

How Similarity Affects the Ease of Rule Application

Ulrike Hahn (sapuh@cardiff.ac.uk)

School of Psychology; Cardiff University
Cardiff CF1 3YG - Wales, UK

Mercè Prat-Sala (M.Prat-Sala@wkac.ac.uk)

Department of Psychology; King Alfred's College
Winchester SO22 4NR - England, UK.

Emmanuel M. Pothos (E.Pothos@ed.ac.uk)

Department of Psychology, University of Edinburgh
Edinburgh EH8 9JZ - Scotland, UK

Abstract

There is good theoretical reason to believe that the application of explicit rules does not proceed in a strictly context free manner, whereby the features stipulated by the rule are simply checked one by one. The fact that specifications of general knowledge seem inherently prone to exception suggests that a more flexible approach is required. One way of balancing the simplicity of rules with the need for flexibility is through the combination of rule application with the monitoring of instance-similarity. As a test of this hypothesis, this paper reports an experiment which examines effects of instance similarity on the speed with which a simple explicit rule can be applied, both as a function of experience with the rule and its complexity.

Introduction

Since its very beginnings Cognitive Science has sought to establish the role of rules in human cognition. However, this work has focussed primarily on *internal*, often implicit, rules. With this we mean rules which are internal to the cognitive system and typically unavailable for conscious inspection, and which have no external, public manifestation. Hence this work has had little to say about how we reason with external, explicit rules such as legal rules or explanatory rules provided in educational settings.

This state of affairs is highly unsatisfactory last but not least because evidence for internal rules has been harder to come by than cognitive scientists originally thought, with any claim for rule-based behaviour typically countered by alternative –most frequently similarity-based- explanations, and supposedly supporting data. By contrast, the existence of vast numbers of external rules is beyond doubt. How we reason with these should thus be a central concern.

We present here an experimental investigation of the way in which explicit rules are applied, examining specifically the role of similarity in this process.

General Background

Past research within cognitive psychology has sought evidence for internal rules in a wide range of domains, from language, to implicit learning, reading and problem-solving. Each of these areas has seen intensive debate between proponents of rule-based accounts and proponents of alternative, similarity-based explanations (see e.g., Hahn & Chater, 1998a for an overview). For example, categorization might construed as the application of rules the learner has abstracted during learning (e.g., “If it is furry, four legged and barks, then the creature is a dog”) or similarity comparison to known exemplars or a prototype (e.g., “This creature is so similar to Lassie, that it must be a dog”). In this way, rules and similarity have typically been viewed in opposition.

However, detailed computational considerations –which draw lessons from Artificial Intelligence – suggest that neither purely rule-based reasoning, nor purely similarity-based reasoning are optimal or even feasible strategies for real-world tasks (Hahn & Chater, 1998a, 1998b; Oaksford & Chater, 1991). Such considerations lessen the plausibility of any cognitive account that seeks to explain human performance solely in terms of one or the other.

However, similarity- and rule-based reasoning differ in their respective strengths and weaknesses. Hence, computational considerations suggest that ‘blends’, which combine the strength of both, are an extremely interesting class of account (Hahn & Chater, 1998a, 1998b). This is reflected in a recent interest in hybrid experts systems within AI which encompass both rule and similarity-based components (e.g., Rissland & Skalak, 1991).

Furthermore, it is highly suggestive that law, next to science the single most elaborate and explicit system we have developed for dealing with every-day life, displays both similarity- and rule-based reasoning in the form of precedent and statute. While legal systems differ regarding the relative weight they place on each of these factors (e.g., the Anglo-American legal tradition emphasises similarity to past cases, and the continental tradition emphasises rules),

the 'blend' of both is a robust finding for all western legal systems.

Together, these considerations suggest that research might profitably turn toward studying the potential interplay between rules and similarity in human thought.

The particular way we propose to do this also turns away from the focus on internal rules to external rules which, as the examples of both law and educational instruction show, are of indisputable significance in human cognition.

Previous Research

In comparison to the wealth of research examining either rules or similarity (see Hahn & Chater, 1998a for references), the body of previous experimental research examining a possible *interplay* of rules and similarity is tiny (e.g., Ross, Perkins & Tenpenny, 1990; Allen & Brooks, 1991; Nosofsky, Palmeri and McKinley, 1994). Of this work, only two studies consider explicit, external rules that are given to participants by the experimenter -Nosofsky, Clark & Shin, 1989 and Allen and Brooks (1991). Both find effects of exemplar similarity in the context of rule application. This is first positive evidence for a routine interaction between similarity and the application of explicit rules as might seem desirable in the light of computational considerations. Further examination, however, is required. In particular, the Allen and Brooks (1991) study does not provide the most robust test of the hypothesis that similarity generally influences rule application. Allen and Brooks provided participants with an explicit rule by which a set of training stimuli could be perfectly classified. Despite the fact that the rule was perfectly predictive, i.e. sufficient for classification, and that participants were both aware of the rule and the instruction to use it, their performance on novel items was significantly affected by the test items degree of similarity to items seen during training. However, the nature of the rule used raises worries about the robustness and generality of their findings.

Crucially, the explicit rule used in the experiment defined a prototype. The use of similarity might conceivably be an (artefactual) result of this, rather than a property of rule application in general. Specifically, the rule used in the Allen & Brooks study had the form "if X has 3 of the 5 features{a,b,c,d,e}, then X is a category member". In other words, the rule specifies a so-called **m-of-n concept** and thus is formally equivalent to a prototype (n, i.e. an item with *all* relevant features present) and a threshold (m, i.e. the number of matching features required) which determines the degree of similarity to the prototype which items must have in order to be category members (Langley, 1994; Hahn & Chater, 1998a). This equivalence makes Allen and Brooks' similarity effects rather less surprising. The rule effectively defines a prototype plus similarity threshold, thus virtually suggesting the general use of similarity to participants. Thus, the fact that even similarities irrelevant from the perspective of the rule influenced classification, Allen and Brooks' finding, might have arisen only because the nature of the rule pointed toward similarity in the first place.

To establish more generally a role for similarity in rule application we need to repeat the basic Allen and Brooks study with a rule that does not involve the specification of a prototype. It is crucial to see what happens if the rule is something like "choose all symmetric patterns," or "all patterns that have an even number of corners". Does one still find effects of similarity to training items on classification speed and accuracy with such a rule?

The second limitation of the Allen and Brooks' study concerns rule complexity. The rule used by Allen and Brooks is a fairly complicated one, so that participants might conceivably have had difficulty operationalising it. Thus Allen and Brooks similarity effects might be the result of alternative, "fall-back" strategies on the part of the participants or of errors arising from failure to use the rule, and not -as we would like to see- the result of genuine interaction between rule and similarity-based processing.

This suggests that it is crucial to look also at asymptotic performance, (that is performance after several experimental trials), to see whether influences of previously seen exemplars disappear as participants' accuracy with the rule grows.

Our study seeks evidence for a constructive interplay of similarity and rule application which avoids these limitations. To this end we conduct a simple replication with a different type of rule, varying rule complexity, and monitoring asymptotic performance.

Experimental Investigation

Our materials were simple geometric shapes as depicted in Figure 1 below.

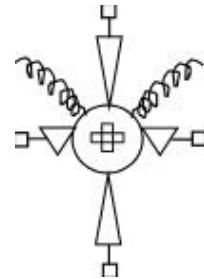


Figure 1: an example of a training material used in the experiment.

Category membership was governed by a simple, explicit rule "*is an A if it has an upside down triangle at the side*" in one-feature rule condition, and by a more complex rule, which made reference to three features -"*is an A if it has an upside down triangle at the sides, a cross in the centre, and a curly line at the top*" - in the complex rule condition. Neither of these constitute an m-of-n concept as did Allen & Brooks' rule. Testing and comparing both a simple and a more complicated rule would further allow us to ascertain the robustness of any putative similarity effect. Specifically, it would allow us to examine whether the similarity effects in the Allen & Brooks study were simply down to participants difficulties in operationalising the rule and, as a result, seeking out alternative, easier strategies.

To the same end, we also monitored participants' performance over four blocks of training –each of which comprised 24 exposures to the test stimulus items. This would allow us to examine whether similarity effects vanished as participants became more and more proficient with the rule.

At the heart of the study is its similarity manipulation. This was achieved by manipulating the additional features not referenced by the rule. Participants received a set of training items to illustrate the rule they had been provided with. The actual experimental test items were either high in similarity or low in similarity to these illustrative training items. The high-similarity test items differed from the training items in one feature; this feature was not referenced by the rule and consequently was irrelevant to the categorization task at hand. The low similarity items differed from the illustration items in 3 features, again features irrelevant to the application of the rule. Because the similarity manipulation concerned only stimulus aspects which were irrelevant to the rule, a difference in the way the high- and low-similarity items were treated would indicate an influence of similarity on the rule-based classification process, despite these differences in similarity being completely irrelevant to the application of the rule. This would provide evidence for an automatic monitoring and processing of similarity information in the context of rule application even where there were no task demands to necessitate this.

To make the classification task a genuine one, it was necessary to have both test items that complied with the rule and ones that didn't. The non-compliant test items, again, were either high- or low- similarity to the illustrative test items, differing from these items in a single feature which contradicted the rule, in the case of the high-similarity non-compliant items, or in a rule-feature and two further irrelevant features in the case of the low-similarity, non-compliant items. As a consequence, similarity effects could emerge both where the rule was applicable and where it did not apply.

The central prediction of this study was that there should be a difference in the speed with which the rule was applied between high- and low-similarity items. Specifically, compliant items which were also high in similarity to the training items should be classified more quickly than compliant items which were low in similarity. Differences should also emerge between the non-compliant high- and low-similarity items, though the direction of the difference is harder to predict here; low-similarity items might be rejected more quickly, but advantages at the decision-making stage are likely to be offset to greater or lesser extent by increased costs of processing a less familiar image. In addition, we would expect an effect of training on reaction time, and that the three feature rule should be slower to apply than the one feature rule because it requires a more complete scan of the stimulus, although these two predictions have no bearing on our experimental question.

With regards to similarity effects, we would also expect differences in the amount of errors elicited by high- and low-similarity items. Low-similarity compliant items, and

high-similarity non-compliant items expected to be more error prone than their counterparts.

Method

Participants

Ninety-one undergraduate students from the University of Bangor (Wales) participated in this experiment as an extra credit option in a Psychology course. 45 participated in the simple rule condition while 46 took part in the complex rule condition.

Materials

The stimuli were line drawings of geometric shapes that varied in 6 aspects: Body, Side/Ears, Top/Bottom, Inner, Antenna, and Hair. There were six alternative realisations of each of these aspects. We generated a total of 108 stimuli: 12 training items and 96 test items. Figure 1 gives an example of a training material used in the experiment.

The 96 test items were composed of four different types of items which formed to four conditions: 24 high-similarity (C-High) and 24 low-similarity (C-Low) (items that complied with the rule), and 24 high-similarity (N-High) and 24 low-similarity (N-Low) (items that did not comply with the rule). High-similarity items differed from (one of the items of) the training set in one value. Low-similarity items differed from (one of the items of) the training set in three values. For the items that complied with the rule, the differed value was always an irrelevant feature. For the items that did not comply with the rule, the differed feature was a value of the rule for the high-similarity items, and a value of the rule plus two irrelevant values for low-similarity items.

Apparatus

The experiment was controlled by the ERTS software run on a PC computer. The pictures were presented as black line drawings on a white background on a (640x480) VGA monitor. Participants used a two-key response pad attached to the ERTS EXKEY-logic connected to the computer to express their response. One of the key was labelled YES and the other was labelled NO. Participants used their dominant hand to press the YES key. Participants' key responses and time elapsed from the presentation of an item to the participants' response were recorded.

Procedure

Participants were tested individually. They were seated in a quiet booth at a comfortable viewing distance in front of a monitor. They received the instructions displayed on the monitor. The instructions told them that they would see a series of pictures of abstract objects. Some of the pictures corresponded to 'good objects' and some to 'bad objects'. Participants' task consisted of pressing the key labelled YES if the abstract object was a 'good object' and press the key labelled NO if the abstract object was a 'bad object'. Before the initiation of the experiment, participants were presented with 10 trials with either YES or NO displayed

on the screen to familiarise them with pressing the YES/NO keys. Participants were given a rule (see above) to determine whether an object was a good object or not. Immediately after the written description of the rule, participants received a graphic example of the rule.

The experimental session was divided into 5 blocks: there was 1 training block and 4 test blocks. During the training phase participants were given 36 trials made up of 12 training items seen 3 times each in a random order. Participants controlled the speed of presentation of each of the training items using the space bar.

After the training phase participants were given four blocks of test. Each test block consisted of 24 items selected randomly from the total of 96 test items, with the only constraint that each test block contained 12 items that comply with the rule (6 high-similarity items and 6 low-similarity items) and 12 items that broke the rule (6 high-similarity do non-compliant items and 6 low-similarity non-compliant items). The test items of each block were presented in a different random order to each participant.

After each of test block participants were reminded of the rule and were shown the 12 items from the training phase (this time they saw each training item only once). This was done with the aim to reinforce rule application. As during the training phase, this repeated training phase was controlled by the participants using the space bar. Participants were asked to respond as quickly as possible without compromising accuracy.

Results

We begin with the *error* analysis. In the case of the simple, 1 feature rule, 123 (2.84%) from a total of (96x45) 4320 responses were errors, that is the participant either pressed the YES key when the NO key was appropriate according to the rule or the other way round (see Table 1). For the complex, 3 feature rule, the data of 4 participants was not included because they failed to respond using the correct key for all the trials of at least one of the blocks in one of the conditions. From a total of (96x42) 4032 responses, 305 (7.56%) were errors were the participant pressed the YES key when the NO key was expected or the other way round (see Table 1).

Table 1: Errors for both 1 and 3 feature rules. “C” indicates items which comply with the rule, “N” items which violate it, “High”- and “Low” refer to the degree of similarity to the test items.

	C-High	C-Low	N-High	N-Low
Errors 1	26	33	36	36
Errors 3	46	61	92	107

We begin with the analysis of the simple, 1 feature rule. Though there were slightly more errors on the low similarity items which complied with the single rule, as we had expected, a paired t-test comparing the proportion of errors made by each subject in across high vs. low similarity comply did not reach significance ($t(44) = 1.26$, $p = 0.1$, one-tailed). The expectation that on the non-rule compliant items there should be more errors on the high similarity

items than on the low similarity non-compliant items was not born out at all, with equal numbers of errors in both cases. In summary, the error analysis for the simple rule revealed no effects of similarity on rule application.

A different picture is presented by the error data for the complex, 3 feature rule. Again, there were more errors on the low-similarity compliant items than on the high-similarity compliant items as expected, but here the difference is statistically significant ($t(41) = 2.10$, $p < 0.02$, one-tailed). Again, the expectation that, of the non-compliant items, it should be the high similarity ones which elicit more errors was not born out in that participants made *more* errors on the low similarity, non-compliant items (Table 1); however, this unexpected difference, tested post hoc, was not significant ($t(41) = 1.48$, $p = 0.115$, two-tailed). In summary, at least for those items which comply with the rule, a significant effect of similarity on rule application emerged.

Finally, a comparison of the error proportions for the simple and complex rule revealed the expected higher level of errors in the complex rule condition: two 2-way Analysis of Variance (ANOVA) (one for the compliant items and another for the non-compliant items) with similarity (High vs. Low) within subject and rule (simple vs. complex) between subject analysis showed a main effect of rule complexity for the non-compliant data ($F(1,85) = 18.56$, $p < 0.0001$) but not for the compliant data ($F(1,85) = 2.21$).

In summary, the analysis of the error data revealed similarity effects for the complex rule, but not the simple rule. We next ask whether this finding is confirmed by participants' reaction times.

Reaction Time Analysis. We begin with the simple, 1 feature rule condition. For each participant we calculated the mean reaction time of response for each condition per block, giving us 16 data points per participant. These data point were transformed into their logarithm. This formed the bases for our analyses. Table 2 shows the mean reaction time per condition and block across all participants.

Table 2: mean reaction time across participants for one-feature rule.

	C-High	C-Low	N-High	N-Low
Block 1	565	587	662	626
Block 2	541	540	548	521
Block 3	514	515	530	583
Block 4	534	513	551	528

We analyse the compliant and the non-compliant items separately, not only because we expect different patterns of result, but because two different hands were used for YES and NO responses and therefore the RT's from the dominant hand are not directly comparable to those of the non-dominant hand.

A two-way ANOVA (fully within) was performed on the reaction time data for the *compliant* items. The variables were block (1-4) and similarity (high vs. low). A main effect of block was found ($F(3,132) = 7.15$, $p < 0.0001$). No main effect of similarity was found ($F(1,44) = .083$) and there was no interaction ($F(3,132) = .102$).

A further two-way ANOVA (fully within) was performed on the reaction time data for the *non-compliant* items. Again a main effect of block was found ($F(3,132) = 29.32, p < 0.0001$). No main effect of similarity was found ($F(1,44) = .804$) but the interaction was significant ($F(3,132) = 8.92, p < 0.0001$).

Repeating these analyses for the complex rule condition, we calculated, for each participant, the mean reaction time of response for each condition per block, giving us 16 data points per participant. These data points were transformed into their logarithm to minimise individual differences. This formed the bases for our analyses. Table 3 shows the mean reaction time per condition and block across all participants.

Table 3: Mean reaction time across participants for complex-rule condition.

	C-High	C-Low	N-High	N-Low
Block 1	1166	1276	1056	1166
Block 2	1018	1077	815	837
Block 3	934	976	723	736
Block 4	902	939	732	780

A two-way ANOVA (fully within) was performed on the reaction time data for the *comply* condition. The variables were block (1-4) and similarity (high vs. low). A main effect of block ($F(3,123) = 13.67, p < 0.0001$) and of similarity ($F(1,41) = 18.42, p < 0.0001$) was found. However, the interaction did not reach significance ($F(3,123) = 1.58$).

A further two-way ANOVA (fully within) was performed on the reaction time data for the *non-compliant* items. Again a main effect of block ($F(3,123) = 69.69, p < 0.0001$) and of similarity ($F(1,41) = 12.78, p < 0.001$) was found. However, the interaction did not reach significance ($F(3,123) = 1.72$).

Finally, comparing the simple and the complex rule condition, participants were significantly faster classifying on the basis of the simple rule, as predicted for both the compliant condition ($F(1,85) = 173.93, p < 0.0001$) and the non-compliant condition ($F(1,85) = 121.80, p < 0.0001$).

We illustrate the implications of the above analyses of reaction times with reference to two graphs plotting the mean reaction times for both kinds of rules, and high and low similarity items, with one graph each for the compliant (Figure 2) and the non-compliant items (Figure 3).

The main interest of the experiment, lies in potential differences between high- and low-similarity items, both as a function of rule complexity and of the amount of training.

As can be seen from Figure 2, the difference in reaction time for high and low similarity items is much greater for the complex rule than it is for the simple rule.

The difference also does not seem to change much with practice. These informal observations are confirmed in the above analyses in that there is no consistent effect of similarity in the case of the simple rule (the ANOVA found no main effect of similarity, nor any interaction between similarity and block) while there is an effect of similarity for the complex rule. However, this effect does not change significantly with practice, that is across blocks (the

ANOVA showed a main effect of similarity, but no interaction).

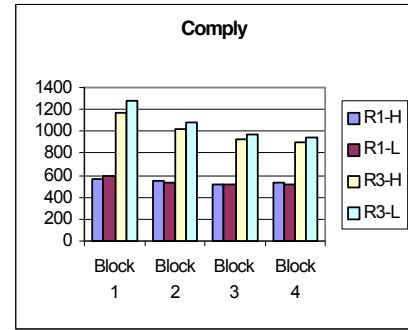


Figure 2: Plotted along the x- axis is the number of blocks, plotted along the y axis is the mean reaction time. R-1 refers to the simple rule, R-3 to the complex, and H and L stand for high and low similarity to training items.

The non-compliant items deviate from this picture only slightly. As can be seen from Figure 3, the magnitude of the difference between high- and low-similarity is again greater for the complex rule, which also shows no clear effect of practice. This is confirmed by the above statistical analysis: for the non-compliant items in the complex rule condition there is, once again, a significant main effect of similarity, but no interaction between similarity and block. The data for the non-compliant items of the simple rule are not quite as clean showing an anomalous increase on block three. As a result of this increase, there is a significant interaction between similarity and block, but as before with the compliant items, there is no main effect of similarity. Crucially, the failure to show a consistent similarity effect observed in the compliant items is thus replicated in the non-compliant items.

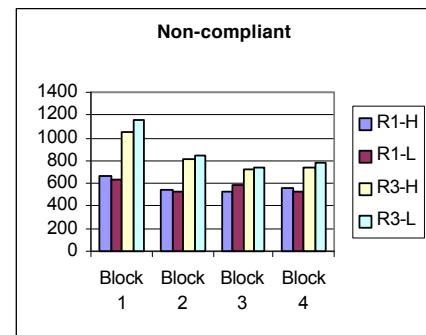


Figure 3: Plotted along the x- axis is the number of blocks, plotted along the y axis is the mean reaction time. R-1 refers to the simple rule, R-3 to the complex, and H and L stand for high and low similarity to training items.

In summary, then, we find similarity effects only for the complex rule, not the simple rule – a result which corresponds with the findings from the error analysis. Furthermore, there is no evidence that this similarity effect goes away with increased practice.

Discussion

The most important result of this experiment is that it manages to replicate the previously observed intrusion of similarity on rule-application even though there was nothing about the rules themselves which in any way made similarity seem relevant.

The similarity effect was apparent both in the error patterns and in participants reaction times and the finding that the similarity effects (where present) did not seem to abate with increased practice in applying the rule further underscore the generality of the observed interplay between rule-application and similarity assessment.

However, we failed to find a consistent effect of similarity in the simple rule condition. One must be cautious in interpreting such a null-effect in a single study in that an effect might well be observable given slight changes in experimental procedure or a considerably larger sample. However, the failure to find an effect does suggest that the role of similarity is not independent of rule-complexity and in that sense neither ubiquitous nor completely automatic. In contrast, the similarity effect found on the complex rule is very robust.

Why might one find effects of similarity in the context of rule application at all, and can our differential results for rules of different complexity be linked into an explanation? The weakness of rule-based reasoning lies in the fact that it is so exceedingly difficult to come up with perfectly predictive rules. Most regularities seem to be prone to countless exceptions, or hold only relative to certain background assumptions which are virtually impossible to capture. For example, the rule "birds fly" is true by and large, but there are some birds which don't. Making rules more specific doesn't eliminate the problem: "robins fly" seems true enough, but, again, will be false if the robin in question has broken its wing, has its feet stuck in concrete, or has eaten too many worms... The potential exceptions are endless and there seems to be no clean cut way of ruling them out in advance. These difficulties have dogged rule-based approaches within Artificial Intelligence which were once held to lead to "expert behaviour" within decades. The frame problem, the difficulties encountered by the naive physics project (Hayes, 1979), and the difficulties in formalising defeasible inference (e.g., Reiter, 1980) are testimony to these difficulties (see also Pickering & Chater, 1995, for discussion).

Our assumption is that similarity might go some way toward alleviating these difficulties, thus allowing human beings to harness the undoubted power and clarity provided by explicit rules. Specifically, tracking the similarity of a potential candidate instance of the rule to previously encountered instances may provide a means whereby one is alerted to potentially deviating circumstances, where the rule -though seemingly applicable- does not, in fact apply. The fact that a novel instance seems dissimilar from previous instances might be a clue to the fact that it should actually be treated differently. This, we would argue, is one of the reasons why an interplay between rules and similarity seems profitable, and hence, is pervasive in large scale legal systems such as the law (see Hahn & Chater 1998a and 1998b for fuller discussion). If this is true, tracking

similarity in the context of rule application would seem useful in most contexts, but we might expect greater reliance on similarity under conditions of increased uncertainty. One way of looking at the increase in rule complexity is as an increase -for participants- in uncertainty. This also seems compatible with Nosofsky, Clark and Shin's (1991) observation of similarity effects despite an explicit rule as their task required fine-grained perceptual distinctions along continuous valued dimensions, which it seems unlikely could be achieved exclusively by the rule. However, further experimentation is required to establish whether the uncertainty reduction is indeed central. The most obvious experimental path to pursue is that of increasing uncertainty in the case of the simple rule, for example by introducing exceptions, thus making the classification task "noisy", or by making the differences in appearance even more extreme. In the meantime, what our results suggest is that instance similarity has some role to play in rule application – a role which needs both further clarification and a satisfactory explanation.

Acknowledgements

The work reported in this paper was funded by a grant from the Nuffield Foundation.

References

Allen, S. & Brooks, L. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, 120, 3-19.

Hahn, U and Chater, N. (1998a). Similarity and Rules: Distinct? Exhaustive? Empirically Distinguishable? *Cognition*, 65, 197-203

Hahn, U. & Chater, N. (1998b). Understanding Similarity: a Joint Project for Psychology, Case-Based Reasoning and Law. *Artificial Intelligence Review*, 12, 393-429.

Hayes, P. (1979). The naive physics manifesto. In D. Michie (Ed.), *Expert systems in the Micro-electronic Age*. Edinburgh University Press, Edinburgh.

Langley, P. (1994) Machine Learning. Addison Wesley

Nosofsky, R., Clark, S. & Shin, H. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 282-304.

Nosofsky, R., Palmeri, T. & Mckinley, S. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.

Oaksford, M. & Chater, N. (1991). Against Logacist Cognitive Science. *Mind and Language*, 6, 1-38.

Pickering, M. and Chater, N. (1995). Why cognitive science is not formalized folk psychology. *Minds and Machines*, 5, 309-337.

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81-132.

Rissland, E. & Skalak, D. (1991). CABARET: rule interpretation in a hybrid architecture. *International Journal of Man-Machine Studies*, 34, 839-887.

Ross, B., Perkins, S. & Tenpenny, P. (1990). Reminding-based Category Learning. *Cognitive Psychology*, 22, 460-492.