

Relational Learning

A. Clark and C. Thornton

June 3, 2003

‘Trading Spaces’ explores a familiar topic (the limits of simple statistical learning) in what we hope is a rigorous and challenging way. The motivation for the paper was simply the observation that certain types of problem are both frequently solved (by biological learning devices) 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 6 issues and we divide our reply accordingly.

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

Our claim, recall, 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 conjectures about morphology). Some commentators, however, proved unable to surpress 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 “non-monotonic” 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 kind’s 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 section 3 and on the potential role of real world structures and action in section 4). We are, however, in full agreement with the idea that external representations play a very major role in empowering biological learning devices (1ch 7, Clark (in press)).

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. Also Clark, in press), and it is one that we intend to pursue in detail in our future work.

We cannot resist here relating a further example, shown to us by Roger Thompson, 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, Oden and Boyson (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 i.e. 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 re-ifying the relations between basic domain elements renders the problem tractable to on-board biological cognition. We here trade external 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 and Shulz 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.

2. 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’ (see also Marcus) suggestion that “evolution’s gift of an appropriate set of type 2 problem 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, due to 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 2, 3chapt. 4,5). 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 comments 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 re-codings 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 towards addressing the very real worry (Wells) that it is unclear how to choose an achieved re-coding 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 don't 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 comments on Newport in section 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 the [behaviours] that work." Our central notion would become the idea that recoding was only involved in the acquisition of certain 'behaviours that work'. The rest of our story would remain the same.

3. 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

co-ordination 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 and Perruchet (see also section 4 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 3 ch.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’ very pertinent commentary. Like us, Kurtz hopes to account for complex theory driven categorization without “quitting the tangible realm of data and experience”. 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 re-ification of additional *functional* regularities i.e. 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.

These 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 and Perruchet and others. But what we should *not* do, we believe, is to simply 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.

4. 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 and Perruchet, Dominey), analogical reason (Dominey), categorization (Kurtz) and general research on recoding (Haberlandt). Vinter and 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 suggest that the ploys and

stratagems we uncover provide “a neat illustration of how very complex problems can be dealt with by elementary processes drawing information from the basic statistical regularities of the input”. They 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 and 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 “logical reasoning and inference 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 and 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, pleasingly, 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 i.e. 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! For 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 e.g. 4pp 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 etc. 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 conditions) was abandoned in the 50s 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 it. 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 training data.

5. Internal Representation.

All our talk of inner re-codings raised the hackles of some folk 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 as 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 (in press) 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 re-ify 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 “something like connectionist competence (and) something like 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. 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.

6. 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-generalisation problem. The slight novelty in this was indeed the utilisation of parity as a *generalisation* problem. Where parity problems have been used in machine learning, they have typically been presented as memory tasks, i.e., 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’ generalisation problem. Damper’s commentary is valuable because it shows how muddled the thinking behind this judgement can be.

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

In coming to terms with the idea of parity generalisation, 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 generalisation. 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-neighbours 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 generalisation problem.

Interestingly, despite Damper’s objections to our usage of parity generalisation problems, he has no difficulty with the central thesis of the paper, 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 provides 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 manouvres 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 ‘formalisation’ of our framework, certain ambiguities result. However, this conceals the fact that identity-function learning actually has a rather natural characterisation 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 characterised 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 and Hall argue that the Elman network (5) 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 and Hall rely 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 don’t 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.

NEW REFERENCES

Thompson, R, Oden, D and Boyson, S (in press) **Language-naive chimpanzees (pan troglodytes) judge relations between relations in a conceptual matching-to-sample task** JOURNAL OF EXPERIMENTAL PSYCHOLOGY: ANIMAL BEHAVIOR PROCESSES

Hutchins, E (6) **COGNITION IN THE WILD** (MIT Press)

Clark, A (in press) **BEING THERE: PUTTING BRAIN, BODY AND WORLD TOGETHER AGAIN** (mit press)

Clark, A (7) **MICROCOGNITION** (mit press)

References

- [1] Clark, P. and Niblett, T. (1989). The CN2 induction algorithm. *Machine Learning*, 3 (pp. 261-283).
- [2] Karmiloff-Smith, A. (1992). *Beyond modularity: a developmental perspective on Cognitive Science*. Cambridge, Ma.: MIT Press/Bradford books.
- [3] Karmiloff-Smith, A. and Clark, A. (1993). What's special about the development of the human mind/brain?. *Mind and Language*, 8, No. 4 (pp. 569-581).
- [4] Thornton, C. and du Boulay, B. (1992). *Artificial Intelligence Through Search*. Intellect Press.
- [5] Elman, J. (1993). Learning and development in neural networks: the importance of starting small. *Cognition*, 48 (pp. 71-99).
- [6] Kirsh, D. and The Intelligent Use of Space, (1995). *Artificial Intelligence*, 72.
- [7] McCloskey, M. and Cohen, N. (1989). Catastrophic interference in connectionist networks: the sequential learning problem. In G.H. Bower (Ed.), *The Psychology of Learning and Motivation*. New York: Academic Press.