An Examination of the Static to Dynamic Imitation Spectrum

Joe Saunders, Chrystopher L. Nehaniv and 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 consider the issues that arise from an examination of the continuum between two social learning paradigms that are widely used in robotics research: (i) *following* or *matched-dependent behaviour* and (ii) *static observational learning*. We use physical robots with minimal sensory capabilities and exploit controllers using neural network based methods for agent-centred perception of model angle and distance. The robot is first trained to perceive the dynamic movement of a robot model carrying a light source, then the robot learns by observing the model demonstrate a behaviour and finally it attempts to re-enact the learnt behaviour. Our results indicate that a dynamic observation using rotation performs significantly better than static observation. However given the embodiment of the robot a dynamic strategy using both rotational and translational movement becomes more problematic. We give reasons for this, discuss lessons learned for combining these types of social learning and make suggestions for requirements for imitator robots using dynamic observation.

1 Introduction

In this paper we build on our previous research (Saunders et al., 2004) in considering the issues that arise from an examination of the continuum between two social learning paradigms that are widely used in robotics research: (i) *following* or *matched-dependent behaviour* and (ii) *static observational learning*. Our motivation in examining these issues is the belief that an understanding of the mechanisms underlying social learning should be considered as a prerequisite for building adaptive and intelligent robots. We believe that social learning leads to an acceleration of the acquisition of intelligent behaviour (Zentall, 2001; Galef and Heyes, 1996; Dautenhahn and Nehaniv, 2002) with the promise of easier robot task acquisition, increased behavioural complexity and ultimately some form of cultural transmission (Alissandrakis et al., 2003). In this respect we focus on the mechanisms supporting Imitation[1] with experiments with physical robots in an attempt to simplify and focus on key aspects of imitative processes. The background of this paper is an ongoing investigation of social learning and the interaction between both human/robot and robot/robot pairs to understand the social dimension of imitative behaviour. The perspective of both the imitator and the imitatee and the problems of perception and action encountered by both are considered. Our starting point is the different imitator perspectives which are widely applied in paradigms used in robotics imitation research, namely *following* behaviour (Hayes and Demiris, 1994; Billard and Dautenhahn, 1997; Dautenhahn, 1994) and *static observation* behaviour (Kuniyoshi et al., 1994; Gaussier et al., 1997; Bentivegna and Atkeson, 2002; Schaal, 1997; Matarić et al., 1998; Alissandrakis et al., 2003).

From a psychological/ethological viewpoint *following* is more rightly considered as *matched-dependent behaviour* (Zentall, 2001). The imitator observes and immediately matches the behaviour of the model as it is being performed, staying close to the model. For example rats can be trained to follow a lead rat through a maze which they then learn to navigate (Miller and Dollard, 1941). The rats may have no idea of intentionality of the lead rat and can be trained to follow other salient (including non-animal) stimuli, this behaviour is sometimes called *discriminated following*.

[1]We take Thorndike's 1898 classical definition of imitation (Thorndike, 1898) as " learning how to do something by seeing it done" but extended to include non-biological agents (Mitchell, 1987).

Likewise, *static observation* by the imitator who stays at a fixed location is related to the ethological/psychological notion of *observational learning*. Here the behaviour of the demonstrator is copied after it is 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 (Galef and Heyes, 1996), without reference as to whether the demonstrating rat becomes ill or dies. These examples hint at 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 our previous research (Saunders et al., 2004) we considered the extremes of a purely reactive following behaviour and contrasted that against a static observational behaviour using some simple experiments with Khepera miniature robots. Two controllers were designed to allow either a reactive following behaviour or a static observation behaviour. Each robot either followed or statically observed another robot making various geometric shapes over varying terrain. In both cases the robot could learn the observed behaviour and attempt to re-enact it. The model was perceptable by the imitator due to placement a small light bulb on top of the model. No explicit communication was permitted between the model and imitator; in fact the sensory information was basically the perceived brightness of the moving light bulb. The research results from these experiments identified trade-offs that are summarised in the spectrum table shown in figure 1.

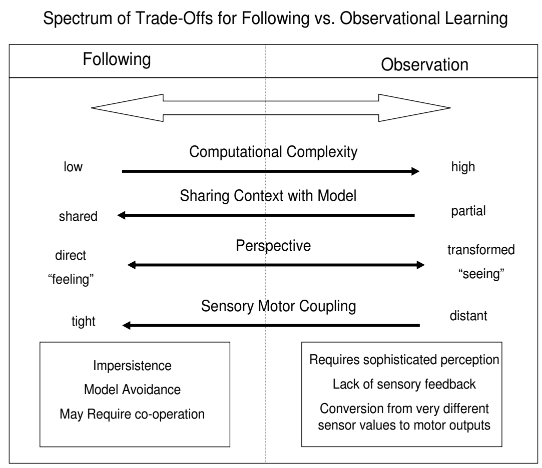

Figure 1: *The table summarises the key aspects revealed by the previous research experiments (Saunders et al., 2004) with extremes of each aspect shown. Comparative costs are shown in the boxes. The current research considers mixed approaches which might allow the balance of these costs and benefits.*

The results indicated that there was a clear trade-off between positional accuracy obtained from static observation and the advantages of direct perception-action coupling available from following. This lack of precision during following we called *impersistence* to reflect the fact that the robot is always reacting to the latest sensor reading and not persisting to meet the goal signalled by the previous reading. We believed that the accuracy available from static observation was unsurprising, given that static observation allows the design of the robot controller to concentrate exclusively on angle and distance perception and apply more complex and engineered methods to this task. We believed that similar complexity in observational systems were also engineered into most other social learning robotics experiments.

The relative simplicity of the following paradigm also hid some key advantages, in that the robot was able to directly map its perceptions against its motor actions. It was thus able to learn much about the environment directly and relatively cheaply. However to achieve positional accuracy, more complex observational algorithms were required, but observation alone was insufficient to completely assess the physical complexities of the environment. We said that there may be an argument for suggesting that observation could be most effective after a following episode, i.e. observation could fine-tune already stored movement patterns. Similarly there may be an appropriate time to 'see' (observe from a distance) as opposed to 'feel' (follow, experiencing the same context) 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 of static observations are made prior to each episode of following behaviour.

Dynamic Observation. In this paper we consider in more detail some of the effects of allowing a more dynamic observational approach. We study the quality of the imitation attempt from the imitator's perspective in two experiments using either an 'observe and rotate' or an 'observe and move' strategy to match the movement patterns of the model. These successively augment static observation, respectively, by adding orienting rotational changes to allow the imitator to track the model (observe and rotate) or by adding rotation and translation.

2 Experimental Overview

For our experiments we use a controller previously designed (Saunders et al., 2004) to investigate the imitation of movements using static observation and extend it to provide a mechanism to investigate dynamic imitation. All experiments are carried out in real-time on physical robots (i.e. simulation is not used) on a desktop in a typical busy academic environment with light levels varying during the day.

Figure 2: *The picture shows the experimental platform. The Khepera acting as a model has a small bulb placed on top of it. The imitator is shown tracking the model which is tracing out a triangle.*

We examine the behaviour of the imitator when imitating various geometric shapes made by the model. We consider intermediate positions in the spectrum table in two experiments to examine the effects of a mixed observation and movement execution strategy. Both experiments involve dynamic observation which combines both observation and movement.

Observe and Rotate. The first experiment extends the static observational perspective by allowing the robot to alter its orientation so as the better exploit it sensory facilities. The embodiment of the Khepera robot is such that the majority of the light sensors are in front of the wheels, with two sensors at the back. The estimation of distance is therefore more accurate when the robot is able to employ all of its front facing sensors as it is receiving more information from the environment. To ensure that these sensors are in an optimal position we program the circular Khepera robot to rotate in place orienting toward the model. The rotation is such that the imitator will attempt to directly face the model if the model's angle with respect to the imitator exceeds a given threshold. However, if the imitator has to rotate to achieve this then all subsequent observations must be converted back to the original reference frame in order to replay the imitation. To achieve this conversion, accuracy in measuring how far the robot has turned is critical to this process. We tested threshold angles of 0, 30, 60 and 90 degrees. In both this and the experiment described below the model was preprogrammed to make 4 geometric shapes. The first was a 10cm radius circle around the imitator, the second a 10cm circle 5cm in front of the imitator. The third and fourth a triangle and T-Shape 5cm in front of the imitator.

Observe and Move. The second experiment allows the robot to record a sequence of observations of the model and then attempt to use a given subset of these observations to imitate the model's movement sequence. Once the imitator has completed this part of the imitation it recommences observing.

In a two-dimensional parameterisation of the spectrum, different social learning mechanisms are given by varying both the number n of *observations* and the number m of *movements* made by the imitator. A single observation is an estimation by the imitator of the model's angle and distance from the imitator. A movement is the transformation and execution by the imitator of observations to motor-commands in order to achieve the same effect.

These mechanisms however present a series of challenges due to the fact that after each movement sequence the robot's memory of previous observations will be from a different perspective from the current observation set. This is because the imitator, after partially replaying the imitation (by transforming a subset of the observed vectors) will find that the remaining observations need to take account of the new observation position. Furthermore, the new observation position may not be optimal for accurate readings, therefore a rotation (as in experiment 1) will be necessary. To then replay the next part of the imitation the effect of the rotation must be reversed and subsequently a transformation of the observations re-performed.

3 Controller

The controller used in both experiments relies on computing the distance and angle from the imitator to the moving model and storing these observation points as a list of two element vectors. Prior to observing, the robot first learns how to measure angles and distance.

Learning to Measure Angles. The robot is first trained to accurately compute the angle of the light from the centre of the imitator. A number of methods were evaluated including using a *light compass* (Nehmzow, 1993), or computing the angle by using *vector summation* of the inputs to each of the light sensors (Arkin, 1998). 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 sensory experience and is 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 it 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 (Hecht-Nielson, 1988)). This is a fully connected feed-forward three layer network. The first layer takes 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 final layer used to output the conversion of these values to angles. Using this technique has a number of advantages. Firstly that the network can be built as the environment is observed, secondly there are no additional training steps i.e. there is no further training of the neural network, 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. 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 (Murphy, 2000).

Learning to Measure Distance. For distance measurements various mechanisms were also assessed. A first approach was to use *triangulation*, exploiting the fact that accurate angle measurement was now possible. The approach measured the light angle from the model, moved the imitator a fixed distance and then read the new angle. This allows the computation of the original distance using the two angles and the travelled distance. However this mechanism was unreliable for two reasons, firstly that, over small movement distances (which minimised errors in the odometry readings from the wheel encoders), the derived angle would be small and tiny errors in the angle measurement would result in an amplified error in the distance computation, secondly if the model was moving, the measurements/movement combination of the imitator could never be fast enough to resolve the position of the model accurately. An alternative method based on *environmental sampling* 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. Using this method gave a reasonable distance accuracy to about 25 cm from the robot at an angle between approximately $30°$ to $150°$ in front of the robot. However, outside these parameters the distance accuracy was very poor.

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. Automatic interpolation between observed values is achieved by setting the middle layer 'winning nodes' to a value greater than 1.

Observing Distance. During the observation phase the angle is computed, followed by magnitude of the vector summation[2], the two values providing the key to the lookup table to yield distance.

Altering the Angle of Observation. In both experiments the robot collects a set of angles/distances

[2] Refer to (Saunders et al., 2004) for details.

from itself to the model whilst the model is moving. The imitator cannot poll its sensors when it itself is moving. Thus in a fixed time period the number of possible observations when the imitator is not moving will be higher than when the imitator is moving. In the first experiment the imitator can either not move or rotate to face the model once a threshold angle has been exceeded (see figure 3). The lower the threshold angle the greater the rotational movement of the robot to face the imitator when the angle is exceeded. The higher the angle, the smaller the rotational movement, but the robot will move more often. In our previous research we had fixed the imitator position and allowed it to observe the moving model. The model was at all times in front of the imitator and therefore within range of angle/distance computation mechanism. We now allowed the model to be both in front of the imitator and at any angle around the imitator. By varying the rotation threshold we can then examine both the effect of rotational movement size and the effect of frequency of movement on the imitation attempt.

Time Averaging and Way Points. In both experiments the recorded observations are smoothed using a simple moving average. The smoothed trajectory is then thresholded to yield a set of way points. This procedure is necessary for two reasons. Firstly to eliminate the effect of noisy observations and secondly to avoid two observation points being too close to one another - this closeness causing large and potentially damaging changes in the robot's motor systems if replayed directly. The imitator uses the derived way points to then imitate the model's trajectories. In the second experiment this procedure is only applied when computing the required movement. Any unused observations (which result from the movement index being less than the observation index - see experiment 2 below), remain unmodified as these may be subject to geometric transformation following the actual movement of the imitator.

4 Experimental Results

In our experiments we compared imitation behaviour on four simple patterns. These were a triangle, a circle enclosing the imitator, a circle observed ahead of the imitator 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 using a more dynamic approach in observational learning.

4.1 Experiment 1 - Dynamic Observation with Rotation Only

Details of Set-up. In each case the imitator is placed at the centre of the experimental platform (shown as point 0,0 on the graphs in figure 4) facing forward (at $90°$ along the positive Y-axis). The model is pre-programmed to move according to the prescribed shape. A threshold rotation angle is then set and the imitation run commenced for a fixed period. The threshold supplies a range of values around the front of the robot. For example, setting a threshold of say $60°$ means that if the imitator perceives the model within a forward range of $60 - 120°$ (see figure 3) no rotation will be applied. If however the model

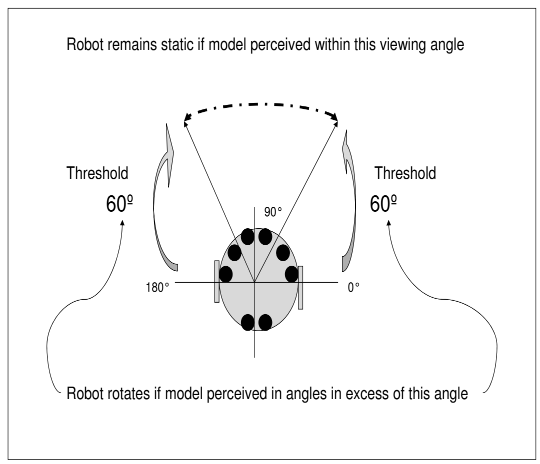

Figure 3: *Rotation Threshold. In this example the imitator will not move at values between* $60°$ *and* $120°$. *Between* $61°$ *and* $121°$ *the imitator will rotate to face the model.*

moves to, say, $50°$ the imitator will rotate so that the model is directly in front of it, and thus be, from the imitator's new perspective, at $90°$. The higher the threshold angle (to the limit of $90°$) the more often the imitator will move to match the model but it will rotate by a smaller amount. If the threshold is set to zero, then the imitator will only move when the model is outside the range $0 - 180°$, however the robot will then rotate by at least one quarter of its circumference.

Results. Figure 4 shows the results from a test with the enclosing circle. The robot is placed facing forward along the positive Y-axis. After the run the

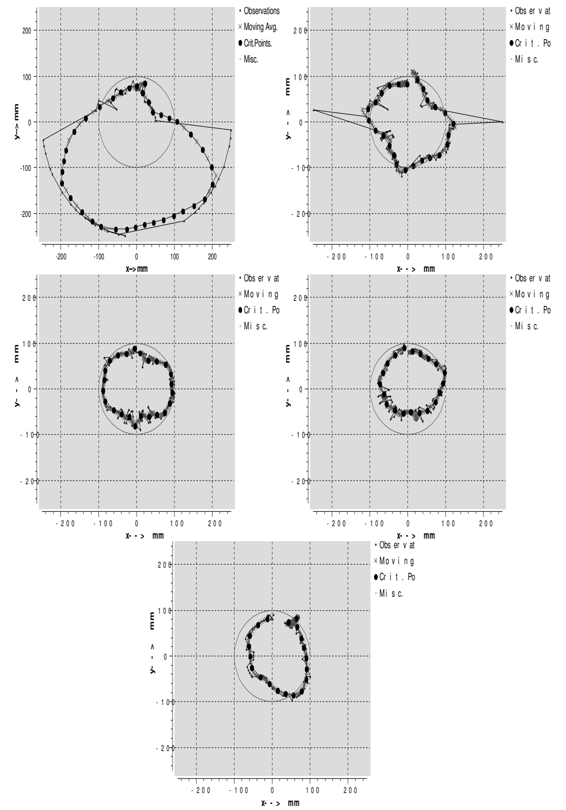

Figure 4: ***Dynamic Observation with Rotation Only.*** *Imitative Behaviours for an enclosing circle. The first diagram shows the result with no rotation, the final four graphs show rotation at* $0°$, $30°$, $60°$ *and* $90°$ *degrees thresholds. The continuous line shows the path of the model, dotted line the imitator observations, crosses the smoothed observation and large dots the way points which are replayed by the imitator.*

imitator robot attempts to re-enact the observed behaviour. The large dots on the graphs show the way points, these being the path that the imitator will take when replaying the imitation. The first graph shows the imitation when no rotation has been applied and thus where only static observation is taking place. As expected at angles outside the angle/distance range the imitation is poor. The second graph shows the first example of a dynamic observation with the imitator moving only when the model moves outside the range $0 - 180°$. Two extreme observation points are shown reflecting the inability of the distance/angle sensor to correctly measure the distance. However, once the $180°$ or $0°$ angle is exceeded the robot turns and starts again to make reasonably accurate readings. On imitation replay the outlying readings are smoothed away. The situation is further improved at $30°$ when the sensory apparatus is always in range but

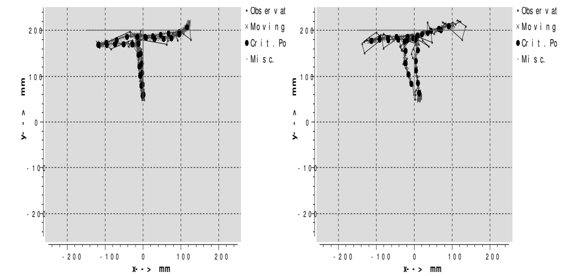

Figure 5: ***Dynamic Observation with Rotation Only.*** *Imitative Behaviours for T-Shape The two graphs show results with rotation at* $60°$ *and* $90°$ *thresholds.*

the number of moves small. However at $60°$ and $90°$ the situation is ambiguous. We would suggest that the imitation is slightly less accurate. This may be due to the increasing effect of odometry errors as the number of moves increases. This is especially true at $90°$ where there would be a small movement for every $1°$ change on the model's position.

Results for the forward circle and triangle (not shown) were less marked, however the robot was subject to less movement due to the constrained angles presented by both shapes. The T-Shape however is more interesting (see figure 5). The nature of the shape meant that at $0°$, $30°$ and $60°$ the imitated trajectories were broadly similar, however at $90°$ the robot was affected again by odometry drift and a similar worsening of readings ensued.

Analysis. These effects show some of the advantages and disadvantages of a tracking mechanism described above. Observing whilst not moving (static observation) has the key advantages of being fast and thus able to make more observations in a given time period (given the sequential nature of the observe/move scenario presented here). There are no odometry concerns as the imitator is not moving and the energy required would be lower than for a moving imitator. The major disadvantage is of course that the model can move into imitator blind spots. The advantages of the tracking imitator (dynamic observation with rotation only) is that blind spots can now be seen, however this is offset by the disadvantages of increasing odometry errors as more movement is carried out, a higher energy cost, and more complex computation as reference frame adjustments are continuously required. However at a particular movement/rotation ratio, which for this robot appears to be around $30°$, there appears to be a point where accuracy is optimised. This suggests an clear strategy - expend energy and computational costs by moving only when

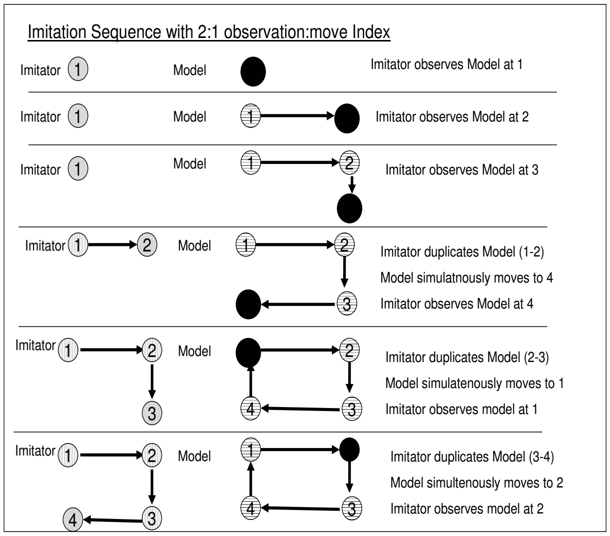

Figure 6: *Analysis of Observation and Movement Index. In the analysis the model describes a square pattern. Here we see the imitator using a observation:movement index of 2:1 and successfully matching it. Similar successful matching will always occur when the movement index is set to 1 regardless of the observation ratio.*

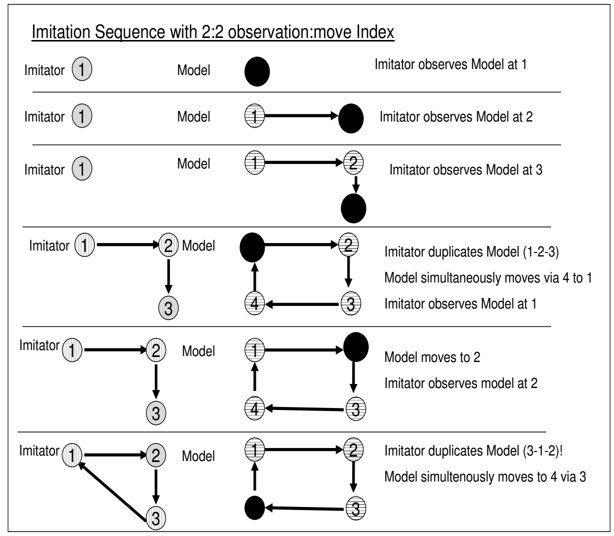

Figure 7: *Analysis of Observation and Movement Index. This example shows the imitator an index of 2:2. When the movement index is set above 1 the imitator always fails to match the pattern.*

not to do so would give incorrect results. Or more simply - keep still until movement is almost necessary (in our case when the model goes beyond the 30° threshold into 'peripheral' vision).

4.2 Experiment 2 - Dynamic Observations: Varying Observation and Movement Cycles

Theoretical Results and Detailed Set-up. This experiment explored how movement and observation might be intermixed. This was attempted by varying the number of look-ahead observations against an equal or smaller number of moves. Thus the robot would first make an initial observation [3] and then subsequently observe for n cycles and then move, based on these observations, m times. This procedure iterated throughout the imitation attempt. Prior analysis of this method, using the imitator and model represented as points (see figures 6 and 7) suggested that accurate imitation may only be possible if the number of moves were set to 1. To simplify the analysis we assumed that the imitator and model moved at approximately the same constant speed. Additionally, due to the control system of the robot, observation and movement execution are not possible in the same time step. We then imagined three scenarios cyclically alternating n observations (o) of model transitions with m moves (x) by the imitator (note: it is not possible to imitate further than our observed sequence, and therefore n is always larger or equal to m). The first scenario was of n observations to 1 move e.g. 1:1 *o-o x o x o x o x* ..., secondly a scenario where there are an equal number of observations and moves but where both are greater than 1, e.g. 2:2 *o-o-o x-x o-o x-x o-o x-x* ... and finally where n is greater than m and both are greater than 1 e.g. 3:2 *o-o-o-o x-x o-o-o x-x o-o-o x-x* .... Figure 7 shows an example of the failure to correctly match the movement pattern when the move index is set higher than 1. This occurs because the imitator has failed to observe one or more critical points in the model's move sequence. The effect is similar to the *impersistence* problem we noted when analysing 'following' behaviour (Saunders et al., 2004), however rather than failure to complete or persist in its goal, as was the case for following, here the problem is one of 'inattentiveness'. The imitator is blind to the moves of the model. This problem occurs at all values of n and m which are larger than 1. Figure 7 shows an example of this when n:m is set to 2:2.

Results. The robot was tested on a series of index values on each geometric shape presented by the

[3] For each move two observation vectors are required, therefore at the start of the run one additional observation is made.

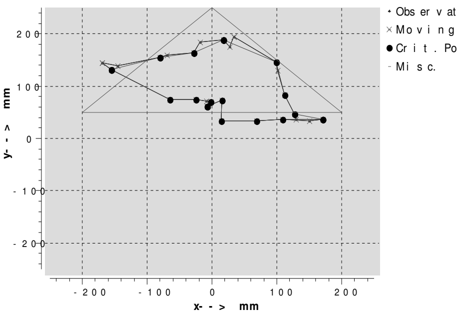

Figure 8: ***Dynamic Observation: varying Observation and Movement cycles.*** *The results show the inability of the imitator to correctly replay the model's path. In this case a triangle shape.*

model. Figure 8 shows an example of the physical robot using a 5:1 n:m index on the triangle shape. The imitator fails to match the model. Similar failures occurred in all attempts with the physical robot on all shapes. This was initially surprising, however the difficulty became clear once the actual imitator movement was considered.

Analysis. The simplicity of the point analysis above hides some crucial implementation issues. For example the robot can only move in the direction of its fixed wheels (i.e. it cannot arbitrarily move sideways), therefore a rotation may be necessary to orient the robot to the correct movement vector signalled from the model. Also, the Khepera has a fixed placement of sensors around its circular wheelbase. In order to correctly 'focus' on the model the robot must be in the appropriate sensor range. Thus the rotation mechanism described in experiment 1 was employed. Therefore in addition to the move or moves calculated from the observations we may have up to 2 additional moves: one to focus the sensors on the model and the other to orient the imitator for its move. Whilst these moves are being carried out the problem of 'inattentiveness' is compounded. Two further issues were also apparent. Firstly, each move is accompanied by a small odometry error. The total error therefore increases as the number of moves increases. Secondly, the smoothing effect of time averaging has little or no effect when attempting to model a small number of moves. This means that unsmoothed noisy observations are replayed leading subsequently to a poor imitation attempt.

We believe that the failure of the imitation is due primarily to the constraints imposed by the embodiment of this particular robot and might be obviated if the sensory apparatus was independent of the actuator mechanism e.g. distance/angle sensors which rotated and were focusable independently of movements of the main robot body. Such a mechanism is in fact used by (Gaussier et al., 1997) in their experiments. In fact there are no known imitating animals whose observation sensors cannot focus at least somewhat independently of the orientation of their bodies.

5 Discussion

In this research we have started to examine some of the practical issues which face an imitator when trying to use a dynamic observational behaviour. Here the problems of perception, perspective and action must be considered. We have greatly simplified the problem domain by restricting the imitative actions to that of replaying geometric shapes and have used simple robots with fixed sensor embodiments and with limited perceptual capabilities. We have previously suggested that a 'following' behaviour, although limited in its imitative accuracy, has the major advantage of computational simplicity and the added value of direct interaction with the environment through proprioceptive polling of its actuators whilst moving. We do not suggest that this opportunity to 'feel' the environment is exclusive to a following strategy and accept that there are alternative and probably better ways to proprioceptively explore the environment. However this strategy has the straightforward merits described above. It is also true that both a follower and an static observer are necessarily out of phase with the model and for this reason it seems that the follower's sensory cues may not be more appropriate than an observer's, however work by (Billard and Dautenhahn, 1997) showed that the these cues are dependent on the distance between a follower and the model and within a critical distance the follower's sensory cues become very relevant. What we describe here is an initial attempt to provide a movement mechanism to an observer in order to combine the advantages of observational accuracy with the feedback obtained from actively exploring the environment. Clearly a simple and modular solution to this task would be to keep to the 'extreme' behaviours and simply apply each strategy in turn e.g. follow-statically observe-follow. One of the aims of this research has been to explore the challenges faced in combining these strategies whilst retaining the positive aspects of both. The experiments themselves are clearly limited as we are constrained both by the sensor embodiment of the robot and its internal control system, but we believe valuable lessons still emerge.

Suggestions for Imitators Dynamically Observing from a Fixed Location. Our first experiment showed that dynamic observation with rotation was successful in that it allowed the model to pass out of view of the imitator and be reacquired. It was superior to static observation alone in this respect and it appeared that the benefit of tracking accuracy could be balanced against the cost of rotation frequency and rotational movement based on a turn threshold. Thus to retain observational accuracy, rotational movement should be limited so that odometry errors are minimised in their effect on the geometric transforms required to replay the imitation. Thresholds near the periphery of vision balance these factors. In robotics the issue of errors from odometry drift is clearly not new, however the literature on robotic observational imitation seem rarely to cite it as being a problem for a moving imitator.

Suggestions for Imitators that Observe and Move. Our second experiment showed that with this particular robot, dynamic observation with movement of the imitator was extremely difficult and failed to replicate with reasonable accuracy the model's path. Our theoretical analysis suggested that the 'inattentiveness' problem may be soluble for a dynamic observational imitator where the movement value is set to unity. This region in our spectrum corresponds with the methods of other research (Wit, 2000) where a single solution to this issue is considered. However the need to make additional movements over and above those required to track the model means that the movement value can never be unity for an embodiment where the sensor orientation is completely fixed for a fixed body orientation. Thus the imitation will be poor.

Possible Solutions. A solution to this might be *independent sensing and actuator mechanisms*. We envisage that such a system would additionally employ independent computation facilities for both mechanisms to allow continuous and parallel calculation of model position. Thus appropriate movement vectors could be sent to the actuators reducing unnecessary movements and the associated additional odometry drift. The sequential nature of the move-sense cycle on our robot may mean that accurate dynamic observation is very difficult, however other control systems employing a *parallel cycle* may provide solutions. There may also be simpler alternatives, for example the model may *repeat the pattern* and the imitator might manage to fill the gaps caused by earlier inattentiveness, or the model might simply *wait for the imitator*.

Even in our own human experience it appears much harder to both partially replay an imitation and observe the model before the model has finished its actions. Animals in fact may have obviated this issue by evolving alternative mechanisms. In this respect the recent neurological evidence of 'mirror neurons' in primates and humans (Gallese et al., 1996) and their role in action perception may play a considerable role in static observational learning with the imitator experiencing perhaps as good a corrrelation to its own behavioural patterns whilst statically observing as when attemping to match movements directly.

Acknowledgements

Many thanks to the reviewers of the extended abstract of this paper for their useful and constructive comments.

The work described in this paper was partially conducted within the EU Integrated Project COGNIRON ("The Cognitive Companion") and funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020.

References

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.

Ronald C. Arkin. *Behavior-based robotics*. MIT Press, Cambridge, Mass., London, England, 1998.

Darrin C. Bentivegna and Christopher 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.

A. Billard and K. Dautenhahn. Grounding communication in situated, social robots. In *Proc. TIMR*, 1997.

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.

K. Dautenhahn and C. L. Nehaniv, editors. *Imitation in Animals and Artifacts*. MIT Press, 2002.

B. G. Galef and C. M. Heyes, editors. *Social Learning in Animals: The Roots of Culture*. Academic Press, 1996.

V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119:593–609, 1996.

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.

Gillian Hayes and John Demiris. A robot controller using learning by imitation. In *Proc. International Symposium on Intelligent Robotic Systems, Grenoble*, pages 198–204, 1994.

Robert Hecht-Nielson. Applications of counterpropagation networks. *Neural Networks*, 1:131–139, 1988.

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.

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.

N. E. Miller and J. Dollard. *Social Learning and Imitation*. Yale University Press, 1941.

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.

Robin R. Murphy. *Introduction to AI Robotics*. The MIT Press, Cambridge, 2000. p.76.

Ulrich Nehmzow. Animal and robot navigation. In Luc Steels, editor, *The Biology and Technology of Intelligent Autonomous Agents*. Springer Verlag, 1993.

J. Saunders, C. L. Nehaniv, and K. Dautenhahn. An experimental comparison of imitation paradigms used in social robotics. In *Proc. IEEE Robot and Human Interactive Communication (ROMAN '04)*, pages 691–696. IEEE, September 2004.

Stefan Schaal. Learning from demonstration. In Michael C. Mozer, Michael I. Jordan, and Thomas Petsche, editors, *Advances in Neural Information Processing Systems volume 9*, pages 1040–1046. MIT Press, 1997.

E. L. Thorndike. Animal intelligence: An experimental study of the associative process in animals. *Psychol. Rev. Monogr.*, 2:551–553, 1898.

J.S. Wit. *Vector Pursuit Path Tracking for Autonomous Vehicles*. PhD thesis, University of Florida, 2000.

T. R. Zentall. Imitation in animals: Evidence, function and mechanisms. *Cybernetics and Systems*, 32 (1-2):53–96, 2001.