

See discussions, stats, and author profiles for this publication at: <http://www.researchgate.net/publication/220907602>

# PADUA Protocol: Strategies and Tactics.

CONFERENCE PAPER · JANUARY 2007

DOI: 10.1007/978-3-540-75256-1\_42 · Source: DBLP

---

CITATIONS

6

---

DOWNLOADS

68

---

VIEWS

62

## 3 AUTHORS:



[Maya Wardeh](#)

University of Liverpool

14 PUBLICATIONS 70 CITATIONS

SEE PROFILE



[Trevor Bench-Capon](#)

University of Liverpool

205 PUBLICATIONS 2,477 CITATIONS

SEE PROFILE



[Frans Coenen](#)

University of Liverpool

207 PUBLICATIONS 1,236 CITATIONS

SEE PROFILE

# PADUA Protocol: Strategies and Tactics

Wardeh, M., Bench-Capon, T., and Coenen, F.

Department of Computer Science,  
The University of Liverpool,  
Liverpool L69 3BX, UK  
{maya,tbc,frans}@csc.liv.ac.uk

**Abstract.** In this paper we describe an approach to classifying objects in a domain where classifications are uncertain using a novel combination of argumentation and data mining. Classification is the topic of a dialogue game between two agents, based on an argument scheme and critical questions designed for use by agents whose knowledge of the domain comes from data mining. Each agent has its own set of examples which it can mine to find arguments based on association rules for and against a classification of a new instance. These arguments are exchanged in order to classify the instance. We describe the dialogue game, and in particular discuss the strategic considerations which agents can use to select their moves. Different strategies give rise to games with different characteristics, some having the flavour of persuasion dialogues and other deliberation dialogues.

## 1 Introduction

In this paper we describe an approach to classifying objects in a domain not governed by strict rules which makes use of a novel combination of argumentation and data mining techniques. Our scenario is that classification is performed by two agents, each of which has their own set of records of past examples recording the values of a number of features presumed relevant to the classification and the correct classification. One of the agents will propose a classification, and a set of justifying reasons for the classification. This proposal is based on the application of an association rule mined from the agent's set of examples to the case under consideration. The classification is the consequent of the rule, and the antecedent gives the justifying reasons. The other agent will then use its set of examples to play "devil's advocate" and attempt to overturn the proposed classification. We call our system *PADUA* (*Protocol for Argumentation Dialogue Using Association Rules*).

This interaction can be viewed as a form of dialogue game (e.g. [5]), based on the exchange of arguments. Dialogue games come in a variety of flavours [11], including *Persuasion*, where each participant tries to persuade the other participant of its own thesis, by offering arguments that support this thesis, and *Deliberation*, in which the participants exchange arguments to reach an agreed decision, with neither of them committed to a particular position at the outset.

Our interaction has aspects of both, with the balance different according to the dialogue strategies employed. *Formal Dialogue Games* [7] are interactions between two or more players, where each player moves by making utterances, according to a defined set of rules known as a *Dialogue Game Protocol*, which gives the set of moves possibly expected after a previous move; choosing the best move among these moves is the *Strategy Problem*.

As mentioned above, the key idea of PADUA is to form arguments directly from some set of records providing examples relating to a particular domain, avoiding any need for expert analysis of the data, or knowledge representation. The repository of background knowledge used by each participant can be considered to be a binary valued data set where each record represents a previous case and each column an attribute taken from the global set of attributes described by the background knowledge. Given this set up we can apply *Association Rule Mining* (ARM) [1] techniques to the data set to discover relations between attributes, expressed in the form of *Association Rules* (ARs). In order to use this information, we follow the notion of presumptive argumentation as the instantiation of argument schemes subject to challenge through characteristic critical questions introduced by Walton [11]. PADUA makes use of a custom argument scheme and associated critical questions. In this paper we shall discuss the strategy problem in PADUA, and the consequences of the strategy used for the dialogue type.

The rest of this paper is organized as following: Section 2 describes the argument scheme and the basic structure of the PADUA protocol. Section 3 gives some necessary background on strategies in dialogue systems. Section 4 discusses in detail the suggested strategy heuristics to be applied in PADUA protocol. Section 5 gives a detailed example of the suggested strategy, and some discussion of the relation between these strategies and dialogue types.

## 2 PADUA Protocol

The model of argumentation we will follow is that of [11] in which a prima facie justification is given through the instantiation of an argument scheme. This justification is then subject to a critique through a number of critical questions which may cause the presumptive conclusion to be withdrawn.

The basic argument scheme is one we have devised for the purpose, Argument from Proposed Rule. The *Premises* are:

1. *Data Premise*: There is a set of examples D pertaining to the domain.
2. *Rule Premise*: From D a Rule R can be mined with a level of confidence greater than some threshold T. R has antecedents A and a conclusion which includes membership of class C.
3. *Example Premise*: Example E satisfies A.
4. *Conclusion*: E is a C because A.

This can be subject to a number of critical questions:

- *Can the case be distinguished from the proposed rule?* If D supports another Rule R2 and the antecedents of R2 subsume those of R, and are satisfied by E, and the confidence of R2 is below T, this suggests that these additional features may indicate we are dealing with some kind of exception to R.
- *Does the rule have unwanted consequences?* If the conclusion of R includes some fact F not satisfied by E, this suggests that R is not applicable to E.
- *Is there a better rule with the opposite conclusion?* If there is another rule R3 which can be mined from D and the antecedents of R3 are satisfied by E, and the confidence of R3 is greater than that of R, this suggests that R is not applicable to E.
- *Can the rule be strengthened by adding additional antecedents?* This does not challenge the classification, but rather the justification for the classification. If there is another Rule R4 which can be mined from D and the antecedents of R4 subsume those of R, and are satisfied by E, and the confidence of R4 is greater than that of R, this suggests that the additional features should be included in the justification of the classification.
- *Can the rule be improved by withdrawing consequences?* This challenges neither the classification, nor the justification, but the rule proposed. If there is another Rule R5 which can be mined from D and the conclusions of R subsume those of R5 and include a feature not satisfied by E, provided the confidence of R5 remains above the threshold, this suggests that the additional features should be excluded from the rule justifying the classification.

This argument scheme and these critical questions form the basis of the PADUA dialogue game.

## 2.1 Dialogue Scenario

The proposed dialogue game consists of two players (the proponent and the opponent) which have conflicting points of views regarding some case ( $C$ ). The proponent claims that the case falls under some class ( $c_1$ ), while the opponent opposes the proponent's claim, and tries to prove that case actually falls under some other class ( $c_2 = \neg c_1$ ). Each player tries to establish its point of view by the means of arguments based on association rules, which are mined from player's own database, using an association rule mining technique as described in [4].

The proponent starts the dialogue by proposing some AR ( $R_1 : P \rightarrow Q$ ), to instantiate the argument scheme. The premises ( $P$ ) match the case, and the conclusion ( $Q$ ) justifies the agent's position. Then the opponent has to play a legal move that would undermine the initial rule proposed by the proponent: these moves are based of the five critical questions described above. As can be seen from the questions, four of these moves involve some new rule. This is mined from the opponent's background database, and represents an attack on the original rule. The turn then goes back to the proponent which has to reply appropriately to the last move. The game continues until one player has no adequate reply. Then this player loses the game, and the other player wins.

## 2.2 PADUA Framework

The formal framework we suggest *Argumentation Dialogue Framework (ADF)* is defined as follows:

$$ADF = \langle P, Attr, C, M, R, Conf, playedMoves, play \rangle \quad (1)$$

Where  $P$ : denotes the players of the dialogue game.  $Attr$ : denotes the whole set of attributes in the entire framework.  $C$ : denotes the case argued about.  $M$ : denotes the set of possible (legal) moves.  $R$ : denotes the set of rules that govern the game.  $Conf$ : denotes the confidence threshold, all the association rules proposed within this framework must satisfy this threshold.  $playedMoves$ : denotes the set of moves played in the dialogue so far, this set of played moves represents the commitment store of the dialogue system under discussion. Finally,  $play$ : is a function that maps players to some legal move.

## 2.3 PADUA Players

Each player in PADUA game ( $\forall p \in P = Pro, Opp.$ ) is defined as a dialogical agent [3]:

$$\forall p \in P : p = \langle name_p, Attr_p, G_p, \Sigma_p, \gg_p \rangle \quad (2)$$

where:  $name_p$ : is the *player (agent) name*, here:  $\forall p \in P$  then  $name(p) \in \{pro, opp\}$ .  $Attr_p$ : is the set of attributes this player can understand.  $G_p$ : is the set of *goals* this player tries to achieve, here  $G_p$  is defined as a subset of the attributes set  $Attr_p$ , i.e.  $G_p$  is the set of attributes (classes) this player tries to prove true.  $\Sigma_p$ : is the set of ARs the player has mined from its background database, hence fore  $p$  is defined as follows:  $\forall p \in P : \Sigma_p = \{r_1 \dots r_m\}$ , where  $r_i = \langle Prem, Con, Conf \rangle$  is an association rule and can be read as  $Prem \rightarrow Con$  with a confidence  $= Conf$ . The elements of  $Prem$  and  $Conc$  are defined as a tuple  $\langle attribute, value \rangle$ , where  $attribute \in Attr_p$ , and  $value$  is the list of values assigned to this attribute in the given rule.  $\gg_p$ : represents the preferences order over  $\Sigma_p$ , a definition of this preference relationship is suggested as  $\gg_p: \Sigma_p \times \Sigma_p \rightarrow \{true, false\}$ , but the exact implementation of this relation may differ from player to player.

## 2.4 PADUA Legal Moves

The set of moves ( $M$ ) consists of 6 possible moves, one based on instantiating the argument scheme and five based on the critical questions against an instantiation. They are identified as follows:

1. *Propose Rule*:  $p$  plays this move to propose a new rule with a confidence higher than some confidence threshold.
2. *Distinguish*: this move is played to undermine a previously played move, as it adds some new premise(s) to this rule, such that the confidence of the new rule is lower than the confidence of the original rule (and/or lower than some acceptance threshold).

3. *Unwanted Consequences*: Here  $p$  suggests that certain consequences (conclusions) of some rule do not match the case under discussion.
4. *Counter Rule*:  $p$  plays this move to propose a new rule that contradicts the previous rule. The confidence of the proposed counter rule should be higher than the confidence of the previous rule (and/or than the threshold  $Conf$ ).
5. *Increase Confidence*:  $p$  plays this move to add some new premises to a previous rule so that the overall confidence rises to some acceptable level.
6. *Withdraw Unwanted Consequences*:  $p$  plays this move to exclude the unwanted consequences of the rule it previously proposed, while maintaining a certain level of confidence.

This defines the formal dialogue game. We now consider the strategies that might be used when playing the game. First we consider some related previous work on developing strategies for formal dialogue games.

### 3 Dialogue Strategies: Background

This section discusses some previous argumentation systems that have considered argument selection strategies:

Moore, in his work with the DC dialectical system [8], concluded from his studies that an agent’s argumentation strategy is best analyzed at three levels:

1. Maintaining the focus of the dispute.
2. Building its point of view or attacking the opponent’s one.
3. Selecting an argument that fulfils the objectives set at the previous two levels.

The first two levels refer to the agent’s strategy, i.e. the high level aims of the argumentation, while the third level refers to the tactics, i.e. the means to achieve the aims fixed at the strategic levels. Moore’s requirements form the basis of most other research into agent argumentation strategies.

In [2] a computational system was suggested that captures some of the heuristics for argumentation suggested by Moore. This system requires a preference ordering over all the possible arguments, and a level of prudence to be assigned to each agent. The strength of an argument is defined according to the complexity of the chain of arguments required to defend this argument from the other arguments that attack it. An agent can have either a “build” or a “destroy” strategy. When using the build strategy (b-strategy), an agent tries to assert arguments the strength of which satisfies its prudence level. If the b-strategy fails, it switches to the destroy strategy (d-strategy), where it tries to use any possible way to attack the opponent’s arguments. The basic drawback of this approach is that computational limits may affect the agent’s choice.

In [6] a three layer system was proposed to model argumentation strategies: the first layer consists of the “default” rules, which have the form (utterance - condition); the higher two layers provide preference orderings over the rules. The system is shown to be deterministic, i.e. a particular utterance is selected

in a given situation every time, but this system still requires hand crafting of the rules.

In [10], a decision heuristic was proposed to allow the agents to decide which argument to advance. The idea behind this work is that an agent should, while attempting to win a dispute, reveal as little of what it knows as possible, as revealing too much information in a current dialogue might damage an agent’s chances of winning a future argument. A new argumentation framework was developed to represent the suggested heuristics and arguments. The main shortcoming of this approach is the exponential complexity of the algorithms used.

## 4 Strategies and Tactics for PADUA

In PADUA, a player  $p \in P$  must select the kind of move to be played, and also the particular content of this move depending on: the thesis this player aims to prove true (or false), the case under discussion, the player’s set of association rules, the amount of information this agent is willing to expose in its move, and the player’s current state in the dialogue. All these factors must be considered in the strategy the player adopts and the tactics applied to implement this strategy.

Table 1 lists the *possible next moves* after each of the legal moves in PADUA protocol. A player must select a single move to play in its turn; moreover every possible next move is associated with a set of possible rules: this set contains the rules that match the selection criteria of the move, i.e. their confidence, premises and conclusion match this move. Except for *unwanted consequences*, the moves introduce a new rule. Proposing a counter rule leads to a switch in the rule being considered, so entering a nested dialogue. The notion of move (act) and

<i>Move</i>	<i>Next Move</i>	<i>New Rule</i>
1	2,3,4	yes
2	1,3,5	yes
3	1,6	No
4	1,2,3	Nested Dialogue
5	2,3,4	yes
6	2,3,4	yes

**Table 1.** Possible Moves

content selection is argued to be best captured at different levels, as suggested by Moore [8]. In [2] the first level of Moor’s layered strategy was replaced with different profiles for the agents involved in the interaction. We also adopt this approach. Here we also add another level to Moore’s structure (level 0) which distinguishes PADUA games into two basic classes. In one players attempt to win using as few steps as possible, i.e. the move’s type and content are chosen so that the played move gives the opponent’s the least freedom to plan its next move. In the other, games that are played to fully explore the characteristics of

the underlying argumentation system, and dialogue game, so here the move’s type and content are chosen so that the played move will restrict the opponent’s freedom to plan its next move to the least extent possible. The layered strategy system we adopt is defined as follows:

- Level 0: *Define the game mood*: i.e. *Win mode* or *Dialogue mode*.
- Level 1: *Define the players (agents) profiles*.
- Level 2: *Choose to build or destroy*: where in a *Build mode* the player tries to build its own thesis, while in a *Destroy mode* the player tries to destroy the opponent’s thesis.
- Level 3: *Choose some appropriate argumentative content*: depending on the tactics and heuristics suggested.

#### 4.1 Agent Profile

In [3], which used arguments based on standard if then rules, five classes of agent profiles were defined as follows:

1. *Agreeable Agent*: Accept whenever possible.
2. *Disagreeable Agent*: Only accept when no reason not to.
3. *Open-minded Agent*: Only challenge when necessary.
4. *Argumentative Agent*: Challenge whenever possible.
5. *Elephant Child Agent*: Question when ever possible.

In this paper we consider only the first two profiles (i.e. agreeable and disagreeable agents), as these attitudes are the most appropriate for the particular argument scheme we are using.

#### 4.2 PADUA Strategy

The function *Play* is defined as follows:

$$Play : P \times M_{poss} \times R_{poss} \times playedMoves \times S \times \rightarrow M \quad (3)$$

Where:  $P$  is the set of game players;  $playedMoves$  is the set of moves played in the dialogue so far; and  $M$  is the set of possible (legal) moves.  $M_{poss}$ : is the set of the possible moves this player can play  $M_{poss} \subseteq M$  (as defined in Table1).  $R_{poss}$ : is the set of legal rules that this agent can put forward in the dialogue ( $R_{poss} \subseteq 2^{\mathcal{E}_p}$ ); this set contains the rules that match the each of the possible moves.  $S$ : is the *Strategy Matrix*, and has the form  $S = [g_m, profile_P, s_m]$  where:  $g_m \in G_m$ : is the game mode, where  $G_m = \{win, dialogue\}$ ,  $profile_P \in Profile_P$ : is the player profile, where  $Profile_P = \{agreeable, disagreeable\}$ , and finally,  $s_m \in S_m$ : is the strategy mode, where  $S_m = \{build, destroy\}$ .

#### 4.3 PADUA Tactics

A set of tactics are suggested to fulfil the strategic considerations discussed above; these concern the best move to play and, where applicable, the content of the chosen move, i.e. the best rule to be put forward in the dialogue.



**Legal Moves Ordering** *Legal moves' ordering* defines the order in which legal (possible) moves are considered when selecting the next move. All games begin with Propose Rule: there are three possible responses to this, and these in turn have possible responses. The preference for these moves depends on whether the agent is following a build or a destroy strategy. In a destroy strategy the agent will wish to discredit the rule proposed by its opponent, and hence will prefer moves such as *unwanted consequences* and *distinguish*. In contrast when using a build strategy an agent will prefer to propose its own rule, and will only attempt to discredit its opponents rule if it has no better rule of its own to put forward. The preferred order for the two strategies is shown in Table2.

Whether players are agreeable or disagreeable will have an influence on whether the agent wishes to dispute the rule put forward by its opponent, and, the nature of the challenge if one is made.

<i>Last Move</i>	<i>Build Mode</i>	<i>Destroy Mode</i>
1	4,3,2	3,2,4
2	1,3,5	3,5,1
3	1,6	6,1
4	1,3,2	3,2,1
5	1,3,2	3,2,1
6	1,3,2	3,2,1

**Table 2.** Possible Moves Preferences

**Agreeable Players** An agreeable player  $ap \in P$  accepts a played rule without challenging it if:

1. An exact match of this rule can be found in its own set of association rule ( $\Sigma_{ap}$ ) with a higher or similar confidence.
2. Can find partial match of this rule in its own set of association rule ( $\Sigma_{ap}$ ), a rule  $r_{pm} \in \Sigma_{ap}$  is considered to be a partial match of another rule  $r \in \Sigma_{ap}$  if it has the same conclusion (consequences) of  $r$ , it's set of premises is a superset of rule  $r$  premises, and all these premises match the case; and finally it has a higher or similar confidence.

Otherwise the agreeable agent challenges the played move, depending whether it wishes to build or destroy using the legal moves preferences shown in Table2 selecting a rule using the following content tactics:

1. *Confidence*: Confidence of moves played by agreeable agent should be **considerably** lower/higher than the attacked rule, otherwise the agent agrees with its opponent.
2. *Consequences*: Consequences always contain a class attribute. Minimum changes to previous move consequences. As few attributes as possible.

3. *Premises*: Premises are always true of the case.  
Minimum changes to previous move premises.  
As few attributes as possible.

**Disagreeable Players** A disagreeable agent accepts a played rule if and only if all possible attacks fail, and so does not even consider whether its data supports the rule; the choice of the attack (i.e legal move) to be played depends on the preferences shown in Table2 and the choice of rule is in accordance with the following content tactics:

1. *Confidence*: Confidence of moves played can be:
  - (a) Considerably different from last move
  - (b) Slightly different from last move.
 The choice of confidence depend on the general mode of game whether it's in a win-mood or a dialogue-mood.
2. *Consequences*: Consequences always contain a class attribute.  
As few attributes as possible.
3. *Premises*: Premises are always true of the case.  
As few attributes as possible.

**Best Move** Table3 brings these considerations together and shows the best move relative to the agent type and the game mode, for each of the move types. For example in win mode an agent will want to propose a rule with high confidence, as one which the opponent is likely to be forced to accept, whereas in game mode, where a more thorough exploration of the search space is sought, any acceptable rule can be used to stimulate discussion.

## 5 Example

Our example domain concerns the voting records of US members of Congress, on the basis of which we wish to classify them according to party affiliation. Although there will be typical Democrat and Republican views on various issues, people may vote against the party line for personal or regional reasons. Some members of Congress may be more maverick than others. Thus, while there is no defining issue which will allow us to classify with certainty, we can argue for a classification on the basis of voting records. The data set we use is taken from [9], and it represents the U.S. House of Representatives members of Congress (in the 1984 US congressional elections) on the 16 key votes identified by the (CQA). The congressional voting records database contains 435 instances, among which (45.2%) are Democrats and (54.8%) are Republicans. The dataset original 17 binary attributes (including the class attribute) were normalized to 34 unique numerical attributes, each corresponds to certain attribute value. This dataset was horizontally divided into two equal size datasets, each of which was assigned to a player in PADUA framework. Rules were mined from this dataset using 30% support, and 70% confidence thresholds.

Best Moves				
	Agreeable		Disagreeable	
	Win mode	Game mode	Win mode	Game mode
<b>propose</b>	High confidence	Average confidence	High confidence	Average confidence
	Fewest attributes	Average attributes	Fewest attributes	Fewest attributes
<b>distinguish</b>	Lowest confidence	Average drop	Lowest confidence	Average drop
	Fewest attributes	Fewest attributes	Fewest attributes	Fewest attributes
<b>Unwanted consequences</b>	If some consequences are not in or contradict the case	Only if some consequences contradict the case	If some consequences are not in or contradict the case	
<b>Counter rule</b>	Average confidence	High confidence	High confidence	Average confidence
	Fewest attributes	Fewest attributes	Average attributes	Fewest attributes
<b>Increase Confidence</b>	Highest confidence	Average increase	Highest confidence	Average increase
	Fewest attributes	Fewest attributes	Fewest attributes	Fewest attributes
<b>Withdraw unwanted consequences</b>	The preferable reply to unwanted consequences attack → selecting criteria is the same of the very last move that led to the unwanted consequences.			

Table 3. Best move content tactics

We have experimented by running several PADUA dialogue games, starting from the same case. The difference between the games lays in the underlying strategy options of each agent that participate in each of these games.

Table4 shows the attributes of the case used in the example.

Case:	[5, 7, 13, 15, 17, 21, 24, 26, 29]		
5:	adoption-of-the-budget - resolution=y.	7:	physician-fee-freeze=n.
13:	anti-satellite- test-ban=y.	15:	aid-to-nicaraguan -contras=y.
17:	mx-missile=y.	21:	synfuels-corporation-cutback=y.
24:	education-spending=n.	26:	superfund-right-to-sue=n.
29:	duty-free- exports=y.		

Table 4. Example Case

As an illustration, we will describe the run with two disagreeable agents playing in win mode, the proponent (*Prop*) using a *build strategy* and the opponent (*Opp*) a *destroy strategy*. *Prop* begins by proposing a rule to justify saying that the member of Congress concerned is a Democrat: **R1**: Democrat because

education-spending=n and **duty-free-exports =y** with a (97.66%) confidence. *Opp* can reply by distinguishing this rule, since adding the premise **aid-to-nicaraguan-contras=y** reduces confidence to 80.43%. *Prop* now proposes a new rule: **R2: Democrat because mx-missile=y and duty-free-exports=y** with a (98.63%) confidence. This rule cannot be distinguished or countered since there is no better rule for Republican and so *Prop* wins.

Note how, *Opp*, being in *destroy mode*, uses first the distinguish move and only proposes a rule if this move cannot be played. In *build mode* *Opp* plays a rule of its own. Note also that the distinction made greatly reduces the confidence, whereas a distinction with a less drastic effect could have been played in game mode. When *Opp* is an *agreeable agent* it would simply accept the proposed rule, as it too can mine the rule with sufficient confidence. Where *Prop* is in *destroy mode*, it responds to the distinction with an increase confidence move, forcing *Opp* to propose a rule of its own.

As would be expected, in *game mode*, longer dialogues are produced. Where the agents are both agreeable, game mode leads to a series of rule proposals until a mutually acceptable one is found. Where *Opp* is in *destroy mode*, *Prop's* proposals will be met by a distinction, and where *Opp* is in *build mode* it will produce counter proposals as long as it can. Where *Prop* is in *destroy mode* it will make use of the unwanted consequences move to refute *Opp's* distinction if possible. Where both agents are disagreeable and in win mode, because the game does not terminate on the proposal of an acceptable rule, this last move, refuting a distinction by pointing to unwanted consequences which cannot be met with a withdraw consequences move, is what ends the game.

## 6 Discussion

Padua provides a way of determining the classification of cases on the basis of distributed collections of examples related to the domain without the need to share information, and without the need for analysis and representation of the examples. The argumentation leads to a classification which, while uncertain, is mutually acceptable and consistent with the different collections of examples.

Different strategies for move selection give rise to dialogues with different characteristics. Using disagreeable agents gives rise to a persuasion dialogue, since the opponent will do anything possible to avoid accepting the proposal. Win mode will lead to the swiftest resolution: game mode between disagreeable agents will lead to a lengthier exchange, and concession may be forced without the best argument being produced. A dialogue between two agreeable agents has the characteristics of a deliberation dialogue in that here the opponent is happy to concede once an acceptable proposal has been made. Win mode may be a very short exchange, since this simply verifies that *Prop's* best rule is also acceptable with respect to the second agent's data set. When game mode is used, the game has the flavour of brainstorming in that more ideas, even some which are less promising, will be explored.

Further work will empirically explore our system to examine the efficiency and quality of classifications and the effect of giving the individual data sets used by the agents particular characteristics. We also intend to explore domains in which classification is into a enumerated set of options rather than binary, and develop an extended version of the game with more than two participants.

## References

1. Agrawal, R., Imielinski, T., and Swami, A. N. (1993). Association rules between sets of items in large databases. In Proc. of ACM SIGMOD Int. Conf. on Management of Data. Washington, D.C, May 1993. pp 207-216.
2. Amgoud, L. and Maudet, N. (2002) Strategical considerations for argumentative agents (preliminary report). In Proc. of 9th Int. Workshop on Non-Monotonic Reasoning (NMR). Toulouse, France, April 2002. Special session on Argument, Dialogue, Decision. pp. 409-417.
3. Amgoud, L., Parsons, S. (2001) Agent dialogues with conflicting preferences. In Proc. of 8th Int. Workshop on Agent Theories, Architectures and Languages. Seattle, Washington, August 2001. pp 1-15
4. Coenen, F.P. Leng, P., and Goulbourne, G. (2004). Tree Structures for Mining Association Rules. Journal of Data Mining and Knowledge Discovery, Vol 8, No 1, pp25-51.
5. Hamblin, C. L.(1970) Fallacies. Methuen, 1970.
6. Kakas, A.C., Maudet, N. and Moraitis, P.(2004) 2nd Int. Workshop on Argument in Multi Agent Systems (Argmas 2004), Springer LNCS 3366 2005.
7. Mcburney, P., Parsons, S. (2002). Games That Agents Play: A Formal Framework for Dialogues between Autonomous Agents. In Jo. of logic, language and information, 11(3), pp 315-334.
8. Moore, D.: Dialogue game theory for intelligent tutoring systems. PhD thesis, Leeds Metropolitan University (1993).
9. D.J. Newman, S. Hettich, C.L. Blake and C.J. Merz (1998). UCI Repository of machine learning databases. <http://www.ics.uci.edu/~mllearn/MLRepository.html>. University of California, Irvine, Dept. of Information and Computer Sciences.
10. Oren, N., Norman, T. J., Preece, A.(2006). Loose Lips Sink Ships: a Heuristic for Argumentation. In Proc. of 3rd Int. Workshop on Argumentation in Multi-Agent Systems(Argmas 2006), Hakodate, Japan, May 2006. pp. 121134. Proceedings available at <http://homepages.inf.ed.ac.uk/irahwan/argmas/argmas06/argmas2006proceedings.pdf>
11. Walton, D. N., Krabbe, E. C. W. (1995). Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. SUNY Press, Albany, NY, USA.
12. Walton, D. N. (1996) Argument Schemes for Presumptive Reasoning. Lawrence Erlbaum Associates, Mahwah, NJ, USA.