Fuzzy Sets and Systems 18 (1986) 315–328
North-Holland

INDUCTION OF INFERENCE RULES FOR EXPERT SYSTEMS

Brian R. GAINES and Mildred L.G. SHAW
Department of Computer Science, University of Calgary, Calgary, Alberta, Canada T2N 1N4

Received February 1985
Revised March 1985

Many of the decision-making processes of experts may be represented in terms of the fuzzy sets that correspond to their analysis of situations in terms of entities and their relationships. The elicitation of these fuzzy sets from experts can be carried out automatically by computer-based interview techniques. From this data the inference relations used in decision may be derived and encoded for an expert system shell. This paper analyses logical cluster-analytic techniques for inducing these inference relationships and gives examples of their applications to elicited data. The relations induced involve trade-offs between fuzzy truth values and the probability of their being correct. An information-theoretic measure of the uncertainty reduction due to a hypothesized relation is used to determine the optimum trade-off and to compare hypotheses.

Keywords: Knowledge engineering, Entail, Repertory grid data, Probabilistic hypotheses, Inductive data, Fuzzy logic, Expert systems.

In dedication to Richard Bellman, his ideas, ideals and inspiration.

1. Introduction

Richard Bellman was a remarkable person in the sheer span and vitality of his writings, illuminating old and new topics alike with fresh viewpoints, so surprising at first and yet so clear and inevitable in retrospect. Those who knew him for dynamic programming and invariant embedding must have been very surprised if they came across one of his papers on 'humor and paradox' [3], 'the evolution of consciousness' [4], 'psychotherapeutic interviews' [9] or 'man, society and the computer' [5]. For Bellman mathematics was not just formalism but:

> "the study of conceptual structures, of their interactions, and of their transformations over time." "My belief is that mathematics affords a new methodology in all of the human sciences." [6]

Bellman's wide-ranging interests, deep mathematical and systemic skills, and willingness to attempt to analyse significant engineering and psychological concepts even if the mathematical tools were inadequate, paralleled similar traits and capabilities in Lotfi Zadeh, and Bellman immediately recognized the importance of the concept of fuzzy sets and its wide-ranging applicability. His

0165-0114/86/$3.50 © 1986, Elsevier Science Publishers B.V. (North-Holland)

1970 paper with Zadeh on 'decision-making in a fuzzy environment' was the foundation of studies of fuzzy decision systems [10]. His 1973 paper with Giertz formally established the rules of fuzzy logic commencing with intuitively sound constraints, and is still widely quoted for its main results [8].

Bellman was an outstanding mathematician and formalist, yet he had an acute sense of wonder at the capabilities of the human mind to solve problems beyond the domain of current mathematics. He remarks in a paper on communication and ambiguity:

> "What is ironic and what saves us ultimately, and allows us to make vital decisions, is not the conscious use of superior logic and sophisticated mathematics, but rather the proto-logic of an animal mind developed by evolution." [7]

He saw the computer as an important tool to aid human decision-making through simulation:

> "It is essential to note that the computer serves in all this as a 'logic machine' exploring the consequences of data and hypotheses of the sort that can be furnished by experts. Different experts, of course, mean different predictions." [6]

He goes on to note how the system can furnish a means of communication between the partisans of differing theories, and a means for an expert to communicate with himself and preserve and develop his ideas.

Bellman can be said to have forseen the development of *expert systems* (ESs) and to have foretold of their role, significance and power. His vision of 1973 was substantiated in the operation of MYCIN [43] and TEIRESIAS [12], and his views on the importance of such inferential 'logic machines' has been vindicated by the widespread scientific and commercial success of ESs [29]. He saw the significance of the intrinsic ambiguity of human reasoning, not as a defect, but as an essential feature of any system coping with a complex world. He also saw the importance of the computer as an aid to this process:

> "to make a large and complex society feasible we need combinations of man and machine in various forms" [7].

In dedicating this paper to the memory of Richard Bellman we are pleased to take this opportunity to report new results on knowledge engineering techniques for expert systems, on the transfer of inferential knowledge from the person to the computer. The results are particularly apt because they bring together fuzzy data and fuzzy, probabilistic and information-theoretic analyses, to derive linguistic inferences that can be used to prime ES shells. The methodology described brings mathematical precision to what have previously been successfully applied, but heuristic, techniques. It does this while retaining the exact inexactitude of the original data. This is mathematics in the role of which Bellman was the supreme master, to represent human reality and human reasoning.

2. Knowledge engineering

Examples of practical ESs and the availability of domain-independent ES shells has generated widespread interest in ES applications. However, what Feigen-

baum terms *knowledge engineering* [13], the reduction of a large body of knowledge to a precise set of facts and rules, has become a major bottleneck impeding the application of ESs in new domains. Hayes-Roth, Waterman and Lenat outline the communication problems impeding the process of transferring expertise into a program [25]. They note that the vocabulary initially used by the expert to talk about the domain with a novice is often inadequate for problem-solving and the knowledge engineer and expert must work together to extend and refine it. They emphasize that one of the most difficult aspects of the knowledge engineer's task is helping the expert to structure the domain knowledge, to identify and formalize the domain concepts.

The problem of knowledge elicitation from a skilled person is well-known in the literature of psychology. Bainbridge has reviewed the difficulties of verbal debriefing and notes that there is no necessary correlation between verbal reports and mental behavior, and that many psychologists feel strongly that verbal data are useless [2]. However, this remark must be taken in the context of experimental psychologists working within a positivist, behavioral paradigm. Other schools of psychology have developed techniques for making use of verbal interaction, for example through interviewing techniques that attempt to by-pass cognitive defenses, including those resulting from automization of skilled behavior.

The results reported here are part of a study of interactive computer aids to the knowledge engineer [21–24, 33, 39–41] based on Kelly's *Personal Construct Psychology* (PCP) [26]. PCP develops a model of the person based on distinction-making through personal constructs that has strong systemic and psychological foundations [19, 27, 36]. The knowledge engineering tools resulting are techniques for eliciting the expert's vocabulary and structuring some aspects of the domain knowledge. Shaw and Gaines have reported experiments on eliciting constructs related to expertise in business record-keeping [39]. Boose has reported success in a wide range of ES developments at Boeing using these techniques to set up an initial system very much more rapidly than is possible with existing interview techniques [11].

3. Personal construct psychology

Kelly developed a theory of human cognition based on the primitive notion of constructs, or dichotomous distinctions, which are 'templets' that a person "creates and then attempts to fit over the realities of which the world is composed" [26]. He proposes that all of human activity can be seen as a process of anticipating the future by construing the replication of events, and revising constructs according to the success of this anticipation. Kelly developed his theory in the context of clinical psychology and hence was concerned to have techniques which used it to by-pass cognitive defenses and elicit the construct systems underlying behavior. In a non-clinical context this is the problem of knowledge engineering discussed above.

Kelly's *repertory grid* is a way of representing personal constructs as a set of distinctions made about elements relevant to the problem domain. In clinical

Table 1. Repertory grid from a manager on staff appraisal

Title: ICI-P130
Elements 10, constructs 14, range 1–5
Purpose: staff appraisal

	1	2	3	4	5	6	7	8	9	10	
intelligent 1	1	1	4	5	3	3	5	2	3	5	1 dim
willing 2	1	2	4	5	1	1	4	3	1	2	2 unwilling
new boy 3	1	2	3	5	4	4	4	1	4	3	3 old sweats
little supervision reqd 4	3	1	4	5	2	1	5	2	2	3	4 need supervision
motivated 5	1	1	4	5	2	2	5	3	3	2	5 less motivated
reliable 6	3	2	2	5	1	1	5	1	2	3	6 not so reliable
mild 7	3	4	5	2	2	3	1	5	4	5	7 abrasive
ideas men 8	1	1	5	4	2	3	1	3	4	4	8 staid
self starters 9	2	1	5	5	1	3	5	3	4	5	9 need a push
creative 10	1	1	5	5	2	3	4	3	4	5	10 non-creative
helpful 11	4	3	4	2	3	5	1	4	5	5	11 unhelpful
professional 12	1	2	3	3	2	1	5	2	4	4	12 less professional
overall rating high 13	2	1	3	4	1	2	5	2	3	4	13 overall rating low
messers 14	2	2	5	4	3	5	1	5	3	1	14 tidy
	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	

psychology this domain will often be personal relationships and the elements may be family members and friends. In the development of expert systems the elements will be key entities in the problem domain such as oil-well sites or business transactions. Repertory grids have been widely used: in clinical psychology [31, 42]; to study processes of knowledge acquisition in education [28]; and to study decision making by individuals and groups in management [30].

Table 1 shows a repertory grid elicited from a manager on the topic of 'personnel appraisal' [30]. The *elements* he is construing are his subordinates and are shown as the columns of a matrix. The *constructs* he is using to make bipolar distinctions between these elements are shown as the columns of the matrix. A rating scale from one through five is being used to enable him to make more precise distinctions in situations where he does not wish to allocate an element definitely to one pole of a construct. A rating of 1 means at the left-hand pole (LHP), and of 5 at the right-hand pole (RHP).

The importance of the repertory grid in clinical psychlogy is that it can be elicited without the client considering the overall grid and the relations between the constructs. Discussion between the therapist and client concentrates on the elements, which are generally actual people or situations, and the ways in which they are distinguished. It is only after the grid has been obtained that the therapist analyses it for relations between constructs of which the client may be unaware and which may be the source of problems.

Similar considerations apply to knowledge engineering. The manager from whom the data of Table 1 was elicited had only to consider people well-known to him and the similarities and differences between them. In doing this he has specified a range of situations and a range of constructs which he applies to them.

He has made overt much of his vocabulary in thinking about these people and exemplified its application. The knowledge engineer can use the grid to come to an understanding of the vocabulary and constructs involved.

It is particularly significant that the processes of repertory grid elicitation and analysis can be automated through interactive computer programs. The expert can work independently of the knowledge engineer for a substantial part of the process. This is clearly important in minimizing the requirement for scarce skills and resources. However, it has other significant advantages also:

(1) It allows the expert to experiment with different approaches to his subject area without any inter-personal pressures. People are often reluctant to discuss the basis of their expertise, not because it is proprietary, but because they do not have an explicit model of it. Experts may need the opportunity to try to create such a model without feeling that they are wasting someone's time or likely to appear foolish.

(2) The on-line computer elicitation of elements and constructs can be very much more effective than that by a person because the computer system can analyse the grid as it is being developed and feed back information to aid elicitation. They key to effective knowledge engineering is to span the relevant domain by encouraging the expert to generate elements and constructs that lead to relevant distinctions. On-line analysis is necessary for this as it is beyond the capabilities of most people.

PLANET is an integrated range of programs which operationalizes Kelly's work and may be used for the interactive elicitation and analysis of repertory grids [32]. In this paper we are not concerned with the elicitation of grids but rather their analysis to derive the inter-relations between the constructs used. The next section shows how the data from an expert, or a group of experts, may be analysed.

4. Logical analysis of repertory grid data

It is conventional to analyse a repertory grid as if it were a set of vectors of numeric data representing the position of each element in the space of constructs. This leads naturally to distance-based cluster-analytic approaches such as principle components analysis, for example Slater's INGRID [44], and hierachical sorting methods, for example Shaw's FOCUS [30]. However, while distance-based cluster analyses can portray two-sided relations between constructs, that construct A is similar to construct B, they cannot derive one-sided inferential relations, that construct A entails construct B. We have previously suggested a logical analysis of construct entailment based on an interpretation of repertory grids as fuzzy predicates [34, 35].

If element e has a rating of r on construct c where the range of ratings is 1 through n, then we say that e has a degree of membership of $(r-1)/(n-1)$ to the right-hand pole (RHP) of construct c, and a degree of membership $(n-r)/(n-1)$ to the left-hand pole (LHP) of construct c. We may regard each pole as representing a logical predicate that defines a fuzzy set and the repertory grid as a

set of such fuzzy sets, one for each predicate. Following the conventions of Gaines [14] we shall use ordinary characters for fuzzy sets, boldface characters for fuzzy variables, and denote the degree of membership of a variable to a set by concatenating it to the right of the set. Thus, x is the fuzzy set defined by predicate x (a LHP or RHP of the repertory grid) and the degree of membership of element $\boldsymbol{e}$ to fuzzy set x is $x\boldsymbol{e}$.

If we take u to be the predicate defined by LHP 9 (L9) in Table 1 and v to be that defined by LHP 13 (L13) then the fuzzy sets defined by the grid are

$$u = (0.75, 1.00, 0.00, 0.00, 1.00, 0.50, 1.00, 0.50, 0.25, 0.00), \tag{1}$$

$$v = (0.75, 1.00, 0.50, 0.25, 1.00, 0.75, 0.00, 0.75, 0.50, 0.25). \tag{2}$$

One distance-based measure of the similarity of two predicates, x and y, is

$$\Box(x \equiv y) = \sum_{\boldsymbol{e}} (x\boldsymbol{e} - y\boldsymbol{e})^2/n \tag{3}$$

where n is the cardinality of the set of elements,

$$n = \sum_{\boldsymbol{e}} 1. \tag{4}$$

This is the covariance measure that is used in the principal-components analysis of repertory grids [44]. Another distance-based measure is

$$\Box(x \equiv y) = \sum_{\boldsymbol{e}} |x\boldsymbol{e} - y\boldsymbol{e}|/n. \tag{5}$$

This is used by Shaw [30] for hierarchical clustering and may be regarded as a fuzzy measure of the distance between two predicates using the sigma-count of a fuzzy set [45].

We have used the symbol for logical necessity, $\Box$, in (3) and (5) to indicate that the measures defined may be regarded as estimators of the necessity of the predicates p and q being equivalent in a fuzzy modal logic. Both the measures defined have proved effective in distance-based clustering of repertory grid data. However, the semantics of both measures are difficult to define in terms of fuzzy logic and linguistic reasoning. The measure of (3) has no known semantics in fuzzy logic and that of (5) involves complex iterated implications [18]. We have suggested an alternative measure of equivalence involving only simple operations in fuzzy logic [34, 35, 37, 38]:

$$\Box(x \equiv y) = \bigwedge_{\boldsymbol{e}} (x\boldsymbol{e} \equiv y\boldsymbol{e}), \tag{6}$$

which, if we treat the logical connectives as those of the Łukasiewicz L_1 form of fuzzy logic, becomes

$$\Box(x \equiv y) = \underset{\boldsymbol{e}}{\text{MIN}} \min(1 - x\boldsymbol{e} + y\boldsymbol{e},\ 1 - y\boldsymbol{e} + x\boldsymbol{e}). \tag{7}$$

This logical formulation is particularly interesting because it extends from the symmetrical equivalence relation used in similarity clustering to the asymmetric entailment relation used for inference in ESs:

$$\Box(x \supset y) = \bigwedge_{\boldsymbol{e}} (x\boldsymbol{e} \supset y\boldsymbol{e}), \tag{8}$$

which for the Łukasiewicz L_1 form of fuzzy logic, becomes

$$\Box(x \equiv y) = \underset{e}{\text{MIN}} \min(1, 1 - xe + ye). \tag{9}$$

We may apply the usual symbols for necessary implication, or entailment:

$$x \rightarrow y \equiv \Box(x \supset y), \qquad x \leftrightarrow y \equiv \Box(x \equiv y). \tag{10}$$

The logical entailment relations of (6) through (10) have been incorporated in the repertory grid analysis program ENTAIL and shown to be similar in clustering to the conventional similarity-based programs while giving additional data about one-sided entailments [20]. Boose has used measures derived from these in his Expertise Transfer System (ETS) and shown it to give rapid prototyping of expert systems for a wide range of applications [11].

However, ENTAIL as originally presented derives all entailments that are *consistent* with the data in the grid, that is any *possible* entailment that is not disconfirmed by the data. This presents two problems:

(a) that no estimate is given of the significance of these entailments, the degree to which the data confirms them;

(b) that there may be few entailments that are totally true yet large numbers that are partially true and the theory so far gives no way of selecting those that are reasonable.

These are partly problems of induction from finite data, partly problems of the material conditional which is confirmed in trivial cases ($p \supset q$ is true if p is false), and partly problems of the logical predicates defined so far not capturing all that may be reasonably said about the data.

In the next section we tackle the first two problems of estimating the significance of entailments. In the following section we enrich the set of predicates used to include probabilistic assertions. Together, these give a formal theory of the induction of entailments from repertory grid data.

5. Significance analysis of repertory grid data

The significance of an assertion may be measured in terms of the reduction of uncertainty that results from that assertion. If we consider the fuzzy set

$$w = (0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00), \tag{11}$$

then the assertions $w \rightarrow v$ and $u \rightarrow v$ both have truth value 1 from (9) yet for the first knowing we gives no reduction in the uncertainty of ve whereas for the second knowing ue gives a great deal of reduction.

We may measure this reduction in uncertainty by defining a measure of the uncertainty about an individual event and then deriving the average uncertainty over all events. Gaines defines the *surprise* at an event to be minus the logarithm of the probability ascribed to that event before it occurred, and shows how surprise minimization may be used to formalize inductive inference [16]. We may apply this to the fuzzy sets defining predicates by examining the distribution of

degrees of membership, the probabilities that a given degree of membership will occur. The distribution, ∇x, of the fuzzy set x may itself be regarded as a fuzzy set of degrees of membership to these probabilities and written

$$\nabla x = \{m \mid p(m) : \exists \boldsymbol{e}, m = x\boldsymbol{e}\} \tag{12}$$

where $p(m)$ is the probability (relative frequency) of the degree of membership m in the fuzzy set x and '|' is used as a separator as has become conventional for fuzzy sets. For the examples given we have

$$\nabla u = \{1.00 \mid 0.3, 0.75 \mid 0.1, 0.50 \mid 0.2, 0.25 \mid 0.10, 0.00 \mid 0.30\}, \tag{13}$$

$$\nabla v = \{1.00 \mid 0.2, 0.75 \mid 0.3, 0.50 \mid 0.2, 0.25 \mid 0.20, 0.00 \mid 0.10\}, \tag{14}$$

$$\nabla w = \{0.00 \mid 1.00\}. \tag{15}$$

The expected surprise for a distribution ∇x is

$$\alpha(x) = \sum_{\boldsymbol{e}} -\log_2(p(x\boldsymbol{e}))/n \tag{16}$$

which is just the Shannon entropy for the distribution.

We may use this formulation to define the conditional surprise as the expected surprise for the distribution given another distribution and a hypothesis about their relationship:

$$\alpha(x \mid y \rightarrow x) = \sum_{\boldsymbol{e}} -\log_2(p(x\boldsymbol{e} \mid y\boldsymbol{e}, y \rightarrow x))/n \tag{17}$$

where $p(x\boldsymbol{e} \mid y\boldsymbol{e}, y \rightarrow x)$ is the probability of $x\boldsymbol{e}$ having its degree of membership given the degree of membership of $y\boldsymbol{e}$ and the truth value of the assertion $y \rightarrow x$. We may interpret this assertion as a constraint on the relation between $x\boldsymbol{e}$ and $y\boldsymbol{e}$ and enables the distribution of $x\boldsymbol{e}$ to be restricted when the value of $y\boldsymbol{e}$ is given. For example, given that $u \rightarrow v = 1$, and given the first value of $u\boldsymbol{e}$ in (1) as 0.75 we can infer that the conditional distribution for the first value of $v\boldsymbol{e}$ in (2) is $\{1.00 \mid 0.4, 0.75 \mid 0.6\}$ which has a greatly reduced uncertainty from the unconditional distribution of (14). The conditional probability is $p(v\boldsymbol{e} \mid u\boldsymbol{e}, u \rightarrow v) = 0.6$ compared with the unconditional probability $p(v\boldsymbol{e}) = 0.3$.

We may now define the uncertainty reduction due to a hypothesis as the difference between the unconditional and conditional expected surprise given the hypothesis

$$\rho(y \rightarrow x) = \alpha(x) - \alpha(x \mid y \rightarrow x). \tag{18}$$

This uncertainty reduction can be used to rank the entailments induced by ENTAIL. Clearly such trivial entailments as $w \rightarrow v$ have zero uncertainty reduction, as do such equally trivial entailments as $v \rightarrow v$ and $v \rightarrow w$, even though all three have truth value 1. The entailment $v \rightarrow u$ is not trivial but its low truth value of 0.5 leads to it providing little constraint and hence a low value of uncertainty reduction.

Table 2 in an ENTAIL analysis of the grid of Table 1 with the truth values and uncertainty reductions (termed 'information') shown together with the entailments expressed linguistically. The column marked 'Prob.' will be discussed in the

Table 2. Entailment analysis of repertory grid on staff appraisal

Entail	Truth	Prob.	Information (Cutoff 35)	
L9→L13	100	100	59	self starters→overall rating high
L1→L9	75	100	58	intelligent→self starters
L10→L8	100	100	57	creative→ideas men
R9→R1	75	100	56	need a push→dim
L1→L10	75	100	52	intelligent→creative
R8→R10	100	100	52	staid→ non-creative
R13→R6	75	100	52	overall rating low→not so reliable
L10→L9	75	100	51	creative→ self starters
R10→R1	75	100	49	non-creative→ dim
L9→L10	75	100	47	self starters→creative
L10≡L1	75	100	47	creative≡intelligent
R1≡R10	75	100	47	dim≡non-creative
R9≡R10	75	100	47	need a push≡non-creative
R13→R9	100	100	46	overall rating low→need a push
R4→R13	75	100	44	need supervision→overall rating low
R5→R4	75	100	44	less motivated→need supervision
R5→R13	75	100	44	less motivated→overall rating low
L1→L3	75	100	42	intelligent→new boy
L6→L13	75	100	42	reliable→overall rating high
L4≡L13	75	100	39	little supervision reqd≡overall rating high
L13≡L6	75	100	39	overall rating high≡reliable
R1→R6	50	100	39	dim→not so reliable
R6≡R13	75	100	39	not so reliable≡overall rating low
R13≡R4	75	100	39	overall rating low≡need supervision
R12→R13	75	100	38	less professional→overall rating low
R13→R12	75	100	38	overall rating low→less professional
L8≡L10	100	70	37	ideas men≡creative
R9→R4	50	100	37	need a push→need supervision
R9→R12	50	100	37	need a push→ less professional
R9→R13	50	100	37	need a push→overall rating low
R10≡R8	100	70	37	non-creative≡staid
L4→L5	75	100	36	little supervision reqd→motivated
L4→L9	50	100	36	little supervision reqd→self starters
L12→L9	50	100	36	professional→self starters
L12≡L13	75	100	36	professional≡overall rating high
L13→L5	75	100	36	overall rating high→motivated
L13→L9	50	100	36	overall rating high→self staraters
R2→R4	75	100	35	unwilling→need supervision
R2→R5	75	100	35	unwilling→less motivated
R2→R13	75	100	35	unwilling→overall rating low
R3→R1	75	100	35	old sweats→dim

next section. It will be noted that ENTAIL checks for equivalence separately from the check for one-way entailment. This is because the uncertainty reduction from the hypothesis of equivalence, $x \equiv y$, cannot be inferred from the separate uncertainty reductions of the logically equivalent pair of hypotheses, $x \rightarrow y$ and $y \rightarrow x$, although it will be at least the sum of these. In order to rank-order the hypotheses the uncertainty reduction for one-sided entailment is doubled for comparison with that of two-sided equivalence.

Note how a hypothesis of high truth value such as R13 → R9 can give less information reduction than one of lower truth value such as L1 → L9. Note also how a one-sided hypothesis, such as L1 → L10, may give greater uncertainty reduction than its two-sided equivalent, L1 ≡ L10 and hence have a higher ranking. Similarly, the two one-sided hypotheses, such as L9 → L10 and L10 → L9 can also occur in preference to the two-sided hypotheses, L9 ≡ L10. The list of hypotheses in Table 2 has been terminated when the uncertainty reduction falls below 35% for a two-sided hypothesis. It can be truncated at any level and the entailments plotted in graphical form. However, since approximate inference is involved the graph will not necessarily be transitive and it is important to preserve the intransitivity. It avoids the pardoxes of approximate reasoning such as those of sorites and falakros [14].

The entailments derived by ENTAIL can be converted directly to rules that may be loaded into a standard ES shell as Boose has shown [11]. The truth value of the rule codes into the truth value represention used by the shell. The rules may be checked by the expert for their content but their prior evaluation for uncertainty reduction makes it unlikely that spurious rules will be generated. The rules form a decision tree expressing the relationship between an expert's data and his conclusions, both personally expressed. For example, the data of Table 1 may be regarded as that of an expert on staff appraisal and one may be concerned with deriving his 'overall rating' (construct 13) from behavioral assessments such as 'intelligent' and 'creative'. The ENTAIL analysis of Table 2 shows that L1, L4, L6, L9, L10 and L12 imply L13, that "intelligent, creative, reliable and professional self-starters requiring little supervision receive a high overall rating", whereas R2, R4, R5, R6, R9 and R12 imply R13, that "being unwilling, less motivated, not so reliable, less professional, needing supervision and needing a push leads to a low overall rating."

6. Probabilistic hypotheses

In our initial analysis of entailment we introduced the notion of entailments being true 'usually' and interpreted this as 'in all except one instance'. The original version of ENTAIL generates entailments that are usually true and these have proved useful with some data where there is a common pattern upset only by an isolated case. The concept of 'usually' can be generalized and formalized within the framework already established and it leads to an interesting two-dimensional analysis of fuzzy data.

We quantify the notion of a hypothesis being true 'usually' by ascribing a probability to the hypothesis measured by the relative frequency with which it is true. For example, using the data of (1) and (2), $u \rightarrow v$ has truth value 1 with probability 1; it is always absolutely true. However, $v \rightarrow u$ is more interesting because it has truth value 0.5 with probability 1, truth value 0.75 with probability 0.9 (fails in one case out of ten), and truth value 1.0 with probability 0.4 (fails in six cases out of ten). Thus evaluation of an entailment is now not a fuzzy truth value but rather a distribution as already defined, a fuzzy set of degrees of

membership of probabilities:

$$v \rightarrow u = \{1.00 \mid 0.4,\ 0.75 \mid 0.9,\ 0.50 \mid 1.0\}. \tag{19}$$

This type of trade-off distribution between the truth value of an assertion and its probability of being true occurs in inductive inference with fuzzy data, for example in inferring structure from fuzzy sequential data [17]. We can evaluate the possible trade-offs by using the uncertainty reduction analysis based on expected surprise already described. The calculation of the conditional probability of an event under a hypothesis takes into account probabilistic assertions by calculating the expected distribution as a mixture of the unconditional distribution with the probability that the hypothesis is false, and the conditional distribution with the probability that the hypothesis is true. It is interesting to note that the correct calculation has to be based on considerations of expected surprise not expected uncertainty; the notion of surprise is more fundamental than that of uncertainty.

As can be seen from Table 2 the preferred hypothesis for L13 → L9 is a truth value of 0.50 with probability 1. However, the hypothesis L8 ≡ L10 is preferred as having a truth value of 1 with probability 0.70. The distribution is:

$$\text{L8} \equiv \text{L10} = \{1.00 \mid 0.7,\ 0.75 \mid 0.9,\ 0.25 \mid 1.0\} \tag{20}$$

and the greatest uncertainty reduction is achieved with a hypothesis that is false in 3 cases.

It is also of interest to note that the notion of 'usually' with the probabilistic interpretation given here may be regarded as a generalization of Atkin's Q-Analysis [1] for fuzzy data [15]. We have previously compared fuzzy Q-Analysis with other distance-based clustering techniques for repertory grids and shown that it generates similar clusters [20]. The formulation of entailment analysis given here subsumes Q-Analysis as a special case of clustering based on two-sided logical equivalence but with probabilistic hypotheses.

7. Conclusions

A method has been given for inductive inference of entailments from repertory grid data allowing combined fuzzy logical and probabilistic hypotheses. The method provides a measure based on expected surprise as a means of estimating the significance of a hypothesis through its reduction of uncertainty. The results of the analysis are fuzzy linguistic inferences that may be interpreted directly to prime expert system shells. The overall methodology provides rigorous mathematical foundations for knowledge engineering based on personal construct psychology techniques.

We note in passing that the mathematical techniques presented in this paper have wider relevance than to repertory grid data. They may be regarded as a general way of inducing and representing structure from fuzzy data giving the degree of membership of abstract entities to certain attributes. They may also be

regarded as non-parametric techniques for the statistical analysis of such data. The fuzzy-truth-value versus probability-of-correctness trade-off seems a fundamental concept of inductive inference and data analysis. It has a heuristic relationship to distance measures that have proved useful for many purposes but do not have a natural interpretation except for parametric data. It is suggested that more general application of the techniques described in this paper will prove valuable.

Acknowledgements

Financial assistance for this work has been made available by the National Sciences and Engineering Research Council of Canada.

References

[1] R.H. Atkin, Mathematical Structure in Human Affairs (Heinemann, London, 1974).
[2] L. Bainbridge, Verbal reports as evidence of the process operator's knowledge, Internat. J. Man–Machine Stud. 11(4) (1979) 411–436.
[3] R. Bellman, Humor and paradox, in: W. Mendel, Ed., A Celebration of Laughter (Mara Books, Los Angeles, 1970) 33–45.
[4] R. Bellman, Hierarchies of control processes and the evolution of consciousness, in: Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability: Volume V (University of California Press, Berkeley, CA. 1971) 311–314.
[5] R. Bellman, Man, society and computer, in: Man and Computer (Karger, Basel, 1972) 229–245.
[6] R. Bellman, Mathematics and the human sciences, in: J. Wilkinson, R. Bellman and R. Garaudy, Eds., The Dynamic Programming of Human Systems: a Social and Historical Analysis (MSS Information Corporation, New York, 1973) 11–18.
[7] R. Bellman, Communication, ambiguity and understanding, Math. Biosci. 26 (1975) 247–356.
[8] R. Bellman and M. Giertz, On the analytic formalism of the theory of fuzzy sets, Inform. Sci. 5 (1973) 149–156.
[9] R. Bellman, C.P. Kell and W.C. Hopgood, Computer vignettes: a research tool for the study of the initial psychotherapeutic interview, J. Cybernet. 1(1) (1971) 19–27.
[10] R. Bellman and L.A. Zadeh, Decision-making in a fuzzy environment, Management Sci. 17(4) (1970) B 141–164.
[11] J.H. Boose, Personal construct theory and the transfer of human expertise, Proceedings AAAI-84 (American Association for Artificial Intelligence, California, 1984) 27–33.
[12] R. Davis and D.B. Lenat, Knowledge-Based Systems in Artificial Intelligence (McGraw-Hill, New York, 1982).
[13] E.A. Feigenbaum, Knowledge Engineering: the Applied Side of Artificial Intelligence, Report STAN-CS-80-812 (Department of Computer Science, Stanford University, CA, 1980).
[14] B.R. Gaines, Foundation of fuzzy reasoning, Internat. J. Man–Machine Stud. 8(6) (1976) 623–668.
[15] B.R. Gaines, V-fuzzy Q-analysis, EES-MMS-QFUZ-76, Department of Electrical Engineering Science, Essex University, Colchester (1976).
[16] B.R. Gaines, System identification, approximation and complexity, Internat. J. General Systems 3 (1977) 145–174.
[17] B.R. Gaines, Sequential fuzzy system identification, Fuzzy Sets and Systems 2(1) (1979) 15–24.
[18] B.R. Gaines, Precise past–fuzzy future, Internat. J. Man–Machine Stud. 19(1) (1983) 117–134.
[19] B.R. Gaines and M.L.G. Shaw, A programme for the development of a systems methodology of

knowledge and action, in: W.J. Reckmeyer, General Systems Research and Design: Precursors and Futures (Society for General Systems Research, 1981) 255–264.

[20] B.R. Gaines and M.L.G. Shaw, New directions in the analysis and interactive elicitation of personal construct systems, in: M.L.G. Shaw, Ed., Recent Advances in Personal Construct Technology (Academic Press, London, 1981) 147–182.

[21] B.R. Gaines and M.L.G. Shaw, Logical foundations of expert systems, Proceedings of IEEE International Conference on Systems, Man and Cybernetics, Halifax, Nova Scotia (1984) 238–247.

[22] B.R. Gaines and M.L.G. Shaw, Foundations of expert systems, in: J. Kacprzyk and R.R. Yager, Eds., Management Decision Support Systems Using Fuzzy Sets and Possibility Theory, ISR Interdisciplinary Systems Research (Verlag TÜV, Rheinland, 1984).

[23] B.R. Gaines and M.L.G. Shaw, Expert systems and simulation, in: G. Birtwistle, Ed., AI, Graphics and Simulation (Society for Simulation Research, La Jolla, CA, 1985) 95–101.

[24] B.R. Gaines and M.L.G. Shaw, Systemic foundations for reasoning in expert systems, in: M.M. Gupta, A. Kandel, W. Bandler and J.B. Kiszka, Eds., Approximate reasoning in Expert Systems (North-Holland, Amsterdam, 1985).

[25] F. Hayes-Roth, D.A. Waterman and D.B. Lenat, Editors, Building Expert Systems (Addison-Wesley, Reading, MA, 1983).

[26] G.A. Kelly, The Psychology of Personal Constructs (Norton, New York, 1955).

[27] J.R. Mancuso and J.R. Adams-Weber, Editors, The Construing Person (Praeger, New York, 1982).

[28] M.L. Pope and M.L.G. Shaw, Personal construct psychology in education and learning, in: M.L.G. Shaw, Ed., Recent Advances in Personal Construct Technology (Academic Press, London, 1981) 105–114.

[29] W. Reitman, Editor, Artificial Intelligence Applications for Business (Ablex, Norwood, NJ, 1984).

[30] M.L.G. Shaw, On Becoming a Personal Scientist (Academic Press, London, 1980).

[31] M.L.G. Shaw, Editor, Recent Advances in Personal Construct Technology (Academic Press, London, 1981).

[32] M.L.G. Shaw, PLANET: some experience in creating an integrated system for repertory grid applications on a microcomputer, Internat. J. Man–Machine Stud. 17(3) (1982) 345–360.

[33] M.L.G. Shaw, Interactive knowledge elicitation, in: Proceedings of CIPS SESSION 84 (Canadian Information Processing Society, Calgary, 1984) 202–208.

[34] M.L.G. Shaw and B.R. Gaines, Externalizing the personal world: computer aids to epistemology, in: R.F. Ericson, Ed., Improving the Human Condition: Quality and Stability in Social Systems (Society for General Systems Research, KY, 1979) 136–145.

[35] M.L.G. Shaw and B.R. Gaines, Fuzzy semantics for personal construing, in: B.H. Banathy, Ed., Systems Science and Science: Proceedings of the 24th Annual Meeting of SGSR (Society for General Systems Research, KY, 1980) 146–154.

[36] M.L.G. Shaw and B.R. Gaines, The personal scientist in the community of science, in: W.J. Reckmeyer, Ed., General Systems Research and Design: Precursors and Futures (Society for General Systems Research, 1981) 59–68.

[37] M.L.G. Shaw and B.R. Gaines, Fuzzy entailment analysis, in: G.E. Lasker, Ed., Applied systems and Cybernetics Volume VI (Pergamon Press, Oxford, 1981) 3010–3014.

[38] M.L.G. Shaw and B. R. Gaines, Eliciting entailment, in: R. Trappl, L. Ricciardi and G. Pask, Eds., Progress in Cybernetics and Systems Research (Hemisphere, Washington, DC, 1982) 425–435.

[39] M.L.G. Shaw and B.R. Gaines, A computer aid to knowledge engineering, Proceedings of British Computer Society Conference on Expert Systems (Cambridge, 1983) 263–271.

[40] M.L.G. Shaw and B.R. Gaines, Deriving the constructs underlying decision, in: H.-J. Zimmermann, L.A. Zadeh and B.R. Gaines, Eds., Fuzzy Sets and Decision Analysis, TIMS Series Studies in the Management Sciences 20 (North-Holland, Amsterdam, 1984) 335–355.

[41] M.L.G. Shaw and B.R. Gaines, Some experience in eliciting knowledge structures for expert systems, in: A.W. Smith, Ed., Proceedings of the Society for General Systems Research International Conference (Intersystems, California, 1984) 605–614.

[42] E. Shepherd and J.P. Watson, Editors, Personal Meanings (John Wiley, London, 1982).
[43] E.H. Shortliffe, Computer-Based Medical Consultations: MYCIN (Elsevier, New York, 1976).
[44] P. Slater, Editor, Dimensions of Intrapersonal Space: Volume 1 (John Wiley, London, 1976).
[45] L.A. Zadeh, The role of fuzzy logic in the management of uncertainty in expert systems, Fuzzy Sets and Systems 11(3) (1983) 199–227.