An argument structure abstraction for bayesian belief networks: just outcomes in on-line dispute resolution

Nial Muecke **Andrew Stranieri**

School of Information and Mathematical Sciences
University of Ballarat,
PO Box 663, Ballarat, Victoria 3353,
Email: [n.muecke, a.stranieri]@ballarat.edu.au

Abstract

There are many different approaches for settling disputes on-line, such as simple email systems, fixed bid systems and intelligent systems. However, to date there have been no attempts to integrate decision support methods into the dispute resolution process for the purpose of supporting outcomes that are consistent with judicial reasoning.

This paper describes how a model of judicial reasoning can be used to assist divorcees with the resolution of property issues online, in a manner that is consistent with decisions a judge would make if the matter was heard in Court. The approach uses an argument based model of the discretionary nature of decisions made by judges in Australian Family Law. This is integrated with a protocol for online dispute dialogue. Predictions of the likelihood of alternates outcomes is achieved with a series of Bayesian Belief Networks.

Keywords: Argument Theory, Decision Support, Just Decisions, On-Line Dispute Resolution, Family Law

1 Introduction

There are many different approaches for settling disputes on-line, such as simple email systems, fixed bid systems and intelligent systems (Tyler, Bretherton & Firth 2003). The majority of these systems use negotiation, mediation or arbitration as the mechanisms with which to resolve disputes(Bonnet, Boudaoud, Gagnebin, Harms & Schultz 2002). In cases involving e-commerce, Online Dispute Resolution (ODR) is often preferred to litigation because it is less expensive and avoids jurisdictional issues across international borders. (Bonnet et al. 2002).

Family-Winner (Zeleznikow & Bellucci 2003, Bellucci & Zeleznikow 2005) provides useful support for disputes that take place in family law concerning custody or property issues following a divorce. It is a sophisticated negotiation system which ensures that the users receve a fair negotiation, however what the system is incapable of is distributing the assets in proportion to what is just and equitable for each of the disputants. Thus any agreement reached has a high probability that it will be seen as unjust. Similar problems arise with other attempts at legal dispute resolution systems such as DEUS(Zeleznikow & Bellucci 2003).

Unlike disputes that occur in commercial settings, in family law there is a greater need for outcomes to be consistent with decisions a Family Court judge would hand down if the dispute was litigated. Negotiating parties in family law disputes are often not equal in bargaining position and, according to (Alexander 1992), sometimes one party will accept a less than just agreement in order to end the conflict, particularly in marriages where spousal abuse has occurred.

An ODR approach will be advanced in this paper that tightly couples a prediction of judicial outcomes into the dispute resolution process. However, due to the discretionary nature of Australian Family Law, determining what a judge sees as a fair decision is not a trivial task. The main statute, the Family Law Act of Australia, provides a judge with a list of relevant factors to consider, such as the the age of both parties and the paramount interests of the child - but the statute is silent on their relative weightings. This issue of representing reasoning in discretionary legal domains is not new and although there has been a great deal of discussion from a legal theoretic perspective there has been less discussion about the appropriate ways to represent knowledge for a model of discretionary reasoning.

Dowrkin(1977) identified three basic types of discretion that exemplified different conclusions inferred in different ways: two types of weak discretion and a type he characterized as strong. Strong discretion characterizes that reasoning which involves the liberty, on the part of the reasoner, to incorporate standards of his or her own choosing. To use his example, an Army sergeant required to select five men exercises strong discretion because no constraints are supplied to guide the decision. Dworkin proposes that this is the nature of the position of the judge in a situation where usual standards or rules do not apply. Weak discretion of type one exists when one's decision is bound by rules that may inherently have variable interpretations, but nevertheless, those rules apply. The Sergeant required to select five experienced men exercises Type 1 discretion because the decision involves interpreting the standard of experienced. The second type of weak discretion exists when a decision is made according to applicable rules and standards but the decision maker's decision stands as final and cannot be appealed. A decision made by an umpire during a football match exemplifies Type 2 discretion.

Although Mc Cormack(1981) suggested discretion is more a continuum rather than distinct categories as proposed by Dworkin, the basic tenent still maintained was that discretion involved the liberty to infer conclusions in various ways. Christie (1986) claimed that the exercise of discretion involves power relationships within a political system but shares the view with Bayles (1990) that discretion involves the ability to reason toward one of a number of possible conclusions.

Knowledge representation schemes that have been used to model discretionary reasoning include the layout of arguments advanced by Toulmin (1969) . The Toulmin structure has been extended by (Yearwood & Stranieri 2006) in their Generic Actual Argument Model(GAAM). DiaLaw Lodder(1999)(Lodder 1999) uses a combination of dialogue and argumentation ideas to represent a legal argument. Other approaches have been used such as hybrid Fuzzy Logic/Neural Networks (Hollatz 1999, Philipps & Sartor 1999) and Bayesian Belief Networks (Davis & Pei 2003, Halliwell, Keppens & Shen 2003).

A model that combines a Generic Actual Argument Model with a Bayesian Belief Network (BBN) will be advanced in this paper. It will be argued that in order to model discretionary reasoning in family law sufficiently well to make predictions and achieve just outcomes in an ODR system, the use of the GAAM or BBN alone will not suffice. Instead a combination of the two is better suited to the task. The next two sections will discuss the GAAM and the BBN respectively in the context of discretionary reasoning. Following this there will be a description of how the two representations can be combined. This will be followed by a sample consultation, to highlight how the proposed model works in an ODR environment.

2 Generic Actual Argument Model

The Generic Actual Argument Model (GAAM) (Yearwood & Stranieri 2006, Stranieri, Yearwood & Zeleznikow 2002) is a structure for capturing expert reasoning. The approach is based on the argument structure proposed by Toulmin (1969). The Toulmin Argument Structure (TAS) Fig 1 was proposed as an alternative to the traditional approach proposed by Aristotle of 'minor premise; major premise; so conclusion'. An example of this is 'Socrates is a man; All men are mortal; so Socrates is a mortal'. Toulmin argued that although the Aristotle approach to arguments was suited to analytical argument it is inadequate for uses with other types of arguments. Take for example an argument that involves uncertainty, 'Mustafa is an Arab; Most Arabs are Muslims; So Mustafa is most likely a Muslim'. Under Aristotle's argumentation approach this type of argument would be dismissed because the link between the premise and conclusion is uncertain. The TAS on the other hand acknowledges that these types of arguments are commonplace and arguments that involve uncertainty are still valid.

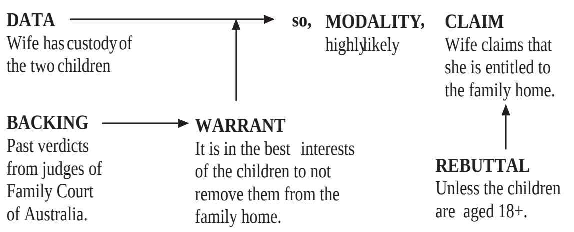

Figure 1: Toulmin Argument Structure, Family Law example

The TAS consists of six parts which each relate to different components of a well formed argument. In Fig 1 these are:

Claim: The claim being advanced. e.g The wife is entitled to the family home.

Modality: The extent to which the claimant believes in the claim they are advancing. e.g highly likely

Data: One or more data items that support the claim. e.g Wife has primary care of the children.

Warrant: For each data item a reason why that data item is relevant to the argument. e.g It is in the best interests of the children to not remove them from the family home.

Backing: For each warrant there may be a backing to strengthen the warrant and justify it's relevance. e.g Past verdicts from judges.

Rebuttal: An argument against the advanced claim or a condition under which the claim becomes invalid. e.g Unless the children are under 18.

.

The argument made in Fig 1 would read: The wife believes that she has a strong claim to the family home. This is because she has custody of the two children from the marriage and she is able to cite precedents from past cases that dictate that the lives of children under 18 should be disrupted as minimally as possible. The data item, that she has primary care of the children, is a claim of another argument that itself has a data item, warrant, modality and backing.

The Generic Actual Argument Model is an adaptation of the TAS. One difference involves the replacement of the warrant in the Toulmin structure with two components; a reason for the relevance of each data item and an inference mechanism that acts as a procedure to infer the claim from the data items. Fig 2 represents part of a generic argument involved in selecting the five experienced patrolmen. In this case the assertion made is that patrolmen A, B, C, D and E are the five selected. This claim is made on the basis of data. The two data items are the pool of patrolmen and the years of service for each patrolman. The reason that the first data item is relevant derives from the Captains orders. The reason that years of service is relevant is that this is the standard for experience that the Sergeant has chosen. The inference mechanism is a function that ranks patrolmen using the standard.

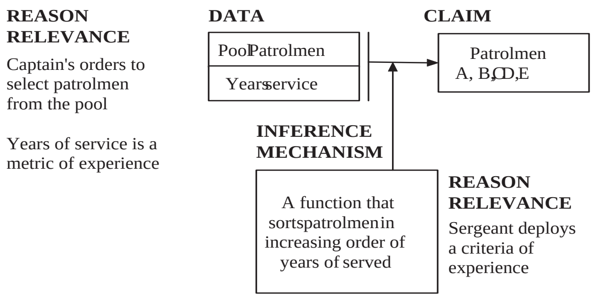

Figure 2: Argument structure for Dworkin's patrolmen

Three ways a decision maker exercises discretion can be discerned from the representation in 2.

- The decision maker has discretion to add (or remove) data item factors
- The decision maker has discretion to use an inference procedure of his/her own choosing to infer a claim value from data item values
- The decision maker has the discretion to leave data items, reasons for relevance, inference procedure, and reasons for the appropriateness of inference procedures implicit.

The GAAM also differs from the TAS by utilizing a variable/value representation of claims. Fig 3 illustrates a generic argument that models the claims the wife is entitled to the family home and is not entitled to the family home. An actual argument is an instance of the generic template. Fig 3 illustrates an actual argument that infers the wife is entitled to the family home. This claim is made using a set of rules known as myRules on claims that the wife has primary care and moving will have enormous adverse effects

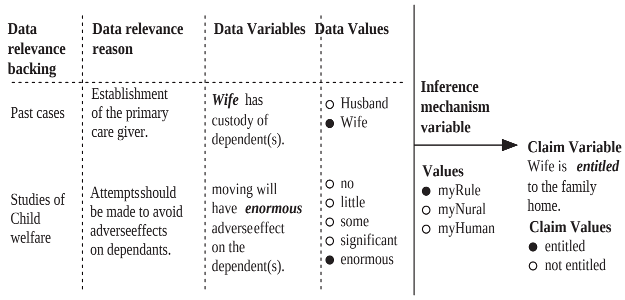

Figure 3: GAAM Family Law example

In Fig 3 an additional difference involves the omission of the rebuttal component of the TAS and the addition of two new items: Inference Mechanism Variable and Inference Mechanism Value described below:

Claim Variable(Claim): The current claim being advanced.

Claim Value: These are different values, one of which will be used to change the claim.

Inference Mechanism Variable: The current inference mechanism being used. This is used to combine all the Data Variables and Data Values for a claim in order to make a prediction as to what the the claim Value should be.

Inference Mechanism Value: A list of the available inference mechanisms.

Data Variable (Data): These are subarguments or factors that are relevant in supporting/determining the claim.

Data Value: These are different values, one of which will be used to change the corresponding data Variable.

Data Relevance Reason (Warrant): This is the reason why the corresponding data Variable is relevant to the claim.

Data Relevance Backing (Relevance): This is some form of backing for the corresponding Data Relevance Reason.

.

(Stranieri & Zeleznikow 2005) described the use of the GAAM for modeling reasoning in family law property proceedings in their Split-Up system. The GAAM uses a tree structure known as an argument tree. The argument tree extracts the claim and data items from a series of interconnected arguments, where each argument takes the form described in Fig 3. Each level of the tree represents a subargument or factor of the higher level. A segment of the Split-Up Family Law tree is provided in Fig 4 which illustrates the tree structure. Claim values, inference mechanisms and reasons for relevance are not included in Fig 4.

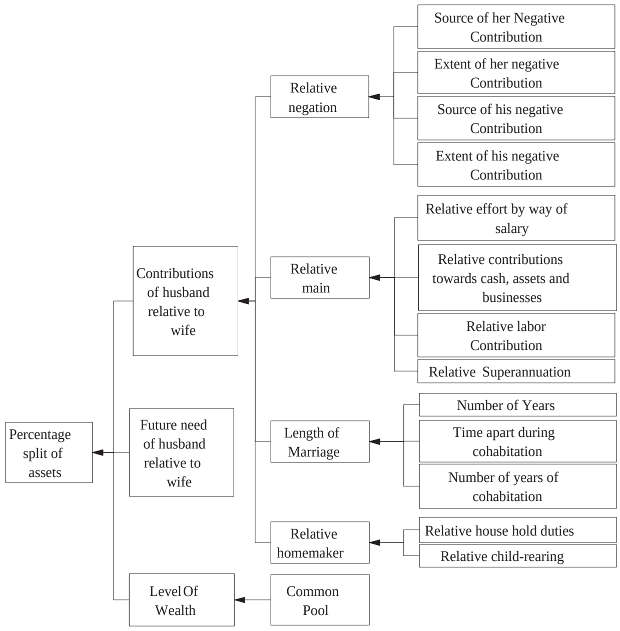

Figure 4: Section of the Split-Up Argument tree. (Stranieri et al. 2002)

The top of the tree is the root node, which represents the argument of division of marital assets. Under the root node are the three most important arguments in determining the division of marital assets. These arguments were elicited from past cases and legislation by family law experts. These nodes are children of the root node. Below each of the three arguments are the important subarguments and so on. For instance, a judge working out the length of a marriage will consider the number of years of marriage, time apart during cohabitation and the number of years of cohabitation prior to marriage. A judge will combine these factors in order to reach, what he or she believes to be a judgement about the length of a marriage.

In total, 93 factors were identified by Stranieri & Zelenzikow (2005) as being relevant by family law experts. These were structured as an argument tree that can be used to infer a percentage split of property at the root node. For a factor to be relevant, family law experts involved in the knowledge acquisition must articulate a reason for its relevance. Skabar (1997) demonstrated that a subset of 17 of these factors could be discovered to yield predictions as accurate as those obtained using the entire set. However, notwithstanding this, the full 93 were insisted upon by experts as necessary for inclusion in order to offer a more complete explanation of the reasoning. To take an extreme example of this, one can imagine a magic dice that correctly infers judicial outcomes yet is clearly inadequate for explaining a decision.

Stranieri & Zelenzikow (2005) trained neural networks by gathering data from past cases heard in the Family Court. The data was used to train a network at each level of the argument tree. Split-Up predicted court outcomes by prompting the user for leaf node values and running inferences up the tree until a claim value at the root node was inferred. This approach is limited to the prediction of a judge's decision given one set of facts; those of the party setting the leaf node values. However, the approach is not satisfactory in an ODR context because both parties must weigh up the strength with which they hold their beliefs against the likelihood that a judge will not agree with their claims. What is required is an approach that can propagate certainty values forward from leaf node to root node for each party. In the next section Bayesian Belief Networks are deployed for this purpose.

3 Bayesian Belief Networks

A Bayesian Belief Network(BBN) is a directed graph where nodes represent factors of a domain problem. Nodes are linked with directional arcs which represent the dependencies of which one node relies on another in order to make uncertainty predictions.

Fig 5 & Table 3 illustrates a BBN for calculating the length of a marriage. A husband or wife assert their claim for a share of the property following divorce. A judge's decision about the length of a marriage(Lm) is made on the basis of three factors; the number of years the parties have been married(Ym), the length of cohabitation prior to the marriage(Cl), if any, and the duration of any time apart during the marriage or cohabitation(Ta).

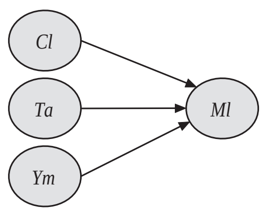

Figure 5: Bayesian network for marriage length

	Node	Values
Ml	Marriage length	{very long, long, short, very short}
Cl	Prior Cohab Length	{very long, long, short, very short}
Ta	Time apart	{great deal, some, none}
Ym	Years married	{0-5,6-10,11-15,16-20,21-25, over 25}

Table 1: Values for Marriage length

Table 3 depicts the belief the husband holds about the length of the cohabitation prior to the marriage. He is 90% certain the prior cohabitation was long but accepts there is a smaller chance that it could be seen as very long. Table 3 illustrates the husband's certainty that the marriage was between 16 and 20 years of length. The husband is also certain there was no time apart during the marriage or cohabitation though this is not illustrated below. The conditional probability provided below utilizes the prior probability in the calculation for each marriage length outcome given the antecedent factors. For example, the prior probability that a marriage is very long given a very long prior cohabitation, no time apart and 16 to 20 years marriage is 90%. Prior probabilities are currently estimated by family law experts but can conceivably be drawn from databases of past decisions.

$$P(Ml|Cl, Ta, Ym)$$

The application of the Bayesian calculation using the husband's beliefs illustrated in Tables 3 & 3 indicate the marriage length is likely to be regarded as very long by a judge with a confidence of 72% and long with a confidence of 28%. The probability that a judge will regard the marriage as short or very short is deemed to be 0%.

Cl	P(Cl)
very long	0.1
long	0.9
short	0.0
very short	0.0

Table 2: Prior probabilities for Prior Cohab Length

Ym	P(Ym)
0-5	0.0
6-10	0.0
11-15	0.0
16-20	1.0
21-25	0.0
over 25	0.0

Table 3: Prior probabilities for Yrs Married

Typically a Bayesian Belief Network is constructed in consultation with domain experts who supply relevant factors and estimate prior probabilities. In the next section an approach that integrates Bayesian Belief Networks into the GAAM is presented.

4 Combining GAAM with BBN

The Generic Actual Argument Model (GAAM) and the Bayesian Belief Network (BBN) are integrated for the ODR approach advanced here. A justification of the advantages of doing this are presented later in this section. Before this is done, the way the two are integrated is described. The GAAM is a structure for modeling legal reasoning in an explicit manner that can be justified. However the model does not specify a procedure for performing inferences. In contrast, the BBN is a method for calculating uncertainty that is also capable of modeling reasoning.

The GAAM provides a structure for modeling a dispute but not for calculating the strength of an argument. In this work, a Bayesian Belief Network (BBN) is embedded into each level of the argument tree. Integrating the GAAM and the BBN is a relatively simple task. In the GAAM an inference procedure is a mapping between child claim values and parent claim values. A BBN is deployed to implement a mapping. At each node of the argument tree where sub nodes exist (example Fig 4) there needs to be a BBN to infer the support (Fig 6). The BBN at each of these points resembles the structure of the GAAM at the same point; for instance in Fig 6 the node to be inferred is the *Relative Homemaker* and the subnodes that are to be used to make the inference are *Relative household duties* and *Relative child rearing*. This structure is similar in both the GAAM and the BBN, however rather than use a single BBN a number of small BBNs are embedded into the GAAM. For example in the *Relative Homemaker* are the three nodes that make up the BBN - but they are only one part of the GAAM. Though the BBNs used for inferring are separate they share a link to each other through the GAAM.

Fig 7 shows the relationship between the GAAM nodes (*Root*, *A*, *B*, *C* and *D*) and the BBN nodes (*r*, *a*, *b1*, *b2*, *c* and *d*). The Links that join GAAM nodes C & D to their counter part BBN nodes c & d represent the user's argument claim being passed from the GAAM to the BBN. Once the values have been passed in to the BBN, the BBN is able to infer the strength of each of the argument claims for node *b1*, the results of which are passed to node B. There the user is asked to consider the results and make a

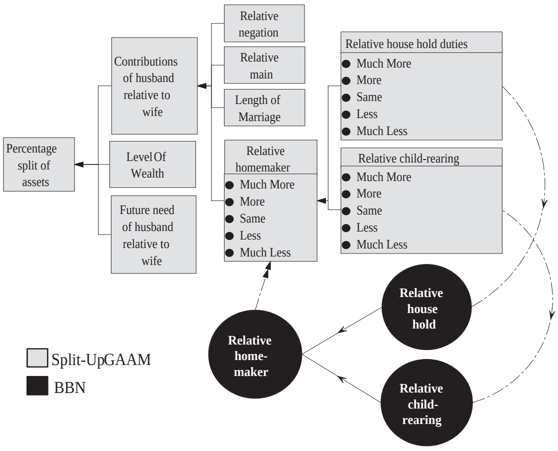

Figure 6: Subsection of the Split-Up Argument tree, including claim values and Inference Belief Network.

selection.

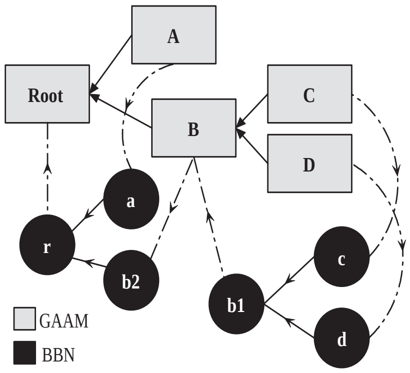

Figure 7: Relationship betwen GAAM and BBN

Both the GAAM and the BBN are tools for modeling reasoning so why not use one or the other? Although they have a common purpose, modeling reasoning, they both have attributes that are not found in the other which can be combined beneficially. The GAAM as illustrated earlier is a structure that is well suited to the representation of domain/expert knowledge, as is the BBN. However the GAAMs ability to not only capture the reasoning but to go further and provide the relevance for the reasoning is lacking in the BBN. Though in some domains this is not overly important it is extremely important in law, as was explained by (Stranieri & Zeleznikow 2005).

The construction of BBN is based on a loosely described protocol whereby an expert's reasoning is mapped out through linking relevant factors. There is no real standard to how factors should relate and it is conceivable that two, or more, different BBNs could be constructed for the same problem. This may or may not affect the inference process, yet domain models constructed with the GAAM are constructed in a particular manner. The linking of related factors requires an expert to articulate a reason and relevance for the link to be justified.

Further, the GAAM can act as a platform for uncertain calculations other than BBN to be used. This enables a greater flexibility when inferences need to be made. It was illustrated earlier that the GAAM by itself specifies no given inference mechanism. This feature of the model allows the BBN to be used. However it is this same feature that allows other inference mechanisms to be included. In this way the GAAM acts as a bridge between different inference mechanisms allowing the different uncertainties inferences to combined.

To use the GAAM without the BBN would result in inferred outcomes that would be of little use to a disputant. On the other hand a useful model can undoubtedly be derived from only using the BBN, however without an explicit and justifiable structure the results would be inadmissible in a legal domain.

In the next section a sample consultation is provided to illustrate the way the GAAM/BBN is a central feature of an ODR approach that encourages decisions that are similar to judicial outcomes.

5 Sample Consultation

The ODR protocol commences once the husband and the wife independently and asynchronously access a web site. They are each presented with the root of the argument, Percentage split of assets(Fig 4) and asked to assert a claim for this argument. The husband (H) claims the split should be 80/20, (80% to H, 20% to W) whereas the wife (W) claims that the split should be in the order of 55/45. Behind the scenes the system compares the claims and informs H and W that they have a disagreement on how the property should be split but does not divulge the respective claims. To do so would be to invite combat and possibly inflate anger.

In order to resolve the root node disagreement the system now prompts H and W to examine and make claims on the next layer of the argument. Both H and W are now presented with those factors; *the contribution of H relative to W*, *the relative needs of H to W* and the *level of wealth in the marriage.* H claims that he contributed much more than W, that W's needs are more than his and the the level of wealth was average. W claims that H contributed more, that she needs more than H and that the level of wealth was average. Again the system compares the claims and informs W and H that there is agreement on two of the arguments but that there is disagreement on the claim to the contributions.

Once again H and W are prompted with the subarguments for the argument in dispute: the relative negative contribution of H to W (negative contributions - exemplified by gambling losses - are those that diminish marital resources), the relative main contribution of H to W, the length of the marriage and the relative home maker contribution of H to W. H claims that: his negative contribution by way of gambling or domestic violence was the same as W (i.e non existent), that he contributed much more than W and that he contributed about the same as W toward the homemaking role.

W though, unsure how to answer the first argument, elects to drill down to the children nodes and have the system infer values for her negative contributions. Now W is given the subarguments and is asked to make claims for each. Where this is a nonleaf argument then all that W would need to do is make the claims as she has done for the other arguments, but as this is a leaf argument this is not the case. In addition to making a claim for each of the arguments W is required to provide some sort of evidence to support her claims. In this example W may have financial records that show large sums of money were withdrawn from a casino in addition to other discrepancies.

Having done this, the system infers the results and suggests to W, that based on her claims and the evidence to support the claims: there is high support a judge will see that H has made a more negative contribution than her, there a low chance that it will be seen that H has made a much more or average negative contribution and is not at all likely to see that H has made the same, less or much less negative contribution. In finding this, W decides to claim that H has made a high negative contribution. W decides that she does not need assistance in deciding claims for the other arguments and claims that: H contributed more than her, that the marriage was of average length and she did contribute more to the home making.

The system once again compares the input from W and H and discovers that there are three arguments that differ. H makes assertions for each of the subarguments as he has done before and provides any evidence for the leaf arguments that he encounters. W also continues in this fashion, however the system informs W that as she has already made claims for the subargument of negative contributions that she is not required to reconsider those claims. If W did decide to reconsider her claims for negative contribution the system will prevent the user from making the change unless they are prepared to submit claims for the subarguments too. As the negative contribution is a leaf node, W can change her claims and or the evidence to support her claims.

Both W and H, having now explored the arguments to the lowest level possible have only to decide if they are happy with the claims suggested by the system or if they would prefer to ignore the system's advice and select a different claim for each argument. To give some incentive for W and H to accept the recommended claims the system provides a support for each claim. The support is a calculation of the likelihood a judge will agree with a particular claim based on the evidence and claims made by both parties. Using the support measure W and H independently work their way back to the root argument where the final recommendation is made, a 60/40 split.

6 Conclusion

An ODR approach was advanced that can incorporate legal fairness. Though the approach can not guarantee 100% fairness, as doing this would move away from ODR into the creation of a virtual judge, it does place a greater emphasis on fairness, which has been lacking in other ODR approches. By examining the way in which a judge reasons and then using this judicial reasoning to structure the dispute, it is hoped that the disputants gain an understanding of the dispute from all positions. Through providing judicially based advice to the disputants on their position, it is believed that disputants will focus more on resolving the dispute than on turning the dispute into a battle of wills. The approach integrates an argumentation based model of reasoning Generic Actual Argument with a Bayesian Belief Network to implement inferences with uncertainty. The prior probabilities derive mainly from past judgements. The next stage in the development of the model discussed in this paper is to implement the ideas as an on-line system. Future research aims to empirically validate the claim that this approach enhances deliberative dialogue in an ODR system.

References

Alexander, R. (1992), 'Mediation, violence and the family', *Alternative Law Journal* **17**(6), 276–299.

Bayles, M. D. (1990), *Procedural Justice. Allocating to Individuals.*, Kluwer, Dordrecht.

Bellucci, E. & Zeleznikow, J. (2005), Trade-off manipulations in the development of negotiation decision support systems, *in* 'Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05)', p. p. 36c.

Bonnet, V., Boudaoud, K., Gagnebin, M., Harms, J. & Schultz, T. (2002), Online dispute resolution systems as web services, *in* 'Proceedings Hewlett-Packard OpenView University Association Workshop'.

Christie, C. G. (1986), 'An essay on discretion', *Duke Law Journal.* pp. 747–78.

Davis, G. A. & Pei, J. (2003), Bayesian networks and traffic accident reconstruction, *in* 'Artificial Intelligence and Law', Vol. 9th.

Dworkin, R. (1977), *Taking rights seriously*, Harvard University Press, Cambridge.

Halliwell, J., Keppens, J. & Shen, Q. (2003), Linguistic bayesian networks for reasoninng with subjective probabolity in forensic statistics, *in* 'Artificial Intelligence and Law', Vol. 9th, pp. 42–50.

Hollatz, J. (1999), 'Analogy making in legal reasoning with neural networks and fuzzy logic', *Artificial Intelligence and Law* **7**, 289 – 301.

Lodder, A. R. (1999), *DiaLaw On Legal Justification and Dialogical Models of Argumentation*, Vol. 42 of *Law and Philosophy Library*, Kluwer Academic Publishers, Dordrecht, Netherlands.

McCormack, D. N. (1981), *H. L. A Hart*, Edward Arnold.. New York:.

Philipps, L. & Sartor, G. (1999), 'Introduction: from legal theories to neural networks and fuzzy reasoning', *Artificial Intelligence and Law* **7**(2 - 3), 115 – 128.

Skabar, A., Stranieri, A. & Zeleznikow, J. (1997), Using argumentation for the decomposition and clarification of tasks for hybrid system development, *in* 'Progress in Connectionist-Based Information Systems: Proceedings of the 1997 International Conference on Neural Information Processing and Intelligent Information Systems', pp. 814–817.

Stranieri, A., Yearwood, J. & Zeleznikow, J. (2002), 'Argumentation structure that intergrate dialectical and non-dialectical reasoning', *The Knowledge Engineering Review* **16**(4), 331–348.

Stranieri, A. & Zeleznikow, J. (2005), *Knowledge Discovery from Legal Databases*, Law and Philosphy Series, Kluwer Academic Publishers.

Toulmin, S. E. (1969), *THE USES OF ARGUMENT*, Cambridge University Press.

Tyler, M. C., Bretherton, D. & Firth, L. (2003), Research into online alternative dispute resolution, Feasibility report, Department of Justice, Victoria.

Yearwood, J. & Stranieri, A. (2006), 'The generic/actual argument model of practical reasoning', *Decision Support Systems* **41**(2), 358–379.

Zeleznikow, J. & Bellucci, E. (2003), Family winner: integrating game theory and heuristics to provide negotiation support, *in* 'Proceedings of Sixteenth International Conference on Legal Knowledge Based Systems', IOS Publications, Amsterdam, Netherlands.