Truth-from-Trash Learning and the Mobot Footballer

Chris Thornton
Cognitive and Computing Sciences
University of Sussex
Brighton
BN1 9QH
UK

Email: Chris.Thornton@cogs.susx.ac.uk
WWW: http://www.cogs.susx.ac.uk/users/cjt
Tel: (44)1273 678856

January 18, 1999

Abstract

As natural resources become less abundant, we naturally become more interested in, and more adept at utilisation of waste materials. In doing this we are bringing to bear a ploy which is of key importance in learning — or so I argue in this paper. In the 'Truth from Trash' model, learning is viewed as a process which uses environmental feedback to assemble fortuitous sensory predispositions (sensory 'trash') into useful, information vehicles, i.e., 'truthful' indicators of salient phenomena. The main aim will be to show how a computer implementation of the model has been used to enhance (through learning) the strategic abilities of a simulated, football playing mobot.

1 Introduction: the dangling-glove learner

Figure 1 illustrates a simple experiment which can be performed using objects readily available in the home. The main steps in the experiment are as follows. (1) Take a glove and stiffen it by insertion of newspaper; (2) secure the glove to a long, weak spring; (4) suspend the spring from a fixed point; (4) invite a sequence of subjects to 'shake hands' with the glove; (5) Observe that the spring-loaded glove now moves up and down with a characteristic 'hand-shaking' motion.

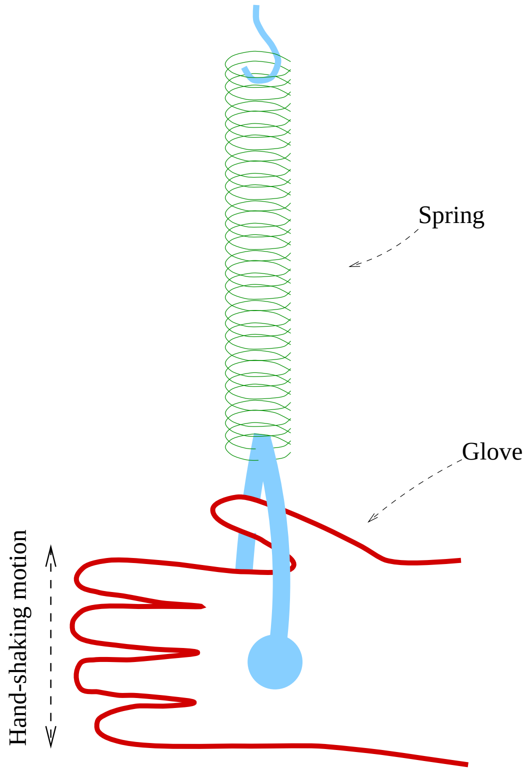

Figure 1: The dangling-glove experiment.

How should we describe the behaviour of the dangling glove in this experiment? One possibility is to say that as a result of the way in which the subjects pulled on the glove, the spring was caused to oscillate for a few moments. This is perhaps the most natural description. But it turns out that, technically, we may also say that, as a result of the physical inputs applied by the three subjects, the dangling glove *learned* to shake hands.

This interpretation seems a little bizarre. But it is, in fact, well justified by the many models of learning (from Machine Learning, Classification, and Connectionism) that view learning as a process in which a dynamic system produces new, useful behaviour, following receipt of environmental stimuli. In such models (e.g., the ubiquitous Backpropagation model of [Rumelhart, Hinton and Williams, 1986]), learning may involve the production of a *representation* of the relevant behaviour. But even in this context our interpretation may apply. We can simply view the modified dynamic behaviour of the glove as a ‘representation’ of the way in which the overall system should generate handshakes. We can even say that in capturing the central tendency of the handshake inputs, the spring has naturally generated a *generalisation* of the relevant behaviour, much in the manner of, say, the LVQ learning method of Kohonen [1988].

This imaginative interpretation of the behaviour of the dangling glove may not be quite within the spirit of the relevant learning models. But the fact that it is not *eliminated* by them raises interesting questions about the role that such models can play in the description of natural learning processes. The suggestion seems to be that for explanatory purposes, attention should shift to more detailed, less permissive models. Many such models already exist, e.g., the ILP [Muggleton, 1992] and constructive induction [Rendell and Seshu, 1990] models. But these often introduce — as an assumed resource — explicit background knowledge of one sort or another, which inevitably limits generality. The aim of the present paper, then, is to present a model of learning which decreases permissiveness without compromising generality. The core of the model will be the idea that learning should be viewed primarily as a *constructive* operation than an *adaptive* one.

2 Learning and the Khepera mobot

Learning may be characterised as the acquisition of new, useful behaviour, cf. [Mitchell, 1977; Kodratoff, 1988]. The process is normally seen as occurring within a particular, concrete agent. However, there is nothing to prevent us from thinking of learning as something which operates at the level of the group or the society. Indeed, evolution is often viewed as a learning process which operates at the level of the species. In these more general situations it is, of course, more difficult to be specific about the boundaries of the agent, and about what sorts of information flow into and out of it.

An example of a simple (concrete) agent is the Khepera robot [K-Team, 1993], see Figure 2. Developed by a Swiss research team, this mobile robot or ‘mobot’ is circular in shape and very small, measuring approximately three inches in diameter. It has a central axis driving two wheels and a trailing castor to keep it upright. Both wheels can be driven forwards or backwards. The mobot can be fitted with various sensory systems. In one configuration it has two forwards-facing infra-red proximity detectors. These produce readings which provide an indication of the distance to the nearest obstruction in a particular direction.

Khepera mobots have played a key role in recent work in artificial life and simulated evolution. At the 1997 European Conference on Artificial Life (held in Brighton), attendees were invited to prepare and enter their own Khepera mobots into ‘football’ competitions. In these events, the aim was for the mobots to push balls towards goal areas while simultaneously preventing oppo-

Figure 2: The Khepera mobot (from the K-team WWW pages).

nent mobots from doing the same thing. (Figure 3 is an extract from the 'call for competitors' WWW page. It shows, on the far left, the basic game scenario. The current Khepera is displayed as a circle with a small disk representing the camera. The viewed Khepera and the ball are shown as a large and a small circle respectively. The three displays on the right show, from left to right, the full image seen by the camera, the image with the viewed Khepera removed and finally the image with the Khepera *and* the ball removed.) The results of the competition were mixed. However, the scenario provides a fertile testing ground for learning-related inquiries.

Consider, for example, the process by which a Khepera mobot might learn to produce sensible attacking moves, i.e., to make sensible decision about when to 'go for the ball' and when to 'hang back.' The learning should produce a disposition in the mobot to try to get possession of the ball only in those situations where sensory inputs indicate that the situation is ripe for doing so. The learning must thus produce a predisposition to produce a certain set of responses to a certain set of external (and possibly internal) stimuli. But what exactly does this mean? And what is involved in the acquisition of such a predisposition?

3 Minimal models and geometric visualisation

To better understand what is required it is helpful to focus on a simple scenario. Consider then the situation in which we have a Khepera mobot equipped with just two light sensors. The behaviour we want it to acquire is the 'attack'

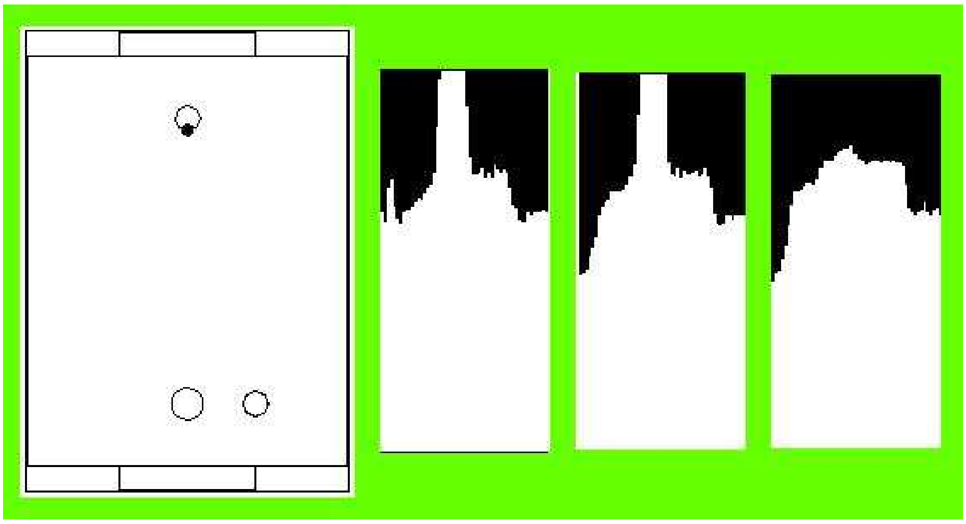

Figure 3: Data for ECAL-97 Mobot Football Competition (Frederic Gruau, WWW document).

behaviour described above. The mobot should acquire a disposition to 'go for the ball' only in situations when the sensory inputs indicate that this action is appropriate. More precisely, we want the mobot to acquire a disposition to produce two distinct actions, namely 'go for the ball' and 'do not go for the ball'.

The advantage of situations in which the agent is in receipt of just two, graded sensory signals (at any one time), is that they are amenable to geometric visualisation. The behaviour to be acquired consists of a set of stimulus responses (i.e., appropriate attack responses to specific combinations of stimuli). These stimulus responses can be visualised as an instantiation of datapoints in a 2-d sensory space, as in Figure 4. The two dimensions of the space here represent the possible levels of input from the two light sensors. Each datapoint represents a stimulus combination in which one of the two actions is appropriate. In the diagram, stimulus combinations appropriate for the 'go for the ball' action have been labelled '1' while stimulus combinations appropriate for the 'do not go for the ball' action have been labelled '0'.

Each datapoint's coordinates are a combination of sensory inputs. Its label is the action which ideally should be produced in response — commonly called the 'target action' or 'target output'. Thus, the diagram shows in pictorial terms which actions are appropriate for which sensory inputs. In effect, it allows one to visualise the pattern of stimulus responses which must be implemented in the

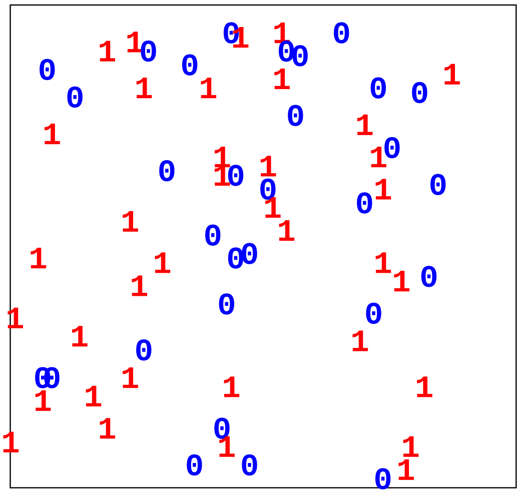

Figure 4: Stimulus-response mapping for attack behaviour.

production of the relevant behaviour.

4 Learning and generalisation

In the context of the 2d diagram, learning means getting the right points in the right places. For an agent with full access to the entire stimulus-response mapping, this can be achieved simply by *memorising* the mapping. The result of this process is the construction in memory of what is known as a **lookup table**. But typically we are not interested in agents which have perfect information of this kind. Rather, we are interested in agents which have to work on the basis of partial information or 'feedback'. In particular we are often interested in agents whose only information about the behaviour is a sample of associations between sensory inputs and target actions. In this situation, to accomplish the relevant learning, the agent must **generalise** from the seen examples to unseen examples. To do this, it must find patterns in the examples which enable accurate guesses to be made about unseen cases.

In terms of the football playing Khepera mobot, a typical learning scenario would involve the Khepera being presented with a set of example stimulus-response associations and then being tested on its ability to produce appropriate responses to other stimulus combinations, i.e., ones not included in the examples. In effect, the Khepera is presented with a subset of the labelled datapoints and then tested on datapoints not included in the sample, see Figure 5.

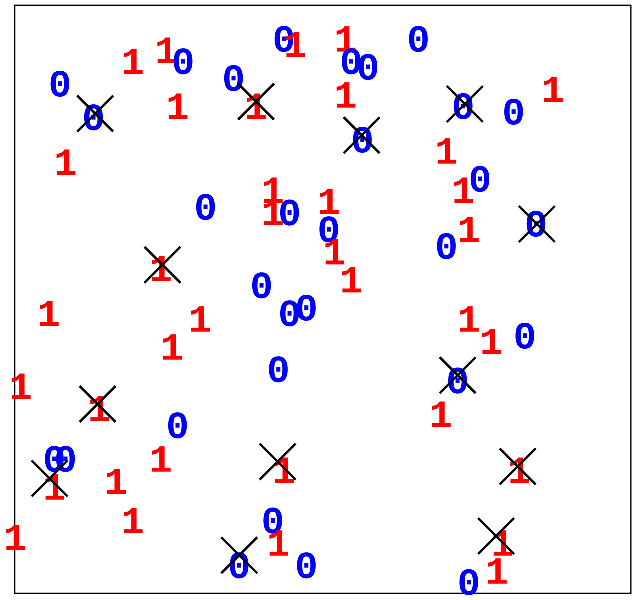

Figure 5: Training set of example S-R associations for ball-attack.

Formally, the two sets of stimulus-response associations (seen and unseen) form a **learning task** or **generalisation problem**. Obtaining a solution to the problem involves producing appropriate generalisations for unseen cases.[1]

5 Learning methods

The close association between learning and generalisation suggests that learning methods need to be successful generalisers. It is no surprise, then, to find that the more frequently encountered learning methods (and models) are those which tend to produce better levels of generalisation. The number of such methods in common currency is surprisingly large. But although each method operates in a slightly different way, when analysed in terms of the geometric visualisation, many are revealed as utilising the same, basic strategy. They all attempt to introduce and/or manipulate simple *boundaries* separating *regions* containing datapoints with the same target action. The rationale here is that once the entire space has been separated out into **uniform** regions, generalisation follows

[1] A common preoccupation of Machine Learning researchers is a learning task called *classification* in which the stimuli are numeric or symbolic values forming a description of some object and the associated actions are correct object identifications.

automatically. The target action for an unseen case can be readily predicted to be the action associated with seen datapoints from the same region. This generic approach is termed **boundary-based** (**BB**) learning, or, less formally, 'fence-n-fill' learning.

6 Rogues gallery of boundary-based learning methods

Boundary-based learners can be divided up into two main groups: methods which add new boundaries and methods which manipulate existing (i.e., predefined) boundaries. These groups can then be subdivided depending on the type of boundary utilised. Some of the best known BB methods can be characterised as follows.

- **PERCEPTRON** [Minsky and Papert, 1988] manipulates a single, linear boundary.
- **ID3** [Quinlan, 1983] introduces an arbitrary number of axis-aligned, extreme boundaries.
- **C4.4** [Quinlan, 1993] follows the same approach as ID3.
- **BACKPROPAGATION** [Rumelhart, Hinton and Williams, 1986] manipulates a fixed number of linear boundaries.
- **LVQ** [Kohonen, Barna and Chrisley, 1990] manipulates a fixed number of spherical boundaries
- **K-NN** [Duda and Hart, 1973] utilises implicit, convex polyhedra surrounding training examples.
- **CASCADE-CORRELATION** [Fahlman and Lebiere, 1990] introduces an arbitrary number of linear boundaries.
- **FOCUSSING** [Bundy, Silver and Plummer, 1985] introduces an arbitrary number of axis-aligned boundaries at user-defined positions.

For further illustrations see [Thornton, 1989].

Of course, in describing a learning method as 'boundary-based' we are not simply saying that the method utilises boundaries. *All* learning method do this in some sense: the utilisation of boundaries thus cannot itself form the basis of any meaningful distinction. The key property in the boundary-based methodology is the utilisation of a *simple* bounding construct — in geometric terms, the utilisation of circles, lines, spheres, planes, hyperplanes and hyperspheres and the like.

Utilising simpler bounding constructs reduces the complexity of the learning process. But there is a hidden and significant cost. The approach succeeds if and only if datapoints with the same action label tend to cluster together in geometrically simple regions. Boundary-based methods effectively pin their hopes on the assumption that datapoints of the same type will cluster together in the same parts of the sensory space. The question is, then, can this assumption can be relied upon in general? Or are there situations in which birds of a feather tend *not* to flock together?

7 Alignment

Different sensory mechanisms respond to different phenomena, i.e., different properties and objects of the environment. But not all sensors respond to all phenomena. Thus there are various relationships a particular sensory mechanism S may have with some particular phenomenon P. These can be characterised in terms of variations in sensor **alignment**:

perfect alignment S explicitly measures or detects P, i.e., signals from S correspond directly to states of P.

perfect non-alignment S does not respond to P in any way.

partial alignment P has an indirect impact on S, i.e., signals from S are affected by states of P but not in any direct, 1-to-1 way

To illustrate these cases, imagine that our footballing Khepera mobot is equipped with a light sensor whose outputs vary monotonically with the amount of light arriving at the sensor surface. With respect to the phenomenon of 'light intensity', the sensor is perfectly aligned. With respect to the phenomenon of 'wind speed' the sensor is perfectly unaligned. And with respect to the phenomenon of 'attack opportunity' (as described above) the sensor has to be considered partially aligned.

Now consider an 'obstacle' sensor. This is perfectly aligned with respect to obstacles, perfectly unaligned with respect to light and partially aligned with respect to 'threat-of-capture' (i.e., the state of play in which an opponent mobot is about to capture the ball).

A partially aligned sensory signal might seem to be much the same as a noisy signal. But alignment and noise are quite different things and the alignment classifications should, in fact, be treated as relating to original, noiseless signals. Thus noise has no relevance to the alignment taxonomy.

8 Salience

For a given behaviour, some properties/objects of the environment are salient and some are not. With respect to a feeding behaviour food may be salient but sand is probably not. With respect to tightrope-walking, gravity is probably salient but UV radiation is probably not. If, in a particular learning scenario, a sensor is perfectly aligned with a phenomenon which is salient for the target behaviour, then particular signals from that sensor will obviously tend to be associated with particular actions. In geometric terms, this means that all the datapoints which belong to particular values (e.g., are in the same row or column) of the relevant dimension in the sensory space will all have the same output label. In a 2d sensory space, if both sensors are perfectly aligned with salient phenomena, then points with the same label will necessarily 'cluster together'.

BB methods, then, are guaranteed to succeed if utilised sensors are aligned with salient phenomena. If the sensors are only partially aligned, clustering is not guaranteed and BB methods are not guaranteed to succeed. If the sensors are perfectly *un*aligned, then *any* learning method should fail since it is attempting to operating without *any* salient information about its environment.[2]

BB learning methods, then, work well if and only if sensors are aligned. In classical Machine Learning, this is expressed by saying that empirical learning methods work well if and only if a 'suitable' input representation is used [Dietterich, London, Clarkson and Dromey, 1982]. But the reliance on perfectly aligned sensors may pose problems for the cognitive scientist.

Complex agents need to be able to learn *many* behaviours which are likely to be contingent upon a wide range of phenomena. Engineers committed to use of BB methods face the prospect of having to equip such agents with large numbers of perfectly aligned sensors. Even if this can be done without irretrievably compromising the agent's viability, there is still the problem of where the sensors are going to come from in the first place. It is not unreasonable to assume that sensory systems for many salient phenomena will remain beyond the 'state of the art' for the forseeable future. Scientists committed to use of BB models encounter a more severe variation of the same problem. Nature tends to exploit general purpose sensory mechanisms (vision, audition, olfaction etc.) which tend to be partially aligned with a wide-range of salient phenomena. Thus explanatory models which rely on the utilisation of perfectly aligned sensors appear to have little hope of achieving full generality.

The implication of this should be that perfectly aligned sensors play a rather limited role in both engineering-oriented and explanation-oriented cognitive science. Unfortunately, due to the widespread utilisation of computer modelling, the opposite seems the case. The researcher who wishes to create an artificial agent (or a model of a natural agent) which learns a behaviour which happens to be contingent upon phenomena not perfectly aligned with any realistic sensory mechanism is likely, as a preliminary exercise, to construct a computer

[2] The distinction between aligned sensory information and partially aligned sensory information is simply the 'sensory' version of the distinction made in [Clark and Thornton, 1997] between statistical and relational data effects. It can also be viewed as a variant of the distinction between statistically independent signals and statistically dependent signals.

simulation. Surprisingly enough, this may appear to demonstrate that the behaviour *can* be successfully learned using a standard (BB) learning method, eg. Backpropagation or C4.5. Papers may be published and readers duly impressed.

However, on closer inspection, it may well turn out that the successful learning performance is really attributable to the fact that the programmer has equipped the simulated agent with special-purpose (i.e., perfectly aligned) sensors, only feasible within the context of computer simulation. The researcher in this case is said to have utilised **magic sensors**. The work has not really demonstrated a realistic way in which the relevant behaviour can be learned. It has merely provided an illustration of the way in which computer simulations can mislead.

Work based on simulations which utilise magic sensors should, by rights, have a low currency in the scientific domain. But there seems to be a general lack of awareness of the key role played by sensory alignment in learning simulations. For many researchers, the question of sensory configuration does not figure in assessments of the sophistication of a learning model or of the complexity of specific problem scenarios. Very often a learning problem is considered difficult if and only if it can be viewed as having something to do with the 'real world'. Artificial learning problems (like learning to do addition) thus tend to be considered 'easy' while real-world problems (like learning to diagnose plant diseases) tend to be considered 'hard'. (For an example of this, see [Holte, 1993].)

But all too often, these assessments are quite at odds with the truth. When boundary-based learning methods are used, the effective complexity of a problem is primarily a function of sensory alignment. It has little or nothing to do with the problem's *context*. A problem involving something that seems to us easy (like learning to distinguish even from odd numbers) may be utterly intractable to a BB learner. Conversely, a problem which seems to us rather difficult (learning to distinguish those states of a jet engine likely to lead to failure) may be trivial. It all depends on how the inputs are presented, i.e., what assumptions are made about the alignment of the learner's 'sensory' mechanisms. This is of course just another version of the old Computer Science ruling that the difficulty of a problem depends on the way in which it is *represented*.

9 Truth from Trash

The utilisation of magic sensors in computer simulations, then, should be firmly avoided. But what alternatives does this leave us with? Are we restricted to modelling learning only in those few contexts where the utilisation of perfectly aligned sensors is realistic? To answer this question we need to look carefully at the way in which partially aligned sensors instantiate stimulus-response spaces. For simplicity, we stick with the original mobot-football 'attack' data (Figure 5).

Note that, at first glance, the two types of datapoint in the 2d sensory space appear to be scattered at random throughout the sensory space. This is what we expect. The sensors are partially aligned with the behaviour in question and there is therefore no reason to expect datapoints of the same type to cluster together. However, on closer inspection we see that the distribution is not entirely without order. There is some local clustering which arises as a result of accidental factors.[3]

The phenomenon which is salient for this behaviour is 'attack opportunity'. This is an abstract property of football scenarios and we therefore consider both of the (light) sensors to be partially aligned. On this assumption we might write off the local clustering as worthless random variation, i.e., noise. But we need to remember that the datapoints themselves are not the only source of information available to the learner. There is also the labelling or 'environmental feedback' which tells the learner which action goes with which stimulus combination. If this information is fully utilised, the noise can be brought under control to form the foundations for an approximation of a 'true' sensor for the salient phenomenon.

The means of doing this is a simple procedure involving three steps:

(1) A BB learning method is used to process the data, i.e, to find regions of uniformly labelled stimuli.

(2) A mapping is constructed from the input space to a **derived space**, which has the effect of drawing same-label regions closer together.

(3) The first two steps are reapplied recursively to the derived data, until derived data are generated which can be satisfactorily processed solely by the BB learning method.

The general effect of this cyclical process is to produce a sequence of **recodings** of the original data, each of which exhibits an increased level of global organisation (i.e., more pronounced clustering). Within the process, uniformly labelled datapoints are caused to 'gravitate' towards each other. But at no stage is there any application of background knowledge or explicit bias. No oracle is queried concerning the pros and cons of different attacking configurations. The operation is 'mindless.' At each point it simply attempts to compress a little more organisation out of the statistical 'trash' which still remains. Hence the name 'Truth from Trash.'

The number of recodings that will need to be derived in a particular case depends on the problem, of course, but also on the complexity and sophistication of the recoding step. With a more sophisticated recoding operation, we can expect a greater enhancement of global organisation to be achieved at each step.

[3] It can be shown theoretically that there is always some local clustering in such spaces unless the problem is 'perfectly relational', i.e., a parity or modulus-addition problem [Thornton, 1996b].

But we inevitably pay a price in terms of bias and operating costs. With a less sophisticated operation, we will expect more recoding steps to be required but that each one will cost us less. Thus there is a kind of depth v. width tradeoff in operation. A sophisticated recoding step means that only a shallow hierarchy of recodings will be required. But the recoding operation itself is rather complex giving us a result which is 'short and fat'. With a more primitive recoding step, a larger hierarchy of recodings will be required but the recoding step is simple, giving us a result which is 'long and thin'.

10 Skeletal Exemplars in Constructive Structures (SECS)

Any method which operates in the incremental manner described above can be deemed a **truth-from-trash** or **TFT** method. But there are many possible implementations. Every pairing between a BB learning method and a plausible recoding strategy provides a unique TFT method. The choice of a particular TFT method is thus something of a problem in itself. In my initial experiments I have opted for a simple approach and used a TFT algorithm which I call **SECS** (**Skeletal Exemplars in Constructive Structures**).

The core of the algorithm is a simple (but apparently novel) boundary-based learning method here termed **Skeletal Exemplars** (**SE**). This operates rather like the k-nearest-neighbours method [Duda and Hart, 1973]. However, while nearest-neighbour methods typically retain in memory the complete set of presented examples, the SE method only retains datapoints whose nearest neighbours have the same label (i.e., action/output). Moreover, it minimises the redundancy of the set by deleting datapoints which are not needed for the purposes of making correct classifications (by the nearest-neighbour rule) of seen data. The datapoints retained are termed 'skeletal exemplars' or just 'exemplars'.

The SECS method uses the SE method combined with a simple recoding operation. In this, the the current datapoints are ordered according to their associated output. (This assumes a predefined ordering on actions, e.g., that actions are represented numerically.) The algorithm then generates derived data in which the first coordinate of each derived datapoint is the linear position within the ordering of the original datapoint's nearest exemplar (which may be the datapoint itself) and whose second coordinate is the proximity of the original datapoint to its nearest exemplar. The general effect is to 'spread' the datapoints out across a 2-dimensional space, according to target action. In a simple, 2-output scenario like the 'attack' problem, the general effect is to recode the data into a kind of 4-part flag arrangement.

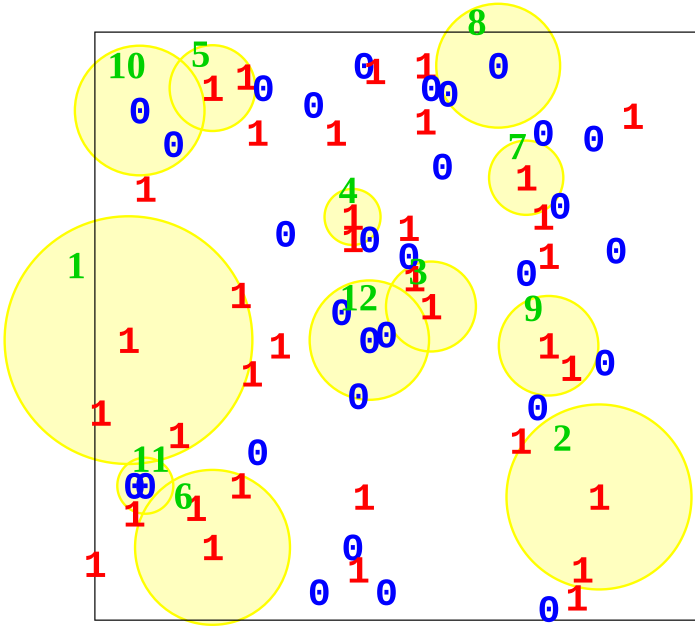

Figure 6: Implicit clustering through derivation of skeletal exemplars.

11 Worked example

The operation of the SECS method can be illustrated using a worked example. To begin with, note how the skeletal exemplars method introduces an implicit clustering of the attack-opportunity training set, see Figure 6. In my current implementation the SE method is implemented as follows. The datapoints are first ordered according to number of 'friends', where a datapoint p is considered to be a friend of datapoint q if it has the same target action/output as p and is closer to p than any datapoint with a different target action. It then extracts datapoints from the ordered training set until either every datapoint has been itself extracted or is a friend of an extracted datapoint. The process also terminates if it is found that the next available datapoint has no friends. The general effect of this procedure is to identify clusters of same-action datapoints and their 'central' member. In Figure 7, each such grouping is shown within a circle centered on the key member. Note how the process has produced an implicit clustering of 11 clusters.

Figure 7 illustrates the way in which the derived clustering performs when used for generalisation. Each circle of friends now has an enveloping outer circle which encloses all those datapoints which have the key (central) member as

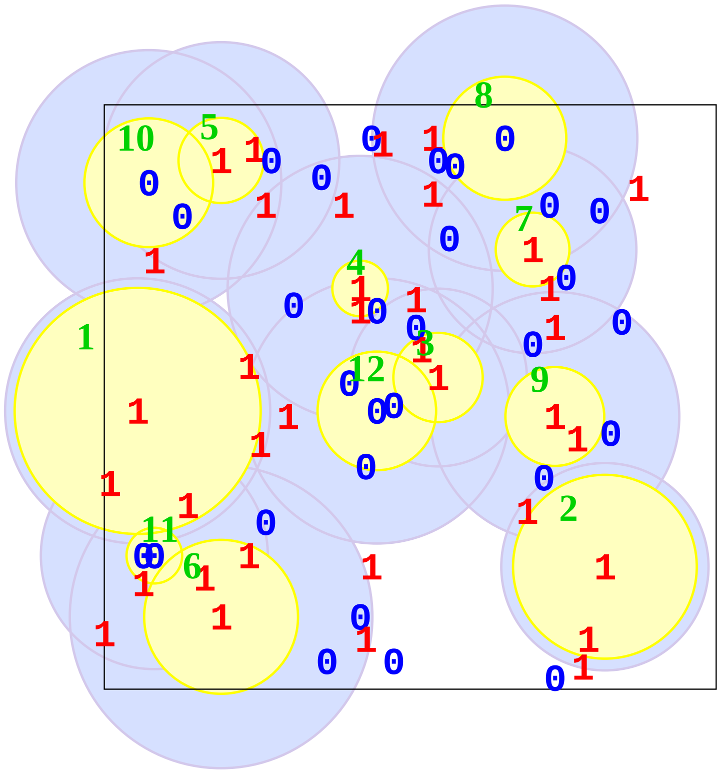

Figure 7: Generalisation performance.

their nearest exemplar. The generalisation performance of the method can thus be derived by counting the number of datapoints whose nearest exemplar has a *different* target action.

Using the same graphic conventions, Figure 8 illustrates the entire processing cycle of the SECS method applied to this dataset. Instead of a single rectangle representing the original sensory space, we now have three rectangles representing the original space and two internal recodings derived during processing. The top-right rectangle is the first derived space and the bottom-left space is the second derived space. Note how the general effect of the recodings is to 'push' positive datapoints (1s) up and to the left and the negative datapoints (0s) down and to the right. In the second derived space, we have a relatively good separation between the two groups. The application of the current BB learning method (skeletal exemplars) thus identifies some exceptionally large, and well populated clusters. The generalisation performance obtained is considered

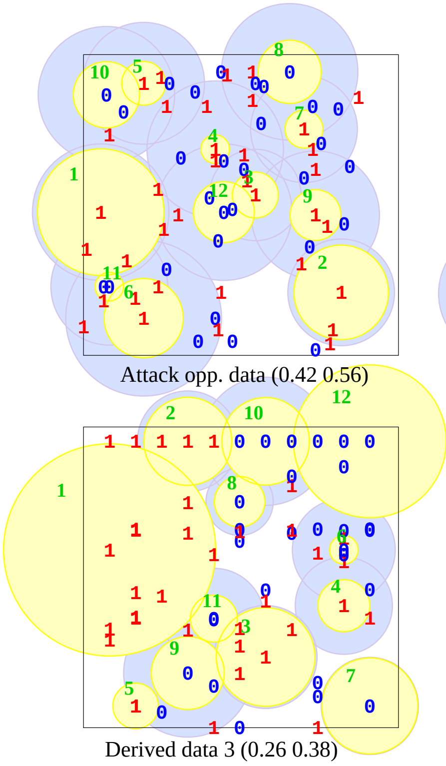

Figure 8: Processing cycle of the SECS method.

satisfactory and processing terminates.

What, then, is the cash value of this rather elaborate learning method? Equipped with a SECS implementation, a Khepera mobot can successfully improve its performance with respect to a higher-level property of the mobot-football game, i.e., a property to which its (light) sensors are only partially aligned. 'Attack-opportunity' is just one of a large number of properties that the truly successful footballing mobot must be sensitive to. Thus the acquisition of this particular functionality in a Khepera is unlikely to make much difference to its overall goal scoring ability. But the value of the approach is that it shows in a mechanistic way how a learning agent might begin to move beyond the 'dictatorship of the stimulus' without resort being made to magical sensory equipment or fine-tuned knowledge bases.

12 The role of the Residual Agent

The TFT model provides a picture of a way in which an agent might learn behaviours contingent upon phenomena not directly sensible by realistic sensory mechanisms, without the need for covert introduction of magical sensory equipment. The main flavour of the idea is the utilisation of the noise or statistical trash which arises at the 'interface' between a particular partial-alignment relationship and a particular behaviour.

Whether the model has any engineering value is yet to be determined. However, it does have a novel explanatory flavour that may make it attractive to those more interested in description and conceptualisation. The nuts and bolts of the model are essentially 'algorithmic' and 'computational'. But the nature of the processes described are sufficiently primitive that they could be re-rendered in a connectionist or neural-networks paradigm. The model is not 'representational' in the classical sense since it makes no use of explicit representational structures (frames, databases, default inheritance, explicit symbols and the like). And yet it does suggest a role for the *process* of representation since it shows how a learning agent can construct internal sensors which measure external, but implicitly-sensed properties of the environment.

As I have previously argued [Thornton, 1996a] these inner sensors are conveniently viewed as **virtual sensors**. Insofar as their signals are *used* by the 'residual agent' (i.e., the parts of the agent not engaged in implementing the virtual sensor) as a sign of an external phenomenon, they have a clear, though slightly counter-intuitive representational status. But to understand the nature of this idea, we must be ready to see the agent as divisible into two parts: the part which implements the sensor and the part which utilises its signals. Once this leap has been made, the TFT model is revealed as providing an interesting route via which representational models of cognition might be grounded in *non*-representational, potentially neural processes.

13 Concluding comment: The Neat-Scruffy Mind

Learning in the TFT model is viewed as something which builds veridical representational signal sources out of what is, in effect, a cascade of kluges. The results of a TFT process are thus something like a 'Rube Goldberg' machine — it works OK in practice but on close inspection, the innards turn out to be a weird assembly of uninterpretable fixes. An interesting property of this viewpoint is the way in which it relaxes the tension between the 'neat' and 'scruffy' philosophies of cognition. It suggests that in certain situations cognisers can be viewed as attempting to develop 'scruffy' means of supporting 'neat' pretences about the ways in which they are coupled to their environment. On this view, if cognition can be said to *be* anything in particular, it might be said to be both neat and scruffy at the same time.

References

[1] Duda, R. and Hart, P. (1973). *Pattern Classification and Scene Analysis.* New York: Wiley.

[2] Mitchell, T. (1977). Version spaces: a candidate elimination approach to rule learning. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (pp. 305-310).

[3] Dietterich, T., London, B., Clarkson, K. and Dromey, G. (1982). Learning and inductive inference. In P. Cohen and E. Feigenbaum (Eds.), *The Handbook of Artificial Intelligence: Vol III.* Los Altos: Kaufmann.

[4] Quinlan, J. (1983). Learning efficient classification procedures and their application to chess end games. In R. Michalski, J. Carbonell and T. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach.* Palo Alto: Tioga.

[5] Bundy, A., Silver, B. and Plummer, D. (1985). An analytical comparison of some rule-learning programs. *Artificial Intelligence, 27,* No. 2 (pp. 137-81).

[6] Rumelhart, D., Hinton, G. and Williams, R. (1986). Learning representations by back-propagating errors. *Nature, 323* (pp. 533-6).

[7] Kodratoff, Y. (1988). *Introduction to Machine Learning.* London: Pitman.

[8] Kohonen, T. (1988). *Self-organization and Associative Memory* (second edition). New York: Springer-Verlag.

[9] Minsky, M. and Papert, S. (1988). *Perceptrons: An Introduction to Computational Geometry* (expanded edn). Cambridge, Mass.: MIT Press.

[10] Thornton, C. (1989). Learning mechanisms which construct neighbourhood representations. *Connection Science, 1,* No. 1 (pp. 69-85).

[11] Fahlman, S. and Lebiere, C. (1990). The cascade-correlation learning architecture. In D.S. Touretzky (Ed.), *Advances in Neural Information Processing Systems 2* (pp. 524-532.). Morgan Kaufmann Publishers, Los Altos CA.

[12] Kohonen, T., Barna, G. and Chrisley, R. (1990). Statistical pattern recognition with neural networks: benchmarking studies. In J.A. Anderson, A. Pellionisz and E. Rosenfeld (Eds.), *Neuocomputing 2* (pp. 516-523). The MIT Press.

[13] Rendell, L. and Seshu, R. (1990). Learning hard concepts through constructive induction. *Computational Intelligence, 6* (pp. 247-270).

[14] Muggleton, S. (Ed.) (1992). *Inductive Logic Programming.* Academic Press.

[15] Holte, R. (1993). Very simple classification rules perform well on most commonly used datasets. *Machine learning, 3* (pp. 63-91).

[16] K-Team, (1993). *Khepera users manual.* Lausanne: EPFL.

[17] Quinlan, J. (1993). *C4.5: Programs for Machine Learning.* San Mateo, California: Morgan Kaufmann.

[18] Thornton, C. (1996a). Re-presenting representation. In D.M. Peterson (Ed.), *Forms of Representation: An Interdisciplinary Theme for Cognitive Science* (pp. 152-162). Intellect.

[19] Thornton, C. (1996b). Parity: the problem that won't go away. In G. McCalla (Ed.), *Proceeding of AI-96* (Toronto, Canada) (pp. 362-374). Springer.

[20] Clark, A. and Thornton, C. (1997). Trading spaces: computation, representation and the limits of uninformed learning. *Behaviour and Brain Sciences, 20* (pp. 57-90). Cambridge University Press.