Gradually Intrusive Argumentative Agents for Diagnosis

Ioan Alfred Letia[1]

Technical University of Cluj-Napoca
Department of Computer Science
Baritiu 28, RO-400391 Cluj-Napoca, Romania
letia@cs.utcluj.ro

Abstract. A new model of agent for acting in processes of diagnosis where more constraints are placed on the performance in real world domains is presented. Preferences on actions and policies are used to induce a preferred course of action by the diagnosis agent. Past experience, gathered from controlled trials, enables the experience agent to offer counter-arguments for improving the diagnosis process. The outcome of actions are introduced in the scenario and agent models to allow argumentation on actions, not just beliefs. We argue that the approach is general and flexible, allowing the natural integration of new knowledgeable agents, while, at the same time, offering human agents transparency of the decision process.

1 Introduction

We have studied the problem of diagnosing the deep venous thrombosis (DVT) [11, 20] by argumentation using preference policies and ontology [13], as a new approach to attain evidence-based medicine with advice from controlled trials. We argue that the approach is general and flexible, allowing the natural integration of new knowledgeable agents, while, at the same time, offering human agents transparency of the decision process.

Although sometimes it is possible to make decisions in diagnosis taking into account just the information available [1], in the cases considered above several roles are involved in the medical decision process, a choice having to be made between the possible conflicting opinions of each role. Additional information may be needed for making that choice, due to the multiple options available.

Automated monitoring of medical protocols has been already tackled with a multi-agent system [2], using a negotiation process to mediate between multiple medical protocols, where role refers to a particular service that can be played by a staff person. A multi-agent environment to support training of diagnostic reasoning and modeling of domains with complex and uncertain knowledge [19] uses Bayesian networks to provide physicians probabilistic reasoning.

The framework described in this paper aims to help automate/model the decision making between different roles that are involved in the medical care process by making use of the knowledge base of each role and also by additional knowledge needed to solve conflicts of opinion between roles [13]. The decision making of agents is more refined than planning in the real world [16] since they are supposed to act in a sensible manner. The contribution of this paper is: (1) a new model of agent using bipolar outcomes in argumentation [3]; (2) an embedding of an action theory in most preferred trajectories [18]; (3) a diagnosis scenario, which is quite pervasive.

We present first the new model of agent with outcomes and action theory embedded, acting by considering outcomes and using argumentation to filter out situations that are not desirable, leaving only those that are preferred. The diagnosis of deep venous thrombosis (DVT) is selected as a running example, where committees of experts have developed rules for acting, and where there are also available controlled trials that capture some experience in the field. The scenario is then presented in a formal manner, which can be easily translated in an implementation, but can also be checked by human agents having proper knowledge. It is followed by a sample of what gradually intrusive acting can offer. After some discussion of related work, a few conclusions are drawn.

2 Agent Model

The knowledge of the *BAGO* agent is represented by the tuple $\langle \mathcal{B}, \mathcal{A}, \mathcal{G}, \mathcal{O} \rangle$ with $\mathcal{B}$: *beliefs* about the world, $\mathcal{A}$: *actions* the agent is capable to carry out, $\mathcal{G}$: agent *goals*, and $\mathcal{O}$: *outcomes* of the actions of the agent and/or events, which are bipolar, that is a tuple $\langle \mathcal{O}+, \mathcal{O}- \rangle$. The beliefs about the world are obtained in the processes of diagnosis by sensing actions, some of which may have a quite high cost and therefore might be acceptable only in exceptional situations.

The real world we consider is dynamic, in the sense that actions performed change certain predicates describing the world. Some of these changes are observable immediately (that is in the next state) while others might have other effects later in time. The real world is also quite often not completely deterministic and observability may itself be reduced in complex scenarios. However, we need to describe the world if we want to have our agents act in it with a reasonable chance of obtaining the desired outcomes. Fluents are used to express aspects of the world captured in the agent's model [15].

The fluents F used in the model have two subsets *inertial fluents* F_I and *non-inertial fluents* F_N, with $F = F_I \cup F_N$ and $F_I \cap F_N = \emptyset$ as in [18]. We express dynamic causality by rules of the form

$$f \xleftarrow{d} a, p_1, \ldots, p_n \tag{1}$$

with $f \in F$, $a \in A$, $p_i \in F$. They are used to describe the changes that take place when action a is executed by the agent in the situation when certain predicates p_i are true.

The static causality rules

$$f \xleftarrow{s} p_1, \ldots, p_n \tag{2}$$

where $f \in F_I$, $a \in A$, $p_i \in F$ express side-effects of actions, enabling us to capture in the action theory properties that are not intended, but which nevertheless occur in the real world when predicates p_i are true.

Executability conditions given by

$$a \xleftarrow{e} p_1, \ldots, p_n \tag{3}$$

with $a \in A$, $p_i \in F$ show that an action can be executed when predicates p_i are true. An action may be executed if the conditions are proper, including the presence of resources and the status of equipment (e.g. not faulty).

2.1 Preferences

The diagnosis process takes place on a time line and can be visualized as in the figure 1. In the state s_i the agents know the history H of the diagnosis process

$$\underbrace{s_0 a_0 s_1 \ldots a_{i-1}}_{\substack{H \\ \downarrow \\ \text{past}}} \; s_i \; \underbrace{a_i s_{i+1} \ldots s_{n-1} a_{n-1} s_n}_{\substack{F \\ \downarrow \\ \text{future}}}$$

Fig. 1. Course of action

and have to decide about the future F, the remaining course of actions so that the proper conclusion is reached.

In general, it is preferable for the agents to come up with a course of actions as short as possible. Another preference can be expressed in rules recommended by experts in the field. Although the goal $\mathcal{G}$ is used as an end result of a course of actions, the negative outcomes $\mathcal{O}-$ are to be avoided, if possible, or otherwise be kept to a reasonable minimum, given the goal. With various actors in such an environment we might also have different roles to accommodate [10] due to the organization in the particular field.

2.2 Acting

We assume that the task of the agents is to maintain some goal $\mathcal{G} = \{\neg o_{n_0}\}$ (the set could include several outcomes, not just one), e.g. to make sure that a certain situation does not occur. In the medical scenarios the agents should not endanger the health of the patient. To achieve this task the agents should find out whether some known condition o_p (or several) is the case for the patient, expressed in terms of the outcomes $\mathcal{O}+ = \{\neg o_p, o_p\}$. Negative outcomes include by default the negation of the goal, but also some other unwanted attributes like invasive procedures (which may be avoided) $\mathcal{O}- = \{o_{n_0}, o_{n1}\}$.

The agents have to make a decision about the selection of the most preferred action a_j when the patient is in the state s_i considering its possible next state s_k that might include an outcome o_k.

$$s_i \stackrel{a_j}{\rightarrow} s_k, \quad o_k \in s_k \tag{4}$$

An acceptable action is one with most preferred positive outcomes and no negative outcomes. Some of these outcomes are known from the causality rules (1), (2), but in some applications these rules cannot cover all the situations encountered in practice due to an incomplete description of the real world.

Controlled trials have been run in such domains to gather experience from the real world. They can be used to make sure that exceptions to the rules produce a reasonable and sensible treatment of the patients, even if they are not covered by the set of rules. If the trials show a state s_{ti} the same with the current state s_i with some negative outcome o_{ni}

$$s_{ti} \stackrel{a_{tj}}{\rightarrow} s_{tk}, \quad s_{ti} = s_i, o_{ni} \in s_{tk} \tag{5}$$

then this fact can be exploited in the decision making process to avoid an undesirable situation. We might also use a state s_{ti} in the trial if it is included in the current state.

$$s_{ti} \stackrel{a_{tj}}{\rightarrow} s_{tk}, \quad s_{ti} \subset s_i, o_{ni} \in s_{tk} \tag{6}$$

2.3 Argumentation

The approaches of planning in the real-world [16] are not acceptable in diagnosis problems of the kind considered in this paper since they are not concerned with preferences and a reduced intrusion. As we have both rules describing the recommended actions in an approximate model of the world as well as some experience gathered in the outcomes of trials we use an agent for each knowledge base. The agents come to a preferred conclusion regarding the decision on acting by argumentation in **extended logic programming** [17]. A proof is a sub-process of acting in which proponents and opponents attack each others arguments by undercuts (attack on argument's premise) and rebuts (attack on argument's conclusion).

An *x/y-dialog* [14] is a finite nonempty sequence of moves $move_i = \{Player_i, Arg_i\}$ (i>1) so that:

1. Player$_i$ = P (Proponent) if i is odd and Player$_i$ = O (Opponent) if i is even;
2. If Player$_i$ = Player$_j$ = P (i $\neq$ j) then Arg$_i$ $\neq$ Arg$_j$;
3. If Player$_i$ = P (i > 2) then (Arg$_i$, Arg$_{i-1}$) $\in$ y, and if Player$_i$ = O (i > 1) then (Arg$_i$, Arg$_{i-1}$) $\in$ x.

This scheme has the advantage of covering a hierarchy of notions of justifiability, which we aim to test.

The beliefs of the agents present in the state s_i are obtained by sensing actions, in diagnosis the results of the tests performed on the patient. The conflicts are resolved by the process of argumentation which may lead to several solutions if more alternatives are available or to no solution, depending on the rules. As our domain of application is a realistic one, we can envisage various uses of the framework: checking the rules for real scenarios, checking if more recent trials provide new results that might be incorporated in the rules, simulation for training, etc.

3 Diagnosis of Deep Venous Thrombosis

Suspected deep venous thrombosis (DVT) is a common condition with lifetime cumulative incidence of 2-5%, which untreated can result in a potentially fatal outcome. Accurate diagnosis of DVT minimizes the risk of complications and averts exposure of patients to the risk of anticoagulant therapy. An agent acting to achieve such goals must find a course of actions that minimizes the risk for the patient at each moment considering the equipment available for investigation, with its positive and negative predictive value. The clinical model for the diagnosis of deep venous thrombosis has been extensively studied, although we use here just the knowledge in [11, 20]. It is therefore a proper challenge problem for realistic acting.

3.1 Rules for diagnosis

Findings that are diagnostic of DVT (quoted from [11]) include the following.

> 1. Venous ultrasonography: Noncompressibility of the common femoral vein or popliteal vein (grade A). Noncompressibility that is confined to the superficial femoral vein, the distal portion of the popliteal vein, or the deep veins of the calf is associated with a lower predictive value (approximately 80%) and should be evaluated with venography (grade B).
> 2. Impedance plethysmography: Abnormal result and a high clinical suspicion of deep venous thrombosis (grade A). An abnormal result of impedance plethysmography combined with a moderate or low clinical suspicion of deep venous thrombosis should be evaluated with venous ultrasonography or venography (grade B).
> 3. Venography: Intraluminal filling defect seen in more than one view (grade A). Nonfilling of the deep veins despite repeated injection of contrast, although highly suggestive of deep venous thrombosis (grade C), must be interpreted in the light of clinical presentation and other investigations (such as results of impedance plethysmography or venous ultrasonography).

Similar rules express findings that exclude DVT (see [11]).

3.2 Controlled trials

An evaluation of D-dimer in the diagnosis of suspected DVT [20] has concluded that DVT can be ruled out in a patient who is judged clinically unlikely to have DVT and who has a negative D-dimer test, and therefore ultrasound testing can be safely omitted in such patients.

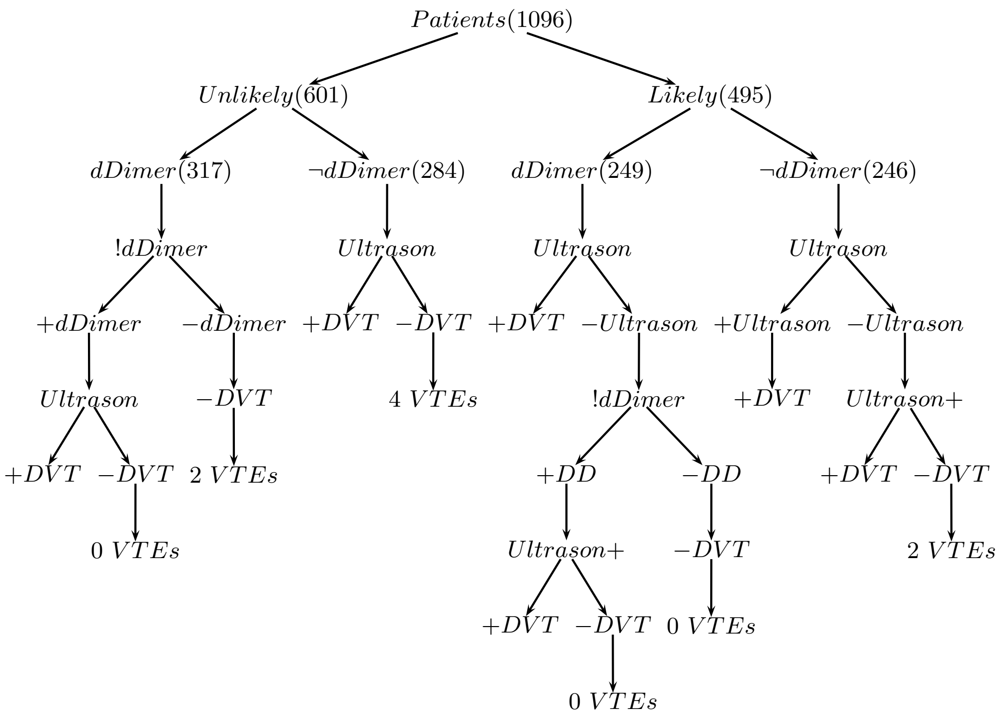

Fig. 2. Patient outcomes in the evaluation trial

The results of this controlled trial are shown in the figure 2 (see [20]). All patients were first evaluated using a clinical model and divided into two groups considered (clinically) unlikely or likely to have DVT. They were then randomly assigned either to undergo ultrasound imaging alone (control group) or to undergo D-dimer testing. Those in the later group then underwent ultrasound imaging (Ultrason) if they had been judged clinically likely or clinically unlikely to have DVT but the D-dimer test (!DD) was positive (+DD).

The primary outcome of this evaluation was the development of a venous thrombo-embolic event (VTE) in patients in whom DVT had initially been ruled out. They offer histories for re-evaluating the diagnosis model or, as we consider here, knowledge for an evaluation agent that can also provide cues for argumentation to the diagnosing agent. For example, we can see that there were two patients judged clinically unlikely who, with a negative D-dimer test (–DD), DVT was ruled out (–DVT) but still developed a venous thrombo-embolic event (2 VTE). For the patients judged clinically likely a second ultrasound imaging (Ultrason+) was performed after one week, which helped to reduce risk in the group with the D-dimer test.

4 Scenario Model

The set of actions $\mathcal{A} = \{c,\ u,\ d,\ v\}$ covers *clinical*, *ultrasonography*, D-*dimer* and *venography*, with the results of these actions expressed as beliefs $\mathcal{B} = \{cl(n),\ cl(l), cl(m), cl(h), ul(dP), ul(dC), ul(n), ul(p), dd(p), dd(n), ve(n), ve(p), ve(s)\}$. The effects of these actions are specified by dynamic causality

$$\mathcal{D}^1 = \begin{cases} d_0 : ve(X) \overset{d}{\leftarrow} v, venograph \\ d_1 : ul(X) \overset{d}{\leftarrow} u, ultrasonograph \\ d_2 : pl(X) \overset{d}{\leftarrow} p, plethysmograph \\ d_3 : cl(X) \overset{d}{\leftarrow} c \end{cases}$$

if the executability conditions are satisfied.

$$\mathcal{E}^1 = \begin{cases} e_0 : v \overset{e}{\leftarrow} suspicious, qualified \\ e_1 : u \overset{e}{\leftarrow} suspicious, qualified \\ e_2 : p \overset{e}{\leftarrow} suspicious \\ e_3 : c \overset{e}{\leftarrow} \end{cases}$$

Venography (v) as action can produce a result (e.g. $ve(p)$ – positive) if a venograph is available, qualified personnel is available and the patient is clinically *suspicious*.

A venous thrombo-embolic event (VTE) should be avoided at any time $\mathcal{G} = \{\neg vte\}$, by deciding whether the patient has DVT or not $\mathcal{O}+ = \{\neg dvt,\ dvt\}$, while also avoiding an invasive procedure, if possible, $\mathcal{O}- = \{invasive,\ vte\}$. Preferences over the situation of the patient are then:

$$\mathcal{P}^0 = \begin{cases} p_0 : \neg dvt \prec dvt \\ p_1 : dvt \prec vte \\ p_2 : invasive \prec vte \end{cases}$$

with static causality expressing suspicion and invasiveness.

$$\mathcal{S}^1 = \begin{cases} s_0 : invasive \overset{s}{\leftarrow} ve(X) \\ s_1 : suspicious \overset{s}{\leftarrow} cl(X), X \neq n \end{cases}$$

It is preferable for the patient not to have DVT, or to have the DVT diagnosed rather than have a venous thrombo-embolic event (VTE), or to perform an *invasive* test like venography rather than have a VTE.

4.1 Rules recommended by experts

For our AIM98 agent we show some rules recommending actions extracted from [11] (and presented in the sub-section 2.1).

$$\mathcal{R}^0 = \begin{cases} a_0 : v \leftarrow ul(sF) \vee ul(dP) \vee ul(dC) \\ a_1 : u \leftarrow pl(a) \wedge (cl(m) \vee cl(l) \\ a_2 : v \leftarrow pl(a) \wedge (cl(m) \vee cl(l) \\ a_3 : u \leftarrow ve(s) \\ a_4 : p \leftarrow ve(s) \end{cases}$$

The rules for a diagnostic of DVT from [11] are

$$\mathcal{R}^1 = \begin{cases} r_0 : dvt(a) \leftarrow ul(cF) \vee ul(p) \\ r_1 : dvt(b) \leftarrow ul(sF) \vee ul(dP) \vee ul(dC) \\ r_2 : dvt(a) \leftarrow pl(a), cl(h) \\ r_3 : dvt(b) \leftarrow pl(a), (cl(m) \vee cl(l)) \\ r_4 : dvt(a) \leftarrow ve(d) \\ r_5 : dvt(c) \leftarrow ve(s) \end{cases}$$

and similarly for those that exclude DVT.

$$\mathcal{R}^2 = \begin{cases} r_6 : \neg dvt(a) \leftarrow ve(n), cl(l) \\ r_7 : \neg dvt(c) \leftarrow dd(n) \\ r_8 : \neg dvt(b) \leftarrow pl(n), dd(n) \end{cases}$$

We also need the strength of the diagnostic rules given by experts as class A, B, C, expressed by preferences.

$$\mathcal{GR}^1 = \begin{cases} gr_0 : dvt(a) \prec dvt(b) \\ gr_1 : dvt(a) \prec dvt(c) \\ gr_2 : dvt(b) \prec dvt(c) \\ gr_3 : \neg dvt(a) \prec \neg dvt(b) \\ gr_4 : \neg dvt(a) \prec \neg dvt(c) \\ gr_5 : \neg dvt(b) \prec \neg dvt(c) \end{cases}$$

5 Gradually Intrusive Acting

Since reasoning about beliefs is different from reasoning about actions [5] we use the outcome as the expected value of actions.

Let us consider the case of a patient judged clinically low to which a plethysmography test revealed abnormality, while the ultrasonography has shown a normal case as shown in the figure 3. The process of argumentation [14, 17] is shown in the figure 4 for a significant patient. The *AIM98Agent* proposes to the *NEJM03Agent* the conclusion that state of the patient excludes DVT (-dvt), but in return the later answers with four thrombo-embolic events in his history (4 VTEs). Therefore the *AIM98Agent* proposes to perform venography but the

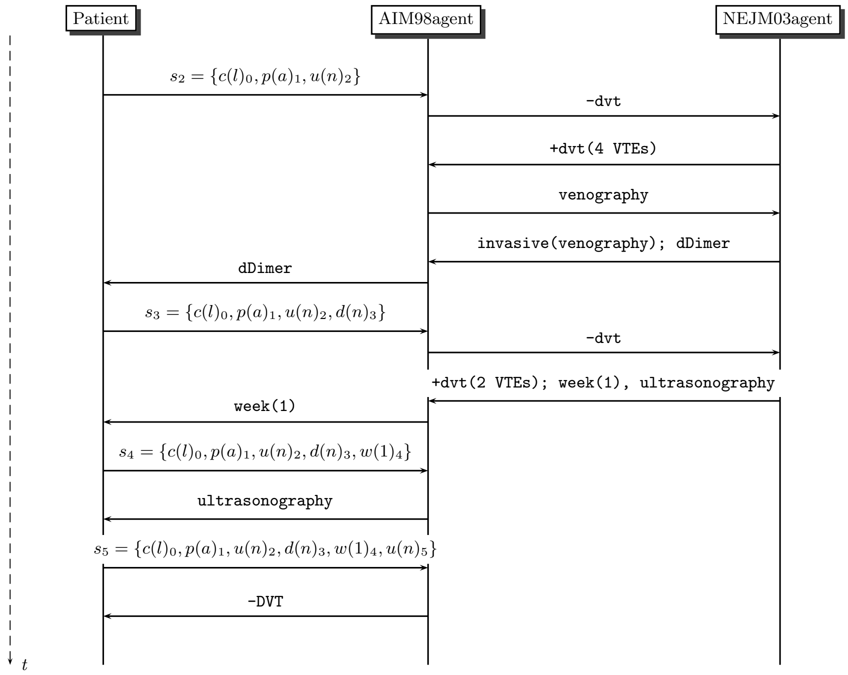

Fig. 3. Sequence diagram illustrating argumentation

NEJM03Agent disagrees on the account that it is invasive and proposes in return D-*dimer*, for which the former agrees. The new state s_4 shows a normal D-*dimer* test and again *AIM98Agent* is tempted to exclude DVT, but the other agent returns cases of VTEs and proposes ultrasonography after one week, which is accepted. After the normal situation in a week's time the conclusion reached by both agents is no DVT.

Let us consider the case of a patient judged clinically low to which a plethysmography test revealed abnormality, while the ultrasonography has shown a normal case. The AIM98Agent is tempted to decide no DVT, but to make sure is requesting the NEJM03Agent's opinion about it. The NEJM03Agent retrieves from its controlled trial history four such cases that have later developed VTEs, and therefore cannot accept this decision. The AIM98Agent proposes venography, but then NEJM03Agent replies that it is invasive and proposes instead d-*dimer*, to which AIM98Agent agrees.

Now in situation s_3, the AIM98Agent decides no DVT, letting NEJM03Agent know about the new proposal. Again, NEJM03 disagrees, this time mentioning two such cases in its historical representation of the controlled trial, and proposing instead ultrasonography after one week. The AIM98Agent agrees and requests the patient to come back after one week to have ultrasonography. On the new state s_5 it decides that the patient does not have DVT.

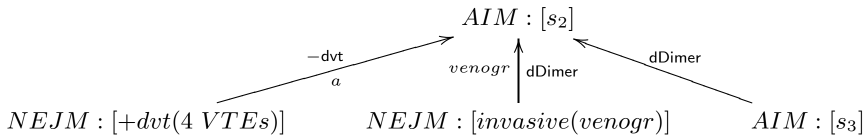

Fig. 4. Sample argumentation tree between the AIM98Agent and the NEJM03Agent

6 Related Work

As can be seen from the above excerpt on reasoning in our scenario to attain the goals and preferred outcomes we cannot use just causal reasoning and planning [6]. We also need a more refined model than the one acceptable in more frequent cases of reaching agreement [12] since the connection between available rules for acting and their outcome is not as clear in the diagnosis case we have considered.

In the previous implementation [13] we used the Gorgias system [10] to perform the argumentation within the AIM98Agent. In the current version an argumentation semantics for extended logic programming [17] with preferences [14] is used. We have also improved the model by explicitly representing the outcome of actions, which contributes to a more clear argumentation process.

Knowledge-base distributed search using teamwork [7] is indeed very convenient when knowledge is distributed as in our case. However, just searching the knowledge-base of the agents is not enough since there is a significant difference between argumentation over action and argumentation over beliefs [4]. Our agents need a theory of action that enables them to act and also sensing actions, otherwise a patient will be tested too many times, sometimes by invasive tests that may endanger life.

Arguments for reasoning about actions and values [8, 9] have been used to capture the results obtained by acting agents in practical reasoning applied to uncertain worlds. A theory of persuasion over action [4] also uses values to enable argumentation between actors with several attacks on proposals for action. We decided to enhance the agent model with preferences and an action theory in order to keep the representation closer to the real world for better transparency to human users.

Reasoning about actions and planning with preferences [18] covers preferences between actions and preferences between final states, but is not enough for the constraints of minimal intrusion that we need in the diagnosis of patients. Although a very nice and powerful theory for other domains, diagnosis with minimal intrusion is more restrictive.

7 Conclusions

The model proposed for both agents and the rules to be applied in practice (recommended by experts) is quite pervasive in the diagnosis domain. Due to uncertainty the link of outcomes to actions is not easy to express, but human agents are still able to make proper decisions in such an environment (if qualified enough). We can use this framework to build artificial agents, but also to evaluate the recommended rules vis-a-vis to the controlled trials (experience).

Since one goal is advising human agents in their decisions on acting in the real world our next step in this line of research will be on how argumentation on the actions could be further refined to better capture their effects on agents' objectives.

8 Acknowledgments

Part of this work was supported by the grant 423-33531 from the National Research Council of the Romanian Ministry for Education and Research. Many thanks are due to the reviewers for the input provided to the improvement of the paper.

References

1. Monica Acalovschi, Dan Blendea, Cristina Feier, Ioan Alfred Letia, Dumitru Dumitrascu, and Ana Veres. Risk factors for symptomatic gallstones in patients with liver cirrhosis: a case control study. *American Journal of Gastroenterology*, 98(8):1856–1860, 2003.
2. T. Alsinet, C. Ansotegui, R. Bejar, C. Fernandez, and F. Manya. Automated monitoring of medical protocols: a secure and distributed architecture. *Artificial Intelligence in Medicine*, 27:367–392, 2003.
3. Leila Amgoud and Souhila Kaci. On the generation of bipolar goals in argumentation-based negotiation. In I. Rahwan, P. Moraitis, and C. Reed, editors, *1st International Workshop on Argumentation in Multi-Agent Systems*, New York, NY, USA, 2004.
4. Katie Atkinson, Trevor Bench-Capon, and Peter McBurney. A dialogue game protocol for multi-agent argument over proposal for action. In *1st International Workshop on Argumentation in Multi-Agent Systems*, New York, NY, USA, 2004.
5. Daniela V. Carbogim, David Robertson, and John Lee. Argument-based applications to knowledge engineering. *The Knowledge Engineering Review*, 15(2):119–149, 2000.
6. James P. Delgrande, Torsten Schaub, and Hans Tompits. Domain-specific preferences for causal reasoning and planning. In D. Dubois and C. Welty, editors, *9th International Conference on Knowledge Representation and Reasoning (KR2004)*, Delta Whistler Resort, Canada, 2004.
7. Jörg Denzinger. Knowledge-based distributed search using teamwork. In V. Lesser, editor, *Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS-95)*, pages 81–88, San Francisco, CA, 1995. AAAI Press.

8. John Fox and Simon Parsons. On using arguments for reasoning about actions and values. In *AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, pages 55–63, 1997.
9. John Fox and Simon Parsons. Arguing about beliefs and actions. In M. Butler, L. Petre, and K. Sere, editors, *Applications of Uncertainty Formalisms*, LNCS 1455. Springer-Verlag, 1998.
10. Antonis Kakas and Pavlos Moraitis. Argumentation based decision making for autonomous agents. In J. Rosenchein, M. Wooldridge, T. Sandholm, and M. Yokoo, editors, *Second International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 883–890, Melbourne, Australia, 2003.
11. Clive Kearon, Jim A. Julian, Toni E. Newman, and Jeffrey S. Ginsberg. Noninvasive diagnosis of deep venous thrombosis. *Annals of Internal Medicine*, 128:663–677, 1998.
12. Sarit Kraus, Katia Sycara, and Amir Evenchik. Reaching agreement through argumentation: A logical model and implementation. *Artificial Intelligence*, 104:1–69, 1998.
13. Ioan Alfred Letia and Monica Acalovschi. Achieving competence by argumentation on rules for roles. In M.-P. Gleizes, A. Omicini, and F. Zambonelli, editors, *Engineering Societies in the Agents World*, LNAI 2934, pages 45–59. Springer-Verlag, 2005.
14. Henry Prakken and Giovanni Sartor. Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, 7:25–75, 1997.
15. Raymond Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, Cambridge, MA, 2001.
16. Stuart Russell and Peter Norvig. *Artificial Intelligence. A modern Approach*. Prentice-Hall, Englewood Cliffs, New Jersey, 2002.
17. Ralf Schweimeier and Michael Schroeder. A parameterised hierarchy of argumentation semantics for extended logic programming and its application to the well-founded semantics. *Theory and Practice of Logic Programming*, 5:207–242, 2005.
18. Tran Cao Son and Enrico Pontelli. Reasoning about actions and planning with preferences using prioritized default theory. *Computational Intelligence*, 20(2):358–404, 2004.
19. Rosa M. Vicari, Cecilia D. Flores, Andre M. Silvestre, Louise J. Seixas, Marcelo Ladeira, and Helder Coelho. A multi-agent intelligent environment for medical knowledge. *Artificial Intelligence in Medicine*, 27:335–366, 2003.
20. Philip S. Wells, David A. Anderson, Marc Rodger, Melissa Forgie, Clive Kearon, Jonathan Dreyer, George Kovacs, Michael Mitchell, and Bernard Lewandowski. Evaluation of D-dimer in the diagnosis of suspected deep-vein thrombosis. *The New England Journal of Medicine*, 349:1227–35, 2003.