

Mediation = Information revelation + Analogical reasoning

Simeon Simoff^{1,3}, Carles Sierra^{2,3} and Ramon López de Màntaras²

¹ School of Computing and Mathematics,
University of Western Sydney, NSW 1797, Australia s.simoff@uws.edu.au

² Institut d'Investigació en Intel·ligència Artificial – IIIA,
Spanish Scientific Research Council, CSIC
08193 Bellaterra, Catalonia, Spain {sierra, mantaras}@iia.csic.es

³ University of Technology, Sydney, Australia

Abstract. This paper presents an initial study of the relevant issues on the development of an automated mediation agent. The work is conducted within the ‘curious negotiator’ framework [1]. The paper demonstrates that mediation is a knowledge intensive process that integrates information revelation and analogical reasoning. The introduced formalism is used to demonstrate how via revealing the appropriate information and reshaping the set of issues of the disputing parties mediation can succeed. The paper presents MediaThor - a mediating agent that utilises past experiences and information from negotiating parties to mediate disputes and change the positions of negotiating parties.

1 Introduction

Negotiation is the process whereby two (or more) individual agents with conflicting interests interact, aiming at reaching a mutually beneficial agreement on a set of issues. Engaging in such interactions is a daily activity — from a simple negotiation on the price of a product we buy at the market to the complicated negotiations in dispute resolutions on the international arena. Whatever is the level of negotiation, during such interactions, participants may need to make concessions in order to reach an agreement [2].

Negotiation is goal-directed in the sense that individual agents involved in a negotiation may — probably will — have agendas of their own. But the agendas of the negotiating agents may be incompatible — there may be no solution that satisfies them all. Further the existence of a solution is unlikely to be known when the negotiation commences [3]. So it may not be useful to consider negotiation as a search problem because the solution space may be empty whilst the negotiating agents may believe that it is not so. If the negotiation is a multi-issue negotiation for which the issue set is open (i.e. it can change at any stage in the negotiation) then the agendas of the individual negotiating agents must necessarily be at a higher level than the issues because the issues are unknown, and may even be issues that ‘had never occurred’ to one of the agents. So for multi-issue negotiation the agendas of the agents cannot in general be a high level goal such as ‘to maximise profit on the deal’ as the deal space is unknown. Environmental conflict resolution is a typical example where conflicts involve many different

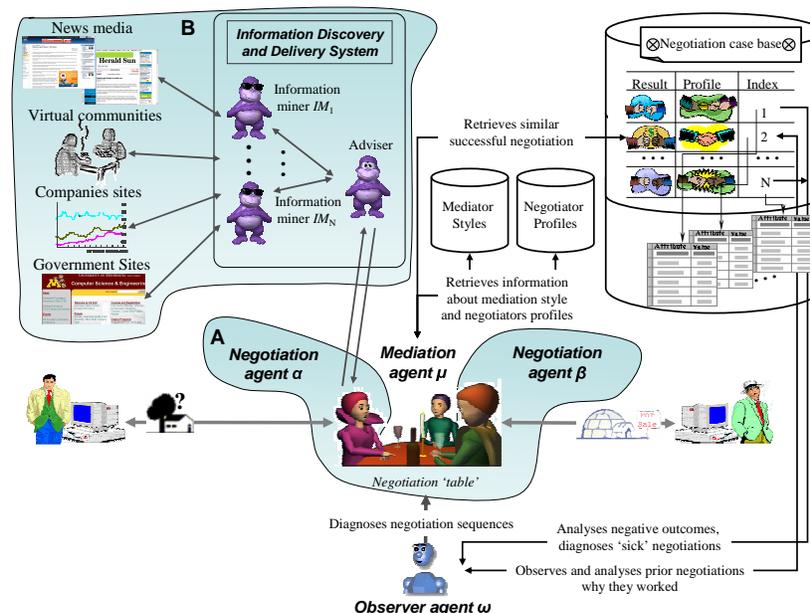


Fig. 1. The design of the 'curious negotiator' and the progress of the research

types of parties, issues and resources and the issue set can change during the process [4].

As a result negotiations may reach a deadlock, taking prohibitively long time without reaching tangible outcomes, or be terminated. This is when in real life the intervention of a mediator can influence the process, facilitating it towards a mutual agreement.

1.1 The 'curious negotiator' approach

The design of the 'curious negotiator' automated negotiation system, outlined initially in [1], is an attempt to address these issues. Figure 1 shows an updated version of the overall design proposed in [1] and the progress of the work.

The 'curious negotiator' is founded on the intuition "it's what you know that matters" and investigates the use of information and information theory, including entropy-based (random worlds) inference, as a foundation for automated negotiation between agents with bounded rationality. The design presented in [1] aimed at exploiting the interplay between contextual information [5] and the development of offers in negotiation conducted in an electronic environment. This contextual information is derived from what happens at the bargaining table and away from it. The work on the negotiation agent (shaded area A in Figure 1) focused on identifying mechanisms and knowledge structures for utilisation information in the negotiation process. Negotiation agent α negotiates with agent β by sending illocutions which represent offers and counter offers. The illocutions are represented in a communication language \mathcal{C} . An example of

such language, where the kernel set of negotiation illocutions is extended with illocutions that enable persuasive negotiation and argumentation, is presented in [6] and [7]. Negotiation agent α also uses an internal language \mathcal{L} for its reasoning.

Negotiation agent α negotiates from a stance that assumes nothing about her opponent's motivations and applies maximum entropy logic to that which it has observed. The basic feasibility of this approach was reported in [8]. The research outcomes under the shaded area A in Figure 1 are known collectively as information-based agency [7], which treats negotiation as an information discovery and revelation process. The information that the agent utilises may come from at least two sources:

- from the ‘negotiation table’, e.g. from all utterances agents make during a negotiation (this is incorporated in the work presented in [8]);
- from external sources, e.g. other deals, news, companies white papers, blogs of virtual communities, and other electronically accessible sources, all of which constitute part of the context in which negotiation happens.

The automation of the discovery, representation and delivery of the necessary information and knowledge to the agents has been the focus of the work on the information discovery and delivery system that services negotiating agents in the ‘curious negotiator’ (shaded area B in Figure 1). Technical aspects of the embedded information mining system include: (i) effective automated techniques for extracting relevant articles from news web sites and their semi-structured representation suitable for further information extraction [9]; (ii) techniques for automatic utilisation of extracted information and providing recommendations to the negotiation agents in requested form (for example, refined exchange rate predictions [10] that utilise information from latest financial news; recommendations on product choice based on information extracted from consumer reviews [11]).

1.2 The missing component

The mechanisms for dealing with negotiations that fail in reaching an agreement, or seemed to be leading to a failure, remain the undeveloped part of the ‘curious negotiator’. It is indicated by the unshaded part in Figure 1, which includes the mediating agent μ , the observer agent ω and their supporting knowledge representation structures.

This paper presents the initial work on the principles of building an automated mediation agent within the ‘curious negotiator’ framework. We explore mediation as a process of intertwined information revelation and analogical reasoning, which incorporates past experiences. To model mediation the work draws from the research in dispute resolution. It specifies the requirements towards the knowledge representation structures supporting mediation. Section 2 looks at mediation, as a knowledge-driven process and explores the changes that information revelation can make to the negotiation space and the outcomes of negotiation. It introduces the notion of ‘mental model’ of participants involved in the process and looks at mechanisms of how these models can be utilised in automated mediation. Section 3 considers some aspects in utilising past experiences and background knowledge in automated mediation. It looks also at the utilisation of information at the diagnosis stage. Section 4 presents the architecture of MediaThor - a

mediation agent that illustrates the computational implementation of mediation mechanisms discussed in the previous sections.

2 Mediation as a knowledge driven process of information revelation.

Contemporary analysts in social and political sciences look at mediation as a process that enables conflict resolution. Mediators are often indispensable in the area of *dispute (or conflict) resolutions*, settling variety of disputes, spanning from conflicts between sovereign nations to conflicts between family members, friends, and colleagues. Successful mediation can make a dramatic difference to the outcome of a negotiation stalemate. For instance, on 14 January 1998 the President of United Nations Security Council issued a statement demanding “that Iraq cooperate fully and immediately and without conditions with the Special Commission in accordance with the relevant resolutions.”⁴ As all UN weapons inspections in Iraq were frozen, during the following month all direct negotiations between the US and Iraq did not reach any agreement and the military conflict seemed unavoidable. The following event sequence illustrates the mediation process: (i) the US authorised the mediation effort; (ii) the UN secretary (the mediator) achieved a possible deal with Iraq; (iii) the UN secretary passed it back to the US; (iv) the US reviewed and accepted the deal. Several months later the conflict escalated, but this time no mediation was sought and military actions started. The mediation made a huge difference in the first dispute resolution.

2.1 Necessary and sufficient conditions for mediation

This example illustrates that mediation as a process involves *information revelation* and part of the mediator’s strategy is guiding the process of information revelation. The following are the necessary (C1, C2) and sufficient (C3) conditions for a mediation to take place:

- *Condition C1*: Negotiating agents α and β are willing to achieve a mutually beneficial agreement;
- *Condition C2*: Negotiating agents α and β are seeking or will accept mediation (in the first case, the awareness about the conflict and the problem with the current state of the negotiation resides with the negotiating agents, in the second case either the mediator agent μ or, if present, the observer agent ω diagnoses the problem);
- *Condition C3*: A mediating agent μ is available (this condition is by default embedded in the ‘curious negotiator’ paradigm).

In the example with the 1998 Iraq crisis, in the second case condition C2 was not present. Conflicts may be a result of a contradiction of interests, as in the example with the 1998 Iraqi crisis, but can be also a result just of a different (but unknown to the disputing parties) perception of the disputed subject.

⁴ <http://daccessdds.un.org/doc/UNDOC/GEN/N98/007/84/PDF/N9800784.pdf>

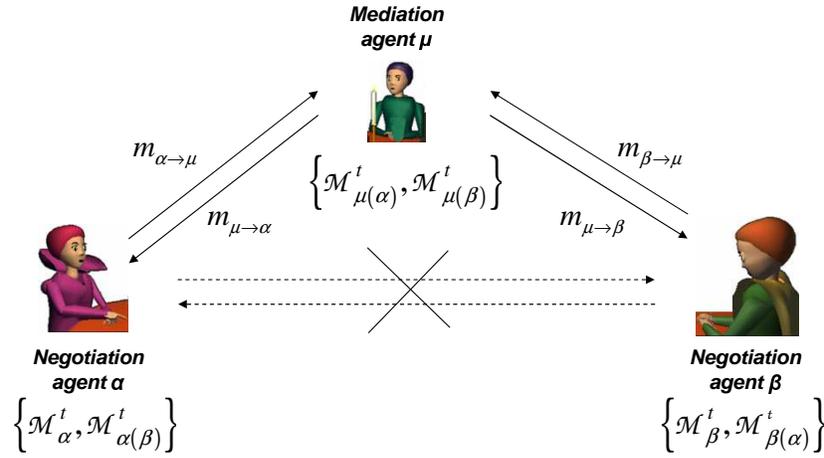


Fig. 2. The mediation agent within the ‘curious negotiator’ framework

2.2 Mediation process within the ‘curious negotiator’ framework

Further we consider the following mediation process, illustrated in Figure 2, where agents α and β are in a deadlock and direct exchange of offers between them has ceased. In a mediation session, α and β interact with messages m only with the mediating agent μ .

\mathcal{M}^t denotes a “mental model” at time t . We use the label “mental model” to denote the view (including related knowledge) of an agent about a dispute, about the views of the other parties on that dispute and the expected outcomes. This knowledge is internal to the agent. Each model is manifested to the other agents through the actions taken by the agent. The label “mental models” has been chosen to emphasise the key role of the mediator in the careful examination of the way negotiation parties have built their views on the disputed issues [12]. It is also in accordance with the view that negotiation can be conceptualised as a problem-solving enterprise in which mental models guide the behaviour of negotiating parties [13]. Further in the text we use the term mental model without quotation marks.

\mathcal{M}_α^t and \mathcal{M}_β^t denote the mental models of agents α and β , respectively. \mathcal{M}_α^t is not known by β and \mathcal{M}_β^t is not known by α . None of them is known by the mediating agent μ . Each of these agents has its own approximations of the mental models of the other agents. $\mathcal{M}_{agent(party)}^t$ denotes the mental model that the *agent* has about another *party*. In particular, $\mathcal{M}_{\alpha(\beta)}^t$ is the mental model of α about β , i.e. about what β wants out of the negotiation; respectively, $\mathcal{M}_{\beta(\alpha)}^t$ is the mental model of β about α , i.e. the position of α in the dispute. Further, $\mathcal{M}_{\mu(\alpha)}^t$ and $\mathcal{M}_{\mu(\beta)}^t$ are the mental models of the mediating agent μ about the positions of α and β in the dispute, respectively.

We use the above formalism to demonstrate some aspects of mediation that need to be taken into account when developing automated mediators. Further we use two

examples - The Orange Dispute [14] and the Sinai Peninsula Dispute to illustrate the role of information revelation and identification of analogy between disputes in order to reshape the set of issues and complete the mediation process.

2.3 The Orange Dispute - reshaping the problem based on additional information

In the Orange Dispute [14], two sisters want the same orange. According to Kolodner [14] “MEDIATOR assumes they both want to eat it and solves the problem by having one sister cut the orange in two and the second chooses her half. When the second sister uses her peel for baking and throws away the pulp, MEDIATOR realises it made a mistake.”⁵

Further we present the two mediation attempts in terms of the agreements reached and the information that can be passed to the mediator. Lets our agent α represent the first sister who wants to have the orange as a desert and agent β represent the second sister who wants to have (only the peel of) the orange for cooking (the recipe requires the whole peel). If our mediation agent μ happens to be the case-based MEDIATOR, then the situation described in the Orange Dispute can be expressed through the mental models of the individual participants in Figure 3, making explicit the wrong assumption (the boxed expressions in Figure 3).

$$\alpha \text{ wants the orange as a dessert} \in \mathcal{M}_{\alpha}^t \quad (1)$$

$$\beta \text{ wants the peel of the orange for cooking} \in \mathcal{M}_{\beta}^t \quad (2)$$

$$\beta \text{ wants an orange} \in \mathcal{M}_{\alpha(\beta)}^{t_{break}} \quad (3)$$

$$\alpha \text{ wants an orange} \in \mathcal{M}_{\beta(\alpha)}^{t_{break}} \quad (4)$$

$$\boxed{\alpha \text{ wants the orange as a dessert}} \in \mathcal{M}_{\mu(\alpha)}^{t_{start}} \quad (5)$$

$$\boxed{\beta \text{ wants the orange as a dessert}} \in \mathcal{M}_{\mu(\beta)}^{t_{start}} \quad (6)$$

Fig. 3. The wrong initial assumption of the MEDIATOR [15] in terms of our mental models (Boxed expressions). This initial assumption (which didn't change as there were no mechanisms for that) caused the failure of that mediator.

In these models t_{break} , and t_{start} indicate the time when negotiation broke and when mediation started, respectively (in the case of the MEDIATOR it has been a one step act). The results of the agreements in terms of the outcomes - *Outcome (agent, issue, result)* are presented in Table 1, where result values are denoted as follows: “+”, “+/-” and “-” for positive, acceptable, and negative, respectively for the corresponding

⁵ MEDIATOR [15] is one of the early case-based mediators. The focus of the work was on the use of case-based reasoning for problem understanding, solution generation, and failure recovery. The failure recovery ability is demonstrated with the Orange Dispute in [14].

agents in terms of the negotiated issue. In the original example [14], the result in the outcome for β should be “+/-” as the second sister still used the peel from her half. Here we added the constraint of the recipe in order to get negotiation about the orange to a halt with an unacceptable “-” result and generate a request for mediation.

The Orange Dispute can be considered an example of a dispute over resource scarcity. The resource in this case has a possible component-based separation (without change of the total amount of available resource) that allows to change the structure of the dispute through mediation, opening the space for a mutually beneficial solution. It exposes two aspects of mediation:

- The difference that a mediator can bring is in exploring the structure of the problem from a broader stance;
- An initial assumption by a mediator can lead to a failure of the mediation effort.

Consequently, we formulate the following postulates for the automated mediator:

- *Postulate P1*: An automated mediator μ should start interaction with extracting more information about the position of the parties on the negotiation;
- *Postulate P2*: An automated mediator should develop an independent “grand view” of the problem, which is more comprehensive than the individual views of α and β , respectively.;
- *Postulate P3*: An automated mediator μ should operate from the initial stance only of conditions C1 and C2.

Starting mediation without initial assumptions means that μ either does not have a model for each of the negotiating agents α and β , or accepts the models $\mathcal{M}_{\alpha(\beta)}^{tbreak}$ and $\mathcal{M}_{\beta(\alpha)}^{tbreak}$ these agents have about each other at the point of requesting mediation. In the case of the Orange Dispute, μ starts mediation with the exit models of α and β :

- $\mathcal{M}_{\mu(\alpha)}^{tstart} = \mathcal{M}_{\beta(\alpha)}^{tbreak}$, i.e. α wants an orange $\in \mathcal{M}_{\mu(\alpha)}^{tstart}$, and
- $\mathcal{M}_{\mu(\beta)}^{tstart} = \mathcal{M}_{\alpha(\beta)}^{tbreak}$, i.e. β wants an orange $\in \mathcal{M}_{\mu(\beta)}^{tstart}$.

This information is not sufficient for mediation, e.g. the uncertainty in the mutual models of α and β , and the model μ are the same. Research in conflict resolution in international relations demonstrates that if a mediator could credibly add information to the system of negotiators this alters the state of the system [16]. Consequently, μ takes steps in order to decrease this uncertainty. In addition, intuitively, it seems worth checking whether both parties have the same understanding of the issues in the dispute, i.e. technically, whether they operate with the same ontology or with compatible ontologies. In the Orange Dispute, μ obtains from each party what the orange is needed

Agent	Agreement clauses	Outcome for α	Outcome for β
α	Cuts the orange into halves	Outcome(α , has orange, +/-)	Outcome(β , has orange, -)
β	chooses one half	Outcome(α , has orange, +/-)	Outcome(β , has orange, -)

Table 1. Outcomes of the Orange Dispute, based on mediation with initial assumption.

for. The Orange Dispute in terms of the mental models of the individual participants in the case of proposed mediation agent is presented in Figure 4. In these models t_{break} , t_{start} and t_{end} indicate the time when negotiation broke and when mediation started and ended, respectively. Note the difference of $\mathcal{M}_{\mu(\cdot)}^{t_{start}}$ for both α and β in Figure 3 and Figure 4. The steps taken by the mediating agent are described in Figure 5 (we do not use a formal illocution based language, but the actions that the language should cater for are shown in italic).

$$\alpha \text{ wants the orange as a dessert} \in \mathcal{M}_{\alpha}^t \quad (7)$$

$$\beta \text{ wants the peel of the orange for cooking} \in \mathcal{M}_{\beta}^t \quad (8)$$

$$\beta \text{ wants an orange} \in \mathcal{M}_{\alpha(\beta)}^{t_{break}} \quad (9)$$

$$\alpha \text{ wants an orange} \in \mathcal{M}_{\beta(\alpha)}^{t_{break}} \quad (10)$$

$$\boxed{\alpha \text{ wants an orange}} \in \mathcal{M}_{\mu(\alpha)}^{t_{start}} \quad (11)$$

$$\boxed{\beta \text{ wants an orange}} \in \mathcal{M}_{\mu(\beta)}^{t_{start}} \quad (12)$$

$$\boxed{\alpha \text{ wants the orange as a dessert}} \in \mathcal{M}_{\mu(\alpha)}^{t_{end}} \quad (13)$$

$$\boxed{\beta \text{ wants the peel of the orange for cooking}} \in \mathcal{M}_{\mu(\beta)}^{t_{end}} \quad (14)$$

Fig. 4. The respective mental models of α , β and μ in the mediation session of the Orange Dispute with our proposed agent.

The Orange Dispute illustrates also another important ability that an automated mediator should possess — the ability to reshape or restructure the dispute, based on the additional information about the models of each party. The outcomes of the restructured Orange Dispute are shown in Table 2.

2.4 The Sinai Peninsula Dispute and its analogy with the Orange Dispute

The ability to reshape the problem is crucial for developing successful automated mediators. The Sinai Peninsula Dispute in the area of international relations shows similar properties to the Orange Dispute. The Sinai Peninsula is a piece of land of about 62,000 square km that separates Israel and Egypt. With its landscape Sinai has a *military value* for either side in terms of mechanised infantry transport or as a shelter for guerrilla

Agent	Agreement clauses	Outcome for α	Outcome for β
α	Peels the orange	Outcome(α , eat, +)	Outcome(β , cook, +)
β	Gets the whole peel	Outcome(α , eat, +)	Outcome(β , cook, +)

Table 2. Outcomes of the restructured Orange Dispute.

1. μ : *ask* α to *send* its ontology of the negotiated item (orange).
2. μ : *ask* β to *send* its ontology of the negotiated item (orange).
3. μ : *compare* ontologies received from α and β .
4. μ : *send* α and β agreed ontology (orange as a fruit which has pulp and peel).
5. μ : *ask* α to *send* μ its preferences on the negotiated item in terms of agreed ontology.
6. μ : *ask* β to *send* μ its preferences on the negotiated item in terms of agreed ontology.
7. μ : *advises* α and β on \mathcal{M}_α^t and \mathcal{M}_β^t based on their preferences
8. μ : *checks* the case base for past cases (resource disputes)
9. μ : *retrieves* resource disputes with divisible components
10. μ : *sends* α and β action separate resource (peel the orange)
11. μ : *tells* α and β to complete negotiation.
12. μ : mediation completed.

Fig. 5. Mediation as information revelation aiming at decreasing uncertainty within the negotiation system

forces. The perceived importance of the territory is evidenced by the fact that Israelis and Egyptians fought in or over the Sinai Peninsula in 1948, 1956, 1967, 1968-1970, and 1973. Since 1967 Sinai had been occupied by Israel. Figure 6 shows a very simplified version of the models of the parties at the initial meeting in Jerusalem, when the negotiations started and halted and the change of the mediators models that lead to the outcomes. For the purpose of this paper we aim to emphasise the high level analogy with the Orange Dispute case (see Figure 4), i.e. the need for a mediator to reframe the problem. In fact, the need for restructuring the problem in order for a mediator to get a bigger picture has been recognised in PERSUADER [2], to resolve labor-management disputes. In recent works [17] the mediator is expected to have a complete knowledge of the solution space.

$$\begin{aligned}
& \alpha \text{ wants security, support for economy, recognition} \in \mathcal{M}_\alpha^t \\
& \beta \text{ wants sovereignty (restored territory), support for economy, security} \in \mathcal{M}_\beta^t \\
& \beta \text{ wants territory and strategic military advantage} \in \mathcal{M}_{\alpha(\beta)}^{tbreak} \\
& \alpha \text{ wants territory and strategic military advantage} \in \mathcal{M}_{\beta(\alpha)}^{tbreak} \\
& \boxed{\alpha \text{ wants territory and strategic military advantage}} \in \mathcal{M}_{\mu(\alpha)}^{tstart} \\
& \boxed{\beta \text{ wants territory and strategic military advantage}} \in \mathcal{M}_{\mu(\beta)}^{tstart} \\
& \boxed{\alpha \text{ wants security, support for economy, recognition}} \in \mathcal{M}_{\mu(\alpha)}^{tend} \\
& \boxed{\beta \text{ wants sovereignty (restored territory), support for economy, security}} \in \mathcal{M}_{\mu(\beta)}^{tend}
\end{aligned}$$

Fig. 6. The respective mental models of α , β and μ in the mediation session of the Sinai Dispute with our proposed agent

Following the initial interaction in Jerusalem, the US President Jimmy Carter initiated a *third-party mediation* effort that culminated in the Camp David accords. For the purposes of this paper we consider a simplified version of the second agreement of the Camp David accords on the future of the Sinai Peninsula. The items in the agreement are presented in Table 3, in a structure, similar to the presentation of the agreements in the Orange Dispute in Table 1 and Table 2. Without getting into the details of the mediation steps, from Table 3 it is evidenced that the initial mutually perceived models and about the need for territory and strategic military advantage have been transformed by the mediation into a Security/Sovereignty trade-off, with economic benefits.

The analogy with the Orange Dispute is in having the initial negotiation framed around a common resource Territory and a similar issue of having strategic military advantage as the main goals that can enable the security. Though both territorial and military components remain on the negotiation table, based on some background knowledge and higher level view of the ontology of the problem, the mediator developed a view of the set of issues and aligned the ontologies of both parties which eventually changed their models: security and restoration may not necessarily be achieved with occupation of a territory or with expensive military presence.

The information injected by the mediator and proposed steps leads to decreasing the differences between perceived mental models $\mathcal{M}_{\alpha(\beta)}^t$ and $\mathcal{M}_{\beta(\alpha)}^t$, and the corresponding actual mental models \mathcal{M}_{β}^t and \mathcal{M}_{α}^t of agents α and β , respectively, i.e. the intervention of the mediator decreases the uncertainty in the negotiation system.

Agent	Agreement clauses	Outcome for α	Outcome for β
α	withdraw its armed forces from the Sinai	Outcome(α , Military, -)	Outcome(β , Territory, +) Outcome(β , Sovereignty, +)
α	Evacuate its 4500 civilians	Outcome(α , Territory, -)	Outcome(β , Territory, +) Outcome(β , Sovereignty, +)
α	Restore Sinai to Egypt	Outcome(α , Territory, -)	Outcome(β , Territory, +) Outcome(β , Sovereignty, +)
α	Limit its forces within 3km from Egyptian Border	Outcome(α , Military, -) Outcome(α , Security, +)	Outcome(β , Security, +)
α	Lost the Abu-Rudeis oil fields in Western Sinai	Outcome(α , Economy, -)	Outcome(β , Economy, +)
β	Normal diplomatic relations with Israel	Outcome(α , Recognition, +)	Outcome(β , Security, +)
β	Freedom of passage through Suez Canal	Outcome(α , Economy, +) Outcome(α , Security, +)	Outcome(β , Security, +)
β	Freedom of passage through nearby waters	Outcome(α , Economy, +) Outcome(α , Security, +)	Outcome(β , Economy, +) Outcome(β , Security, +)
β	Restricted Egyptian forces in Sinai	Outcome(α , Security, +)	Outcome(β , Military, -) Outcome(β , Security, +)

Table 3. The Sinai Peninsula Dispute. α denotes Israel; β denotes Egypt.

3 Utilising past experiences and background knowledge in automated mediation

The American Bar Association defines mediation as a process by which those who have a dispute, misunderstanding or conflict come together and, with *the assistance of a trained neutral mediator*, resolve the issues and problems in a way that meets the needs and interests of both parties.⁶ This definition emphasises the key role of the past experience of the mediator and its unbiased nature. Further, we consider these two aspects, starting with mediator bias.

3.1 Unbiased mediator

The *bias of a mediator* is defined as the presence of a preference towards one of the

- outcomes in the negotiation; or,
- sides involved in the negotiation.

Not having preference towards any of the outcomes of a negotiation means also to keep open all options. For instance, the peace-loving broker's bias towards peaceful solutions makes his or her claims less believable compared to a broker who is indifferent to war or peace [16]. Such bias as a result can decrease the effectiveness of the mediation effort. Protecting automated mediation from introduction of a bias is not seen as a problem.

3.2 Utilising past experiences

Experience is, perhaps, the distinct feature between successful and less successful mediators. Analogical reasoning (CBR + ontology) is an approach to problem solving that emphasizes the role of prior experience during future problem solving (i.e., new problems are solved by reusing and if necessary adapting the solutions to similar problems that were solved in the past) (see [18] for a recent review of the state-of-the-art in the CBR field). From a machine learning point of view, updating the case base is a lazy learning approach (i.e. learning without generalisation). Some aspects of using the past experience by the tandem Mediation and Observation agents have been discussed in [1]. In terms of required case representation, a starting point is the knowledge representation structure for representing negotiation cases, proposed in [19]. This structure needs to be updated for dealing with ontologies. For the mediation, the case based will be linked to the corresponding knowledge base of the mediation strategies used. The case structure now includes a negotiation case as its problem section and the collection of mediation steps, information used and other knowledge, as the solution part of the case.

Important from a computational perspective is the diagnosis stage of the mediation process [20]. The diagnostic function consists of monitoring the progress of negotiation or related interactions intended to settle or resolve disputed issues (Druckman and co-authors [20] refer to [21]). Monitoring provides a long view of unfolding developments,

⁶ <http://www.abanet.org/cpr/clientpro/medpreface.html>

including trends in escalating and de-escalating dynamics. Within the framework of ‘curious negotiator’ we consider this stage as a pre-mediation stage, which is executed by the observer agent ω . To some extent it resembles similarity with OLAP⁷ — the pre-data mining steps in business intelligence, where summary statistics at different levels are generated and later provide guidance to the data mining strategies. Similar to OLAP, monitoring should be able to provide snapshots of the negotiation process at any moment of time at different levels of granularity. The mediator μ should be able to estimate the difference between $\mathcal{M}_{\alpha(\beta)}^t$ and $\mathcal{M}_{\beta(\alpha)}^t$ from the respective actual mental models \mathcal{M}_{β}^t and \mathcal{M}_{α}^t in order to define the intervention time of mediating interventions (if we follow a proactive approach and intervene before negotiation fails).

4 MediaThor: A powerful CBR mediator agent

The architecture of a mediator agent as described in the previous sections has to be based on a clear understanding of the relationships between its models of the agents and the ontology that describes the problem and its solution. The retrieval of a previous case, described in completely different ontological terms, has to rely on some alignment process between the concepts and semantic relations of that past case and the concepts and relations in the problem being solved. Figure 7 shows a representation of four cases. Each case consists of an ontology representing the problem, the agents participating, the goals each agent had, and the solution found. On Figure 7-A and 7-B one can see a representation of the orange and Sinai peninsula disputes. We consider two types of ontology matching in this architecture. A structural matching that aims at connecting the nodes in the ontology from a structural point of view (arrows in blue in the picture — or dark grey for the B&W printers) and from a semantic point of view (arrows in green — or in light grey). The example would account for a perfect matching from a structural point of view and for an almost perfect matching from a semantic point of view (up to the semantic similarity between *required* and *brings*.)

Thus, assuming that the set of all possible ontologies⁸ is noted by \mathcal{O} , we will therefore require that any mediator is provided with a $Match : \mathcal{O} \times \mathcal{O} \rightarrow [0, 1]$ function that determines the level of similarity between any two ontologies (graphs). This function will be a combination of at least structural and semantic similarities between the ontologies. A number of existing solutions can be found in [22]. We will also assume that the ontology of a case is indeed a partial view (subgraph) of a more general ontology. Thus, given an ontology $o \in \mathcal{O}$ we will note by \bar{o} this general ontology. In our example, the orange case could be seen as a view of an ontology of fruits and their usage, while the Sinai case could be seen as a view of an ontology on military affairs. We will therefore associate to each case the ontology from where the case is drawn and moreover we will define a subsumption relation between ontologies, $\sqsubseteq \subseteq \mathcal{O} \times \mathcal{O}$, being $o \sqsubseteq o'$ true when o is a sub-ontology of o' , i.e. a subgraph. Clearly $o_i \sqsubseteq \bar{o}_i$

The concepts in the ontologies will all be understood as resources, and we abuse notation by transforming them into predicates with the following intuitive meaning:

⁷ Online analytical processing.

⁸ In this paper we will take the simplified view that ontologies are labelled directed graphs, with concepts at the nodes and binary relations on the links.

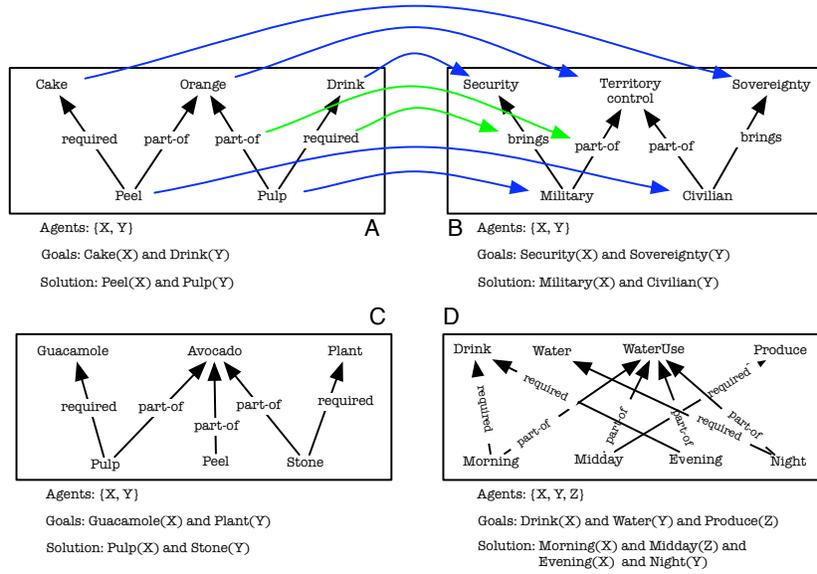


Fig. 7. Four cases represented as: an ontology (a graph), the agents participating and the goals and solution of the problem. Blue arrows represent a structural matching between the ontologies of cases A and B, while green arrows represent a semantic matching between the same cases.

$Security(X)$ means that agent X gets $Security$. The difference between goals and solutions in the cases have to be understood from a planning perspective: a solution is an assignment of resources to agents that logically imply their respective goals.

According to our view above, *MediaThor* will take a very specific view on what $\mathcal{M}_{\mu(\cdot)}^t$ is. In particular we will assume that $\mathcal{M}_{\mu(\alpha)}^t = \langle o_\alpha, g_\alpha, r_\alpha \rangle$ where o_α , g_α , and r_α is the perceived ontology, goals, and reservations of α . We understand by reservations those constraints that α requires the solution to satisfy. *MediaThor* is equipped with a predicate $Sat \subseteq \mathcal{C} \times 2^{\mathcal{R}}$ that is true when a particular case satisfies a set of reservations.

We now follow the steps of a CBR algorithm to describe the architecture of *MediaThor*.

4.1 Retrieval

MediaThor has a memory of past cases $C = \{c_i\}_{i \in I}$ where $c_i = \langle A_i, o_i, G_i, S_i \rangle$, A_i the set of agents involved, o_i an ontology, G_i the goals of the agents participating in the case, and S_i the solution as a set or resource assignments. *MediaThor* also is dealing with agents α and β trying to find a solution to their problem. The model that *MediaThor* has at any moment of time is a partial case where the ontology may be incomplete, and the goals may be wrong. Thus, the current problem is a case c using ontology o .

MediaTor, at retrieval time, looks for cases that are similar to the current case c . *MediaTor* uses a sliding parameter η that filters those cases c_i that are expressed in ontologies o_i that are semantically similar to o up to the value η . By setting $\eta = 1$ we are considering cases exactly over the same ontology, and by decreasing η we consider increasingly farther away cases in terms of semantic similarity.

Figure 7 shows that there is a structural mapping between cases B and C, although in this case C contains a richer structure. Ontology matchers provide a degree of matching but also a function that maps the concepts and labels of one ontology into the concepts and labels of the other. Thus, if the mapping function from B to C is called $f_{B \rightarrow C}$ it is clear that $f_{B \rightarrow C}(o_B) \sqsubseteq o_C$ and that the degree of structural matching between A and C has to be smaller than between A and B. Thus assuming the semantic matching (S-Match) is the same between A and B and between C and B, we can conclude that if B is the current case *MediaThor* would prefer case A to case C at retrieval time.

Thus, for a given η , $\mathcal{M}_{\mu(\cdot)}^t$, current case $c = \langle \{\alpha, \beta\}, o, G_\alpha \cup G_\beta, \emptyset \rangle$ and case base C , the best case $c^* \in C$ is defined as:

$$c^* = \arg \max_{\substack{c_i \in C \\ S\text{-Match}(\bar{o}_i, \bar{o}) \leq \eta \\ Sat(c_i, R_\alpha \cup R_\beta)}} Match(f_{c_i \rightarrow c}(o_i), \hat{o}_i)$$

where

$$\hat{o}_i = \arg \max_{\substack{o' \\ o \sqsubseteq o' \sqsubseteq \bar{o}}} Match(f_{c_i \rightarrow c}(o_i), o')$$

4.2 Reuse

Adaptation is solved by local exploration around the case found while satisfying the reservations of the agents in conflict. The exploration is made by looking for expansions or contractions of the ontology of the case that might increase the matching degree with the current case, and then using the so expanded ontology to generate a solution.

$$o^R = \arg \max_{\substack{o' \\ o^* \sqsubseteq o' \sqsubseteq \bar{o}^*}} Match(f_{c_i \rightarrow c}(o'), o)$$

4.3 Revise

In this step the solution generated in the step above is proposed to the participants. In this way changes in the goal set of the agents or new reservations may appear. Also, the ontology can be refined as the adaptation of the ontology may have introduced elements found irrelevant by the participants. For instance in the case of the water dispute (Figure 7-D) a farmer may say that watering the fields on midday is unfeasible because of sun heat. This reservation might make a solution in which watering is done in the middle of the day unfeasible.

If no success is obtained, another iteration is made.

$$R'_\alpha = f(\text{Dialogue}, R_\alpha)$$

4.4 Retain

MediaThor follows a simple method to memorise new cases. A case will be memorised if there is no other case in the memory of cases that has a similarity degree (using the matching between the ontologies) with it over a threshold θ . That is, if the memory of cases at time t is C^t and the solved case at time t is c then

$$C^{t+1} = \begin{cases} C^t \cup \{c\} & \text{if } Sim(c, c^*) < \theta \\ C^t & \text{otherwise} \end{cases} \quad (15)$$

We will update the ontologies as a fusion/combination of the ontologies of the cases. In that way we can avoid the problem of asking for an ontology which is always annoying. The details of the fusion/combination of ontologies are beyond the scope of this paper.

5 Conclusions

Though during the years there has been some interest in automated mediation [15, 2, 23, 17], the field requires a significant effort in research and development. This paper has presented the initial work on the principles of building an automated mediation agent as part of the ‘curious negotiator’ framework, hence can utilise some of the machinery developed for it, in particular:

- the *information-based agency* [8, 7], which offers mechanisms that allow the interplay between argumentation and information;
- the *information-mining system*, [9, 10], which offer means for automated discovery and (to some extent) delivery of that information to negotiating agents;
- the *electronic/virtual institutions environment* [24, 25], which offers means not only for performing negotiation, but also for collecting the necessary data about the negotiation sessions in order to use it in mediation.

We established that mediation is an information revelation process where analogies (including ones across different domains) can play key role in reshaping the set of issues. The Orange and Sinai disputes demonstrate how through the revealing of the appropriate information and applying analogy to reshape a dispute mediation can succeed. *MediaThor* demonstrated that computationally, the approach requires the specification of the introduced mental models of the agents and the mechanisms for aligning/agreeing on the ontologies of the dispute that they use.

MediaThor also demonstrated that case-based reasoning offers a potential mechanism for the mediator for handling past experiences, though the structure of the case is complex (in comparison to the usually assumed attribute-value structure), extending the already complex structure for representing negotiation cases [19]. Overall, from a knowledge perspective, automating mediation needs to take in account that mediation is:

- a knowledge intensive process, where the mediators utilise their past experiences;

- a process that utilises information from negotiating parties and uses information for changing the positions these parties have on the negotiation table.

As the aim of the mediation is to decrease the difference between what is a perceived model and the actual model. As the mediator utilises information to decrease the uncertainty in the dispute, an automated mediation would require a measure of uncertainty $\mathbb{H}(\mathcal{M}^t)$, allowing to quantify and compare the uncertainty coming from the incomplete knowledge of the mental models of the agents. In terms of the two party mediation in Figure 2, this decrease of uncertainty in mental models should be observable, i.e. $\mathbb{H}(\mathcal{M}_{\mu(\alpha)}^t) < \mathbb{H}(\mathcal{M}_{\beta(\alpha)}^t)$ and $\mathbb{H}(\mathcal{M}_{\mu(\beta)}^t) < \mathbb{H}(\mathcal{M}_{\alpha(\beta)}^t)$. Within the framework of the information-based agency, which follows an information-theoretic approach, such measure should measure the information gain, as the mediator adds such gain. Viewing mediation as a dialogue system (see Figure 2), e.g. the mediator is engaged in a dialogue with each party, points also to the information-theoretic work in dialogue management strategies in conversational case-based reasoning [26]. In terms of an automated mediation system, the mediator should have the mechanism to determine the most informative question to ask at each stage of the interaction to each of the negotiating agents. These issues remain beyond the scope of this paper.

As it may deal with confidential information, mediation requires trust in the mediator from the parties involved, as much of the information about their position negotiating parties would not reveal to the other side. Though this has been beyond the scope of the paper, we are aware of this issue.

In conclusion, we would like to stress that the importance of mediation has been recognised world-wide. It's interesting to note that nowadays mediation skills are taught to students at various levels and schools spanning from elementary schools to university schools, including the Harvard Law School. Hence, the development of an automated mediation system is on the top priority of the research agendas.

Acknowledgements This research has been supported by the Sabbatical Programme of the Spanish Ministerio de Educación y Ciencia Grant SAB2006-0001, the Australian Research Council Discovery Grant DP0557168, the Generalitat de Catalunya Grant 2005/SGR/00093 and by the Spanish Ministerio de Educación y Ciencia Consolider Grant CSD2007-0022.

References

1. Simoff, S.J., Debenham, J.: Curious negotiator. In Klusch, M., Ossowski, S., Shehory, O., eds.: Proceedings of the Int. Conference on Cooperative Information Agents, CIA-2002, Springer, Heidelberg (2002)
2. Sycara, K.P.: Problem restructuring in negotiation. *Problem restructuring in negotiation* **37**(10) (1991) 1248–1268
3. Lewicki, R.J., Saunders, D.M., Minton, J.W.: *Essentials of Negotiation*. McGraw Hill (2001)
4. Franklin Dukes, E.: What we know about environmental conflict resolution: An analysis based on research. *Conflict Resolution Quarterly* **22**(1-2) (2004) 191–220
5. Gomes, A., Jehiel, P.: Dynamic process of social and economic interactions: on the persistence of inefficiencies. Centre for Economic Policy Research, CEPR, London (2001)

6. Ramchurn, S.D., Sierra, C., Godo, L., Jennings, N.R.: Negotiating using rewards. *Artificial Intelligence* **171** (2007) 805–837
7. Sierra, C., Debenham, J.: Information-based agency. In: *Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*, Hyderabad, India (2007) 1513–1518
8. Debenham, J.: Bargaining with information. In Jennings, N.R., Sierra, C., Sonenberg, L., Tambe, M., eds.: *Proceedings Third International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2004*, ACM Press, New York (2004) 664–671
9. Zhang, D., Simoff, S.: Informing the curious negotiator: Automatic news extraction from the internet. In Simoff, S., Williams, G., eds.: *Proceedings 3rd Australasian Data Mining Conference*, 6 - 7th December, Cairns, Australia (2004) 55–72
10. Zhang, D., Simoff, S., Debenham, J.: Exchange rate modelling for e-negotiators using text mining techniques. In: *E-Service Intelligence - Methodologies, Technologies and Applications*. Springer, Heidelberg (2007) 191–211
11. Aciar, S., Zhang, D., Simoff, S., Debenham, J.: Informed recommender: Basing recommendations on consumer product reviews. *IEEE Intelligent Systems* **May/June** (May/June 2007) 39–47
12. Gentner, D., Stevens, A.L., eds.: *Mental Models*. Erlbaum, Hillsdale, NJ (1983)
13. Van Boven, L., Thompson, L.: A look into the mind of the negotiator: Mental models in negotiation. *Group Processes & Intergroup Relations* **6**(4) (2003) 387–404
14. Kolodner, J.: *Case-Based Reasoning*. Morgan Kaufmann Publishers, Inc., San Mateo, CA (1993)
15. Kolodner, J.L., Simpson, R.L.: The mediator: Analysis of an early case-based problem solver. *Cognitive Science* **13**(4) (1989) 507–549
16. Smith, A., Stam, A.: Mediation and peacekeeping in a random walk model of civil and interstate war. *International Studies Review* **5**(4) (2003) 115–135
17. Chalamish, M., Kraus, S.: Automed - an automated mediator for bilateral negotiations under time constraints. In: *Proceedings of the International Conference on Autonomous Agents and Multi Agent Systems, AAMAS07*, Honolulu, Hawaii, USA, IFAAMAS (2007)
18. De Mantaras, R.L., McSherry, D., Bridge, D., Leake, D., Smyth, B., Craw, S., Faltings, B., Maher, M.L., Cox, M.T., Forbus, K., Keane, M., Aamodt, A., Watson, I.: Retrieval, reuse, revision and retention in case-based reasoning. *The Knowledge Engineering Review* **20**(3) (2005) 215–240
19. Matos, N., Sierra, C.: Evolutionary computing and negotiating agents. In: *Agent Mediated Electronic Commerce*. Springer, Heidelberg (1999) 126–150
20. Druckman, D., Druckman, J.N., Arai, T.: e-mediation: Evaluating the impacts of an electronic mediator on negotiating behavior. *Group Decision and Negotiation* **13** (2004) 481–511
21. Zartman, I.W., Berman, M.R.: *The Practical Negotiator*. Yale University Press, New Haven, CT (1982)
22. Giunchiglia F., Yatskevich M., S.P.: Semantic matching: Algorithms and implementation. *Journal on Data Semantics* (2007) 1–38
23. Wilkenfeld, J., Kraus, S., Santmire, T.E., Frain, C.K.: The role of mediation in conflict management: Conditions for successful resolution. In: *Multiple Paths to Knowledge in International Relations*. Lexington Books (2004)
24. Esteva, M.: *Electronic Institutions: From specification to development*. Phd thesis, Technical University of Catalonia, Barcelona (2003)
25. Bogdanovych, A.: *Virtual Institutions*. Phd thesis, Faculty of Information Technology, University of Technology, Sydney, Sydney (2007)
26. Branting, K., Lester, J., Mott, B.: Dialogue management for conversational case-based reasoning. In: *Proceedings of ECCBR04*. (2004) 77–90