

An Experimental Comparison of Imitation Paradigms used in Social Robotics

Joe Saunders, Chrystopher L. Nehaniv, Kerstin Dautenhahn
Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire
College Lane, Hatfield Herts AL10 9AB, United Kingdom
{J.2.Saunders, C.L.Nehaniv, K.Dautenhahn}@herts.ac.uk

Abstract

We study and contrast particular issues arising in two social learning paradigms that are widely used in robotics research: (i) following or matched-dependent behaviour and (ii) static observational learning. Experiments are carried out with physical Khepera robots whose controllers include motor schemas and new neural network based methods for model agent-centred perception of angle and distance. The robots are trained to perceive the dynamic movement of a human or robot demonstrator carrying a light source. The robots learn the behaviour either through perception from a static location or while following. The differences and implications of the results of both the following and observation mechanisms are compared and contrasted.

1 Social Learning Paradigms

Psychologists consider imitation¹ to be one of the key elements in social learning, with social learning leading to the acceleration of the acquisition of intelligent behaviour [24, 8, 7]. If robots could imitate either each other or humans, both singly and in groups, then robot task acquisition would be easier and faster with the possibility of increased behavioural complexity and ultimately some form of cultural transmission [1]. In fact we consider this social dimension to be the key to making robots behave more intelligently [6], an approach inspired from studies of social animals (e.g. apes) and the ‘social intelligence hypothesis’ [5], which proposes social origins for primate intelligence.

In our research program we investigate social learning and interaction between both human/robot and robot/robot pairs to understand the social dimension of imitative behaviour. One part of this research is the need to consider the perspective of both the imitator and the imitatee and study the problems of perception and action encountered by both. In this

¹We take Thorndike’s 1898 classical definition of imitation [22] as “leaning how to do something by seeing it done” but extended to non-biological agents [15].

paper we discuss two imitator perspectives which are widely applied paradigms in robotics imitation research, *following* behaviour [10, 4, 6] and *observation* behaviour [12, 9, 3, 21, 13, 1].

1.1 Following and Observation

From a psychological/ethological viewpoint *following* is more rightly considered as *matched-dependent behaviour* [24]. For example rats can be trained to follow a lead rat through a maze which they then learn to navigate [14]. The rats may have no idea of intentionality of the lead rat and can be trained to follow any other salient (including non-animal) stimuli.

Likewise, *static observation* can be defined via *observational learning*. Here the behaviour of the demonstrator is copied after they are observed carrying it out. Typically the demonstrator and imitator operate within a shared context but at a distance from one another. For example Norway Rats apparently develop food preferences by smelling the breath of a conspecific [8], without reference as to whether the demonstrating rat becomes ill or dies. The mechanism is very simple - perceive what others do and imitate it [19].

To set the scene from a human-human perspective consider the following example: If I were to ask you to imitate my hand movement, and I then drew a triangle in space, it is very likely that you would very easily also draw a triangle. When imitating my movements you would have not attempted to follow my hand with your hand as I was making the triangle shape; you would have probably observed my movements and when you felt that you knew how to reproduce the shape, started your own movement sequence.

Interestingly, a more complex set of movements may result in a mixed approach that combines following and observation. For example, assume initially both hands are palm to palm, then I make a clockwise circle with one hand and an anti-clockwise circle with the other. If I then asked you to imitate my

movements you might initially consider this trivial. However most people when first seeing this demonstration will typically make clockwise circles with both hands and assume that they are correctly imitating the demonstrator. When it is pointed out that an error has been made, the demonstrator will usually repeat the demonstration by slowing down and the imitator will usually start to track each movement by following the demonstrator’s hands. In both of the examples above the demonstrator and imitator will consider the imitation successful if the resulting pattern resembles the original; very fine matching of movements are not considered.

These examples signal some interesting but not widely researched features of imitative behaviour in the relationship between static observation of, and active participation in, an event to be imitated. In the triangle example there is observation followed by imitation, in the second example, observation is usually insufficient to completely duplicate the task and further mechanisms are employed – the demonstrator effectively slowing down and teaching the imitator, and the imitator using a following mechanism to capture the movements.

These different responses highlight a number of issues. Firstly that animals and humans may use a range of mechanisms to imitate, in the two cases above: static observation followed by an imitative response, or an immediate following mechanism in conjunction with coaching from the demonstrator. Secondly, using only static observation seems to be possible in both cases, however in the second case the resulting assemblage of actions tends to be incorrect. Finally the imitator will usually only revert to using the following mechanism as a way to understand and track the movements when observation has failed.

1.2 Paradigm Examples in Robotics

Both the following and observational mechanisms are used in robotics imitation research. For example, [4, 10] used following to replicate a demonstrator’s actions or to improve an imitator’s learning capability. Following is successful since it allows the imitator to more closely share context and experience as the demonstrator performs the behaviour. As the imitator follows it can map its sensory experiences directly to its motor outputs, which are matched to, and depend on, the demonstrator’s actions. It can thus learn the necessary perception-action couplings directly and use them in similar situations in the future without the teacher being present. Following is used in [9] as a means of testing an imitation model based on reducing perceptual errors, the imitative behaviour be-

ing achieved by continually adjusting motor outputs when presented with a difference between perceived states and goal states.

The observational paradigm is also extensively used, often employing complex vision processing (e.g. [12]). Both physical and simulated robots are used in [3] to statically observe a continuously changing game (air hockey or marble maze). Observation is used in [20] to recognise what to imitate in attempts to build robots that imitate people.

From the imitator robot’s viewpoint these are very different perspectives and it is highly likely that improved robot-human interaction may be possible if the characteristics and key issues involved in these learning paradigms were better understood. We begin this process by offering some observations comparing the social learning paradigms in an implementation study with Khepera robots.

2 Framework

An imitation framework for describing the sequences of *actions*, *states* and/or *effects* can be described using the *correspondence problem* formalisation [17].² The correspondence problem formalism allows the model’s given action/state/effect and the desired action/state/effect in the imitation sequence to be compared using metrics or measures of dissimilarity; the number of matched states/actions/effects, i.e. the fineness of the imitation attempt can be described as a measure of granularity.

Our focus is to compare the following and observational paradigms. To do this we simplify the context of the imitation and restrict the actions/states and effects. In the experiments we use a physical Khepera-1 robot as the imitator. The model is perceived by means of a bright light. This light can be mounted on another Khepera or alternatively manipulated by a person. The goal for the imitator is simply to replicate movements of the model. Actions are limited to two possibilities: either turn by a given angle and/or move a given distance. States are described as the perceived vector from the centre of the robot to the light. For example figure 1 shows the state/action sequence derived if the model’s actions were to describe a triangle.

Two controllers were designed to investigate the relationships between imitation of movements using static observations and imitation of movements by following. All experiments were carried out in real-time on physical robots (i.e. simulation was not used) on a desktop in a typical busy academic environment with

²Note that in order to remove the intentionality from the term *demonstrator* it is normally replaced with the term *model*.

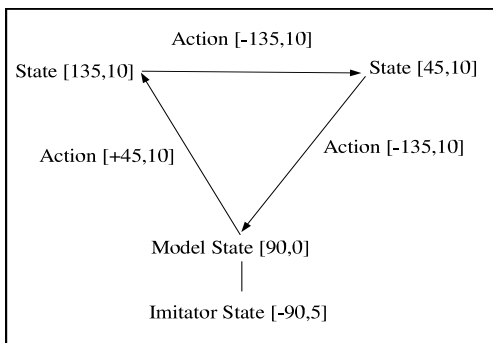


Figure 1: *States and Actions from Model's Initial Perspective. States are specified by (degrees θ , distance r). The triangle shows each state achieved by executing each action. The model is initially facing forward at 90° . The imitator is shown placed behind the model at -90° . Action can be interpreted as the necessary movement ($\Delta\theta, \Delta r$) to achieve the given state (θ, r).*

light levels varying during the day. We examine the behaviour of the imitator when imitating triangles, circles and letter shape patterns.

2.1 Controller for a Following Robot

We used a ‘Motor Schema Vector Fields’ methodology [2] to facilitate following behaviour. In this approach following the model is achieved with an *attractive* vector³, robot movement is achieved with a *uniform* vector⁴ and obstacle avoidance achieved with a *repulsive* vector.⁵ Each vector has two elements, angle and magnitude that are functions of sensor readings. The sum of the vectors is computed locally from the agent perspective at each time step and result in a single vector. The angle of this resultant vector is then used to control the orientation of the robot with the vector magnitude converted to wheel speeds of the robot. Thus as the robot observes a moving light source the perceptions afforded by the sensors are converted to motor outputs.

The attraction vector is designed so that the robot will orient itself towards the light bulb, the uniform vector provides a forward movement for the robot and the repulsive vector ensures that the robot does not

³ $(\theta_1, r_1) = (\theta_0, 1 - \frac{r_0}{k})$, where

$$(\theta_0, r_0) = \sum_{i=1}^n (LightSensorAngle(i), \frac{LightSensorReading(i)}{LightSensorMax(i)})$$

⁴ $(\theta_2, r_2) = (90^\circ, gain)$

⁵ $(\theta_3, r_3) = \sum_{i=1}^n (-IRSensorAngle(i), \frac{IRSensorReading(i)}{IRSensorMax(i)})$

where summation is polar, n is the number of Sensors, k is set to n times the maximum values of the light sensors, $gain$ is the magnitude applied to the uniform vector to facilitate forward movements and sensor angles are with respect to the imitator's body and hence fixed.

bump into the model (either the human hand or the other robot). This summation effectively makes the robot follow the light and keep a fixed distance from the model.

2.2 Controller for a Statically Observing Robot

The observation experiment is designed so that the imitator does not move, but learns as it observes the model. It achieves this by computing both the distance and angle from itself to the moving model and stores these observation points as a list of two element vectors. Prior to observing, the robot must first learn how to measure angles and distance.⁶

Learning to Measure Angles. The robot is first trained to accurately compute the angle of the light (held by a person or a robot) from the centre of the imitator. A number of methods were evaluated including using a *light compass* [18], or computing the angle by using *vector summation* of the inputs to each of the light sensors [2]. However both of these methods were not accurate and suffered from incorrect readings especially when none of the robot's sensors were directly facing the light. A new method, which we call *environmental sampling*, was grounded in sensors and to some extent nearer to a biological solution: the robot is allowed to learn about light angles simply by observing them. As the Khepera is a circular robot it rotates in a circle in the presence of the model. It detects when the circle is complete by polling its wheel encoders and stopping when the appropriate value has been exceeded. (During the turn it reads its light sensors every 200ms. A robot turning at 8mm/s would typically poll its sensors 65 times.) As the speed of the turn is constant the time interval between readings can thus be converted to an angle. Each of the sensor readings are then normalised. This has two effects, firstly that of making distant readings of angle equivalent to closer readings, and secondly allowing these values to be loaded directly as weights into a neural network (a counter-propagation network [11]). This is a fully connected feed-forward three layer network. The first layer being the normalised input of the 8 light sensors, the number of middle layer neurons is set to the number of times the robot was able to poll its sensors and the third layer outputting the conversion of these values to angles. Using this technique has a number of advantages. Firstly that the

⁶Clearly a more sophisticated technology, e.g. laser scanners, may have provided an alternative method to calculate both angle and distance to the model. However our approach, of using minimal methods and limited sensors, shows that imitation, perceptual matching and shared context can be achieved without sophisticated sensory hardware.

network can be built as the environment is observed, secondly there are no training steps, thirdly the size of the network is directly related to the internal rotation speed, sensor modality and sensor polling time of this particular robot and finally that the method is partially resilient to sensor failure.

Learning to Measure Distance. A method based on *environmental sampling* (which proved superior to *triangulation*) was used for the angle computation, the light sensors being summed as vectors as the robot turned. This exploited the fact that sensors directly facing the light would have a larger effect on the vector magnitude than those further away. The robot was trained by rotating at increasing 1cm distances from the light source. The vector magnitude was then held in a lookup table indexed by angle and distance.

Following these procedures the robot can compute both angle and distance without further training.

Observing Angles. After the learning phase is complete the network operates by feeding a normalised sensor vector to the input layer and receiving the angle from the output layer. The network is thus operating as a pattern matching mechanism. There are some biological observations which may show similar (though not equivalent) mechanisms in animals. For example young bees appear to record the image of their hive from many angles and positions around it: they fly in and out of the hive varying their circular flight path each time [16].

Observing Distance. During the observation phase the angle is computed, followed by magnitude of the vector summation⁷, the two values providing the key to the lookup table to yield distance. This mechanism allowed reasonably accurate distance measurements between around 10° to 170° in front of the robot to a distance of around 30cm from the light source.

During the observation experiment the robot first collects a set of angles/distances from itself to the model which are then time averaged. Critical changes based on density of the observed points are extracted and implied velocities are calculated. The imitator uses the critical points and implied velocities to then imitate the model’s trajectories and velocities.

3 Results

In our initial experiments we compared behaviour on three very simple imitations. These were a triangle, circle and the letter T. The triangle was chosen because of the sharp changes of direction at each vertex, the circle because of its continuous shape and the

letter T because of the need to reverse direction and remap the shape. We emphasise that our goal was not to design robots that perfectly imitate geometric shapes but rather investigate relevant aspects of the imitation attempt in the two learning paradigms.

3.1 Following

Observations of robot behaviour using the ‘following’ controller highlighted some issues with using a simple reactive architecture, one of which is a familiar problem for path following mobile robots [23, 9]. This is where the robot fails to follow the path of the imitator with precision due to the tight reactive cycle between the sensors and the motors. In our experiments this is shown in figure 2: for the triangle and circle the robot either cut the triangle corners or inscribed a smaller circle inside the model circle. The T-Shape exhibited this cutting problem and also showed the imitator *avoiding the model* when the model started to reverse towards it when inscribing the T-shape. These problems arise due to what we call *impersistence*, i.e. the inability to sustain appropriate actions, and is a consequence of the imitator always reacting to the latest perception vector whilst effectively ‘forgetting’ the previous perceptions. Also the reaction of the imitator’s motors always lags behind that of the sensors (for example the imitator may have partly turned towards the model when a new reactive cycle begins).




Figure 2: *Following Behaviours for triangle, circle and T shape. The graphics represent typical behaviours of the following mechanism. Dotted line is the model, bold line is the imitator.*

The *impersistence problem* in robotics appears to be unsolved, but partial solutions have been attempted by providing further dampening to the system, by delaying the imitator response, slowing down the model or alternatively using ‘vector’ or ‘pure’ pursuit methods (see [23]) although both the latter methods rely on accurate measurement of a distant ‘goal’ point. A novel solution proposed by [9] uses an additional camera which can move independently of the robot body controlled through a biologically inspired neural network. However precise imitation, both when and after following, is difficult and the solutions can limit the

⁷ $(\theta_4, r_4) = \Sigma_{i=1}^n (LightSensorAngle(i), LightSensorReading(i))$, with the resultant magnitude r_4 held in a lookup table indexed by robot turn angle and distance from the model.

behaviours of both model and imitator severely. Indeed, there is a deep issue relating what an imitator’s sensors are telling it and what it is actually doing when following. A further issue is that as the imitator reacts, its reference frame is changing in relation to the model. Calculation of these changes is not possible from using the reactive vector alone. This means the robot cannot retain a memory of the imitation from its sensor perceptions; it can however do so by storing its motor outputs (i.e. the velocity of each wheel during each time-cycle). Thus it imitates what it ‘feels’ rather than what it ‘sees’. However this can be beneficial. As part of the experiment a series of small strips of tape were placed on parts of the imitator’s path. A Khepera initially moving at constant velocity will typically slow down when moving over these strips. However, when following and in order to maintain contact with the imitator the reactive vector will produce an increase in motor outputs as the model moves away at constant velocity from the imitator. This change then becomes part of the memory of the imitation. Thus the imitator manages to imitate more faithfully by reference to its internal state rather than an external observed state.

3.2 Static Observation

A similar set of tests for the static observation experiment are shown in figure 3 below. The results appear promising with the robot imitating the model with reasonable accuracy, with vertices accurately tracked and little evidence of the impersistence problem. However a number of points should be borne in mind. Firstly the performance of the tracking mechanisms depends crucially on how the observed data is filtered. Time averaging and a significant event extraction are both based on thresholds. The choices of threshold have to be carefully chosen to achieve this level of accuracy. Secondly, the method required to




Figure 3: *Observation Behaviours for triangle, circle and T shape. Dotted line is imitator’s imitation attempt. Unbroken line is unfiltered observation. Based on 100 observations at 100ms intervals, 10 point moving average.*

compute the path is complex as compared to the rel-

atively simple vector summation methods used in the following experiment. Finally and most importantly, by using observation alone the imitator would have no way of learning any parts of the imitation which were not addressable from observation alone. For example, the static imitator only records a constant velocity when observing the model traversing the tape strips discussed in ‘Following’ above; when imitating it fails to increase its motor outputs over the strips to maintain velocity. It seems that this need to modify motor outputs may only be obtained from experiencing the situation.

4 Discussion

We have shown, using some simple experiments, different aspects of robotic perception and performance which may be encountered when attempting to use imitation to match movements made by either robots or people. Our results illustrate that each of the two often used paradigms, following and static observation, raise different issues. The table shown in figure 4 illustrates the various dimensions and trade-offs that we observed in the course of our experiments. The examples are simple in that no explicit communication is permitted between the model and imitator, in fact the sensory information is basically the perceived brightness of a light bulb. The results indicate that there is a clear trade-off between positional accuracy obtained from static observation and the advantages of direct perception-action coupling available from following. This is probably unsurprising, given that static observation is certainly the more complex and engineered method in these and most other robotics experiments.

The relative simplicity of the following paradigm also hides some key advantages, in that the robot is able to directly map its perceptions against its motor actions. It is thus able to learn much about the environment directly and relatively cheaply. However to achieve positional accuracy, more complex observational algorithms are required, but observation alone may be insufficient to correctly assess the physical complexities of the environment. There might be an argument for suggesting that observation can be most effective after a following episode, i.e. observation can fine-tune already stored movement patterns.

Animals use both seeing and feeling in imitation and there may be an appropriate time to see (observe) as opposed to feel (follow) in social learning. A mixed approach may be valuable, this approach corresponding to intermediate positions or switching in the spectrum table shown. One could imagine for example cases where the observation is less static e.g. several follow-observe-follow cycles, or where a series

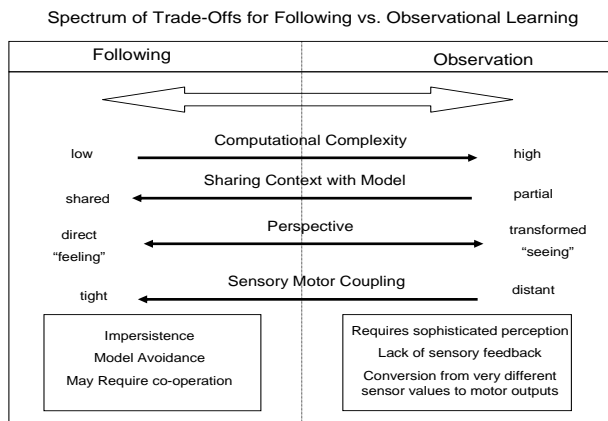


Figure 4: The table summarises the key aspects revealed by the experiments with extremes of each aspect shown (see text). Comparative costs are shown in the boxes. Mixed approaches might allow the balance of these costs and benefits.

of static observations are made prior to each episode of following behaviour. The next stage in this research will be to look at these mechanisms more closely by comparing points along the continuum from following to observation with mixed strategies. We also intend to add a further aspect to the table in looking at simple, non-symbolic feedback using re-inforcement signals between the model and imitator allowing the model to indicate success of the imitation attempt.

We remark as a final consideration, that there are many examples in human life where observation alone would lead to disastrous consequences. So, if you have never skied before please don't attempt to ski down a steep mountain by only watching others!

References

- [1] A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Synchrony and perception in robotic imitation across embodiments. In *Proc. IEEE CIRA '03*, pages 923–930, 2003.
- [2] R. C. Arkin. *Behavior-based robotics*. MIT Press, Cambridge, Mass., London, England, 1998.
- [3] D. C. Bentivegna and C. G. Atkeson. Learning how to behave from observing others. In *Proc. SAB'02-Workshop on Motor Control in Humans and Robots: on the interplay of real brains and artificial devices, Edinburgh, UK, August, 2002*, 2002.
- [4] A. Billard and K. Dautenhahn. Grounding communication in situated, social robots. In *Proc. TIMR*, 1997.
- [5] R. W. Byrne. *The Thinking Ape, Evolutionary Origins of Intelligence*. Oxford University Press, 1995.
- [6] K. Dautenhahn. Trying to imitate – a step towards releasing robots from social isolation. In P. Gaussier and J.-D. Nicoud, editors, *Proc. From Perception to Action Conference, Lausanne, Switzerland*, pages 290–301. IEEE Computer Society Press, 1994.
- [7] K. Dautenhahn and C. L. Nehaniv, editors. *Imitation in Animals and Artifacts*. MIT Press, 2002.
- [8] B. G. Galef and C. M. Heyes, editors. *Social Learning in Animals: The Roots of Culture*. Academic Press, 1996.
- [9] P. Gaussier, S. Moga, J. P. Banquet, and M. Quoy. From perception-action loops to imitation processes: A bottom-up approach of learning by imitation. In *Socially Intelligent Agents*, pages 49–54. AAAI Press, Technical report FS-97-02, 1997.
- [10] G. Hayes and J. Demiris. A robot controller using learning by imitation. In *Proc. International Symposium on Intelligent Robotic Systems, Grenoble*, pages 198–204, 1994.
- [11] R. Hecht-Nielsen. Applications of counterpropagation networks. *Neural Networks*, 1:131–139, 1988.
- [12] Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: Extracting reusable task knowledge from visual observations of human performance. *IEEE Trans. Robot. Automat.*, 10:799–822, November 1994.
- [13] M. J. Matarić, V. Zordan, and Z. Mason. Movement control methods for complex, dynamically simulated agents: Adonis dances the macarena. In *Proc. Autonomous Agents '98*, pages 317–324. ACM Press, 1998.
- [14] N. E. Miller and J. Dollard. *Social Learning and Imitation*. Yale University Press, 1941.
- [15] R. W. Mitchell. A comparative-developmental approach to understanding imitation. In P. P. G. Bateson and P. H. Klopfer, editors, *Perspectives in Ethology 7: Alternatives*, pages 183–215. Plenum Press, 1987.
- [16] R. R. Murphy. *Introduction to AI Robotics*. The MIT Press, Cambridge, 2000. p.76.
- [17] C. L. Nehaniv and K. Dautenhahn. Like me? - measures of correspondence and imitation. *Cybernetics and Systems*, 32(1-2):11–51, 2001.
- [18] U. Nehmzow. Animal and robot navigation. In L. Steels, editor, *The Biology and Technology of Intelligent Autonomous Agents*. Springer Verlag, 1993.
- [19] J. Noble and P. M. Todd. Imitation or something simpler? modeling simple mechanisms for social information processing. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*, pages 423–439. MIT Press, 2002.
- [20] B. Scassellati. Imitation and mechanisms of joint attention: A developmental structure for building social skills. In C. L. Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 176–195. Springer Lecture Notes in Artificial Intelligence, Volume 1562, 1999.
- [21] S. Schaal. Learning from demonstration. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems volume 9*, pages 1040–1046. MIT Press, 1997.
- [22] E. L. Thorndike. Animal intelligence: An experimental study of the associative process in animals. *Psychol. Rev. Monogr.*, 2:551–553, 1898.
- [23] J. Wit. *Vector Pursuit Path Tracking for Autonomous Vehicles*. PhD thesis, University of Florida, 2000.
- [24] T. R. Zentall. Imitation in animals: Evidence, function and mechanisms. *Cybernetics and Systems*, 32(1-2):53–96, 2001.