PISA - Pooling Information from Several Agents: Multiplayer Argumentation from Experience

Maya Wardeh, Trevor Bench-Capon and Frans Coenen
Department of Computer Science
The University of Liverpool
UK

Abstract In this paper a framework, PISA (*Pooling Information from Several Agents*), to facilitate multiplayer (three or more protagonists), "argumentation from experience" is described. Multiplayer argumentation is a form of dialogue game involving three or more players. The PISA framework is founded on a two player argumentation framework, PADUA (*Protocol for Argumentation Dialogue Using Association Rules*), also developed by the authors. One of the main advantages of both PISA and PADUA is that they avoid the resource intensive need to predefine a knowledge base, instead data mining techniques are used to facilitate the provision of "just in time" information. Many of the issues associated with multiplayer dialogue games do not present a significant challenge in the two player game. The main original contributions of this paper are the mechanisms whereby the PISA framework addresses these challenges.
Keywords: Multiplayer argumentation by example.

1 Introduction

In many situations agents need to pool their information in order to solve a problem. For example in the field of classification one agent may have a rule that will give the classification, but that agent may be unaware of the facts which will enable the rule to be applied, whereas some other agent does know these facts. Individually neither can solve the problem, but together they can. One method to facilitate information sharing is to enable a dialogue between the two agents. Often this dialogue takes the form of a persuasion dialogue where two agents act as advocates for alternative points of view. A survey of such approaches is given in (Prakken 2006). The systems discussed by Prakken suppose that agent knowledge is represented in the form of belief bases, essentially a set of rules and facts. In consequence dialogue moves are strongly related to knowledge represented in this form. A typical set of moves for the systems in (Prakken 2006) are:

- *Claim P*: *P* is the head of some rule
- *Why P*: Seeks the body of rule for which *P* is head
- *Concede P*: agrees that *P* is true
- *Retract P*: denies that *P* is true
- *P since S*: A rule with *P* head and *S* body.

In (Wardeh et al, 2007, 2008) we introduced an alternative basis for such persuasion dialogues to enable what we termed arguing from experience to solve classification problems. Here the agents do not have belief bases, but only a database of previous examples. When presented with a new case the agents use data mining techniques to discover associations between features of the case under consideration and the appropriate classification according to their previous experience. We argued that this has several advantages:

- Such arguments are often found in practice: many people do not develop a theory from their experience, but when confronted with a new problem recall past examples;
- It avoids the knowledge engineering bottleneck that occurs when belief bases must be constructed;
- There is no need to commit to a theory in advance of the discussion: the information can be deployed as best meets the need of the current situation;
- It allows agents to share experiences that may differ: one agent may have encountered types of case that another has not.

The moves made in arguments based directly on examples contrast with those found in persuasion dialogues based on belief bases, and have a strong resemblance to those used in case based reasoning systems, e.g. (Ashley 1990) and (Aleven 1997). The work in (Wardeh et al 2007, 2008) describes an implementation allowing argumentation from experience of this sort, called PADUA (Protocol for Argumentation Dialogue Using Association Rules). PADUA, however, is, like the systems of (Prakken 2006), restricted to just two agents. In this paper we extend the approach to allow a number of agents to participate in the discussion.

The rest of this paper is orgnaised as follows: *Section 2* will motivate the need for the extension to more than two players, and discuss some of the key issues involved in doing so. *Section 3* will briefly describe the distinctive argument moves made in PADUA and *Section 4* will recall the two player PADUA protocol. *Section 5* describes the changes made to the two player version to allow multiple participants. Finally *Section 6* offers some concluding remarks.

2 Need for Multiparty Dialogue

The focus of argumentation dialogues in work on agents has largely been limited to two party dialogues. Typically these dialogues have been adversarial. In practice, however, such dialogues can take place with more than two parties. The two party scenario is not always the ideal way to handle certain situations. For some classification problems there may be a set of possibilities, and it is desirable to allow each possibility to have its own advocate, so that each possibility can be given fair consideration. Examples are numerous. In debates within political institutions

the representatives of different parties and factions may look at the same issue from several different angles, and each propose different solutions. In medicine, several diseases may share similar symptoms. Hence doctors may confuse one with the other and only by consulting and arguing with each other can the right diagnosis can emerge. At an everyday level, consider a group of friends, each with their own favourite restaurant, deciding where to eat. By allowing these parties to all take part in the dialogue they can pool their experience, increase their chances of winning over their opponents, and of the correct solution being reached.

The above real life examples not only illustrate the frequency and importance of multiparty argumentation dialogues, but also highlight several issues that must be taken into consideration when trying to model and implement multiparty dialogues in general and argumentation dialogue in particular. There are several different models for multi-party dialogue: formal meetings, informal meetings, bulletin boards, seminars, brainstorming sessions and so on. Some of the most important issues arising from this variety of models are discussed in (Dignum and Vreeswijk 2004). For each of these issues, choices must be made to yield a particular flavour of dialogue. The most important features for our purposes here are:

1. *System openness*: multiparty argumentation can either be closed, which means that the dialogue starts with *n* players and continues with *n* players until it is terminated, so that new participants are not allowed to join the dialogue once it has started, and players cannot leave the dialogue while it is in progress. Open systems are rather the opposite as players in such systems are free to join in or leave whenever they want.
2. *Players' roles***:** in 2-party argumentation dialogues the roles of players are limited to one proponent and one opponent, while in multiparty argumentation dialogues the situation is more complicated. We may have several proponents and several opponents of the thesis. Alternatively we may have several participants each with their own distinct option. Also, some parties within the dialogue can take a neutral point of view, or stand as mediator between the opposing parties (other dialectical roles may also be taken into consideration). Also, linguistically speaking, in two player dialogues one (and only one) player can speak per turn (be the speaker) while the other listens (be the listener or hearer). In multiparty dialogues there can be more than one hearer per turn, and the roles of those hearers may also vary (e.g. addressee, auditor, over-hearer etc…). Besides, one can argue that there can be more than one speaker per turn, since in real life people may start talking at the same time, or interrupt each other, or even start yelling at each other --- the loudest winning the situation!
3. *Addressing*: this is also a solved issue in 2-party dialogues, which one can picture as two players facing each other so when one talks the other realizes this and starts to listen! In multiparty dialogue the issue of deciding to whom a player is addressing his speech is not as clear. Different strategies can be used to solve this problem. These include: public broadcasting where all the players listen to what the speaker is saying; or targeted broadcasting of the *speech act* to some players (but not all of them); or just addressing the speech act to one

particular player. In some situations the possibility of private communication can play an important strategic role.

4. *Turn taking*: In 2-party dialogues "turn switching" is straightforward, the listener becomes the speaker when the current speaker finishes. This is not the case in multiparty dialogue, as more than one player may wish to talk at the same time and they all may request to take the next turn. In argumentation the decision whether the turn can be given to some player, or to more than one player at the same time, or whether the turn passes from one player to another in a certain order can greatly influence the final result of the argumentation process. Additional turn taking can give rise to significant fairness issues.
5. *Termination*: The termination of 2-party argumentation dialogues happens once one of them has convinced the other (or once one of them has run out of things to say). In multiparty argumentation the dialogue may be terminated either when *all* the other players are convinced or once the *majority* of them are. Another issue regarding termination is that sometimes players may fail to convince each other and could end up playing for ever; therefore there should be a mechanism to end such dialogues. Finally in some scenarios the game may end without one single player winning the majority of other players' votes: in these cases there should also be a mechanism to determine the winner of the game or simply by allowing ties to take place.

In the following we shall first recapitulate the current framework of the PADUA protocol and then describe the developments we have made in order to be able to host multiplayer argumentation dialogue games.

3 Arguing from Experience

In this section we consider what speech acts will be used of dialogues attempting to argue from experience. One field in which arguing on the basis of precedent examples is of importance is common law. Important work has been carried out by, amongst others Ashley (Ashley 1990) and Aleven (Aleven 1997). What has emerged from this work is there are three key types of move:

- Citing a case
- Distinguishing a case
- Providing a Counter Example

We will discuss each of these in turn, anticipating the next section by indicating in brackets the corresponding speech acts in the PADUA protocol described in section 4. One difference with case based reasoning that should be noted is that in argumentation from experience the whole database is used rather than identifying a single case as a "*precedent*". Unlike legal decisions the authority comes from the

frequency of occurrence in the set of examples rather than endorsement of a particular decision by an appropriate court.

In legal argument, citing a case involves identifying a previous case with a particular outcome which has features in common with the case under consideration. Given the things in common, the suggestion is that the outcome should be the same. Applied to argumentation from experience in the case of classification, the argument is something like: *in my experience, typically things with these features are Cs: the current example has those features, so it is a C* (propose rule[1]). The features in common are thus presented as reasons for classifying the example as *C*, justified by the experience of previous examples with these features. Distinguishing is one way of objecting to this, giving reasons why the example being considered does not conform to this pattern. It often involves pointing to features present in the case which make it atypical, so that the "typical" conclusions do not follow. For example the feature may indicate an exception: *although typically things with these features are Cs, this is not so when this additional feature is present* (distinguish). As an example, swans are typically white, but this is not so for Australian swans. Another form of distinction is to find a missing feature that suggests that the case is not typical: *while things with these features are typically Cs, Cs with these features normally have some additional feature, but this is not present in the current example* (unwanted consequences). A third kind of distinction would be to supply a more typical case: *while many things with these features are Cs, experience would support the classification more strongly if some additional feature were also present* (increase confidence).

Thus we have three types of distinction, with differing strengths: (i) that the current example is an exception to the rule proposed; (ii) that there are reasons to think the case untypical, and so that it may be an exception to the rule proposed; and (iii) that the confidence in the classification would be increased if some additional features were present. In all cases, the appropriate response is to try to refine the proposed set of reasons to meet the objections and thus accommodate the exception. Confidence in arguments is important: arguments from experience are usually associated with a degree of confidence: experience will suggest that things with certain features are often/usually/almost always/without exception *Cs*. This is also why dialogues to enable experience to be pooled are important: one participant's experience will be based on a different sample from that of another's. In extreme cases one person may have had no exposure to a certain class of exceptions: a person classifying swans with experience only of the Northern hemisphere needs this to be supplemented with experience of Australian swans. In less extreme cases, it may only be the confidence in the classification that varies.

---

[1] Note that "rule" in this context is not an axiom of a theory. Rules are merely proposed as a way of deciding the current case. They function in the same way as the decision rules or tests proposed in oral argument in US Supreme Court hearings, designed to serve as the basis for further probing (Ashley et al 2007) to see whether the test can be used.

Counter examples differ from distinctions in that they do not attempt to cast doubt on the reasons, but suggest that there are better reasons for believing the contrary. The objection here is something like: *while these features do typically suggest that the thing is a C, these other features typically suggest that it is not* (counter rule). Here the response is either to argue about the relative confidence in the competing reasons, or to attempt to distinguish the counter example. Thus a dialogue supporting argument from experience will need to accommodate these moves: in the next section we will describe how they are realized in the PADUA protocol.

4 PADUA Protocol

PADUA (*Protocol for Argumentation Dialogue Using Association Rules*) is an argumentation protocol designed to enable participants to debate on the basis of their experience. PADUA has as participants agents with distinct datasets of records relating to a classification problem. These agents produce reasons for and against classifications by mining association rules from their datasets using data mining techniques (Agrawal 1993, Goulbourne et al 1999, Coenen et al. 2004). By "*association rule*" we mean that the antecedent is a set of reasons for believing the consequent. In what follows $P \rightarrow Q$ should be read as "P are reasons to believe Q". A full description of PADUA is given in (Wardeh et al 2008).

PADUA adopts six dialogue moves, related to the argument moves identified in the previous section:

1. *Propose Rule*: allows generalizations of experience to be cited, by which a new association with a confidence higher than a certain threshold is proposed.
2. *Attacking moves:* these pose the different types of distinction mentioned above:
   - *Distinguish*: When a player p plays a *distinguish* move, it adds some new premise(s) to a previously proposed rule, so that the confidence of the new rule is lower than the confidence of the original rule.
   - *Counter Rule*: is very similar to *propose rule* and is used to cite generalizations leading to a different classification
   - *Unwanted Consequences*: Here the player p suggests that certain consequences (conclusions) of the rule under discussion do not match the case under consideration.

3. *Refining moves:* these moves enable a rule to be refined to meet objections:
   - *Increase Confidence*: a player p adds one or more premise(s) to a rule it had previously played to increase the confidence of this rule.

– *Withdraw unwanted consequences*: a player *p* plays this move to exclude the unwanted consequences of the rule it previously proposed, while maintaining a certain level of confidence.

The PADUA protocol defines for each of those six moves a set of legal next moves (i.e. moves that can possibly follow this move). Table 1 summarizes PADUA protocol rules, and indicates whether a new rule is introduced.

Move	Label	Next Move	New Rule
1	Propose Rule	3, 2, 4	Yes
2	Distinguish	3, 5, 1	No
3	Unwanted Cons	6, 1	No
4	Counter Rule	3, 2, 1	Nested dialogue
5	Increase Conf	3, 2, 4	Yes
6	Withdraw Unwanted Cons	3, 2, 4	Yes

Table 1. The protocol legal moves

5 PISA

Extending the PADUA system based on two players to support many players is not a trivial task, as there are a number of problems that need to be resolved, in particular:

1. PADUA is a closed system with exactly two players;
2. Limited players' roles: Roles in the present PADUA systems are very restricted, in virtue of the simplicity inherent in two player games. Thus decisions have to be made about turn taking, and the addressing issues discussed above;
3. The lack of a powerful control structure: The existing PADUA implementation is not provided with any sophisticated control structure, as in two player games everything is simple, from turn taking which is obvious as no choice as to who will have the next turn is needed, to addressing which basically follows the model of face to face conversations, to game termination, which occurs when a player runs out of moves. This simplicity of two player games does not apply to multiplayer situations, which needs more careful surveillance and enforcement to guarantee that everything goes according to the design choices made with respect to the player roles and the conduct of the dialogue.

Therefore it is very important to resolve these issues before extending PADUA to support multiparty argumentation dialogues. The changes to be introduced to PADUA system should not affect the basic protocol structure (i.e. the moves and

rule mining), but rather complement the existing PADUA protocol with a control structure that makes it possible to organize the players and their turns within the games, and to identify the termination conditions for those games. Therefore we must first decide on the type of dialogue we will aim to produce, in terms of the issues identified in section 2. We will refer to this multiplayer version as PISA (*Pooling Information from Several Agents*).

The dialogue we will model is one where there are a range of options for classification, and each of the indefinite number of participants is the advocate of one of these options. Additionally there will be one agent, the chairperson, who will not be the advocate of any position, but rather manage the dialogue and facilitate communication between the advocates. This style of dialogue thus determines the roles of the players: a chairperson, and, for every option, one player acting as its advocate. Alternative models, where we might have several supporters for a particular option are no less valid, but we need to fix on one style of dialogue. The dialogue will be open, in that a participant (other than the chair) may enter or leave when they wish. For turn taking, we adopt a structure with rounds, rather than a linear structure where a given agent is selected as the next speaker. In each round, any agent who can make a move can do so: the chair then updates a central argument structure, and another round occurs. This is not perhaps the most usual structure for human meetings, but it can be found in some board games such as *Diplomacy*. We believe that the structure is particularly appropriate in order to achieve fairness in our situation where every advocate is playing for themselves, and has to regard every other advocate as an opponent (even though they may temporarily focus their efforts on a particular opponent). For addressing, every move after the first attacks a move of some other agent: that agent can be regarded as the addressee of that move, and the others as auditors. The game will terminate when no agent makes a contribution for two rounds (to ensure that they have really finished and not withheld a move for tactical reasons) or after some limiting number of rounds have been played. The model is essentially that of a facilitated discussion, with the chairperson acting as facilitator. We will now discuss the realization of this model, and the choices summarized above, in more detail in the following sub-sections.

5.1 Control Structure

The suggested control structure (Figure 1) can be pictured as a meeting room in which players can be seated and equipped with a blackboard like structure on which players can place their arguments (moves). The meeting is guided by a chairperson responsible for: organizing the dialogue, monitoring the players, controlling the turn taking procedure and applying the protocol rules. There is no distinction between the players other than their points of view regarding the case un-

der discussion. When a new game commences the chairperson randomly chooses one player to start the dialogue; in the meeting room scenario this player (P_1) is given the first seat on the meeting table; the rest of the players are seated randomly around the table and given according names ($P_2 \ldots P_n$). Then P_1 proposes a new rule and pastes it on a black board, this is called the first argumentation round (R_1). The game continues in the same manner, and in each of the following rounds all the players who can and wish to attack any of the arguments played in the previous turn(s) are allowed to place their arguments (moves) on the black board. The suggested facilitated discussion scenario enjoys the following advantages:

1. It increases the flexibility of the overall PISA system: by assigning the majority of protocol surveillance to the chairperson the system gains great flexibility. For example the system can be switched between closed and open by applying a few limited changes to the chairperson, while the rest of the players remain unaffected.
2. It is a very simple structure: there is no complicated turn taking procedure involving a choice of the next player, allowing the internal implementation of the players to be kept as simple as possible.
3. It provides a fair dialogue environment: the organizational configuration of the dialogue is neutralized by restricting the control tasks to the chairperson who is not allowed to take sides in the dialogue. This means that no one can be privileged to speak while the others must remain silent.

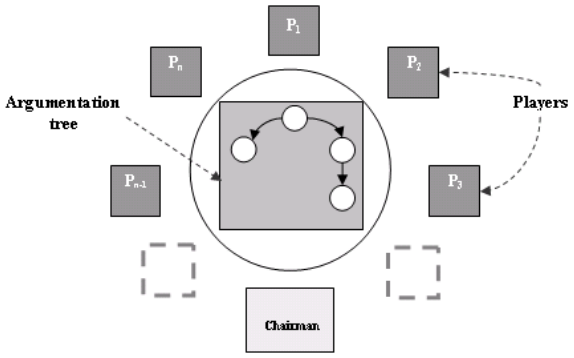

Fig. 1. The Suggested Structure

5.2 Turn Taking Policy

There is no strict turn taking procedure in PISA; any player who can make a legal move can participate in any round. There is no limitation on the number of players that can participate in a round, but each player is limited to one move per round:

the reason behind this restriction is mainly to simplify the dialogue, and contain the growth of the argument tree and repetition of moves within reasonable bounds.

This turn taking policy gives great freedom for players to apply their strategy in the way they think best to win the game; it also simplifies the game. Besides, if a strict turn taking procedure was applied it may affect the final result of the game, for example when the turn goes back to some player, the arguments against its proposition may be so numerous that deciding which one to defend against may be difficult. Also a player may not be able to defend itself against some attacks in its own turn, but might be able to do that after some other players have moved, In general the turn taking technique adopted here enjoys the following advantages:

1. It allows for turn skipping: Players who do not want to take part in a particular round can remain silent. (Although they can only remain silent for a limited number of rounds otherwise they will eventually lose).
2. Participants can play their attacks/counter attacks as soon as they want and they do not have to wait for an opportunity to contribute.
3. This way we solve the justice issue that rises from the fact that the turn taking process requires that some agent be favoured with the chance to speak. With turn taking a player can lose a game just because it is not given the opportunity to present its case at a favourable time. In other words all the players, under this policy, have an equal chance to win the game.

5.3 Game Termination

The chairperson terminates a PISA game when two rounds have passed and the argument state has not changed (i.e. no player has anything to say). The reason behind waiting two rounds is that sometimes some players (who are still capable of arguing) may choose to skip some rounds for strategic reasons. Therefore, if a round is passed without any move being played, the chairperson has to warn all the players that if they have anything to say the next round would be the last chance for them to say if they wish to prevent the game from ending.

Termination as described above is called "*legal termination*", but there are also cases in which the game should be exceptionally terminated (i.e. *exceptional termination*); the chairperson has the authority to terminate the game if any of the following events occur:

1. If only one player remains active after all the other players have withdrawn (in which case the surviving player wins the game).
2. If the game has taken more than *n* rounds (assuming that if the parties could not agree in n rounds were *n* is sufficiently large then they will not ever agree). In

this case no one wins the game. Of course, the value of *n* can differ from game to game, but it should allow for a large number of rounds.

5.4 Roles of the Players

Roles in PADUA are limited due to the inherent simplicity of two-player games, but in PISA more attention has to be paid to roles, and more importantly the way these roles change from round to round. The main distinction in players' roles is between attackers and defenders and between speakers and listeners as follows:

1. *Attacker(s) vs. defender(s):* while players are most certainly defenders of their point of views, they can take different positions regarding other players' proposals, therefore each player can decide whether to attack or defend other players' arguments. Enabling players to defend the arguments of other players (supposedly, and in the long term, their opponents) may be of strategic importance within the game. For example weak players may join forces to defeat stronger players, or a player may help another player to defeat a common opponent that it cannot defeat on its own. Once that opponent has been eliminated the remaining players can then attack one another.
2. *Speaker(s) vs. Listener(s) (addressee(s))***:** in the first round of PISA there is only one speaker (P_1) while the rest of the players are addressees (the chairperson may be considered as an auditor). In all the subsequent rounds there are *s* speakers (where *s* <= *m* number of players) where *s* is the number of the players participating in the given round. Once the speakers are done with their moves the addressees of the round are defined as the players whose arguments were attacked in this round and the rest of the players (i.e. those who did not participate and were not attacked in the given round) are assumed to be auditors.

5.5 Argumentation Tree

In PISA the notion of *Argumentation Tree* is used to describe the data structure used to represent the arguments played in each round, and the attack relations between those arguments. This tree acts as a mediating artifact for the dialogue as described in (Olivia et al 2008). Hence it is similar to other tree structures used in the literature, but this particular tree differs from other data structures because:

1. It consists of arguments played by more than one player
2. It implements three colours to mark the status of the arguments played so far.
3. It contains two types of links: explicit links (edges) and implicit links.

The *PISA Argumentation Tree* data structure consists of:

1. *Nodes*: represent the moves (i.e. arguments) played so far; each node has a number representing the player (the full structure of the node is described in the following subsection). The colour of the nodes changes from round to round as defined in the following Table 2:
2. *Links*: represent the explicit attack relationship between the nodes.
3. *Green Confidence*: is a global value (a real number) associated with the tree representing the highest confidence of the undefeated nodes (representing moves of types: 1, 4, 5 and 6 in Table 1). This value helps in defining the implicit attack relationships across multiple rounds (instead of physically representing those relationships as links in the tree).

Colour	**Meaning**	**Shifting to**
Green	The node is undefeated in the given round, and was added by one of the moves 1, 4, 5 and 6	(to red) If attacked by at least one undefeated node.
Red	The node is defeated in the given round.	(To green) if all attacks against a node are successfully defeated and the original node colour was green. (To blue) if all attacks against a node are successfully defeated and the original node colour was blue.
Blue	The node is undefeated in the given round, but nodes of this colour do not enforce new rules, their sole purpose is simply to undermine an argument played by another player using moves 2 and 3.	(To red) If attacked by at least one undefeated node.

Table 2. The Argumentation Tree Colours

When player P_i plays some move (m), this move must satisfy a number of conditions in order to be added as a node to the argumentation tree, otherwise the move will be rejected. The conditions of acceptance are as follows:

1. Move m is added to the tree if and only if it changes the colouring of the tree (consequently a player can not add an argument that enhances its position in the game if such a move does not change the tree colouring).
2. A player can put forward one move only per round (deciding which rule to play is strategy issue).
3. Moves (1, 4, 5, 6) implicitly attack all the other (1, 4, 5, 6) moves played by other players which have lower confidence.
4. Moves (2, 6) affect only the nodes they directly attack.
5. Moves (1, 4, 5, 6) explicitly attack the (2, 6) nodes they are associated with (if any).
6. Participants should not play moves that weaken their position (i.e. a player should not put forward moves that change the colouring of the argumentation

tree such that another player would take the lead). This condition holds when a player tries to attack blue nodes that were originally made to attack an argument proposed by other players.

5.6 Winner Announcement

Once a game has terminated, the chairperson consults the argumentation tree to determine the winner. The winner should satisfy one of the following rules:

1. If all the green nodes on the graph belong to the same player then this player wins the game.
2. If the green nodes on the graph belong to more than one player then the player which has played the green node with the highest confidence wins the game (as this node implicitly attacks the other nodes).
3. If there are no green nodes and all the blue nodes were played by the same player then this player wins the game.

Unfortunately not all the games end up with one clear winner, as there are cases where the decision as to who won the game may not be obvious. For example: If there is more than one green node with the same confidence on the argumentation tree at the end of the game belonging to different players and no green nodes with higher confidence. Also if the game ends without any green nodes on the argumentation tree and the blue ones belong to more than one player.

The first case is considered a *strong tie situation*, as the players have actually proposed classifications within the game. One possible solution may be starting a new game between the tying parties only and see how this game ends. But there is nothing to guarantee that this game will not also end up with a tie. In this case the chairperson may be forced to announce a tie (after the second game or after g number of mini games).

The second case is considered a *weak tie situation*, as the tied players did not actually have any proposed classifications at the end of the game. In such cases starting a new mini game may be of great benefit, but with the condition that the players should propose as many reasons for their classification as they can this time.

7 Conclusions

In this paper, PISA, a multiplayer argumentation from experience framework has been described and illustrated. The main original contribution of the paper is the

mechanisms whereby the framework addresses the many challenges found in multiplayer dialogue games which are either not present in the two player game or are not of significance in the two player game. Of particular note is the control structure used in PISA, the turn taking policy, the approach to game termination and the definition of the roles of the players allowing them to adopt differing strategies. The supporting argumentation tree data structure is also significant. Current evaluation of the PISA implementation has provided the authors with encouraging results indicating that in its current form PISA is a genuinely useful tool to support agent based argumentation. The evaluation has also indicated areas for further investigation. These include the development of alternative game strategies, the potential for players to form dynamic groups and alliances and further variants on the current PISA model.

References

1. Agrawal R, Imielinski T, Swami A.N et al (1993). Association rules between sets of items in large databases. In: Proc. of ACM SIGMOD Int. Conf. on Management of Data, Washington, 207-216.
2. Aleven V (1997). Teaching Case Based Argumentation Through an Example and Models. PhD thesis, University of Pittsburgh, Pittsburgh, PA, USA.
3. Ashley K. D (1990). Modeling Legal Argument. MIT Press, Cambridge, MA, USA.
4. Coenen F. P, Leng P, Goulbourne G et al (2004). Tree Structures for Mining Association Rules. In: Journal of Data Mining and Knowledge Discovery, Vol 8, No 1, 25-51.
5. Dignum F and Vreeswijk G.A.W (2003). Towards a test bed for multi-party dialogues. In Advances in Agent Communication, F. Dignum, Ed. Springer Verlag. Berlin, 121-230.
6. Goulbourne G, Coenen F. P, Leng L et al (1999). Algorithms for Computing Association Rules Using a Partial Support Tree. In: Proc. of ES99, Springer, London, UK, 132-147.
7. Olivia, E, Viroli, M, Omicini and McBurney et al (2008). Argumentation and artifact for dialogue support. In Proc. ArgMAS 2008, the 5th Inter Workshop on Argumentation in Multi-Agent Systems, AAMAS 2008, Lisbon, Portugal.
8. Prakken H (2006). Formal systems for persuasion dialogue. In: The Knowledge Engineering Review 21, 163-188.
9. Wardeh, M, Bench-Capon, T, Coenen, F.P et al (2008). Arguments from Experience: The PADUA Protocol. To appear in Proc. COMMA'2008.
10. Wardeh, M, Bench-Capon, T, Coenen, F.P et al (2007). PADUA Protocol: Strategies and Tactics. In Proc. ECSQARU 2007, 9th European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, LNAI 4724: 465-476.