Imprecise belief updating for uncertain user modelling

J.M.Rossiter, T.H.Cao, T.P.Martin, and J.F.Baldwin
Department of Engineering Mathematics
University of Bristol
Bristol BS8 1TR, UK
Phone: +44-117-9289734
Fax: +44-117-9251154
Jonathan.Rossiter@bris.ac.uk

Abstract

In this paper we present a new approach to incremental user recognition in fuzzy environments where user classification is updated within an epistemological model. We extend Einhorn and Hogarth's anchor and adjustment method ([7]), itself derived from a study of human behaviour, from the point value representation of belief and evidence to the case where belief and evidence are imprecise. We represent imprecision by intervals of belief in the range $[0, 1]$.

Keywords: interval belief updating, user modelling, object oriented, uncertain class hierarchies, Fril++.

1 Introduction

In recent years user modelling has become a major topic of academic and commercial research. This focus has been driven by a combination of two factors; firstly, the construction of huge databases of information about our daily lives, and secondly, by the desire of organisations to use this data to understand the people they deal with and hence to improve their services.

The real world, and in particular human behaviour, is not often conducive to crisp study. Rather, human behaviour is more often about degrees of action than absolutes. For example, take the statement 'a television viewer spends most evenings watching low quality television'. In this case the words 'most' and 'low quality' capture, to some extent, the generalising qualities of this statement. These qualities push the study of user behaviour and user modelling in the direction of techniques that can capture real world uncertainties such as the concepts 'most' and 'low quality'.

In this paper we present a new approach to incremental user recognition in fuzzy environments where user classification is updated within an epistemological model. We extend Einhorn and Hogarth's anchor and adjustment method ([7]), itself derived from a study of human behaviour, from the point value representation of belief and evidence to the case where belief and evidence are imprecise. We represented imprecision by intervals of belief in the range $[0, 1]$.

2 A simple user recognition example

Let us take the example where we classify food consumers into one of the classes *candy-eater*, *cookie-eater* or *cake-eater*. We may wish to represent these consumer classes in a Fril++ ([2],[3],[4],[10],[5]) uncertain class hierarchy as shown in Figure 1. In the same way a hierarchy can be constructed for the food these consumers eat, as shown in Figure 2.

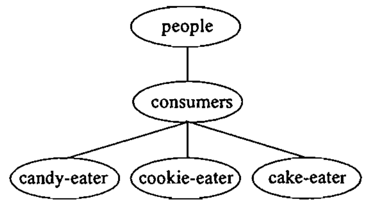

Figure 1: A consumer class hierarchy

The consumer classes also define the prototypical behaviours of these consumers through the following statements:

a *candy-eater* eats lots of candy most of the time
a *cookie-eater* eats lots of cookies most of the time
a *cake-eater* eats lots of cake most of the time

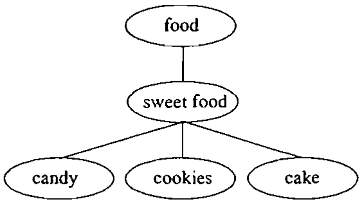

Figure 2: A food class hierarchy

In a simple representation we could use a conditional probability interval to represent the 'most' qualifier. For example, if we find that the statement 'eats lots of candy most of the time' is true for eight or more cases out of every ten candy-eaters, then we can assign an interval $[0.8, 1]$ to the conditional probability $Pr(eats|lots\ of\ candy)$. This approach gives us the following Fril++ class definition for the *candy-eater* class:

((**class** *candy-eater* **inherits from** *consumer*)
(**has property**
((eats X)(X.isa candy)(X.quantity *lots*)) : (0.8 1)))

where ((eats X)(X.isa candy)(X.quantity *lots*)) : (0.8 1) defines a Horn clause with (eats X) as the head and ((X.isa candy)(X.quantity *lots*)) as the body and $Pr(head|body)$ is defined by the interval [0.8, 1]. The conclusion (eats X), given some instance X, is calculated using Jeffrey's rule ([8]).

Now consider a new food consumer u who makes a decision whether or not to eat food x and we wish to determine u's membership in the classes *candy-eater*, *cookie-eater* and *cake-eater*. The only information we have is the decision that u made with respect to eating food x. We can determine memberships by comparing u's decision with the decision that would be made by a prototypical member of each of the classes *candy-eater*, *cookie-eater* and *cake-eater*, given food x. Food x itself may be an uncertain member of any or all of the classes *candy*, *cookie* and *cake*. For example, if x is a sweet iced biscuit then x is clearly a member of the class *cookie*, but may also have non-zero membership in the class *candy*. In a Fril++ environment this is an object recognition problem where the memberships are calculated using an inverted interval form of Jeffrey's rule as presented in [5]. Memberships are expressed as support intervals in the range $[0, 1]$.

The remainder of this paper is concerned with the case where we wish to update u's membership in the classes *candy-eater*, *cookie-eater* and *cake-eater* as u chooses whether or not to eat each item of food in the ordered stream $x_1, \ldots, x_n$.

3 Interval anchor and adjustment belief updating

We now consider Einhorn and Hogarth's anchor and adjustment method of updating belief in light of new evidence ([7]). Where this differs from probabilistic approaches such as Bayesian updating is in the epistemological background to this approach. Einhorn and Hogarth have done much work on trying to understand how humans update beliefs ([6], [7]). Their conclusions are that humans update their belief S in hypothesis h in a way that is greatly dependent on the degree to which they previously held a belief in h. That is, $S_k = f(S_{k-1}, s(x_k))$, where $s(x_k)$ is the support for h derived from new evidence x_k, and f is some belief updating function. Equations (1) and (2) define Einhorn and Hogarth's anchor and adjustment belief updating function for cases where belief S is a point value.

$$S_k = S_{k-1} + \alpha S_{k-1}(s(x_k) - R) \quad \text{for } s(x_k) \le R \quad (1)$$
$$S_k = S_{k-1} + \beta(1 - S_{k-1})(s(x_k) - R) \quad \text{for } s(x_k) > R \quad (2)$$

R is a reference point for determining if the evidence $s(x_k)$ is positive or negative, and typically $R = 0$ or $R = S_{k-1}$. α and β are constants which define how sensitive the model is to negative or positive evidence respectively. Notice that updating under negative evidence is in proportion to the previous positive belief (S_{k-1}) and updating under positive evidence is in proportion to the previous negative belief $(1 - S_{k-1})$.

New let us consider the case where current belief S is defined by an interval $[S^-, S^+]$, a subinterval of $[0, 1]$. We can simply substitute the lower and upper bounds S^- and S^+ into equations (1) and (2) to give equations (3) to (6).

$$S_k^- = S_{k-1}^- + \alpha S_{k-1}^-(s(x_k) - R) \quad \text{for } s(x_k) \le R \quad (3)$$
$$S_k^- = S_{k-1}^- + \beta(1 - S_{k-1}^-)(s(x_k) - R) \quad \text{for } s(x_k) > R \quad (4)$$
$$S_k^+ = S_{k-1}^+ + \alpha S_{k-1}^+(s(x_k) - R) \quad \text{for } s(x_k) \le R \quad (5)$$
$$S_k^+ = S_{k-1}^+ + \beta(1 - S_{k-1}^+)(s(x_k) - R) \quad \text{for } s(x_k) > R \quad (6)$$

In equation (5) negative belief updating is in proportion to S^+. However, for the interval belief $[S^-, S^+]$ we can only say for definite that positive belief for a hypothesis is S^-. Consequently we must update equation (5) in light of this definitive. Likewise we must update equation (4) in light of the definite negative belief $(1 - S^-)$. We should also consider when $s(x_k)$ and R are represented by intervals $[s^-(x_k),\ s^+(x_k)]$ and $[R^-,\ R^+]$ respectively. Updated versions of equations (3) to (6) are shown in equations (7) to (10) respectively.

$$S_k^- = S_{k-1}^- + \alpha S_{k-1}^-(s^-(x_k) - R^-) \quad \text{for } s^-(x_k) \le R^- \quad (7)$$
$$S_k^- = S_{k-1}^- + \beta(1 - S_{k-1}^+)(s^-(x_k) - R^-) \quad \text{for } s^-(x_k) > R^- \quad (8)$$
$$S_k^+ = S_{k-1}^+ + \alpha S_{k-1}^-(s^+(x_k) - R^+) \quad \text{for } s^+(x_k) \le R^+ \quad (9)$$
$$S_k^+ = S_{k-1}^+ + \beta(1 - S_{k-1}^+)(s^+(x_k) - R^+) \quad \text{for } s^+(x_k) > R^+ \quad (10)$$

R^- is a reference point for determining if the lower bound of belief derived from the presented evidence is positive or negative with respect to the lower bound of previous belief and R^+ is the corresponding reference point for the upper bound. Of course, the interval restriction $S^- \le S^+$ imposes extra constraints on the values of R^-, R^+, α, and β above those for the point value case.

3.1 Examples of interval updating

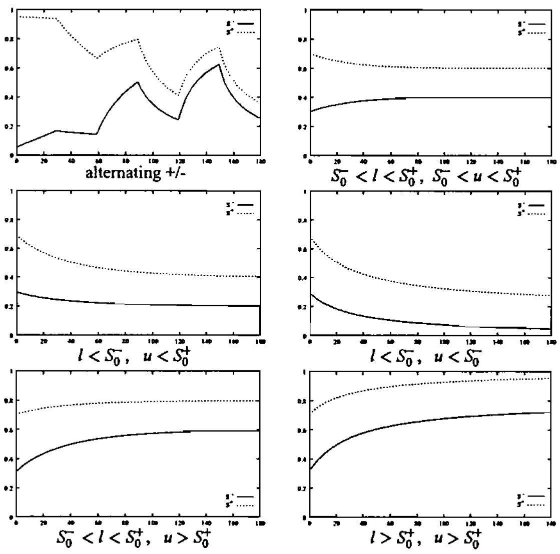

Figure 3: Belief updating example cases

Figure 3 shows the interval anchor adjustment belief updating approach for some simple cases. In these figures $[S_{k-1}^-,\ S_{k-1}^+]$ defines the previous belief in some hypothesis and $[l_k,\ u_k]$ defines the belief in the hypothesis derived from the latest piece of evidence. The first figure is especially interesting. It shows the imprecise belief updating approach when $[S_0^-,\ S_0^+] = [0.05,\ 0.95]$ and evidence is presented alternately in blocks of positive evidence followed by blocks of negative evidence, i.e. $[++++----++++\cdots]$, and where positive evidence is defined by the interval $[0.8,\ 0.9]$ and negative evidence is defined by the interval $[0.1,\ 0.2]$. Notice how the updated belief interval converges on the belief derived from the latest piece of evidence.

4 A user modelling example

Let us now consider a simple user modelling example using an iterated prisoner's dilemma testbed ([1], [9]). In this problem we have a population of 10 prisoners each of whom exhibits one of the 5 possible types of behaviour shown in Table 1. The prisoner population is shown in Table 2. At each turn pairs of prisoners are interviewed by the police. Each prisoner can either cooperate with his partner or defect. Differing combinations of cooperating and defecting between pairs of prisoners results in differing prison sentences. Using the techniques described in this paper we will attempt to classify prisoners into prisoner behaviour classes in much the same way as we would classify users into user classes.

Behaviour	Description
cooperative	always cooperate with opponent
uncooperative	always defect against opponent
tit-for-tat	cooperate unless last opponent defected
random	equal random chance of defect or cooperate
respd	defect unless the last 6 opponents chose to cooperate

Table 1: Behaviour classes

Individual	Behaviour before 60^{th} round	Behaviour after 60^{th} round
1	random	cooperative
2	random	uncooperative
3	cooperative	tit-for-tat
4	cooperative	respd
5	uncooperative	random
6	uncooperative	respd
7	tit-for-tat	cooperative
8	tit-for-tat	random
9	respd	tit-for-tat
10	respd	uncooperative

Table 2: The prisoner population

A population of ten prisoners was created and a game of 75 rounds was initiated. Each round involved picking pairs of prisoners at random from the population until non were left and for each pair recording the behaviours they exhibit (defect, cooperate, etc). From the past history of each player, and using the techniques described earlier ($\alpha = \beta = 0.3$), they were classified into the five behaviour classes. The winning class was taken as the class in which minimum membership (i.e. the lower bound of the membership interval) was greatest. If the winning class matched the actual class in Table 2 then the classification was recorded as a success.

To recreate the situation where user behaviour changes, after 60 rounds the behaviours of all 10 prisoners was changed, as shown in the third column of Table 2. After this point the game was continued for 15 rounds. We compared classification results using the interval anchor and adjustment belief updating method with the FILUM method described in [9]. The whole process was repeated five times and the mean of the results was take.

Interval a-a method	FILUM method
63.6%	63.3%

Table 3: Classification results before 60^{th} round

Interval a-a method	FILUM method
57.3%	22.2%

Table 4: Classification results after 60^{th} round

As can be seen from Table 3 classification results before the 60^{th} round (the point of behaviour change) are quite similar between the two wethods. After the 60^{th} round, however, there is a marked difference in the results, with a large fall in the performance of the FILUM approach. The effect can be seen most clearly in Figure 4 where the FILUM approach is slow to react to the change of behaviour of prisoner number 3 from cooperative to tit-for-tat.

5 Conclusions

In this paper we extend Einhorn and Hogarth's epistemological anchor and adjustment belief updating approach to the case where evidence and belief are expressed as intervals within $[0,1]$. We show how this approach is ideally suited to incrementally updating user classifications in uncertain user modelling. We present preliminary results using an iterated prisoner's dilemma testbed.

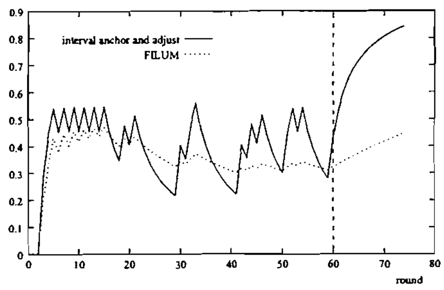

Figure 4: Membership of prisoner 3 in class tit-for-tat

References

[1] R.Axelrod "The evolution of cooperation" New York: Basic Books, 1985

[2] J.F.Baldwin and T.H.Cao and T.P.Martin and J.M.Rossiter, "Towards Soft Computing Object-Oriented Logic Programming", Proceedings of FUZZ-IEEE 2000

[3] J.F.Baldwin and T.H.Cao and T.P.Martin and J.M.Rossiter, "Implementing Fril++ for Uncertain Object-Oriented Logic Programming", Proceedings of IPMU2000

[4] J.F.Baldwin and T.P.Martin, "Fuzzy classes in object-oriented logic programming", Proceedings of the 4th IEEE International Conference on Fuzzy Systems, pp. 1358-1364

[5] T.H.Cao and J.M.Rossiter and T.P.Martin and J.F.Baldwin "Inheritance and recognition in uncertain and fuzzy object-oriented models", Proceedings of NAFIPS 2001

[6] H.J.Einhorn and R.M.Hogarth "Ambiguity and uncertainty in probabilistic inference", Psychological Review, 1985, Vol 93, pp. 433-461

[7] R.M.Hogarth and H.J.Einhorn "Order effects in belief updating: the belief-adjustment model", Cognitive Psychology, 1992, Vol 24, pp. 1-55

[8] R.C.Jeffrey "The logic of decision" McGraw-Hill, 1965

[9] T.P.Martin "Incremental Learning of User Models - an Experimental Testbed", Proceedings of IPMU2000

[10] J.M.Rossiter and T.H.Cao and T.P.Martin and J.F.Baldwin "A Fril++ Compiler for Soft Computing Object-Oriented Logic Programming", Proceedings of Sixth International Conference on Soft Computing, IIZUKA2000, pp. 340-345