Arguments from Experience: The PADUA Protocol

Maya WARDEH Trevor BENCH-CAPON and Frans COENEN
Department of Computer Science, University of Liverpool
Liverpool L69 3BX, UK

Abstract. In this paper we describe PADUA, a protocol designed to enable agents to debate an issue drawing arguments not from a knowledge base of facts, rules and priorities but directly from a dataset of records of instances in the domain. This is particularly suited to applications which have large, possibly noisy, datasets, for which knowledge engineering would be difficult. Direct use of data requires a different style of argument, which has many affinities to case based reasoning. Following motivation and a discussion of the requirement of this form of reasoning, we present the protocol and illustrate its use with a case study. We conclude with a discussion of some significant issues highlighted by our approach.

Keywords. Argumentation, Dialogue Games, Classification

1. Introduction

One application of argumentation which has received a good deal of attention is as the basis of a dialogue between two participants, in which one participant is trying to persuade another of the truth of some claim. In some variations, both participants are trying to persuade the other, and in others the participants are not so committed to a point of view, so that the dialogue takes on the characteristics of deliberation rather than persuasion. A thorough survey of a number of systems can be found in [18]. In this work Prakken identifies the speech acts typically used in such dialogues:

- *claim* P (assert, statement, ...). The speaker asserts that P is the case.
- *why* P (challenge, deny, question, ...). The speaker challenges that P is the case and asks for reasons why it would be the case.
- *concede* P (accept, admit, ...). The speaker admits that P is the case.
- *retract* P (withdraw, no commitment, ..). The speaker declares that he is not committed (anymore) to P. Retractions are real retractions if the speaker is committed to the retracted proposition, otherwise it is a mere declaration of non-commitment (e.g. in reply to a question).
- P *since* S (argue, argument, ...). The speaker provides reasons why P is the case. Some protocols do not have this move but require instead that reasons be provided by a claim P or claim S move in reply to a why move (where S is a set of propositions). Also, in some systems the reasons provided for P can have structure, for example, a proof tree or a deduction.

- *question P* (...). The speaker asks another participant's opinion on whether P is the case.

These moves seem to suppose that the participant's knowledge is organized in a certain way: namely as a set of facts and, typically defeasible, rules of the form $fact \rightarrow conclusion$. Thus *why P* seeks the antecedent of a rule with P as consequent; *P since S* volunteers the antecedent of some rule for P, and the other questions suggest the ability to pose a query to a knowledge base of this sort. Prakken's own instantiation of this framework [16] presupposes that the participants have *belief bases* comprising facts, defeasible rules, and priorities between rules. That the participants are presupposed to be equipped with such belief bases doubtless derives in part from the context in which these approaches have been developed. The original example of the approach was probably Mackenzie[12] who was interested in exploring a particular logical fallacy. The take up in Computer Science has largely been by people working in knowledge based systems and logic programming, so that the form of the belief base is a natural one to assume. The result, however, is that the debate takes place in a context where the participants have knowledge (or at least belief), and the dialogue serves to exchange or pool this knowledge. In this way persuasion takes place in the following ways:

- One participant supplies the other with some fact unknown to that participant, which enables the claim to be deduced;
- One participant supplies the other with some rule unknown to that participant, which enables the claim to be deduced;
- An inconsistency in one participant's belief base is demonstrated, so that a claim, or an objection to a claim is removed.

At least one of these must occur for persuasion to happen, but in a complicated persuasion dialogue all three may be required. This necessitates certain further assumptions about the context: that the beliefs of the participants are individually incomplete or collectively inconsistent. Although the participants have knowledge, it is defective in some way, and corrected or completed through the dialogue.

While persuasion dialogues of this form do take place, others take a different form, involving the sharing not of *knowledge*, but of *experience*. In this situation the participants have not analysed their experiences into rules and rule priorities, but draw directly on past examples to find reasons for coming to a view on some current example. One classic example of such reasoning is found in common law, especially as practiced in the US, where arguments about a case are typically backed by precedents. Even where decisions on past cases are encapsulated in a rule, the *ratio decendi*, the particular facts are still considered and play crucial roles in the argument. In informal everyday argument also the technique is common: "the last time we did this, that happened". Given the prevalence of such arguments, it is worthwhile to address the requirements for such dialogues and how they differ from the traditional persuasion dialogues described in [18].

Quite apart from the widespread use of arguments from experience of this sort there are compelling pragmatic reasons for investigating such arguments. The formation of effective belief bases requires a good deal of, typically expensive and skilled, effort. The so called knowledge engineering bottleneck has bedevilled the practical implementation of knowledge based systems throughout their history. If we see the dialogue system as a way of adding value to existing systems, we will find that there are very few suitable belief bases available. In contrast, there are many large datasets available, with each record in the dataset representing a particular case, a particular experience. This provides an extensive amount of experience to draw on, if we can find a way of deploying it through argumentation.

In the context of these arguments, typically all of the facts regarding the case under consideration are available at the outset. Thus this source of incompleteness, resolved through belief based persuasion dialogues, is not present. Nor can the rules be incomplete or give rise to inconsistency: there are no rules. In such arguments persuasion occurs not through one participant telling the other something previously unknown, but rather because the experience has been incorrectly or unsafely generalised to the current case, or because - importantly - experience differs from participant to participant, and one participant may have encountered an untypical or over narrow set of examples. For example, generalising on experiences confined to the Northern hemisphere, one might conclude that a bird was white on being told that it was a swan, but should be open to persuasion by another participant with experience of Australia also.

2. Arguing From Experience

Having seen a need to model arguments from experience, we now need to consider what speech acts will be typical of such dialogues, and see how they differ from those typical of the belief based persuasion dialogues identified by Prakken. One field in which arguing on the basis of precedent examples is Law. Important work has been carried out by, amongst others Rissland, Ashley and Aleven [4,2]. 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.

Citing a case involves identifying a previous case with a particular outcome which has features in common with the case under consideration. Given these things in common, the suggestion is that the outcome should be the same. Applied to argument from experience regarding the classification of an example, the argument is something like: *in my experience, typically things with these features are Cs: this has those features, so it is a C* (propose rule). 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, by saying 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 exhibit 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). Suppose we were considering a duck billed platypus: while we might classify it as mammal on the basis of several of its features, we would need to consider the objection that mammals are typically viviparous. 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 can have three types of distinction, with differing forces. The first argues that the current example is an exception to the rule proposed; the second that there are reasons to think the case untypical, and so that it may be an exception to the rule proposed.; the third argues no more than that 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, for example to accommodate the exception.

The point about confidence is important: arguments from experience are typically associated with some degree of confidence: our 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. In extreme cases this may mean that one person has 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.

Counterexamples differ from distinctions in that they do not attempt to cast doubt on the reasons, but rather to 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.

Another lesson from work on reasoning with legal precedent is the importance of intermediate concepts e.g. [5]. The point is analogous to the difficulty in classifying examples of XOR using a single layer perceptron [14]. No simple classification rule for XOR over two variables can be produced using only the truth functions of the inputs. Rather we must produce the intermediate classifications "and" and "or" and then classify in terms of these ("or" and not "and"). So too, with law: some features used in classifying cases are not simple facts of the case, but rather classifications of the applicability of intermediate concepts on the basis of a subset of the facts of the case. Dialogues representing arguments from experience must therefore be able to accommodate a degree of nesting, where first the satisfaction of intermediate concepts is agreed, and then used in the main classification debate.

3. The PADUA Protocol

In this section we describe PADUA (*Protocol for Argumentation Dialogue Using Association Rules*) 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/against classifications by mining association rules from their datasets using data mining techniques [1,7,9]. By “association rule” we mean no more than 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”. Six of the dialogue moves in PADUA relate to the argument moves identified above. One represents citing a generalization of experience, three pose the different types of distinction mentioned above, one enables counter examples to be proposed, and one enables a rule to be refined to meet objections.

Formally, PADUA is defined drawing on various elements from the different systems suggested in [3,13,17], as the following tuple:

$$< L_t, L_c, A, DP, \varphi, K, L, E, P, O, S > \tag{1}$$

where:

1. L_t: The topic language of PADUA dialogue game, containing the following elements:
   (a) $I = \{i_1, i_2, .., i_n\}$: the set of items. Each item $i \in I$ has a set of possible values V_i.
   (b) D = the set of database transactions, each transaction $T \in D$ is a subset of the items in I such that $T \subseteq I$. A transaction T satisfies a set of items $X \subseteq I$ if and only if $X \subseteq T$.
   (c) Association rules written as $ar(P \rightarrow Q, conf)$:
      i. P: rule’s premises.
      ii. Q: rule’s conclusion.
      iii. Each element $e \in P \cup Q$ is a tuple <*name, value*> where name is an item $i \in I$, and $value \in V_i$ is the value of this item in this association rule.
      iv. $P \cap Q = \phi$.
      v. $conf$: rule confidence, which means that $conf$% of the transactions in D that contains P contains Q also (i.e. the conditional probability of Q given P).
2. L_c: The communication language including:
   (a) *Speech Acts SA = {propose rule, distinguish, unwanted consequences, counter rule, increase confidence, withdraw unwanted consequences}* where:
      i. *propose Rule*: stands for citing examples in PADUA system by which a new rule with a confidence higher than a certain threshold is proposed. *counter Rule* is very similar and is used to cite counter examples in the same way.
      ii. *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. *increase Confidence* is very similar, except that it increases the confidence of an original rule.
      iii. *unwanted Consequences*: Here the player p suggests that certain consequences (conclusions) of the rule under discussion do not match the studied case. *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.

(b) *Moves*: each move $m \in M$ (set of all moves) is defined as a tuple <*sa, content*> such that:

   i. $sa \in SA$: is the move speech act (or type).
   ii. *content*: is an association rule matches the speech act of the move (except when *sa* = *unwanted Consequences* then $content = U \subset I$ (the set of unwanted consequences)).

(c) *Dialogue Moves*: a dialogue move $dm \in DM$ (the set of all dialogue moves) is defined as a tuple <*S,H,m,t*> such that:

   i. $S \in Ag$ is the agent that utters the move, given by $Speaker(dm) = S$
   ii. $H \subseteq Ag$ denotes the set of agents to which the move is addressed, given by a function $Hearer(dm) = H$
   iii. $m \in M$ is the move, given by a function $Move(dm) = m$.
   iv. $t \in DM$ is the target of the move i.e. the move which it replies to, given by a function $Target(dm) = t$. $t = \phi$ if M does not reply to any other move (initial move).

(d) *PADUA Dialogues*: defined as a set of finite dialogues, denoted by $DM^{<\infty}$ the set of all finite sequences from L_c. For any dialogue $d = \{dm_1, .., dm_n\}$, the speech act of the first move (dm_1) is a propose rule.

3. $A = \{a_1, ..a_n\}$: The set of participants (players). Each player in PADUA game is defined as:

$$\forall a \in A \;\; a =< name_a, I_a, C_a, \Sigma_a > \tag{2}$$

Where:

(a) $name_a$:the player (agent) name.
(b) I_a:the set of items this player can understand (i.e. the items included in the player's database).
(c) C_a: the set of classes this player tries to prove that the discussed cases fall under. Each class $c \in C_a$ is a tuple <*name, value*> where name is an item $i \in I$, and $value \in V_i$ is the value this item the player tries to proove it holds.
(d) Σ_a: is a representation of the player's background database enables this player to mine for the suitable association rules as needed.

4. DP: Is the *dialogue purpose* of PADUA games, defined as the resolution of conflicting opinions about the classification of an instance $\varphi \subseteq I$, for example in the case of two players (the proponent pro and the opponent opp), the proponent may claim that the case falls under some class $c_1 \in C_{pro}$, while the opponent opposes the proponent's claim, and tries to prove that case actually falls under some other class $c_2 \in C_{opp}$ such that $c_2 \neq c_1$.
5. φ: The instance argued about i.e. the dialogue subject.
6. $K \subseteq L_t$: The *dialogue context* containing the knowledge that is presupposed and must be respected during a dialogue. The context is assumed consistent and remains the same throughout a dialogue.
7. L: The *logic* for L_t .

8. E: The *effect rules* for L_c, specifying for each move $md < p, S, m, t > \in DM$ its effects on the commitments of the participants. We will not discuss effect rules in detail here as they do not relate directly to the subject of this paper.
9. P: A protocol for L_c specifying the legal moves at each stage of a dialogue. P is defined formally as the function: $P \quad M \rightarrow 2^M$, where M is the set of dialogue acts (moves). Table 1 lists the possible next moves after each move in PADUA protocol.
   (a) *Termination Rules*: in this version of PADUA, the dialogue ends when a player fails to play a legal move in its turn, in this case, this particular player loses the game while the other player wins it.
   (b) *Turn taking Rules*: The current PADUA game applies a simple turn taking rule, in which each player is allowed to play exactly one move (speech act) in its turn. The turn taking in PADUA shifts uniformly to all the agents in the dialogue.
10. O: The *Outcome rules* of PADUA dialogues define for each dialogue d and instance φ the winners and losers of d with respect to instance .
11. S: The *Strategy function*

Move	*Label*	*Next Move*	*New Rule*
1	Propose Rule	3, 2, 4	yes
2	Distinguish	3, 5, 1	yes
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

3.1. Nested Dialogues

PADUA allows for dialogues to be nested so that a number of secondary dialogues take place to solve the disputes over some intermediate classifications, before the main dialogue over the main classification starts. To formalize this concept a Control Layer is implemented into the PADUA system. The aim of this layer is controlling the arrangements of the main and secondary dialogues; this layer also manages the communication among the players of every dialogue, to cover the cases in which some players are engaged only in some dialogues, and not in all of them. This layer has been kept as simple as possible, mainly because PADUA dialogues are basically of a persuasive nature. The formalization of the PADUA control layer is defined in the terms of the following components:

1. *Players*: is the set of players engaged in all the PADUA dialogues controlled by this layer.
2. Gs: set of PADUA secondary dialogue games. Each $gs \in Gs$ is defined as an instance of PADUA framework.
3. gm: PADUA main dialogue game, defined as an instance of PADUA framework.
4. *start*: a function that begins a certain PADUA dialogue game, start($gs \in Gs$) begins a secondary dialogue game, while start(gm) begins the main dialogue.

4. Example

To illustrate experimentally the kinds of dialogues produced by PADUA, we applied PADUA to a fictional welfare benefit scenario, where benefits are payable if certain conditions showing need for support for housing costs are satisfied.This scenario is intended to reflect a fictional benefit Retired Persons Housing Allowance (RPHA), which is payable to a person who is of an age appropriate to retirement, whose housing costs exceed one fifth of their available income, and whose capital is inadequate to meet their housing costs. Such persons should also be resident in this country, or absent only by virtue of "service to the nation", and should have an established connection with the UK labour force. These conditions need to be interpreted and applied [6]. We use the following desired interpretations:

1. Age condition: "Age appropriate to retirement" is interpreted as pensionable age: 60+ for women and 65+ for men.
2. Income condition:"Available income" is interpreted as net disposable income, rather than gross income, and means that housing costs should exceed one fifth of candidates' available income to qualify for the benefit.
3. Capital condition: "Capital is inadequate" is interpreted as below the threshold for another benefit.
4. Residence condition: "Resident in this country" is interpreted as having a UK address.
5. Residence exception: "Service to the Nation" is interpreted as a member of the armed forces.
6. Contribution condition: "Established connection with the UK labour force" is interpreted as having paid contributions in 3 of the last 5 years.

These conditions fall under a number of typical conditions' types: conditions (2 and 3) represent necessary conditions over continuous values while conditions (4 and 5) represent a restriction and an exception to the the applicant's residency, condition (1) deals with variables depending on other variables and condition (6) is designed to test the cases in which it is sufficient for some n out of m attributes to be true (or have some predefined values) for the condition to be true

A major problem with benefits such as this is that they are often adjudicated by a number of different offices and exhibit a high error rate due to the misunderstanding of the legalisation. This yields large data sets which contain a significant number of misclassifications, the nature of which varies from office to office. To test how PADUA can cope with this situation artificial RPHA benefits datasets (each comprises of 12,000 records) were generated to mimic different systematic misapplications of the rules, such that one does not consider the exceptions to the residency condition (i.e. only UK residents are considered valid candidates for housing benefits), while another interprets the "established connection with the UK labour force" as having paid contributions in 3 of the last 6 years rather than 5. The purpose of this test was to find out whether the proposed dialogue game helps in correctly classifying examples and henceforth correctly interprets them, even when the two agents are depending on (completely or partially) falsely classified examples, this could facilitate the sharing of best practice between offices. Each dataset was assigned to a PADUA player, corresponding association rules were mined from these sets using a 70% confidence threshold for both players, and PADUA was applied to different sets of examples each of which focuses on an exception of one of the six conditions mentioned above.

In the example discussed below the applicant is a male aged around 70 years, a UK resident who satisfies all the entitlement conditions except that he had paid contributions to the UK labour force in three out of the last 6 years (namely last year, the year before that and 6 years before), this is the case φ argued about between the game players ($A = \{proponent, opponent\}$) ; the datasets we use are the ones described in the last paragraph. The dialogue purpose DP is to decide whether this applicant is entitled to housing benefit or not, where the proponent says he does not($C_{proponent} = \{(entitles, no)\}$) noy while the oppoenet thinks he does ($C_{opponent} = \{(entitles, yes)\}$).

The dialogue starts with the proponent proposing the rule (`R1: contr y5= not paid -> entitles= no`) with a confidence= 73.14%, the opponent then tries to distinguish this rule in the light of to its own experience. For the opponent the rule (`R2: contr y4= not paid, contr y5= not paid -> entitles= no, capital > 3000`) holds with confidence = 2.34% only. This is true because the opponent uses an incorrect interpretation based on its own data, in which the sixth contributions year in considered. This last move is defeated by the proponent by the unwanted consequences attack (`capital>3000` does not hold). The opponent then proposes a counter rule (`R4: age>=65, residence= UK, gross disposable income <20%, 2500 < capital <3000 -> entitles= yes`) with 77.11% confidence, but the proponent can successfully distinguish this rule by emphasizing the fact that the candidate has not paid the contribution fees in the fifth year. The dialogue then progresses in a similar way with the proponent focusing on the unpaid contributions and the opponent trying to get away from this topic in accordance with their own interpretation. For example the proponent proposes the rules: (`R13: contr y3= not paid, contr y5= not paid -> entitles= no`) (88.77% confidence), (`R21: gender= male, contr y3=not paid, contr y5= not paid -> entitles= no`) (confidence = 89.39%). Finally the proponent puts forward the rule (`R23: contr y3= not paid, contr y4= not paid, contr y5= not paid ->entitles= no`), with a confidence of 89.39%, and this rule successfully exposes the mistake in the case under discussion, as by playing this rule the proponent manages to indicate the three years in which the contributions were not paid. The opponent tries then to distinguish this rule by manipulating its premises so it plays the rule (`R24: gender= male, contr y3= not paid, contr y4= not paid, contr y5= not paid -> entitles= no, contr y2= not paid`) in which the confidence falls to 37.89%, but again the opponent's move is defeated by the unwanted consequences attack (the second year contribution is actually paid). The dialogue ends here as the opponent fails to defeat the rule R23 and the proponent wins the game, and the candidate is classified as not entitled to the housing benefits. This game takes 24 moves.

Unfortunately when n out of m attributes are needed to decide whether a condition is satisfied or not, like contribution years in our example it is not always the case that the classification process will run correctly. It is more reliable to allow for an intermediate nested dialogue over the contribution years factor, which gives as a result the status of the contribution condition (true or false) before a main dialogue takes place over the eligibility of the applicant. For example if, we take the case $\varphi 2$ of a male applicant that satisfies all the conditions except for the contribution condition as he paid only the contribution fees of the third, fourth and the sixth years, and apply the one-dialogue

PADUA to this case between the same proponent and opponent as in the last example, the proponent fails to correctly classify the candidate status even after a very exhaustive 30 step dialogue in which contribution years are considered as independent factors and thus the classification is affected directly, as can be shown by some of the rules played in the dialogue:

```
R1-proponent-Propose Rule: contr y5= not paid -> entitles= no
confidence= 73.14.
R23-proponent-Propose Rule: gender=male, contr y2= not paid,
contr y5= not paid -> entitles= no
confidence=87.69.
R29-proponent-Propose Rule: residence=UK, contr y1= not paid,
contr y2= not paid, contr y5= not paid -> age>=65, entitles= no
confidence=95.31
R30-opponent-Counter Rule: age>=65, residence=UK, contr y3= paid,
net disposable income <20%, capital <2500 -> entitles= yes
confidence=96.82%
```

The latter rule is the final rule in the dialogue, as the proponent fails to defeat it using any of the valid PADUA attacks. Figure 1 shows how, by applying two dialogues (nested and main) to the same case, the proponent becomes able to win the game; and that by winning the nested dialogue over contribution years first, can then apply the result of that dialogue to the main dialogue.

Nested Dialogue	**Main Dialogue**
(1) – proponent - Propose Rule {contribution y1= not paid, contribution y5= not paid} --> {contribution = no} confidence=74.71 (2) – opponent - Distinguish {contribution y1= not paid, contribution y3= paid, contribution y5= not paid} --> {contribution = no} confidence=30.0 (3) – proponent - Increase Confidence {contribution y1= not paid, contribution y2= not paid, contribution y3= paid, contribution y5= not paid} --> contribution = no} confidence=100.0 (4) – opponent - Distinguish {contribution y1= not paid, contribution y2= not paid, contribution y3= paid, contribution y5= not paid, contribution y6= paid} --> {contribution = no} confidence=30.0 (5) – proponent - Increase Confidence {contribution y1= not paid, contribution y2= not paid, contribution y3= paid, contribution y4= paid, contribution y5= not paid, contribution y6= paid} --> {contribution = no} confidence=100.0 -- proponent wins	(1) – proponent - Propose Rule {contribution = no} --> {age>=65, entitles=no} confidence=94.0 (2) – opponent - Distinguish {gender = male, contribution = no} --> {age>=65, entitles=no, 2500<capital <3000} confidence=18.85 (3) – proponent - Unwanted Consequences {gender = male, contribution = no} --> {age>=65, entitles=no, 2500<capital <3000} confidence=18.85 -- proponent wins

Figure 1 Nested and Main Dialogues

5. Discussion

PADUA provides a means for agents to engage in discussion about a classification on the basis of raw data, unmediated by knowledge representation effort, to present this data in the form of rules. PADUA necessarily has significant differences from the existing protocols designed to argue about knowledge represented as rules, and the resulting dialogues have a flavour akin to dialogues related to case based reasoning in law. The protocol is particularly applicable to domains in which there are large volumes of data available, where it would prove unrealistic to craft a knowledge base. PADUA can thus complement rule based protocols, since its performance is actually enhanced by large volumes of data, whereas, for example, the work of [8], which used dialogue to generate a rule based theory, can only be applied to comparatively small datasets. Moreover PADUA is ideal for applications with several distributed datasets generated from different samples, since it can exploit and reconcile any systematic differences in the underlying data available to the dialogue participants. Also the work suggested in [10] to generate defeasible and strict rules using asocciation rule mining techniques is limited to small datasets, other restrictions are forced on the datasets used in this work such as they should have no missing values and that all values are correctly recorded.

As it can be viewed as a dialogue game, there is also the question of what strategies and tactics the participants should adopt. Some preliminary work has been done on this [20], there it was shown that the participants can, for example, be represented as cooperative or adversarial. The reported experiments confirm that different strategies give rise to different flavours of dialogue. Some have the flavour of persuasion dialogues, others of deliberation dialogues, demonstrating how these distinct types of dialogue, identified by Walton and Krabbe [19], can be realised in the same protocol when different strategies are used. Further experiments will explore questions relating to how strategies impact on the quality of decisions and the quality of justifications. In [15] an argumentation framework for learning agent is proposed: this framework is similar to PADUA in taking the experience, in the form of past cases, of agents into consideration and focusing on the argument generation process. Yet, the suggested protocol applies learning algorithms techniques, while PADUA implements simpler association rule mining techniques to produce arguments. Also the protocol in [15] is designed for pairs of agents that collaborate to decide the joint solution of a given problem, while PADUA can be applied in variety of situations including persuasion, deliberation and classification.

An important topic of discussion in recent work on reasoning with cases in law is the notion of intermediate predicates (see [5] and [11]). In [11] the important distinction is made between intermediate predicates which are truth functionally determined by some base level predicates, and those for which there is no simple truth functional relationship. For these latter kind of intermediate predicates, it may be necessary to first agree their application before deciding the main question. This is accommodated in PADUA through the possibility of nested dialogues, and the improvements gained were illustrated by an example in the previous section. While this does require some degree of domain analysis to identify and organize the intermediate predicates, so as to form what is termed in IBP [5] a “logical model” of the domain, this analysis is at a high level and, as in IBP, does not require the consideration of individual cases. Once identified this “logical model” can be used by the control layer of PADUA to set the agenda for the dialogue.

Future work will next focus on a set of empirical experiments using a variety of datasets interpreted using a range of misinterpretations and misinterpretations mixtures to further examine how PADUA can reconcile them. For example we wish to understand how much noise can be tolerated. We also intend to extend PADUA to more than two players as we expect interesting dialogues to come out of such applications, and this is a typical need in the scenarios to which PADUA was applied. Moreover in situations where cases can be classified into more than two categories, adding more players to the game, so that each possibility can have its own advocate, provides a promising solution to such classification problems.

References

[1] R. Agrawal, T. Imielinski, A.N. Swami: Association rules between sets of items in large databases. In: Proc. of ACM SIGMOD Int. Conf. on Management of Data, Washington, (1993), 207-216.

[2] V. Aleven: Teaching Case Based Argumentation Through an Example and Models. PhD thesis, University of Pittsburgh, Pittsburgh, PA, USA. 1997

[3] L. Amgoud, S. Belabbès, and H. Prade: A formal general setting for dialogue protocols. In: Proc. of AIMSA'06 12th Int. Conf. on Artificial Intelligence: Methodology, Systems, Applications, Varna, Bulgaria, (2006), 13 - 15.

[4] K. D. Ashley: Modeling Legal Argument. MIT Press, Cambridge, MA, USA, 1990.

[5] K. D. Ashley and S. Brüninghaus: A Predictive Role for Intermediate Legal Concepts. In: Bourcier D. (ed) Proceedings of Jurix 2003, IOS Press: Amsterdam (2003), 153-162.

[6] T.J.M. Bench-Capon: Knowledge Based Systems Applied To Law: A Framework for Discussion. In: T.J.M. Bench-Capon (ed), Knowledge Based Systems and Legal Applications, Academic Press, (1991), 329-342.

[7] F. P. Coenen, P. Leng and G. Goulbourne: Tree Structures for Mining Association Rules. In: Journal of Data Mining and Knowledge Discovery, Vol 8, No 1, (2004), 25-51.

[8] A. Chorley, T. J. M. Bench-Capon: AGATHA: Using heuristic search to automate the construction of case law theories. Artificial Intelligence and Law 13(1), (2005), 9-51.

[9] G. Goulbourne, F. P. Coenen and P. Leng: Algorithms for Computing Association Rules Using A Partial-Support Tree. In: Proc. of ES99, Springer, London, UK, (1999), 132-147.

[10] G. Governatori and A. Stranieri. Towards the Application of Association Rules for Defeasible Rules Discovery. In jurix, Amesterdam (2001). 63-75.

[11] L. Lindahl and J. Odelstad: Normative positions within an algebraic approach to normative systems. In Journal of Applied Logic, 17 2, (2005).

[12] J. Mackenzie: Question-begging in non-cumulative systems. In: Journal of Philosophical Logic 8, (1979), 117-133.

[13] P. McBurney and S. Parsons: Games That Agents Play: A Formal Framework for Dialogues between Autonomous Agents. In: Journal of logic, language and information, 11(3), (2002), 315-334.

[14] M Minsky and S. Papert: Perceptrons: An Introduction to Computational Geometry. Cambridge MA: MIT Press, (1969).

[15] S. Ontañón and E. Plaza. Arguments and Counterexamples in Case-Based Joint Deliberation. In ArgMAS Hakodate, Japan (2006). 36-53.

[16] H. Prakken: On dialogue systems with speech acts, arguments, and counterarguments. In: Proc. of the 7th European Workshop on Logic for Artificial Intelligence, no. 1919 in Springer Lecture Notes in AI, Berlin. Springer Verlag, (2000), 224-238.

[17] H. Prakken: Coherence and flexibility in dialogue games for argumentation. In: Journal of Logic and Computation 15 , (2005), 1009-1040.

[18] H. Prakken: Formal systems for persuasion dialogue. In: The Knowledge Engineering Review 21, (2006), 163-188.

[19] D. N. Walton and E. C. W. Krabbe: Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. SUNY Press, Albany, NY, USA, (1995).

[20] M. Wardeh, T. J. M. Bench-Capon and F. P. Coenen: PADUA Protocol: Strategies and Tactics. In Proc. ECSQARU , 9th European Conf. on Symbolic and Quantitative Approaches to Reasonong with Uncertainty, LNAI 4724, (2007), 465-476.