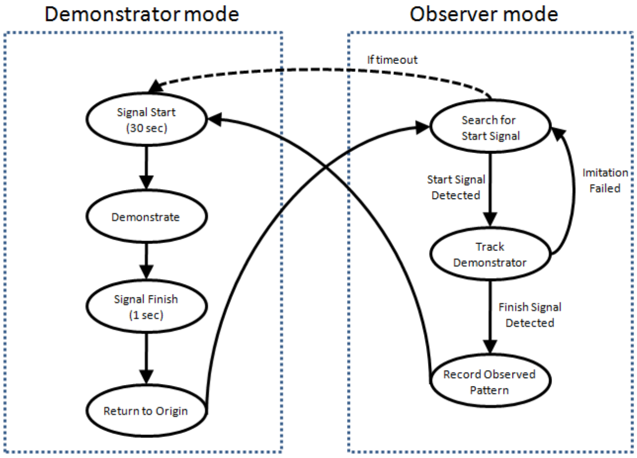


Figure 5: Finite state machine of the controller of the robots. The robots are programmed to copy each others' movement trajectory as they keep changing their roles to demonstrator and observer. To prevent a deadlock with all the robots searching for a start signal, two of the robots (Robot A and Robot B) are programmed to time out and enter '*Signal Start*' state after completing two complete scans of the arena in '*Search for Start Signal*' state (the dashed arrow).

served patterns, these paths cannot easily become dominant. Second, we see that as patterns evolve during multiple cycles of imitation, some paths that are able to be copied with high quality emerge and propagate among robots. In this run, pattern 27 has this property. Fig. 8 shows the evolution of pattern 27. It is a descendant of the original equilateral triangle trajectory, and there are 5 intermediate copies between the original triangle and pattern 27. At each copy, the pattern is modified by the imitating robot. Finally pattern 27 emerges and a sharp increase in quality of imitation can be observed after this point ($Q_i > 0.94$ for all of its descendants). What makes this pattern and its descendants easily copiable? First, short moves are more prone to error, as a small mistake in perception can cause them to vanish; a pattern that can be copied with high quality typically does not include short moves. Second, the length of each move varies at each subsequent copy. Although estimating the relative size and position of the demonstrator robot is straightforward image processing, it is error-prone, because of the relatively low resolution of each robot's image sensor. A move directed towards or away from the observing robot can only be detected if it causes a perceptible change in the size of the demonstrator robot, i.e. a detectable change in number of pixels in the image of the demonstrator. At each copy, the observing robot stores what it infers from the demonstration, as perceived from its relative position and perspective. Therefore, the patterns tend to evolve into ones that can be more easily imitated. Fig. 9 shows pattern 27 and its descendants. As can be seen, there is a high level of similarity between these patterns. At the end of the run, pattern 27 and its descendants form a cluster of similarly-shaped patterns in the robots' memories. Fig. 10 shows the average $Q_i$ value for this experiment in comparison with the average $Q_i$ value for the cluster formed by pattern 27's descendants. As can be seen, although the distance between the robots is 1 m, the average $Q_i$ value is slightly low; around 0.83. This can be explained by the fact that some low quality imitations occur during the evolution of patterns. A sharp increase in $Q_i$ value can be observed after a pattern emerges that is more robust to uncertainties in the robot's sensors and the imitation process: the average $Q_i$ value for the cluster that is formed by the descendants of pattern 27 is 0.96.

## Imitation with Limited Memory

In the previous set of experiments, we showed that certain patterns, those that are more robust to uncertainties in the real robots' sensors and actuators and the estimation process of imitation, can emerge and evolve during multiple cycles of imitation. As these emergent patterns can be copied with high quality, their descendants have similar, inherited characteristics. As a result, clusters of highly copiable patterns are formed in the memories of the robots. These clusters may grow larger with subsequent cycles of imitation if, by chance, members of these clusters are selected for demonstration. We now show that with a limited memory, these emergent patterns and their copies can become dominant. In the third set of experiments, an example run with limited memory is presented, in which an emergent pattern and its highly similar descendants become dominant. Here robots have a limited memory, in which they can store only the most recent 5 patterns observed. When the memory is full and a new pattern observed, the oldest pattern in memory is replaced with the new pattern. Fig. 11 shows the pat-

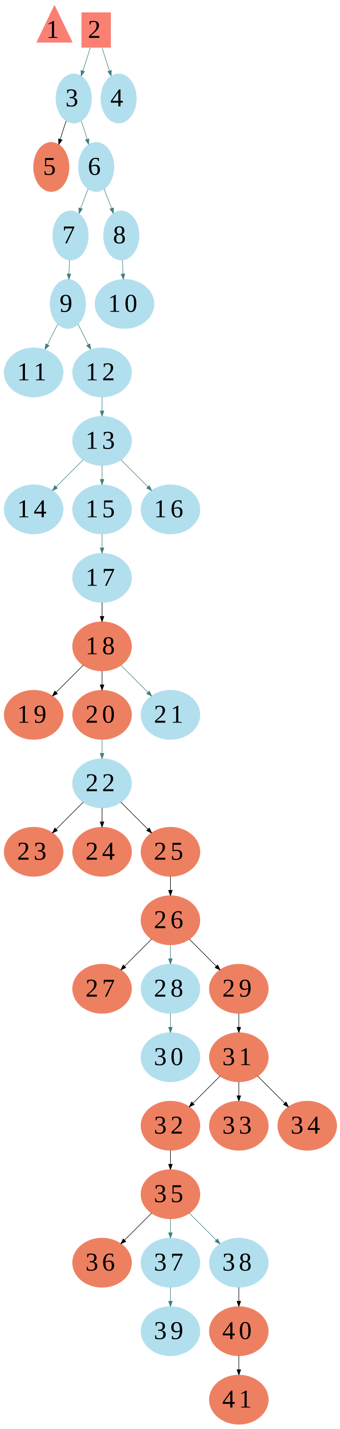


Figure 6: Pattern evolution map for a 4 robot experiment with no memory. Each node in the figure represents the demonstration of a movement pattern. If a pattern is demonstrated and imitated, the new copy of that pattern is linked to it by an arrow. For instance, pattern 2, the original square, was demonstrated by Robot A and was copied by two robots. The new (child) copies of pattern 2 are patterns 3 and 4. If the copy is of high quality (i.e. $Q_i >= 0.85$), then the node has a dark colour.

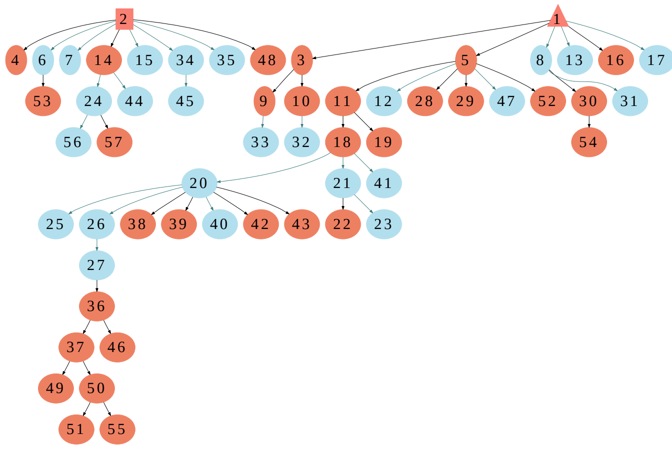


Figure 7: Pattern evolution map for a 4 robot experiment with unlimited memory. Initial movement patterns are a triangle (1) and a square (2).

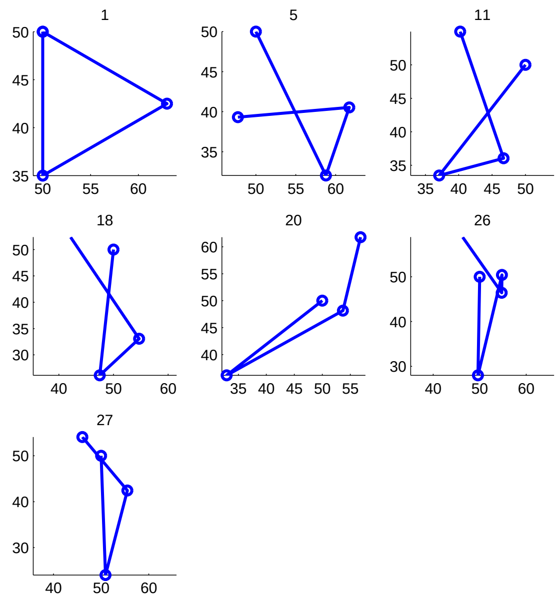


Figure 8: Evolution of pattern 27 in Fig. 7. Pattern 27 is a descendant of the original equilateral triangle pattern. By following the imitation links on the pattern evolution map for this experiment, we can see that there are 5 intermediate copies between the original triangle and pattern 27: the patterns numbered 5, 11, 18, 20, 26. All of these patterns, starting with the original triangle and ending with pattern 27, are shown here in order. All axis are marked in cm.

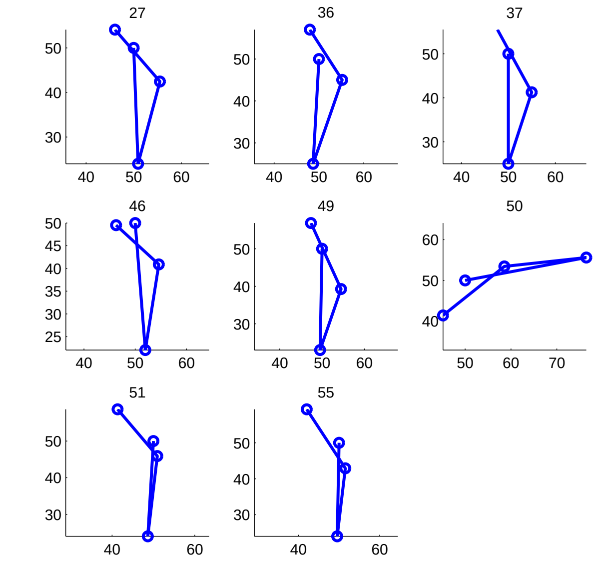


Figure 9: The descendants of pattern 27 in Fig. 7. Starting with pattern 27, its descendants (patterns 27, 36, 37, 46, 49, 50, 51, 55) are shown in order. All axes are marked in cm.

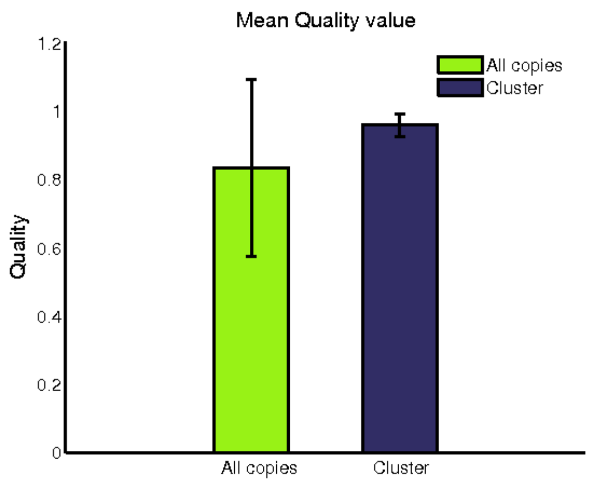


Figure 10: Average $Q_i$ value for all imitation events in the experiment shown in Fig. 7 (All copies) and average $Q_i$ value for the cluster formed by pattern 27's descendants (Cluster), with 95% confidence intervals

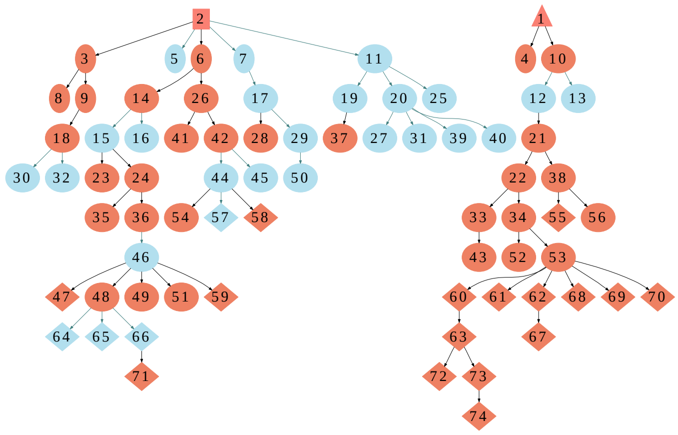

Figure 11: Pattern evolution map for experiment with limited memory. The 20 patterns in the memory of all 4 robots at the end of the experiment are highlighted as diamonds.

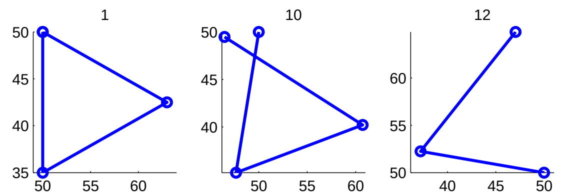


Figure 12: Evolution of pattern 12 in Fig. 11. There is an intermediate copy (10) between the original triangle and pattern 12. All axes are marked in cm.

tern evolution map from an experiment with these settings in which 72 successful imitations were completed in 60 minutes. In this run, a V-shaped pattern, pattern 12, emerged and all of its descendants are high quality copies. Fig. 12 shows the evolution of this path and Fig. 13 shows some of its high-quality descendants. At the end of this run, 12 of the 20 patterns in the memory of all 4 robots are descendants of this pattern. Since the robots randomly choose which pattern to demonstrate, there is now a 60% chance that one of the descendants of pattern 12 will be demonstrated again. Once it is selected and copied, the new copy is itself likely to be a high quality copy and so similar to pattern 12. This process will then increase the percentage of patterns in the memory that are similar to pattern 12. We conclude therefore that with limited memory, patterns robust to uncertainties that emerge are more likely to become dominant.

## Conclusion and Discussion

In this work, we have used real robots to model social interactions between artificial agents, in particular learning by imitation. We have shown that variations in the real robots' sensors allow certain behaviours to emerge and evolve during multiple cycles of imitation. These evolved movement

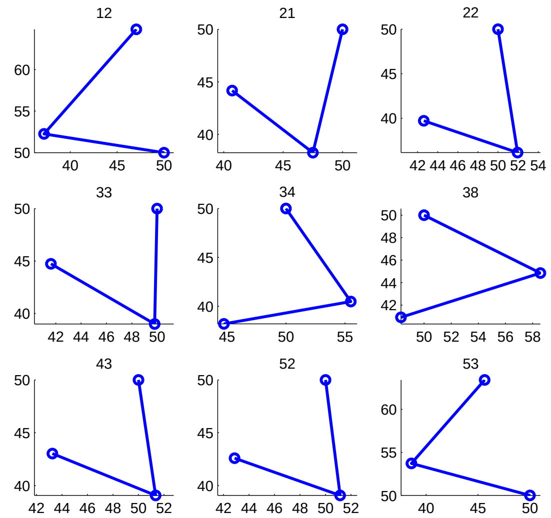


Figure 13: The descendants of pattern 12 in Fig. 11. Starting with pattern 12, some of its high-quality copy descendants (patterns 12, 21, 22, 33, 34, 38, 43, 52, 53) are shown in order. All axes are marked in cm.

patterns are more robust to the uncertainties of the real robot embodied imitation process and so they can be imitated with high fidelity. As the robots share a similar perceptual context and embodiment, they are able to – in effect – agree on the structure of the movement patterns that can be transferred between them.

We have experimentally tested three cases with different sizes of robot memories: no memory, unlimited memory and limited memory, in order to test the hypothesis that memory size will effect the likelihood of dominant movement patterns emerging. In the no memory case, the evolution of movement patterns is extremely sensitive to any instance of poor quality imitation, which means that the original movement patterns very quickly deteriorate. In the unlimited memory case, patterns emerge that can be easily copied but are less likely to then become dominant, as the number of patterns in the robots' collective memory grows larger with each new imitation cycle. However, in the case with limited memory, these evolved patterns can become dominant if they and their descendants are, by chance, chosen for demonstration. For simplicity of analysis, we have carried out our limited memory experiments with a small memory size (5 patterns per robot). We conjecture that with a larger (but still limited) memory, multiple patterns that can be imitated with high fidelity can emerge and form clusters of similarly-shaped patterns in the robots' collective memory. In this way, the robots can collective evolve an ensemble of patterns that can be copied between them with high fidelity. Here the imitated patterns are not linked to a task or an environmental context. However it seems possible, and testable using the embodied approach outlined in the paper, that associating imitation with behaviours that have utility could lead to the emergence of non-verbal communication between robots.

## Acknowledgements

This work is supported by EPSRC research grant: EP/E062083/1.

## References

Alissandrakis, A., Nehaniv, C. L., and Dautenhahn, K. (2004). Towards robot cultures? learning to imitate in a robotic arm test-bed with dissimilar embodied agents. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, 5(1):3–44.

Bakker, P. and Kuniyoshi, Y. (1996). Robot see, robot do : An overview of robot imitation. In *AISB96 Workshop on Learning in Robots and Animals*, pages 3–11.

Billard, A. (1999). Imitation: A means to enhance learning of a synthetic proto-language in an autonomous robot. In *Imitation in Animals and Artifacs*, pages 281–311. MIT Press.

Dautenhahn, K., Nehaniv, C. L., and Alissandrakis, A. (2003). Learning by experience from others - social learning and imitation in animals and robots. In Kühn, R., Menzel, R., Menzel, W., Ratsch, U., Richter, M. M., and Stamatescu, I. O., editors, *Adaptivity and Learning: An Interdisciplinary Debate*, pages 217–241. Springer Verlag.

Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.

Liu, W. and Winfield, A. F. T. (2011). Open-hardware e-puck linux extension board for experimental swarm robotics research. *Microprocessors and Microsystems*, 35(1).

Mataric, M. J. (2000). Getting humanoids to move and imitate. *IEEE Intelligent Systems and their Applications*, 15(4):18–24.

Meltzoff, A. N. and Moore, M. K. (1992). Early imitation within a functional framework: The importance of person identity, movement, and development. *Infant Behavior and Development*, 15(4):479–505.

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J. C., Floreano, D., and Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In *9th Conference on Autonomous Robot Systems and Competitions*, page 59:65.

Scassellati, B. (1999). Knowing what to imitate and knowing when you succeed. In *AISB'99 Symposium on Imitation in Animals and Artefacts*, pages 105–113.

Steels, L. and Kaplan, F. (2001). Aibo's first words: The social learning of language and meaning. *Evolution of Communication*, 4(1):3–32.

Winfield, A. F. T. and Erbas, M. D. (2011). On embodied memetic evolution and the emergence of behavioural traditions in robots. *Accepted for Journal of Memetic Computing*.

Zentall, T. R. (2001). Imitation in animals: Evidence, function, and mechanisms. *Cybernetics and Systems*, 32:63–96.