Mixed-Initiative Argumentation: Group Decision Support in Medicine

Chee Fon Chang[1], Andrew Miller[2], and Aditya Ghose[1]

[1] Decision Systems Lab and Centre for Oncology Informatics
School of Computer Science & Software Engg
University of Wollongong, NSW, Australia
[2] Centre for Oncology Informatics
Illawarra Health and Medical Research Institute
University of Wollongong, NSW, Australia
{c03,amiller,aditya}@uow.edu.au

Abstract. This paper identifies ways in which traditional approaches to argumentation can be modified to meet the needs of practical group decision support. Three specific modifications are proposed. Firstly, a framework for accrual-based argumentation is presented. Second, a framework for outcome-driven decision rationale management is proposed that permits a novel conception of mixed-initiative argumentation. The framework is evaluated in the context of group decision support in medicine.

Keywords: Abstract Argumentation System, Group Decision Support, Knowledge Management System.

1 Introduction

In this paper, we propose a framework for mixed-initiative argumentation, which interleaves "winner determination" in the style of classical argumentation with decision identification by the user, coupled with the recording of decision rationale. We therefore present a spectrum, with classical argumentation performing "forward reasoning" model at one extreme, and decision rationale recording in the "reverse justification" model at the other extreme. Mixed-initiative argumentation represents the middle ground. The best means of obtaining a machinery for recording and managing decision rationale is to "invert" the machinery for decision generation. In other words, in a group decision making setting, we need to ask the following question: what inputs to a group decision "generation" system would have generated the selected decision? These inputs then constitute the rationale for the selected decision. Argumentation provides a basis for determining the set of "winning" arguments in settings where multiple points of view need to be accommodated, such as group decision support. Very little exists in the literature on rationale management in group decision support. Rationale management is an important question in a variety of settings. Recording decision rationale can help ensure consistency across a sequence of decision that can be

P. Kostkova (Ed.): eHealth 2009, LNICST 27, pp. 43–50, 2010.

used to justify other decisions. This is also a valuable pedagogical tool. These aspects of rationale management are critical in the medical domain.

Argumentation theory is concerned primarily with reaching conclusions through logical reasoning, starting with certain premises. Argumentation theory is therefore concerned with acceptability, and not necessarily any notion of truth or agreement. In argumentation theory, the notion of conflict is generally represented by either *attack* or *defeat* relations [1]. The use of argumentation in medical decision support is not new. In [2], the authors investigated on the collaborative decision-making and communicative discourse of groups of learners engaged in a simulated medical emergency. In [3], the authors introduced the use of arguments for decision support and advocated the need for decision support systems to support more than single, isolated decision making as most decisions are made in context of extended plans of action. However, the proposed system fails to exploit the full potential of argumentation and does not allow for rationale management or the evolution of rules and preferences. In [4], the authors provides a brief insight to a body of work centred on applications of argumentation in biomedicine. Our work differ by taking a mixed-initiative approach to elicit the background knowledge required in the group decision making as well as rationale management. To motivate our work, let us examine several extracts from a medical group decision session. The discussion involves several medical specialists (Surgeon (S_1,S_2,S_3), Radiation Oncologist (RT_1,RT_2)) debating on the best treatment for a patient with early stage superficial unilateral larynx cancer.

Disease Definition: Larynx Cancer Early Superficial Unilateral	
S_1	: (A_1) My opinion is to take out the patient's larynx. This is has the best cure rate of 99%.
S_2	: (A_2) I agree, taking out the patient's larynx would provide the best cure potential.
S_3	: (A_3) I also agree, taking out the patient's larynx would provide the best cure potential.
RT_1	: (A_4) But if you take out the patient's larynx, the patient will have no voice.
RT_1	: (A_5) However, if you use radiotherapy, there is a 97% cure rate from the radiotherapy and about 97% voice quality, which is very good. The 3% who fail radiotherapy can have their larynx removed and most of these will be cured too.

The above example illustrates several important issues. Firstly, the need for accrual in argumentation. Within argumentation, "accrual" generally refers to the grouping of arguments to support or refute a particular opinion [5]. To highlight our point, let us focus on three key arguments. A_4 forms the basis of an attack on the argument A_1. When just considering these two arguments alone, it maybe difficult to determine which course of action is the most appropriate. Now, let us consider the argument: A_1 in conjunction with the argument A_5. Again, it maybe difficult to determine which choice is a more appropriate action to take. However, when we consider all three argument together, it is clear that the best course of action is to perform radiotherapy before taking out the patient's larynx. Secondly, the ability to strengthen arguments by repetition. To illustrate our point, let us focus on the arguments: A_1, A_2, A_3. Although these three arguments do not enlighten the discussion with any additional information, it is conceivable that in a human debate situation, the number of arguments is sufficient enough to overwhelm any suggestion of the contrary. However, we are not advocating that we should always strengthen a position simply by providing multitude of identical arguments. Performing such task should be informed by additional information such as source's expertise or credibility. Finally, the importance of the information sources during argumentation. If we consider the accrual of identical arguments as a reflection of the norms of a community, then it is conceivable that the first course of action would be to take out the patient's larynx. However, if the specialist RT_1 has special insight or knowledge not shared with the others specialist (e.g. the specialist is the ONLY radiation oncologist in the group), therefore might occupy a somewhat privileged position. It is then possible that the arguments made by this particular specialist may carry more weight. In this example, we motivate that the credibility of the individual presenting the argument is important. Using this notion of credibility, we can infer a preference ordering on the arguments.

Let us consider another snippet extending from the same discussion.

S_2 : My opinion is also that the patient should have a hemi-laryngectomy. This will give a cure rate is as good as radiation therapy.

S_3 : I agree, performing a hemi-laryngectomy would give a cure rate as good as radiotherapy.

RT_1 : Yes, I have performed many hemi-laryngectomies, and when I reviewed my case load, the cure rate was 97%, which is as good as that reported internationally for radiotherapy.

RT_2 : I agree, however, you fail to take into account the patient's age. Given the patient is over 75, operating on the patient is not advisable as the patient may not recover from an operation.

RT_1 : Yes, however, in this case, the patient's performance status is extremely good, the patient will most likely recover from an operation. (i.e. the general rule does not apply)

Notice that the above example illustrates an interesting phenomenon. In this particular instance, the specialist RT_1 did not disagree with the correctness of the presented facts and conclusion in the argument presented by RT_2, but rather the applicability of the underlying inference rule that is used to construct the argument. This phenomenon is defined by [1,6] as "undercut". In this situation, the argument presented by RT_1 is more specific. This indicates that there exist some exceptions to the general decision rules that are context dependent. Furthermore, this also indicates that a revision of the general attack relation should be performed. Finally, let us consider another snippet extending from the same discussion.

S_2 : Reviewing our past case decisions, evidence suggest that the we have always performed a hemi-laryngectomy, hence my preference is to do the same.

S_3 : I agree, however, there is some new medical literature reporting that the voice quality after a hemi-laryngectomy was only 50% acceptable and the reporting institution was the North American leaders in hemi-laryngectomy, hence we should perform radiotherapy.

This example illustrates an attack on the user preference. Similar to the previous example, attacks on the user preference are generally context sensitive. This example also illustrates that an argumentation system should evolve over time, it can accumulate past decision as justification for future decisions (similar to that of a legal common law system). However, it is clear that in some instances, we wish to overrule past precedent. In most argumentation and decision support systems presented in the literature, the systems are relatively static. Most systems are open to new facts, however, have difficulties handling changing rules and preferences.

In the next section, we will present a preference-based accrual abstract argumentation framework (PAAF) modeled on works of Dung[7] and Bench-Capon[8]. Although Bench-Capon[8] presented a value-based argumentation framework, in our work, we view these values as merged preferences of the audiences. Section 3 will illustrate the use of this framework in a medical group decision system.

2 Formal Framework

A preference-based accrual abstract argumentation framework is a triple: $\langle AR, attacks, Bel\rangle$ where AR is a set of arguments, $attacks$ is a binary relation on AR. $Bel = \langle V, \leq, \Phi\rangle$ where V is a set of abstract values, $\leq$ is a total order on V, Φ is a total mapping function which maps elements of 2^{AR} to elements of V. Given $\alpha, \beta \in AR$. For readability, we will denote $attacks(\alpha, \beta)$ to mean α attacks β. Similarly, given $v_1, v_2 \in V$, we will denote $pref(v_1, v_2)$ to mean $v_1 < v_2$ or v_1 is preferred to v_2.

Given the abstract framework, we will now define a notion of a *conflict-free* set of arguments. A *conflict-free* set of arguments is simply a set of arguments where arguments in the set do not attack each other. Let $PAAF = \langle AR, attacks, Bel\rangle$, a set of arguments S is said to be **conflict-free** if and only if $\neg(\exists\alpha\exists\beta((\alpha \in S) \wedge (\beta \in S) \wedge attacks(\alpha, \beta)))$. The notion of acceptability is defined with respect to a set of arguments. An argument is acceptable to a set of arguments (in other words accepted into the set) if the set of arguments attack any arguments that attack the joining argument.

Given $PAAF = \langle AR, attacks, Bel\rangle$ and $\alpha \in AR$. α is **acceptable** with respect to a set of arguments S (denoted as $acceptable(\alpha, S)$) if and only if $\forall\beta((\beta \in AR) \wedge attacks(\beta, \alpha)) \rightarrow \exists\gamma((\gamma \in S) \wedge attacks(\gamma, \beta)))$. Given the notion of conflict-free and acceptability, we are now able to define a notion of admissibility. Admissibility is simply defined as a set of arguments that are conflict-free and that defends itself from all attacks from arguments outside the set by attacking them. A conflict-free set of arguments S is **admissible** if and only if $\forall\alpha((\alpha \in S) \rightarrow acceptable(\alpha, S))$. The admissible sets are ordered based on the $<$ ordering imposed on the abstract value V and by utilising the mapping function Φ, the admissible sets are assigned an abstract value.

For any given admissible set of arguments S, S is an **preferred extension** if there is no admissible set $S' \subseteq AR$ s.t. $pref(\Phi(S'), \Phi(S))$ Note that our definition of a **preferred extension** deviates from that traditionally defined in [7,8]

Given two systems, $PAAF_1 = \langle AR, attacks_1, Bel\rangle$ and $PAAF_2 = \langle AR, attacks_2, Bel\rangle$, $attacks_1$ and $attacks_2$ are *A-consistent* if and only if $(attacks_1 \subseteq attacks_2) \vee (attacks_2 \subseteq attacks_1)$.

Given two systems, $PAAF_1 = \langle AR, attacks, Bel_1 \rangle$ and $PAAF_2 = \langle AR, attacks, Bel_2 \rangle$, where $Bel_1 = \langle V, \leq, \Phi_1 \rangle$ and $Bel_2 = \langle V, \leq, \Phi_2 \rangle$, Φ_1 and Φ_2 are *Φ-consistent* if and only if $\nexists \alpha, \beta \subseteq AR s.t (\Phi_1(\alpha) \leq \Phi_1(\beta) \wedge \Phi_2(\beta) \leq \Phi_2(\alpha))$. The above formula allows us to detect inconsistencies between two different attack and preference relations as well as performing context sensitive revision on the attack and preference relations.

Typical usages of mixed-initiative interaction are in scenarios consisting of multiple machines cooperating or collaborating to perform tasks or activate coordination, such as, in distributed planning in multi-agent systems. Several different levels of mixed-initiative are presented in [9]. Our mixed-initiative argumentation framework falls into the "Sub-dialogue initiation" category where in certain situations the system initiate a sub-dialogue to ask for clarification which may take several interactions. Hence, the system has temporarily taken the initiative until the issue is clarified.

Our system initiates the mixed-initiative interaction by first generating a preferred extension and asking the user for verification. If the user agrees with the decision, the system terminates. If the user disagree with the generated preferred extension, a process of query and answer occurs. Requesting the user to validate arguments and provide additional arguments such that when the system recomputes, an agreement occurs. Each decision is then recorded for future reference. This is an iterative process which interleaves "winner determination" in the style of classical argumentation and decision identification by the user, coupled with the recording of decision rationale.

During the "reverse justification" interaction, any combination of three possible types of modification can occur: a new fact is introduced via a new argument, a revision on the attack relation via a new argument using the notion of a-consistency or a revision on the preference mapping via a new argument using the notion of Φ-consistency. In this mode, prioritized revision occurs on the system where arguments, attack and preference relations from the user are imposed onto the system. When the user presents a set of winning arguments, four outcomes are possible during this comparison.

Firstly, the set of winning arguments identified by the user is a preferred extension with respects to the system. In this instance, no modifications are required and hence no additional arguments or rationales are required.

Secondly, the set of winning arguments is not a preferred extension with respect to the system, however, an addition of a new fact or collection of facts will allow the system to generate the identified preferred extension. In this situation, a new argument or a set of arguments is inserted. This new argument represents a reason for supporting the conclusion, hence the preferred extension constitutes a decision rationale.

Thirdly, in the case where the inconsistency lies between the attack relation, a new argument or sets of arguments that eliminate the inconsistency between the attack relations is required. To perform this, we will refer to two functions. Firstly, a function that extracts from the set of arguments the subset that is *relevant* in relation to the attack relation. In particular, the subset of arguments that makes reference to the particular attack relation within its sub-structure. Thus, $rel_{attacks}(S)$ represents that subset of a set of arguments S that make reference to the offending *attacks* relation. Secondly, a function that will extract from the sub-structure of each argument the encoded attack relation. Thus, $ext_{AR}(S)$ represents a set of binary relations extract from the set of arguments S with respect to some set of arguments AR. For example, assume that we have a set of arguments $AR = \{\alpha, \beta \gamma\}$, an attacks relation consisting of $\{attacks(\alpha, \beta)\}$ and a set of arguments $S = \{\alpha, \beta, \delta\}$ and within the substructure of δ the attack relation $\neg attacks(\alpha, \beta)$. $ext_{AR}(rel_{attacks}(S))$ will return $\{\neg attacks(\alpha, \beta)\}$. We will then augment the existing attack relation with that extracted from the arguments.

Similarly, in the case where the inconsistency lies between the preference mapping function, a new argument or sets of arguments that eliminate the inconsistent between the mapping function is required. For any specific instance, we eliminate the inconsistencies between mapping functions by substituting the existing mapping function with one that is encoded in the arguments. We will again refer to two functions. Firstly, a function that extracts from the set of arguments the subset that is *relevant* in relation to the preference mapping function. In particular, the subset of arguments that makes reference to the mapping function within its sub-structure. Thus, $rel_{\Phi}(S)$ represents that subset of a set of arguments S that make reference to the mapping function. Secondly, a function that will extract from the sub-structure of each argument the encoded mapping function. Thus, $ext_{\Phi}(S)$ represents a set of binary relations extract from the set of arguments S with respect to some set of arguments Φ. We will then replace the existing preference mapping function with that extracted from the arguments. Two key benefits exist in such an approach. Firstly, we are able to evolve an existing argumentation rather than constructing a new argumentation system with new attack and preference relations. This allows for the reuse of arguments, attack and preference relation. Secondly, we are able to address "traceability" issues as the system accumulates justifications as it evolves from one instance to the next, hence allowing us to manage rationale over the life of the system. This allows the process of argumentation to form the basis for rationale management.

3 Medical Group Decision Support System

Utilising a Web 2.0 philosophy, we have constructed a web enabled medical group decision support system utilising Asynchronous JavaScript and XML (AJAX) with a back-end repository. HyperText Markup Language (HTML) and Javascript are used to build the user interface and controls the interaction with the web server. Hypertext Preprocessor (PHP) is used to build the reasoning engine to perform back-end computation of the arguments. MySQL is used as the database repository. The benefits of this appraoch are platform independence, portability, scalability and accessibility.

The prototype was presented to several oncologists and a "head-and-neck" session was simulated. A "head-and-neck" session is where groups of oncologists meet to discuss treatment therapy for cancer cases in the head to neck region. During this session, a typical larynx cancer case was discussed. Treatment analyses are performed over 5 categories. These categories (in order of importance) are as listed: survival, control, physical toxicity, psychological toxicity and clinician's choice. These categories are addressed in stages starting from the most important to the least. Argumentation is performed at each stage and final recommendation is based on the accrual of all arguments over all the stages. Each stage can be viewed as a decision-making cycle where decision made affects the available choices for the next cycle. Given a case description, the system presents a possible recommendation (if one exists). Specialists are then asked if the recommendation is acceptable. If the recommendation is not acceptable, the system asks the specialist to select a recommendation and justify it with arguments, with which the system then recomputes a new recommendation. If the recommendation does not coincide, the system presents its findings and asks for more justifications. This process is iterated until the recommendation of the system coincides with the specialist's choice.

In Figure 1, we present the argument modification interface. The users are allowed to add, delete and modify the arguments associated with a particular treatment choice in the forward learning mode. The changes require that the clinician provide the strength of the evidence and a literature references if used. In essence, by associating the argument with the treatment choice, the user has provided justification for the particular treatment.

In Figure 2, we present the resulting output, which illustrates the recommended decision for each facet of a given sequence of decisions. Each facet has different priority (if two treatments have identical cure and control rates, the one with lower physical toxicity is preferred) and the final treatment choice is computed using these preferences. This figure also illustrates the ability for the user to validate the recommendations and subsequently activate the second of the two learning modes where the user disagrees with the recommendation.

Fig. 1. Arguments

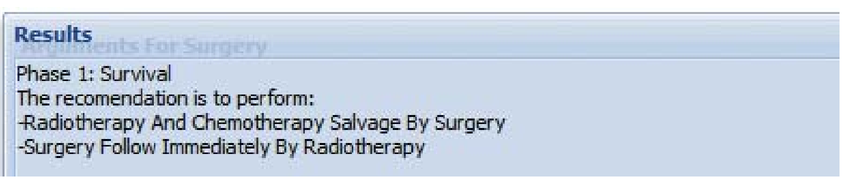

Fig. 2. Recommendations

The general response to using the tool has been positive. The specialists found that the tool is useful both as a practical tool for a trained specialist as well as a teaching aid for medical registrar. The specialists also felt that the tool has huge potential, however, since the tool is still in its' infancy there are still several practical issues that needs addressing. During the trial, several issues were identified. These issues falls into two categories: usability of the user interface and performance of the argumentation engine.

During an original execution of the tool, we found that when computing recommendations, the tool can take sometime to return a decision (in some cases, several hours). This issue was address by limiting the scope to the available 9 therapy choices rather than allowing the system to compute all therapy choices (including non-existing ones).

4 Conclusion

We have identified ways in which traditional approaches to argumentation can be modified to meet the needs of practical group decision support, presented a tool and described its evaluation in the context of group decision support for medical oncology.

References

1. Prakken, H., Vreeswijk, G.: Logics for Defeasible Argumentation, 2nd edn. Handbook of Philosophical Logic, vol. 4, pp. 218–319. Kluwer Academic Publishers, Dordrecht (2002)
2. Lu, J., Lajoie, S.P.: Supporting medical decision making with argumentation tools. Contemporary Educational Psychology 33(3), 425–442 (2008); Collaborative Discourse, Argumentation, and Learning
3. Glasspool, D., Fox, J., Castillo, F.D., Monaghan, V.E.L.: Interactive decision support for medical planning. In: Dojat, M., Keravnou, E.T., Barahona, P. (eds.) AIME 2003. LNCS (LNAI), vol. 2780, pp. 335–339. Springer, Heidelberg (2003)
4. Fox, J., Glasspool, D., Grecu, D., Modgil, S., South, M., Patkar, V.: Argumentation-based inference and decision making–a medical perspective. IEEE Intelligent Systems 22(6), 34–41 (2007)
5. Verheij, B.H.: Reason based logic and legal knowledge representation. In: Carr, I., Narayanan, A. (eds.) Proceedings of the 4th National Conference on Law, Computers and Artificial Intelligence, University of Exeter, pp. 154–165 (1994)
6. Pollock, J.L.: Defeasible reasoning. Cognitive Science 11(4), 481–518 (1987)
7. Dung, P.M.: An argumentation-theoretic foundation for logic programming. Journal of Logic Programming 22(2), 151–171 (1995)
8. Bench-Capon, T.J.M.: Persuasion in practical argument using value-based argumentation frameworks. Journal of Logic and Computation 13(3), 429–448 (2003)
9. Allen, J.F.: Mixed-initiative interaction. IEEE Intelligent Systems 14(5), 14–23 (1999)