

Case-based Learning from Proactive Communication

Santiago Ontañón* and Enric Plaza

IIIA - Artificial Intelligence Research Institute
CSIC - Spanish Council for Scientific Research
Campus UAB, 08193 Bellaterra, Catalonia (Spain)
{santi,enric}@iiia.csic.es

Abstract

We present a proactive communication approach that allows CBR agents to gauge the strengths and weaknesses of other CBR agents. The communication protocol allows CBR agents to learn from communicating with other CBR agents in such a way that each agent is able to retain certain cases provided by other agents that are able to improve their individual performance (without need to disclose all the contents of each case base). The selection and retention of cases is modeled as a case bartering process, where each individual CBR agent autonomously decides which cases offers for bartering and which offered barterers accepts. Experimental evaluations show that the sum of all these individual decisions result in a clear improvement in individual CBR agent performance with only a moderate increase of individual case bases.

1 Introduction

Distributed case-based reasoning (CBR) aims at studying and applying CBR techniques for situations where distributed resources are present [Plaza and McGinty, 2006]. Current approaches range from one CBR agent using several case bases [Leake and Sooriamurthi, 2001] (or several agents using one case base) to multiple agents having individual case bases [Plaza and Ontañón, 2001], and applications span from classification [Plaza and Ontañón, 2001] and planning [McGinty and Smyth, 2001] to engineering applications [Watson and Garding, 1999].

Our goal in this paper is studying learning opportunities derived from communication among different CBR agents possessing individual case bases. Communication among agents offers the possibility of using the other agents as an additional source of information. Typically, a CBR system learns by retaining the cases they solve during the *Retain stage* [Aamodt and Plaza, 1994]. The presence of cases owned by other CBR agents presents the opportunity of acquiring some of those cases to improve an individual CBR agent performance.

An alternative approach might be to centralize all data. However, there are several reasons for not following this path always (i.e. centralization is feasible and efficient on some situations but not on others). First, since we assume data distributed over several sites (each with a case base C_i), centralization means basically sending all data to everyone's site. In this way, every CBR agent will have all known cases on its case base ($C = C_1 \cup \dots \cup C_n$); this scenario is simple but hardly efficient nor scalable, and may be unfeasible in situations where the different sites would agree to partially disclose some data but not all data all the time.

However, simply having more cases does not assure better CBR performance [Smyth and Keane, 1995], and we need to use case base maintenance techniques or retain policies to select the subset of cases that improves CBR performance [López de Mántaras *et al.*, 2006]. Thus, when using the centralized approach, after receiving all cases from the other agents each agent will proceed to *purge* its individual case base to obtain an individual *reduced case base* $C^R \subset C$. In fact, a CBR agent A_i would have only needed to acquire the cases in the set $C^R - C_i$ to achieve this situation. Therefore, all that is needed is for A_i to acquire from the other agents a subset of cases ($C^R - C_i$ or a similar one) such that allows a real improvement of A_i 's individual performance.

Therefore, a more intelligent strategy is to find a subset of cases to be acquired from the other agents that is equivalent (for the purposes of improving a CBR agent performance) to the $C^R - C_i$ set. Our proposal is developing a *distributed* and *decentralized* strategy that achieves exactly this effect: instead of a centralized process we propose a protocol of cooperation (that respects the agents autonomy of decision) based on communication and bartering. This approach is also feasible in situations where the different agents involved are willing to share only part of their data, a situation where the centralized approach is not feasible.

Communication may allow the agents to discover how they can help one another. However, we need to define which kind of communication protocol is needed to facilitate the agents discovering which part of their respective data should be exchanged and with whom — without recourse to simply disclosing and transferring all data to all agents. The next section (§2) presents our proposal for a communication process aimed at discovering the individual informational deficits and the possible sources to overcome them.

¹Current address: Georgia Institute of Technology, Atlanta, GA 30332/0280

Case bartering provides an equitable and practical way to reach agreements on what they effectively share once the individual information deficits and the possible sources to overcome them have been identified. Section 4 presents an interaction protocol for case bartering and an individual policy for the agents to decide on the contents of each barter exchange. An experimental evaluation is discussed in Section 5, and the paper ends with a conclusions section.

2 Learning from Communication

Centralized CBR only considers learning from the problems the CBR system solves, i.e. the Retain stage only considers the cases “owned” by the system. In a distributed CBR scenario, however, a CBR agent may learn from the cases solved by another agent. For this purpose, it is necessary to define a way to determine which cases of other CBR agents are of interest for an individual CBR agent to retain in addition to those it has already in the case base. In this paper we consider that the cases of interest are those that if present in an individual case base would improve the individual performance of that CBR agent. In this section we will present the process of communication (for information exchange) and bartering (for case exchange) proposed for improving the individual performance of CBR agents.

Previously, Ontañón and Plaza [2002] used case bartering to improve performance in multiagent case-based reasoning. However, their goal is that of improving performance for situations where the individual case bases were *biased* (when the contents of an individual case base is not a good sample of all cases it is said to be *biased*). In this approach, the communication process was simple and efficient: the only exchanged information that were *class-distribution statistic* data that indicated which agents have more (or respectively less) cases of a particular class than the average, information that is then used to barter cases of a specific class from agents that have more than average to those that have less than average. Therefore, Ontañón and Plaza [2002] did not address the issue of improving a CBR agent performance in general circumstances, since in the absence of bias no bartering would take place.

In order to enrich the communication process we will use the notion of *justified prediction*, introduced at [Ontañón and Plaza, 2003] as a useful tool for multiagent learning. Using justified predictions, an agent communicates to another agent not just the solution of a given problem (the *prediction*) but also a symbolic description of the aspects that were important in determining the solution for that problem. Figure 1 shows a graphical representation of the justification given by a CBR agent in the domain of marine sponges identification (used later in §5). This justification means: “The problem P belongs to the class *Hadromerida* because it has no *Gemmules*, the *spiculate skeleton* does not have a uniform length and the *megascleres* (in the *spiculate skeleton*) have a *tylostyle smooth form*”. This justification was obtained using the CBR technique LID [Armengol and Plaza, 2001], but eager learning methods can also produce symbolic justifications — e.g. a decision tree can produce as justification the collection of branch conditions satisfied by the problem at hand.




Figure 1: Example of symbolic justification returned by LID in the marine sponges identification task.

The main goal of a *proactive communication process* is to establish an interaction among CBR agents that allows to determine which information deficits each individual agent has; once these deficits are determined the case bartering process can proceed with the involved CBR agents possessing the necessary knowledge to offer and accept profitable case exchanges (see the case selection policy explained in §4.2). However, determining those deficits is not obvious, since an agent cannot determine its individual deficits.

Our proposal is then that an agent determines the deficits of other agents. Imagine, for instance, that an agent A_i sends three cases from its case base (c_1, c_2, c_3) and stripping them from their known solutions sends the three problems ($c_1.P, c_2.P, c_3.P$) to an agent A_j . Agent A_j solves the three problems and sends back to A_i three justified predictions ($J_{A_j}^{c_1}, J_{A_j}^{c_2}, J_{A_j}^{c_3}$). Then, A_i compares the predictions with the known solutions and determines, for instance, that A_j has solved correctly problems $c_1.P$ and $c_2.P$ but has given an incorrect solution to $c_3.P$; thus, A_i has discovered a deficit of A_j , one of which A_j is necessarily unaware of. Moreover, A_i is surely in a disposition to help A_j to improve its performance by repairing this deficit. The direct solution would be to offer case c_3 to agent A_j in exchange of some other case of A_j that would help A_i (case bartering).

In fact, since agents are exchanging justified predictions, A_i can find several cases that would be interesting for A_j to exchange with. A_i can examine its individual case base for those cases that satisfy the symbolic description sent in the justified prediction $J_{A_j}^{c_3}$ and that have the same solution as c_3 — since any of those cases are likely to help A_j in avoiding similar errors in the future. Therefore, we can design a *proactive communication process* with which CBR agents can find out which deficits the other agents have, and determine the cases they own that can be useful for other individual agents; engaging this process insures to each participating agent that the other agents will detect its own deficits and determine useful cases for him. The next sections formalize these ideas defining justified predictions (§3) and specifying a collaboration strategy for a system of CBR agents (§4).

3 Justified Predictions

In this section we will formally define the concept of justified prediction and some related notions required to clearly

present our case bartering interaction protocol. We will use the following notation for cases. A case $c = \langle P, S \rangle$ is a tuple containing a case description $P \in \mathcal{P}$ and a solution class $S \in \mathcal{S}$. We will use the dot notation to refer to elements inside a tuple. e.g., to refer to the solution class of a case c , we will write $c.S$. Specifically, we define a justified prediction as:

Definition 3.1 A Justified Prediction is a tuple $J = \langle S, D, P, A \rangle$ where agent A considers S the correct solution for problem P , and that prediction is justified a symbolic description D such that $J.D \sqsubseteq J.P$.

Thus, a justified prediction provides a way in which an agent A can explain or justify, using a symbolic description $J.D$, the reason why A predicts solution $J.S$ for problem $J.P$. The symbolic description has to subsume the problem ($J.D \sqsubseteq J.P$) because $J.D$ has to be a generalization of $J.P$ such that includes only those aspects of the problem that have determined predicting $J.S$ as the solution. For example, Figure 1 shows a symbolic description that is a generalization of (subsumes) a problem description (generated by the LID CBR technique [Armengol and Plaza, 2001] on the sponges data set used in §5).

Moreover, when an agent A_i receives an incorrect justified prediction J_{A_j} from another agent A_j , A_i can examine this justification, and determine whether it has some cases in its local case base C_i that can contradict J_{A_j} , i.e. those cases that would be useful to A_j to repair its knowledge deficit. These cases are *counterexamples* of the justified prediction J_{A_j} .

Definition 3.2 A counterexample of a justified prediction J is a case c such that $J.D \sqsubseteq c.P \wedge c.S \neq J.S$, i.e. a case c subsumed by the justification $J.D$ that has a solution different from the predicted solution $J.S$.

Using these definitions, we can now proceed to present the case bartering interaction protocol.

4 The PCCL Collaboration Strategy

We present now a collaboration strategy that supports a proactive communication process and a case bartering protocol.

Definition 4.1 The Proactive Communication Case-based Learning (PCCL) Collaboration Strategy is a tuple $\langle I_{PCCL}, D_{CS} \rangle$, where I_{PCCL} is the PCCL interaction protocol (§4.1) and D_{CS} is the Case Selection decision policy (§4.2) used to decide which cases to offer to a given agent in exchange of other cases.

The PCCL interaction protocol does not separate a proactive communication stage from a case bartering stage; instead, PCCL interleaves proactive communication and case bartering into a single process.

4.1 Case Bartering Interaction Protocol

The main idea of the case bartering interaction protocol is that each time an agent A receives a new case c to be retained, A uses c to find deficits of other agents. Specifically, A sends them the problem $c.P$ and looks for counterexamples for any incorrect justified prediction received from them. These counterexamples form the *refutation sets*. A refutation set is a record that contains a set of counterexamples found for




Figure 2: Refutation set of an agent A_1 .

an incorrect justification received from some agent. Agents keep collecting refutation sets and each time that two agents have counterexamples to send each other (i.e. that an agent A_i has some refutation sets for another agent A_j and at the same time A_j has refutation sets for A_i), they will exchange cases to improve their individual case bases. Figure 2 shows how an agent A_1 creates a refutation set R after receiving from another agent A_2 an incorrect justified prediction J at a given time t . A_1 checks in its casebase C_1 for those cases subsumed by the explanation $J.D$ (grey area), finding cases c_2, c_3 and c_4 ; the refutation set R incorporates those cases (c_2 and c_4) with correct solution S_1 .

Specifically, let $\mathcal{A} = \{A_1, \dots, A_n\}$ be a multi-agent system where each agent A_i keeps a collection $I_i = \{R_1, \dots, R_n\}$ of refutation sets. When an agent A_i receives a case c to be retained, it proceeds as follows:

1. A_i asks the rest of the agents to make a justified prediction for $c.P$.
2. Every agent A_j that has received $c.P$ generates a justified prediction $J_{A_j}^c$ and sends it back to A_i .
3. A_i now looks for counterexamples to all the received justified predictions that are wrong (i.e. those which predict a solution different than $c.S$). For each wrong justified prediction $J_{A_j}^c$ for which a non empty set of counterexamples C is found, add a refutation set $R = \langle A_j, t, C \rangle$ to I_i (where t is the current time).
4. A_i finally retains c into its case base C_i .
5. For each agent A_j for which A_i has some non empty refutation sets in its collection I_i of refutation sets:
 - (a) A_i uses its D_{CS} decision policy to construct the belying set $B_{i \rightarrow j}$ (explained in §4.2).
 - (b) A_i informs A_j that A_i offers a number $\#(B_{i \rightarrow j})$ of cases to barter with A_j .
 - (c) Then A_j uses its D_{CS} decision policy to construct the belying set $B_{j \rightarrow i}$.
 - (d) If $B_{j \rightarrow i} = \emptyset$ then, A_j has no cases to offer to A_i , and thus A_j sends a message to A_i rejecting the




Figure 3: A belying set $B_{i \rightarrow j}$ by agent A_i to agent A_j .

bartering offer. Otherwise, ($B_{j \rightarrow i} \neq \emptyset$) agent A_j sends a message to A_i accepting the bartering offer and informing that has a number $\#(B_{j \rightarrow i})$ of cases to barter with.

- (e) When a bartering offer is accepted both agents select a subset of cases from their belying sets with exactly $\min(\#(B_{j \rightarrow i}), \#(B_{i \rightarrow j}))$ cases. Then those cases are actually bartered (i.e. a copy of each bartered case is sent in a message), and the bartered cases are removed from the stored refutation sets in I_i and I_j .

The next section explains the decision policy used in the protocol at step 5.a to select the cases offered to barter with another specific CBR agent.

4.2 Case Selection Decision Policy

The *Case Selection* decision policy D_{CS} is in charge of selecting a set of cases that would be offered to a given agent A_j in exchange of other cases.

Intuitively, when an agent A_i uses its D_{CS} decision policy to select cases, the process is the following one: A_i has collected counterexamples in the refutation sets $I_i = \{R_1, \dots, R_m\}$ during the past iterations of the PCCL interaction protocol. Moreover, since the stored counterexamples are cases that may help other agents to improve their predictions (since those cases bely some of the other agents' wrong predictions), they are the cases that A_i will offer to other agents.

Specifically, given an agent A_i that wants to offer cases to another agent A_j , D_{CS} works as follows:

1. $B_{i \rightarrow j} := \emptyset$.
2. For each $R_k \in I_i$ such that $R_k.A = A_j$
 - (a) Select a case $c_k \in R_k$ such that $c_k \notin B_{i \rightarrow j}$.
 - (b) If c_k exists then $B_{i \rightarrow j} := B_{i \rightarrow j} \cup c_k$.
3. $B_{i \rightarrow j}$ is the set of cases to be offered to agent A_j .

We call $B_{i \rightarrow j}$ the *belying set* because this is the set of cases that contradict the justifications received from A_j that were found incorrect. Figure 3 illustrates this process, and shows how an agent A_i generates the belying set $B_{i \rightarrow j}$ for another agent A_j by selecting one case of each of the refutation sets it has collected for A_j .




Figure 4: Individual average accuracy with only individual learning (IL) and with learning from communication (LFC).

5 Experimental Evaluation

In this section we empirically evaluate our learning from communication framework. We have made experiments in three different data sets: *sponges*, *soybean*, and *zoology*. The *Sponges* data set is a marine sponges identification task, contains 280 marine sponges (of class *Demospongiae*) represented as a relational cases and they have to be identified as pertaining to one of three different orders (*Astrophorida*, *Hadromerida* or *Axinellida*). *Soybean* is a standard data set from the UCI machine learning repository with has 307 examples pertaining to 19 solution classes, why the *Zoology* data set is also from UCI and has 101 examples pertaining to 7 solution classes.

In order to evaluate PCCL, we have used a 5 agent system. In an experimental run the data set is divided into training set (90% of the cases) and test set (10%), a 25% of the training cases are initially distributed among the 5 agents without replication, i.e. there is no case shared by two agents. The remaining training cases are sent to the agents one by one. When an agent receives a training case, the agent uses PCCL to find opportunities of learning from communication with other agents. Moreover, periodically, test cases are sent to the agents to evaluate their improvement in classification accuracy as they receive more training cases.

Figure 4 shows the learning curves for a 5 agents system in the sponges, soybean and zoology dataset. The vertical axis shows the individual classification accuracy of the 5 agents on average, and the accuracy values are evaluated (over the horizontal axis) with different percentage of training cases




Figure 5: Average number of cases retained with only individual learning (IL) and with learning from communication (LFC).

learnt by an individual agent (from 25% to 100%). The IL plot shows the average individual accuracy for learning only from the training set, while the LFC plot shows the average individual accuracy for learning from the training set and also from communication (using the PCCL collaboration strategy). We can see that the additional case learning obtained from case bartering clearly improves individual accuracy in the three data sets. Moreover, this improvement is achieved with a relatively modest number of additional cases (compared to what a “large” number of cases would be sharing in a scenario where all cases are shared). Finally, the difference in accuracy is statistically significant for the three data sets according to a t-test with $p = 0.05$.

Figure 5 shows the evolution of the average individual case base size of the CBR agents both for IL and LFC: the horizontal axis represents percentage of training cases learnt and the vertical axis shows the number of cases retained in an individual case base on average. The three plots of Figure 4 also show that the proactive communication process increases the speed at which an individual CBR agent achieves higher accuracy values. Moreover, the accuracy with case bartering keeps getting higher as the training set increases, showing that the cases learnt from communication are useful for the individual CBR agent. Thus it is clear that with more cases (acquired via proactive communication) the average individual accuracy improves.

Let us now compare the accuracy achieved using case bartering with that achieved by the (random) cases in the training set. For instance, in the zoology data set accuracy is 82.18 with 17.74 retained cases on average (at 70%) using case bartering, while without bartering accuracy is only 80.73

with 18.18 retained cases (at 100%); that is to say, with a smaller number of cases the accuracy is higher — and this effect is due to the selection of cases realized in PCCL. Another example, in the sponges data set, is that accuracy with case bartering is 82.07 with 50.4 retained cases (at 55%) while accuracy without case bartering is 79.21 with 50.4 retained cases (at 100%); i.e. the same effect is present. In the soybean data set this effect is not as apparent: the reason is that accuracy rapidly degrades when a CBR agent has few cases, and in this situation any new case improves accuracy. Summing up, the results in the three experiments show that the cases acquired via case bartering are useful not just as being new cases, but they are useful specifically for the individual for which they have been selected in the PCCL collaboration strategy. That is to say, the case bartering protocol (§4.1) together with the Case Selection decision policy (§4.2) achieve the selection and effective retention of a *small and individualized* set of cases that clearly improve the case base of each CBR agent.

The effect of achieving more *compact* case bases (for a given level of accuracy) can only be explained as a result of three factors. Firstly, a “good selection” (in the sense of being *individualized*) of cases is obtained by the D_{CS} decision policy; otherwise more compact case bases would not achieve higher accuracy levels. Secondly, a “small enough” set of cases is selected by the D_{CS} decision policy; otherwise the size of case bases could grow by means of bartering a much larger amount of cases. Thirdly, the bartering protocol is effective in transforming individual goals (improving one’s accuracy) into a mutually beneficial strategy.

Finally, let us focus on step 5.e of the interaction protocol, where the number of cases effectively bartered is determined. Step 5.e takes the minimum cardinality of the two belying sets $B_{j \rightarrow i}$ and $B_{i \rightarrow j}$ assuring the barter exchanges cases on a 1 for 1 basis. At first sight, this may seem too restrictive, since if agent A_i offers 5 cases to A_j who offers back 1 case then only 1 case is bartered and the other 4 (that might be useful for A_j) are “lost.” In order to evaluate this assumption, we performed the same experiments changing the step 5.e of the interaction protocol in a way that as long as the two agent have some cases to barter they are exchanged regardless of their numbers —i.e. we allow $n : m$ barterers. In principle, this is not unreasonable: traditional barter does not formally enforce a 1 : 1 exchange of goods since the utility of goods for each barterer may be quite different. This may also be true for case bartering, e.g. for a 1 : 3 exchange the single case may be very useful for the recipient.

The experiments showed that this looser $n : m$ exchange policy is not better than the strict 1 : 1 exchange policy. Specifically, the experiments showed in the three data sets that $n : m$ bartering achieves roughly the same accuracy as 1 : 1 bartering but the average individual case base size appreciably increases. Thus, each CBR agent retains more cases from the proactive communication process but they are redundant, since they are not increasing the average individual accuracy. This situation is an instance of the principle mentioned in the Introduction section: simply retaining all cases is not always the best retain policy. If we examine more deeply the reason for this effect, we notice that the selection

policy D_{CS} only selects one case from each refutation set. The reason is that, in principle, one counterexample is enough to prevent an agent A_j to make the same error again. Sending more than one counterexample to an agent A_j for the same justified prediction will not add extra value to A_j , and will (unnecessarily) increase its case base. The experimental results confirm that one counterexample is enough since the accuracy improvement is roughly the same. We can then conclude that the PCCL collaboration strategy (case bartering protocol plus case selection policy) exchange an adequate number of cases and not more than those necessary.

6 Conclusion

The content of a case base in a CBR system is traditionally determined by applying retention policies and/or case base management techniques [López de Mántaras *et al.*, 2006]. However, these traditional approaches assume that there is a single source of case acquisition (the problems solved by the CBR system), while we have presented a situation where communicating with other CBR agents is a second source for case acquisition. We have presented a proactive communication process, by which a CBR agent sends problems to other agents and acquires their justified answers; this process is proactive because it is engaged to acquire the information that will later be useful in case bartering.

During the proactive communication process, CBR agents not only solve the received problems, they use *justified predictions* to express the aspects they took into account in making that prediction. The symbolic description that is contained in a justified prediction is what allows an agent to determine not only when another agent fails, but the *particular deficit* that caused that failure. Comparing the symbolic justification with its own cases, a CBR agent can discover which cases (*counterexamples*) can be useful for another agent to avoid a similar failure in the future. Notice that it is the existence of *justified predictions* (introduced at [Ontañón and Plaza, 2003]) that allows such a fine-grained analysis and thus supports later a more focused bartering of cases process resulting in an overall improvement with only a relatively small number of cases exchanged.

Previous work on case bartering [Ontañón and Plaza, 2002] used a much less informed decision policy (ICB), and would not improve the performance of CBR agents in our experimental setting. ICB is useful when CBR agents have biased case bases —i.e. less (resp. more) than average number of cases of a certain class. Biased CBR agents have lower accuracy than unbiased ones and ICB guides a process of bartering cases that allows the CBR agents to eliminate all or most of their bias (achieving the “nominal” accuracy). However, PCCL has been experimented with in this paper with already unbiased CBR agents, and the improvement is therefore on top of the “nominal” accuracy. The use of justified predictions is precisely the tool that allows to detect the cases that need to be bartered.

Acknowledgments This research has been partially supported by the MID-CBR (TIN 2006-15140-C03-01) and ProMusic (IC2003-07776-C02-02) projects.

References

- [Aamodt and Plaza, 1994] Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *Artificial Intelligence Communications*, 7(1):39–59, 1994.
- [Armengol and Plaza, 2001] E. Armengol and E. Plaza. Lazy induction of descriptions for relational case-based learning. In *12th European Conference on Machine Learning*, pages 13–24, 2001.
- [Leake and Sooriamurthi, 2001] D. Leake and R. Sooriamurthi. When two case bases are better than one: Exploiting multiple case bases. In D.W. Aha and I. Watson, editors, *Proceedings of the 4th International Conference on Case-Based Reasoning*, pages 321–335. Springer, 2001.
- [López de Mántaras *et al.*, 2006] R. López de Mántaras, D. McSherry, D. Bridge, D. Leake, B. Smyth, S. Craw, B. Faltings, M.L. Maher, M. Cox, K. Forbus, M. Keane, A. Aamodt, and I. Watson. Retrieval, reuse, revise, and retention in CBR. *Knowledge Engineering Review*, 20:215–240, 2006.
- [McGinty and Smyth, 2001] L. McGinty and B. Smyth. Collaborative case-based reasoning: Applications in personalized route planning. In D.W. Aha and I. Watson, editors, *Proceedings of the 4th International Conference on Case-Based Reasoning*, pages 362–376. Springer, 2001.
- [Ontañón and Plaza, 2002] Santi Ontañón and Enric Plaza. A bartering approach to improve multiagent learning. In *1st International Joint Conference in Autonomous Agents and Multiagent Systems*, 2002.
- [Ontañón and Plaza, 2003] Santiago Ontañón and Enric Plaza. Justification-based multiagent learning. In *Proc. 20th ICML*, pages 576–583. Morgan Kaufmann, 2003.
- [Plaza and McGinty, 2006] Enric Plaza and Lorraine McGinty. Distributed case-based reasoning. *Knowledge Engineering Review*, 20:261–265, 2006.
- [Plaza and Ontañón, 2001] Enric Plaza and Santiago Ontañón. Ensemble case-based reasoning: Collaboration policies for multiagent cooperative cbr. In I. Watson and Q. Yang, editors, *In Case-Based Reasoning Research and Development: ICCBR-2001*, number 2080 in LNAI, pages 437–451. Springer-Verlag, 2001.
- [Smyth and Keane, 1995] Barry Smyth and Mark T. Keane. Remembering to forget: A competence-preserving case deletion policy for case-based reasoning systems. In *Proceedings of IJCAI-95*, pages 377–382, 1995.
- [Watson and Gardingen, 1999] I. Watson and D. Gardingen. A distributed case-based reasoning application for engineering sales support. In T. Dean, editor, *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, volume 1, pages 600–605. Morgan Kaufmann Publishers Inc., 1999.