An Argumentation Framework for Learning, Information Exchange, and Joint-Deliberation in Multi-Agent Systems

Santiago Ontañón*and Enric Plaza
Artificial Intelligence Research Institute (IIIA-CSIC)
Campus UAB, 08193 Bellaterra, Catalonia (Spain),
{santi,enric}@iiia.csic.es

June 8, 2011

Abstract

Case-Based Reasoning (CBR) can give agents the capability of learning from their own experience and solve new problems, however, in a multi-agent system, the ability of agents to collaborate is also crucial. In this paper we present an argumentation framework (AMAL) designed to provide learning agents with collaborative problem solving (joint deliberation) and information sharing capabilities (learning from communication). We will introduce the idea of CBR multi-agent systems ($\mathcal{M}$AC systems), outline our argumentation framework and provide several examples of new tasks that agents in a $\mathcal{M}$AC system can undertake thanks to the argumentation processes.

1 Introduction

Case-Based Reasoning (CBR) [1] can give agents the capability of learning from their own experience and solve new problems [19]. Moreover, in a multi-agent system, the ability of agents to collaborate is crucial in order to benefit from the information known by other agents, both during learning and problem solving. In this paper we will present an argumentation framework designed for learning agents (AMAL), and show that agents can use it to learning and problem solving. On the one hand we will show that individual agent's learning can be enhanced through *learning from communication*, and on the other hand we will show that individual problem solving can also be enhanced by *joint deliberation*.

Learning agents are capable of learning from experience, in the sense that past examples (situations and their outcomes) are used to predict the outcome

*Part of this work was done while the first author was at the Georgia Institute of Technology, Atlanta, USA for the situation at hand. For example, agents might be able to predict the species of a given animal from observing the animal features thanks to reasoning about past observations of animals. In general, in this work, we focus on any kind of predictive classification tasks. However, since individual agents experience may be limited, individual knowledge and prediction accuracy is also limited. Thus, learning agents that are capable of arguing their individual predictions with other agents may reach better prediction accuracy after such an argumentation process. Specifically, joint deliberation involves discussion over the outcome of a particular situation or the appropriate course of action for a particular situation.

Existing argumentation frameworks can be classified in two groups: *abstract argumentation frameworks*, derived from Dung's seminal work [7], or logical argumentation frameworks [4]. Argumentation frameworks focus on how to assess, given a set of arguments, which arguments are acceptable, and which of them are defeated by other arguments. However, these argumentation frameworks do not focus on how are arguments *constructed*, i.e. where do they come from. Moreover, all additional pieces of knowledge required for argumentation, like the attack relation among arguments, or any required preference relation are assumed to be given in advance. In this paper, we focus on an *Argumentation-based Multi-Agent Learning* (AMAL) framework where both arguments and preference relation relation are learned from experience.

Specifically, we consider a scenario with agents that (1) work in the same domain using a shared ontology, (2) are capable of learning from examples, and (3) communicate using an argumentative framework. We present a case-based approach to address both: how learning agents can generate arguments from examples, and how they can define a preference relation among arguments based on examples. Our agents use case-based reasoning (CBR) [1] to learn from past experience, represented as a set of examples or *cases* (where a case is a situation and its outcome) in order to predict the outcome of a new situation. We propose an argumentation protocol inside the AMAL framework to support agents in reaching a joint prediction over a specific situation or problem — moreover, the reasoning needed to support the argumentation process will also be based on cases. Finally, we present several applications where the argumentation framework can be useful. First we will show how using argumentation agents can achieve joint deliberation, and we'll see how agents can act as committees or as an information market. Then we will show how agents can use argumentation as an information sharing method, and achieve effective learning from communication, and information sharing among peers.

This paper is an extended version of the results informally presented in [14]. The paper is structured as follows. Section 2 introduces our multi-agent CBR ($\mathcal{M}$AC) framework. After that, Section 3 briefly describes our argumentation framework. Section 4 presents several applications of the argumentation framework, and finally Section 5 presents related work. The paper closes with related work and conclusions sections.

2 Multi-Agent Case-Based Reasoning Systems

[Figure 1 about here.]

A *Multi-Agent Case Based Reasoning System* ($\mathcal{M}\mathsf{AC}$) $\mathcal{M} = \{(A_1, C_1), ..., (A_n, C_n)\}$ is a multi-agent system composed of $\mathcal{A} = \{A_i, ..., A_n\}$, a set of CBR agents, where each agent $A_i \in \mathcal{A}$ possesses an individual case base C_i, as illustrated in Figure 1. Each individual agent A_i in a $\mathcal{M}\mathsf{AC}$ is completely autonomous and each agent A_i has access only to its individual and private case base $C_i = \{c_1, ..., c_m\}$ consisting of a collection of cases. CBR methods solve new problems by retrieving similar problems stored in a case base, where each *case* is a previously solved problem. Once a set of problems has been retrieved, the solution to the problem at hand is computed by reusing the solution contained in the retrieved cases (adapting or combining those solutions if needed). The newly solved problem might be incorporated into the case base as another case.

Agents in a $\mathcal{M}\mathsf{AC}$ system are able to individually solve problems by using case-based reasoning. In this paper we will limit our selves to analytical tasks, where solving a problem means to identify a particular *solution class* among a set of possible solutions. For example, diagnosing a patient with the right disease, classifying a customer in the right risk category for a loan, etc.

CBR gives agents the capability to individually learn how to solve these kinds of tasks from experience, however, in a multi-agent system where each agent is exposed to different experiences we would like agents to collaborate and make use of information known by other agents. However, we are not interested in complete information sharing, but in a selective information sharing that only shares the information that is needed for the task at hand, thus keeping the amount of information each agent knows and has to share manageable.

The AMAL framework presented in this paper complements $\mathcal{M}\mathsf{AC}$ systems by allowing agents to perform joint deliberation (solve classification tasks in a collaborative way) and learning from communication.

3 Argumentation-Based Multi-Agent Learning: AMAL

The AMAL argumentation framework is based on the idea that, when CBR agents solve new problems, they can provide a *justification* of the predicted solution. These justifications can then be used as arguments. The kinds of arguments that CBR agents can generate are thus based on justifications and cases. For example, in a medical domain, a justification could be "I predict patient A has pneumonia because he coughs and has night sweat." In the following sections we will define this idea of justifications, and then define the different components of the AMAL framework: the set of argument types that agents can use, a preference relation based in cases, and finally the AMAL argumentation protocol.

3.1 Justified Predictions

[Figure 2 about here.]

The basis of the AMAL framework is the ability of some machine learning methods to provide *explanations* (or *justifications*) for their predictions. We are interested in justifications since they can be used as arguments. Most of the existing work on explanation generation focuses on generating explanations to be provided to the user. However, in our approach we use explanations (or justifications) as a tool for improving communication and coordination among agents.

In particular in the AMAL framework agents use CBR as their learning and problem solving method. Since CBR methods solve problems by retrieving cases from a case base, when a problem P is solved by retrieving a set of cases $c_1, ..., c_n$, the justification D will contain the relevant information from the problem P that made the CBR system retrieve that particular set of cases, i.e. it will contain the relevant information that P and $c_1, ..., c_n$ have in common, but no other case in the case-base does. More formally:

Definition 3.1 *A justification D built by an agent A_i to justify a prediction S for a problem P, solved by retrieving a set of cases $C_P = \{c_1, ..., c_n\} \subseteq C_i$ is a symbolic description D such that $D \sqsubseteq P$ and $\forall c_j \in C_P : D \sqsubseteq c_j$ and $\forall c_j \in C_i \wedge c_j \notin C_P : D \not\sqsubseteq c_j$, i.e. D is satisfied ($\sqsubseteq$) by P and by all the retrieved cases, but by no other case in the case base of A_i.*

So, when an agent solves a problem providing a justification for its solution, it generates a *justified prediction*. A *Justified Prediction* is a tuple $J = \langle A, P, S, D\rangle$ where agent A considers S the correct solution for problem P, and that prediction is justified by a symbolic description D. Justifications can have many uses for CBR systems [16, 18]. In this paper, we are going to use justifications as arguments, in order to allow learning agents to engage in argumentation processes.

For instance, Figure 2 shows a real justification generated by LID [3], a CBR method capable of generating justifications, after solving a problem P in the domain of marine sponges identification. In particular, Figure 2 shows how when an agent A_1 receives a new problem to solve (in this case, a new sponge to determine its order), the agent uses LID to generate a justified prediction using the cases in the case base of the agent. The justification shown in Figure 2 can be interpreted saying that "the predicted solution to problem P is hadromerida because the smooth form of the megascleres of the spiculate skeleton of the sponge is of type tylostyle, the spiculate skeleton of the sponge has no uniform length, and there is no gemmules in the external features of the sponge".

3.2 Arguments and Counterarguments

For our purposes an *argument* α generated by an agent A is composed of a statement S and some evidence D supporting S as correct. In the context of

$\mathcal{M}$AC systems, agents argue about predictions for new problems and can provide two kinds of information: a) specific cases $\langle P, S\rangle$, and b) justified predictions: $\langle A, P, S, D\rangle$. Using this information, we can define two types of arguments: justified predictions, and counterexamples:

- A *justified prediction* α is generated by an agent A_i to argue that A_i believes that the correct solution for a given problem P is $\alpha.S$, and the evidence provided is the justification $\alpha.D$.
- A *counterexample* c is a case that contradicts an argument α. Thus a counterexample is a counterargument, one that states that a specific argument α is not always true, and the evidence provided is the case c that is a counterexample of α.

.

In our framework, justified predictions can also be used as counterarguments. A *counterargument* β is an argument offered in opposition to an argument α. Counterarguments can be either justified predictions or counterexamples. For example, a justified prediction $\beta = \langle A_j, P, S', D'\rangle$ generated by an agent A_j with the intention to rebut an argument α generated by another agent A_i, that endorses a solution class S' different from that of $\alpha.S$ for the problem at hand is a counterargument to α. The only restriction we impose for a justified prediction to be a counterargument is that $D \sqsubseteq D'$, i.e. the counterargument must be more specific than the original argument.

[Figure 3 about here.]

For example Figure 3 shows a justified prediction, generated as a counterargument to the justified prediction shown in Figure 2. Notice that the justification in the counterargument is more specific than the justification in the original argument.

3.3 Case-Based Preference Relation

[Figure 4 about here.]

A specific argument provided by an agent might not be consistent with the information known to other agents (or even to some of the information known by the agent that has generated the justification due to noise in training data). This means that other agents might have cases in their case-bases which contradict arguments generated by other agents. Therefore, it is possible that different agents generate justified predictions for the same problem which predict different solutions. For that reason, we will define a preference relation over contradicting justified predictions based on cases. Basically, we will define a *confidence* measure for each justified prediction and the justified prediction with the highest confidence will be the preferred one. In the absence of further evidence (e.g. counterarguments or counterexamples attacking some of the justified predictions), when agents are faced with competing arguments for a given problem, the preference relation can be used to determine which argument to accept.

The idea behind case-based confidence is to count how many of the cases in an individual case base *endorse* a justified prediction, and how many of them are counterexamples of it. The more the endorsing cases, the higher the confidence; and the more the counterexamples, the lower the confidence. Specifically, an agent estimates the confidence of an argument as:

$$\mathsf{C}_{A_i}(\alpha) = \frac{Y_\alpha^{A_i}}{1 + Y_\alpha^{A_i} + N_\alpha^{A_i}}$$

where $Y_\alpha^{A_i}$ are the set of cases in the case base of A_i that endorse α and $N_\alpha^{A_i}$ is the set of its counterexamples in the case base of A_i, defined as follows:

- $Y_\alpha^{A_i} = |\{c \in C_i|\ \alpha.D \sqsubseteq c.P \wedge \alpha.S = c.S\}|$ is the number of cases in the case base of A_i *subsumed* by the justification $\alpha.D$ that belong to the solution class $\alpha.S$,
- $N_\alpha^{A_i} = |\{c \in C_i|\ \alpha.D \sqsubseteq c.P \wedge \alpha.S \neq c.S\}|$ is the number of cases in the case base of A_i *subsumed* by justification $\alpha.D$ that do *not* belong to that solution class.

Figure 4 illustrates the individual evaluation of the confidence of an argument, in particular, three endorsing cases and one counterexample are found in the case base of agents A_i, giving an estimated confidence of 0.6.

Moreover, we define the *joint confidence* of an argument α as the confidence computed using the cases present in the case bases of all the agents in the group:

$$\mathsf{C}(\alpha) = \frac{\sum_i Y_\alpha^{A_i}}{1 + \sum_i \left(Y_\alpha^{A_i} + N_\alpha^{A_i}\right)}$$

In AMAL, agents use this joint confidence as the preference relation: a justified prediction α is preferred over another one β if $C(\alpha) \geq C(\beta)$.

3.4 The AMAL Argumentation Protocol

The main idea behind joint deliberation is to follow the problem solving procedure used by committees of humans, using a two stage process: in a first stage (deliberation), agents expose their point of view and argue about it; in a second stage (voting), agents take into account all the previously exposed arguments to cast a vote and decide on a final solution for the problem at hand. If all agents reach an agreement during deliberation, there is no need for voting. The AMAL interaction protocol outlined in this section exactly models this process (for a more formal description, see [17]).

The interaction protocol of AMAL allows a group of agents $A_1, ..., A_n$ to deliberate about the correct solution of a problem P by means of an argumentation process. If the argumentation process arrives to a consensual solution, the joint deliberation ends; otherwise a weighted vote is used to determine the joint solution. Moreover, AMAL also allows the agents to learn from the counterexamples received from other agents. Thus, letting the agents learn from communication during deliberation.

The AMAL protocol consists on a series of rounds. At each round, each agent holds one single justified prediction as its preferred prediction. In the initial round, each agent generates its individual justified prediction for the current problem P and uses it as its initial preferred prediction.

Then, at each round t each agent agent has a chance to rebut the prediction made by any of the other agents. The protocol uses a token passing mechanism so that agents (one at a time) can send counterarguments or counterexamples if they disagree with the prediction made by any other agent. Specifically, each agent is allowed to send one counterargument or counterexample each time it gets the token (notice that this restriction is just to simplify the protocol, and it does not restrict the number of counterargument an agent can sent, since they can just be delayed for subsequent rounds). When an agent receives a counterargument or counterexample, it informs the other agents if it accepts the counterargument (and changes its prediction) or not. Agents take that decision based on the preference relation: when the received counterargument is preferred to the currently held argument, the counterargument is accepted, otherwise it is not, and the agent will try to generate a counterargument to it. Thus, agents have also the opportunity to answer to counterarguments when they receive the token, by trying to generate a counterargument to the counterargument.

When all the agents have had the token once, the token returns to the first agent, and so on. If at any time in the protocol, all the agents agree or during the last n rounds no agent has generated any counterargument, the protocol ends. Moreover, if at the end of the argumentation the agents have not reached an agreement (an agreement is reached when the arguments that all the agents are holding at a particular round endorse the same solution), then a voting mechanism that uses the confidence of each prediction as weights is used to decide the final solution. Specifically, if H_t is the set of arguments that all the agents are holding in the last round t of the protocol, the final solution is defined as:

$$S = \arg\max_{S_k \in \mathcal{S}} \sum_{\alpha_i \in H_t | \alpha_i.S = S_k} C(\alpha_i)$$

Moreover, notice that agents can learn from the counterexamples received from other agents during an argumentation process. As we will show in the next section, the counterexamples received by a particular agents are those ones that are in contradiction with the agent's predictions, and thus the ones where the different agents in the committee disagree about. It is well known from the *active learning* field in machine learning [5], that those are precisely the examples which can better improve the problem solving performance of agents. Thus, learning from counterexamples exchanged during deliberation is a very effective way to improve performance.

Notice that the main difference between the AMAL argumentation framework and existing argumentation framework such as Dung's [7], is that in existing frameworks arguments are given, and the goal is to decide which arguments to accept. In AMAL, arguments are not given, but generated from examples by the agents, and argumentation is used as a communication framework for agents to decide when to change their minds concerning the prediction for a specific problem P.

4 Applications of AMAL

The AMAL argumentation framework gives agents in a $\mathcal{M}$AC system two new capabilities: joint deliberation and learning from communication. In this section we will present an evaluation of those two capabilities, in addition to a third evaluation where agents use AMAL as an "information sharing" mechanism.

4.1 Joint Deliberation

[Figure 5 about here.]

To evaluate the joint deliberation capabilities of agents using AMAL we designed the following experiment. A machine learning data set is divided in two disjoint sets: a training set and a test set. The training set distributed among 5 agents without replication (the training set is split in 5 disjoint parts and each agent only has access to one of them). Then, one of the agents is given a problems from the test set (not in the training set) and is asked to solve it. Such agent will engage in an argumentation process with some other agents in the system about the correct solution for each problem. We compare how accurate the prediction is using argumentation with respect to traditional voting mechanisms, and also study how much the number of agents that take part in the argumentation affects the prediction accuracy. The agents in our experiments use LID to solve problems and generate justifications.

We performed experiments with two different data sets: *soybean* (a propositional data set from the UCI machine learning repository) and *demospongiae* (a complex relational data set also from the UCI machine learning repository). The soybean data set has 307 examples and 19 solution classes, while the sponge data set has 280 examples and 3 solution classes. In the testing stage, problems in the test set are sent randomly to one of the agents, and its goal is to predict the correct solution.

We ran experiments using 2, 3, 4, and 5 agents respectively (in all experiments each agent has 20% of the training data, since the training is always distributed among 5 agents). Thus, in our experiments with the soybean data set, each agent has about 55.26 cases each, and in the demospongiae dataset, each agent has about 50.4 cases each.

Figure 5 shows the result of those experiments. For each number of agents, three bars are shown: *individual*, *Voting*, and AMAL. The individual bar shows the average accuracy of individual agents' predictions; the voting bar shows the average accuracy of the joint prediction achieved by voting but without any argumentation; and finally the AMAL bar shows the average accuracy of the joint prediction using argumentation. The results shown are the average of 5 10-fold cross validation runs.

Figure 5 shows that collaboration (voting and AMAL) outperforms individual problem solving. Moreover, as we expected, the accuracy improves as more agents collaborate, since more information is taken into account. We can also see that AMAL always outperforms standard voting, proving that joint decisions are based on better information as provided by the argumentation process.

For instance, the joint accuracy for 2 agents in the sponge data set is of 87.57% for AMAL and 86.57% for voting (while individual accuracy is just 80.07%). Moreover, the improvement achieved by AMAL over Voting is even larger in the soybean data set. The reason is that the soybean data set is more "difficult" (in the sense that agents need more data to produce good predictions). These experimental results show that AMAL effectively exploits the opportunity for improvement: the accuracy is higher only because more agents have changed their opinion during argumentation (otherwise they would achieve the same result as Voting).

4.2 Learning from Communication

[Figure 6 about here.]

[Table 1 about here.]

In order to evaluate the learning from communication capabilities of agents using AMAIL , we run the following additional experiment. Using the same scenario as the previous experiment, we distributed 25% of the training set among the five agents; after that, the rest of the cases in the training set is sent to the agents one by one (each case sent at random to one agent); when an agent receives a new training case c, it has several options: 1) the agent can discard it, 2) the agent can retain it, or 3) the agent can use it for engaging in an argumentation process. This last option means that the agent takes the new case c, consisting on a problem P and its solution S, and starts a deliberative agreement process to try to predict the solution of P (ignoring the fact that the solution is already known). This is basically used to create an opportunity for learning from communication.

We compared the evolution of the individual classification accuracy of agents that perform each one of these 3 options. Figure 6 contains three plots, where NL (not learning) shows accuracy of an agent with no learning at all; L (learning), shows the evolution of the individual classification accuracy when agents learn by retaining the training cases they individually receive (notice that when all the training cases have been retained, the accuracy should be equal to that of Figure 5 for individual agents); and finally LFC (learning from communication) shows the evolution of the individual classification accuracy of learning agents that also learn by retaining those counterexamples received during argumentation (i.e. they learn both from training examples and counterexamples received during argumentation).

Figure 6 shows that if an agent A_i learns also from communication, A_i can significantly improve its individual performance with just a small number of additional cases (those selected as relevant counterexamples for A_i during argumentation). For instance, in the soybean data set, individual agents have achieved an accuracy of 70.62% when they also learn from communication versus an accuracy of 59.93% when they only learn from their individual experience. The number of cases learnt from communication depends on the properties of the data set: in the sponges data set, agents retained only very few additional cases, and significantly improved individual accuracy; namely they retain 59.96 cases in average (compared to the 50.4 cases retained if they do not learn from communication). In the soybean data set more counterexamples are learnt to significantly improve individual accuracy, namely they retain 87.16 cases in average (compared to 55.27 cases retained if they do not learn from communication). Finally, the fact that both data sets show a significant improvement points out the adaptive nature of the argumentation-based approach to learning from communication: the useful cases are selected as counterexamples, and they have the intended effect.

Table 1 shows the average amount of cases that each agent has at the end of the experiments reported in Figure 6. The table shows that in some data sets, like the sponges one, agents retain only very few additional cases (59.96 versus 50.4) and that their individual accuracy has improved significantly. This fact indicates that the argumentation process provides a useful framework for learning from communication, finding which cases are specifically useful for each particular agent.

4.3 Information Sharing

Finally, a third use of the AMAL framework is for information sharing. To evaluate this capability, we performed some experiments in *prediction markets*. Prediction markets, also known as *information markets*, are an alternative to voting systems. The goal of a prediction market is to aggregate information based on a *price signal* emitted by the members of a group. The advantage of the price signal is that it encapsulates both the information and the preferences of a number of individuals. In this approach, the task of aggregating information is achieved by *creating a market*, and that market should offer the right *incentives* for the participating people or agents to disclose the information they possess.

Prediction markets provide agents with an incentive to provide accurate predictions (since they receive some bonus if they provide the right answer), therefore, it is rational for agents to consult with other agents, before casting their votes. Thus, we can distinguish two phases: an information gathering phase, where agents consult with some of their acquaintances, and a joint deliberation phase, where agents cast their votes for particular solutions, together with a price signal (the price signal can be seen as how much money the agent bets into the predicted solution, which is proportional to the reward the agent will get if its prediction is correct).

In this experiment we used AMAL as a framework for information sharing, and to evaluate it, we designed the following experiment: we split the training set among 8 agents, and each agent in the system had a small set of acquaintances with which it will share information before participating in the market. Their acquaintances are determined according to a social network, where each node is an agent and links represent which agents are acquaintances of which other agents. To perform information sharing, an agent does the following: it first generates its own individual prediction for the problem at hand using its local case base, and then it starts a one-to-one argumentation process with one of its acquaintances. The outcome of this argumentation is a more informed prediction than the original one. Using that prediction as a starting point, the agent engages in another one-to-one argumentation process with its next acquaintance, and so on. After each argumentation process, the resulting prediction is stronger and stronger since it takes into account information known by more agents (without the agents having to share their case bases). The resulting prediction is cast by the agent as its vote in the prediction market, and the joint confidence (computed during the argumentation processes) of that prediction is used to compute its price signal (the higher the confidence, the higher the price signal). In particular, in our experiments we used the following formula to compute the price signal of an agent A_i: $M \times C(\alpha)$, where M is the maximum amount an agent can bet in the prediction market.

[Table 2 about here.]

[Figure 7 about here.]

We have performed experiments with 0 to 5 acquaintances and logged the prediction accuracy of the market. Figure 7 shows three of the networks we used in our experiments, specifically the networks for 1, 2 and 3 acquaintances. The prediction accuracy of each individual agent, and also the average money reward received by each agent per problem when agents can bet between 0 and 100 monetary units per problem, i.e. $M = 100$, and all the agents that predicted the right solution split all the money that every agent bet (plus a 10% bonus).

Table 2 shows that information exchange is positive both for the individual agents and for the market as a whole. We can see that the more acquaintances an agent has, the higher its individual prediction accuracy. For instance, agents with 0 acquaintances have an accuracy of 74.21% while agents with 1 acquaintance have an accuracy of 83.99%, and when they have 5 acquaintances, their accuracy is increased to 88.21%. Moreover, the predictive accuracy of the market increases from 89.71% when agents do not perform information exchange, to above 91% when agents have more than 1 acquaintances.

Another effect we can observe is that the reward that the agents obtain increases when they perform information exchange, starting in 10.35 monetary units per problem when they do not perform information exchange, and going up to about 12 when agents have 2 or 3 acquaintances. It is interesting to notice that the performance of the prediction market doesn't increase linearly with the performance of the individual agents. In fact, the more accurate the individual agents get, the more correlated their individual predictions are, and thus there is less difference between their individual predictions and the prediction of the market as a whole. This is a well known effect in machine learning (known as the *ensemble effect* [6]), or in economics (related to the Condorcet Jury Theorem). Therefore, if the reward signal that the agents get was only related to its individual accuracy, agents might be interested in their classification accuracy to a point were the correlation is too high, and then the market would not achieve it's optimal accuracy. The reward signal takes this into account, and rewards the agents when the market as a whole has high accuracy.

Concerning information exchange, the experiments show that individual and market accuracy improve. This means that the agents make a more informed prediction, and thus that AMAL is effective in providing agents with enough information to correct previously inaccurate predictions.

5 Related Work

Concerning CBR in a multi-agent setting, the first research was on "negotiated case retrieval" [20] among groups of agents. Our work on multi-agent case-based learning started in 1999 [11]; later Mc Ginty and Smyth [12] presented a multi-agent collaborative CBR approach (CCBR) for planning. Finally, another interesting approach is *multi-case-base reasoning* (MCBR) [10], that deals with distributed systems where there are several case bases available for the same task and addresses the problems of cross-case base adaptation. The main difference is that our $\mathcal{M}$AC approach is a way to distribute the *Reuse* process of CBR (using a voting system) while *Retrieve* is performed individually by each agent; the other multi-agent CBR approaches, however, focus on distributing the *Retrieve* process.

Research on MAS argumentation focus on several issues like a) different models of argumentation, b) logics, protocols and languages that support argumentation. The two main models of argumentation are 1) that of Dung [7] and its derivatives, such as weighted argumentation frameworks [8], which focus on abstract argumentation mechanism, and 2) that of Simari et al. [4], which focus on logical models of argument based on defeasible logics. Other models, for instance based on BDI have also been proposed [22]. Although argument selection is a key aspect of automated argumentation (see [21] and [22]), most research has been focused on preference relations among arguments. In our framework we have addressed both argument selection and preference relations using a case-based approach.

Finally, concerning argumentation-based machine learning, Fukumoto and Sawamura [9] propose a new theoretical framework for argumentation-based learning, where they focus on what is the belief status of an agent after receiving a new argument. The main difference with our work is that they perform a theoretical analysis of the belief revision problem after receiving an argument, whereas we are concerned with the full problem of how to generate arguments, evaluate them, and learn from them, all based on learning from examples. Amgoud and Serrurier [2] propose an argumentation framework for classification where both examples and hypothesis are considered as arguments in the same way as in our framework. However, in their framework they focus on how to extract valid and justified conclusions from a given set of examples and hypothesis, where as in our framework we are concerned with how those hypothesis are also generated. Moreover, they only focus on the single agent situation. Other work has tried to improve the performance of machine learning methods by combining them with argumentation techniques, for example, Možina et al. [13] introduced the idea of argumented examples to improve the reduce the space of the hypothesis space and help producing more meaningful hypothesis.

6 Conclusions

While Case-Based Reasoning (CBR) can give agents the capability of learning from their own experience and solve new problems, in a multiagent setting, agents might need to communicate and collaborate with each other. In this paper we have presented an argumentation-based framework for multi-agent learning, AMAL, that allows a group of learning agents to perform joint deliberation, learning from communication and information sharing.

The main difference of this work with other works on computational argumentation is that in our framework, by combining learning techniques with argumentation, we focus on how can agents *generate* arguments from experience. In that sense, our work does not focus only in defining an argumentation framework, but on closing the loop of how do agents learn, generate arguments from experience, and communicate those arguments to other agents with the purpose of solving problems or of further improve their learning.

The main contributions of this work are: a) an argumentation framework for learning agents, where agents generate arguments from experience; b) a case-based preference relation over arguments, based on computing an overall confidence estimation of arguments; and c) an argumentation-based approach for learning from communication. Additionally, we reported empirical evaluations of the performance of AMAL in a collection of machine learning tasks.

Our future work follows two different paths. First, we plan to explore the situations where we have heterogeneous agents that use different learning methods to generate arguments, and we also plan to explore more realistic the effect of having non-trustable agents, that do not always reveal their truth information. Second, AMAL allows agents that use lazy learning techniques to perform joint deliberation; our second line of future work is the integration of argumentation techniques with eager inductive learning techniques. Work on this second line of research has already started, and we are currently working on an framework called AMAIL [15], focusing on inductive learning techniques rather than lazy learning ones.

Acknowledgements Support for this work came from projects MID-CBR TIN2006-15140-C03-01, and Agreement Technologies CSD2007-0022.

References

[1] A. Aamodt and E. Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *Artificial Intelligence Communications*, 7(1):39–59, 1994.

[2] Leila Amgoud and Mathien Serrurier. Arguing and explaining classifications. In Simon Parsons, Iyad Rahwan, and Chris Reed, editors, *Argumentation in Multi-Agent Systems*, volume 4946 of *Lecture Notes in Artificial Intelligence*, pages 164–177. Springer-Verlag, 2007.

[3] E. Armengol and E. Plaza. Lazy induction of descriptions for relational case-based learning. In *ECML'2001*, pages 13–24, 2001.

[4] Carlos I. Chesñevar and Guillermo R. Simari. Formalizing Defeasible Argumentation using Labelled Deductive Systems. *Journal of Computer Science & Technology*, 1(4):18–33, 2000.

[5] David A. Cohn, Les Atlas, and Richard E. Ladner. Improving generalization with active learning. *Machine Learning*, 15(2):201–221, 1994.

[6] T.G. Dietterich. Ensemble methods in machine learning. In J. Kittler and F. Roli, editors, *First International Workshop on Multiple Classifier Systems*, Lecture Notes in Computer Science, pages 1 – 15. Springer Verlag, 2000.

[7] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[8] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. Inconsistency tolerance in weighted argument systems. In *AAMAS '09: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, pages 851–858, 2009.

[9] Taro Fukumoto and Hajime Sawamura. Argumentation-based learning. In Nicolas Maudet, Simon Parsons, and Iyad Rahwan, editors, *Argumentation in Multi-Agent Systems*, volume 4766 of *Lecture Notes in Artificial Intelligence*, pages 17–35. Springer-Verlag, 2006.

[10] D. Leake and R. Sooriamurthi. Automatically selecting strategies for multi-case-base reasoning. In *ECCBR'2002*, pages 204–219. Springer Verlag, 2002.

[11] Francisco J. Martín, Enric Plaza, and Josep-Lluis Arcos. Knowledge and experience reuse through communications among competent (peer) agents. *International Journal of Software Engineering and Knowledge Engineering*, 9(3):319–341, 1999.

[12] Lorraine McGinty and Barry Smyth. Collaborative case-based reasoning: Applications in personalized route planning. In I. Watson and Q. Yang, editors, *ICCBR*, number 2080 in LNAI, pages 362–376. Springer-Verlag, 2001.

[13] Martin Možina, Jure Žabkar, and Ivan Bratko. Argument based machine learning. *Machine Learning*, 171:922–937, 2007.

[14] Santi Ontañón and Enric Plaza. Learning, information exchange, and joint-deliberation through argumentation in multi-agent systems. In *On the Move to Meaningful Internet Systems: OTM 2008 Workshops*, volume 5333 of *Lecture Notes in Computer Science*, pages 150–159. Springer-Verlag, Berlin, Heidelberg, 2008.

[15] S. Ontañón and E. Plaza. Multiagent inductive learning: an argumentation-based approach. In *International Conference on Machine Learning (ICML 2010)*, pages 839 – 846, 2010.

[16] Santi Ontañón and Enric Plaza. Justification-based multiagent learning. In *ICML'2003*, pages 576–583. Morgan Kaufmann, 2003.

[17] Santi Ontañón and Enric Plaza. Learning and joint deliberation through argumentation in multi-agent systems. In *Proceedings AAMAS 2007*, pages 971–978. ACM, 2007.

[18] Enric Plaza, Eva Armengol, and Santiago Ontañón. The explanatory power of symbolic similarity in case-based reasoning. *Artificial Intelligence Review*, 24(2):145–161, 2005.

[19] Enric Plaza and Santiago Ontañón. Learning collaboration strategies for committees of learning agents. *Journal of Autonomous Agents and Multi-Agent Systems*, 13:429–461, 2006.

[20] M. V. Nagendra Prassad, Victor R. Lesser, and Susan Lander. Retrieval and reasoning in distributed case bases. Technical report, UMass Computer Science Department, 1995.

[21] K. Sycara S. Kraus and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence Journal*, 104:1–69, 1998.

[22] N. R. Jennings S. Parsons, C. Sierra. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8:261–292, 1998.

List of Figures

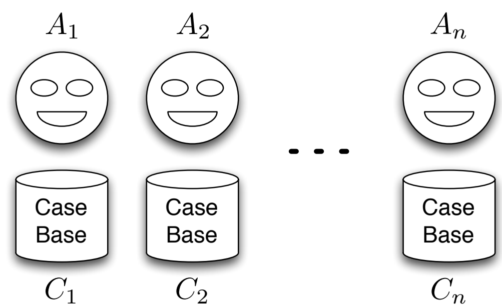

Figure 1: A Multi-Agent Case Based Reasoning System, composed of a collection agents with their respective case-bases.

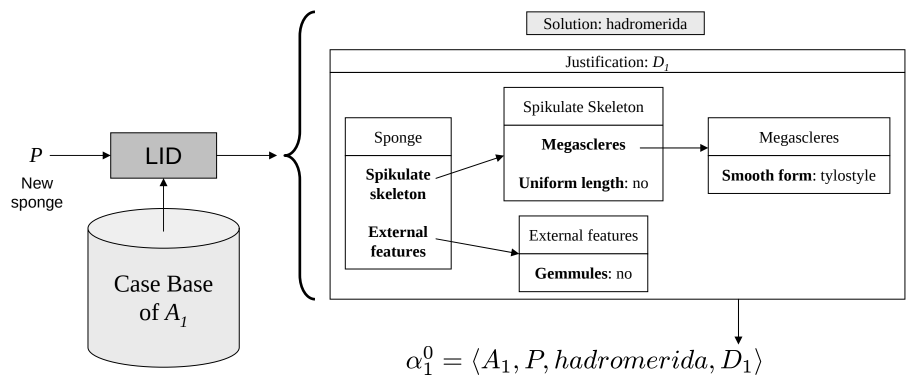

Figure 2: Example of a real justification generated by LID, a CBR method capable of generating justifications, in the marine sponges data set.

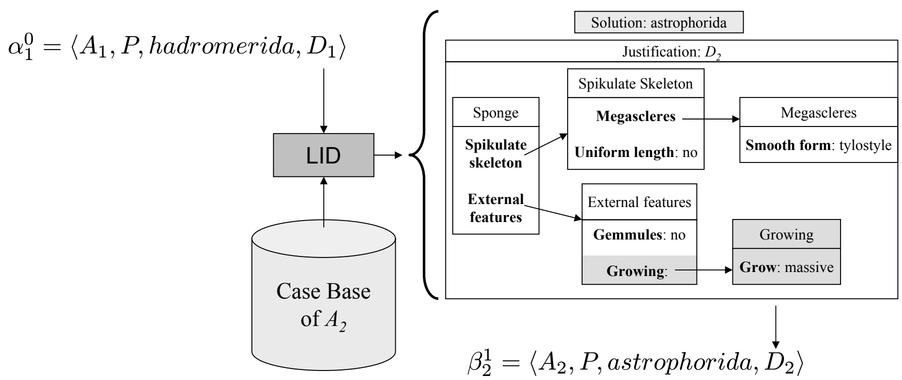

Figure 3: Counterargument generated by LID, to the argument in Figure 2.

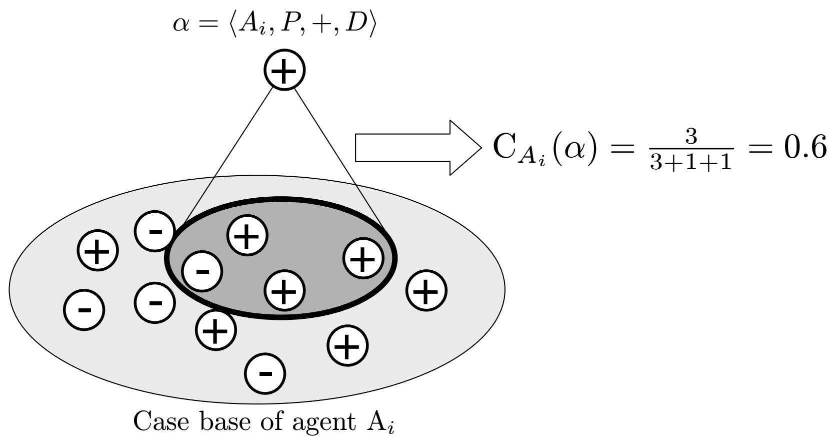

Figure 4: Confidence of arguments is evaluated by contrasting them against the case bases of the agents.

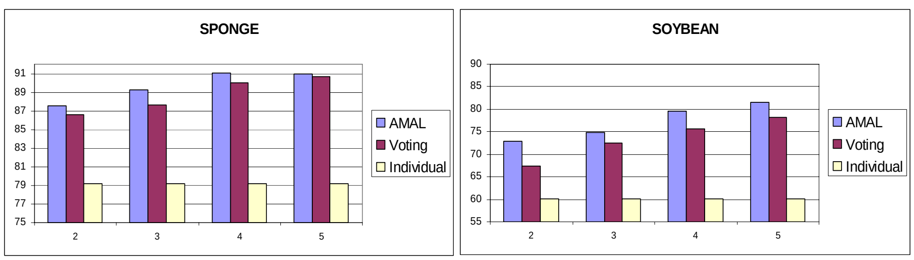

Figure 5: Individual and joint accuracy for 2 to 5 agents.

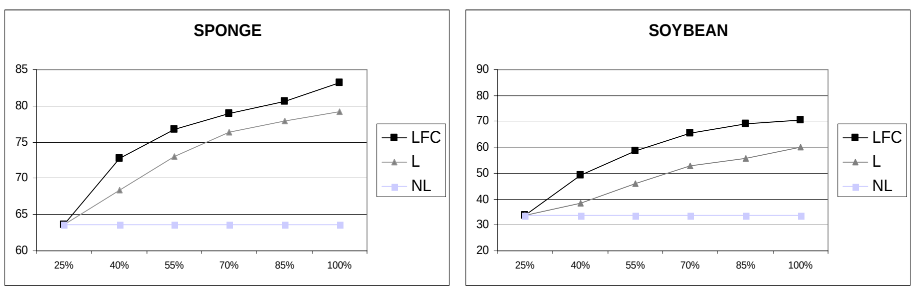

Figure 6: Learning from communication resulting from argumentation in a system composed of 5 agents.

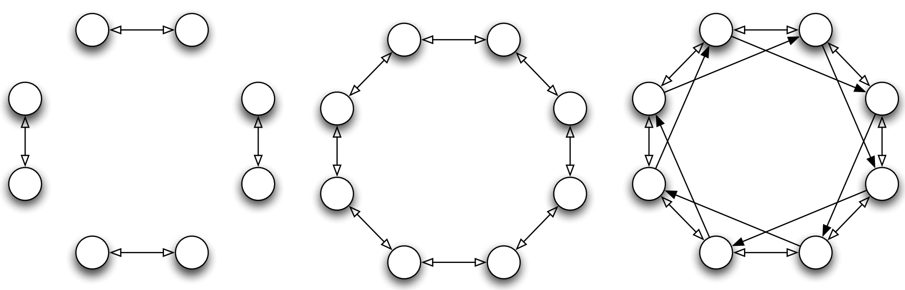

Figure 7: Three of the social networks used in our experiments, with 1, 2 and 3 acquaintances.

List of Tables

Table 1: Cases retained from communication.

SPONGE			SOYBEAN		
NL	L	LFC	NL	L	LFC
12.7	50.4	58.96	13.82	55.26	87.16

social network	*market accuracy*	*individual accuracy*	*average reward*
0 acquaintances	*89.71%*	*74.21%*	*10.35*
1 acquaintances	*90.57%*	*83.99%*	*11.42*
2 acquaintances	*91.29%*	*86.63%*	*12.14*
3 acquaintances	*91.14%*	*87.64%*	*11.94*
4 acquaintances	*91.07%*	*88.16%*	*11.85*
5 acquaintances	*91.21%*	*88.21%*	*11.93*

Table 2: Prediction markets accuracy with information exchange with varying number of acquaintances in the sponge dataset.