

Using Case-Based Reasoning and Principled Negotiation to provide decision support for dispute resolution

Davide Carneiro · Paulo Novais · Francisco Andrade ·
John Zeleznikow · José Neves

Received: 24 August 2010 / Revised: 7 July 2011 / Accepted: 4 August 2012 /
Published online: 18 October 2012
© Springer-Verlag London 2012

Abstract The growing use of Information Technology in the commercial arena leads to an urgent need to find alternatives to traditional dispute resolution. New tools from fields such as artificial intelligence (AI) should be considered in the process of developing novel online dispute resolution (ODR) platforms, in order to make the litigation process simpler, faster and conform with the new virtual environments. In this work, we describe UMCourt, a project built around two sub-fields of AI research: Multi-agent Systems and Case-Based Reasoning, aimed at fostering the development of tools for ODR. This is then used to accomplish several objectives, from suggesting solutions to new disputes based on the observation of past similar disputes, to the improvement of the negotiation and mediation processes that may follow. The main objective of this work is to develop autonomous tools that can increase the effectiveness of the dispute resolution processes, namely by increasing the amount of meaningful information that is available for the parties.

Keywords Online dispute resolution · Expert system · Case-Based Reasoning · Information retrieval

1 Introduction

In the words of Abraham Lincoln: “Discourage litigation, persuade your neighbor to compromise where you can. Point out to them how the nominal winner is often the loser...in

D. Carneiro · P. Novais (✉) · J. Neves
CCTC/Department of Informatics, University of Minho, 4710-057 Braga, Portugal
e-mail: pjon@di.uminho.pt

D. Carneiro
e-mail: dcarneiro@di.uminho.pt

F. Andrade
Law School, University of Minho, Braga, Portugal

J. Zeleznikow
School of Management and Information Systems, Victoria University, Melbourne, Australia

expenses and waste of time.” This is one of the sentences that better describe the need for an alternative to traditional courts and the appearing of alternative dispute resolution (ADR) [10]. ADR includes mechanisms that aim to solve disputes without recurring to the traditional judicial process, that is, courts. ADR methods are nowadays broadly used by both the legal system and the parties involved as the first step in the path toward the dispute resolution. There are even countries in which parties are encouraged or required to try an alternative way of dispute resolution before advancing into a court.

One of the factors that definitively boosted the acceptance of ADR is the current workload of courts. Often, it takes many years for a case to receive a final hearing. Thus, a new approach is required to diminish the number of cases that actually need to be resolved by traditional court processes. Parties, in particular, also look at ADR tools in search for lower expenses, more privacy and less exposure than provided by traditional trials. Another advantage is that the entity that will decide a given case can be agreed on by the parties instead of being mandatorily assigned, which in a certain way increases the confidence and overall satisfaction of the disputing parties in the result obtained.

An alternative way of solving disputes arising out of the contractual performance in virtual environments is that of ODR. This new paradigm allows the moving of the traditional ADR methods from a physical to virtual place [8]. Hence, parties not only reduce litigation but gain a simple and efficient way of dealing with disputes, saving both temporal and monetary costs [27,42]. Several methods of ODR may be considered, from negotiation and mediation to variations in arbitration or modified jury proceedings [22].

Indeed, the growing use of information and communication technologies (ICT) in the commercial arena leads to an urgent need to find alternatives to dispute resolution. The traditional paper-based courts and arbitration, designed for the industrial era, are now outdated. In fact, Civil and Commercial Law are still turned toward a type of society and a way of doing business that is related to a mostly industrial society and the traditional use of the paper as the primary support tool. Given the rising use of Information Technology, electronic means for resolving disputes need to be considered. This should include emerging tools from fields such as artificial intelligence (AI) that are used for effective idea and strategy generation. This will result in a more creative process, rich in ideas and solutions, that will ultimately lead to simpler, faster and more effective dispute resolution processes.

In this sense, the use of software agents that embody intelligent techniques can become particularly interesting considering their ability to recognize and evaluate facts, positions and relevant information. Moreover, the use of these agents, as building blocks of distributed platforms that support ODR processes, has to be considered. Of major importance in this context are systems that are able to simulate and calculate outcomes of disputes, establish negotiation paths to achieve the desired objectives and warn the parties about the possible alternatives to an agreement.

There is a variety of ODR systems, including legal knowledge-based systems and systems that help settle disputes in an online environment [14]. Second-generation ODR, in which ODR systems act as an autonomous agent [33] are a most appealing way of solving disputes. Also interesting here is the view of Katsh and Rifkin concerning the role of technological tools. The authors see technology that works with the mediator or arbitrator as the fourth party in the dispute, together with the two disputing parties and the third neutral [26]. In fact, there has been a recent tendency to foster the intervention of software agents, acting either as decision support systems [8] or as real electronic mediators [33]. Surely, this latest role for software agents could benefit from the use of AI techniques such as Case-Based Reasoning (CBR) and information and knowledge representation models. Models of the description of the fact situations of the factors relevant for their legal effects allow the agents

to be supplied with both the static knowledge of the facts and the dynamic sequence of events [33].

Evidently, representing facts and events is not sufficient for a dispute resolution. In order for the software agent to perform actions of utility for the resolution of the dispute, it also needs to know the terms of the dispute and the rights or wrongs of the parties [33]. Moreover, it is of utter importance that the software agent is able to foresee the legal consequences of the said facts and events. We thus need to consider whether agents can evaluate the position of the parties and present them with useful proposals, taking into a consideration which of the two parties would have a higher probability of being penalized or supported by a judicial decision of the dispute and, therefore, who would be more or less willing to make concessions in their claims [49]. The ability to understand the position of the parties is vital for the successful involvement of software agents in the process. To do so, it is mandatory for the software agent to have the characteristics of consistency, transparency, efficiency and enhanced support for dispute resolution, in order to allow it to replicate the manner in which decisions are made and thus make the parties aware of the likely outcome of litigation [36]. That is to say, software agent intervention in ODR procedures should take into account the alternatives, for the parties, to an ODR negotiated agreement.

Thus, in such a context, it would be interesting to consider some useful and well-known ADR concepts such as the BATNA, or Best Alternative to a Negotiated Agreement. This concept denotes the best scenario possible if the negotiation process fails, that is, if the case was to go into a court, what would be the best outcome? Knowing their BATNA helps parties to achieve an outcome using negotiation that is hopefully better than it would be if achieved through litigation [17]. The importance of this concept is well expressed by Goldberg et al. in [21]. In fact, if a party is unaware of what results could be obtained if the negotiations are unsuccessful, he runs the risk of: entering into an agreement that he would be better off rejecting; or rejecting an agreement he would be better off entering into.

Following the same line of thought, the WATNA, which denotes the worst possible outcome along a litigation path, should also be considered [32]. It can also be quite relevant in complementing Principled Negotiation [44] with a justice- or rights-based approach, thus leading to a calculation of the real risks that parties will face in judicially determined litigation, imagining the worst possible outcome. In the same sense, considering the MLATNA—Most Likely Alternative to a Negotiated Agreement—may also be useful as a way to show parties the alternative that is most likely to happen if the negotiated process fails. These three concepts are definitively of utter importance for the parties to have a clear general picture of what the alternatives are, thus supporting wiser decisions. Finally, one last important concept is the one of ZOPA—Zone of Potential Agreement—which denotes the range of the possible outcomes, that is, what may happen at the end of the process, and is delimited by the WATNA and BATNA [29].

Considering the usefulness of these concepts in ODR, it becomes obvious that these systems will go much further than just transposing ADR ideas into ODR environments. They should actually be guided by judicial reasoning, getting disputants to arrive at outcomes in line with those a judge would reach [50]. Despite there being difficulties to overcome at this level, the generalized use of software agents as decision support systems points out the usefulness of following this path.

Taking all of this into consideration, the work described in this paper focuses on two key points. First, the use of AI techniques that not only enable contextualized information retrieval but also allow the system to evolve according to the results of this retrieval [30]. Particularly, we will be analyzing the role that CBR can play not only in the information retrieval task but also on improving other processes. Second, we will be looking at the use of traditional

techniques, specifically negotiation and mediation, as a way to improve the efficiency of dispute resolution. These two trends merge to fulfill a unique objective: to provide contextualized information and guidance processes to disputant parties so that they can solve their dispute out of courts.

1.1 Related work

Proving the validity and relevancy of this research field, a group of interesting projects with the objective of giving birth to dispute resolution mechanisms can be pointed out. In this section, we describe some of the projects which intersect with this work.

Rule-based legal decision-making system (LDS) dates from 1980 and was one of the first negotiation support systems to be developed [46]. The domain of this system was liability law. This field of law holds responsible product distributors and manufacturers for the injuries their products may cause. The system created embodied the skills and knowledge of a human expert in the shape of antecedent-consequent rules. The project had the objective of formalizing the decision-making processes of attorneys and claims adjusters involved in product liability litigation in the shape of rule-based models so that the effects that changes in legal doctrine have in settlement strategies and practices could be studied. The authors formalized the strict-liability concept on ROSIE language, so that the defendant could or could not be considered liable.

A more recent project focusing on providing support for decisions is EXPERTIUS. This is a decision support system that advices Mexican novice judges and clerks upon the determination of whether the plaintiff is or is not eligible for granting him/her a financial pension (on the basis of the “feeding obligation”) and if so upon the determination of the amount of that pension [11]. The system has three main modules: a tutorial module, an inferential module and a financial module. The tutorial module guides the user through the accomplishment of several tasks. The inferential module evaluates evidence based on weights that the user assigns to each piece of evidence. It determines which prepositions are defeated and which prevail. At last, the financial module assists the user on the calculus of the value of pensions according to some criteria.

For reasoning in these terms, EXPERTIUS has an extensive way of representing the knowledge about the several parameters. Judicial expert knowledge was represented as having three interrelated levels: one for representing the expert knowledge, one for representing the decisions internal to each procedural stage as regulated by procedural law and a third one that corresponds to the dialogical confrontation pattern of the case that arises simultaneously to the decisions taken in the intermediate level.

The knowledge embodied in the system was divided into five layers. The formal doctrine contained rules from the legislation and common law. The informal principles contained rules that are not explicitly expressed in the law but are generally agreed upon by legal practitioners. Under the strategies layer, the authors coded the methods used by legal practitioners to accomplish a given goal. The subjective considerations layer contains rules that anticipate the subjective responses of people involved in legal interactions. At last, the secondary effects layer contained rules that describe the interactions between rules. The authors concluded that despite the number of rules needed for formalizing the law and the strategies, the rule-based model was feasible and suited for this particular domain.

In the field of negotiation support, a major work is the Family_Winner project, which integrates game theory and heuristics in the field of family law, more specifically in disputes arising from divorce processes [53]. The research is largely based on the Adjusted Winner algorithm [53], an algorithm that merges insights from the fields of game and decision theory.

This algorithm aims at the distribution of a group of items or issues by the parties, according to a mechanism of point allocation. In this process, each party quantifies how much he/she wants the item and whoever values more a given item receives it.

Still in the negotiation support, another project worth mentioning is SmartSettle [41,43]. It is an online negotiation system that can be described as a generic tool for decision-makers. It is intended for parties with conflicting objectives that wish to reach a formal agreement. For achieving it, tools that help to define the interests and possible trade-offs are used, with an emphasis on recognizing the parties' satisfaction in the search for optimal global solutions. This platform can be used to solve problems relating to family, insurance, real estate, labor management, contract negotiations, among others.

When comparing the work presented in this paper with the related work mentioned in this section, the main difference that can be noted is that, contrary to the most common approach, we do not aim to provide support for dispute resolution in court. We rather aim for providing valuable information for the parties and eventually the neutrals, so that better and more informed decisions can be taken. While doing so, we also expect disputing parties from going into court, by making them aware of the possible consequences, fostering cooperative off-court processes. Nevertheless, the tool presented can still be used throughout a complete dispute resolution process, ranging from the initial providing and compilation of information to the definition of an outcome. Another major difference is that in this work, there is a close integration between the phase of generating information and the phase of actual dispute resolution through traditional means. That is, the information that is compiled to inform the disputing parties is also used to support and help define the mediation or negotiation process that may follow.

2 Agents and negotiation

As mentioned before, there is a particular interest in implementing ODR tools that behave the same way a human would. Particularly interesting is the use of software agents to implement social interaction processes, specifically negotiation. Generally, negotiation is classified as distributive or integrative [45].

In distributive negotiation, one looks at the problem as something that can be divided and distributed by the parties in an attempt to maximize their satisfaction. An example scenario is the classical winding up of a company in which the assets of the company are sold and the proceedings collected are used to discharge the liabilities. In game theory, this is known as a zero-sum game. Two important concepts here are the ones of utility and resistance [7]. Utility denotes the value that a given item has to a party while resistance denotes the willingness of a party to change the utility of an item. A good negotiator usually tries to convince the other party that certain items do not have the value that they are given. The negotiator will succeed if the opponent has a low resistance in that item, and if he does so, it will be easier for the negotiator to win that item or he will, at least, be in a better position for the rest of the negotiation process [25]. Accordingly, one can define utility functions that help to understand how each party values the items being distributed, therefore predicting possible outcomes and the evolution of the negotiation process [54].

In integrative negotiation, the problem is expected to have more solutions than the ones visible at first sight. In these types of problems, the parties try to bring to the table as much interests as possible so that there are more and more valuable items with which to negotiate. When the parties are increasing the value of what they put in the table, they take into account their interests which include the needs, fears, concerns and desires. This type of negotiation

is also known as interest-based, as parties try to combine their interests and find common points in which they are satisfied. By doing so, more satisfactory outcomes are achieved by both parties. This makes integrative negotiation more desirable than distributive.

A good example of illustrating this difference is an old story about a brother and a sister both wanting the last orange. To solve the dispute, their mother initially gives both children half the orange. When she asks them their goals, she finds that the brother wants the orange for the juice to drink while the sister wants the rind to flavor an orange cake. So the mother gives the sister the rind and the brother the rest and they both get what they desire [18]. The outcome is optimum. An important concept in this field is the one of Pareto efficiency [20]. In this case, the solution obtained with the integrative approach would have been a Pareto-efficient solution as no better solution could have been achieved.

A key factor in dispute resolution is thus, as seen in the previous example, to identify the interests of each party so that their positions can be better understood. Such a process can eventually lead to what is known in game theory as a win–win game, that is, all the parties are better at the end of the negotiation process than when it started.

2.1 Principled Negotiation

One of the most effective methods for resolving conflicts is through negotiation. Specifically, Principled Negotiation was developed by the Harvard Negotiation Project [44] and focuses on the notion that parties look for mutual gains. When interests conflict, Principled Negotiation advocates parties arrive at a ruling that is independent of the beliefs of either side. Principled Negotiation puts forward five key points:

1. Separate the people from the problem. Personal matters or stronger personalities can often influence the outcome of dispute resolution processes. This is especially significant in scenarios in which the disputant parties have or had a personal relationship, such as in family law disputes. Solving disputes by leaving aside personal affairs can avoid biasing a solution toward a more influencing party.
2. Focus on interests, not on positions. It is important for a party to understand what their interests are rather than blindly defending a position that may not be that solid. By isolating the reasons why a position is most appealing, participants in a negotiation will increase the chance of achieving agreement.
3. Invent options for mutual gain. Although parties have divergent objectives or interests, they might still find positions that encompass mutual gain. Once interests have been ranked to determine the relative importance of each, a range of options are discussed before deciding on an outcome. Next, the negotiators need to invent options for mutual gain.
4. Insist on objective criteria. There are negotiation scenarios that cannot be treated as a win–win situation. In such cases, unbiased independent evaluations should provide an outcome or solution that both parties will agree on.
5. Know your Best Alternative to a Negotiated Agreement. When two parties negotiate, they aim at achieving a better result than would otherwise occur. Being unaware of what results one could obtain if the negotiations are unsuccessful, one runs the risk of:
 - (a) Entering into an agreement that would be better off rejecting; or
 - (b) Rejecting an agreement that would be better off entering into.

All the work presented in this paper is guided by these five principles. They were thus present throughout the development process of UMCourt, an ODR platform, which will be described in the following section.

3 UMCourt

UMCourt is being developed at the University of Minho, in the context of the TIARAC project (*Telematics and Artificial Intelligence in Alternative Conflict Resolution*). This project aims to explore the potential of the combined use of new technologies with AI techniques to Law in the domain of ADR [3]. UMCourt consists of a Multi-Agent System, based on open standards and technologies. It aims at more than a simple legal information management tool, being its core a group of agents that embody intelligent techniques for problem solving. The resulting platform is able to provide services such as determining possible outcomes and respective likelihood, proactively providing relevant information in real time, automatic creation of legal documents, determination of the major concepts in ADR for each dispute, among others. Law experts work in close cooperation with the development team of UMCourt so that the resulting platform addresses the above challenges.

These challenges are often faced in the legal domain as well, namely the definition of standards for legal information representation or the formalization of argumentation. Ashley, in [4], even argues that “*Perhaps, the most interesting opportunity for the judicial and AI & Law communities (...) is in cooperating in the design of standards for the presentation of factual descriptions and discussions of law in case opinions.*” There is also the need to provide this information to the parties in an intuitive fashion as there is, not infrequently, a pronounced difficulty in understanding legal information. This evidently constitutes a disadvantage that may be aggravated if the party does not know how to efficiently work with the ODR tool. There is thus an obvious interest in this research field from both legal and informatics practitioners.

3.1 Portuguese labor law and alternative dispute resolution

The work described here was developed in the context of the Portuguese labor law, focusing especially on the relation between employers and employees [24]. Particular attention is paid to the scenario of an employee being fired, a case in which litigation will certainly occur.

In firing an employee, a variety of considerations need to be examined. First of all, litigation is a quite slow process, especially in Portugal, and this results in several disadvantages for both parties [15]. In a dynamic economy, it is important, both for employers and employees, to have supporting mechanisms to settle the disputes in a fast and efficient way. ADR has long been considered and practiced, mainly negotiation by Lawyers. Indeed, the Portuguese Labour Code expressly recognizes the possibility of the parties submitting their disputes to arbitration (art. 564^o Labour Code) [31]. But it must be recognized that the normal way of solving this kind of disputes is through conciliation by the Judge, just before trial, and after a long procedure.

Under legal systems such as the Portuguese, in the case of an employee being fired, a huge deal of legal parameters have to be considered: the possibility of a “just cause for dismissal” being declared by the court, the existence (or not) of a valid and legal procedure of dismissal, the validity (or not) of the dismissal, the antiquity of the worker in the company, supplementary work, night work, justified or unjustified absence from work, the possibility of dismissal being accepted without indemnities or of it being accepted but accompanied by indemnities that could range from a very low to a very high amount of money [15,31]. Either the company performing the sacking can pay a small amount of money, thus making it worthwhile to fire the employee, or it may be placed in the position of being forced to pay a considerable sum. For the worker, the amounts involved are significant. Being fired without good indemnities may lead to the loss of a job and no payment as compensation. But it could, on the other side, become beneficial for him to be fired, provided that he receives a

significant cash offering. For the parties in a labor conflict, the calculation of the possibilities of an outcome of litigation is vital. Thus, it is vital for both parties to know in advance their best and worst alternatives to a negotiated agreement.

3.2 The architecture of UMCourt

The conflict resolution platform is being developed according to the traditional client-server model. The human user is able to interact with the platform by means of a standard web browser. This constitutes the client. Additionally, mobile versions of the interface have been developed, which allow users to access a more reduced set of functionalities, while on the move. The server relies on the Multi-Agent System paradigm [50]. Under this paradigm, a group of intelligent agents share information and (possibly) cooperate to accomplish a common objective. Wooldridge and Jennings define a weak notion of agent in [48] as an agent that shows autonomy, social ability, reactivity and proactiveness. A stronger notion of agent may involve features like mobility, veracity, benevolence or rationality. Such characteristics can thus significantly enrich platforms developed under this paradigm.

However, more than just choosing a technology or a paradigm, one needs a methodology. Wooldridge et al. propose a methodology in [49] for the development of agent-based applications that begins with a high-level definition of the system in terms of their roles. Following this methodology, we have defined four main groups of agents according to their roles: main, secondary, interface and control. Agents in group main have important decisory roles, typically empowered with autonomy and enhanced communication skills. Agents in group secondary have as main role to support the life cycle of the remaining agents by providing low-level but vital services. Agents that belong to the interface group are used to act as interface between a user in the web page and the Multi-Agent System. Finally, agents in group control are responsible for the correct life cycle of the whole platform. Given the scope of the paper, a thorough description of the agents is not provided. We will rather focus on the agents directly involved in the work described in the following sections. This is detailed in Table 1. For a detailed description of the architecture, please see Carneiro et al. [13].

Regarding the agent platform, an analysis of several frameworks has been conducted. Given their characteristics as well as our past experience, the choice has been in favor of the Java Agent Development Framework (JADE) platform [6]. JADE is a software framework that significantly facilitates the development of agent-based applications in compliance with the FIPA specifications [16].

Before advancing into further details, let us now make a first high-level description of the main functionalities implemented by the described software agents. In these functionalities there is a close integration between the information retrieval and its use on the posterior processes (e.g., utility evaluation, negotiation, mediation), independently of the purpose. Some of this functionalities are simple ones (i.e., it takes only one agent and one iteration to implement it) while others are complex (i.e., it may take the interaction of more than one agent or more than one iteration of the same agent to implement it). A detailed description of these functionalities is given:

- *Parsing and Indexing*: In an ODR platform, information from a dispute that took place in the real world must be acquired. This generally includes, in a first instance, the parsing of documents to determine their syntactic validity. Afterward, cases may be indexed so that they can be more efficiently retrieved. This platform is able to automatically parse cases by means of a SAX parser (Simple API for XML), determining their validity against the defined Schemas. Moreover, cases are also indexed in the database, enabling several tasks to be performed quicker.

Table 1 The main agents that build up the architecture, given the focus of this paper, with a description of their roles

Agent name	Role
Coordinator	This agent coordinates the remaining agents in order to implement higher-level tasks through the concatenation of simpler ones. In order to do so, the agent contains knowledge about the high-level tasks defined in terms of finite state automata. A typical scenario is the coordinator receiving external requests for given tasks, which it will answer by combining tasks from other agents
Database	The database agent implements a transparent bridge to the database. Thus, when the remaining agents need to interact with the database, they do it by means of requests to this agent and not directly. This has a main advantage of isolating the database from the platform
Evaluator	This agent collects information about key operations of the platform so that later its performance can be assessed. Based on the information collected and on a description of the operations, the agent is also able to issue recommendations for the improvement in the performance
NearestNeighbor	This agent implements a nearest neighbor algorithm with the objective of defining a measure of similarity between two cases
TemplateRetrieval	This agent implements several information retrieval algorithms, from which the remaining agents can choose, depending on their objectives
Utility	The utility agent is able to determine the utility of a given case for a given party, that is, to measure how good a given case is to a party. Typically, this is used to determine the utility of past known cases to a party which is now dealing with a similar one
Blackboard	This agent implements a mediation and a negotiation algorithm as well as a shared space for message exchange. It is responsible for controlling the evolution of the processes
Party	The Party agent is the virtual representation of a disputing party in the platform. Therefore, different instances of this agent may be present at one time, each one representing a different party. This agent also establishes the bridge between user and platform, that is, the user requests actions or sends messages through this agent

- *Information Retrieval*: Information retrieval is a major concern in any knowledge-based domain. The legal one is evidently not an exception. It is especially important as it is the basis for a wide range of operations and decisions. In that sense, the UMCourt platform considers several information retrieval methods that allow for important information to be retrieved in a context-aware fashion.
- *Case Similarity Assessment*: From the point of view of a disputing party, the assessment of the similarity of his case with a past known case is very important. This will allow him to determine, for example, the likeliness of an outcome, based on an observation of the outcomes of past similar cases. Likewise, legal practitioners can also use such functionalities for training or for supporting their decisions. Thus, in UMCourt, we have implemented several methods for similarity assessment.
- *Utility Computation*: Of major importance is also the ability to compute the utility of a given case. The issue here is to determine how good a given outcome is for a party, that is, if a given outcome was to occur, how much would a party win or lose. This is essential for a disputing party to take good and rational decisions.
- *Generating Important Information*: As mentioned above, there is a wide range of information that can be important for the parties at the time of taking decisions. The ability to compile all this information is thus of value and was considered in the development of UMCourt.

The use of these more basic functionalities enables or fosters the implementation of higher-level ones, which significantly improve the experience of end-users, whether they are the disputing parties or the legal practitioners.

- *Case-Based Reasoning*: The use of CBR techniques can be of importance for both the provision of information to the users and for the evolution of the platform itself. In fact, platforms based on the CBR paradigm can evolve in a dynamic and autonomous way, learning with the own correct and incorrect decisions. Moreover, these platforms can select past cases that may relevant to solve a given problem and are thus of value for the disputing parties as well.
- *Negotiation and Mediation*: Simply suggesting a solution or an outcome for a solution is generally not enough, that is, parties will most likely not agree on it. There is thus the need to let the parties change that solution or negotiate a new one in which they agree. It is thus important to consider successful dispute resolution methods such as negotiation or mediation. In UMCourt, this is done with a particular emphasis on the transparent integration with the information that is retrieved, that is, the same information that is used to inform the disputing parties is also used to guide negotiation and mediation processes.

4 Case-Based Reasoning in UMCourt

As stated above, the estimation of possible outcomes is an important problem as it is the base for a supported decision and planning process in dispute resolution. In this sense, there are several approaches that can be considered. Utility functions have considerable applications in the legal domain (see for example [5, 19, 35]). A good example is the above-mentioned Family_Winner system, in which Bellucci and Zeleznikow have applied this technique to family mediation [9], after they observed that an important way in which family mediators encourage disputants to resolve their conflicts is through the use of compromise and trade-offs [55].

Artificial intelligence techniques have also been used to address the same challenge. James Popple used a simple Expert System to provide advice in fields of case law that have been specified by a legal expert in a specification language defined by the author [34]. Split-Up combines rule-based reasoning with neural networks in order to build a system to assist parties involved in property settlements in a divorce process in Australia [52]. UMCourt is framed in this group of systems that apply AI techniques to ODR.

In the Case-Based Reasoning paradigm, cases represent past experiences stored in memory. Each case contains the description of the problem that was encountered (the initial state of the world), the solution adopted to deal with it, and the outcome (the final state of the world). As this case memory grows, the system gains more knowledge about the problem and the possible different ways to deal with it. In each new problem that is faced, past cases are analyzed in search of similarities, and solutions of similar cases are adapted to be used to solve the new problem. At the end of this process, it is possible to learn with the success or failure of the application of the solution to the new problem.

Choosing a case-based approach was not an arbitrary decision. In fact, law itself relies on the concept of case and implements a very similar concept: the legal precedent [28]. This concept is defined by the Oxford dictionary as “*a previous case or legal decision that may or must be followed in subsequent similar cases.*” There are legal domains in which precedents are divided into two types: binding and persuasive precedents. The first type is

generally the result of a process in a higher court and means that all lower courts must honor it. The persuasive precedent, on the other hand, arises from cases that are decided in lower courts and does not represent an obligation. However, in Portugal, as in civil law countries in general, this separation does not exist as all the precedents are simply persuasive. Judges may follow a precedent decision, but they are also entitled to decide otherwise, since judges are only bound to apply the legal norms and interpretation differences between judges may occur [17,56].

A CBR model would, at first sight, be more suited to be used in common law: a law that is developed through decisions of courts. This, however, does not mean that it cannot be used in a civil law context, such as the Portuguese one, in which laws are written by a legislature's enactment [56]. It is our conviction that CBR is an appropriate method for such a problem-solving domain as the cases and their outcomes are often very similar to known past occurrences. This conviction is shared and supported by several authors. Ashley, in particular, poses the question: "*should researchers in a civil law jurisdiction pursue work on implementing AI and Law models of case-based legal reasoning in a civil law context?*." He then answers the question with a conclusive "*the answer may well be, Yes!*" [4]. In that sense, throughout this section, we will be guided by five main questions synthesized in [38] and, by answering to each one of them, define our agent-based CBR system, namely:

1. What makes up a case? How can we represent case memory?
2. How is memory organized? What are the indexing rules?
3. How does memory change? How do the case memory and indexing rules change over time?
4. How can we adapt old solutions to new problems? What are the similarity metrics and modification rules?
5. How can we learn from mistakes? What are the repair rules?

We commence by developing some basic concepts that will later be useful for understanding the CBR process model and the role of each of the presented agents.

4.1 Preliminary considerations

Let us start with a high-level description of the CBR process as used in our specific domain of application. There is a database of cases, which embodies information from legal disputes settled by courts, with their respective outcomes. The CBR process is mainly intended to look at these cases and infer what would happen if the two parties currently involved in the dispute would have the matter resolved by a court. This basically consists in selecting the most similar cases and presenting their outcomes to be analyzed by the disputing parties. More than that, the system compiles information that is information for the decision-making processes, including the previously mentioned BATNA, WATNA, ZOPA and MLATNA. With this information, parties can build their own image of the whole problem and of the possible paths and respective consequences.

At the end of this process, parties may choose to accept the proposed outcome or they may reject it. If they accept it, the value of the proposed case is increased, denoting that it was successfully applied. Alternatively, if the parties reject the proposed outcome, that value is decreased and, should the parties advance to court, a new case is learned by the system containing all the information that had already been provided by the parties and the outcome from the court. Otherwise, parties can also make use of a traditional dispute resolution method (e.g., negotiation, mediation), integrated in the UMCourt platform.

4.1.1 Parsing and indexing

In an indexing process, property values are assigned to cases so that they can be easily identified according to those properties. Therefore, an important part of this process is the one of selecting the right properties for indexing the cases. Evidently, this is, once more, domain-dependent. In our case, we index the cases according to:

- Background information;
- Norms addressed by parties (i.e., legal information);
- Objectives of the parties;
- Date of the case;
- Outcome.

This allows agents to make requests to the database such as “all the cases in which norm X is addressed”, “all the cases that are within a given time frame” or “all the cases with a given outcome.” Our objective in indexing the cases has essentially to do with efficiency issues. It would be a rather time-consuming task to constantly search all the files for some parameters. By indexing them, this task is simplified to querying the database for the identifiers of the cases that match the desired criteria.

Indexing is often looked as a challenge to overcome in CBR systems, mainly when the indexing method is manual as it tends to become a slow and hard task. Automated methods are thus preferred. In the case of this work, indexing is an automated process. An indexing agent frequently monitors a predetermined folder in the system and indexes new files as they are added. However, the validity of their structure must first be determined.

This agent therefore starts by using a SAX parser to parse the content of the new file and checks whether it is valid according to the current Schemas. If it is, the case is indexed. The indexing process is as simple as reading the appropriate contents from the parser result and adding the identifier of the case to the appropriate tables in the database. However, only the information that important for the retrieval of the cases is stored in the database while the remaining is kept in the files. All the information needed for indexing the case has been previously provided by the parties, while filling in the information requested, or is automatically inferred by the system (e.g., dates, norms addressed), ensuring that the indexing process itself is a completely automated one.

Additionally, a record of failed cases is also maintained, in which each case is associated with the reasons for failure. These cases are intended to be looked at when the CBR process fails in the search for reasons for the failure, as will be seen below. Moreover, when the Retriever agent is retrieving the cases from memory, it may also look at this record before taking actions that influence the parameters of the search (e.g., changing the similarity threshold), in order to prevent errors that have already occurred in the past.

Once the information is indexed, a wide range of operations become possible. As an example, it is possible to mine the database for association rules [51], which are then used to classify the cases (Fig. 1) [40]. This is useful for retrieving cases, as will be seen below, by selecting all the cases that belong to the same category or group, that is, all the cases for which a given rule mined is true.

It is also possible to quickly build different interpretations of the information, suited for specific applications. One of these interpretations is the representation of the information as vectors of binary entries. This is a fairly simple algebraic model for representing text documents in which, instead of using textual fields, a case is represented as a vector. Specifically, in the context of this work, a case is seen as a vector V of binary entries, in which each entry $i < N$ corresponds to a fixed descriptor from the descriptor vector D of size N . Thus, the




Fig. 1 Indexed cases can be organized by means of association rules: all the cases for which a given rule is true belong to the same category

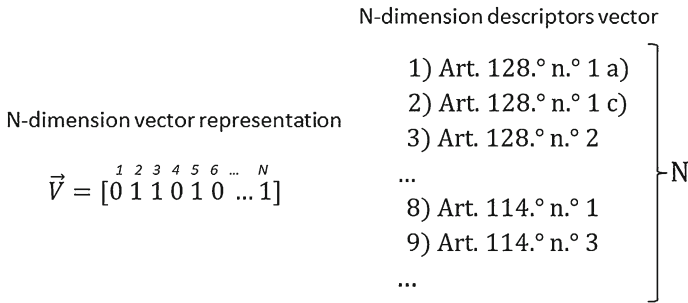


Fig. 2 N-dimension vectorial representation of a case

Fig. 3 Representation of a database of cases as a matrix

$$DB = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & \dots & N \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ \dots \\ N \end{matrix} & \begin{bmatrix} 1 & 0 & 1 & \dots & 1 \\ 0 & 0 & 1 & \dots & 0 \\ 0 & 1 & 1 & \dots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 1 & \dots & 1 \end{bmatrix} \end{matrix}$$

value of each binary entry denotes the presence or absence of that descriptor on the case. Descriptors denote important components of a case (e.g., legal norms, objectives of the party, winner of the case). Thus, one can look at a vector which represents a case and, considering the descriptors vector D, determine which information is or is not present on the case (Fig. 2).

Basically, this representation of a case allows seeing which norms are addressed by each party, which are their objectives and what is the outcome. It is thus a very concise way of representing all this information, demanding very few resources to handle and to store. Following the same line of thought, a database with m cases in which each case is described by N descriptors can be represented as an m-by-N matrix in which each line is a vector representing a case (Fig. 3).

4.1.2 Case structure

In the CBR paradigm, a case is the basic unit of information. It represents a past experience and is thus a contextualized piece of knowledge. CBR allows us to estimate the outcome of a current experience by looking at a past similar experience (that took part in an equivalent context) and its respective outcome. The definition of the information that should be contained in the case is therefore a major task in the development of a CBR model. Generally, the information contained in a case can be organized into three distinct categories [2]:

- *The Problem*: The description of the problem, including the context of applicability, the initial state of the world and all other information that may help to define the problem.
- *The Solution*: A description of the list of steps that were taken in order to solve the problem. The description should be exhaustive enough to be applied again in an autonomous fashion and abstract enough to be adapted to new problems.
- *The Outcome*: The consequences of the application of the described solution, that is, the final state of the world.

Whereas these categories do depend on the problem domain, there may be applications which consider only the problem and the solution (in which it is possible to derive solutions to new problems) while others consider the problem and the outcome (making it possible to estimate outcomes to new problems). In our particular case, we consider the three categories, namely (Table 2):

Being more specific, cases are precisely defined using the Extensible Markup Language standard (XML). In that sense, we formally define the structure of the case using the XML Schema language and, consequently, the cases are XML files stored in the file system. Using XML allows us to take advantage of other related specifications, namely XML Encryption, XML Signature or XSL Transformation. Given the size of the XML Schemas that define the cases, these are not detailed here.

Table 2 The information that makes up a case

Category	Information type	Description
Problem	Background	Basic information about the parties and the dispute such as party's personal information and location, dispute starting date, witnesses
	Objectives	A list of the initial objectives of each party toward the dispute, that is, the expected outcome
	Legal	Legal information such as the laws and norms addressed by the parties and witnesses to support their claims or the guilty statement
	Dates	All the important dates of the case
Solutions	List of actions	A list of the actions performed by the parties in order to achieve the outcome. Generally, these actions comprise trade-offs
Outcome	Outcome description	A list of items that describe the outcome in terms of indemnities to be paid, among others
	Value	A value denoting the percentage of successful applications of this case in the dispute resolution process

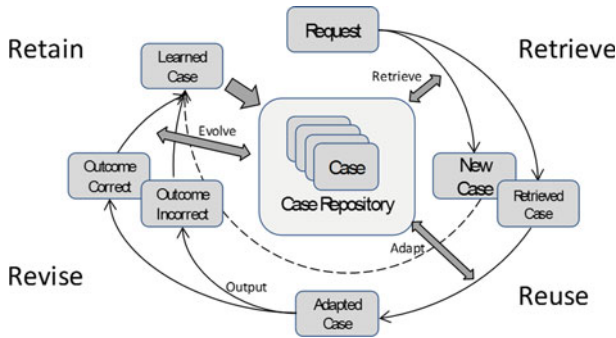


Fig. 4 The CBR model implemented

4.2 The CBR process model

We have, until now, described how the experiences are acquired, organized and stored in the system. However, purely storing information is not enough. We need to define the processes that acquire and use this information. To define this process, we have looked at the work of Aamodt and Plaza [1] and Riesbeck and Bain [37] and made the necessary revisions in order for it to be used in our domain. It is a four-step model composed by the Retrieve, Reuse, Revise and Retain phases (Fig. 4).

4.2.1 Retrieve

The process of retrieving may be triggered under different scenarios. It consists of selecting the cases from the case memory that are of relevance for solving a problem, by looking at their indexes. Unlike database searches that target a specific value in a record, retrieval of cases from the case base must be equipped with heuristics that perform partial matches, since in general there is no existing case that exactly matches the new case [47]. There are several techniques for retrieving cases. Our approach is a hybrid one and consists of two different phases [12]. In the first phase, a pre-selection of cases is performed. On the second phase, the pre-selection is evaluated, eventually more refined using similarity functions, and the cases are sorted. In each of these phases, two different algorithms can be combined. Each of these algorithms is analyzed in terms of their performance in the Results section.

In the pre-selection phase, either a template retrieval or a classification algorithm can be used. The main objective of these two algorithms is to narrow the search space so that the algorithms of the evaluation phase perform quicker. Either algorithm has pre-selection rules, intended to ensure that the result of the pre-selection complies with given parameters (e.g., number of cases retrieved). In that sense, in order to comply with these rules, the algorithms may run more than once per search, dynamically making adjustments to the search settings.

The template retrieval algorithm basically pre-selects all the cases that may be similar. This can be done as we know a priori which type of cases have the possibility of being similar and which ones do not. As an example, cases that do not address the same norms cannot be similar as they address different legal domains. In that sense, template retrieval works much like SQL queries. Given a case, a set of cases are retrieved from the database with the guarantee that they have certain characteristics (e.g., address at least one of the norms addressed by the new case).

The classification algorithm, on the other hand, uses association rules mined previously to determine which category a case belongs to. An example rule could be “if party1 uses norms A and B, then party1 wins.” Thus, the algorithm determines in the first place which category the new case belongs to and then pre-selects all the cases of that category.

In the evaluation phase, additional operations are performed on the pre-selected cases, such as determining their similarity and ordering them according to it. Similarly, in this phase, two different approaches can be used: one based on a nearest neighbor algorithm and another one based on the cosine similarity concept. The basic idea is that any of these algorithms is only applied to the reduced set of pre-selected cases instead of applying them to the whole case base. Thus, the algorithms start by determining a similarity measure between each pre-selected case and the new case.

In the first approach, the nearest neighbor shown in Eq. 1 is used.

$$\frac{\sum_{i=1}^n W_i * f\text{sim}_i(\text{Arg}_i^N, \text{Arg}_i^R)}{\sum_{i=1}^n W_i} \tag{1}$$

In this equation,

- n —number of elements to be considered to compute the similarity;
- W_i —weight of element i in the overall similarity;
- $f\text{sim}$ —similarity function for element i ;
- Arg —arguments for the similarity function representing the values of the element i for the new case and the retrieved case, respectively N and R .

In this algorithm, the weights are at this moment determined by a law expert, based on the importance that according to his experience each of the components of the similarity measure has. However, in the future, we intend to let the system change these values dynamically, looking at past iterations, in an attempt to select the most appropriate weights for each case.

We now detail the information of the case that is considered to be relevant for the computation of the similarity with the nearest neighbor approach, that is, the components. According to our scope of application, we consider three types of information: the objectives of each party, the norms addressed by each party and by the eventual witnesses, and the date of the dispute. Both the norms addressed and the objectives are lists of elements. The similarity function consists in comparing two lists (Eq. 2). The similarity is higher when the two lists have a higher percentage of common members. As for the date, the similarity function verifies whether the two dates are within a given time range; the similarity is higher when the two dates are closer.

$$f\text{sim}_{\text{list}} = \frac{L_N \cap L_R}{n}, n = \begin{cases} |L_N|, & |L_N| \leq |L_R| \\ |L_R|, & |L_N| > |L_R| \end{cases} \tag{2}$$

In the second approach, a similarity measure based on the concept of cosine similarity is used [39]. This method uses the binary vector representation of the information and is based on the notion that the similarity between two vectors can be determined by finding the cosine of the angle between them. Given two vectors of attributes A and B, with N entries each, the cosine similarity, θ , is determined as shown in Eq. 3. The resulting value of similarity of this equation ranges from 0 to 1, with 0 being the smaller value of similarity and 1 meaning that they have exactly the same information.

$$\text{sim} = \frac{A \cdot B}{\|A\| \|B\|} = \frac{\sum_{i=1}^N A_i * B_i}{\sqrt{\sum_{i=1}^N (A_i)^2} * \sqrt{\sum_{i=1}^N (B_i)^2}} \tag{3}$$

This second method of computing similarity is quite simpler and faster as it uses the vectors of binary entries. However, contrary to the previous similarity function, it does not allow to assign weights to the several components of the case. This may or may not be a disadvantage, depending on the scope of application.

In the evaluation phase, another operation takes place which is independent of the method used to compute the similarity: the determination of the utility of each case. This value of utility quantifies how much each case is desirable for each party. After having access to these values, a party can thus compare the retrieved cases in terms of their similarity and their utility and make a personal evaluation of potential advantages versus potential likeliness. The process of computing the utility is relatively simple. Basically, as stated before, each stored case has a solution and an outcome. The outcome denotes the result of applying the solution in the context of that case. The application of a solution to different problems generates different outcomes. The key idea is to apply the solution of each retrieved case to the current problem and determine its value of utility. As a simple example, one could have in a past case:

Outcome = worker receives indemnity of 2500
*Solution = antiquity*25*d_wage*

This example denotes a case in which a worker received an indemnity of 2500, that was determined by the formula depicted in the solution, defined by Decree of Law (DL) 7/2009 (Portuguese laws). As, in that case, the worker had an antiquity of 5 years and a daily wage of 20, the value of the outcome was obtained applying the formula with these values. Now, let us consider that this case was retrieved because it was similar to the new case and that, in this new case, the worker has an antiquity of 15 years and a daily wage of 17. For this worker, the application of the same solution would mean an outcome of 6375. This adds the last information that is compiled in this phase.

Independently of the method used, the output of this phase is a list of similar cases ordered by their degree of similarity. This information can then be used by the disputing parties, neutrals or the platform itself in a wide range of applications. Figure 5 shows the textual output of a typical execution of the retrieve process, combining the template retrieval with the nearest neighbor approaches. This figure highlights the process of refining the pre-selection with successive iterations and changes in the search parameters. Figure 6 depicts the interactions that took place between the several agents in this specific example.

All the retrieve process is coordinated by the Coordinator agent, which has the knowledge to interact with the remaining agents, request the necessary tasks and handle the results. Hence, when a new retrieve request arrives to the Coordinator (typically coming from a user interface or another agent), it starts a new Complex Action (CA). A CA is an action that will use more than one Simple Action. In the case of the retrieve process, the CA includes the following SAs: (1) Pre-Selection, handled by the TemplateRetrieval agent; (2) Database Query, handled by the Database agent; (3) Computation of Similarities, handled by the NearestNeighbor agent; (4) Computation of Utilities, handled by the Utilities agent; and (5) Return Results, handled by the Coordinator agent. Note that some of these actions may repeat and that some must be sequential while others may be executed in parallel (e.g., computation of utility and similarity). So, in this case, the Coordinator starts a new case retrieval by requesting a pre-selection of cases from the TemplateRetrieval.

In this request, the Coordinator includes only a single rule: to start the pre-selection by considering the article of each norm. The TemplateRetrieval then formulates a database request containing a SQL sentence and sends it to the Database agent, which returns the result. The TemplateRetrieval analyses the result from the Database and determines that the

```

New Pre-Selection task arrived: < 1291886422498, 1286439756546 >
****TemplateRetrieval: Requesting for cases that match the rules: (
    depth: article
)
2010/12/09 09:20:23 : Connected to Database
****Database: Received request to query database from TemplateRetrieval@TIARAC1:1099/JADE
2010/12/09 09:20:23 : Connected to Database
****Database: Returning query result...
****TemplateRetrieval: Received a pre-selection of 77 cases
****TemplateRetrieval: Rule violated: "Max cases". Deciding new pre-select rules...
****TemplateRetrieval: Changed search depth to: number
****TemplateRetrieval: Requesting for cases that match the rules: (
    depth: number
)
****Database: Received request to query database from TemplateRetrieval@TIARAC1:1099/JADE
****Database: Returning query result...
****TemplateRetrieval: Received a pre-selection of 77 cases
****TemplateRetrieval: Rule violated: "Max cases". Deciding new pre-select rules...
****TemplateRetrieval: Changed search depth to: item
****TemplateRetrieval: Requesting for cases that match the rules: (
    depth: item
)
****Database: Received request to query database from TemplateRetrieval@TIARAC1:1099/JADE
****Database: Returning query result...
****TemplateRetrieval: Received a pre-selection of 28 cases
****TemplateRetrieval: SUCCESS: Returning results to agent( agent-identifier :name Coordi
****Coordinator: Received results from pre-selection of task: 1291886422498
****Coordinator: Requesting computation of similarities for task 1291886422498
****NearestNeighbour: Received request for similarities: 1291886422498
****NearestNeighbour: Sending list of computed similarities to: ( agent-identifier :name
****Coordinator: Received results from computing similarity of task: 1291886422498
****Coordinator: Returning result of Get Similarity Request
****Coordinator: Ending sub task 1291886422498
****Coordinator: Updating task 1291886422370
****Coordinator: Starting a new Utilities task: 1291886425356
****Coordinator: Requesting computation of utilities for task 1291886425356
****Utility: Received Compute Utility request from ( agent-identifier :name Coordinator@T
****Utility: RECEIVED REQUEST FOR 28
****Utility: Sending results from Utility request to ( agent-identifier :name Coordinaton
****Coordinator: Received results from computing utilities of task: 1291886425356

```

Fig. 5 The output of the execution of a typical case retrieval process

rule that establishes the maximum amount of cases is being violated. Thus, it reformulates the request to the Database in order to get fewer cases. This happens once again until it receives a pre-selection of cases that is in accordance with the rules. At this point, the TemplateRetrieval returns the pre-selection of cases to the Coordinator, which will request the computation of the similarity and the utility values of each pre-selected case to the corresponding agents. In this context, the similarity denotes the extent to which a retrieved case is similar to the current case, and the utility quantifies how much the user would win or lose if the solution of the given case would be applied to his case. When these values are returned, the Coordinator is finally able to compile the result of the case retrieval task and return it to the requesting entity.

4.2.2 Reuse

In the Reuse (or adapt) phase, solutions are adapted to match the characteristics of the new case. In this work, solutions are, as stated before, lists of steps that the parties took in order to achieve the verified outcome (typically trade-offs) in a previous case. This is structured information that can thus easily be changed. Therefore, in order to adapt a case that requires to do so, the system starts by looking at the information of its solution. It then searches for differences with the new case. It is in these different points that the adaptation will take

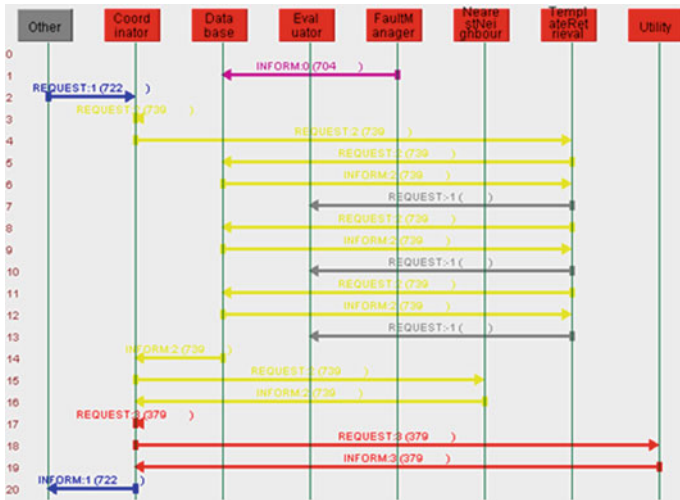


Fig. 6 Visualization of a typical communication between agents for retrieving similar cases

place. Several actions can be performed, namely replacing basic information (e.g., names, addresses, dates, places), omitting steps that do not apply in the context of the new case or adding new steps that are needed. The key idea is to change the cases retrieved so that they can be understood by the disputing parties, under their context.

4.2.3 Revise

In this phase, a solution and the corresponding justification is presented to the parties, after which their behavior is analyzed. In order to do so, the cases and their likeliness to occur, according to the characteristics of the current case, are presented to the users in a graphical fashion (Fig. 7). In this graphical representation, the parties can see the several cases that were retrieved in the form of the small colored circles. By clicking in these circles, an interface pops up that shows details about the corresponding case and about the reasons for its selection. These cases span the space between the BATNA and WATNA of each party, so that they can be compared in terms of their utility for each party. In order to determine the values of the BATNA and WATNA, the system uses logic rules defined after the Portuguese labor law. These are rather simple rules that establish the values of eventual indemnities and other parameters of the outcome, based on concepts like worker antiquity, work hours, extra hours, night work, among others.

Let us take, as an example, the Portuguese labor law domain, as depicted in DL 7/2009 (Portuguese laws), considering a scenario in which a worker wants to end the labor contract claiming that the employer did not pay the last three salaries. According to Article 394th, nr. 2 a, the lack of regular payment of the salary constitutes a just cause for a worker to end the contract. Moreover, Article 394th, nr. 1 states that when there is a just cause, the worker can immediately end the labor contract. The first question is thus to determine whether the lack of payment exists or not, and thus, of a just cause for ending the contract. Assuming that this has been proved, let us try to determine the best and worst scenarios, from the point of view of the worker. The most important norms are found in Article 396th, numbers 1, 3 and 4. Number 1 states that, if Article 394th is true (there is just cause for ending contract),

the worker is entitled to 15–45 days of salary plus indemnity for each year of contract. It also states that this value varies according to the degree of wrongfulness of the employer and that the total indemnity paid to the worker should not be inferior to three salaries plus indemnity. However, number 3 states that the indemnity paid can be higher whenever the worker suffered property damage or other damage, of higher value. Finally, number 4 states that, in the cases of a temporary employment contract, the value of the indemnity cannot be smaller than the value of the salaries that would be received until the end of the contract. Thus, the resemblance between legal norms and rules allows formalizing the computation of the BATNA and WATNA in the form of IF–THEN rules. The example addressed here is presented below.

A simplification of the rules that allow the computation of the BATNA and WATNA values according to the Portuguese labor law is given below. This example code considers only the case in which a worker ends the contract with a just cause. *M_SALARY* denotes the monthly salary; *D_SALARY* denotes the daily salary; *M_REMAINING* denotes the months remaining until the end of the temporary contract; *+VARIABLE* denotes an unknown value, higher than *VARIABLE*.

```

Def_Rule 396
if RULE_394 then
  WATNA := 3 * (M_SALARY + SENIORITY)
  if TEMPORARY_CONTRACT then
    if WATNA < M_REMAINING *(M_SALARY + SENIORITY) then
      WATNA := M_REMAINING *(M_SALARY + SENIORITY)
  if WATNA < 15 * (D_SALARY + SENIORITY) then
    WATNA := 15 * (D_SALARY + SENIORITY)
  BATNA := 45 * (D_SALARY + SENIORITY)
  if BATNA < DAMAGE then
    BATNA := +DAMAGE

```

However, in Fig. 7, besides looking at the retrieved cases and the BATNA and WATNA values, a disputing party can analyze additional information. Clusters are created that are intended to show some grouping between the cases, allowing a first analysis from a higher point of view. The cases belonging to the same cluster are shown depicted with the same color. The interface also shows, for each cluster, the mean values of similarity and utility. A linear regression is also drawn, which may help the disputing party to determine in which region an outcome is more likely or more valuable, according to the utility. This regression is also used to depict the MLATNA: The region in which the line of the regression is green instead of black is the most likely region for the outcome. Finally, using this interface, it is also possible to delimit the ZOPA, the Zone of Potential Agreement. This concept was proposed by Raiffa [36] and denotes the intersection of the range of possible outcomes of the parties, that is, what can happen.

4.2.4 Retain

This CBR process can be used as it was presented until now, as a tool to compile useful information, or it can be integrated in a higher-level process (Fig. 8). Thus, at the end of the last phase, one of several things may happen. The parties may choose to accept the proposed outcome by committing to implement the steps in the solution. In this case, the system increases the value of the case proposed, indicating its successful application to solve a case with these specific characteristics. Alternatively, the parties may not agree to accept the proposed solution. In this case, the parties can make use of a negotiation module that




Fig. 7 The prototype of the interface showing the user the compiled information. This interface shows several cases retrieved from the case base, displayed on a Cartesian plane. Only the cases more similar to the target case are selected. The horizontal axis represents the utility, while the vertical one represents the similarity. Each colored dot represents a case and different colors indicate that cases belong to different clusters. Cases can be clicked on to access further information

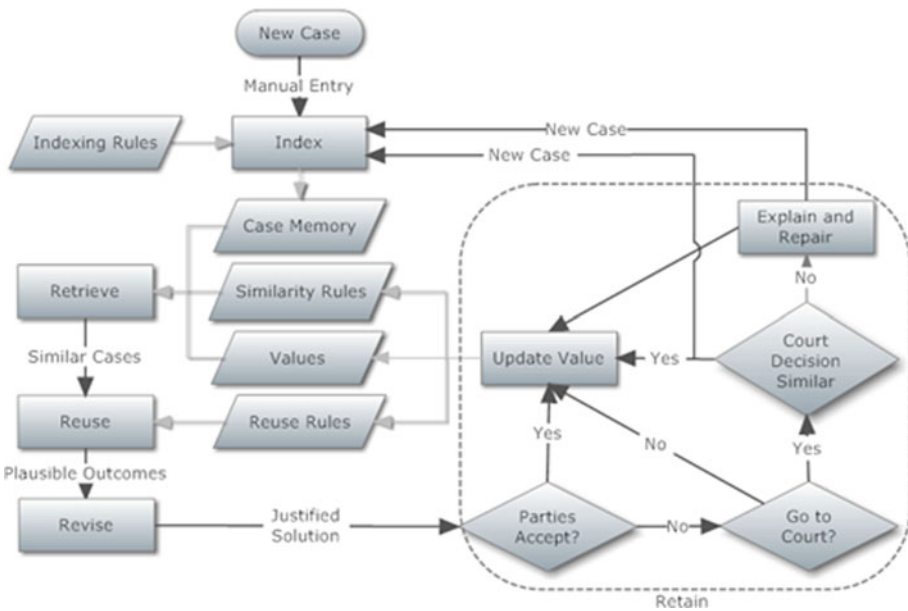


Fig. 8 A Flowchart depicting the integration of the CBR model presented in a higher-level complete process of dispute resolution. Rectangles represent processes, parallelograms represent information structures, and rhombuses represent decisions

allows them to search for alternative solutions by modifying proposed solutions in order to maximize the hypothesis of agreement. Moreover, the parties can decide to stop the alternative resolution process and either advance to a court or chose another dispute resolution method. In this case the process is considered to have failed.

If they decide not to go to a court (e.g., by choosing other methods for conflict resolution), the value of the case proposed is decreased, denoting that it was not applied successfully. If this occurs, the process ends. In order to gather more information that can help to improve the system, the parties may be asked to complete a questionnaire stating why, in their opinions, the process failed. This information will later be useful for improving the system by changes in the process model, in the way information is provided, among others. Alternatively, if the parties decide to go to court, the process continues by waiting on the decision of the court. If the decision concurs with the solution that was proposed by the system, the value of the proposed case is increased and the new case that resulted from the court is added to the case base. This, in fact, constitutes another innovation of the UMCourt platform, as the only information needed to complete the data of the case is the outcome. All the remaining information has already been provided by the parties when using the system. This makes the addition of cases a simple and fairly automated process. If, however, the outcome of the trial differs from the proposed solution, the system will decrease the value of the proposed case and will try to find reasons for the failure of the CBR process. In this situation as well, a new case is added to the case base.

In the case in which the parties want to work the solution of the dispute, using as a starting point the proposed solution, they can use the negotiation or the mediation module embedded in the platform. This is described in detail in the Sect. 5.

5 Using negotiation and mediation to improve CBR efficiency

As mentioned above, the CBR process may fail, namely if the parts do not agree on the outcome proposed. In such scenarios, negotiation or mediation are additional paths that can be followed to search for alternative solutions. This requires the active participation of the parties. These two approaches have been implemented, supported by the same CBR process presented, following a blackboard approach. Therefore, there is a shared space to which the parties publish and debate proposals. This is implemented by the Blackboard agent, which is able to receive and interpret messages from the parties, and take decisions according to the content of those messages in order to guide the process to a successful conclusion (Fig. 9).

5.1 Setting up the process

The mediation and negotiation processes build partially on the previously described functionalities. In that sense, before the actual start of the process, some information is generated that will help define the whole process. This information includes the already mentioned BATNA and WATNA. These values are particularly interesting during a negotiation or mediation as, besides helping parties have a clear picture of the possible outcomes, they might also be used as a way to put pressure on the other party, especially in dispute resolution procedures that allow the choice of going to court [15].

The remaining values mentioned before (e.g., ZOPA, MLATNA) as well as the ordered list of similar cases will also be used as described ahead. Once all this information is compiled, parties may register on the platform by means of their personal agents and the mediation process is ready to start. How this process takes part is detailed below.

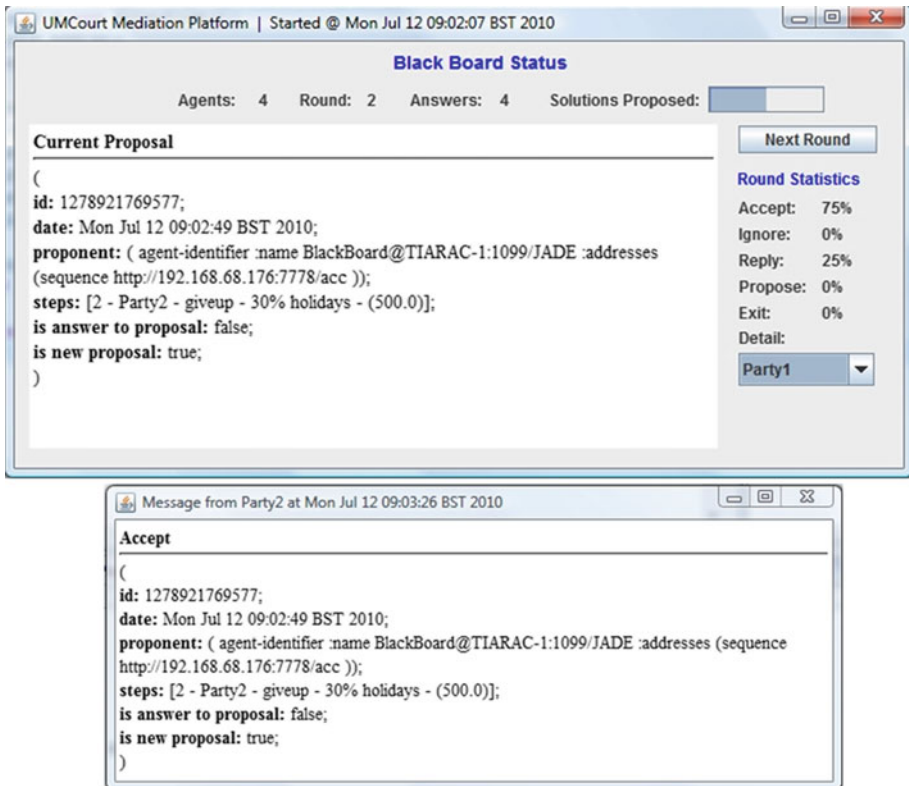


Fig. 9 A simple visual interface for showing the black board status (*upper figure*) and the content of a reply (*lower figure*)

5.2 Mediation

The objective of the mediation process is to suggest alternative solutions to the parties in order to overcome impasses in scenarios in which the parties cannot agree on a given solution. Moreover, each solution that is proposed can be changed by the parties in an iterative process that allows them to draw the most mutually satisfactory solution. This means that this approach is suited for cases in which the two disputant parties have a behavior in which they effectively propose solutions and want to actively participate on the definition of the outcome. Specifically, parties that have a collaborating or compromising conflict style [44] should use this approach. All this process is managed by an electronic mediator represented by the blackboard: It receives messages from all the parties, does the necessary computation and publishes a new proposal to be analyzed. These messages may include proposing new solutions, refusing or agreeing with the current proposal, ignoring a proposal, among others. Furthermore, the blackboard has a list of solutions taken from past known cases, ordered in decreasing order of similarity, that is, the MLATNA is at the head of the list.

After the parties register, the process starts with the blackboard publishing the first proposal for solution, which is the solution given by the MLATNA. Whenever a new proposal is published by the blackboard, each party can perform one of several actions. The party can accept the proposal denoting that it agrees with it or, otherwise, it can refuse it. If the party

wants to propose a new solution that consists of the modification of the present one, it can make the intended changes to the current proposal. Alternatively, the party may also suggest a completely new solution. Finally, one party may also choose to ignore the current proposal or it may leave the mediation process, ending it.

This iterative process is entirely controlled by the blackboard agent and is organized into sequential rounds (Fig. 10). Each round starts with the blackboard publishing a new proposal. This proposal may have several origins. It can come from the list of similar cases, or it may be computed from the previous answers of the parties. In this case, if only one party answered with either a new proposal or a modified proposal, that will be the next solution to be proposed by the blackboard. Otherwise, the blackboard will search for similarities between the several proposals of the parties and compute a new solution based on those similarities. If, however, the proposals are divergent, the blackboard will propose the following solution from the list of similar cases, if any. This process ends when all the agents send an agree message or at least one sends an exit message. Moreover, the process also ends when the parties do not agree on a solution before the system suggesting all the solutions from the list of similar cases, that is, the system runs out of known solutions and the parties did not agree on any of the suggested ones.

The description of the algorithm is presented below.

Algorithm Mediation is

input: List of cases with solutions, C

 List of parties, P

output: A solution for the dispute

```

round := 0          (identifies the round)
msgSet := []       (a list of received messages)
agree := 0         (the number of agents that agree)
exit := 0          (the number of agents that exit)
proposal := []     (the current proposal)
i := 0             (the case C being currently proposed)

while (agree < length(L) and exit < 1)
  if (round = 0)
    proposal := Ci
  else if (msgSet contains "ReplyTo" or msgSet contains "Propose")
    set := ReplyTo in msgSet U Propose in msgSet
    proposal := intersect set
  else i++
    if (i = Length(C))
      return null
    else proposal := Ci
  publish proposal
  msgSet := []
  for each party in L
    msg := receive from party or timeout
    msgSet := msgSet U msg
  agree = count "agree" in msgSet
  exit = count "exit" in msgSet
if (exit > 0)
  return null
else
  return proposal

```





Fig. 10 Several steps of the mediation process model




Fig. 11 Representation of the information needed for guiding a negotiation process, consisting of several sessions with respective value of similarity

5.3 Guided negotiation

This process has a similar objective to the mediation process but tries to achieve it in a different fashion. It has been developed having in mind conflict scenarios in which the parties exhibit essentially an avoiding conflict style [44], that is, one or more parties are reluctant or have difficulties in generating possible solutions. In that sense, the responsibility of generating solutions falls entirely on the system to which the parties can answer in three ways: They either ignore, agree or do not agree with the current proposal.

The negotiation process develops around a list of previous negotiation processes provided by the CBR mechanism that will be used to guide the parties through their current process (Fig. 11). Each negotiation process is defined by a list of events or steps in which each one represents a sequential stage in the negotiation process.

Each step contains the initial state (denoting the initial conditions), a proposed solution (a list of actions or steps that one or more parties must perform, usually trade-offs), the answer of each party to the proposed solution and the final state (denoting the final conditions). The process starts with the blackboard publishing the first negotiation step of the most similar negotiation session, which is determined using a similarity function from the described CBR algorithm. At this point, parties reply to the published proposal. When answering positively, the parties will be stating that they agree to implement the steps that the negotiation step describes and that they accept the corresponding outcome. When answering negatively, they state that they agree neither with the proposed solution nor with the outcome.

At the end of each round, that is, when all the messages from the parties are received by the blackboard, the content of these messages is analyzed. At this point, the negotiation can follow three distinct paths. If all the agents agreed on the current proposal, the negotiation process successfully terminates as a solution that satisfies all parties has been achieved. Otherwise, if the majority of the agents agree with the current proposal, the blackboard will propose the next step on the same negotiation process on the following round. Alternatively, if the majority of the agents do not agree with the current event, the blackboard will fall back and select the following negotiation process on the list and start again from its first step.

This denotes that the majority of the parties do not agree with the current negotiation path and a new one will be selected. This process repeats until either all the parties agree on a solution, the system reaches the last step of the last negotiation process in the list or the majority of the agents do not agree on an event of the last negotiation process. The whole negotiation process is depicted in Fig. 12 and in the algorithm described below.

```

Algorithm Mediation is
: List of previous mediation processes, L
         List of parties, P
output: A solution for the dispute

round := 0      (identifies the round)
msgSet := []   (a list of received messages)
agree := 0     (the number of agents that agree)
exit := 0      (the number of agents that exit)
proposal := [] (the current proposal)
i := 0         (the mediation process being currently used)
j := 0         (the current step being proposed)

while (agree < length(L) and exit < 1)
    mediation := Li
    solution := mediationj
    publish proposal
    for each party in P
        msg := receive from party or timeout
        msgSet := msgSet U msg
        agree = count "agree" in msgSet
        exit = count "exit" in msgSet
if (exit > 0)
    return null
if (agree = Length(P))
    return solution
else if (agree > Length(P)/2)
    j++
    else i++
    j := 0
if (j = Length(mediation))
    j := 0
    i++
if (i = Length(L))
    return null
round++

```

6 Results

Let us now describe the main results of the work presented in this paper. UMCourt is being tested and assessed in a setting with the characteristics that follow. Given the complexity of the Portuguese labor law, for test purposes we are dealing with a restricted set of norms. In that sense, we are considering a group of 36 norms of the DL 7/2009, from February

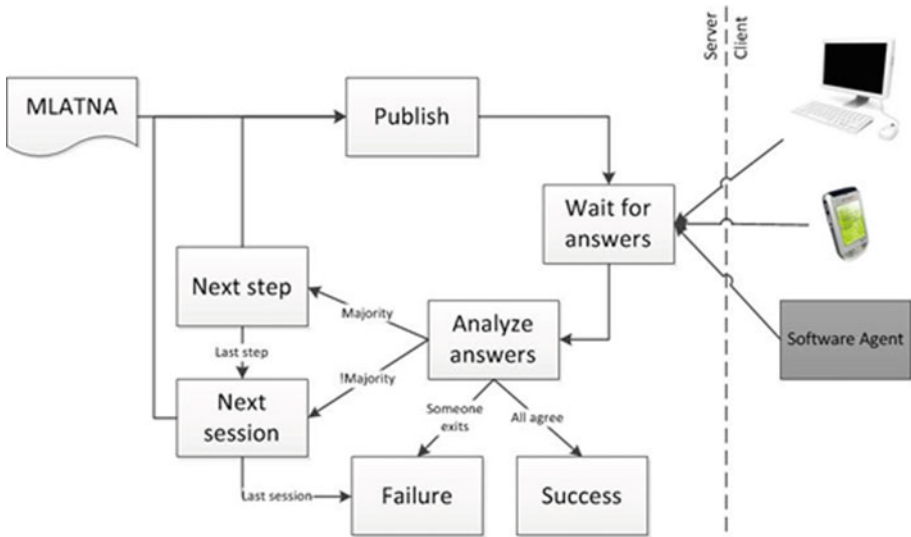


Fig. 12 The process of guiding a negotiation coordinated by the blackboard

Statistics of the case base:

Article 118, Funções desempenhadas pelo trabalhador, 6%
 Article 117, Efeitos de falta de título profissional, 4%
 Article 129, Garantias do trabalhado, 38%
 Article 128, Deveres do trabalhador, 44%
 Article 126, Deveres gerais das partes, 6%

Total cases in the database: 127

Fig. 13 A description of the database detailing which aspects of the Portuguese labor law are addressed by the cases and in which proportion. Article 118 concerns functions performed by the employee; Article 117 is about the effects of the lack of professional title; Article 129 concerns employee’s rights; Article 128 depicts employee’s obligations; and Article 126 concerns the general obligations of the parties

12, 2009. These norms were selected because they are generally present in most of the labor-related disputes and address the following domains: (1) functions performed by the employee; (2) effects of the lack of professional title; (3) employee’s rights; (4) employee’s obligations and (5) general obligations of the parties. The database used in these tests contains a total of 127 indexed cases (Fig. 13). The representation of these cases in vectors of binary entries was generated previously by the system and was also made available, in the form of a file.

6.1 Efficiency

During this process, we have mainly focused on collecting data about two main subjects: efficiency of the different algorithms and platform self-assessment. Hence, concerning the first subject, multiple iterations of the algorithms were executed, under different settings,

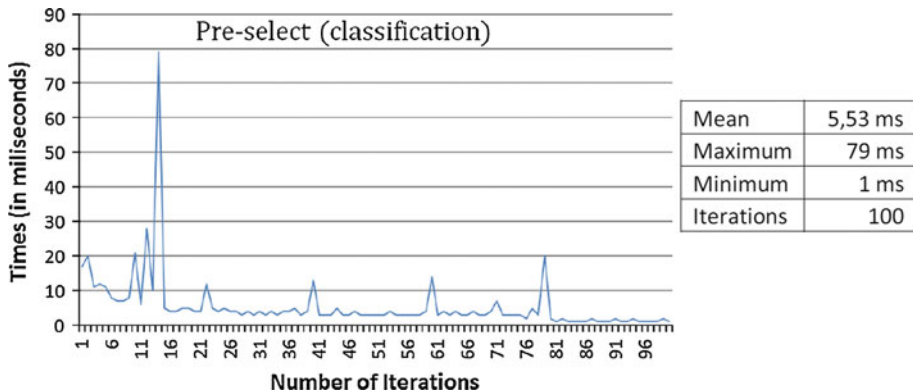


Fig. 14 Results for the classification algorithm of the pre-selection phase

and their times of execution measured. This will be useful to compare them in terms of their execution times. These tests thus evaluate the efficiency of each algorithm presented. We acknowledge that the most important factor in the efficiency of the algorithms is the size of the case base. As depicted in Fig. 13, the database contains at the moment 127 cases decided in court, provided by legal experts. The more cases there are, the higher the efficacy of the algorithms. However, efficiency will necessarily decrease. Nevertheless, the algorithms presented here have linear complexity. In that sense, the efficiency of these algorithms will decrease proportionally to the size of the database.

6.1.1 *Pre-select (classification)*

The data depicted in Fig. 14 concern a total of 100 iterations of the pre-selection algorithm that uses the association rules to retrieve cases according to their classes. To test it, in each iteration, a random case was provided to the algorithm. The algorithm thus had to analyze the cases, determine which rules applied in the cases and then pre-select, from the file containing the vectors of all the cases, the ones that belonged to the same category of the case provided. Analyzing Fig. 14, it is possible to conclude that this is a considerably fast process. This is essentially due to the fact that: (1) the representation of the data as vectors had been previously performed and (2) cases are already indexed according to the rules they comply with. Therefore, once all this information is made available, this is a highly efficient algorithm for information retrieval. Another factor contributing to this is that the data are stored locally. This is possible given the small size of the data when stored according to the Vector Space model, constituting another advantage of this method.

6.1.2 *Pre-select (template retrieval)*

To test this algorithm, several pre-selections were requested, using a random case. To perform the pre-selection, the algorithm must analyze the case and then interact with the database, requesting all the cases that match a given criteria. This algorithm was tested using a local instance of a database and a remote one. Looking at Fig. 15, the first 30 values correspond to the tests in which a remote instance was used, while the remaining correspond to the use of the local instance. There is a visible difference between the two scenarios. This difference is aggravated by the fact that: (1) the database is not a dedicated one; (2) being distributed,

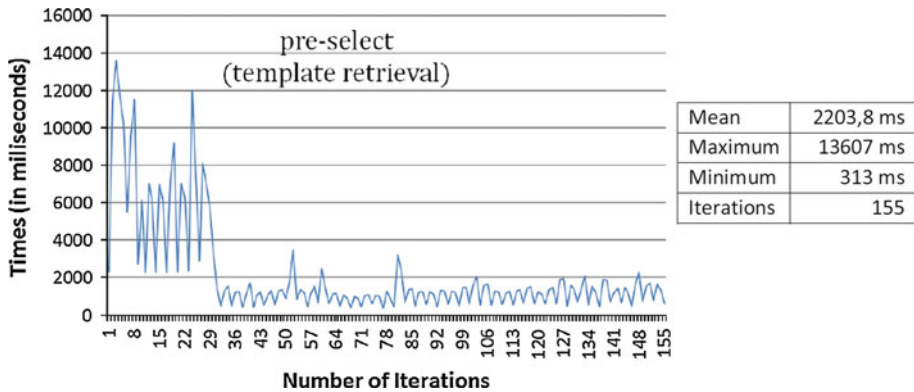


Fig. 15 Results for the template retrieval algorithm of the pre-selection phase

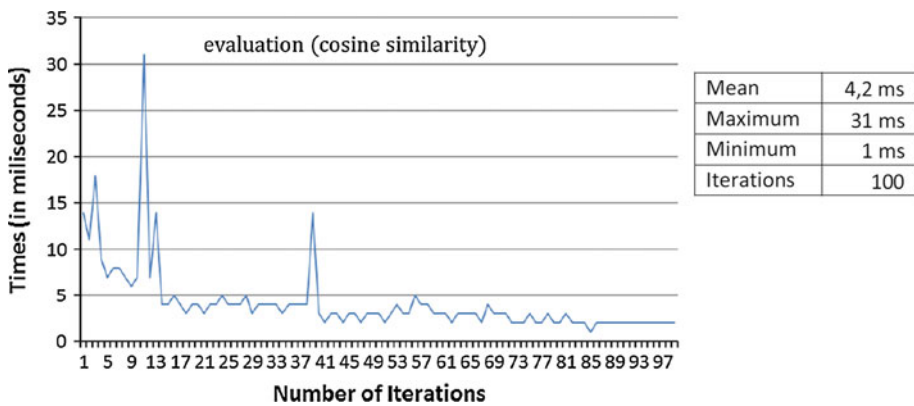


Fig. 16 Results for the cosine similarity algorithm of the evaluation phase

the system depends on external factors like the velocity of the internet connection; and (3) in order to satisfy the pre-selection rules, the algorithm may need to make several iterations (requests to the database). In this sense, one may improve the algorithm by choosing the best pre-selection rules, minimizing the number of requests to the database. This is addressed further ahead. When compared with the previous algorithm, this shows significantly higher times of execution, mainly due to the interaction with a remote database.

6.1.3 Evaluation (cosine similarity)

This algorithm uses the cosine similarity formula described before to determine the similarity between two cases. To test it, the algorithm was provided a case and a list of cases with the objective of determining the similarity of each case in the list with the isolated case. The results depicted in Fig. 16 shows that this is a relatively fast way of computing the similarity, mainly because it only deals with binary values. However, a major disadvantage of this algorithm is that it does not allow assigning weights to the different components of a case.

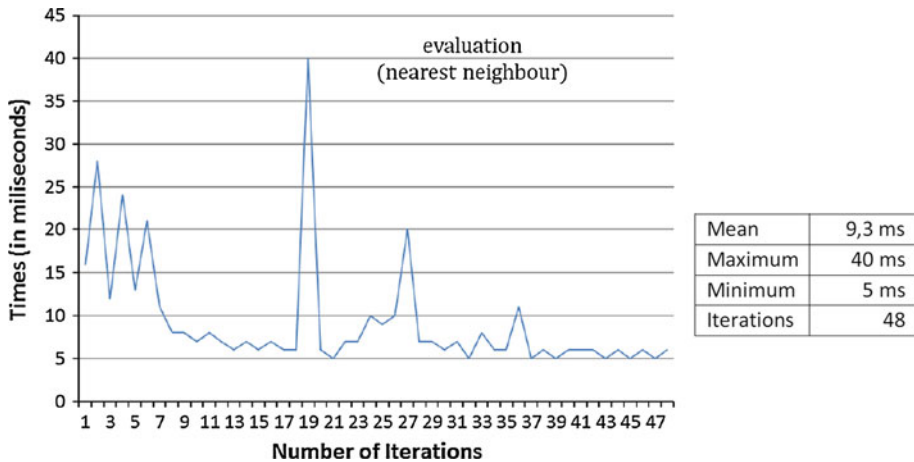


Fig. 17 Results for the nearest neighbor algorithm of the evaluation phase

6.1.4 Nearest neighbor

To test this algorithm, we have proceeded similarly to the previous one. Looking at the collected data depicted in Fig. 17, it is possible to conclude that the nearest neighbor algorithm performs slightly slower. As both algorithms have linear complexity, this poorer performance can be attributed to the fact that this algorithm deals with several types of variables (e.g., integers, strings, floating points) rather than binary ones. However, this algorithm allows for weights to be assigned to the different components of the similarity function, allowing an evaluation that might be closer to the one performed by a human expert.

6.1.5 Get complete info

In this test, the objective was to determine the efficiency of a request of all the information regarding a case, that is, giving a random case to the platform, how much time does it take to compile all the possible information for the user. This includes, as described before, pre-selecting and evaluating cases, computing the BATNA, WATNA, MLATNA, ZOPA as well as building the visual representation of the information. For this purpose, the algorithms described above were randomly selected to be used. Both local and remote requests were made, which is reflected in the execution times, similar to the previously presented results. Thus, the execution time of a complete information request depends mostly on which algorithms are selected. The resemblance of the graph depicted in Fig. 18 with the one depicted in Fig. 15 is also not a coincidence as, given the potentially high values of its execution times, it has a considerable influence on the overall performance.

6.2 Efficacy

More than the efficiency of the platform, one must also analyze its efficacy, that is, it is not enough to perform a given task quickly, but it must also be done correctly. Thus, regarding the second subject being tested, the platform keeps record of some key actions as well

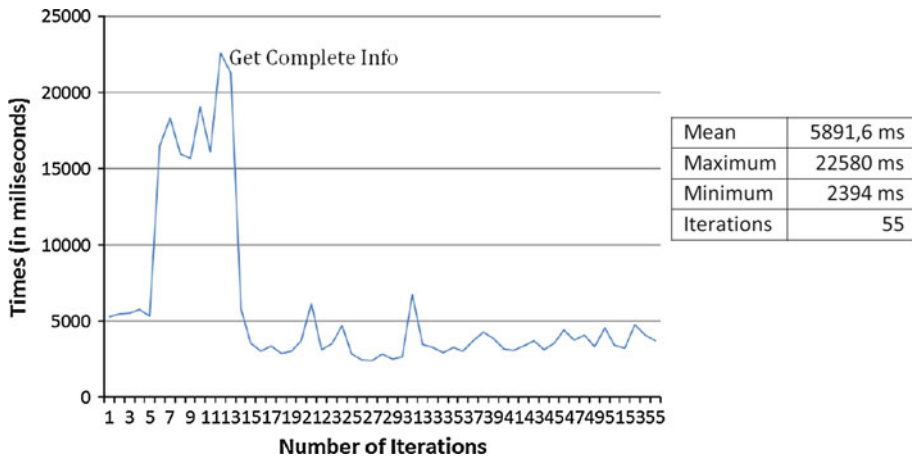


Fig. 18 Results for the process that compiles a complete set of information for a given dispute

as their results. This allows, in a first instance, to determine which actions fail the most. Then, together with a description of the possible problems, eventual causes and eventual solutions, the platform is able to provide advice about what parameters to change in order to potentially improve its efficiency. Our objective is that, in the long term, the platform applies autonomously these recommendations. However, this still requires further validation as it is a very sensitive topic.

A typical output of a self-assessment request is shown in Fig. 19. In this example, the platform is assessing the performance of the pre-select action. First of all, it provides information about the amount of times that the action failed or succeeded. Here, failing means that the pre-selection violates some rule (e.g., regarding the number of cases) and must be reformulated and re-run. On the other hand, succeeding means that the pre-selection respects all the rules. In the example, this action is failing in 66 % of the cases, corresponding to a total of 543 cases. Following, the platform points out the top reasons for failure as well as their frequency. In this example, the pre-selection fails mostly because too many cases are being pre-selected. Other minor reasons include not enough cases being pre-selected or reaching a state in which it is not possible to satisfy all rules. This happens when it is not possible to manage the pre-selection settings with enough precision or when the pre-selection rules are too strict.

Finally, given this, the platform points out several possible actions that might be used to address the described problems. In this case, three actions are suggested: (1) changing the rule that establishes the maximum amount of cases that should be pre-selected (this would actually decrease the number of errors but might not be good for who deals with all the cases later); (2) changing the initial search depth (this is more advisable as changing the initial search parameters might lead to a better result faster); and (3) changing the rule that establishes the minimum cases that should be selected.

7 Conclusions and lessons learned

In this paper, we have presented the underlying architecture of UMCourt as well as its main functionalities, after having described the legal domain being addressed. Throughout the

```

****Evaluator: Indexed cases in the database: 127
****Evaluator: Evaluation of System Performance:
****Evaluator: Action: preselect
****Evaluator: Successfull iterations: 33% (271)
****Evaluator: Failed iterations: 66% (543)
****Evaluator: Top reasons for failure
****Evaluator: Code: 1; Description: Too many cases were pre-selected ; 61% (501)
****Evaluator: Code: 0; Description: Not enough cases were pre-selected ; 2% (23)
****Evaluator: Code: 2; Description: Not possible to satisfy all rules ; 2% (19)
****Evaluator: Recommended actions:
****Evaluator: Change Max Cases rule (33%)
****Evaluator: Change initial search depth (33%)
****Evaluator: Change Min Cases rule (33%)

```

Fig. 19 The output of a self-assessment request for the pre-select action

work, we placed a special emphasis on the challenge of determining the possible outcomes for the cases, based on the observation of past cases, having as motivation the concept of legal precedent. This task is supported by a CBR algorithm which was also described. The main advantage of this algorithm lies, from our point of view, in the fact that it can be used as it was presented (a tool to compile useful information for the parties, the neutral or the platform itself) or it can be integrated in a higher-level process which includes the parties going into litigation or choosing another dispute resolution method, such as mediation or negotiation.

The CBR algorithm was detailed in terms of its several phases. During the description of each of these phases, we have also highlighted another key point of this work: Several techniques can be combined to implement the algorithm. This will lead to different possibilities and different results. It is thus up to who uses the algorithm to decide, depending on the type of information needed or the domain of application, which technique to use.

Another key point presented is the close integration between the CBR algorithm and the negotiation and mediation processes. In fact, it is this same algorithm that provides all the information that is necessary to inform the parties and for the platform to act as a neutral and guide these processes. Moreover, human neutrals can also make use of the information compiled in order to get knowledge about past similar cases so that more rational decisions can be made. And, as pointed out before, this can be useful in both common and civil law systems.

7.1 On the use of Case-Based Reasoning in the legal context

The use of CBR for conflict resolution is a natural approach in common law contexts, although its use in civil law systems also makes sense, as described above. Moreover, negotiators and mediators in general rely on their past experiences in order to take better decisions. These were in fact the reasons that supported our decision to follow a case-based approach. Traditionally, two main drawbacks are associated with CBR [1].

On the one hand, case-based approaches may suffer from inefficiency, mostly when the size of the database grows. Concerning this subject, it was our decision to develop several methods so that their performances could be tested in order to determine the more efficient ones. Moreover, we also acknowledge that conflict resolution is generally an asynchronous process, which takes place over several hours or even days. Concerning this subject, we believe to have succeeded as the methods described here have relatively low times of execution and the context of application does not mandatorily demand for higher performance.

On the other hand, the efficiency and efficacy of case-based approaches is also known to depend directly on the number and characteristics of the cases in the database. And this we acknowledge to be the more serious challenge to overcome. The number of cases, as it is evident, influences the performance of the algorithms but also influences its efficacy. Moreover, case-based approaches are highly domain dependent, that is, cases take place in a given legal context and are not easily adapted to other domains as the norms are different. In our experiment, we used a database whose cases focused mostly on employee's rights and employee's obligations (articles 129 and 128 of the Portuguese labor law, respectively), although some cases also addressed other articles. Our experiment consisted in conflict resolution scenarios in the labor law domain, set up by users, involving the issues addressed by the five articles (and respective numbers and items) addressed by the database (Articles 118, 117, 126, 128 and 129). The users were students and teachers of master courses on Law and on Informatics in our institution. We concluded that disputes involving employee's rights and obligations had many more solutions proposed by the platform than disputes involving other issues, diminishing the success rate of the resolution of the second ones.

In that sense, our main conclusion in this subject is that a pure case-based approach can be quite effective in generating solutions for a conflict resolution in scenarios in which there are enough past cases. In that sense, the domain of the database should be explicitly defined in order to define the domain of the conflict resolution platform. However, a hybrid approach can be used in which CBR is supported by other tools that can generate solutions when a case-based approach alone is not enough. On the one hand, it is possible to rely on parties themselves to generate solutions, although this approach tends to fail when parties are unable or unwilling to do it. On the other hand, other technological tools can be used for the generation of solutions. Namely, we are now working on genetic algorithms to create solutions in the scenarios in which a case-based approach is not enough to guarantee satisfactory results. The main advantage of this approach is that it depends only on the rules of each specific domain, which are needed to ensure the validity of the solutions. Given that, the algorithm will be able to generate a wide range of solutions from which the most relevant ones will be selected to complement the case-based approach.

7.2 On the use of negotiation/mediation for conflict resolution

The use of negotiation or mediation for conflict resolution is indeed one of the most effective ways of solving disputes out of court. Nevertheless, from our experiments, we learned that these processes, by themselves, may not be enough. One of the first issues we had to deal with was a consequence of the fact that most of our users from the informatics field had very little to no knowledge about Portuguese labor law or about conflict resolution at all, as many of the parties involved in conflict resolution do. In that sense, the resolution process often failed because of a clearly unrealistic view of their chances in the dispute. Moreover, these users generally had no idea about the possible outcomes for each side and could not really evaluate how good or bad a given solution was.

Our first conclusion about the use of these alternative methods is that negotiation or mediation alone are not enough. There is, more than anything else, the need for tools that can effectively inform the parties about the possibilities, so that they can take better and more informed decisions. Namely, we found it crucial for parties to know their best and worst possible scenarios, the most likely one as well as some past cases that can be used as learning examples, providing a notion of reality. Our decision was thus to implement this compilation of information before the actual process started, so that the parties could gain a realistic view

about their conflict. After implementing this process, we noticed that parties would converge more quickly to a solution that was realistic and in line with the solutions retrieved from the past similar cases. This increased the success rate of the conflict resolution process.

Another issue that was detected was that the success rate of the process depended on the attitude of the parties, that is, when our users were actively creating solutions and collaborating for the resolution of the dispute, the process was more likely to succeed than when one or more users were not or could not create solutions, limiting their actions to replying affirmatively or negatively to the solutions proposed. In that sense, we developed the two processes depicted in Sect. 5: one targeted at parties that are able and willing to create solutions for the resolution of the conflict and the other for parties that are not. Evidently, the challenge here is still the one of accurately determining the conflict style of each party (e.g., cooperative, collaborative, avoiding). In that sense, we are now developing an automated and non-invasive method that is able to classify the conflict style of a party based on their behavior during the process (e.g., Is the party proposing solutions? Are the solutions proposed selfish? Is the party simply answering positively or negatively?). Based on that, the platform will be able to determine the conflict style in real time and will be able to choose the best possible method as well as to adapt during the process.

In conclusion, the approaches presented in this work have as main innovations: the enrichment of the information retrieval process with several alternatives, allowing a user (or the platform) to select the one that most suits their needs and a close integration between the information retrieval methods and the posterior processes (e.g., negotiation, mediation, trial) in a single platform in a transparent way. This not only saves resources and time as it increases the organization of the information but also makes it easier for parties and practitioners to use it. All this will result in richer conflict resolution processes that, by being supported by information and decision support systems, will lead to better and more mutually satisfactory outcomes.

Acknowledgments This work is funded by National Funds through the FCT—Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) within project PEst-OE/EEI/UI0752/2011. The work of Davide Carneiro is also supported by a doctoral grant by FCT (SFRH/BD/64890/2009).

References

1. Aamodt A, Plaza E (1994) Case-based reasoning: foundational issues, methodological variations, and system approaches. *AI Commun* 7(1):39–59
2. Alterman R (1989) Panel discussion on case representation. In: Proceedings of the second workshop on case-based reasoning, Pensacola Beach
3. Andrade F, Barbieri D, Carneiro D, Novais P (2010) Artificial intelligence applications in ODR: the UMCourt project. In: Proceedings of the 17th annual northwest dispute resolution conference, University of Washington, USA
4. Ashley KD (2004) Case-based models of legal reasoning in a civil law context. In: International congress of comparative cultures and legal systems of the instituto de investigaciones jurídicas
5. Behrman BW, Davey SL (2001) Eyewitness identification in actual criminal cases: an archival analysis. *Law Hum Behav* 25(5):475–491
6. Bellifemine F, Poggi A, Rimassa G (2008) *Developing Multi-agent Systems with JADE*. Springer, Berlin
7. Bellucci E, Zeleznikow J (2001) Representations of decision-making support in negotiation. *J Decis Syst* 10(3–4):449–479
8. Bellucci E, Lodder A, Zeleznikow J (2004) Integrating artificial intelligence, argumentation and game theory to develop an online dispute resolution environment. In: Proceedings of the ICTAI-2004—16th IEEE international conference on tools with AI, pp 749–754
9. Bellucci E, Zeleznikow J (2006) Developing negotiation decision support systems that support mediators: a case study of the family winner system. *J Artif Intell Law* 13(2):233–271

10. Brown H, Marriott A (1999) ADR principles and practice. Sweet & Maxwell, London
11. Cáceres E (2008) EXPERTIUS: a Mexican judicial decision-support system in the field of family law. In: Francesconi EBE, Sartor G, Tiscornia D (eds) Legal knowledge and information systems. IOS Press, Amsterdam, pp 78–87
12. Carneiro D, Novais P, Andrade F, Zeleznikow J, Neves J (2009) The legal precedent in online dispute resolution. In: Legal Knowledge and Information Systems, Guido Governatori (eds) Proceedings of the Jurix 2009—the 22nd international conference on legal knowledge and information systems, IOS press, pp 47–52. ISBN 978-1-60750-082-7
13. Carneiro D, Novais P, Costa R, Neves J (2010) Enhancing the role of multi-agent systems in the development of intelligent environments. Advances in intelligent and soft computing, vol 71. Springer, pp 123–130. ISBN 978-3-642-12432-7
14. De Vries BR, Leenes R, Zeleznikow J (2005) Fundamentals of providing negotiation support online: the need for developing BATNAs. In: Proceedings of the second international ODR workshop, Wolf Legal Publishers, Tilburg, pp 59–67
15. Fernandes AM (2005) Direito de trabalho, Almedina, (in Portuguese)
16. Foundation for Intelligent physical Agents: FIPA ACL message structure specification (2002) available at <http://www.fipa.org/specs/fipa00061>, accessed in August 2012
17. Fisher R, Ury W (1991) Getting to YES: negotiating agreement without giving. In: Penguin books, ISBN 978-0140157352
18. Follett MP (1940) Constructive conflict. In: Metcalf HC, Urwick IL (eds) Dynamic administration: the collected papers of Mary Parker Follett. Harper, New York, pp 30–49
19. Foxall G (2004) What judges maximize: towards an economic psychology of the judicial utility function. *Liverp Law Rev* 25:177–194
20. Fudenberg DA, Tirole J (1983) Game theory, Chapter 1, Section 2.4. MIT Press
21. Goldberg SB, Sander FE, Rogers N et al (2003) Dispute resolution: negotiation, mediation and other processes. Aspen Publishers, New York
22. Goodman JW (2003) The pros and cons of online dispute resolution: an assessment of cyber-mediation websites. *Duke Law Technol Rev*, pp 1–16
23. Hasan M, Salem S, Zaki M (2010) SimClus: an effective algorithm for clustering with a lower bound on similarity. *Knowl Inf Syst*. doi:10.1007/s10115-010-0360-6
24. INCM (2009) Código do Trabalho (in Portuguese)
25. Jennings N, Faratin P, Lomuscio A et al (2001) Automated negotiation: prospects, methods and challenges. *Group Decis Negot* 10(2):199–215
26. Katsch E, Rifkin J (2001) Online dispute resolution—resolving conflicts in cyberspace. Jossey-Bass Wiley Company, San Francisco
27. Klaming L, Van Veenen J, Leenes R (2008) I want the opposite of what you want: summary of a study on the reduction of fixed-pie perceptions in online negotiations. In: Expanding the horizons of ODR, Proceedings of the 5th international workshop on online dispute resolution, Firenze, Italy, pp 84–94
28. Landes WM, Posner RA (1976) Legal precedent: a theoretical and empirical analysis. *J Law Econ* 19:249–307
29. Lewicki R, Saunders D, Minton J (1999) Zone of potential agreement. *Negotiation*, 3rd edn. Irwin-McGraw Hill, Burr Ridge
30. Lynda T, Mohand B, Mariam D (2010) Evaluation of contextual information retrieval effectiveness: overview of issues and research. *Knowl Inf Syst* 24(1):1–34
31. Martinez PR, Monteiro L, Vasconcelos J, et al. (2008) Código do trabalho, Almedina, (in Portuguese)
32. Notini J (2005) Effective alternatives analysis in mediation: “BATNA/WATNA” analysis demystified. Available at <http://www.mediate.com/articles/notini1.cfm>. Last accessed August 2012
33. Peruginelli G, Chiti G (2002) Artificial Intelligence in alternative dispute resolution. In: Proceedings of the workshop on the law of electronic agents—LEA
34. Popple J (1996) A pragmatic legal expert system. Ashgate, Dartmouth
35. Posner RA (1993) What do judges and justices maximize. *Supreme Court Econ Rev* 3:1–41
36. Raiffa H (1982) Art and science of negotiation. Harvard University Press, Cambridge
37. Riesbeck C, Bain W (1987) A methodology for implementing case-based reasoning systems, Lockheed
38. Slade S (1991) Case-based reasoning: a research paradigm. *AI Mag* 12:42–55
39. Steinbach M, Tan PN, Kumar V (eds) (2005) Introduction to data mining. Pearson Addison Wesley, Boston
40. Tang M, Zhou Y, Li J, Wang W, Cui P, Hou Y, Luo Z, Li J, Lei F, Yan B (2010) Exploring the wild birds’ migration data for the disease spread study of H5N1: a clustering and association approach. *Knowl Inf Syst* 1–25. doi:10.1007/s10115-010-0308-x

41. Thiessen EM (1993) ICANS: an interactive computer-assisted multi-party negotiation support system. In: PhD dissertation, School of Civil and Environmental Engineering, Cornell University, Ithaca
42. Thiessen EM, McMahon JP (2000) Beyond win-win in cyberspace. *Ohio State J Disput Resolut* 15:643
43. Thiessen EM, Fraser K (2003) Mobile ODR with smartSettle. In: *Proceedings of the UNECE Forum on ODR*
44. Thomas K, Kilmann R (1974) Conflict and conflict management. Available at <http://www.kilmann.com/conflict.html>, accessed in August 2012
45. Walton PRE, McKersie RB (1965) *A behavioral theory of labor negotiations*. McGraw-Hill, New York
46. Waterman DA, Peterson M (1980) Rule-based models of legal expertise. In: *Proceedings of the first national conference on AI*, Stanford University
47. Watson I, Marir F (1994) Case-based reasoning: a review. *Knowl Eng Rev* 9:327–354
48. Wooldridge M, Jennings NR (1995) Intelligent agents: theory and practice. *Knowl Eng Rev* 10:115–152
49. Wooldridge M, Jennings NR, Kinny D (2000) The Gaia methodology for agent-oriented analysis and design. *Auton Agents Multi-Agent Syst* 3(3):285–312
50. Wooldridge M (2002) *An introduction to multiagent systems*. Wiley, London
51. Wu X, Kumar V, Quinlan JR et al (2007) Top 10 algorithms in data mining. *Knowl Inf Syst* 14(2008)1:1–37
52. Zeleznikow J, Stranieri A (1995) The split-up system: integrating neural networks and rule-based reasoning in the legal domain. In: *Proceedings of the 5th international conference on artificial intelligence and law*, ACM, College Park, pp 185–194
53. Zeleznikow J, Bellucci E (2003) Family-Winner: integrating game theory and heuristics to provide negotiation support. In: *Proceedings of sixteenth international conference on legal knowledge based, system*, pp 21–30
54. Zeleznikow J, Bellucci E (2004) Building negotiation decision support systems by integrating game theory and heuristics. In: *Proceedings of the IFIP international conference on decision support systems*. Monash University, Melbourne, Australia
55. Zeleznikow J, Bellucci E, Schild UJ, Mackenzie G (2007) Bargaining in the shadow of the law—using utility functions to support legal negotiation. In: *Proceedings of the 11th international conference on artificial intelligence and law*, ACM, pp 237–246
56. Zweigert K, Kötz H (1998) *An introduction to comparative law*, 3rd edn. Clarendon Press, Oxford

Author Biographies



Davide Carneiro is a researcher at the CCTC, in the Department of Informatics, University of Minho, Braga, Portugal. He is currently enrolled in a joint Doctoral Programme in Computer Science of three top Portuguese Universities. He develops scientific research in the field of Artificial Intelligence, joining Multi-Agent Systems, Genetic Programming and Case-Based Reasoning applied to the Legal field and Ambient Intelligence. He has participated in several research projects in the fields of Ambient Intelligence and Online Dispute Resolution. He is the co-author of several publications in his field of interest, including book chapters, international journals and conference and workshop papers. He has been a member of Program and Local Organization Committees of conferences in his field in the last years. He has been awarded the TLeIA08—a National Award for Artificial Intelligence projects by the Portuguese Artificial Intelligence Association and in 2009; he has been awarded an Academic Merit Scholarship by the Portuguese Government.



Paulo Novais is a professor of Computer Sciences at the Department of Informatics, in the University of Minho, Braga, Portugal, and researcher at the CCTC. He received a Ph.D. in Computer Sciences from the same university in 2003 and his Habilitation in Computer Sciences in 2011. He develops scientific research in the field of Artificial Intelligence, namely Knowledge Representation and Reasoning, Machine Learning and Multi-Agent Systems, with applications to the areas of Law and Ambient Intelligence. He has led several research projects sponsored by Portuguese and European public and private Institutions and has supervised several Ph.D. and MSc students. He is the co-author of over 100 book chapters, journal papers, conference and workshop papers and books. He has been chair of the Organising or Scientific Committee of several international conferences in the fields of Artificial Intelligence and Ambient Intelligence. He is vice president of APPIA, the Portuguese Association for Artificial Intelligence.



Francisco Andrade is a Professor of Law at the Law School in the University of Minho. He received his Ph.D. in Juridical Privatistic Sciences from the same University in 2009 and Msc in Electronic Information Management from Information Studies Department, University of Sheffield, UK. He develops scientific research in the field of Law and Informatics, specially concerning electronic contracting, online dispute resolution and data protection. He is the Director of the Master in Law and Informatics at Universidade do Minho. He is the Director of the Directive Committee of the courses on “Privacy and Data Protection” and “Electronic Commerce and Electronic Contracting” at Universidade do Minho. He is the co-author of several book chapters, journal papers, papers in International conferences and workshops. He is also a member of APPIA, the Portuguese Association for Artificial Intelligence and of APDI—Portuguese Association for Intellectual Law.



John Zeleznikow Director of Laboratory for Decision Support and Dispute Management, at the Victoria University, is an internationally recognized expert, with over 30 years of research experience including the last 15 years working on negotiation decision support and machine learning in law. His publications include three books, 50 refereed journal articles (including the Harvard Negotiation Law Review) and 150 refereed conference articles. He has received over \$7million in competitive research grants, significant international collaboration and industry funding, supervised 11 Ph.D. students to completion and conducted over twenty radio and television appearances. He has worked in the university sector for 38 years, with extended experience in USA, France, Netherlands, United Kingdom, Poland and Israel. He has also worked with Victoria Legal Aid, Relationships Australia, the New Zealand government and the Victorian Institute of Sport.



José Neves is professor of Computer Science at Minho University, Portugal, since 1983. Jose Neves is the Deputy Director of the Division for Artificial Intelligence (AI). He received his Ph.D. in Computer Science from Heriot Watt University, Scotland, in 1983. His current research interests relate to the areas of knowledge representation and reasoning and evolutionary intelligence, aiming to construct dynamic virtual worlds of complex symbolic entities that compete against one another in which fitness is judged by one criterion alone, intelligence, here measured in terms of a process of quantification of the quality of their knowledge, leading to the development of new procedures and other metaheuristics and their application in complex tasks of optimization and model inference in distinct areas, namely in the healthcare arena (e.g., machine learning in an intensive care unit environment, ambient intelligence).