

# Argumentation for Decision Support

Katie Atkinson<sup>1</sup>, Trevor Bench-Capon<sup>1</sup> and Sanjay Modgil<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Liverpool, Liverpool, UK  
{katie, tbc}@csc.liv.ac.uk

<sup>2</sup> Advanced Computation Lab, Cancer Research UK, Lincoln's Inn Fields, London, UK  
sm@acl.icnet.uk

**Abstract.** In this paper we describe an application based on a general approach towards modelling practical reasoning through defeasible argumentation. The purpose of the paper is to show how the incorporation of an argumentation component can add value to a collection of existing information agents. The example application is a system for reasoning about the medical treatment of a patient. An agent, called the *Drama* agent, orchestrates a number of information sources to supply a set of arguments on the basis of which the decision regarding treatment can be taken. We describe the general approach and its instantiation for this application, and illustrate the operation of the system with a running example.

## 1 Introduction

We describe an application based on a general approach towards modelling practical reasoning through defeasible argumentation<sup>3</sup> to show how an argumentation component can add value to a collection of existing information resources. The example application is a system for reasoning about the medical treatment of a patient. We assume that a number of information sources, representing different areas of medical knowledge and facts about individuals, and different policies and perspectives relevant to the problem are available. The focus of this paper is the *Drama* (for Deliberative Reasoning with Arguments about Actions) agent which orchestrates these contributions in argumentation terms, and comes to a decision based on an evaluation of the competing arguments. We begin by describing our general approach to such deliberative reasoning. Section 2 will give an overview of the application. Section 3 will describe how the general approach is used in the particular application using a representative example of such deliberation, and section 4 will discuss the potential advantages of the approach.

A general approach to persuasive and deliberative reasoning about action has been presented in [1]. First a presumptive justification for a course of action is found. This takes the form of an instantiation of the argument scheme AS1:

**AS1** In the circumstances R  
    We should perform action A  
    Whose effects will result in state of affairs S  
    Which will realise a goal G  
    Which will promote some value V.

---

<sup>3</sup> For a comprehensive survey of logical models of argument see [3].

AS1 is an enrichment of the *Sufficient Condition Scheme*, one of the presumptive argument schemes for practical reasoning proposed by Walton [9]. This enrichment allows us to additionally distinguish between: the consequences of the action (S); the *desired* consequences of the action (G); the *reason* why those consequences are desired (V).

These distinctions have been found to be crucial in some applications, including reasoning with legal cases and reasoning about political decisions [1]. The importance of making these distinctions will be further shown in the example.

Next the presumptive justification must be subjected to a critique. Associated with an argument scheme are a number of characteristic critical questions, which could lead to the justification being defeated. This critique will also identify alternative actions for consideration. In [1] sixteen critical questions associated with AS1 are identified. Two sample critical questions are: *are the current circumstances in fact R?* and *are there alternative ways to achieve the goal G?* Each critical question has associated with it preconditions for making a counter argument. Thus an agent could question the truth of the circumstances if it believed that they were other than R, or could suggest an alternative action if it believed that it would also realise G. Preconditions in terms of the beliefs and desires of an agent are given in [1]. For each critical question whose preconditions are satisfied, one or more arguments attacking the original justification can be produced. These arguments may in turn be subject to the same process of critical questioning to generate counter arguments.

When the set of arguments and counter arguments have been produced, it is necessary to consider which of them should be accepted. In order to do this the arguments *Args* together with the binary *Attack* relation on *Args* are organised into an *argumentation framework* (*Args, Attack*), as introduced by Dung in [5]. In [5] an argument *A1* is always *defeated* by an attacker *A2*, unless *A2* can itself be defeated. This is appropriate to reasoning about beliefs, but when reasoning about actions we are, to a certain extent, free to choose what we will attempt to bring about. Thus we may choose to reject an attacker even if it cannot be defeated, provided we regard the purpose motivating the attacked argument as more important. For example, an argument *A2* that a particular drug is expensive attacks an argument *A1* for prescribing the drug. However, we may none the less choose to prescribe it (accept the argument *A1*) if that would serve a purpose (promote a value) we rate more highly than expense. To accommodate the notion of the value promoted by the acceptance of an argument we use an extension of Dung's framework, *value-based argumentation frameworks* [2]. The idea is that we are given an *argumentation framework*, a set of values *V*, a function *val* mapping each argument to a value  $v \in V$ , and a set of *audiences* *a* (in the sense of [6]) to which the arguments are addressed. Each audience is represented (as in [2]) by a strict partial ordering  $>_a$  on the values. Then, for a given audience  $>_a$ , *A2* defeats *A1* iff *A2* attacks *A1* and  $val(A2) \not>_a val(A1)$ . Note that by this definition, if *A2* attacks *A1* and both promote the same value then the attacks always succeeds as a defeat.

For a given audience and defined defeat relation, we can then determine which arguments in *Args* are acceptable by determining the preferred extension. The preferred extension is the maximal (under set inclusion) subset *S* of *Args* such that no two arguments in *S* defeat each other, and all arguments *A* in *S* are acceptable with respect to *S*,

i.e., for any argument  $A$  in  $S$ , if  $A$  is defeated by an argument  $A'$  that is not in  $S$ , then there exists an argument in  $S$  that defeats  $A'$ . The preferred extension thus represents the maximal consistent set of acceptable arguments with respect to the argumentation framework and a given audience or value ordering. Cycles in the same value give rise to multiple preferred extensions (e.g.,  $A_1$   $A_2$  attack each other and promote the same value). In building the framework it is necessary to resolve cycles in a single value by expressing a preference for one of the arguments in the cycle, based on considerations other than the value: for example that the action is intrinsically preferred.

It has been shown in [2] that an efficient algorithm exists for the computation of the preferred extension of such a framework for a given value ordering, once the cycles have been resolved. We may therefore compute the preferred extensions corresponding to the possible value orderings to discover the dialectical status of the arguments in the framework. When evaluated the arguments may have a unique status, or their status may be dependent on the value ordering. In this latter case the agent may either commit to a particular value ordering, or determine the value ordering in the course of its deliberation [4]. Thus the Drama agent will deliberate on a course of action by:

- obtaining a presumptive justification for some course of action,
- generating any counter arguments to the course of action by posing critical questions, where each such counter-argument is itself subject to counter-arguments (posing of critical questions),
- selecting the course of action by organising the resulting arguments into an argumentation framework, and calculating the preferred extension corresponding to its ranking of values.

In the next section we will discuss the particular application which we will use to exemplify our approach in this paper.

## 2 Deliberative Reasoning about Medical Treatment

Clinical guidelines promote best practice in clinical medicine by specifying the selection and sequencing of medical actions for achievement of medical goals. There is a large body of research into computational support for authoring and enactment of clinical guidelines [8]. Authoring tools support specification of a guideline in some suitable knowledge representation formalism. This specification can then be executed in a specific clinical context so as to enforce compliance with the best practice encoded in the guideline. The authored guidelines need to be specified at a level of abstraction that enables enactment in any number of contexts. It is at execution time that the context dependent choice of specific medical actions must be made.

For example, a guideline may indicate that treatment of a patient recovering from myocardial infarct (heart attack) requires realisation of the treatment goals: treat pain; treat sickness; prevent blood clotting. It is at execution time that one must take into account the specific context in order to choose which precise action should be chosen for realising each of these goals. Examples of contextual factors that influence the decision include:

- information about the specific patient being treated, e.g., administration of a particular drug for preventing blood clotting may for safety reasons be contraindicated by a patient’s clinical history,
- concomitant treatments, e.g., the efficacy of a drug for preventing blood clotting may be reduced by drugs being administered for a gastrointestinal condition,
- local resource constraints, e.g., budget constraints at the local hospital may indicate a preference for one drug over another,
- local organisational policies, e.g., the local health authority may have evidence based preferences for one drug over another.

We propose that it is through deliberative argumentation of the type described in this paper that one can model how these contextual factors can be brought to bear on the decision as to what is the most appropriate treatment action in a given situation. In particular, by structuring a recommendation for action as an argument instantiating argument scheme AS1, one can effectively account for the influence of contextual factors on the decision making process; i.e., by instantiating AS1’s critical questions. Furthermore, the complexity and diversity of the contextual knowledge and reasoning suggests that the information required to be considered may best be distributed across a number of information sources.

In the example below, the medical knowledge cited is for illustrative purposes only: we make no claims for it either as a model of the medical domain, or as a representation of the state of the art of medical systems. Our purpose is only to show how value can be added by the addition of an argumentation agent capable of reasoning with multiple perspectives and drawing on a range of sources.

### 3 Application to the Medical Domain

In our application we locate all argumentation knowledge inside the Drama agent, and the other resources are conventional knowledge and database systems. In particular these other resources are independent of values. The other resources that the Drama agent will interact with in our example are shown in Table 1. Some will contain generic medical knowledge, while others are specific to the organisation.

**Table 1. Resources in the Drama System**

Resource	Type	Scope
Treatment KB	Knowledge Base	Generic Medical Policy and Knowledge
Policy KB	Knowledge Base	Organisation Specific Knowledge
Safety KB	Knowledge Base	Generic Medical Knowledge
Patient DB	Database	Patient Specific Information
Cost KB	Knowledge Base	Organisation Specific Knowledge
Efficacy KB	Knowledge Base	Specific Medical Knowledge

Following the general approach, the Drama agent will use critical questions to identify the information required to generate arguments. In any particular application a char-

acteristic set of the critical questions will be pertinent (see [1]). In this application we assume that all resources have a common vocabulary, and any information given can be accepted as true. Given these assumptions there are five critical questions pertinent to this particular application:

- CQ1: Are there alternative ways of realising the same effects?
- CQ2: Are there alternative ways of realising the same goal?
- CQ3: Are the assumptions on which the argument is based true?
- CQ4: Does performing the action have a side effect which demotes some other value?
- CQ5: Will the action have the effects described?

The Drama agent now constructs an argumentation framework by instantiating AS1 and posing these critical questions. We will illustrate the operation of the system with a running example of a patient whose health is threatened by blood clotting. The framework begins with the null option - do nothing (EA0). The purpose of this is similar to the assumption of the negation of the desired goal in refutation resolution: extensions of the resulting argument frameworks will be acceptable only if they do not contain this argument. The goal of preventing blood clotting is now issued to the *Treatment KB*.

The Treatment KB is one among a number of treatment knowledge bases, each of which is specialised for recommending treatment actions for a medical speciality. In our example, the Treatment KB is specialised to the cardiac domain. For the purposes of this example we will suppose that the Treatment KB, and the other KBs, use simple Prolog rules and are capable of solving a goal and returning a proof trace. The Treatment KB might include (P denotes the patient in question):

```
prevent_blood_clotting(P):-
    reduce_platelet_adhesion(P).
prevent_blood_clotting(P):-
    increase_blood_clot_dispersal_agents(P).
reduce_platelet_adhesion(P):-
    not contraindicated(aspirin,P),
    prescribe(aspirin,P).
reduce_platelet_adhesion(P):-
    not contraindicated(chlopidogrel,P),
    prescribe(chlopidogrel,P).
increase_blood_clot_dispersal_agents(p):-
    not contraindicated(streptokinase,P),
    prescribe(streptokinase,P).
```

It will therefore be able to return the information that blood clotting can be prevented by reducing platelet adhesion, which can, assuming aspirin is not contraindicated, be achieved by prescribing aspirin. The Drama agent can use this information to instantiate AS1, thus providing a justification for this action:

**EA1** *Assuming no contradictions, we should prescribe aspirin, which will reduce platelet adhesion, preventing blood clotting, and so is an efficacious course of action.*

This argument has to be subjected to a critique to ensure that there are no better alternatives. The Drama agent will go through its repertoire of critical questions. Posing CQ1 will ask for alternative solutions to reduce platelet adhesion from the Treatment Agent and elicit the information that chlopidogrel will also reduce platelet adhesion. Asking CQ2 will seek further solutions from the Treatment Agent for preventing blood clotting and will identify the alternative course of action of administering streptokinase, which has the same goal of preventing blood clotting, but via a different effect of increasing the blood's production of agents that disperse clots. These are formed into two arguments, EA2 and EA3:

**EA2** *Assuming no contradictions, we should prescribe chlopidogrel, which will reduce platelet adhesion, preventing blood clotting, and so is an efficacious course of action.*

**EA3** *Assuming no contradictions, we should prescribe streptokinase, which will increase blood clot dispersal agents, preventing blood clotting, and so is an efficacious course of action.*

These three arguments all mutually attack one another, giving rise to the argumentation framework shown in Figure 1.




Figure 1: Initial Argumentation Framework.

Any of EA1, EA2 or EA3 would serve to defeat EA0, 'do nothing'. However, they are in mutual conflict. As they all relate to the same value (which means that the preferred extension is empty for all audiences), there is a free choice between them. They can be chosen according to intrinsic preferences regarding the goal or the actions themselves. The Drama agent therefore contacts the *Policy KB* to see what the preferences of the organisation are.

The *Policy KB* contains organisation specific information to determine preferences between goals, effects and actions. Any criteria could be used here. For the purposes of the example we will assume that the *Policy KB* prefers the effect 'reduce platelet adhesion' as a means by which the goal can be realised, since the effect of increasing blood clot dispersal agents has potentially more undesirable side-effects. Hence, the *Policy KB* will favour actions with the former effect over actions with the latter effect. This, however, does not discriminate between aspirin and chlopidogrel. Again many criteria are possible: it could depend on local stocks held, or a local preference for generic drugs. Here we will assume that cost is the basis for preference and that aspirin is cheaper than chlopidogrel.

We represent these preferences in our argumentation framework by removing the attacks of the unfavoured actions, so that EA1 is no longer attacked, and EA3 no longer attacks EA2 (see Figure 2, ignoring the dotted arrows indicating arguments submitted later on). Now EA1 will form the preferred extension of this framework, and so its action is currently the best candidate. There remain, however, some further critical questions that can be asked of EA1.

EA1 assumed that aspirin was not contraindicated. CQ3 instructs us to test this assumption. This is the role of the *Safety KB*. The Safety KB has knowledge of contraindications of the various drugs, and the reasons for the contraindication. The Safety KB might contain:

```
contraindicated(D,P):-  
    risk_of_gastric_ulceration(D,P).  
risk_of_gastric_ulceration(D,P):-  
    increased_acidity(D),  
    history_of_gastritis(P),  
    not acid_reducing_therapy(P).  
increased_acidity(aspirin).
```

When contacted by the Drama agent it will use this knowledge, together with patient specific information obtained from the *Patient DB* to inform the Drama agent that since the patient has a history of gastritis, aspirin is contraindicated because its acidity may result in gastric ulceration. The Drama agent will form this into an argument motivated by the value of safety. Note that because each of the information sources represents a particular perspective on the problem, the Drama agent may ascribe a motivating value to the argument on the basis of its source.

**EA4** *Where there is a history of gastritis and no acid reducing therapy, we should not prescribe aspirin, which would cause excess acidity, which would risk ulceration, and so is unsafe.*

When EA4 is added to the argumentation framework (arrow 1 in Figure 2), EA4 attacks EA1. Assuming that safety is preferred to efficacy, EA4 defeats EA1 and so EA2 replaces EA1 in the preferred extension.

Assuming EA2 cannot be attacked by CQ3, the next critique follows from CQ4. Efficacy is not the only value: any action must be acceptable within the cost constraints of the organisation. Answering this critical question is the province of the *Cost KB*. This KB will have knowledge of the budgetary constraints on treatment, and will compare the cost of the proposed treatment with these constraints. Suppose that chlopidogrel exceeds these limits. At the minimum this is simply a query as to whether the cost of the treatment exceeds a given threshold, posed to a database of treatment costs. The Drama agent can now form the argument EA5:

**EA5** *Where cost of chlopidogrel is £N, we should not prescribe chlopidogrel, which would cost £N, exceeding our budget, which demotes the value of financial prudence.*

Adding EA5 (arrow 2 in Figure 2) means that EA2 is defeated if cost is preferred to efficacy. This now means that EA3 is in the preferred extension. Since it is unchallenged there is an obligation to critique the proposal to prescribe streptokinase, by returning to CQ3 and CQ4. Suppose that streptokinase is not contraindicated, and that it falls within the cost constraints. There remains CQ5, and we must now investigate whether streptokinase will be effective for the particular individual we are treating. The *Efficacy KB* will contain specific data from clinical trials and past cases indicating the efficacy of actions with respect to treatment goals for particular patient groups. Perhaps (and this is simply an illustrative conjecture on our part) the efficacy of streptokinase has been found to depend on age. The Efficacy KB may then contain rules such as:

```
effectiveness(P, streptokinase, prevent_blood_clotting, 90):-
    age(P,A),A < 50.
effectiveness(P, streptokinase, prevent_blood_clotting, 30):-
    age(P,A),A > 49.
acceptable(P, Treatment, prevent_blood_clotting):-
    effectiveness(P, Treatment, prevent_blood_clotting, E),E > 75.
```

Together with particular patient data obtained from the Patient DB, the Efficacy KB passes this information to the Drama agent which expresses it as EA6:

**EA6** *Where patient is aged 72, we should not prescribe streptokinase, as the likelihood of success is 30%, which is below the required threshold, which demotes efficacy.*

We add EA6 to the framework, as shown in Figure 2 (arrow 3). EA3 is attacked by an argument with the same value and so is defeated. If safety is preferred to efficacy then EA1 is defeated by EA4. If cost is preferred to efficacy then EA2 is defeated by EA5. This would mean that EA0 would be included in the preferred extension as all its attackers are defeated. However, as stated from the outset, this is unacceptable as the patient's health is then in jeopardy. There are two possibilities: either we must re-order our values so that efficacy is preferred to one of safety or cost (respectively making EA1 or EA2 preferred), or else we must find an argument with which to defeat the attackers of one of EA1–3 and so reinstate one of our actions.

Suppose we re-order the values so as to prefer efficacy to at least one of the other values i.e., we must choose whether we disregard safety or cost. The choice will depend on the particular circumstances: it may be that the Drama agent is allowed to exceed budget if necessary, in which case efficacy will be preferred to cost and chlopidogrel will be prescribed. But if the cost constraint is rigid, there may be no better option than to disregard the contraindications and risk using aspirin, believing the complications to be less threatening than the immediate danger.

These hard choices can, however, be avoided if we can succeed in defeating one of the attacking arguments. We therefore run through our critical questions with respect to the arguments currently in the preferred extension of the framework. CQ3 can be posed with respect to EA4, as it is predicated on an assumption that there is no acid reducing therapy prescribed to the patient. We may therefore return to another Treatment KB and attempt to find such an acid reducing therapy. This will supply the knowledge that a



proton pump inhibitor (a particular type of acid reducing therapy) will have the desired effect. We can form this into EA7:

**EA7** *Where there are no contraindications, prescribing a proton pump inhibitor, will prevent excess acidity, removing risk of ulceration, promoting the value of safety.*

The complete argumentation framework after the addition of EA7 is shown in Figure 2:




Figure 2: Final Argumentation Framework showing all critiques.

Of course, EA7 is now subject to critical questioning. Assuming, however, that there are no alternatives, that it is not contraindicated, within budget and likely to be effective, the argument gathering stops here as we have now exhausted our critical questions. We compute the preferred extension by first including the arguments with no attackers: EA5, EA6 and EA7. EA7 defeats EA4 because they are motivated by the same value. This means that EA1 can be included, as its only attacker is defeated. EA1 thus defeats EA2 and EA3, again because they are motivated by the same value, and also excludes EA0, as desired. Note that in this case we need express no value preferences: the preferred extension is the same irrespective of value order. From this we conclude that aspirin is the preferred treatment, and should be recognised as such by any audience.

## 4 Discussion

The system for deliberative reasoning described above has a number of worthwhile features:

- It models deliberation using a model of argument with presumptive justification subject to critique, which has been developed to capture a number of features of practical reasoning observed in the philosophical [7] and informal logic literature [9]. These include the defeasible nature of putative solutions, the importance of perspectives (values), and the potential for context dependent orderings on perspectives to accommodate different audiences.
- This model is effected inside a single agent: the other components in the system can therefore be conventional knowledge and database systems, simplifying their participation in other systems. If, however, more sophisticated resources, such as autonomous agent systems, are available, these can be used by the Drama agent without modification.

- The various perspectives which need to be considered when making a medical decision are kept separate, and it is made explicit from which perspective the various arguments derive. This means that the perspectives can be given their due weight, but discounted if necessary.
- Each of the information sources used by the Drama agent is dedicated to the provision of particular information, has no need to consider every eventuality, and plays no part in the evaluation. This simplifies their construction and facilitates their reuse in other applications.
- Distinction can be made between information sources which are generic and those which are particular to a specific organisation or individual.
- Critiques are made only as and when they can affect the dialectical status of arguments already advanced. This means that all reasoning undertaken is of potential relevance to the solution.
- Patient information is made available only as and when needed.

The combination of the use of a well motivated model of deliberation, use wherever possible of conventional and generic components, and the ability to make flexible and context dependent decisions, provides, we believe, an approach to reasoning about decisions based on several information sources (such as is the case in medicine) that has considerable potential.<sup>4</sup>

## References

1. K. Atkinson. *What Should We Do?: Computational Representation of Persuasive Argument in Practical Reasoning*. PhD thesis, Department of Computer Science, University of Liverpool, Liverpool, UK, 2005.
2. T. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13 3:429–48, 2003.
3. C. Chesñevar, A. Maguitman, and R. Loui. Logical models of argument. *ACM Computing Surveys*, 32(4):337–383, 2000.
4. S. Doutre, T. Bench-Capon, and P. E. Dunne. Explaining preferences with argument position. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI 2005)*, pages 1560–1561, 2005.
5. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
6. C. Perelman and L. Olbrechts-Tyteca. *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press, Notre Dame, IN, USA, 1969.
7. J. R. Searle. *Rationality in Action*. MIT Press, Cambridge, MA, USA, 2001.
8. S. W. Tu and M. A. Musen. Representation formalisms and computational methods for modeling guideline-based patient care. In M. Mussen M. Stefanelli B. Heller, M. Loffler, editor, *Proceedings of First European Workshop on Computer-based Support for Clinical Guidelines and Protocols*, Leipzig, Germany, 2000. IOS Press.
9. D. N. Walton. *Argument Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.

<sup>4</sup> Trevor Bench-Capon and Sanjay Modgil acknowledge partial support received from the European Commission through Project ASPIC (IST-FP6-002307). We would like to thank the anonymous reviewers for their comments and suggestions. We are also grateful to Peter McBurney for his contributions to our work on argumentation and practical reasoning.