

Learning How to Do Things with Imitation

Aris Alissandrakis, Christopher L. Nehaniv, and Kerstin Dautenhahn

Adaptive Systems Research Group
Department of Computer Science
University of Hertfordshire
College Lane
Hatfield, Hertfordshire
AL10 9AB, United Kingdom

{a.alissandrakis, c.l.nehaniv, k.dautenhahn}@herts.ac.uk
<http://homepages.feis.herts.ac.uk/~nehaniv/ASRG>

Abstract

In this paper we discuss how agents can learn to do things by imitating other agents. Especially we look at how the use of different metrics and sub-goal granularity can affect the imitation results. We use a computer model of a chess world as a test-bed to also illustrate issues that arise when there is dissimilar embodiment between the demonstrator and the imitator agents.

Hayes 1996, Billard and Dautenhahn 1997). Imitation mechanisms are shown to improve performance in trajectory planning compared to standard model-based reinforcement algorithms in software agents (Price and Boutilier 2000). See (Dautenhahn and Nehaniv, forthcoming) for an introduction on research on imitation across disciplines.

Introduction

Imitation is one of the most important mechanisms whereby knowledge may be transferred and skills acquired between agents (both biological and artificial). It requires two or more agents sharing a context that allows one agent to imitate another. Imitation involves reproducing observed actions and/or effects in an appropriate context. A comprehensive treatment of imitation needs to examine the actions, effects, context and goals of the participating agents. An interdisciplinary framework treating imitation formally is presented in (Nehaniv and Dautenhahn, in press).

Most of the imitation literature is by animal researchers and psychologists, but they do not generally agree on how to define imitation, what constitutes evidence of imitation, or the extent to which extent research has proved evidence of imitation for various species, e.g. (Noble and Todd 1999). In experimental psychology, the Associative Sequence Learning theory (Heyes and Ray 1999) is an attempt to provide a testable and predictive framework for imitation although it does not address the effects on the environment.

The scientists that study imitation are concerned with theoretical issues such as: specification of mechanisms by which observed and executed movements are matched (learning to imitate) (Nehaniv and Dautenhahn, in press), the structure of knowledge transferred (Byrne and Russon 1998), and how imitation might be exploited as a means of acquiring knowledge (learning by imitation) (Demiris and

Dissimilar Embodiments

A fundamental problem when learning to imitate is to create an appropriate (partial) mapping between the actions afforded by particular embodiments to achieve corresponding effects by the demonstrator and the imitator (Nehaniv and Dautenhahn, 1998 & in press). For similar embodiments, this seems to be straightforward (although it actually involves deep issues of perception and motor control). But once we drop the assumption that the agents belong to the same "species", i.e. having similar bodies and an equivalent set of actions, the problem becomes more difficult and complex. Even among biological agents of the same species, individual differences in issues of perception, anatomy, neurophysiology, and ontogeny can create effectively dissimilar embodiments. A close inspection of similar artificial agent embodiments can yield similar conclusions due to issues like individual sensor and actuator differences (hardware) or the particular representations and processing that these agents employ (software). In some cases it may even be desirable to have different kinds of agents in the learning process, for example human-artificial in the area of programming by example (Cypher 1993, Lieberman 2000). So, although this appears to be a special case, at closer examination it turns out the dissimilar bodies case is more general than it may have seemed at first. In fact, methods that work for dissimilar bodies can also be applied to the simpler case of similar ones.

What to Imitate

The structure of knowledge transferred in learning how to do things by imitation is another fundamental problem. The interacting agents might have different goals and receive different rewards. The imitator must first examine whether the demonstrator's actions are beneficial or relevant to his own tasks and then extract a possible behaviour to imitate.

A distinction between different modes of imitation can be made. Low level imitation of replicating the same actions, or high level imitation of deriving the purpose and performing (not necessary the same) actions to that end.

The purpose of the imitation relates to the degree of its success. Using the example of imitating a dance instructor, one might choose to imitate only the end result, e.g. reach a specific location on the dance floor. This can be done either by following the path that the instructor used or maybe ignoring it (partially or totally). The imitator can also try to do the dance steps that the instructor performed while moving on that path. In this example choosing to ignore either the path or the actions that need to be performed in the exact order of sequence will result in poor imitation according to most dance judges. This is especially true because as external observers they can be unaware of the goals and priorities that the imitator has chosen. Using a different granularity, i.e. distinguishing which are the important aspects of the demonstrated performance and dividing them into a number of sub-tasks, can produce different results.

In the case of imitating the painting of a wall, the situation is quite different. The net result, i.e. covering the entire wall with paint can be achieved in numerous ways and the order of paintbrush strokes in the sequence is not important. Neither is replicating the exact actions; the use of a brush of different width or type than the one used in the demonstration is not important here, nor is the use of a ladder to make reaching the higher part of the wall easier. Replicating exactly the same hand movements as the demonstrator, but being far from the wall without the brush making actual contact with the surface will result in failure due to misinterpretation of the desired effect.

Using the same example we can also point out that even if the entire wall is finally covered, any paint stains on the floor or the rest of the room will undermine the overall result. Therefore, taking into account the effects on the environment can affect the success of the imitation.

The observed actions' point of reference - whether it is absolute or relative - poses another important issue. If any of the actions performed by the demonstrator are not mirror-symmetric, choosing among the possible versions will affect the imitation result, e.g. if I want to imitate somebody that raises her right hand should I raise my right or left hand?

A Chess World Model

The issues discussed above are mostly ignored or simplified in the majority of the current research on imitation. Our work addresses most of them, using as one test-bed an agent model implemented in the Swarm simulation system. The study of these fundamental imitation issues is done without restricting to any specific application domain, illustrating instead general principles in learning to do things using imitation in adaptive systems.

The model world consists of a chessboard, on which the different chess pieces engage in imitation behaviours. Movement rules for each kind of chess pieces are different, so effectively they represent different embodiment instances. A chess piece, acting as a demonstrator, performs a sequence of moves that in turn each of the chess piece types must try to imitate. The imitators apply a metric to internally evaluate possible actions for a successful imitation. Using the values from the metric, they choose a best corresponding (sequence of) moves to achieve a sequence of successive sub-goals. They are given the task of trying to imitate the demonstrator's behaviour at a particular level of granularity, which determines the sub-goals (see below).

Although the model uses elements of chess, the intention in doing so is to help illustrate some issues in the study of imitation using a familiar context, and not to consider aspects of the game of chess itself. The different chess pieces use simple, well-defined moves in a discrete world and as such they are very well suited for the study of imitation in simple dissimilar embodiments.

The simple imitation technique used in the model tries to reduce the value of a metric between the outcome of a possible action, i.e. the distance between the square that the piece will land on if that action is taken, and a desired effect, i.e. reaching another sub-goal square on the board. In order to know what to imitate, the imitator observes¹ a sequence of displacements, each relative to the previous one, which combined with a starting point result in the demonstrator's progression on the board. From this it extracts a sequence of sub-goals, desired displacements, in accordance with a particular granularity of attempted imitation. Then, using the same origin, the imitator tries to achieve similar displacements and reach the same sub-goal squares as the demonstrator did, by trying out possible moves from its repertoire and performing the most appropriate ones for each of the sub-goals, in sequence. If at some point along the sequence it is not possible to

¹ The model does not consider deep issues of perception. It is assumed that the imitators instantly obtain all the required data as soon as the demonstrator completes the demonstration, noise-free. However granularity of the imitation determines the salience of cumulative displacements, i.e. it determines which displacements (or equivalently squares) are considered as sub-goals.

further minimise the value of the metric by performing any of the possible moves, focus is transferred to the next point.

The algorithm could be effectively outlined by the following lines of pseudocode:

IMITATION ALGORITHM PSEUDOCODE

1. Observe demonstrator's moves.
2. Convert perceived moves to a sequence of sub-goals for imitation depending on granularity.
3. For each sub-goal displacement (x,y) :
 - 3a. Choose a possible move that maximally decreases the metric distance to (x,y) .
 - 3b. Repeat step 3a until sub-goal is reached or there is no move that reduces further the distance to sub-goal.
4. Move to next sub-goal in sequence, if any, and repeat step 3, else stop.

The algorithm is effectively a greedy search without any elements of planning.

The Different Pieces and Movement Rules

The different chess pieces that are used in the model are the following:

The *Rook* can only move either horizontally or vertically any number of chessboard squares in each move.

The *Bishop* can only move diagonally any number of chessboard squares in each move, but as a result cannot visit squares of opposite color.

The *Queen* can move along any direction any number of chessboard squares, effectively combining the movement rules for the Rook and the Bishop, but without the limitation of the latter.

The *King* can move like the Queen but at only a single chessboard square per move.

The *Knight* has the most complex rules, moving to the diagonally opposite end of any 3×2 (or 2×3) rectangle (see for example Fig. 3 bottom). It is considered that it actually instantly jumps to that location, so no other chessboard squares are visited and therefore the Knight cannot be blocked by intervening pieces.

In this model the *Pawn* was not considered due to its severe limitations (e.g. not being able to move backwards) that make it improbable to even remotely imitate any complex movement sequence.

Although having distinctive movement abilities, it is reasonable to expect from pieces of different type to be able to move in ways resembling one another, effectively constructing a correspondence between their different embodiments.

Metrics and Sub-goal Granularity

For the imitators to evaluate each potential action, three different types of metric were used in the model. These are

well known from mathematical analysis e.g. (Rudin 1976). Of course other possible metrics (e.g. non-geometric, behavioural metrics) could be used, including very complex ones (Nehaniv and Dautenhahn, in press).

The measures of distance between any two squares (x_1,y_1) and (x_2,y_2) on the chessboard used are:

Hamming metric (or *one-norm*):

$$|x_1 - x_2| + |y_1 - y_2|$$

Euclidean distance (or *two-norm*):

$$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$$

Infinity norm:

$$\max\{|x_1 - x_2|, |y_1 - y_2|\}$$

The different metrics provide different notions of distance as visualised in Fig. 1 below.

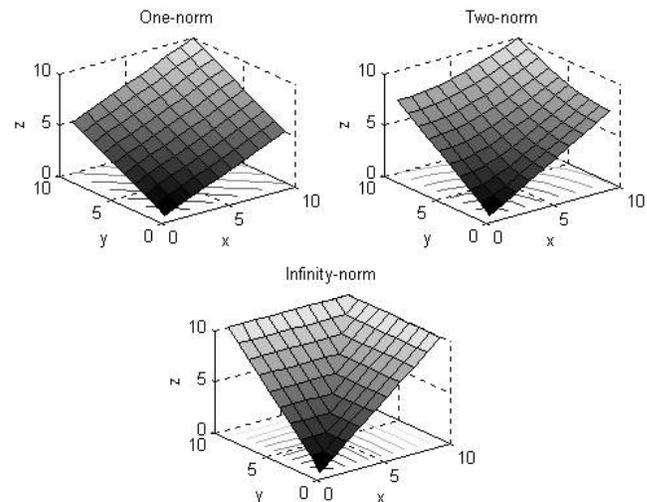


Figure 1. The different metrics measuring the distance between $(0,0)$ and (x,y) . The vertical axes are normalized.

The salience of sub-goals in a movement sequence observed by the imitator induces structure on the perceptual data it receives. Three types of granularity were used:

If the sub-goals consist only of the total replacement of the entire sequence, this effectively provides the end point reached by the demonstrator, and therefore the imitator will ignore everything else and just try to reach the *end goal* square.

If the sub-goals contain the total displacements caused by each of the moves in sequence, this can be considered as the *trajectory* of the demonstrator. The imitator will sequentially try to reach the endpoints of each of the demonstrator's moves.

If the sub-goals include for each individual move of the demonstrator the sub-displacements showing all the squares that were visited in order during that move, this effectively gives the *path* travelled. The imitator will try to visit

sequentially each one of the squares in the path with a separate move.

All the cases above also include the starting point of reference that will also be used by the imitator.

Using the Model

An exhaustive series of experiments were done using all the combinations of different metrics and granularity, with each of the different chess piece types attempting then to imitate the demonstrator. Detailed results will appear in a forthcoming publication due to length restrictions. The results shown here illustrate the most important observed phenomena.

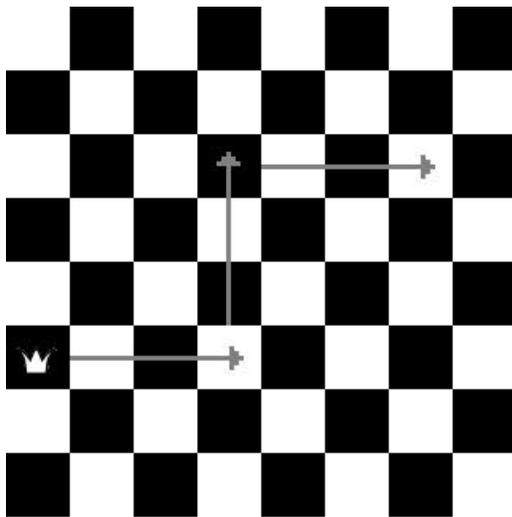


Figure 2. Sequence of moves performed by the demonstrator.

The sequence that was used for imitation in the experimental runs of the model is shown in Figure 2. The demonstrator is assumed to be a Queen, although it could also be a Rook since only horizontal and vertical moves are involved. This demonstration example is used throughout the rest of the paper.

In many cases, more than a single corresponding solution is found, especially by the Bishop and Knight. However often restrictions due to differences of embodiment make it possible only to attain some of the sub-goals in the sequence. For example in Fig. 3, the Bishop (top) cannot reach the final square in the demonstrator sequence as this is of opposite color. But it actually manages to sequentially visit all the same color squares in the path and in doing so performs well: it stays as close to the path as is possible, given the specified granularity and its embodiment limitations. The Knight (bottom) is trying to imitate at trajectory level granularity, and achieves only one out of three of the sub-goals, thus performing more poorly as well as not reaching the final square. However its trajectory does generally match the one of the demonstrator Queen and it at least does get within distance 1 of each sub-goal.

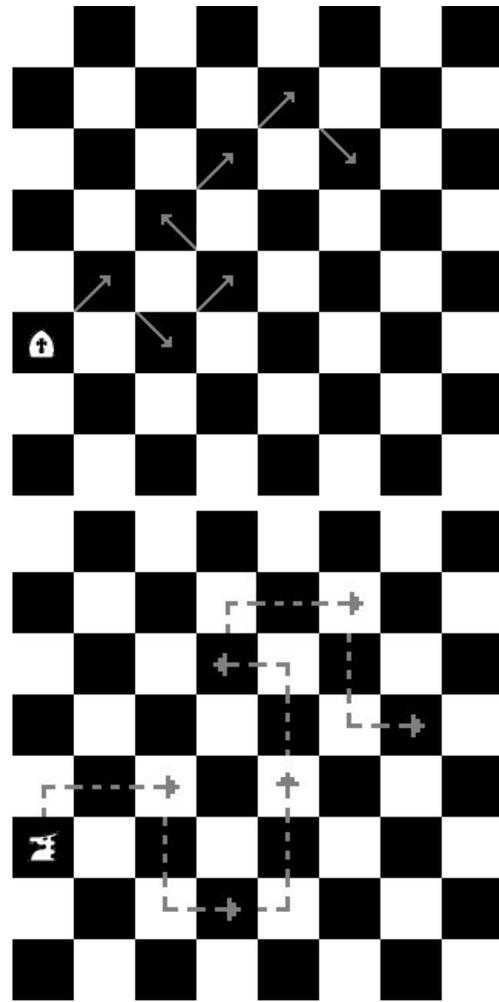


Figure 3. Bishop imitating demonstrator path using Euclidean metric (top) and Knight imitating demonstrator trajectory using Hamming metric (bottom).

Observations on Metrics

The choice of metric can greatly affect the outcome of the imitation attempt. For example in Fig. 4, due to its embodiment, using the Hamming metric (top), the King is unable to move more than one square at a time, so it has to break down the sub-goal into a number of moves, when trying to imitate at trajectory level granularity. Still the trail resembles the one of the demonstrator. When trying to achieve the same sub-goals but using the infinity norm metric (bottom), the King actually deviates from the trail. But this is accepted as it still manages to satisfy all of the trajectory sub-goals, despite the qualitative differences from the demonstrator's behavior.

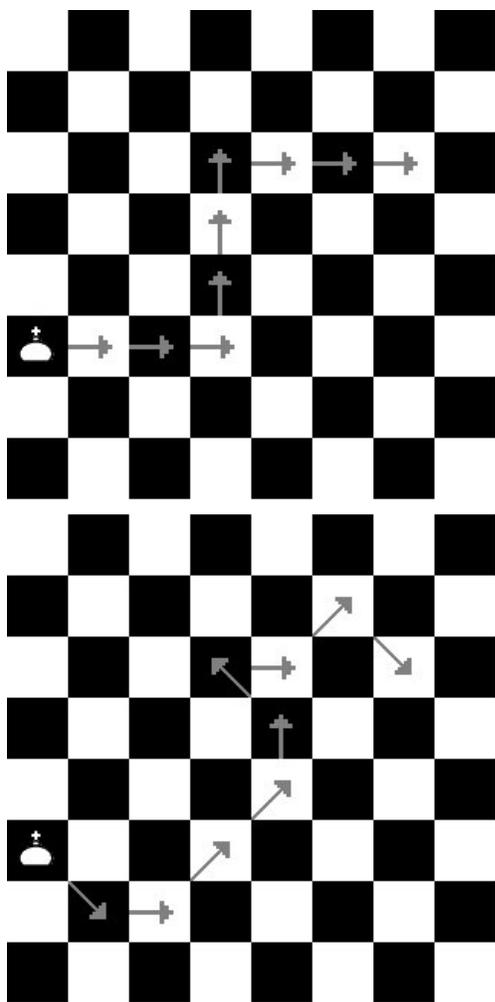


Figure 4. King imitating demonstrator trajectory using one-norm (top) and infinity norm (bottom).

In some cases the imitator can even get “blocked” i.e. unable to get closer to a sub-goal by any single move. We observe this for example for the Bishop and the Knight in Fig. 3 above.

	Metric		
	Hamming	Euclidean	Infinity norm
Rook	---	---	$x=y$
Bishop	$(\pm 1, 0)$ or $(0, \pm 1)$	$(\pm 1, 0)$ or $(0, \pm 1)$	$x=0$ or $y=0$
Queen	---	---	---
King	---	---	---
Knight	$(\pm 1, 0)$ or $(0, \pm 1)$	$(\pm 1, 0)$ or $(0, \pm 1)$	$(\pm 1, 0)$ or $(0, \pm 1)$ or $(\pm 1, \pm 1)$

Table 1. All possible situations, given by a desired displacement (x,y) , where no possible move can reduce the value of the metric. Here (x,y) is the displacement of the next sub-goal relative to the current location.

Table 1 above shows all possible cases in which such “getting stuck” can happen. Although this is caused because the algorithm examines only *single moves*, rather than *sequences of moves*, it illustrates again that the choice of metric can be crucial for the imitative success.

Observations on Granularity

Depending on the chosen granularity of the imitative behaviour, *very* different aspects of that behaviour are imitated as seen clearly for example in Fig. 5 below.

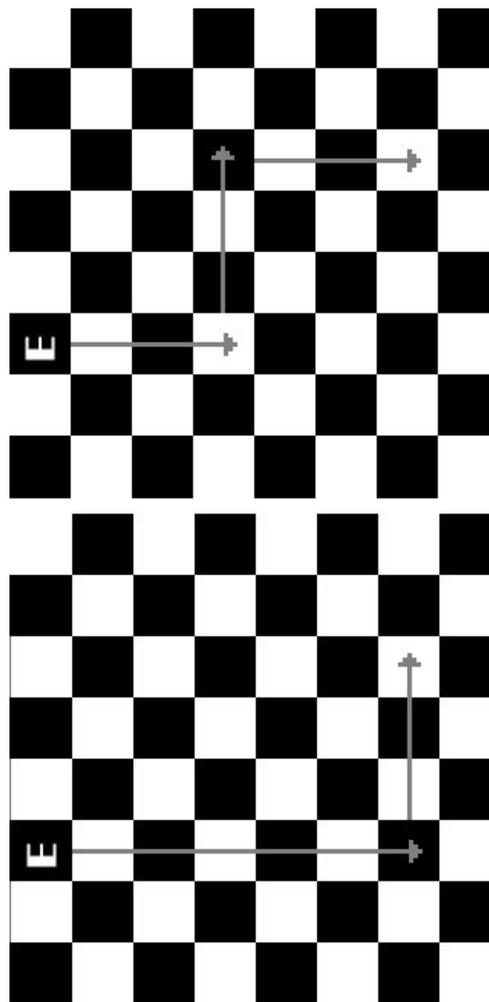


Figure 5. Rook imitating demonstrator at trajectory (top) and end result (bottom) levels of granularity using the Euclidean metric.

The Rook at a lower level of granularity (bottom) ignores parts of the trail, seeking to satisfy only the end result, instead of visiting all the end-points that it would if it were considering a higher granularity (top).

Discussion

The results from the chess world model indicate how by exposing the imitator to the demonstrator’s behaviour we

can build up partial solutions to the correspondence problem, i.e. mapping the actions afforded by particular embodiments to achieve similar effects. The result may resemble the original in various degrees, which depend on defining what is “similar” according to the metric used. Granularity is also seen to play an important role. A partial solution to the correspondence problem, if reused, can be viewed as a structured case-based reasoning method for learning how to imitate.

Even for the same embodiment with loose sub-goal granularity the result might not “look the same” but it can satisfy the metric used. For example considering a Queen that moves on a straight line, an imitator Queen might break the single move down into a number of smaller straight moves. Conversely, another Queen could imitate, using a single move, a demonstrator Queen that takes two moves, to achieve horizontal and vertical displacements if these are of equal length.

The current algorithm does not examine beyond the first layer of possible moves i.e. a sequence of more than one moves when considering minimising the metric, and as a result there are cases where the imitator is getting stuck, unable to achieve the desired effect (see Table 1). This could be overcome by applying already known correspondences (based on previous experience) instead of trying to solve the problem every time. For example, although the Knight is unable to immediately reduce the value of the metric for an adjacent square of different color, it could use a combination of three moves to reach exactly that square. That combination cannot be found here as a solution to a past problem due to the greedy nature of the algorithm; the metric is bound to momentarily increase while performing the sequence, although the overall result would eventually be achieved. This can be discovered by examining the history of the moves done in respect to the results (desired or not) that were achieved. Alternative corresponding moves could be accessed in case that e.g. an obstacle is blocking the way, making some choices illegal. A correspondence formally is a relation, not just a function (see Nehaniv and Dautenhahn, in press).

The algorithm at present also does not examine whether a move will satisfy anything else than the current objective. As a result the Queen will require the same number of moves to travel a straight line as the King would, even though a single move would be enough. But when considering the path type granularity for example, making a move that at the same time satisfies more than one sub-goal would be desirable.

The above aspects will be further addressed in forthcoming publications.

References

- Billard, A., and Dautenhahn, K. 1997. Grounding Communication in Situated, Social Robots, In *Proceedings TIMR, Manchester, Towards Intelligent Mobile Robots TIMR UK 97*. Manchester University.
- Byrne, R. W., and Russon, A. E. 1998. Learning by Imitation: A Hierarchical Approach. *Behavioural and Brain Sciences 21*: 667-684.
- Cypher, A. ed. 1993. *Watch What I Do: Programming by Demonstration*. MIT Press.
- Dautenhahn, K. 1994. Trying to Imitate- a Step Towards Releasing Robots from Social Isolation. In *Proceedings From Perception to Action Conference*, Lausanne, Switzerland, September 7-9, 1994, IEEE Computer Society Press, ISBN 0-8186-6482-7, edited by P. Gaussier and J-D. Nicoud, 290- 301.
- Dautenhahn, K., and Nehaniv, C. L. eds. *Imitation in Animals and Artifacts*. MIT Press. Forthcoming.
- Demiris, J., and Hayes, G., 1996. Imitative Learning Mechanisms in Robots and Humans, In *Proceedings of the 5th European Workshop on Learning Robots (Bari, Italy, July 1996)*, 9-16.
- Heyes, C. M., and Ray, E. D. 1999. What is the Significance of Imitation in Animals?, *Advances in the Study of Behaviour 29*.
- Nehaniv, C. L., and Dautenhahn, K. 1998. Mapping between Dissimilar Bodies: Affordances and the Algebraic Foundations of Imitation. *Working Notes of Agents in Interaction - Acquiring Competencies Through Imitation Workshop at Autonomous Agents '98 (Minneapolis, USA -- May 1998)*.
- Nehaniv, C. L., and Dautenhahn, K. 2001 (In press). Like Me? — Measures of Correspondence and Imitation. *Cybernetics and Systems: An International Journal*. <http://homepages.feis.herts.ac.uk/~nehaniv/likeme/>
- Lieberman, H. ed. 2000 Special section on "Programming by Example", *Communications of the Association For Computing Machinery*, March 2000, 43(3):72-114.
- Noble, J., and Todd, P. M. 1999. Is it Really Imitation? A Review of Simple Mechanisms in Social Information Gathering. In *Proceedings of the AISB'99 Symposium on Imitation in Animals and Artifacts*: 65-73. Society for the Study of Artificial Intelligence & Simulation of Behaviour.
- Price, B., and Boutilier, C. 2000. Imitation and Reinforcement learning in agents with heterogeneous actions. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML-2000)*.
- Rudin, W. 1976. *Principles of Mathematical Analysis*, third edition. New York: McGraw-Hill.
- Swarm development group homepage: <http://www.swarm.org>