UEcho: A Model of Uncertainty Management in Human Abductive Reasoning

Hongbin Wang*(WANG.190@OSU.EDU)
Todd R. Johnson**(JOHNSON.25@OSU.EDU)
Jiajie Zhang*(ZHANG.52@OSU.EDU)
*Department of Psychology
**Division of Medical Informatics, Department of Pathology
Center for Cognitive Science
The Ohio State University
Columbus, OH 43210

Abstract

This paper explores the uncertainty aspects of human abductive reasoning. Echo, a model of abduction based on the Theory of Explanatory Coherence (Thagard, 1992a), captures many aspects of human abductive reasoning, but fails to sufficiently manage the uncertainty in abduction. In particular, Echo does not handle belief acquisition and dynamic belief revision, two essential components of human abductive reasoning. We propose a modified Echo model (UEcho), in which we add a learning mechanism for belief acquisition and a dynamic processing mechanism for belief revision. To evaluate the model, we report an empirical study in which base rate learning serves as a testbed for belief acquisition and the order effect serves as a testbed for belief revision.

Introduction

People live in an uncertain world. When an event happens, the meaning or the implication of the event is seldom completely clear at the very beginning. Uncertainty results in belief and belief guides decisions and actions (Schmitt, 1992). Therefore, accurately evaluating the nature of a belief is one of the fundamental tasks that people have to face in both everyday reasoning and scientific discovery.

Abduction, a distinct type of inference from deduction and induction, is a form of reasoning that infers causes or explanations from effects (see Fann, 1970 for an introduction to the theory of abduction). Since in abduction premises do not guarantee the truth of conclusions due to incomplete information, uncertainty and thus belief play an important role in human abductive reasoning.

The research presented in this paper represents an attempt to explore the essence of uncertainty in human abductive reasoning, and to capture it in a modified Echo model.

This paper is organized as the following three parts. First, some theoretical developments of abduction, uncertainty, and Echo are reviewed. The conclusion is that although in general Echo is a good candidate for modeling human abductive reasoning, in order to handle the uncertainty aspects of abduction, some modifications are needed. In the second part, we describe how to modify Echo to perform uncertainty and belief related functions. The model is evaluated by comparing its performance with empirical data. In the third part, future research plans are discussed.

Abduction

Abduction is an essential component of many tasks, including medical diagnosis (Feltovich, Johnson, Moller, and Swanson, 1984), scientific discovery (Thagard, 1989), and discourse comprehension (Kintsch, 1988). The key task of abduction is to find a best explanation of a set of observations (Peng & Reggia, 1990; Josephson & Josephson, 1994).

Abduction can be represented in the following general form:

The surprising fact C is observed,
But if A were true, C would be a matter of course;
Hence, there is a reason to suspect that A is true.

Clearly, different from deduction, where the conclusion necessarily follows from the premises, in abduction the conclusion does not follow from the premises with necessity. That is, given a set of observations, many hypotheses (or conjunctive hypotheses) can be formed, each of which may have different degrees of plausibility. In general, how do we select a best one?

Many researchers (e.g., Josephson & Josephson, 1994; Thagard, 1992a; Paul, 1993) distinguish two components of an abductive reasoning process. That is, abduction is a process that includes both hypothesis generation (forming a set of plausible hypotheses) and hypothesis evaluation (choosing a best one). Note that this distinction does not imply that abduction is a clean 2-stage process. The fact that people do not exhaustively generate all possible hypotheses indicates that the two components can happen simultaneously.

Uncertainty

Uncertainty is inevitable at all levels-of humans' interaction with their environment. At the lowest level, biological processes are never clear-cut and without noise. At the cognitive level, uncertainty results from inadequate information sources, limited information processing capacity, or ambiguities in natural language.

An important aspect of uncertainty research is how to quantify or measure the uncertainty[1,2]. One classical ap-

[1] Non-quantitative approaches to uncertainty will not be discussed here.

proach is probability theory in general and the Bayesian approach in particular (see Robert, 1994). In the Bayesian approach, uncertainty can be represented by a probability, (P, a number between 0 and 1), interpreted as degree of belief based on all available knowledge. When new evidence comes, belief is updated or revised[3] based on Bayes' Theorem, which is generally regarded as the normative model of belief revision. This approach has a sound theoretical foundation and is easy to understand, thus it has been successfully applied to a wide range of domains. An alternative approach to uncertainty management is Dempster-Shafer (D-S) theory (e.g., Shafer, 1976), which represents the belief about a proposition as an interval [*Bel*, *Pl*]. *Bel* is interpreted as the degree of belief (or credibility) and *Pl* is the degree of plausibility. The interval between the two represents ignorance, which can not be represented by probability theory. Both *Bel* and *Pl* are numbers between 0 and 1, and the two are related in the following way: $Pl(\mathrm{H})=1-Bel(\sim\mathrm{H})$. That is, the degree of plausibility of H is the complement of the credibility of not H. Probability theory is claimed to be a special case of D-S theory, when the [*Bel*, *Pl*] interval degenerated into a point, i.e., $Bel=Pl=P$. It has been shown that D-S theory has its limitations. It is harder to understand and implement than probability theory. In the meantime, it assumes that the set of hypotheses is exclusive and exhaustive and requires an independent body of evidence, which is usually unrealistic.

A number of studies try to identify different types of uncertainty based on the psychological sources of uncertainty. For example, Kahneman and Tversky (1982) distinguish external uncertainty from internal uncertainty. External uncertainty refers to disposition or randomness of external events, which are something people cannot control. An example is the outcome of throwing a coin. External uncertainty can be assessed in either a relative frequency mode or by subjective judgment of a single event. Internal uncertainty refers to ignorance, which results from incomplete knowledge. An example is that one is unsure if Beijing is the capital of China.

Clearly, human abductive reasoning involves both external and internal uncertainty. Therefore, adequately managing the uncertainty is a critical component of human abductive reasoning.

Echo

The Theory of Explanatory Coherence (TEC hereafter; Thagard, 1989, 1992a) is claimed to be a theory of abduction. More precisely, however, it should be called a theory of hypothesis evaluation since it ignores the hypothesis formation part of abduction. According to this theory, the best explanation is the one with the most explanatory coherence based on all current hypotheses, evidence, and explanatory relations.

[2] We do not tackle the problem of imprecision measurement studied in fuzzy sets theory and possibility theory (see Dubois & Prade, 1988 for distinctions between imprecision and uncertainty).

[3] Although people may sometimes distinguish belief updating and belief revision (see Wang, 1993), we use them interchangeably in this paper.

From the perspective of TEC, propositions P and Q cohere if there is some explanatory relation between them. More specifically, propositions P and Q cohere if any one of the following conditions is satisfied:

- P is part of the explanation of Q;
- Q is part of the explanation of P;
- P and Q are together part of the explanation of R; or
- P and Q are analogous in the explanations they respectively give of some R and S.

There are several underlying principles in TEC that provide general guidelines for explanatory coherence evaluation. Some important ones include *symmetry* (if P and Q cohere, Q and P cohere; if P and Q incohere, Q and P incohere), *explanation* (if P1...Pm explain Q, then for each Pi in P1...Pm, Pi and Q cohere; for each Pi and Pj in P1...Pm, Pi and Pj cohere; the degree of coherence is inversely proportional to the number of propositions P1...Pm.), *data priority* (propositions that describe the results of observations have a degree of acceptability of their own), *competition* (if P and Q both explain a proposition, and if P and Q are not explanatory connected, then P and Q incohere), and *acceptability* (the acceptability of a proposition P in a system S depends on its coherence with the propositions in S). Constrained by these principles and acting as a whole, the system tries to pursue the highest explanatory coherence at the system level rather than at the individual proposition level.

Echo is a connectionist implementation of TEC. In Echo, propositions (both data and hypotheses) are represented by nodes. Coherence relations are represented by excitatory links and incoherence relations are represented by inhibitory links. Node activation represents the node's degree of coherence with other propositions in the network. The system updates itself based on parallel constraint satisfaction. Once the system settles down, the best explanation consists of the nodes with highest activation values.

Theoretically, Echo satisfies some critical constraints in abduction. For example, it simultaneously handles several important criteria in hypothesis evaluation, including explanatory breadth (the model prefers a hypothesis that explains more); simplicity (the model prefers a simpler hypothesis); being explained (the model prefers a hypothesis which itself is explained); data reliability (the credibility of an observation also depends on its coherence in the system); and analogy (analogous hypotheses are coherent). Empirically, Echo has also acquired much experimental support (e.g., Thagard, 1989, 1992b).

Something is missing, however. Given the above discussions about abduction and uncertainty, it is clear that Echo does not have enough power to handle various uncertainty aspects in abduction. First, Echo does not handle belief acquisition. Since Echo does not learn from its experience, it has no background knowledge necessary to determine the degree of belief for any given hypothesis and the strength for any given connection between hypothesis and evidence. As a result, it assumes that all hypotheses are equally probable and the strengths of all connections are fixed[4] when

[4] Although you can vary the strength of a particular explanatory relation by specifying a number between 0 and 1, the strength is fixed thereafter.

they are pre-entered into the system. This is not the case in human abductive reasoning. People rarely treat all proposed hypotheses equally and would often assign different prior beliefs to different hypotheses and connections based upon previous experiences. Second, Echo does not handle dynamic belief revision. Echo assumes that all evidence is available at the very beginning, therefore, no belief revision is necessary. This is not a realistic assumption either. Evidence usually does not come all at once. People need to revise their opinions as more and more information becomes available. Since Echo does not do belief revision, it misses this dynamic aspect of abduction.

Modifying Echo to Handle Uncertainty

In this part of the paper, we propose a modified version of Echo (UEcho, for Uncertainty Echo), designed to handle the uncertainty aspects in abduction. In particular, we add a learning mechanism to handle belief acquisition, and a dynamic processing mechanism to handle belief revision. To evaluate UEcho, we report an empirical study in which the base rate learning serves as a testbed for belief acquisition and the order effect serves as a testbed for belief revision.

Before we do that, let us first briefly review some relevant findings and theories in the base rate learning and the order effect literature.

Base Rate Learning and Order Effects

Base rate information describes the statistical properties of the environment in general. It clearly plays an important role in human reasoning. However, literature about base rate acquisition and use is controversial. Although for some researchers, "the genuineness, the robustness, and the generality of the base-rate fallacy" (Bar-Hillel, 1980, p215) suggest that human beings are programmed to systematically and stubbornly ignore base rate information when making judgments under uncertainty, others disagree (Kahneman & Tversky, 1973; for a general discussion on human heuristics and biases, see Kahneman, Slovic, & Tversky, 1982). Systematic investigations have been carried out to determine the conditions under which the base rate fallacy appears or disappears. First, it has been shown that, in some circumstances, people can automatically and accurately acquire and use base rate and frequency information (Hasher & Zacks, 1979, 1984; Medin & Edelson, 1988). Second, it has also been argued that people do not ignore the base rate. Rather, they may simply misaggregate the relevant information (Edwards, 1968). Third, the base rate fallacy disappears when information is presented to subjects in frequency format instead of single-event probability format (Gigerenzer & Hoffrage, 1995; Cosmides & Tooby, 1996). Finally, from an ecological point of view, natural sampling theory (Kleiter, 1993) claims that there is a condition in which it is rational to neglect base rate information since the base rate actually does not enter into the estimated Bayesian probability. This condition is called "natural sampling" which refers to the situation where the structure of the environment is sequentially learned through experience. Ecologically this is the most natural way people acquire frequency information.

Another closely related phenomenon is the order effect. Given that people naturally acquire information sequentially, should the order of information presentation result in a difference in the final result? In other words, will it be possible that people prefer hypothesis H1 when given evidence A first and evidence B second, and prefer hypothesis H2 when given evidence B first and evidence A second? Indeed, this kind of order effect is a fairly robust finding in the human reasoning and judgment literature (see Hogarth & Einhorn, 1992; Schlottmann & Anderson, 1995, for reviews). Hogarth and Einhorn (1992) analyzed the various features of tasks in which order effects occur (or do not occur) and proposed a belief-adjustment model of belief revision. The main assumption of the model is that people use an anchoring and adjustment heuristic — people adjust a current belief (the anchor) on the basis of how strongly new information confirms or disconfirms this belief. Therefore, according to this model, people do not consider all available information each time they need to revise their opinion; instead, they adjust their current belief in the direction of the new information. The order effect results from the differential weighting of the new information. More specifically, the adjustment for negative evidence is bigger when the current anchor is large than when it is small; and the adjustment for the positive evidence is bigger when the anchor is small than when it is large. Therefore, the step-by-step evaluation of beliefs for mixed positive and negative evidence produces a recency effect: the final evaluation of belief is mainly determined by the last evidence item. However, the step-by-step evaluation of beliefs for consistent evidence (all positive or all negative) produces a primacy effect.

Given the above overview, how can Echo be modified to do belief acquisition and belief revision? In particular, how can Echo be modified to model base rate learning and the order effect? We address these issues in the following two sections.

Belief Acquisition

Ideally, the weights of the connections between evidence nodes and hypothesis nodes in Echo should reflect the summarized probability of co-occurrence. However, in Echo, the weights are fixed according to two parameters, *EXCITATORY_WEIGHT* (EW) and *INHIBITORY_WEIGHT* (IW). To learn from experience the weights between node pairs must be modified based on learning experience, so that their magnitude can be tuned to the statistical structure of the environment and reflect the corresponding probability of co-occurrence. One central issue is to select an appropriate weight-updating rule.

Our choice for weight updating is the Rescorla-Wagner (R-W) rule. This choice is made based on the following considerations. First, numerous studies have shown that the R-W rule is good at learning from experience and has been successfully applied to a wide range of domains from behavioral specification to connectionist modeling (see Gluck & Bower, 1988; also see Miller, Barnet, & Grahame, 1995 for a review). Although it has been shown that the R-W rule has its limitations, such as its incapacity of handling complex non-linear problems (see Barto, 1990), we choose it because 1) it demonstrates that a learning mechanism can be added to Echo; and 2) it is adequate to deal with the simple learning task used in this research. Second, on the surface it seems that Echo and the R-W rule are not compatible in the sense that Echo emphasizes symmetry and the R-W rule is directional. We argue that a directional learning rule is appropriate since 1) it is a natural mapping to the causal relations we are modeling; and 2) Even in Echo symmetry is not absolute: it is the hypothesis nodes "cause" the evidence nodes but not vice versa.

The application of the R-W rule in UEcho is illustrated in Figure 1 (where F and H are hypotheses, and R and ID are evidence that can be either positive or negative).

The system runs in the following manner. Given an episode, after it settles down, UEcho evaluates its current belief (0 to 1) about a hypothesis based on its activation (-1 to 1) according to a logistic function. After that, UEcho receives the actual truth-values of the hypotheses and calculates the desired weight changes based on the R-W rule. Finally, UEcho updates its connection weights. The weights are bounded by [IW, EW]. The relative position of the final connection weights in the whole possible range (from EW to IW) will reflect the probability of co-occurrence.

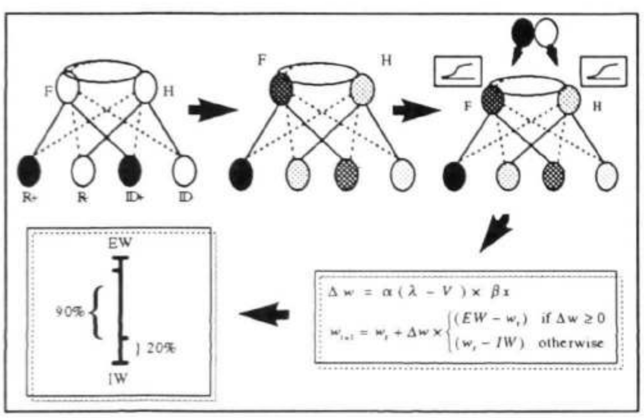

Figure 1: Belief Acquisition in UEcho

Belief Revision

Echo in its original format is operating in an End-of-Sequence (EoS) mode. In other words, the system runs itself after all the necessary hypotheses and evidence are given. In order for UEcho to model order effects, we modified Echo to operate in a Step-by-Step (SbS) mode. That is, when some new evidence is available, UEcho needs to incorporate the new evidence and continue updating itself. Furthermore, in order to account for the full range of order effects, UEcho somehow needs to encode the presentation order of the evidence and distinguish evidence based on their different presentation positions.

How does UEcho distinguish two evidence items that do not come into play at the same time? Since in Echo every evidence item is associated with a Special Evidence Unit (SEU, which has a fixed activation value of 1), a simple decay mechanism is natural and sufficient to denote the different presentation time. That is, while the most recent evidence is given a full connection (designated by *DATA_EXCITATION* or DE) to the SEU, earlier evidence is given a connection to SEU with a decayed weight. The rate of decay depends upon both d, a new parameter, and the time interval since the evidence was presented. In current implementation, DE decays exponentially, that is, given the time interval since the evidence was presented is b, DE is determined by:

$$DE_{t+1} = DE_t \times d^b$$

When d is equal to 1, UEcho will operate similarly as in the original EoS manner. When d is equal to 0, UEcho will completely ignore previous evidence and act in a memoryless manner.

An Empirical Study

Zhang, Johnson, and Wang (1996) have reported part of the experimental work discussed here. Though new results have been obtained, the experimental paradigm is the same. For completeness, we briefly introduce the experimental task (for details, see Zhang, Johnson, & Wang, 1996), then report the combined results.

The experiment was implemented on the CIC (Combat Information Center) simulator developed by Towne (1995). The task is to decide whether a contact, which is seen on the ship's radar area, is friendly or hostile. To make the decision, subjects need to collect relevant evidence, such as speed, altitude, route, and verbal (radio) identification.

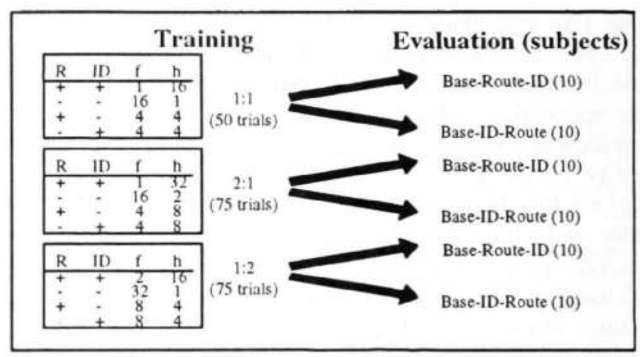

Figure 2: The design and procedure used in the experiment.

The experimental design is illustrated in Figure 2. Two factors are controlled in a 3x2 between-subjects design. The first factor is friendly-to-hostile base rates, which can be 1:1, 2:1, or 1:2. The second factor is the order in which evidence was presented in the questionnaire, either R- followed by ID+, or ID+ followed by R-. In the text and diagrams, R+ and R- correspond to on or off a commercial air route, and ID+ and ID- correspond to the response given to a radio request for identification (either response or no response claiming to be a commercial plane).

The experiment proceeds in two stages. First, subjects acquire background belief information by natural sampling. More specifically, subjects perform the task in many trials. In each trial, subjects see a different contact. They collect route (which is either R+ or R-) and identification (which is either ID+ or ID-) information about the contact then decide whether they think the contact is friendly or hostile. They then receive feedback about the true identity of the contact, which ends the trial and immediately begins the next trial.

After they finish all the training trials, they enter the second stage. In this stage, they fill out a questionnaire that presents some evidence sequentially about a contact and requires belief evaluation after each new piece of evidence is revealed. The same design and procedure were used to test the UEcho model.

Results and Discussion

The results are reported separately for belief acquisition (training) and belief revision (questionnaire). However, empirical results are combined together with the UEcho modeling results for easy comparison.

Training Performance The binary decisions of all training trials by each subject were transformed into conditional probabilities, which were then averaged across the 20 subjects in each base rate group. The results are shown in Figure 3, together with the corresponding Bayesian values (calculated from training trial distributions) and UEcho training results.

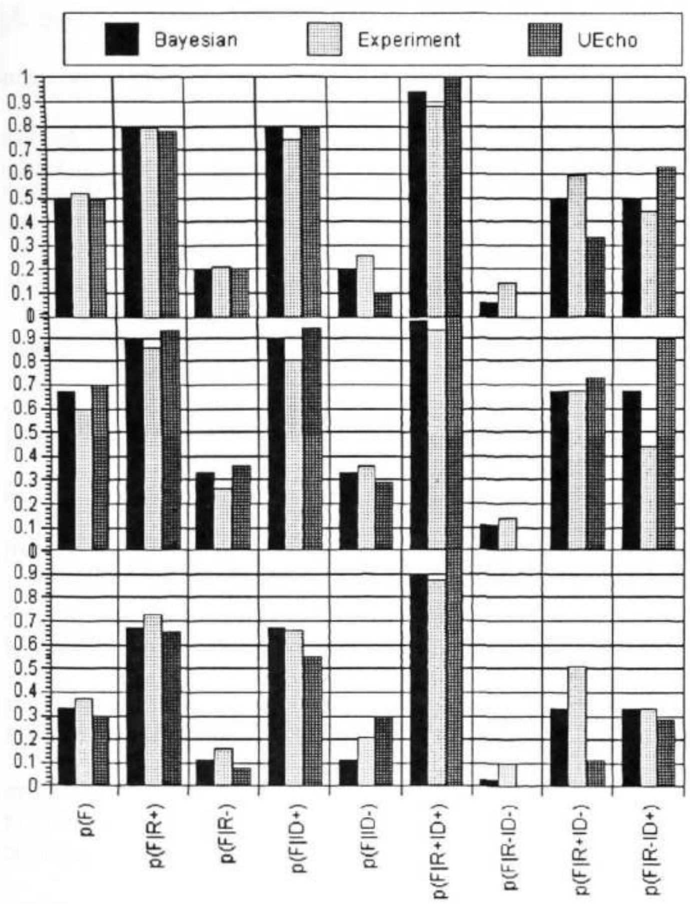

Figure 3: Training performance for 1:1, 2:1, and 1:2 conditions (top-down).

Statistics show that subjects can correctly acquire most of the probabilities. In addition, UEcho does reasonably well capturing the statistical structure in the training trials. One obvious deviation of UEcho from both the Bayesian values and the empirical results is UEcho's relatively extreme views under the R+ID+ condition, where UEcho suggests p(F)=1, and the R-ID- condition, where UEcho suggests p(F)=0. Lack of noise in the input may be one of the reasons for this.

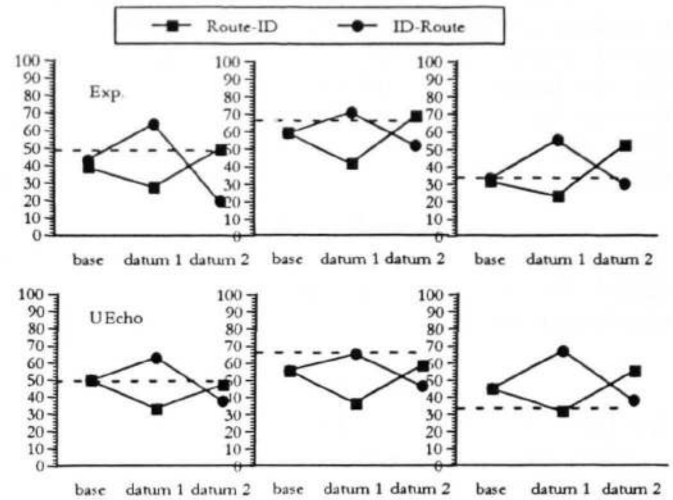

Figure 4: Belief evaluation results for 1:1, 2:1, and 1:2 conditions (left-right). The dashed lines indicate the Bayesian values of the base rate and final belief.

Belief Evaluation The results of belief evaluations after the training phase are shown in Figure 4. First, both subjects and UEcho show fairly accurate base rate judgments, which suggests that they can correctly assess and report the base rate information as long as the information is presented in the natural sampling manner. Second, despite accurate base rate information, fish-tail figures suggest that the order effect occurs under all conditions. More specifically, both UEcho and the experiment show a recency effect: the final decision is more determined by the direction of the last evidence item.

Conclusions and Future Plans

Abduction is a distinct type of hypothetical reasoning which infers something *may-be*. The empirical work here shows that human abductive reasoning indeed involves uncertainty and thus requires belief operations. Echo, a proposed model of abduction, accounts for many aspects of human abductive reasoning, but does not incorporate uncertainty management. We described a modified version of Echo (UEcho), which has the potential to handle the uncertainty aspects of abduction. It was shown that UEcho does a fairly good job modeling belief acquisition and dynamic belief revision, two critical components in human abductive reasoning.

As far as managing uncertainty in abduction is concerned, UEcho has advantages and disadvantages. On the one hand, the learning and dynamic mechanisms in UEcho allow it to adapt to its environment and learn from its experience. This is preferable compared to some other uncertainty management models, such as Bayesian belief networks and graphical belief networks, which usually do not learn. The adaptive and parallel satisfaction nature of UEcho also makes it capable of modeling human heuristics and biases, thus it is more preferable than normative models. On the other hand, it has been mentioned that Echo, and also UEcho, can only perform hypothesis evaluation and selection but not hypothesis formation, which makes it incomplete as an abduction model. It would be interesting to see how a complex hypothesis can be formed by merging several strongly connected simple hypotheses. Moreover, at present, UEcho does not handle higher-order uncertainty nor the distinction between belief and plausibility in the D-S theory sense. These disadvantages suggest that further investigations are needed.

Acknowledgements

This work is funded by Office of Naval Research Grant No. N00014-95-1-0241 and a summer fellowship from the Center for Cognitive Science at The Ohio State University.

References

Bar-Hillel, M. (1980). The base rate fallacy in probability judgments. *Acta Psychologica, 44*, 211-233.

Barto, A.G. (1990). Connectionist learning for control: An overview. In W.T. Miller, III, R.S. Sutton, & P.J. Werbos (Eds.), *Neural networks for control.* Cambridge, MT: MIT Press.

Cosmides, L, & Tooby, J. (1996). Are human good intuitive statisticans after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition, 58*, 1-73.

Dubois, D. & Prade, A. (1988). *Possibility Theory: An Approach to Computerized Processing of Uncertainty.* New York, NY: Plenum Press.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal Representations of human judgment.* New York: Wiley.

Fann, K.T. (1970). *Peirce's Theory of Abduction.* The Hague, Martinus Nijhoff.

Feltovich, P. J., Johnson, P. E., Moller, J. H. & Swanson, D. B. (1984). LCS: The role and development of medivcal knowledge in diagnostic expertise. In W. J. Clancey & E. H. Shortliffe (Eds.), *Readings in medical artificial intelligence.* Reading: Addison-Wesley.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency format. *Psychological Review, 102(4)*, 684-704.

Gluck, M.A., & Bower, G.H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General, 117(3)*, 227-247.

Hasher, L., & Zacks, R.T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General, 108(3)*, 356-388.

Hasher, L., & Zacks, R.T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologists, 39(12)*, 1372-1388.

Hogarth, R.M. & Einhorn, H.J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology, 24*, 1-55.

Josephson, J.R., & Josephson, S.G. (1994). *Abductive Inference: Computation, Philosophy, Technology.* Cambridge, NY: Cambridge University Press.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and Biases.* Cambridge, NY: Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 273-251.

Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition, 80*, 273-251.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review, 95*(2), 163-182.

Kleiter, G.D. (1993). Natural sampling: Rationality without base rates. In G.H. Fischer & D. Laming (Eds.), *Contributions to Mathematical Psychology, Psychometrics, and Methodology.* New York: Springer-Verlag.

Medin, D.L., & Edelson, S.M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General, 117(1)*, 68-85.

Miller, R.R., Barnet, R.C., & Grahame, N.J. (1995). Assessment of the Rescorla-Wagner model. *Psychological Bulletin, 117(3)*, 363-386.

Paul, G. (1993). Approaches to abductive reasoning: An overview. *Artificial Intelligence Review, 7*, 109-152.

Robert, C. P. (1994). The Bayesian choice: a decision-theoretic motivation. New York: Springer-Verlag.

Schlottmann A., & Anderson, N. H. (1995). Belief revision in children: Serial judgment in social cognition and decision-making domains. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21(5)*, 1349-1364.

Schmitt, F. (1992). *Knowledge and belief.* New York, NY: Routledge.

Shafer, G. (1976). *A mathematical theory of evidence.* Princeton, NJ: Princeton University Press.

Sutton, R.S., & Bargo, A.G. (1981). Toward a modern theory of adaptive networks: Expectation and Prediction. *Psychological Review, 88(2)*, 135-170.

Thagard, P. (1989). Explanatory Coherence. *Behavioral and Brain Sciences, 12(3)*, 435-502.

Thagard, P. (1992a). *Conceptual revolutions.* Princeton, NJ: Princeton University Press.

Thagard, P. (1992b). Adversarial problem solving: Modeling an opponent using explanatory coherence. *Cognitive Science, 16(1)*, 123-149.

Towne, D. (1995). *CIC: Tactical Decision Making (Version 2.0).* Unpublished manuscript. Behavioral Technology Laboratories, University of Southern California.

Wang, P. (1993). Belief revision in probability theory. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence* (pp. 519—526). San Mateo, CA: Morgan Kaufmann Publishers.

Zhang, J., Johnson, T.R., & Wang, H. (1996). Order effects and frequency learning in belief updating. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.