THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Relational learning re-examined

Citation for published version:

Thornton, C & Clark, A 1997, 'Relational learning re-examined', *Behavioral and Brain Sciences*, vol. 20, no. 1, pp. 83-83. <https://doi.org/10.1017/S0140525X97440025>

Digital Object Identifier (DOI):

[10.1017/S0140525X97440025](https://doi.org/10.1017/S0140525X97440025)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Behavioral and Brain Sciences

Publisher Rights Statement:

© Thornton, C., & Clark, A. (1997). Relational learning re-examined. *Behavioral and Brain Sciences*, 20(1), 83-83doi: 10.1017/S0140525X97440025

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Evolution's gift is the right account of the origin of recoding functions

Andrew Wells

Department of Social Psychology, The London School of Economics and Political Science, London WC2A 2AE, England. a.j.wells@lse.ac.uk

Abstract: Clark & Thornton argue that the recoding functions which are used to solve type-2 problems are, at least in part, the ontogenetic products of general-purpose mechanisms. This commentary disputes this and suggests that recoding functions are adaptive specializations.

Clark & Thornton (C&T) have enhanced our understanding of the nature of problem spaces with their distinction between type-1 and type-2 problems. Type-1 problems are solved directly. Type-2 problems are solved indirectly via recoding functions which make manifest otherwise hidden regularities in input domains. Recoding reduces type-2 problems to type-1. This is a valuable insight and the statistical framework within which the distinction is formalized is both apt and informative. C&T suggest that type-2 problems are commonplace in a wide variety of domains but are regularly solved despite their computational intractability, which results from the infinity of potential recoding functions.

A central question, therefore, concerns where the recoding functions used to solve type-2 problems come from, and it is here that C&T's analysis is problematic. One possibility, which is currently at the heart of much promising work in evolutionary psychology (Barkow et al. 1992), is that recoding functions are task-specific, evolved adaptations. C&T are clearly unsympathetic to this idea, which they call "heavy-duty nativism" and they accept only that it "is no doubt sometimes plausible" (Abstract).

The reasons for C&T's hostility to the evolutionary solution are not made clear, but they appear to be uncomfortable with the idea of a large number of adaptively specialized type-2 problem solving devices. One of their goals is, therefore, to show how more general-purpose mechanisms and processes might be used in the solution of type-2 problems. Their preferred account of recoding functions builds them via a two-stage ontogenetic process which retains a central role for associative learning. C&T argue that the first stage consists of associative learning over type-1 problem domains which results in the achievement of specific representational states. In the second stage the achieved states can be used more generally on a trial and error basis as potential recoding functions for arbitrary problems which might thus be reduced from type-2 to type-1.

The trouble with C&T's ontogenetic scheme is that it does not solve the problem and hence there is no reason to prefer it to the phylogenetic account of the origins of recoding functions. Let us suppose that associative learning can modify the connectivity of a module or a subnet, as hypothesized by C&T, to realize a specific function which solves a type-1 problem. Let us further suppose that the module can then be made more widely accessible for use by other input domains, perhaps in one of the ways suggested by Rozin (1976). It is hard to see what advantage this confers. The difficulty with type-2 problems is that the space of potentially applicable recoding functions is infinite. All that the first stage of the ontogenetic process can achieve is to make one of these recoding functions available to the problem solver. But unless problem spaces have related structures, the second, trial and error, stage of the process would be of no value, because the probability is vanishingly small that the acquired recoding function would just happen to reduce an independent type-2 problem in a useful way. C&T appear to have no way to avoid this conclusion because they accept that the principle relating achieved representations to problem spaces is chance. "Each such recoding may just happen to reduce a problem that was previously type-2." (sect. 4, para. 6).

C&T's reluctance to accept a phylogenetic account of the origins of recoding functions is all the more curious in the light of their enthusiasm for trading computation for representation. Given that type-2 problems "permeate biological cognition right down to its roots" (sect. 5, para. 1) it is clearly the case that selective

pressure would be exerted in favour of mechanisms which instantiated more powerful representations and thus solved problems faster or more accurately than those that did not. It is a computational truism that special purpose machines are faster and more efficient than general purpose machines and it is also evident that natural selection preserves mechanisms which offer selective advantage with respect to specific problems. Evolution's gift of an appropriate set of type-2 problem-relevant recoding biases is exactly what we ought to expect.

Authors' Response

Relational learning re-examined

Chris Thornton^a and Andy Clark^b

^a*Cognitive and Computing Sciences, University of Sussex, Brighton, BN1 9QH, United Kingdom;* ^b*Philosophy/Neuroscience/Psychology Program, Washington University in St. Louis, St. Louis, MO 63130. chris.thornton@cogs.sussex.ac.uk; andy@twinearth.wustl.edu*

Abstract: We argue that existing learning algorithms are often poorly equipped to solve problems involving a certain type of important and widespread regularity that we call "type-2 regularity." The solution in these cases is to trade achieved representation against computational search. We investigate several ways in which such a trade-off may be pursued including simple incremental learning, modular connectionism, and the developmental hypothesis of "representational redescription."

The target article explores a familiar topic (the limits of simple statistical learning) in what we hope is a rigorous and challenging way. Its motivation was simply the observation that certain types of problem are both frequently solved (by biological learning devices) and yet appear highly intractable from a statistical point of view. These intractable (so-called "type-2") scenarios are ones in which the learner must identify *relations* among raw input elements rather than associations. The puzzle is: how is it possible for limited biological agents to negotiate such statistically impenetrable problem domains? The answer is (we claim) that short of being provided with antecedent search-space-shrinking knowledge (in which case the problem does not arise) the only hope lies in a "bag of tricks" approach that exploits general strategies for pressing maximal effect from those rare cases in which, by chance, a useful re-coding has been found. Re-coding is essential since it is a process that can take a relational property and turn it into a bona fide higher level element in a new space in which previously complex and elusive properties (such as relations between relations) appear as simple patterns (relations).

This thesis, we concede, can seem by turns trivial (of course higher order relational learning is tough!), wildly speculative (surely there are more direct ways of solving this kind of learning problem?), over-technical (did we really need statistics to make our point?), and under-technical (just how precise is the type-1/type-2 distinction anyway?). It is to the great credit of the commentators that, pretty much without exception, they responded constructively, repeatedly underlining our central theme and offering a wealth of useful suggestions and links to other bodies of work. Their responses bear mainly on six issues and we divide our Response accordingly.

R1. Is there still a Grand Ploy waiting to be discovered?

Our claim was that no general algorithm can exist for the systematic discovery of type-2 regularities in unrestricted domains. The most nature can do is to press maximal utility from whatever re-codings are found by chance or by simple search in less problematic domains, or to adjust the problem space itself so as to better exploit the existing biases of the human learning device (as in the Newport [1990] conjectures about morphology). Some commentators, however, proved unable to suppress a laudable optimism and felt (**Berkeley, Haberlandt**) that some more powerful and general mechanism might yet be available. Thus Berkeley, while appearing to be in agreement with much of our thesis, suggests that backpropagation networks using non-monotonic units can in fact deal with the type-2 parity-generalization scenario which we refer to in our paper. He cites a number of simulation results which back this up. We have no difficulty with this proposal and would only comment that the use of such “nonmonotonic” units equips the learning method in question with an implicit recoding ability and that this just happens to be appropriate for the problem domain he concentrates on, namely parity generalization. Thus Berkeley effectively demonstrates that a suitably biased type-2 method can solve a type-2 problem. Such a demonstration, however, is in no way suggestive of the re-coders grail: a fully general algorithm that achieves type-2 learning whatever the domain.

Several commentators (**Oberlander, Stufflebeam**, and, to some extent, **Kurtz**) suggested that the nearest thing that nature provides to such a general Grand Ploy may be the use (by a lucky few evolved creatures) of a variety of *external* representational systems such as language, maps, graphs, and other kinds of real world structure. This is a powerful and important suggestion, and one that we merely touched upon in our original treatment (see our comments on Dennett in sect. 3 and on the potential role of real world structures and action in sect. 4). We are, however, in full agreement with the idea that external representations play a very major role in empowering biological learning devices (Clark, 1989, Ch. 7; Clark 1997).

We found **Oberlander's** thoughtful and constructive commentary of special help in this regard. Oberlander develops a number of very compelling examples of ways in which we can simplify inner computations by adding structure to the local environment. This is a theme whose time has clearly come, for it is surfacing again and again in recent and influential work on so-called embodied and embedded cognition (see e.g., Hutchins 1995 and Clark 1997), and it is one that we intend to pursue in detail in our future work.

We cannot resist relating a further example, shown to us by Roger Thompson (personal communication), that seems perfectly to illustrate this theme. It concerns the ability of certain chimpanzees (pan troglodytes) to use experience with external tokens to enable them to perform higher order matching tasks that they would otherwise find impossible. The basic result, described at length in Thompson et al. (in press) is that when trained to associate a relational feature of some inputs (e.g., the feature of sameness) with an arbitrary external token (such as a plastic heart), the chimps can go on to learn to perform a higher order task (matching relations *between* relations) that would otherwise defeat them. Thus they become able to judge of two

pairs of objects – such as two identical shoes and two identical cups – that the pair of pairs is an instance of the sameness relation at a higher level, that is, sameness in respect of sameness, each pair being itself an instance of the basic relation of object level sameness. This task of matching relation between relations is, we think, a clear instance of a type-2 learning scenario. But one in which, as predicted by **Oberlander**, the existence of external tokens capable of reifying the relations between basic domain elements renders the problem tractable to on-board biological cognition. We here trade externally provided props and structures against expensive and perhaps even intractable episodes of inner computation.

In addition to the dedicated seekers after a Grand Ploy, some commentators suggested useful additional locally effective props and stratagems that might be added to our bag of tricks. **Szilas & Shultz** note the virtues of cascade correlation networks and suggest that a greater use of between network connections may do much to reduce the need for whole network copying and to overcome mismatches of input size during episodes of analogical reasoning. We agree that these and other technical tricks may help explain in detail how codings developed in one domain get to be transferred to others in which they may, at times, reduce type-2 complexity to type-1 tractability. The basic strategy however is still simply the re-use of achieved representation – it is trading spaces just as we envisaged it.

R2. The role of evolution

One contentious move in our original treatment was to avoid reliance on what we (perhaps unadvisedly) termed “heavy-duty nativism.” Many otherwise sympathetic commentators (**Bullinaria, Wells, Elton, Dartnall**) felt this to be a too hasty dismissal of a potentially rich source of re-coding functions. With this, however, we have no argument. Our move was rather a strategic one, designed to focus attention on the problematic (but surely inevitable?) residual range of cases in which evolution has not already done our re-coding work for us. We thus accept **Wells's** (see also **Marcus**) suggestion that “evolution's gift of an appropriate set of type-2 problem-relevant recoding biases is exactly what we ought to expect,” at least as far as various evolutionarily central learning functions are concerned. But if evolution is to be the *only* source of such re-codings, the lenses of human thought and science must be much weaker and narrower than we had supposed. It seems implausible, to us, to thus limit the space of humanly possible thought (though it surely *has* limits – just not ones directly set by an evolved set of recoding functions). Hence our desire was to explore any other strategies that might be available to us on an ontogenetic *or* cultural-evolutionary time scale.

An interesting suggestion, from **Bullinaria**, is that learning algorithms that build in a few simple and biologically plausible constraints may show improved performance on many problems that would otherwise involve intractable search. Such constraints include assumptions of symmetry between certain weights and the assumption that local information is more likely to matter than distal information. Such fixes and biases constitute, it seems to us, some very plausible ways in which a thrifty nature might subtly bias learning systems so as to promote the successful learning of specific skills in ecologically normal settings (see Karmiloff-

Smith 1992; Clark 1993, Chs. 4 and 5). But once again our primary target lies elsewhere in any residue of cases that must be dealt with by ontogenetic or cultural-evolutionary means.

Similar remarks apply to **Skokowski's** admonition to look more closely at the architectural inheritance of our biological brains. Of course, as **Dartnall** nicely points out, this is hardly an all-or-nothing matter. Our ontogenetic forays into type-2 space must be primed and rooted in early episodes of thought and learning that exploit quite evolutionarily basic mechanisms for apprehending and responding to our world. Like Dartnall, we envisage a cascade of information-processing activity in which evolution's gifts (**Wells, Skokowski**) and cultural and ontogenetic luck and hard labor come together. Getting straight about their relative contributions and complex interactions is, we think, one of the most pressing tasks facing contemporary cognitive science. The type-1/type-2 distinction is intended as a heuristic tool in the service of just such an endeavor.

If, however, we live in a world in which evolutionarily unanticipated type-2 learning scenarios are regularly encountered, the possibility arises (**Bullinaria, Dominey**) that we may evolve if not fully general, at least multi-purpose built-in strategies to support such learning. One plausible contender for such a built-in ploy is whatever on-board machinery supports the process of analogical reasoning. Analogical reason, as noted in our original treatment, provides an open-ended means of re-using achieved recodings so as to view a new problem space through a highly structured lens. Dominey's helpful and illuminating commentary presents a convincing and powerful demonstration of our basic thesis and clearly shows how analogical transfer can at times help to overcome the problem. The central example involved learning sequences not related by surface structure but only by abstract underlying relational structure. Bullinaria shows both the (anticipated) failure of type-1 learning and the success of an augmented model that applies filters (see also **Dartnall**) transferred from other learning experiences. We were especially impressed with Bullinaria's demonstration that the filters for re-use could be selected by a type 1 process of sequence recognition, as this goes some way toward addressing the very real worry (**Wells**) that it is unclear how to choose an achieved recoding for use in a new domain.

A very different set of issues comes to the fore in **Elton's** interesting and provocative comments concerning some important differences between evolutionary and ontogenetic learning scenarios. Elton claims that since problems and solutions can co-evolve it is misleading to think of the issue (over evolutionary time) as one of finding a kind of pre-determined target mapping. Instead, he says, it is a matter of finding a kind of behavioral profile that works and then sticking to it. We agree but we do not see that this re-statement in any way undermines our project. First, because our principal focus is, as we have said, on individual and perhaps cultural-evolutionary learning. Second, because we already anticipated the role of co-evolution in our original treatment (see our discussion of Newport [1990] in sect. 3). And third, because there is really nothing in our framework that commits us to the view that learning or evolution is anything like passing an exam. In fact our entire argument could be reformulated using Elton's own notion that "creatures stick with [behaviours] that work." Our central notion would become the idea that recoding was

only involved in the acquisition of certain "behaviors that work." The rest of our story would remain the same.

R3. Statistics and theories

In pursuing arguments predicated upon the limitations of simple kinds of statistically driven learning, we expose ourselves to the rapid rejoinder that there are simply more things on heaven and earth . . . specifically, what about theory-driven thought, explicit, conscious reflection and the like? Thus **Leiser** reminds us that advanced learning answers to requirements of coherence and coordination and suggests that we should attend more closely to the peculiar dynamics of theory-formation. **Memmi**, likewise, suggests that the presence of rich theoretical contexts and background knowledge deeply inform our advanced searches for recoding biases for relational learning. And **Vinter & Perruchet** (see also sect. R4 below) highlight the role of explicit, conscious reflection in going beyond simple statistical learning.

The general debate here, between theory-based and statistics-based conceptions, is addressed from a connectionist perspective in Clark 1993, Chapter 5. Our general feeling, however, is that the distinction, though clearly important, is easily overplayed. For we still need some account of the *origin* of the theoretical pictures that thus inform subsequent reasoning. And that origin, as far as we can see, can involve only some combination of innate biases and the fruits of an incremental cascade of statistically driven learning.

One crucial link between "mere statistics" and explicit human theorizing is powerfully displayed in **Kurtz's** very pertinent commentary. Like us, Kurtz hopes to account for complex theory-driven categorization without "quitting the tangible realm of experience and data." This involves, Kurtz suggests, going beyond mere perceptual similarity without losing the solid statistical foundations that are, we believe, the root of all learning. And this in turn involves the recognition and reification of additional *functional* regularities, that is, abstract features that unite disparate instances via the common actions they evoke, the common goals they relate to, and so on. In this way the idea of incrementally constructed, statistically-based feature spaces phases, rather naturally, into the idea of something akin to a theory based take on the various domains of human activity. Our goal is thus not to sever statistical and theory-driven learning but to understand various levels of theoretical understanding as themselves the fruits of an incremental sequence of episodes of type-1 learning, augmented by various tricks involving the re-use of achieved representational resources.

The latter stratagems will often include the uses of analogical reason, and the roles of culture, language, and conscious reflection as stressed by **Memmi, Leiser, Vinter & Perruchet**, and others. But what we should *not* do, we believe, is simply to invoke "theories and background knowledge" as the sufficient answer to the hard question, How is type-2 learning possible at all? For such an invocation is ultimately unexplanatory, trading a problematic chicken for an unexplained egg. Instead, we need to see how such theoretical knowledge can arise from real confrontation with environmental data and how it can be maximally exploited in future learning.

R4. Cognitive psychology

One of the most fruitful and exciting outcomes of the *BBS* process was, for us, the discovery of a quite unexpected wealth of links and connections between our (machine learning based) work and ongoing research in various areas of cognitive psychology, such as implicit learning (**Vinter & Perruchet, Dominey**), analogical reason (**Dominey**), categorization (**Kurtz**), and general research on re-coding (**Haberlandt**). Vinter & Perruchet, in particular, reveal very promising links between some of our claims and work on so-called implicit learning (learning without conscious or linguistic reflection). They agree that the ploys and stratagems we uncover help reveal how very complex problems can be dealt with by elementary, statistics-based processes and note that the course of such learning as predicted by our model is in good accord with experimental research (their own and others) on implicit learning in human subjects. They worry, however (and this ties in with the theory/statistics issues mentioned above) that we fail to address more explicit modes of thought and hence fail to do justice to the full spectrum of human learning.

Our reason for thus stopping short is – as noted above – that we really do believe that in a certain sense there really is no other kind of learning to be had. Such learning, at the very least, lies at the heart of all existing algorithms capable of learning about a domain and not equipped with heavy, task-specific initial biases. We do not deny, of course, that acquired knowledge can be used precisely so as to induce biases that will in effect take the system beyond the domain of visible statistical features of the inputs. Indeed, this is exactly our point: that the only way to go “beyond” the statistics is to use knowledge, itself acquired through earlier instances of type-1 learning (or else innate) so as to re-shape the space for future learning.

Vinter & Perruchet seem to suggest, in addition, that there may be special features of conscious thought and reflection that enable us to do more than simply re-shape the space for learning. Such features would include, for example, the use of “processes which rely on the specific power of conscious thought.” Once again, however, we fear a chicken and egg scenario. Our goal is to understand how biological agents can come to wield the knowledge to which such powers may be applied. We do concede, however, that certain aspects of very high level thought look to lie beyond the scope of our treatment. Thus **Vinter & Perruchet** (also **Dartnall**) mention the human ability not just to know a recoding function but to know that we know it. Such knowledge is not of the world so much as of the ways in which we know the world. This certainly does seem like an important ability though the extent to which it figures in daily problem solving is perhaps open to doubt. Whether such top level, meta-reflective capacities merely represent the culmination of a cascade of processes of type-1 learning and re-deployment of achieved representation (as we suspect) or rely on the operation of some wholly different faculty (perhaps tied up with conscious thought) is an important topic for further research. It is, of course, very likely that different neurological structures play a role in type-1 learning and type-2 re-coding (see, e.g., **Dominey**’s comments on the role of the frontostriatal system in type-1 learning). But this is consistent with our claim that the combination of these strategies is effectively all that nature can provide.

Halford’s useful and interesting commentary suggests that the type-1/type-2 distinction can help to clarify several issues in the development of children’s understanding of mathematics. Halford then makes the important point that not all re-codings are reversible, that is, that it may not be possible to re-create the original data from the recoded data. To avoid this loss of potentially useful information, he suggests a technique that involves representing the input as multiple distinct dimensions that are then processed together. The basic idea sounds interesting, but we found the details of the suggestion elusive. One worry is that, in Halford’s own example, the probabilities are based on complete entries in the target mapping. But a complete listing would here constitute a direct recapitulation of the training set – a fact which seems to reduce the technique to the use of a look-up table. In any case, it seems to us that, in this case, one simply cannot have one’s code and eat it! In a sense the information-losing properties of the recoding process are crucial since they power the simplification and data compression that in turn lead to the properties of improved search and generalization that the whole process is designed to support. The only real hope in this area, it seems to us, lies in the use of what Karmiloff-Smith (see her 1992, pp. 21–24) once termed conservative redescription – a process in which re-codings are generated but the original representations remain intact and available for use in certain contexts.

In general, then, we were especially pleased to discover these rich links between our themes and treatment and ongoing work in cognitive psychology. One commentator, however (**Ohlsson**), felt that our treatment amounted to a reversion to a discredited behaviorist vision of psychology – one which concerned itself not with the understanding of inner mechanisms but only with patterns of stimulus and response. Here (as with **Elton**) it seems we may have misled by our use of the vocabulary of target mappings, input-output mappings, and so on. But of course we do not wish to claim that no important and contentful inner states mediate between inputs and outputs. Indeed, our whole argument is devoted to displaying the sheer complexity and importance of the search for fruitful inner transformations to be applied to raw input patterns. When Ohlsson says that “the idea that what is learned [by human learners] is an input/output mapping (or a set of stimulus-response connections) was abandoned in the 1950s because people began taking the generativity of human cognition seriously,” we are in complete agreement! We are perplexed that Ohlsson sees our paper as in any way disputing this. Our central aim was, in fact, to argue that interesting forms of learning involved not the acquisition of stimulus-response conditions but rather the construction of complex recoding structures (possibly under incremental learning regimes), which would then provide the basis for generative knowledge enabling the learner to go beyond any presented data.

R5. Internal representation

All our talk of inner re-codings raised the hackles of some commentators who seem a little leery of the very idea of internal representation (**Stufflebeam**) or who wanted at least to suggest some possible alternative mechanisms (both internal and external) for achieving the same kinds of result (**Dartnall, Oberlander**). We have already endorsed

the suggestion (**Oberlander, Stufflebeam**) that external structures may sometimes contribute mightily to successful type-2 learning. Stufflebeam seems, in addition, to want to cast doubt on the idea that inner states properly thought of as representation have any role to play in the process at all. We demur, but this is a large and lively debate that we cannot hope to do justice to here – see Clark (1997) for a defense of a kind of modest representationalism consistent with our claims. We would, however, comment that to whatever extent a system creates inner states that effectively reify relational features of the inputs that carry adaptively significant information (as where the chimpanzees learn to match higher order sameness), it is hard to see why we should refuse to label such states internal representations.

Dartnall draws a useful distinction between this (weak) notion of internal representation and one more closely tied to the idea of conscious thought and reflection. Dartnall usefully locates our discussion as part of a larger research program whose goal is to understand the transition between connectionist competence and structured thought. In this context, he suggests that the early stages of our “re-coding” cascade may be more fruitfully conceived in terms of a sequence of increasingly powerful ways of accessing the knowledge rather than in terms of re-codings of the knowledge itself. This is an interesting idea and one that seems to invite a slightly different perspective on Karmiloff-Smith’s notion of representational re-description – a notion that provided much of the motivation and inspiration for the present treatment [see multiple book review of Karmiloff-Smith’s *Beyond Modularity* BBS 17(4) 1994]. We are not yet convinced, however, that this distinction (between changing access and changing knowledge) is deep and well-defined. But it certainly suggests some directions for future research, perhaps using some concrete computational models to refine our currently hazy intuitions.

R6. Technicalia

A number of commentators made useful technical suggestions concerning the description of the type-1/type-2 distinction, and raised important questions about the relations between the parity considerations and generalization and the example of the Elman net. Thus **Damper** (see also **Gaskell**) worries that holding back even a single pattern on the classical (2 variable, XOR) parity problem simply makes the problem insoluble (the machine would need to read our minds to know the intended function) as the learning algorithm lacks sufficient data. He concludes that it must be wrong to link parity learning to issues about generalization. But let us step back a little here and review our strategy in more detail. In the paper we suggested that backpropagation learning does not provide a ready-made solution to the problem of type-2 scenarios and backed this up with a demonstration that backpropagation reliably fails on some forms of parity-generalization problem. The slight novelty in this was indeed the utilization of parity as a generalization problem. Where parity problems have been used in machine learning, they have typically been presented as memory tasks, that is, learning methods have been required to acquire complete mappings. One of the justifications put forward for this approach is the idea that parity constitutes an “unfair” generalization problem. Damper’s commentary

is valuable because it shows how muddled the thinking behind this judgment can be.

Damper implies that parity cannot be a generalization problem because parity mappings exhibit neutral statistics, that is, chance-level output probabilities. This observation was a fundamental component in our own presentation. But it demonstrates not that parity problems are ungeneralizable but merely that they cannot be generalized on the basis of statistical effects.

In coming to terms with the idea of parity generalization, it is useful to turn attention away from the familiar XOR case towards higher-order cases. In 4-bit parity there are 16 cases. The removal of a single case leaves 15 cases as a basis for generalization. Somehow this does not seem quite so unreasonable. It may also be helpful to consider the two-spirals problem in which the learning algorithm must learn to correctly assign a 2-d point to one of two interlocking spirals in an image. The problem is parity-like since nearest-neighbors in the input space always have opposite classifications. And indeed, the statistics of a typical training set are usually nearly neutral. And yet, as far as we are aware, this problem has never been treated as anything other than a generalization problem.

Despite **Damper**’s objections to our use of parity generalization problems, he has no difficulty with our central thesis that learning problems which require recoding present a special and cognitively important case.

Turning to the type-1/type-2 distinction itself, **Chater** argues that the distinction is simply ill-defined and hence will confuse rather than clarify matters. Notice, however, that we took some care, in section 1 of the paper, to stress that “there is no obvious operational definition for the class of type-2 problems.” However, by proceeding on the basis that such a definition exists, and indeed that our paper was supposed to provide it, **Chater** has arrived at a number of interesting though strictly speaking irrelevant observations. His approach involves taking our analysis of the ways in which supervisory feedback can provide justifications for assignments of particular probabilities to particular outputs as a formal definition of a problem class. He shows that this soon leads to nonsensical results and enables dubious maneuvers such as the adding of “dummy variables” so as to change the “formal” characterisation of a particular problem.

These arguments may be of interest to the computational learning theorist. However, they completely miss the point of our paper and in some cases actually mislead. For example, **Chater** contrives to show that if we try to give identity-function learning a classification using his “formalization” of our framework, certain ambiguities result. However, this conceals that fact that identity-function learning actually has a rather natural characterization as a type-2 operation within our framework.

Assume that the input for the learner is based on two variables – one representing the input to the identity function and the other representing the output – and that the target output for the learner is a value which shows whether the input forms a valid application of the identity function (i.e., whether the two input values are the same). In this scenario the learner’s guessing of a particular output cannot be justified on the basis of observed frequencies in the training data since every input value is unique. However, were we to recode the learner inputs by applying an identity recognition function to them, we would produce a

recoding of the problem which could be solved in exactly that “statistical” way. Thus identity function learning is naturally characterized in our framework as requiring recoding and hence is type-2.

Golden, in his gracious and intriguing commentary, offers an amendment to our account of type-1 learning. He suggests that the category can be broken down into sub-cases using a “parametric, model-based approach” and that this may help avoid some potential problems. Alas, we are not sufficiently familiar with the background material for this proposal to properly judge its value or necessity. We agree, however, that there may well exist other (perhaps more elegant) ways of picking apart the kinds of cases we wish to distinguish, and we look forward to Golden (forthcoming) to learn more about the methods he has in mind.

Finally, some questions were raised concerning the role of the Elman network in our discussion. Thus **Marcus** argues that the Elman network (Elman 1993) fails to illustrate our central point as it does not, after all, learn about higher order regularities in the data set. The evidence that Marcus relies on, however, concerns only the failure of such a net to generalize an abstract structure (an X is an X) when presented with a case involving a totally novel “filler” (a dax is an . . .). It may be that such extensions simply require more than the kind of grammatical knowledge that the data set makes available. In any case, it does not follow from this kind of failure that the original network does not acquire grammatical knowledge that is higher order in the sense we require. For the successful network did indeed proceed by first identifying lower level grammatical features and then going on to learn about regularities involving relations between these lower level features. In fact, this is exactly what the incremental batching/staged memory manipulations were designed to encourage. **Gaskell** seems to worry that the need for such manipulations renders the network unsuitable for our purposes. We do not see why: our point here is simply that the desired performance is achieved only by the prior isolation (by whatever means) of a “building block” set of regularities which then mediate between the raw data and the target mapping so as to shrink the search space to a manageable size. The network thus trades prior learning against subsequent search.

Taken together, the various commentaries have done much to advance (and where necessary, to unsettle) our thinking about the nature of learning and the canny ways in which biological cognizers may trade achieved representation against potentially infinite computational search. Our treatment, we readily concede, can easily appear either trite or wildly speculative. Few disagree with the central tenet (re-coding matters!). Few will concede our central claim (that the maximal exploitation of the fruits of simple learning or chance penetrations into type-2 space is the best nature can provide). We hope, however, to have fueled the fires of debate. And we thank all the commentators for their thoughtful and genuinely constructive responses. We have learnt a lot, and we look forward to trading it against our future search in computational space.

References

Letters “a” and “r” appearing before authors’ initials refer to target article and response, respectively.

Agre, P. E. (1993) The symbolic worldview: Reply to Vera and Simon. *Cognitive Science* 17(1):61–69. [RSS]

- Agre, P. E. (1995) Computation and embodied agency. *Informatica* 19(4):527–35. [RSS]
- Anderson, J. A. (1995) *An introduction to neural networks*. MIT Press. [KH]
- Anderson, J. R. (1983) *The architecture of cognition*. Harvard University Press. [KH]
- Armstrong, S. L., Gleitman, L. R. & Gleitman, H. (1983) What some concepts might not be. *Cognition* 13:263–308. [KK]
- Ausubel, D. (1963) *The psychology of meaningful verbal learning*. Grune & Stratton. [SO]
- Barkow, J. H., Cosmides, L. & Tooby, J. (1992) *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford University Press. [AW]
- Baum, E. B. & Haussler, D. (1989) What size net gives valid generalization? *Neural Computation* 1:151–60. [JAB]
- Baum, E. B. & Wilczek, F. (1988) Supervised learning of probability distributions by neural networks. In: *Neural information processing systems*, ed. D. Z. Anderson. American Institute of Physics. [RID]
- Baxter, J. (1995) Learning internal representations. In: *Proceedings of the Eighteenth International Conference on Computational Learning Theory*. ACM Press. [NS]
- Becker, S. & Cun, Y. (1988) *Improving the convergence of back-propagation learning with second-order methods* [CRG-TR-88-5]. University of Toronto Connectionist Research Group. [aAC]
- Berkeley, I., Dawson, M., Medler, D., Schopflocher, D. & Hornsby, L. (1995) Density plots of hidden unit activations reveal interpretable bands. *Connection Science* 7(2):167–86. [ISNB]
- Brooks, R. (1991) Intelligence without representation. *Artificial Intelligence* 47:139–59. [RSS]
- Bullinaria, J. A. (1994) Modeling reading, spelling and past tense learning with artificial neural networks. *Brain and Language* (in press). [JAB]
- Campbell, D. (1974) Evolutionary epistemology. In: *The philosophy of Karl Popper*, ed. P. Schillp. Open Court. [aAC]
- Carey, S. & Spelke, E. (1994) Domain-specific knowledge and conceptual change. In: *Mapping the mind: Domain specificity in cognition and culture*, ed. L. A. Hirschfeld & S. A. Gelman. Cambridge University Press. [DL]
- Caruana, R. A. (1994) Multitask connectionist learning. In: *Proceedings of the 1993 Connectionist Summer School*. Erlbaum. [NS]
- Chi, M. T. H. (1992) Conceptual change within and across ontological categories. In: *Cognitive models of science*, ed. R. N. Giere. University of Minnesota Press. [DL]
- Chomsky, N. (1959) Review of Skinner’s “Verbal behavior.” *Language* 35:26–58. [SO]
- Chomsky, N. & Miller, G. (1963) Introduction to the formal analysis of natural language. In: *Handbook of mathematical psychology*, vol. 2, ed. R. D. Luce, R. R. Bush & E. Galanter. Academic Press. [aAC]
- Churchland, P. M. (1995) *The engine of reason, the seat of the soul*. MIT Press. [aAC]
- Churchland, P. & Sejnowski, T. (1992) *The computational brain*. MIT/Bradford. [aAC]
- Clark, A. (1989) *Microcognition*. MIT Press. [rAC]
- (1993) *Associative engines: Connectionism, concepts and representational change*. MIT/Bradford. [aAC]
- (1997) *Being there: Putting brain, body and the world together again*. MIT Press. [RSS, rAC]
- Clark, A. & Karmiloff-Smith, A. (1993) The cognizer’s innards: A psychological and philosophical perspective on the development of thought. *Mind and Language* 8:487–519. [aAC]
- Clark, A. & Toribio, J. (1994) Doing without representing? *Synthese* 101:401–31. [RSS]
- Cloete, I. & Ludik, J. (1993) Increased complexity training. In: *Proceedings: New trends in neural computation*, ed. J. Mira, J. Cabestony, & A. Prieto. Lecture notes in computer science, no. 686. Springer-Verlag. [NS]
- Dawson, M. & Schopflocher, D. (1992) Modifying the generalized Delta Rule to train networks of nonmonotonic processors for pattern classification. *Connection Science* 4(1):19–31. [ISNB]
- Denett, D. (1991) *Consciousness explained*. Little, Brown. [aAC]
- (1994) Labeling and Learning: Commentary on Clark and Karmiloff-Smith. *Mind and language* 8:544–548. [aAC]
- Dienes, Z. & Berry, D. (in press) How implicit is implicit learning? *Psychonomic Bulletin and Review*. [AV]
- Dietterich, T., London, B., Clarkson, K. & Dromey, G. (1982) Learning and inductive inference. In: *The handbook of artificial intelligence*, vol. 3, ed. P. Cohen & E. Feigenbaum. Morgan Kaufmann. [aAC]
- Dominey, P. F. (1995) Complex sensory-motor sequence learning based on recurrent state-representation and reinforcement learning. *Biological Cybernetics* 73:265–74. [PFD]
- Dominey, P. F. (in press) An anatomically structured sensory-motor sequence

- learning system displays some general linguistic capacities. *Brain and Language*. [PFD]
- Dominey, P. F., Arbib, M. A. & Joseph, J. P. (1995) A model of cortico-striatal plasticity for learning oculomotor associations and sequences. *Journal of Cognition and Neuroscience* 7(3):311–36. [PFD]
- Dominey, P. F., Ventre-Dominey, J., Broussolle, E. & Jeannerod, M. (1995) Analogical transfer in sequence learning: Human and neural-network models of fronto-striatal function. *Annals of the New York Academy of Sciences* 769:369–73. [PFD]
- Dominey, P. F., Ventre-Dominey, J., Broussolle, E. & Jeannerod, M. (in press) Analogical transfer is effective in a serial reaction time task in Parkinson's disease: Evidence for a dissociable sequence learning mechanism. *Neuropsychologia*. [PFD]
- Dreyfus, H. L. (1992) *What computers still can't do: A critique of artificial reason*, rev. ed. MIT Press. [RSS]
- Elman, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition* 48:71–99. [arAC, JAB, RID, MGG, GFM, DM, KH, NS, PS, TD]
- Eysenck, M. W. & Keane, M. T. (1995) *Cognitive psychology: A student's handbook*. Erlbaum. [SO]
- Fahlman, S. & Lebiere, C. (1990) *The cascade-correlation learning architecture* [CMU-CS-90-100]. Carnegie-Mellon University, School of Computer Science, Pittsburgh, PA 15213. [aAC]
- Fahlman, S. E. & Lebiere, C. (1990) The cascade correlation learning architecture. In: *Advances in neural information processing systems 2*, ed. D. Touretzky. Morgan Kaufman. [NS]
- Fodor, J. A. (1975) *The language of thought*. Harvard University Press. [RSS]
- (1995) The folly of simulation. In: *Speaking minds: Interviews with twenty eminent cognitive scientists*, ed. P. Baumgartner & S. Payr. Princeton University Press. [RSS]
- Funt, B. V. (1980) Problem solving with diagrammatic representations. *Artificial Intelligence* 13:210–30. [JO]
- Gagne, R. M. (1962) The acquisition of knowledge. *Psychological Review* 69:355–65. [SO]
- Gallant, S. I. (1994) *Neural network learning and expert systems*. MIT Press. [KH]
- Gallistel, R. (1994) Interview. *Journal of Cognitive Neuroscience* 6(2):174–79. [aAC]
- Gentner, D., Rattermann, M. J. & Forbus, K. D. (1993) The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology* 25:524–75. [DL]
- Gick, M. & Holyoak, K. J. (1980) Analogical problem solving. *Cognitive Psychology* 12:306–55. [NS]
- (1983) Schema induction and analogical transfer. *Cognitive Psychology* 15:1–38. [NS]
- Gold, R. (1987) *The description of cognitive development: Three Piagetian themes*. Oxford University Press. [DL]
- Goldberg, D. (1989) *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley. [aAC]
- Golden, R. M. (1988) A unified framework for connectionist systems. *Biological Cybernetics* 59:109–20. [RMG]
- (forthcoming) *Mathematical methods for neural network analysis and design*. MIT Press. [RMG, rAC]
- Goldstone, R. L. (1994) The role of similarity in categorization: Providing a groundwork. *Cognition* 52:125–57. [KK]
- Gomez, R. L. & Schaveneveldt, R. W. (1994) What is learned from artificial grammars? Transfer tests of simple association. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20(2):396–410. [PFD]
- Goswami, U. (1991) Analogical reasoning: What develops? A review of research and theory. *Child Development* 62:1–22. [DL]
- Haberlandt, K. F. (1990) Expose hidden assumptions in network theory. *Behavioral and Brain Sciences* 13:495–96. [KH]
- Haberlandt, K. F., Graesser, A. C., Schneider, N. J. & Kiely, J. (1986) Effects of task and new arguments on word reading times. *Journal of Memory and Language* 25:314–22. [KH]
- Halford, G. S. (1993) *Children's understanding: The development of mental models*. Erlbaum. [GSH]
- Halford, G. S., Wilson, W. H. & Phillips, S. (submitted) Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. [GSH]
- Hare, M., Elman, J. & Daugherty, K. (1995) Default generalization in connectionist networks. *Language and Cognitive Processes* 10:601–30. [GFM]
- Hertz, J., Krogh, A. & Palmer, R. G. (1991) *Introduction to the theory of neural computation*. Addison-Wesley. [RID]
- Hinton, G. & Sejnowski, T. (1986) Learning and relearning in Boltzmann machines. In: *Parallel distributed processing: Exploration of the microstructures of cognition*, vols. 1 & 2, ed. D. Rumelhart, J. McClelland & the PDP Research Group. MIT Press. [aAC, NC]
- Hutchins, E. (1995) *Cognition in the wild*. MIT Press. [rAC]
- Jackson, G. M., Jackson, S. R., Harrison, J., Henderson, L. & Kennard, C. (1995) Serial reaction time learning and Parkinson's disease: Evidence for a procedural learning deficit. *Neuropsychologia* 33(5):577–93. [PFD]
- Jacobs, R., Jordan, M. & Barto, A. (1991a) Task decomposition through competition in a modular connectionist architecture: The what and where visual tasks. *Cognitive Science* 15:219–50. [aAC]
- Jacobs, R., Jordan, M., Nowlan, S. & Hinton, G. (1991b) Adaptive mixtures of local experts. *Neural Computation* 3:79–87. [aAC]
- Jarvella, R. J. (1979) Immediate memory and discourse processing. In: *The psychology of learning and motivation: Advances in research and theory*, ed. G. H. Bower. Academic Press. [KH]
- Karmiloff-Smith, A. (1979) *A functional approach to child language*. Cambridge University Press. [aAC]
- (1992a) Nature, nurture and PDP: Preposterous development postulates? *Connection Science* [special issue on philosophical issues in connectionist modelling] 4(3/4):253–70. [arAC, TD, NS]
- (1992) *Beyond modularity: A developmental perspective on cognitive science*. MIT Press. [DM]
- Karmiloff-Smith, A. & Clark, A. (1993) What's special about the development of the human mind/brain? *Mind and Language* 8(4):569–81. [aAC]
- Keil, F. C. (1994) The birth and nurturance of concepts by domains. In: *Mapping the mind: Domain specificity in cognition and culture*, ed. L. A. Hirschfeld & S. A. Gelman. Cambridge University Press. [DL]
- Kintsch, W. (1988) The use of knowledge in discourse processing: A construction-integration model. *Psychological Review* 95:163–82. [KH]
- Kirsh, D. (1991) When is information explicitly represented? In: *Information, thought and content*, ed. P. Hanson. UBC Press. [aAC]
- Kirsh, D. (1995) The intelligent use of space. *Artificial Intelligence* 72:1–52. [aAC]
- Kirsh, D. & Maglio, P. (1994) On distinguishing epistemic from pragmatic action. *Cognitive Science* 18:513–19. [aAC]
- Krogh, A. & Hertz, J. A. (1992) A simple weight decay can improve generalization. In: *Advances in neural information processing systems*, ed. J. E. Moody, S. J. Hanson & R. P. Lippman. Morgan Kaufman. [JAB]
- Kurtz, K. J. (1996) *Category-based similarity*. Poster presented at the Eighteenth Annual Conference of the Cognitive Science Society, San Diego, CA. [KK]
- Larkin, J. H. & Simon, H. A. (1987) Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science* 11:65–100. [JO]
- LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W. & Jackal, L. (1989) Back propagation applied to handwritten zipcode recognition. *Neural Computation* 1:541–51. [aAC]
- Leiser, D. (1996) Computational emergence and developmental emergence: A sobering survey. *Intellectica* 22:203–220. [DM]
- Lewicki, P., Hill, T. & Bizot, E. (1988) Acquisition of procedural knowledge about a pattern of stimuli that can not be articulated. *Cognitive Psychology* 20:24–37. [AV]
- Lindsay, R. K. (1988) Images and inference. *Cognition* 29:229–50. [JO]
- MacKay, D. J. C. (1992a) Bayesian interpolation. *Neural Computation* 4:415–47. [JAB]
- (1992b) A practical Bayesian framework for backpropagation networks. *Neural Computation* 4:448–72. [NC]
- Marcus, G. F. (submitted a) Rethinking eliminative connectionism. [GFM]
- (submitted b) What does it take to get a connectionist model to generalize a low frequency default? [GFM]
- Marcus, G. F., Brinkmann, U., Clahsen, H., Wiese, R. & Pinker, S. (1995) German inflection: The exception that proves the rule. *Cognitive Psychology* 29:186–256. [GFM]
- Marr, D. (1982) *Vision*. Freeman. [aAC]
- McClelland, J. L. & Rumelhart, D. E. (1988) *Exploration in parallel distributed processing: A handbook of models, programs, and exercises*. MIT Press/Bradford. [RID]
- Miller, G. A. (1956) The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 63:81–93. [SO, KH]
- Minsky, M. & Papert, S. (1988) *Perceptrons: An introduction to computational geometry*, 3rd ed., expanded. MIT Press. [aAC, ISBNB]
- Murphy, G. L. & Medin, D. L. (1985) The role of theories in conceptual coherence. *Psychological Review* 92:289–316. [KK]
- Newport, E. (1990) Maturation constraints on language learning. *Cognitive Science* 14:11–28. [arAC, TD]
- Nissen, M. J. & Bullemer, P. (1987) Attentional requirements of learning: Evidence from performance measures. *Cognition Psychology* 19:1–32. [PFD]
- Nowlan, S. N. & Hinton, G. E. (1992) Simplifying neural networks by soft weight sharing. *Neural Computation* 4:473–93. [JAB]

- Ohlsson, S. (1996) Learning from performance errors. *Psychological Review* 2:241–62. [SO]
- Ohlsson, S. & Rees, E. (1991) The function of conceptual understanding in the learning of arithmetic procedures. *Cognition and Instruction* 8:103–79. [SO]
- Perruchet, P. & Gallego, J. (in press) A subjective unit formation account of implicit learning. In: *How implicit is implicit learning?* ed. D. Berry. Oxford University Press. [AV]
- Perruchet, P., Gallego, J. & Savy, I. (1990) A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. *Cognitive Psychology* 22:493–516. [AV]
- Piaget, J. (1971) *Biology and knowledge*. Chicago University Press. [DM]
- (1974) *Adaptation vitale et psychologie de l'intelligence*. Hermann. [DL]
- Port, R. & Van Gelder, T. (1995) *Mind as motion: Explorations in the dynamics of cognition*. MIT Press. [RSS]
- Quinlan, J. (1986) Induction of decision trees. *Machine Learning* 1:81–106. [aAC]
- (1993) *C4.5: Programs for machine learning*. Morgan Kaufmann. [aAC]
- Reber, A. S. (1993) *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. Oxford University Press. [KH]
- Reeves, L. M. & Weisbert, R. W. (1994) The role of content and abstract information in analogical transfer. *Psychological Bulletin* 115:381–400. [DL]
- Rendell, L. (1989) A study of empirical learning for an involved problem. In: *Proceedings of the Eleventh Joint Conference on Artificial Intelligence*. Morgan Kaufmann. [aAC]
- Rescorla, R. A. & Wagner, A. R. (1972) A theory of Pavlovian conditioning: The effectiveness of reinforcement and nonreinforcement. In: *Classical conditioning 2: Current research and theory*, ed. A. H. Black & W. F. Prokasy. Appleton-Century-Crofts. 64–69. [RID]
- Richard, M. E. & Lippman, R. P. (1991) Neural network classifiers estimate Bayesian a posteriori probabilities. *Neural Computation* 3:461–83. [NC]
- Rozin, P. (1976) The evolution of intelligence and access to the cognitive unconscious. In: *Progress in psychobiology and physiological psychology*, ed. J. Sprague & A. Epstein. Academic Press. 245–80. [AW]
- Rudy, J. W. (1991) Elemental and configural associations, the hippocampus and development. *Developmental Psychobiology* 24(4):221–36. [GSH]
- Rumelhart, D. E., Durbin, R., Golden, R. & Chauvin, Y. (1995) Backpropagation: The basic theory. In: *Backpropagation*, ed. D. E. Rumelhart & Y. Chauvin. Cambridge, MA: MIT Press. Bradford Books. [KK]
- Rumelhart, D. E., Hinton, G. & Williams, R. (1986) Learning representations by backpropagating errors. *Nature* 323:533–36. [aAC, ISNB, KH, KK]
- Schultz, T. R., Schmidt, W. C., Buckingham, D. & Mareshal, D. (1995) Modeling cognitive development with a generative connectionist algorithm. In: *Developing cognitive competence: New approaches to process modeling*, ed. T. Simon & G. Halford. Erlbaum. [DL]
- Searle, J. R. (1990) Is the brain a digital computer? *APA Proceedings* 64(3):21–37. [RSS]
- Sejnowski, T. & Rosenberg, C. (1987) Parallel networks that learn to pronounce English text. *Complex Systems* 1:145–68. [aAC]
- Sejnowski, T. J. & Rosenberg, C. R. (1986) *NETalk: A parallel network that learns to read aloud (Report No. JHU/EECS - 86/01)*. Johns Hopkins University. [KH]
- Stenning, K. & Oberlander, J. (1995) A cognitive theory of graphical and linguistic reasoning: Logic and implementation. *Cognitive Science* 19:97–140. [JO]
- Stufflebeam, R. S. (1995) Representations, explanation, and PDP: Is representation talk really necessary? *Informatica* 19(4):599–613. [RSS]
- (1997) *Whither internal representations? In defense of antirepresentationalism and other heresies*. Doctoral dissertation, Washington University, St. Louis, MO. [RSS]
- Sun, R. & Bookman, L. (1994) *Computational architectures integrating neural and symbolic processes*. Kluwer. [DM]
- Sutton, R. S. & Barto, A. G. (1981) Towards a modern theory of adaptive networks: Expectation and prediction. *Psychological Review* 88:135–70. [RID]
- Szilas, N. & Ronco, E. (1995) Action for learning in nonsymbolic systems. In: *Proceedings of the European Conference on Cognitive Science*. INRIA. [NS]
- Tetewsky, S., Shultz, T. & Buckingham, D. (1993) Reducing retroactive interference in connectionist models of recognition memory. Canadian Society for Brain, Behaviour, and Cognitive Science. Third Annual Meeting. [NS]
- Tetewsky, S. J., Shultz, T. R. & Takane, Y. (1995) Training regimens and function compatibility: Implications for understanding the effects of knowledge on concept learning. In: *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Erlbaum. 304–309. [NS]
- Thagard, P. (1989) Explanatory coherence. *Behavioral and Brain Sciences* 12:435–502. [DL]
- Thagard, P., Holyoak, K. J., Nelson, G. & Gochfeld, D. (1990) Analogical retrieval by constraint satisfaction. *Artificial Intelligence* 46:259–310. [PFD]
- Thompson, R., Oden, D. & Boyson, S. (in press) Language-naïve chimpanzees (pan troglodytes) judge relations between relations in a conceptual matching-to-sample task. *Journal of Experimental Psychology: Animal Behavior Processes*. [rAC]
- Utgoff, P. (1986) *Machine learning of inductive bias, vol. 15 (Kluwer International Series in Engineering and Computer Science)*. Kluwer. [aAC]
- Valiant, L. G. (1984) A theory of the learnable. *Communications of the ACM* 27:1134–42. [NC]
- Van Gelder, T. (1995) What might cognition be if not computation? *Journal of Philosophy* 92(7):345–81. [RSS]
- Vinter, A. & Perruchet, P. (1994) Is there an implicit level of representation? *Behavioral and Brain Sciences* 17:730–31. [AV]
- Vosniadou, S. & Brewer, W. F. (1987) Theories of knowledge restructuring in development. *Review of Educational Research* 57:51–67. [DL]
- White, H. (1989) Some asymptotic results for learning in single hidden-layer feedforward network models. *Journal of the American Statistical Association* 84:1003–13. [RMG]
- Wilson, A. & Hendler, J. (1993) Linking symbolic and subsymbolic computing. *Connection Science* 5:3–4. [DM]
- Wisniewski, E. J. & Medin, D. L. (1994) On the interaction of theory and data in concept learning. *Cognitive Science* 18:221–81. [KK]