

Truth-from-Trash Learning and the Mobot Footballer

Chris Thornton

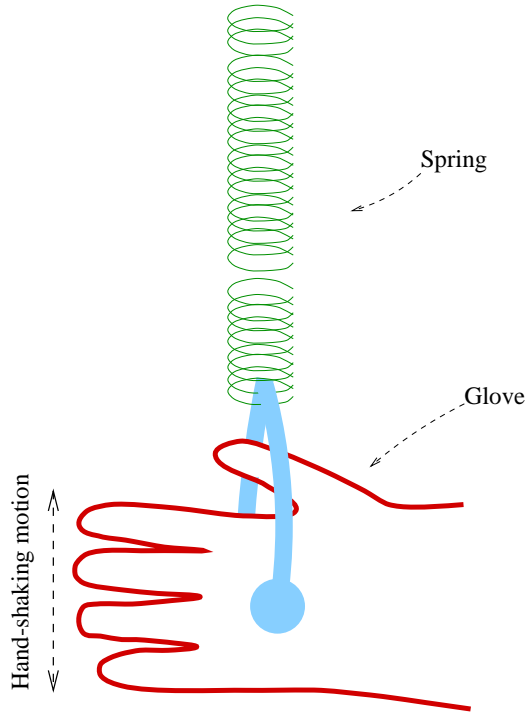
Cognitive and Computing Sciences
University of Sussex
Brighton
BN1 9QH
UK

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



‘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 sugges-



Figure 2:




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

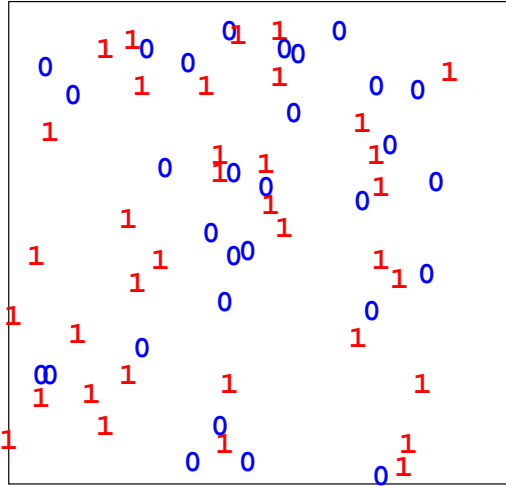



Figure 4:



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, ex-

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 Saliency

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 *unaligned*, 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 foreseeable 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

²The 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.

Note that, at first glance, the two types of datapoint in the 2d sensory space

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

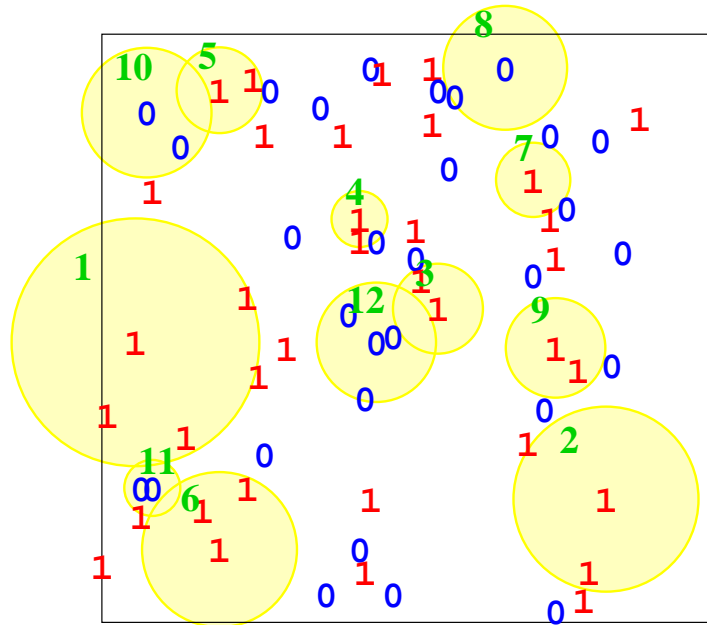


Figure 6: Implicit clustering through derivation of skeletal exemplars.

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,

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.