Creativity and Runaway Learning

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

Email: Christopher.Thornton@firenet.uk.com
Tel: (44)1273 678856

May 29, 2003

1 Introduction

[text slightly updated from the version that appears in the Dartnall volume]

This chapter adopts a strictly *non-empirical* strategy. Rather than offering a model of the experimental data, such as it exits, the aim will be to carry out a logical analysis of the operational characteristics of basic learning procedures and to use this analysis to tease out some interesting facts about the relationship between creativity and learning.

The key idea to be worked out is that our ability to be creative rests partly on our ability to *learn*. I will argue that certain creatve processes (although certainly not all) may be viewed as learning processes running 'out of control.' The notion of a connection between learning and creativity is nothing new, of course. Many authors have identified close relationships between creativity, abduction and analogy, See for example (1), (2), (3) and (Humel and Holyoak, this volume). But this paper will try to put some new flesh on the relationship by showing that a particular type of creativity may be understood in terms of a particular type of learning. The aim will be to show that what one might term intellectual or scientific creativity is naturally viewed as a form of empirical learning extended beyond the normal boundaries of objectivity. As mentioned, the argument will be based, not on any empirical model but rather on a logical analysis of the learning task. The reliability of the conclusions drawn will thus be primarily a function of their analytical coherence.

2 Learning analysed as a task

Definitions and characterisations of the process(es) we call 'learning' vary widely. Clearly there are many different forms of the process each of which may operate on a characteristically distinct basis. But we can cut away some of the complexity by focussing attention, not on the process itself, but on the *task* that it addresses.

All forms of learning involve some form of behaviour acquistion. In any given behaviour — which may range from a simple motor skill to an abstract cognitive transition — certain 'actions' must be produced in certain 'situations'. And it is, of course, these contingency relationships which give the behaviour its distinct character. The problem which is solved in a learning event, therefore, is the establishment of a certain set of contingengy relationships, or, in computational terms, the implementation of a certain set of IF-THEN assocations.

At this level of abstraction, no judgements need to be made about the objects involved in the contingency relationships.[1] Any particular contingency relates to some particular set of antecedent objects. But it can do so in two different ways. The contingency may relate to an *absolute* property of the antecedents or to a *relative* property. Putting it another way, the IF-THEN rules may test for absolute properties (e.g., does antecedant A have property P?) or it may test for relative properties (e.g., do antecedant A and B show relationship R?).

What this tells us is that learners may follow two different strategies. They may pursue a **relational learning** strategy, in which case the identified contingencies are tagged to relational properties of antecedant objects. Or they may pursue a **non-relational learning** strategy, in which case the contingencies associate are tagged to non-relational properties of antecedant objects. Interestingly, this conclusion is obtained without anything being known about how learning works *in practice*. The choice of relational and non-relational strategies of operation is a property of learning tasks and is thus *necessarily* faced by any learning procedure, regardless of its implementation or biological origin.

As a matter of fact, it turns out that practical learning methods, insofar as they have been identified and characterised, *do* tend to divide up into relational and non-relational groupings (4). Relational learning is, in fact, generally considered to be an independent subfield of machine learning research focussing on discovery-oriented learning procedures. The analytical conclusion relating to the relational/non-relational bifurcation thus do seem to carry over into the empirical domain.

[1] As 'behaviourists' we might assume that the antecedent objects are sensory stimuli and the consequent objects are motor outputs. As 'cognitivists' we might say that the antecendents objects are elements of internal (mental) state and the consequent objects representational constructs.

3 The need for recursive relational learning

A key observation for the argument to be developed is that the identification of a relationship within certain data effectively *recodes* those data. For example, imagine that the data are integer values and the relationship identified is *greater-than.* The identification offers us a recoding of the data in which the original set of integer values are transformed into a set of truth values, these being the results of evaluating the greater-than relationship within the given data.

Relational learning is thus a process which *potentially* generates new data. And like the snake which faces the possibility of eating its own tail, the relational learner has the choice of applying *itself* to the data that it, itself, generates. We can say that it definitely will do so. But we can easily envisage situations in which 'recursion' of this form will be advantageous.

For example, imagine that the antecedant objects (in some behaviour) exist at one or more levels of description *below* that of the consequent objects. (This would occur, for example, if the antecedant objects were retinal stimuli and the consequent objects were actions involving real-world objects). On the assumption that an object at one level of description is made up of a set of *related,* lower-level objects, we know that contingencies in this situation necessarily associate consequent objects with some *relational* property of the antecedant objects. If there are several levels of description interposed between antecedant and consequent levels, then we know that the contingencies associate consequent objects with a hierarchy of relational effects, i.e., exactly the structure which can only be identified through a process of *recursive* relational learning.

The implications of this are quite interesting. If learning ever confronts a situation in which antecedant and consequent objects are at different levels of description — and this is, arguably, the *expected* situation for agents using low-level stimuli for purposes of dealing with macroscopic objects — recursive relational learning of some sort is a *necessity.*

To summarise the last few paragraphs:

(1) Learners may utilise relational or non-relational strategies.

(2) If learning agents tend to use low-level stimuli for dealing with high-level objects and events, then the level of description of antecedant objects will typically be several levels lower than that of consequent objects. In this case, learners must necessarily utilise a recursive, relational strategy.

The bottom line, here, is that *realistic* learners will tend to pursue the recursive relational strategy. We should therefore expect to see such learners engaging in an internal process of *self-stimulating data generation.* But in executing this strategy, learners may be doing much more than mere learning.

Note that a learner pursuing any sort of relational strategy has to be able to identify relationships. This type of learner therefore must utilise 'knowledge' of some sort concerning possible relationships. Succeeding at recursive relational learning involves applying the *right* relationships in the *right* way. Successful learners will therefore be those which attempt to identify those relationships which are 'really there', i.e., actually instantiated in the learner's sensory data stream. The learning is then tantamout to the incremental discovery of successive levels of description — a kind of constructive representation-building operation.

Further to this, it is natural to assume that a learner of this type will continue the learning process up to the point at which the antecdent and consequent levels of description have been brought into alignment. At this point the relevant contingencies effectively cease to have a relational basis and there would seem to be no point in continuing. However, termination at this point is not *inevitable*. From the theoretical point of view there is no reason why the process should not conintue on. In practice, it may even be difficult to determine when full alignment has been achieved, meaning that the termination criterion will be impossible to apply.

But letting the process run on is likely to have some unexpected effects. The generation of internal data will continue to instantiate successive levels of description but these levels of description will now become increasingly divorced from reality. Initially, the discontinuity will be relatively modest. There will be some initial, post-termination, level of description containing objects made up of components — themselves real objects — tied together within familiar and realistic relationships. However, as the data-generation process continues, the familiarity and realism of the objects populating new levels of description will gradually diminish. Eventually a point will be reached where the process is effectively 'fantasising', i.e., building constructs out of 'imaginary' objects tied together in 'unwordly' relationships.

What is happening, here, is that our learning process is starting to runaway with itself. It is beginning to flesh out the possible hiearchical constructs which can be built out of known objects and relationships. Put another way, the process is beginning to explore the space of possible ways of reconstructing reality. This looks to be a form of activity which is, in some sense, 'creative'. And in fact, when we come to examine the dynamics more carefully we find that the process is, indeed, a fairly plausible instantiation for the creativity model of authors such Boden (5) who depict creativity in terms of the exploration of 'conceptual spaces'.

Boden views creativity as an activity involving the development and exploration of new conceptual spaces. This is effectively the task performed by recursive, relational learning in the runaway phase. The learning strategy thus meets Boden's creativity criterion fairly well, the implication being that our process model has a split personality. On the one hand it is a learning procedure dealing with certain types of problematic behaviour. On the other hand it is a creative procedure pursuing the exploration and development of new conceptual spaces. The learning personality is built-in by design but the creative personality is a simple consequence of the fact that termination in recursive learning is a strictly contingent issue.

4 Example

Let me illustrate the creative effects of runaway relational learning with an example. Consider the Poker dataset shown in Table 1. Each row in this

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10		x11
(1)	13	2	2	3	8	2	2	1	2	4	-->	4
(2)	8	3	6	4	6	2	8	3	8	1	-->	5
(3)	12	1	5	2	3	3	3	2	3	1	-->	4
(4)	13	4	13	3	8	2	8	1	8	3	-->	5
(5)	9	3	10	1	11	2	12	1	13	4	-->	7
(6)	10	4	10	3	1	3	1	4	10	2	-->	5
(7)	13	4	11	4	11	3	13	4	13	4	-->	5
(8)	9	2	4	2	5	2	13	2	10	2	-->	6
(9)	7	4	12	4	12	2	4	2	12	1	-->	4
(10)	13	2	8	2	1	3	1	3	1	4	-->	4
(11)	10	3	10	1	5	2	13	2	10	2	-->	4
(12)	13	4	3	4	4	1	3	4	3	4	-->	4
(13)	11	2	8	4	4	4	4	2	4	4	-->	4
(14)	11	3	11	4	13	1	13	1	13	3	-->	5
(15)	2	3	2	1	2	1	2	2	1	4	-->	9
(16)	8	2	2	2	9	2	11	2	13	2	-->	6

Figure 1: Simplified poker data.

table is an input/output pair representing a particular hand in the card game Poker. Each odd-numbered variable contains the face value for a particular card, represented as a number (where 10=jack, 11=queen, 12=king and 13=ace). The adjacent even-numbered variable holds the corresponding suit value (1=hearts, 2=spades, 3=clubs and 4=diamonds). Values of the output variable represent the rank of the hand, using the following scheme.

```
1 nothing
2 two of a kind (a pair)
3 two pairs
4 three of a kind
5 full house
6 flush
7 straight
8 straight flush
9 four of a kind
10 royal flush
```

Let us think about the way in which a relational learner might tackle these data. Initially, it might try different ways of applying relevant relationships to the data. For example, the learner might search for instantiations of the *equality* relationship. As a first step, the equality of all 10 input values might be considered. In the data at hand, there are, in fact, no cases in which all the input values are identical. Thus the application evaluates to false in all cases. A next step might be to try applying the equality relationship across subsets of variables. This produces slightly more interesting results. Cases 8 and 16 turn out to both exhibit equality among the same selection of variables. In fact, every even-numbered variable in both data has the same value.

The learner might also investigate the possibilities of applying the equality relationship in a more flexible way. In particular, it might see whether there are cases in which arbitrary selections of *values* from particular data satisfy the relationship. This approach might lead to the discovery that each of the cases 9, 10, 11, 12 and 13 contains at least three identical values, i.e., three values which mutually satisfy the equality relationship. (It might also reveal that in each case the relevant values come from the odd numbered variables.)

These initial experiments with equality reveal that even using a single relationship in a single application protocol, there are many different results that can be produced. A tendency to focus on equality relationships among specific variables tends to lead to the identification of effects corresponding to the existence of *flushes*, since in these cases the values exhibiting equality will always be sited in the same variables. A tendency to consider equality among arbitrary collections of values, on the other hand, leads more directly to the identification of n-of-a-kind hands.

In a recursive learning process the effects of any bias (i.e., search strategy used) are cumulative. A bias which leads to the identification of flushes will — in the next round — lead to the identification of hands which are built out of flushes, e.g., *royal flushes*. A bias favouring n-of-a-kind hands, conversely, will tend to lead to the identification of hands which are built out of n-of-a-kind hands, e.g., *full houses*. The learner with one bias thus discovers one set of phenomena in the data, while the learner with a different bias finds a different set. In poker the possibilities are particularly numerous. A learner whose bias gave it a predisposition to consider integer *sequences*, for example, would be led to discover objects such as *straights* and *runs*.

Relational learners thus always have a hand in the creation of their own sensory phenomena. To some degree, they 'worlds' they inhabit are self-built. The original data start out as the most influential factor. But as the learning proceeds, the accumulated consequences of the learner's own bias become increasingly dominant. The learner is led in a particular direction and discovers a certain subset of relational effects in the data. A given relational learner confronted with the Poker data will thus identify a certain *subset* of the full range of Poker hands, these being the hands that may be constructed using relationships available to the learner.

But let us imagine that we have an *ideal* learner which succeeds in generating the full range of known Poker hands, and ask what will happen if the learning is not terminated at this point. The process of searching for relational effects continues on. But now it applies to internal data generated in previous iterations and the objects over which relationships are sought are identified *patterns* of card combination. The inevitable consequence is the generation of imaginary poker hands: 'full-houses' comprising three-card straights and two-card runs; 'runs' involving ascending sequences of pairs; 'flushes' involving alternating suits. And so on. The mechanical procedure used to discover genuine Poker hands now begins to generate imaginary hands.

What is happening, in effect, is that the process is starting to explore the space of possible 'Poker-like' hands. It is doing this in a relatively unguided way. But there is an element of selectivity, in the sense that new objects are explored contingently depending on their grounding in the original data. This selection process might be viewed as a rudimentary form of 'aesthetic filtering' which prioritises new interpretations of reality over pure fantasy.

Obviously the example is a very simple one. But it should be sufficient to show how the creative potential of runaway learning may be cashed out in a reasonably realistic context. It should also clarify what type of creativity is actually generated by the process.

5 Concluding comments

The paper has presented a logical analysis of learning which suggests that the learning process has the potential to make a transition from a characteristically objective to a characteristically subjective mode. Once past this transition point, the products of the process have an increasingly creative character while the dynamics become increasingly interpretable in terms of the conceptual-exploration activity offered by Boden and others as a model of generic, intellectual creativity. According to this view, then, learning and creativity are to be treated as extreme points on a single dimension. And we should begin to think of both in terms of a generic, constructive procedure whose operations are initially 'close to the ground' but later achieve something akin to 'free flight'.

From the technological point of view the argument suggests that those wishing to devise artificial, creative agents should pay close attention to work going on in the learning sciences. The scientific implications are also worthy of note. The argument provides us with the beginnings of what looks to be a theoretically well-grounded process model for certain types of creative action. Whether this model can be provided with any empirical support remains to be seen.

6 References

Koestler, A. (1). THE ACT OF CREATION. London: Hutchinson.

Langley, P., Zytkow, J., Simon, H. and Bradshaw, G. (3). The search for regularity: four aspects of scientific discovery. In R. Michalski, J. Carbonell and T. Mitchell (Eds.), MACHINE LEARNING: AN ARTIFICIAL INTELLIGENCE APPROACH: VOL II (pp. 425-469). Los Altos: Morgan Kaufmann.

Boden, M. (6). THE CREATIVE MIND: MYTHS AND MECHANISMS. London: Weidenfeld and Niicolson.

Schaffer, S. (7). Making up discovery. In M.A. Boden (Ed.), DIMENSIONS OF CREATIVITY (pp. 13-52). MIT Press.

Clark, A. and Thornton, C. (8). Trading spaces: computation, representation and the limits of uninformed learning. BEHAVIOUR AND BRAIN SCIENCES, 20 (pp. 57-90). Cambridge University Press.

References

[1] Koestler, A. (1964). *The Act of Creation*. London: Hutchinson.

[2] Schaffer, S. (1994). Making up discovery. In M.A. Boden (Ed.), *Dimensions of Creativity* (pp. 13-52). MIT Press.

[3] Langley, P., Zytkow, J., Simon, H. and Bradshaw, G. (1986). The search for regularity: four aspects of scientific discovery. In R. Michalski, J. Carbonell and T. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach: Vol II* (pp. 425-469). Los Altos: Morgan Kaufmann.

[4] 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.

[5] Boden, M. (1990). *The Creative Mind: Myths and Mechanisms*. London: Weidenfeld and Niicolson.

[6] Elman, J. (1990). Finding structure in time. *Cognitive Science*, *14* (pp. 179-211).

[7] Holsheimer, M. and Siebes, A. (1994). Data mining: the search for knowledge in databases. Technical report CS-R9406, CWI.

[8] Thornton, C. (1997). Separability is a learner's best friend. In J.A. Bullinaria, D.W. Glasspool and G. Houghton (Eds.), *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representations* (pp. 40-47). London: Springer-Verlag.