

Learning to Make Decisions Under Uncertainty: The Contribution of Qualitative Reasoning

Richard Cooper

Department of Psychology
Birkbeck College
University of London
Malet St., London, WC1E 7HX
R.Cooper@psyc.bbk.ac.uk

John Fox

Advanced Computation Laboratory
Imperial Cancer Research Fund
Lincoln's Inn Fields
London, WC2A 3PX
jf@lif.acl.icnet.uk

Abstract

The majority of work in the field of human judgement and decision making under uncertainty is based on the use and development of algebraic approaches, in which judgement is modelled in terms of mathematical choice functions. Such approaches provide no account of the mental processes underlying decision making. In this paper we explore a cognitive model (implemented within COGENT) of decision making developed in order to account for subject performance on a simulated medical diagnosis task. Our primary concern is with learning, and empirical results on human learning in the modelled task are also reported. Learning in the computational model shares many qualitative features with the human data. The results provide further support for cognitive (i.e., non-algebraic) approaches to decision making under uncertainty.

Introduction

Mainstream approaches to Judgement and Decision Making (JDM) aim to predict how the decisions that people make are affected by different conditions. Normative algebraic techniques (e.g., Expected Utility Theory: Lindley, 1985) have long been dominant in the field. These approaches generally describe decision making in terms of conditional probabilities of outcomes and their associated expected utilities. However, many features of real human decision making are non-normative. Such features have been accounted for by the introduction of a variety of heuristics and biases (e.g., Tversky & Kahneman's (1974) availability and representativeness heuristics). The resultant approaches are (arguably) adequate for predicting statistical regularities across human decision making.

However, because algebraic techniques do not address decision making from an information processing perspective, they cannot take account of situational factors (e.g., specific task requirements) and cognitive factors (e.g., memory limitations and specific subject strategies). Consequently, such theories are restricted in their ability to account for the detailed structure of the decision making process. This situation is further exacerbated by the reduction of individual subject knowledge and experience to conditional probabilities.

In order to address these shortcomings several authors have argued for the development of models and theories of the processes underlying JDM (e.g., Fox, 1980; Beach, 1990; Busemeyer, Hastie, & Medin, 1995; Fox & Cooper, 1997). In addition to addressing the above difficulties, such approaches are able to bring the results and techniques of research in

other areas of cognitive psychology (e.g., memory, perception, problem solving, etc.), where issues of process and representation are routinely considered, to bear on decision making, thus bridging the gap which exists between JDM and much of the rest of the discipline.

The account of JDM adopted here takes an information processing perspective, and assumes that decision making can be understood in terms of general purpose (knowledge-lean) rules acting on domain-specific knowledge. Given this view (support for which follows), one step in developing a fully explicit account of JDM is to provide an account of expert (i.e., knowledge-rich) performance. Such an account, however, only addresses the question of *what* domain-specific knowledge is acquired. It does not address the questions of *how* and *when* that domain-specific knowledge is acquired. For this, a processing account must include an account of learning.

A concrete foundation for the current work is provided by a simulated medical diagnosis task. In this task subjects are required to diagnose a hypothetical patient. Subjects are presented with an initial symptom (e.g., the patient is vomiting) and allowed to query the presence of certain other symptoms before giving a diagnosis. Feedback is given on the final diagnosis allowing subjects to learn the task. Performance is measured in terms of diagnostic accuracy (i.e., the percentage of trials on which the diagnosis offered by a subject is the disease used to generate the symptom pattern) and in terms of the number and order of symptoms queried in coming to a diagnosis. Learning is seen in both an increase in diagnostic accuracy and a decrease in the number of symptoms queried before a diagnosis is offered. The task, though considerably simpler than real-world medical diagnosis, offers a number of complexities not seen in many standard learning and rule induction tasks (e.g., Wason's (1960) 2-4-6 task).

In the following section we summarise our previous empirical and computational work using the above diagnosis task. We then present empirical results on learning the task, followed by a computational model of an idealised learner. The idealised learner's behaviour shares many qualitative features with the subject data, providing strong support for both the rule-based account of decision making and the symbolic account of learning. We conclude with a discussion of some issues raised by the learning model and directions for future research.

Background

The diagnosis task used to carry the current work was first introduced by Fox (1980) in an attempt to explore hypotheses about the interaction of knowledge and memory processes in decision making under uncertainty. The task is based loosely on clinical diagnosis. On each trial the subject is presented with one of five symptoms. This symptom is understood to be the presenting symptom of a patient. The subject knows that the patient is suffering from one of five diseases. After being told about the presenting symptom, the subject can ask about the presence or absence of any of the other symptoms, in any order, and can offer a diagnosis at any point. The selection of the presenting symptom, and the answers to any questions asked, are determined by reference to a set of conditional probabilities (cf. table 1).

In the experiment reported by Fox (1980), the subjects were first-year medical students and the diseases were real medical conditions. The conditional probabilities of symptoms given diseases, however, were chosen for experimental convenience. Subjects were able to learn the task surprisingly well, achieving a diagnostic accuracy of over 80% on the fourth block of 25 trials. (Chance performance on this task is 20% and perfect accuracy is not reliably possible due to the inherently probabilistic nature of the symptom/disease associations.) Subjects achieved this accuracy by asking, on average, 2.12 (out of a possible 4) questions. In addition, clear preferences for initial question selection were found.

The primary aim of the study was to provide data with which to compare a family of algebraic models of the subjects' performance with an information processing, knowledge-based account. Within the information processing account, rules are used to infer suspected diseases from presenting symptoms, and then to infer from suspected diseases what further symptoms might be expected. The symptom to be queried is determined from this set, according to either a discriminating strategy (i.e., choose a symptom which one would expect to be present given one suspected disease but absent given another suspected disease) or a verifying strategy (i.e., focus on one disease and check that each of its typical symptoms are present). The order of recall of symptom/disease associations (which is critical to the order of questioning and hence the final diagnosis) was determined from a memory task which was interleaved with blocks of the diagnosis task. Symptom/disease associations which were more quickly confirmed/disconfirmed were assumed to be recalled before symptom/disease associations which were less quickly confirmed.

Computer simulations of both the algebraic models and the information processing model were performed. The algebraic models provided a reasonable fit to the subject data, but the fit between the information processing model and the subject data, on both diagnostic accuracy and patterns of initial question selection, was found to be superior.

This work provides evidence for the sufficiency of qualitative reasoning in decision making under uncertainty. Although probabilistic/frequentistic information is implicit in the model via the availability of knowledge in memory, the model differs from algebraic accounts in that it is purely de-

terministic and it does not explicitly manipulate probabilities.

Despite the success of this work in modelling "expert" performance on the diagnosis task, neither the empirical work nor the modelling work address the issue of learning. From the empirical side, the subjects, being medical students, were clearly not learning the domain-knowledge from scratch. From the modelling side, the computational tools available when the work was originally performed lacked sufficient power to enable additional progress.

Computational modelling tools and techniques have advanced considerably since the information processing model of diagnosis was first developed, and Fox & Cooper (1997) present a reconstruction of the model within the COGENT modelling environment. COGENT is a tool designed for the development of functionally modular models in the box and arrow style (cf. Cooper & Fox, 1997). COGENT allows models to be specified in terms of interacting processes and buffers, so that, in the diagnosis model for example, working memory can be modelled more clearly as a distinct information store (see Fox & Cooper, 1997; Cooper, 1996; Cooper & Franks, 1996, for examples and further details).

Within the reconstruction, the features of memory retrieval claimed to be critical by Fox (1980) (i.e., that more determinate disease/symptom associations will be accessed before less determinate disease/symptom associations) are captured by buffer access functions provided by COGENT. Significantly, manipulation of the retrieval functions supports the original conclusion: the fit between model and subject performance disappears with alternate retrieval assumptions.

Fox & Cooper (1997) present two further models based on the COGENT reimplementations which demonstrate that assumptions about the specific knowledge representation are not critical. These models show that the same behaviour can be obtained from a simple propositional rule system (as used by Fox (1980)), a first-order diagnosis system, or a generic decision procedure.

The first-order diagnosis system and generic decision procedure abstract the domain-specific knowledge of the diagnosis task from the rules employed in making specific decisions. An important residual issue remains, however. How do people actually acquire domain-specific knowledge? Until an answer can be given to this question, the first-order and generic models are cast into doubt.

A Learning Study

In order to investigate learning within the diagnosis task a further previously unpublished study was performed by Fox when the original model was developed. In contrast to the original study, hypothetical diseases — deptinnitis, malengitis, ritengitis, tepittitis and parontitis — were employed, thus preventing any biases from subjects' previous experience.

Two conditions were investigated. In the first condition the relationship between diseases and symptoms was "sparse", in that relatively few symptoms were associated with each disease. In the second condition the relationship was "dense": on average more symptoms were associated with each disease. The matrices of the conditional probabilities for each condition are given in table 1. (Thus, from table 1a, it can be seen

	Deptinnitis	Malengitis	Ritengitis	Tepittitis	Parontitis
Stiffness	1.00	1.00	0.00	0.00	0.00
Vomiting	0.00	0.00	0.50	1.00	0.00
Headache	1.00	0.25	1.00	0.00	0.25
Earache	0.00	0.00	0.50	0.00	1.00
Pyrexia	1.00	0.25	1.00	0.25	0.00

a) Probability of symptoms given a disease (sparse condition)

	Deptinnitis	Malengitis	Ritengitis	Tepittitis	Parontitis
Stiffness	0.00	0.00	1.00	1.00	1.00
Vomiting	1.00	1.00	0.50	0.00	1.00
Headache	0.00	0.75	0.00	1.00	0.75
Earache	1.00	1.00	0.50	1.00	0.00
Pyrexia	0.00	0.75	0.00	0.75	1.00

b) Probability of symptoms given a disease (dense condition)

Table 1: The conditional probabilities of each symptom (given a disease) used to generate hypothetical cases.

that the probability of a “patient” having headache given that the patient has malengitis is, in the sparse condition, 0.25.) Note that the dense and sparse conditions are symmetrical in that the conditional probability of a symptom given a disease in one condition plus the corresponding conditional probability in the other condition is always 1.00. The two conditions are therefore logically equivalent: the probability of a patient having disease D with symptoms V , W , and X , but not Y and Z in the sparse condition is the same as the probability of a patient having disease D without symptoms V , W , and X , but with Y and Z in the dense condition. This manipulation was performed in order to investigate the general finding that people tend to focus on positive data. It was anticipated that, despite the symmetry of the conditions, the manipulation would affect the learning speed and accuracy.

The task consisted of three blocks each of twenty five trials. On each trial, a disease was chosen at random (subject to the constraint that each disease appear five times in each block). Once a disease was selected for a particular trial, a presenting symptom was generated with reference to the conditional probability matrix. The probability of a symptom being selected as a presenting symptom for a disease was proportional to probability of the symptom given the disease. Subjects were then able to query the presence or absence of further symptoms until offering a diagnosis.

Eight subjects took part in each condition. (16 subjects in total.) Their diagnostic accuracy and mean number of questions asked over each block were recorded. Summary statistics are shown in table 2.

Looking first at the results of the sparse condition, it can be seen that there is a substantial increase in diagnostic accuracy between the first and second blocks, but no change between the second and third blocks. The increase between blocks 1 and 2 is significant (Wilcoxon matched-pairs signed-rank test, $T = 4$, $n = 8$, $p < 0.05$). In contrast, the number of symptoms queried shows a continued decrease over all blocks. The decrease between blocks 1 and 2 approaches significance (Wilcoxon, $T = 6$, $n = 8$, $p \approx 0.055$), while the decrease

	Block 1	Block 2	Block 3
Accuracy	44% (16%)	70% (28%)	70% (16%)
Queries	3.18 (0.40)	2.48 (1.54)	2.04 (0.86)

a) Subject Learning Data (sparse condition)

	Block 1	Block 2	Block 3
Accuracy	44% (20%)	76% (28%)	70% (20%)
Queries	3.66 (0.72)	3.36 (0.60)	2.82 (0.48)

b) Subject Learning Data (dense condition)

Table 2: Subject learning data in both conditions. Data is given as median (interquartile range).

between blocks 2 and 3 is significant (Wilcoxon, $T = 3$, $n = 8$, $p < 0.02$).

Similar comments apply to the results from the dense condition. Again diagnostic accuracy is significantly greater on block 2 than on block 1 (Wilcoxon, $T = 0$, $n = 7$, $p < 0.01$), and the number of symptoms queried decreases between blocks 1 and 2 and blocks 2 and 3 (though the decreases are not significant — more data would seem to be required: Wilcoxon, $T = 3$, $n = 6$, $p < 0.08$, and Wilcoxon, $T = 1$, $n = 4$, $p < 0.10$, respectively).

Interestingly, diagnostic accuracy is not affected by the sparse/dense manipulation, but the number of symptoms queried is, being significantly fewer in the sparse condition across corresponding blocks (Wilcoxon/Mann-Whitney: Block 1: $w_s = 36.5$, $n_1 = n_2 = 8$, $p < 0.001$; Block 2: $w_s = 50.5$, $n_1 = n_2 = 8$, $p < 0.05$; Block 3: $w_s = 48$, $n_1 = n_2 = 8$, $p < 0.025$)

A Model of Learning

The model of learning which we have developed builds additionally on the existing COGENT model of “expert” performance on the diagnosis task. It was developed in order to test both the original model — Is it consistent with processes required

for learning? — and the COGENT modelling environment — Does it provide sufficient additional modelling power, beyond the tools available when the initial production system model of the task was developed, to account for the learning data?

The Basic Model

Figure 1 shows the functional modules (as they appear in the COGENT specification) of the diagnosis model with learning. The diagrammatic representation of the model of expert performance is the same except that it lacks *Learning Mechanism* and its associated arrows.

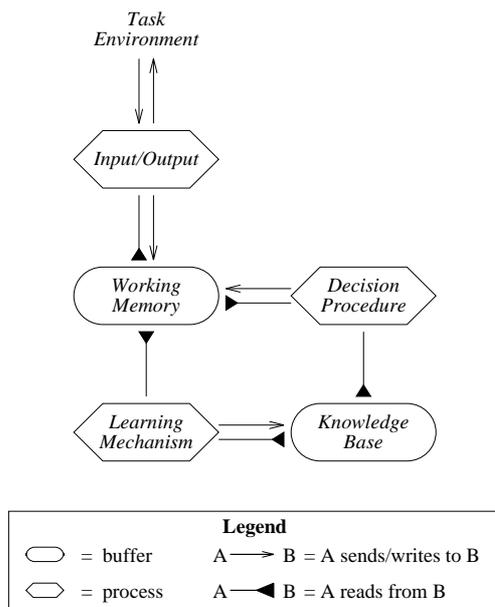


Figure 1: Box/Arrow diagram of the model with learning

The function of each of the boxes is as follows. *Task Environment* (the detail of which is not shown) generates subject data, presents it to the rest of the model, answers queries concerning the presence/absence of symptoms, and records all protocols. It is not part of the cognitive model but is implemented within COGENT so as to automate the data presentation and analysis. *Input/Output* models the subject's perceptual/articulatory processes. Messages from *Task Environment* trigger additions to *Working Memory* (e.g., adding information about the presence of a symptom), and the existence of appropriate elements in *Working Memory* trigger generation of articulatory output (e.g., a query about a symptom). *Working Memory* is a passive data store in which information about the current case is stored and manipulated. There is no decay of, or limit on, the information stored here. *Decision Procedure* is a set of inference rules which modify *Working Memory*, implementing the basic diagnostic strategy. The rules are generic and first-order, but instantiated for the particular task by the subject's beliefs about symptoms and diseases (which are stored in *Knowledge Base*). In the model of "expert" performance, *Knowledge Base* is initialised with

information of symptom/disease associations based on subjects' responses to the interleaved memory task in the original study (Fox, 1980), ordered according to response times, such that symptom/disease associations which are more quickly confirmed/disconfirmed are accessed before those which are less quickly confirmed/disconfirmed.

Processing in the expert model is initiated when a presenting symptom is received by *Input/Output* from *Task Environment*. The corresponding information (e.g., `told(vomiting, present)`) is immediately added to *Working Memory*. This triggers a rule in *Decision Procedure* which augments *Working Memory* with the set of diseases which the presenting symptom suggests (ordered according to accessibility in *Knowledge Store*). The presence of suspected diseases in *Working Memory* prompts recall of their associated symptoms (through a second *Decision Procedure* rule). At this stage of processing *Working Memory* might contain the following elements (where the left column shows the cycle number on which the element appeared in the buffer):

```

5: expected(tepittitis, pyrexia, present).
5: expected(tepittitis, vomiting, present).
5: expected(ritengitis, vomiting, present).
4: diseases(tepittitis, suspected).
4: diseases(ritengitis, suspected).
3: told(vomiting, present).
  
```

If at this (or any) stage, there exists a symptom which is explicitly expected to be present given one disease but absent given another, then a discrimination rule will trigger a query over that symptoms' value (`present/absent`). This is not the case in the situation shown above. Instead a more general rule queries the first symptom (that has not previously been queried) retrieved from the knowledge base, in this case pyrexia. The prompt for *Input/Output* to ask the question appears in *Working Memory* on cycle 6:

```
6: query(pyrexia, present).
```

Several cycles later the response to this query appears in *Working Memory* (via *Input/Output*):

```
9: told(pyrexia, absent).
```

This counts against `tepittitis`, but, more importantly, the subject has (incorrectly) inferred from previous experience that the presence of `vomiting` and the absence of `pyrexia` imply that the patient has `ritengitis`. This is thus the diagnosis suggested:

```
10: diagnosis_is(ritengitis).
```

Adding Learning

One advantage of a first order model (over a purely propositional one) is that it separates task knowledge (of which subjects start with none) from strategic knowledge. If we assume that subjects' strategic knowledge does not change throughout the task (not in itself an unproblematic assumption), then the improvement in subjects' performance across blocks can be attributed entirely to the accumulation/modification of task-specific knowledge. The separation therefore makes clear what the subject must learn during the task.

The first-order knowledge store contains information in three forms:

`suggests(Symptom, Disease)`: This is used to determine which diseases should be suspected given the presence of a specific symptom. If `Symptom` is known to be present then `Disease` will be suspected.

`association(Disease, Symptom, Value)`: This is used to determine which symptoms to expect when considering the possibility that the patient has a given disease. `Value` is either present or absent.

`pattern(Disease, SymptomList)`: This is used in actually making a diagnosis. If `Disease` is suspected and the symptom configuration specified in `SymptomList` is known to hold then `Disease` will be offered as the diagnosis.

Learning Mechanism contains rules which specify how each form of information is learnt.

Learning triggering symptoms The model uses two rules to learn `suggests` clauses, both triggered by feedback on a diagnosis appearing in *Working Memory*:

1. `suggests` clauses corresponding to symptoms known to be present for the actual disease are added;
2. `suggests` clauses corresponding to symptoms known to be absent for the actual disease are deleted.

In order to force the model to initially consider all diseases, it is initially assumed that all symptoms suggest all diseases. This is reflected in the initial state of *Knowledge Base*. As learning proceeds, many symptom/disease pairs are deleted from the knowledge base, and the ones that remain are reordered such that the most recently observed associations are retrieved first. This reordering, which is performed automatically through properties specified on the COGENT buffer that implements *Knowledge Base*, is consistent with the results of Fox (1980).

Learning expected symptoms When feedback on a diagnosis is received:

1. if the symptom is present a `present` association is recorded;
2. if the symptom is absent an `absent` association is recorded;
3. if the symptom is unknown then any existing association is deleted.

These rules are naive in that they do not merge existing knowledge with knowledge which may be inferred from the current case. Logically, it might be more correct, for example, to infer no association between a disease and a symptom only when we have seen the disease both with and without the symptom. However, it appears that such intelligent rules are not required in order to simulate subject performance.

It is assumed initially that all diseases and symptoms are positively associated. There is little motivation for this assumption, except that some associations need to be present in order for the model to function. An alternate approach would be to include further special purpose rules. However, we believe that the assumption is not critical, as the above learning rules ensure that it only has an effect on the first few trials.

Learning symptom configurations It is information about disease/symptom configurations (i.e., `pattern` clauses) that

has the largest effect on the model’s performance (both in terms of diagnostic accuracy and number of queries before a diagnosis). Initially, there is no `pattern` information, and diagnoses are based purely on a guessing strategy. On receiving feedback of the correct diagnosis, all known symptoms are collected into a list which is merged with any existing patterns for the disease. If no patterns are present, the list is added as a possible symptom configuration for the disease. Otherwise, if some other pattern exists for the disease, and that pattern differs from the new pattern by just one symptom (which is known to be present in one pattern and absent in the other), then that symptom is removed from the existing pattern, yielding a shorter, generalised, pattern. If the model makes an incorrect diagnosis, and that diagnosis was based on an existing pattern (i.e., it wasn’t a guess), then the incorrect pattern is removed from *Knowledge Base*.

Results

As noted above, *Task Environment* simulated the generation and presentation of stimuli and responses to the model, using the same procedure as was used in the initial study. Using this, the model was run eight times for each condition (sparse/dense). Each run consisted of 3 blocks of 25 trials (as in the initial study) and on each block the mean number of symptoms queried and diagnostic accuracy were recorded (as in the initial study). Summary statistics are shown in table 3.

	Block 1	Block 2	Block 3
Accuracy	56% (0%)	74% (12%)	78% (8%)
Queries	2.92 (0.84)	1.72 (0.28)	1.64 (0.24)

a) Model Learning Data (sparse condition)

	Block 1	Block 2	Block 3
Accuracy	68% (8%)	96% (4%)	96% (8%)
Queries	3.96 (0.08)	3.88 (0.04)	3.82 (0.08)

b) Model Learning Data (dense condition)

Table 3: Model learning data in both conditions. Data is given as median (interquartile range).

The data were analysed using the same non-parametric tests as were used to analyse the subject data. In the sparse condition, the increase in diagnostic accuracy between blocks 1 and 2 is significant (Wilcoxon matched-pairs signed-rank test, $T = 1, n = 6, p < 0.05$), as is the decrease in symptoms queried (Wilcoxon, $T = 0, n = 8, p < 0.005$). This fits with the subject data. However, unlike the subject data there is no significant difference between diagnostic accuracy and symptoms queried between blocks 2 and 3. In the dense condition the pattern of significance follows exactly that of the subject data. Diagnostic accuracy increases between blocks 1 and 2 (Wilcoxon, $T = 0, n = 8, p < 0.005$) but not between blocks 2 and 3 (Wilcoxon, $T = 10.5, n = 7, p > 0.05$), whereas symptoms queried decreases significantly across all blocks (Blocks 1 and 2: Wilcoxon, $T = 3, n = 7, p < 0.05$; Blocks 2 and 3: Wilcoxon, $T = 4, n = 8, p < 0.05$). Both diagnostic accuracy and symptoms queried are significantly

less in the sparse condition than in the dense condition.

Discussion

Although the simulation data differs quantitatively from the subject data, the model exhibits a number of qualitative similarities to the subjects. Firstly, both subjects and the model are affected by the dense/sparse manipulation. Recall that there is no logical difference between the tasks. The diagnosis model treats positively and negatively information differently, however, and this has a very significant impact upon learning.

Data from the sparse condition closely reflect subject performance, although the reduction in symptoms queried over the second and third blocks was not significant. Learning is, if anything, slightly too effective, with the model tending to query fewer symptoms than the subjects. However, the model's fit with some subjects is impressive: subject 6, for example, was 68% accurate in the final block, querying on average 1.70 symptoms.

With regard to the dense condition, the number of symptoms queried decreases across blocks, but not as sharply as in the subject data, and diagnostic accuracy increases to ceiling by the second block. By querying most symptoms the model is able to perform almost perfectly. Closer inspection of the subject data reveals that this pattern of performance did occur: subject 4 was 92% accurate in the final block, querying on average 3.24 symptoms.

There are many aspects of the model which raise further questions. Alternatives exist for many of the learning rules, for example. It remains to be seen how sensitive learning is to modification of these rules. Could similar behaviour result from different rules? If so, are there "critical" features of the learning rules which lead to performance similar to that of our subjects? There was considerable inter-subject variability. Can individual differences be accounted for in terms of different learning strategies?

In a first attempt to look at some of these issues we have considered one variation on the pattern learning rule: this rule, on receiving feedback, finds the largest subset of symptoms shared between the current instance of the disease and any previous instance, and uses that as the basis for future diagnoses. Although superficially sensible, this rule yields particularly bad learning performance, with the model's accuracy being reduced to little more than chance, and with the number of symptoms queried being reduced to near zero. Curiously, this pattern of performance was also observed in several subjects. Were such subjects adopting poor learning strategies, or were motivational factors influencing their performance?

It may also be possible to account for more of the variance in the subject data by the introduction of performance factors. The model as it stands is a competence model. All buffers, for example, are assumed to be perfect storage devices. In the light of this fact it is not surprising that the model generally does slightly better than subjects. (What is perhaps surprising is how well some subjects are able to do the task.) However, a second line of current work is examining the effect of decreasing the efficiency of learning (both by adding probabilistic firing to the learning rules, such that not all rules fire in all situations, and by adding decay to *Knowledge Base*).

Although this appears promising in accounting for many of the more average subjects, we are aware that such an approach needs methodological care in order to avoid charges of data fitting.

Conclusion

Our initial concern was with the development of cognitive, rather than algebraic, models of decision making under uncertainty. We have reported data on learning in a diagnosis task and presented a model of learning in the domain (developed as an extension to an existing model of the task). The resultant model is very encouraging, demonstrating many of the learning effects seen in the subject data. The inclusion of learning further supports the cognitive claims of the original model, and strengthens our argument for the role of cognitive models within the field of judgement and decision making. In addition, the model provides a further demonstration of the power of the COGENT modelling environment, without which the current work would have been vastly more difficult.

Acknowledgements

We are grateful to Peter Yule and two anonymous reviewers for comments on an earlier draft of this paper.

References

- Beach, L. R. (1990). *Image theory: Decision making in personal and organizational contexts*. John Wiley, Chichester, England.
- Busemeyer, J., Hastie, R., & Medin, D. L. (Eds.). (1995). *Decision making from a cognitive perspective*. Academic Press, San Diego, CA.
- Cooper, R. P., & Fox, J. (1997). COGENT: An interactive graphical environment for cognitive modelling. Unpublished: Manuscript available.
- Cooper, R. (1996). Perseverative subgoal in a production system model of problem solving. In *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, pp. 397–402. San Diego, CA.
- Cooper, R., & Franks, B. (1996). The iteration of concept combination in Sense Generation. In *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, pp. 523–528. San Diego, CA.
- Fox, J. (1980). Making decisions under the influence of memory. *Psychological Review*, 87, 190–211.
- Fox, J., & Cooper, R. (1997). Cognitive processing and knowledge representation in decision making under uncertainty. In Scholz, & Zimmer (Eds.), *Qualitative Theories of Decision Making*.
- Lindley, D. V. (1985). *Making decisions* (2nd edition). John Wiley, Chichester, England.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12, 129–140.