Universidade do Minho
Escola de Engenharia

Davide Rua Carneiro

An agent-based architecture for online dispute resolution services

The MAP Doctoral Program in Computer Science of the Universities of Minho, Aveiro and Porto

universidade de aveiro

Universidade do Minho

U.PORTO

July 2013

An agent-based architecture for online dispute resolution services

Davide Rua Carneiro

UMinho | 2013

Universidade do Minho
Escola de Engenharia

Davide Rua Carneiro

An agent-based architecture for online dispute resolution services

The MAP Doctoral Program in Computer Science of the Universities of Minho, Aveiro and Porto

universidade de aveiro

Universidade do Minho

U.PORTO

A thesis submitted at the University of Minho for the degree of Doctor of Philosophy in Informatics (PhD)
under the supervision of
Paulo Jorge Freitas de Oliveira Novais
and
José Carlos Ferreira Maia Neves

July 2013

DECLARATION

Nome: Davide Rua Carneiro

Endereço Electrónico: dcarneiro@di.uminho.pt
Telefone: 00351 934 335 822

Número do Bilhete de Identidade: 12964271

Título da Tese: An agent-based architecture for online dispute resolution services

Orientadores:
Paulo Jorge Freitas Oliveira Novais
José Carlos Ferreira Maia Neves

Ano de Conclusão: 2013

Ramo de Conhecimento: Informática

Universidade do Minho, Braga, July 31, 2013

Davide Carneiro

Davide Carneiro: *An agent-based architecture for online dispute resolution services,* The MAP Doctoral Program in Computer Science of the Universities of Minho, Aveiro and Porto, PhilosophiæDoctor (PhD), © September 2009 - September 2013

This work was partly funded by National Funds through the FCT - Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) within projects PEst-OE/EEI/UI0752/2011 and TIARAC - Telematics and Artificial Intelligence in Alternative Conflict Resolution (PTDC/JUR/71354/2006).

The work of Davide Carneiro was supported by a doctoral grant by FCT (SFRH/BD/64890/2009).

SUPERVISORS
Paulo Novais
José Neves

LOCATION
Braga

TIME FRAME
September 2009 - September 2013

Science is the attempt
to make the chaotic diversity of our sense-experience
correspond to a logically uniform system of thought.

— Albert Einstein

ABSTRACT

Conflicts are a natural consequence of our daily social interactions and should be regarded as opportunities to improve some aspect, condition or flaw. In order for conflicts to have positive outcomes, not only from an economical view but also concerning the maintenance of good interpersonal-relationships, tools are needed that can explain to the parties the inner mechanisms of the conflict resolution process, its restrictions and its rules. Only well-informed parties can take good and realistic decisions and better understand the others'.

In this thesis, the most important aspects in a conflict resolution process concerning the aforementioned are identified, with the objective of designing a tool that can effectively support the parties from the beginning to the end of the conflict. The resulting tool, UMCourt , has as main objective to support decisions by providing the right information in the right moment to the right stakeholders.

Specifically, several problems are addressed that include the definition of a suitable agent-based architecture, the building of important knowledge and the support in the negotiation process, either by generating solutions or by analysing the behaviour of the parties.

This last issue is addressed in more detail in this thesis. In fact, the most serious drawback that was identified in current conflict resolution methods is their complete disregard for contextual and subjective information about the parties: the trend has been, for many years, to focus on the objective aspects of the conflict.

The main contribution of this thesis is a vision on conflict resolution that goes the other way around: asides from objective information, decisions should also be based on contextual features such as our level of stress, body language, attitudes or our conflict handling style. We rely on this information on a daily basis to communicate efficiently. It results only logical that it should be included in conflict resolution methods that rely so heavily on communication.

The approach put forward relies on the analysis of the individual's behaviour in order to infer such context information. Disputant parties and, in particular, mediators and negotiators, can better understand the state of the participants and take better decisions (e.g. make a pause, understand how a party is affected by an issue). This is particularly important when online dispute resolution methods that rely on cold and impersonal communication technologies (often constituting a barrier to efficient communication) are used.

RESUMO

Os conflitos são uma consequência natural das nossas interações sociais e devem ser vistos como oportunidades para melhorar determinados aspetos, condições ou mesmo falhas. Para que tenham resultados positivos, não só do ponto de vista económico mas também do ponto de vista das relações interpessoais, são necessárias ferramentas que expliquem às partes as particularidades do processo, as suas restrições e as suas regras. Apenas partes bem informadas podem tomar decisões realísticas e melhor entender as decisões dos restantes.

Nesta tese, os aspetos mais importantes num processo de resolução de conflitos são identificados, com o objetivo de definir uma ferramenta que possa, efetivamente, suportar as partes do início ao fim do conflito. A ferramenta resultante, designada UMCourt , tem como principal objetivo suportar decisões fornecendo a informação certa no momento certo às entidades certas. Especificamente, vários problemas são atacados que incluem a definição de uma arquitetura de software adequada, a construção de conhecimento e o suporte à negociação, quer através da geração de soluções quer através da análise comportamental das partes.

Este último tópico é tratado em mais detalhe nesta tese. De facto, a limitação mais significante que foi identificada nos atuais métodos de resolução de conflitos é a negligência da importância dos fatores contextuais e da informação subjetiva acerca das partes: a tendência tem sido no sentido de se focarem apenas nos aspetos objetivos.

A principal contribuição desta tese é a de uma visão do processo de resolução de conflitos que aponta no sentido oposto: para além da informação objetiva, as decisões devem também ser baseadas em aspetos contextuais tais como o nível de stress, a linguagem corporal ou o estilo de lidar com o conflito. Enquanto indivíduos, baseamo-nos nestes aspetos diariamente para comunicar de forma eficiente. É portanto lógico que tal informação seja incluída em métodos de resolução de conflitos que se baseiam de forma tão clara na comunicação.

A abordagem proposta baseia-se na análise comportamental de cada indivíduo para aquisiçao da informação de contexto. As partes em conflito e, em particular, os mediadores e negociadores, podem entender melhor o estado de todos os participantes e tomar melhores decisões (e.g. fazer uma pausa, perceber como uma parte é afetada por uma questão). Isto é especialmente importante quando são usados métodos de resolução de conflitos em linha baseados no uso de tecnologias de comunicação frias e impessoais, que geralmente configuram elas próprias um obstáculo à eficiência da comunicação.

ACKNOWLEDGEMENTS

Obtaining a PhD degree is generally seen as a solitary process, in which the candidate spends several years studying and specializing on some state of the art topic, focusing on his own personal achievements. A different approach was followed on this PhD: one focused not only on personal achievements but also on establishing connecting points with the surroundings and on sharing lessons learned with the peers. Thus being, many people had an influence on the work described in this thesis. Here I would like to acknowledge those who had a particularly significant impact.

First of all I must mention Paulo Novais and José Neves, who more than supervisors became friends. They provided the constant support and conditions needed for reaching this stage of my academic life. Paulo Novais, thank you for your constant support and motivation, for once again pushing me to go further. José Neves, thank you for your experience, ideas, conversations and good moments.

I am also especially grateful to three other friends who in some way alleviated or supported my work. Francisco Andrade, who guided me through the legal meanders, and facilitated my first steps in the intersection between computers and laws. John Zeleznikow that, despite being on the other side of the planet, always had an insightful opinion on our work. Marta Poblet, for her visits to our institution and fruitful conversations.

I must also mention my colleagues and friends of the ISLab, who make up a good environment to work in, namely Ângelo Costa, Tiago Oliveira and Marco Gomes, as well as many other friends who, I'm sure, will understand the fact that their names do not appear here.

I am also grateful to José Carlos Castillo and António Fernández-Caballero who welcomed me in Albacete and showed me not only the good research that is done there but also the good sights, culture and gastronomy. Thank you for the good moments.

Finally, I would like to dedicate a few lines to my family, who has always been supportive and encouraging, giving me the conditions to go further. Thank you Manuel, Helena and Cláudia.

Elsa, you are the last in the acknowledgements but the first in my heart. Thank you for your constant support, your critiques, your ideas and the cheering up smile on the toughest moments.

With all my gratitude,
Davide Carneiro

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LISTINGS

ACRONYMS

ADR	Alternative Dispute Resolution
AI	Artificial Intelligence
BATNA	Best Alternative to a Negotiated Agreement
CBR	Case-based Reasoning
DSS	Decision Support System
IE	Intelligent Environment
IM	Instant Messaging
IT	Information Technologies
KR	Knowledge Retrieval
MAS	Multi-agent System
MDA	Model Driven Architecture
MLATNA	Most Likely Alternative to a Negotiated Agreement
ODR	Online Dispute Resolution
OWL	Web Ontology Language
RbS	Rule-based Systems
SIEJ	Sistema de Informaçao das Estatísticas da Justiça
TIARAC	Telematics and Artificial Intelligence in Alternative Conflict Resolution
UML	Unified Modelling Language
WATNA	Worst Alternative to a Negotiated Agreement
XMI	XML Metadata Interchange
ZOPA	Zone of Possible Agreement

Part I

A NOT SO SHORT PREAMBLE

This thesis is organized into three main parts. The first one establishes the foundations that frame the research work presented in the following ones. It starts by analysing the current state of legal systems, describing some alternative methods for conflict resolution, both traditional and technology-based. However, even these alternative methods encompass disadvantages, which must be considered. After analysing them, the manuscript continues by devising how Artificial Intelligence techniques could possibly be used to alleviate some of the problems identified. Part I ends by defining the research questions addressed during the PhD work, their motivation and their main goals.

1 INTRODUCTION

> The artist is nothing without the gift,
> but the gift is nothing without work.
>
> — Emile Zola

The art of solving conflicts is quite a classic and old one, as old as conflicts themselves. Conflicts are natural and emerge as a consequence of our complex society, in which individuals focus on the maximization of their own gain, sometimes disregarding the others'. The concept of conflict and its resolution has traditionally been addressed by Social Science. However, in the last decades this field also started to be approached by Information Science. The intersection of these two fields is of large interest as it combines all the established theory about conflict resolution with computational power, new methodologies and support tools.

Despite the negative connotation generally associated to conflicts, these should not be regarded as necessarily negative. The fact is that conflicts are universal and, sooner or later, will have to be faced. In that sense, there is a need for conflict resolution processes that can minimise, manage and hopefully settle them. The ideal conflict resolution process is one in which the parties are better at the end of the course of action than they were at the beginning. Unfortunately, not all disputes take or may take this path.

Although the term conflict/dispute resolution is from time to time used to describe Alternative Dispute Resolution (ADR) (Brown and Marriott, 2012) methods, in a broad mode it marks out a process in which two or more parties engage in order to situate their differences. In that sense, dispute resolution can be divided into two main tendencies: Judicial Dispute Resolution and Alternative Dispute Resolution.

The most common form of Judicial Dispute Resolution is litigation, opposing a plaintiff and a defendant. The legal system has the coercive power to enforce an outcome. This means that at the end of the process, the parties are bound to the decision of the court, although in some cases parties may appeal to other instances. Outcomes are decided either by a judge, a jury or a combination of both, taking into account the facts presented by the parties and the application of The Law. All these processes are very formal and are defined by rules established by a legislature.

Alternatively, extra-judicial dispute resolution methods are being adopted as a first attempt to solve a case before advancing to litigation. These methods include arbitration, collaborative law, mediation, among others. These can also be used independently by parties, with the assistance of institutions that do not make part of the judicial system. The steady development of this trend on dispute resolution is due to a perception of greater flexibility, lower costs and faster outcome when compared to litigation. Parties will also often refer as advantages the increased privacy and fewer formalities.

This chapter contextualizes the remaining ones and frames the scope of this thesis. A brief overview of the current state of the legal systems is performed. Afterwards, the concept of "conflict" is defined and common conflict handling styles are described. Finally, alternative traditional and technology-supported methods to conflict resolution are introduced.

1.1 CURRENT STATE OF JUDICIAL SYSTEMS

Litigation is generally seen as an inefficient process, in which the most sound disadvantages are the high costs and long lasting for the way out of cases. Courts are currently clogged due to the amount of cases, most of which of relatively low values. In this thesis, the Portuguese judicial system is analysed in detail and can be used as an example. Several performance indicators of the Portuguese judicial system are put forward in this section, compiled from the Portuguese legal statistics consultation system available online - SIEJ[1].

If we analyse the average duration of the processes in courts of first instance between 2001 and 2010, it is possible to conclude that there was a significant increase, being the larger one in Civil Justice, from 20 months in 2001 to 29 months in 2010. This trend was continuous until 2007 (33 months), after which it slowly inverted. In average, considering the several domains under analysis, each process takes 10 to 15 months until resolution. Figure 1 depicts the data considered.

The number of pending processes awaiting litigation is increasing.

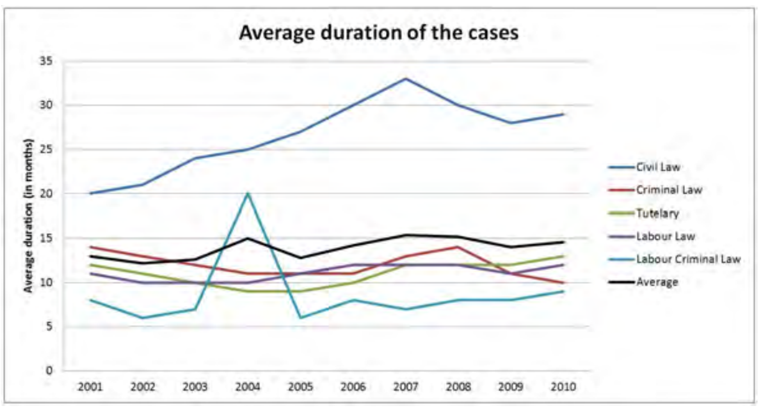

Figure 1: Average duration of the processes in Portuguese courts of first instance between 2001 and 2010.

1 The web-site of the SIEJ is available online at http://www.siej.dgpj.mj.pt <Acessed on November 2012>

This relatively large duration of processes has as main consequence an increase in the number of pending processes. In fact, the congestion rate for the Portuguese judicial courts in 2009 was of 253.3%, i.e. for each process that was concluded, 235.3 were pending[2]. According to the same source, there are currently more than 1.7 million pending processes. This represents an increase of 2.2% compared to the numbers from 2011, in which there were 1.666.348 pending processes. The trend is not recent: in 2007 and 2008 the number of pending processes grew from 1.541.239 to 1.504.553. Contradictorily enough, Portugal was in 2008 the third country in the European Union[3] with the best judge/attorney-to-population ratio, only bested by Poland and Germany (de L'Europe, 2010). According to the same study, the current pending cases would take 430 days to conclude, assuming no new processes.

One could argue that processes are nowadays more complex and include much more information or that defendants have more rights and attorneys have a larger range of action. However, the truth is that nowadays the most trivial act of our life is likely to end in a court, independently of the issue or amount of money involved. Moreover, courtrooms are not the most efficient environment for conflict resolution: they are generally highly competitive milieus, in which parties and their representatives will blindly pursue the maximization of their own personal profit, without any regards for the other party's interests. This constitutes an obstacle for the achievement of a mutually satisfactory outcome, increasing the dissatisfaction of the parties and consequently the number of appeals, contributing to a slower and more inefficient judicial system.

1.2 CONFLICTS AND CONFLICT HANDLING STYLES

Recently, Harvard evolutionary biologist Professor David Haig, brought forward a controversial theory defending that conflicts in our life start while we are still in the womb (Haig, 2010). According to the author, these parent-offspring or offspring-offspring conflicts may result in a wide range of behavioural and psychological disorders, including Tourette syndrome, autism or depression. Conflicts are present in our lives since our conception to our passing. We live with and solve conflicts on a daily basis, frequently in a natural and unconscious way.

Generally, a conflict or a dispute can be seen as a struggle between interdependent parties. In order for a conflict to be felt, parties must be aware of incompatible goals, scarce resources, and interference

2 Source: Portuguese Directorate-General for Justice Policy - Ministry of Justice, PORDATA.

3 The following countries were considered in this study: Germany, Austria, Belgium, Denmark, Scotland, Spain, Finland, France, England and Wales, Italy, Norway, The Netherlands, Poland, Portugal, Sweden and Switzerland.

from others in achieving their goals (Burton, 1990). Moreover, conflicts can be of different types, including conflicts of interests or values, structural conflicts, conflicts of data or relationship conflicts.

Each individual has a personal style of facing a conflict scenario.

Our individual style of dealing with a conflict must be seen as having a preponderant role in the outcome of a conflict resolution process, especially on those in which parties interact directly (e.g. negotiation, mediation). Ultimately, it is acceptable to state that the outcome of the process will largely depend on the conflict resolution style of each party and on the interaction of the styles of the parties. Different approaches can be followed to formalize the way that we respond to conflicts. A well-known definition was presented by Kenneth Thomas and Ralph Kilmann, who encoded the way that we react under a conflict into five different modes (Kilmann and Thomas, 1977). To define these modes, the authors take into consideration the individual's assertiveness, which denotes how much a party tries to satisfy his own interests, and the cooperativeness, which denotes to which extent the party is willing to satisfy the other's interests. Figure 2 depicts the distribution of the conflict handling styles in term's of the individual's cooperativeness and assertiveness. The five different conflict resolution styles defined are as follows:

- Competing – A party that shows this uncooperative style aims at maximizing his own gain, with a consequent minimization of the other's. Usually, a competing individual will use his ability to argue, his rank, his social status or whatever advantageous position that he can have to show dominance over the other party. This is thus a power-oriented style;

- Accommodating – An accommodating party will show a behaviour that can be classified as the opposite of a competing one in the sense that he will be cooperative. It may happen that an accommodating party will even neglect his own gain, thus maximizing the one of the others, in order to achieve a solution. Thus, it may be said that there is an element of self-sacrifice. Generally, such a party will tend to show generosity or charity, will be understanding and will easily obey other's orders or desires even if they represent a drawback;

- Avoiding – An individual that shows an avoiding behaviour is most likely not dealing with the conflict as he usually satisfies neither his own interests nor those of the other party. Common behaviours in this conflict style include diplomatically sidestepping or postponing some issue or even withdrawing from threatening or unpleasant situations;

- Collaborating – On the opposite side of avoiding is the collaborative behaviour. This is a cooperative style in which the party shows the willingness to work with the other party in order to find solutions that can be interesting for both. This implies that the party is interested in finding what the fears and desires of the other are and might even try to explore a disagreement in order to learn from other's insights;

- Compromising – A compromising party will generally try to find a fast and satisfactory solution that can be interesting for both parties. This conflict style can be seen as an intermediary one between the competing and the accommodating. A compromising party is generally willing to split the differences between two positions, to exchange some concessions or to seek middle-ground solutions.

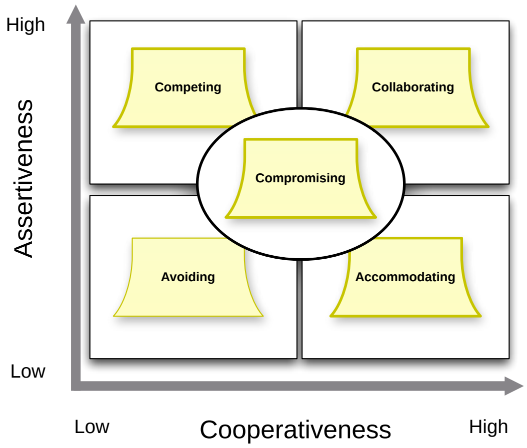

Figure 2: Distribution of the conflict handling styles defined by Thomas and Kilmann in terms of the individual's assertiveness and cooperativeness.

It is known that each individual is able to use several or even all of the conflict styles, depending on factors such as past experiences, temperament or present situation. Therefore, it is not possible to characterize one individual as having a single conflict-handling style. Still, people tend to use some styles more than others, generally due to the personality traits. The knowledge about the conflict-handling style may be of interest, either from the point of view of a human mediator or even from the point of view of a conflict resolution platform as it is possible to predict, to some extent, the evolution of the conflict resolution process according to the personal conflict handling styles.

1.3 ALTERNATIVE DISPUTE RESOLUTION

If we are looking for understanding the need for alternative methods of dispute resolution, we might as well look at a sentence by Abraham Lincoln which is still quite up to date:

Discourage litigation, persuade your neighbour

to compromise where you can. Point out to them
how the nominal winner is often the loser...
in expenses and waste of time

— Abraham Lincoln

Alternative Dispute Resolution refers to mechanisms that aim to solve disputes without recurring to the traditional judicial process

Alternative Dispute Resolution refers to mechanisms that aim to solve disputes without recurring to the traditional judicial process, i.e. litigation in courts. Online Dispute Resolution refers to the use of these mechanisms in a technological context, either supported by technology or under a virtual computational environment. Historically, alternative methods have faced some resistance but they have eventually become used by both the legal system and the parties involved as the first stage to solve a dispute. There are even countries in which parties are encouraged or required to try some kind of alternative method before advancing into court.

Some main alternative forms of conflict resolution can be identified. Negotiation (Raiffa, 1985) is a collaborative and informal process by means of which parties communicate and, without any external influence, try to achieve an outcome that can satisfy both. Negotiation is widely used in the most different fields such as legal proceedings, divorces, parental disputes or hostage situations, just to name a few. There are many ways of organizing the several negotiation techniques. From the perspective of (Walton, 1991), negotiation can be classified as a distributive or integrative approach.

In the former one, the items in dispute are seen as something that can be divided and distributed by the parties in an attempt to maximize their satisfaction. In the integrative one, the problem is expected to have more solutions than the ones visible at first sight. In these types of problems parties try to bring to the table as much interests as possible so that there are more valuable items with which to negotiate. This type of negotiation is also known as interest-based as the parties try to combine their interests and find points in which the interests of both are satisfied. By doing so, more mutually satisfactory outcomes are achieved. This makes integrative negotiation more desirable than distributive.

Another form of conflict resolution is mediation (Brown and Marriott, 2012), in which the parties in dispute are guided by a 3rd neutral and independent entity who tries to guide the process to an outcome that may satisfy both parties. In this approach, like in negotiation, both parties decide about the outcome instead of it being imposed by the non-aligned one, but with its added assistance. The non-aligned is chosen by the parties and has no authority for deciding on the outcome of the dispute but only for guiding and assisting them throughout it. This should be done by maintaining the parties focused on the subject of the dispute and by facilitating all the interaction and communication between them. Mediators are hence essential as their skills and aptitudes may represent the success or failure of the dispute resolution process.

Arbitration (Bennet, 2002) must also be mentioned: a method in which the two parties also use the help of a 3rd independent and neutral entity for solving a dispute but, contrary to mediation, this entity has no active role on helping the parties throughout all the process. In this approach the non-aligned party, the arbitrator, simply hears the parties and, based on the facts presented, takes a decision without influencing the parties during their presentations. Traditionally, the outcome of an arbitration process is binding, i.e., there is a final enforceable award that the parties will respect. However, arbitration can also be non-binding.

Finally, another alternative process is conciliation, which is run by a conciliator that meets the parties separately and aims at the resolution of their differences. The conciliator should be an expert with skills that allow him/her to lower tensions between the parties as a first stage for finding common ground. Then, he/she should communicate effectively with the parties in order to understand all the issues that generated the dispute. After doing this, the conciliator should be able to provide technical assistance to the parties as needed, so that they may have access to all the important information in order to take better decisions. Finally, the conciliator should explore all the potential paths for solutions and, at the end, achieve a negotiated settlement.

1.4 ONLINE DISPUTE RESOLUTION

On-line activities, such as the use of e-commerce sites like amazon.com and ebay.com, have led to the development of on-line disputes. It makes sense that if a transaction occurs online, then disputants are likely to accept online techniques to resolve their differences. New ways of dispute resolution are hence appearing, so that the disputant parties neither need to travel nor to meet in courtrooms or in front of arbitrators or mediators. Different forms or methods of alternative dispute resolution for electronic environments have been pointed out by legal doctrine. As a result, we can now speak of Online Dispute Resolution (ODR) as any method of dispute resolution in which wholly or partially an open or closed network is used as a virtual location to solve a dispute (Katsh and Rifkin, 2001).

A relevant issue, in a quick look, will be to inquire in what way (and to what point) traditional mechanisms such as negotiation, mediation or arbitration can be transplanted or adapted to the new telematic environments, taking advantage of all the resources made available by the newest information and communication technologies. In order to develop intelligent and efficient techniques to support Online Dispute Resolution, one needs to integrate Artificial Intelligence based problem solving techniques with Online Dispute Resolution ones. This information can be considered from two different perspectives: on the one hand, as a tool to help the parties and the decision makers to obtain the best possible results in solving commercial disputes and, on the other hand, considering a new way of autonomous dispute resolution through the use of autonomous and intelligent software, supported by a knowledge base and decision capabilities. Thus, it will be important to consider the many alternatives for dispute resolution that arise from Artificial Intelligence models and techniques (e.g., Argumentation, Game Theory, Heuristics, Intelligent Agents and Group Decision Systems) as pointed out by (Peruginelli, 2002) and (Lodder and Thiessen, 2003).

ODR systems rely on technology to support alternative conflict resolution processes

Contrary to previous approaches, in Online Dispute Resolution it must be considered not just the disputant parties and the eventual third party (mediator, conciliator, arbitrator) but also what Ethan Katsh and Janet Rifkin call "The fourth party", i.e., the technological elements involved. An important element of this "fourth party" will obviously be the emergence of expert systems and intelligent software agents empowered to help the parties and the mediator/arbitrator in reaching a satisfactory solution. And as (Lodder, 2006) already refers, it must be considered as well a "fifth party", i.e., the service providers, those who provide and deliver the technological elements. All this is turning ODR in a quite new and somewhat complex (but eventually quite fast, cheap and advantageous) way of interaction and of solving conflicts.

The goal of AI research in this field is to attain a technological threshold, resulting in computational systems that are indeed the third party. In this sweeping approach, there is no human intervention on the outcome or in guiding the parties to a specific situation. There is, on the other hand, a system that performs that major role. This is usually known as an electronic mediator or arbitrator. It should have skills for communicating with the parties and understanding their desires and fears and have the ability to decide on the best strategy to be followed in each possible scenario. This is evidently the most challenging approach to follow since the cognitive abilities of a human expert, as well as the ability to perceive emotions and desires, are not easy to implement in an artificial system. On the other hand, there is an inherent risk in letting machines take binding decisions that influence our lives (Lodder and Zeleznikow, 2010).

ODR systems can be categorized according to the function that machinery may play (Peruginelli, 2002). First generation ODR systems describe the systems that are used in our moment in time. The main idea behind these systems is that the human beings remain the central pieces in the planning and decision making processes. Computational tools are evidently used, but they are seen as no more than equipment, without any autonomy or a major role in the course of action. In this kind of ODR systems the main technologies used are instant messaging, forums, video and phone calls, video conference, mailing lists, and more recently, Video Presence. Agent-based technologies may be used but have no active part or autonomy. These systems are common nowadays and are usually supported by a web page. They represent a first necessary step before the consideration of those that may be more autonomous, a characteristic that may be achieved through the use of intelligent systems.

The second generation of ODR systems is essentially defined by a more effective use of technical tools. It is no longer used for the mere job of putting the parties into contact and/or making access to information easier. It goes beyond that and it is used for idea generation, planning, strategy definition and decision making. In that sense, it can be said that second generation systems extend its first generation with new intelligent and autonomous artefacts. This new generation relies and is supported by technologies that allow for a regular connectivity among all the entities involved. However, by using innovative technologies on top of this communication layer, it is possible to present services with more added value.

For the implementation of such services, one can look at fields as diverse as Artificial Intelligence, Mathematics or Philosophy. In the intersection of these fields one can find a range of technologies that will significantly empower the previous generation of ODR tools, namely artificial neural networks, intelligent software agents, case-based reasoning mechanisms, methods for knowledge representation and reasoning, argumentation, learning, and negotiation. Thus, we move forward from a paradigm in which reactive communication tools are used by parties to share information, to a virtual environment in which ODR services proactively assist the disputant parties.

Therefore, it is clear that the involvement of different areas of research, namely the ones of Artificial Intelligence, may contribute to develop ODR processes that will deal with other sorts of problems, namely complex multi-party, multi-issue, and multi-contract ones. Using such approaches will also result easier to develop processes that mimic the cognitive processes of human experts, leading to more efficient ODR tools.

Table 8 in Appendix A details a number of commercially available ODR providers and classifies them according to some key features. A previous survey by Melissa Tyler exists, conducted in 2004, in which 115 ODR sites were identified (Tyler, 2005). In that sense, an extensive list of sites was not included in this thesis. Instead, the most important ODR providers were selected, in order to compile a list that can show the wide range of fields in which ODR tools can be used. Moreover, several new sites were identified and included. In this compilation, the isolated use of certain telematic means (such as video-conferencing) is not regarded as an advantage. As an example, there are systems that claim to allow video-conferencing between the parties although in reality they simply arrange the meeting time while parties must use an external video-conference tools. On the other hand, systems that create virtual environments in which such telematic means are embedded and used seamlessly by parties are valued in this analysis.

1.5 *pro et contra*

ODR is a relatively new approach to dispute resolution since the technologies it builds in are also recent. It has however known a fast growth in its use given its advantages, mostly in the United States and Australia. ODR tools are generally intuitive and easy to use, being friendlier than courtrooms. Most of the times, clients interact with intuitive interfaces that hide all the complexity of the laws and formalities behind these processes. This increases the willingness of the parties to solve their disputes online as they feel that it is a more transparent and controlled process. As these systems are available 24/24, parties can submit documents and evidence and use the services at any time.

By having the possibility to communicate synchronously and asynchronously, these tools can be used by parties to communicate even when the time zones are different or when it is impossible to agree on a specific time to meet. It is possible for the parties to communicate with each other and with the third party from virtually any part of the world. This reduces eventual costs with transports and accommodation. Still regarding costs, these tools are generally cheaper and faster than common litigation since they are based on inexpensive and available technologies and avoid all the costs associated to courts.

Another important characteristic is that by being behind an interface instead of being in front of a judge, the environment is less intimidating for the parties. In these conditions people tend to be less afraid of talking. At the same time, because the parties are not in contact personally, they tend to be more focused on the subject under discussion that on the opponent itself, more focused on the facts than on the personal feelings. There are even some extreme cases in which the parties refuse to sit together at the same table so this is often the only peaceful way of trying to solve the dispute. Furthermore, when parties are not in the presence of other individuals that might have a suppressor or inhibitor effect (such as the relations in the cases of domestic violence), they also tend to speak more freely and without fearing the consequences.

When using these tools, parties tend to reflect more and end up taking better decisions than when the processes are carried out in person. This is due to the fact that parties have more time to think about what they are going to do or say than if they were in the presence of their interlocutors. This is even more notorious when asynchronous communication mechanisms like forums or email are being used. It is also easier for all the parties to manage and access the information of the case as everything is available online all the time. The creation of formal documents is also easier since there may be digital assistants that help the parties through the process or that even create the documents automatically. Communications between the parties are also stored so that they can later be analysed and can even act as evidence for posterior phases.

Summarizing, ODR tools can be a more accessible, fast, economic and transparent way of solving disputes. However, the factors that in some cases constitute an advantage of ODR may also constitute disadvantages. If, in some scenarios, moving from a paper based process to an online one brings it closer to people, there are others in which that does not happen. This is true for some population groups that do not have access or that do not have the proper training for using these tools, which results in inequalities. Training can be provided in such cases but the time spent to do it could delay the processes.

Furthermore, it is true that in most of the developed countries the technologies needed for ODR are cheap and available. However, on the other hand there are developing countries in which this is not necessarily true and might lead to higher costs. This may also represent an inequality and inhibit some parties from using an ODR tool. Another disadvantage related with technology concerns security. There are scenarios in which storing the information online means that there is more security. Nevertheless, there are also scenarios in which it means the opposite. If one party has access to the authentication information of the other, it becomes easy to access and change confidential information of the other party. In such cases, it gets even easier than if the information was stored in a paper-based support. This raises the problems associated to online identity: when we communicate with someone online, how can we be sure that it is really with that person that we are interacting? How do we know that someone hasn't stolen the identity of the person and is using it?

Another major disadvantage has to do with online communications. If on the one hand one does not feel intimidated by not speaking in front of a jury or an opponent and may speak truer, it also becomes easier to lie online since there is not the presence of an authority. One even more important issue is the one of the body language. (Mehrabian, 1980b) states that most of the meaning that we derive from a face-to-face conversation comes from other aspects than the words spoken, namely the tone of voice, the loudness, the facial expressions and body gestures. Evidently, all this important context information that makes up a conversation is lost when technologies like IM, email, forums or similar are used. The best solution for this problem would be to use more recent technologies like video conference or TelePresence, which are more similar to a face-to-face conversation, although that is not always possible. The key idea is to place the emphasis on transmitting how things are said instead of just what is said.

Despite the described disadvantages, the use of ODR tools is increasing because the technologies used are becoming cheaper and more available and because the tools themselves are becoming more intuitive and easy to use. In a general way, there is the sense that the disadvantages are supplanted by the numerous and significant advantages.

1.6 SUMMARY

Judicial Systems are becoming slow and unresponsive as the number of pending processes grows in an unprecedented way. This happens despite the increase in the number of attorneys and judges. The main reason is that nowadays every act in our life is likely to end in court, despite its importance or value. However, the litigation environment is also to blame since it does not favours cooperative work. It is a rather competitive environment, in which parties focus on the maximization of the own gain, no matter what. This constitutes an obstacle to successful outcomes, increasing the number of appeals and the time needed to end the process.

In order to address some of these drawbacks Alternative Dispute Resolution methods emerged that have as main objective to create a cooperative environment in which parties willingly and actively participate in the process. Such methods include negotiation, mediation or arbitration and are generally seen as cheaper and faster alternatives to litigation. More recently, these methods started to be used with the support of Information Technologies, allowing for parties to interact, synchronously or asynchronously, despite their physical location. Moreover, such tools, in their most basic form, allow to create, edit exchange and manage legal documents. More evolved tools may even generate ideas, propose solutions and strategies or provide important information for the resolution of the process. Table 7 in Appendix A summarizes the main differences between litigation and alternative methods.

2 ARTIFICIAL INTELLIGENCE AND ODR

The important thing in science is not so much to obtain new facts as to discover new ways of thinking about them.

— Sir William Lawrence Bragg

Artificial Intelligence techniques emerged in the middle of the last century as a way to emulate features generally associated to intelligence, in machines. It is probably the most ambitious and one of the most controversial research fields in Computer Science as it aims for machines that can reason at the same or at a better level than Humans. In this section the role that Artificial Intelligence can play in the development of better ODR tools is analysed. Specifically, several particularly interesting sub-fields of the Artificial Intelligence umbrella are described in detail, in terms of the advantages they can encompass. The section ends with an overview of the present and some thoughts about the possible future of the intersection of AI and The Law.

2.1 ARTIFICIAL INTELLIGENCE

When looking for a definition for the term Artificial Intelligence, we can think on one of the first ones to be proposed by John McCarthy, an American computer and cognitive scientist and one of the most influent scientists in AI research:

Artificial Intelligence is the science and engineering of making intelligent machines

— John McCarthy, American scientist (1927 - 2011)

This definition is interesting in the sense that it says many things and nothing special at the same time: it is quite generic since it depends on the definition of "intelligence", which is not consensual. In humans, intelligence is generally defined in function of a set of capabilities including, but not limited to, learning, self-awareness, communication, reasoning, abstract though or planning. The Merriam-Webster dictionary defines intelligence as "the ability to learn or understand or to deal with new or trying situations" whereas the Cambridge dictionary defines it as "the ability to learn, understand and make judgements or have opinions that are based on reason". Intelligence has also been studied in animals or in plants, where a reduced set of these capabilities also depicts intelligent behaviours.

Thus, in search for a more precise definition, it results logic that Artificial Intelligence can be defined as machines that exhibit capabilities usually associated to humans. Classical research problems in this field focus on the ability of machines to reason and take decisions, to acquire and use knowledge, to plan strategies, to perceive patterns (e.g. images, languages, sounds) and the ability to move and manipulate objects (Russell and Norvig, 2009). These abilities shall be used with the objective of maximizing some own gain so that the machine is better at the end of its use than it was at the beginning. Doing so demonstrates that the machine is not only able to recognize its present situation but also to establish an objective towards the improvement of some aspect, draw a plan of action and execute it.

Artificial Intelligence can be seen as the study and design of software agents that perceive their environment and take actions that maximize its success

Hereupon, Artificial Intelligence can be seen as the study and design of software agents that perceive their environment and take actions that maximize its success. Important words here are "autonomy" (the machine should take decisions without the intervention of other entities) and "proactivity" (the machine should take the initiative of changing the state of the world).

The ultimate objective of research in this field is known as Strong AI. This concept denotes the moment in time when computers will have cognitive skills similar or even better than humans (Kurzweil, 2000). The expression was coined by John Searle when trying to distinguish between two different hypotheses of artificial intelligence: one that could only act like it thinks and one that really thinks.

Searle named the first one as the weak AI hypothesis and the second one as strong AI hypothesis (Searle, 1980). The pursuit for this not yet achieved objective of AI is based on the concepts of intelligent agent (Wooldridge, 2002) and Multi-agent Systems. Strong AI as seen by Searle has not yet been achieved. There is also no consensus about the minimum set of abilities for a machine to be considered intelligent at a Human level: it is not enough to be better than Humans at performing some tasks.

In 1950, Alan Turing proposed a test later named after him that would remain as the best-known test for determining whether a computer system is or is not exhibiting an intelligent behaviour (Turing, 1950). For testing the ability of a machine to imitate our way of thinking, Turing proposed the following test: place a person and a computer in two isolated rooms and a second person, the judge, in a third isolated room. The judge exchanges written messages with the other two participants. When the judge cannot consistently tell which one of the two interlocutors is the human and which one is the computer, then the computer is considered intelligent.

Artificial Intelligence has been growing and is now an umbrella for many application areas. We may even deal with applications of Artificial Intelligence in a daily basis without noticing it. This apparent pervasiveness is related with the fact that we get used to computers showing intelligent behaviours and we take them as granted, not noticing that they are applications of AI techniques. This is called the AI effect. Hogan describes this effect in an intuitive way in (Hogan, 1998):

At the outset of a project, the goal is to entice a performance from machines in some designated area that everyone agrees would require "intelligence" if done by a human. If the project fails, it becomes a target of derision to be pointed at by the sceptics as an example of the absurdity of the idea that AI could be possible. If it succeeds, with the process demystified and its inner workings laid bare as lines of prosaic computer code, the subject is dismissed as "not really all that intelligent after all".

— James P. Hogan, British science fiction author (1941 - 2010)

In fact, the work of AI researchers is often considered ungrateful as the developments, after completely formalized and known, tend to lose the mystic and are looked as a development of any other computational area. Companies that use AI techniques but do not mention it when they expose or sell their work fearing rejection are also common. This can be due to the frenzy that AI has developed on the early years, creating high expectations but taking much time to achieving them, leading to some disappointment.

Nevertheless, AI research has led to the development of many technologies that are nowadays used, most of the times on the shadow of the big systems they build. These technologies are in general used to optimize work in knowledge-based domains, to make products easier to use through the adoption of intelligent interfaces or to automate tasks.

2.2 HOW CAN AI IMPROVE ODR

It is a fact that information systems are being intensively used in virtually every domain, and the legal one is no exception. However, in the legal domain the potentialities of computers are not being fully exploited, being relegated to basic back office tasks such as text processing, billing, agenda management or communication. Nevertheless, the importance of technology in this field will slowly increase as AI techniques develop.

Historically, the beginning of research on AI & The Law can be attributed to Bruce Buchanan and Thomas Headrick when, in 1970, they published the paper "Some Speculation About Artificial Intelligence and Legal Reasoning" (Buchanan and Headrick, 1970). In the interim, research in this field began to increase, with the appearance of some international conferences, associations and journals, revealing the growing interest of the scientific community for this area.

In this section some of the branches of AI research with potential interest for conflict resolution are analysed. Particular emphasis is put on determining how each one can be used to improve current approaches. Table 9 in Appendix A highlights the most interesting features of each of the sub-fields analysed here.

Decision Support Systems.

With the constant growth of the amount of information present in the decision processes, the need for tools to provide support has also grown. Indeed, the new economy, along with increased competition in today's complex business environments, takes the companies to seek complementarities in order to increase their competitiveness and reduce risks (Bonczek et al., 1981). Under this scenario, planning takes a major role in a company's life. However, effective planning depends on the generation and analysis of ideas (innovative or not) and, for this reason, the idea generation and management processes become a crucial tool in present days. The tools used may range from simple systems for compiling useful information from raw data, to more complex ones that make suggestions on the best strategy to be used or the fairest outcome.

Decision support systems may be used in virtually any knowledge based environment and the legal domain is not an exception (Turban, 1993). In the legal arena, these are known as legal decision support systems. However, as its use is still recent, there are no advanced implemented systems. Nevertheless, the ones that have been developed so far have something in common (Zeleznikow and Hunter, 1994): they are rule-based. There are several reasons for this: rule-based systems are generally easy to understand and implement, there are many tools for building rule-based systems and many legal concepts and processes can be modelled using rules. These rules are instructions of the type IF condition THEN conclusion, that is, if certain conditions are verified, one or more conclusions will be true.

Considering the complexity of the legal domain, legal decision support systems can be quite useful, specifically if one considers the significant amount of information that parties and neutrals must analyse in more complex cases in order to take decisions. Without using any supporting tools, doing it is an inefficient and time-consuming process.

Such systems have the ability to analyse relevant facts input by the parties as well as legal information such as norms or past known cases in order to make simple legal decisions. There are some fields in which decision support systems have been more significantly used. Social security systems use them to help practitioners deciding if an unemployed individual should or should not receive a benefit. Banks use them in order to more efficiently decide if a client should be granted a loan. Insurance companies use decision support systems when deciding on the amount of an indemnity to be paid to an insured.

In any of these cases, as well as in the legal domain, results are generally supervised by human experts. Decision support systems are therefore not automated systems that issue outcomes. They are systems that, based on important information, issue justified recommendations and compile information that can be useful for the decision makers.

One example of application of decision support systems in the legal domain is the "Split Up" system (Zeleznikow and Stranieri, 1995). This is an intelligent decision support system that makes predictions about the distribution of marital property following divorce in Australia. Its main purpose is to assist judges, registrars of the Family Court of Australia, mediators and lawyers. Split Up operates as a hybrid system, combining rule-based reasoning with neural networks. A more recent example in the same legal domain can be found in (Abrahams and Zeleznikow, 2010), in which an agent-based negotiation decision support system for the Australian family law is presented.

Expert Systems

According to (Susskind, 1988), Expert Systems can be defined as computer programs that have been constructed in such a way that they are capable of functioning at the level of (and sometimes even at a higher standard than) human experts in given fields. In that sense, such systems are designed, trained and fine-tuned by humans and must embody a depth and richness of knowledge that allow them to perform at such level (Hayes-Roth et al., 1983).

Training can be performed using information from past known cases and respective decisions provided by human experts. On the other hand, these systems can also learn while they are used, generally with the supervision of an expert that makes adjustments according to the input, expected output and verified output. Similarly, (Harmon and King, 1985) defines Expert Systems as intelligent computer programs that use knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution. Both definitions share one common idea: Expert Systems try to mimic the human expertise and knowledge in a given domain (Jackson, 1990). In that sense, it is correct to say that the knowledge necessary to perform such high level tasks as well as the inference procedures used can be seen as models of expertise of the best human experts in the field.

Expert Systems also represent a change in the programming paradigm (Forsyth, 1986). While traditional computer programs are seen as procedures applied to data, Expert Systems are seen as inference engines applied to knowledge. In that sense, the two new major modules are a rich knowledge base and a powerful inference engine. However, if we want to be more specific, four main components can be identified in a fully functional Expert System (Greinke, 1994): (1) the knowledge acquisition module, (2) the knowledge base, (3) the inference engine and (4) the user interface. In that sense, a fully functional Expert System is expected to be able to deal with information relating to a specific problem domain, analyse it and generate knowledge, and then take actions and decisions based on that knowledge. It is expected that these actions resemble the ones that a human practitioner would take in a similar scenario.

Expert systems can be found in a wide range of domains, including the fields of accounting, medicine, process control, financial service, production or human resources. There are several factors that led to a growing use of such approaches. On the one hand, Expert Systems can automate simpler tasks, releasing human experts to other higher level tasks or, eventually, allowing companies to reduce costs. On the other hand, the huge amount of information that practitioners in certain fields must deal with renders nearly impracticable a purely human workforce. This is also true in the specific case of the legal domain. In fact, legal practitioners can no longer deal with the increasing number of disputes and the information that each one requires, resulting in the well-known slowness in judicial systems. In that sense, Expert Systems can be a tool that helps legal practitioners deal with huge amounts of information, automating simpler tasks and, ultimately, allows them to work more efficiently.

On the whole, particularly constructive when considering the legal domain is the ability of an Expert System to detail the reasons for a specific analysis or recommendation, i.e., its ability to explain its actions. In order to do so, legal expert systems generally allow the assignment of weights to factual data on a case. This may trigger additional actions, such as comparing a given case to the cases stored in the knowledge base, producing outcomes based on similarity metrics. However, once again, these outcomes should be regarded as merely informational, i.e., legal expert systems should be used, for example, to help judges to deal more rapidly with the cases, providing guidance based on a model of the legal domain that applies, including the norms, facts and past cases.

Currently, complex Expert Systems of this level are not yet established. There are even authors that argue that these are not only difficult to implement but also unnecessary. Following this simplistic approach, (Popple, 1996) presented SHYSTER, a simpler, pragmatic approach in which the utility of a legal expert system is evaluated by reference, not to the extent to which it simulates a lawyer's approach to a legal problem, but to the quality of its predictions and of its arguments. In fact, most of the legal expert systems currently at use are fairly simple implementations, focusing, for example, on automated drafting of complex legal documents. In such systems, users are generally guided through a series of interfaces with questions, while receiving practice tips or support about the legal domain or strategies. Nevertheless, other authors argue that a purely rule-based approach is inappropriate if the Expert System is to be of use to a lawyer. (Popple, 1990) concludes that a better approach is obtained when rule-based methods are combined with case-based ones.

Knowledge-based Systems

Knowledge is an abstract term that represents a collection of specialized facts, procedures, and judgement conventions. There are many types of knowledge and many different ways of acquiring it. First, knowledge can come from a single source or it can be compiled from several sources. Depending on the domain, it can be compiled from human experts (e.g. observing the behaviour of a law practitioner), sensors (e.g. a domotic environment), pictures (e.g. medical imaging), maps (e.g. finding a path), flow diagrams or historic context, just to name a few. Depending on the type and source of information, several techniques for knowledge acquisition can be used, namely human observation, scanners, pattern matching, pattern recognition or intelligent agents.

Being the field of AI & The Law a knowledge-based one, the subject of Knowledge Representation (KR), is a very important one. According to (Sowa, 1999), KR is a multidisciplinary subject that applies theories and techniques from three other fields: logic, ontologies and computation. Logic provides structural formalisms and rules of inference. It is used for determining the validity of arguments, as well as their redundancy or contradiction, allowing for the test of the validity of a given outcome. Ontology defines the domain or context being dealt with, as well as the objects that subsist in the domain of discourse. Without ontologies, words have no meaning or value, are minimal sets of characters. Computation denotes the needed working out support for dealing with these questions on the processing systems sphere. Without this, KR would be an attitude problem as it has been until recently.

In a few words, KR is concerned with formalizing our way of thinking, i.e., how to represent a given domain with symbols. Considering the complexity of legal knowledge, the development of systems that are able to formally model knowledge is highly desirable (Brachman and Levesque, 2004). These are the so-called Knowledge-based Systems.

Considering the legal domain, such systems are important, in a first stage, to define a model that can deal with a heterogeneous group of information. This group may include, but is not limited to party's information, norms, past cases, facts or arguments. These systems are also essential when such information must be stored digitally in a way that allows for fast and efficient retrieval. When knowledge is stored in a formal and well-defined way, it allows for the development of automated processes that can, for example, interpret the validity of logic arguments. The development of systems that can efficiently cope with huge amounts of knowledge is in fact one important advance that legal practitioners can take advantage of, in order to deal with the current increasing number of disputes. Moreover, knowledge-based systems can be designed to deal with either statute law or case law (Popple, 1996). The main motivation behind the use of Knowledge-based Systems in the legal arena is its capability of representing norms and judgement under uncertainty. In fact, systems can be developed that can produce new facts or conclusions based on knowledge.

In general, five modules can be identified in a fully functional Knowledge-based System: (1) the user interface; (2) the explanation facility that details the inference mechanism regulating the system outcome; (3) the database with the factual information; (4) an inference mechanism that decides on which rules are to enforce and how they are prioritized; and (5) the knowledge acquisition module that is responsible for (possibly in an automatic way) the acquisition of knowledge from the outside world.

Intelligent Interfaces

Lawyers currently face a problem that has already been pointed out in this thesis, i.e., the ever growing amount of information that must be considered in legal problems, either in statute law or case law. On the one hand, in statute law, new statues and treaties are making its way, making legal analysis more complex. On the other hand, case law is faced with more and more disputes, which generate an exponential increase in legal rulings. The main reason for this is the informatization of the society, which not only adds to the information available but also increases the number and diversity of the disputes that must be solved.

Ironically, the same technologies that led to the exponential availability of information also show the way to the development of tools to deal with this information, i.e, technologies such as Expert Systems, Decision Support Systems or Knowledge-based Systems are now available that can assist practitioners. Nevertheless, the adoption of these useful tools has been rather slow, wasting their theoretical advantages. A possible cause for this drawback is pointed out by (Matthijssen, 1995), when the author states that in legal information systems the interface-modalities do not shield the users of these systems from the internal organization of the data and the additional workload associated to the processing mechanisms, i.e., legal practitioners find a conceptual gap in this process, once they work in a given way, and legal information systems are either designed to work in a different way or are designed very closely to the internal structure of the system, providing no abstract interpretation of the decision process.

This gap can be filled with the development of the so-called Intelligent Interfaces. Therefore, it will be useful for the developers to be aware of the way practitioners solve legal problems. Using this information, intelligent interfaces can be developed that reflect the knowledge domain of the practitioner rather than the structure of the stored data. The main objective here is that practitioners can focus on the actual content of the legal concepts rather than on how these concepts are translated and stored in legal information systems. Intelligent Interfaces are very singular, i.e., besides making the bridge between humans and computers, they present additional features.

The eternal problem addressed, present on the legal domain (as well as in any knowledge-based domain), is the one of efficient and effective retrieval of data (generally from a database). When the methods for retrieving information are much related to the structure of the data, it becomes harder for a non-expert to perform efficient searches. According to (Matthijssen, 1999), to a large extent these problems can be attributed to the limitations of the traditional Boolean query mechanism used in text databases, which is difficult for users to operate. Using Intelligent Interfaces, it is possible to develop different forms of abstraction, at the user level, that make it possible to personalize methods to access data, regardless its structure.

In order to implement this behaviour, an intelligent interface needs specific data about the legal domain that is being addressed, as well as models for the representation of legal knowledge, its rules and processes. Additionally, such interfaces can also take into consideration personal preferences or user roles. Then, they can act as an intelligent intermediary between the user and the database. Using such interfaces, practitioners can make use of a more intuitive and powerful tool to analyse and organize information. Possible applications include the structured publication of large amounts of information, automated organization of data according to a given criteria or automated search.

Another interesting area of application is one in which the user is not completely sure of how to search or what to search. A search request may be incomplete, incorrect or inaccurate and the interface is responsible for assisting the user in reformulating the request or trying to guess what the user intentions are in terms of search. In order to fulfil these goals, the interface must be adaptive, proactive, anticipate the needs of the user and be able to explain its actions.

Search engines like Google or Yahoo can also be seen as intelligent interfaces. In fact, they often do successful searches although we misspell the search terms, or suggest similar words or concepts in order to make our search more accurate. Intelligent Interfaces also filter the information, deciding which is closer to what the user is looking for and which is useless. In order to do this, context information is taken into account (e.g. legal domain, past experiences, domain of expertise of the practitioner). One particular case of application of Intelligent Interfaces in the legal domain is the one of the intelligent tutors, aimed at teaching or training its users in a given area. Two examples are LITES - an intelligent tutoring system for legal problem solving in the domain of Dutch Civil law (Span, 1993) and (Ashley and Aleven, 1991), where an intelligent tutoring system for teaching law students to argue with cases is described.

Case-Based Reasoning

Case-based Reasoning (CBR) can be described as a problem solving methodology that relies on past experiences and its data to make present choices (Kolodner, 1992; Aamodt and Plaza, 1994). The key assumption is that if a new problem is similar to an old one, it will have a similar outcome. This procedure is commonly observed in humans and is intrinsically related with our learning processes. As an example, let us consider that some time ago an individual left the home with a cloudy sky, and the clouds turn into rain and he got wet. A few days later, before leaving his house, the same individual looks at the sky and, as it is cloudy, takes an umbrella with him.

In general, this process involves the ability to compare two scenarios (or cases) and admit that if they are similar, they will have an identical outcome. The first task is thus to select among all the features that describe a case, which are the ones that are useful to determine the similarity between two cases. Failing to do so will lead to the impossibility of dealing with all the attributes that define a given universe and their range of possible values. Continuing with the previous example, the individual could take the decision of taking or not taking an umbrella based on different factors: the day of the weak, the weather forecast, the current weather conditions, the clothes worn and/or the distance to the destination. While some of these factors make sense (e.g. the current weather conditions, the weather forecast) considering the nature of the problem, others are completely irrelevant. The first challenge is therefore to select which attributes to consider, according to the problem domain (or universe of discourse).

It is also essential to enquire the relative significance of each of the problem attributes. In our previous example, it makes sense to consider both the current weather conditions and the weather forecast. However, a different weight might be given to the weather forecast attribute if the individual is more worried about the evolution of the weather conditions and not so much concerning the immediate ones. This factor may however be different, depending on the time the individual will be outside or its confidence on the weather forecast.

Generally, a CBR process is organized in four sequential phases: Retrieve, Reuse, Revise and Retain (Kolodner, 1993).

In the first phase, the problem is analysed and the cases that are relevant (i.e. similar enough) are retrieved from memory, possibly ordered according to a value of similarity. This measure of similarity depends on the problem domain but generally consist of on a difference of the sums of the different attributes that characterize the cases.

In the Reuse phase, the solution from the previous case is mapped to the target problem, which may involve adapting the solution to some specific requirement of the new problem. This phase is necessary since, in general, there is no known case that exactly matches the attributes of the new one. In theory, it would be possible to retrieve as many cases as needed to cover all the different attributes. However, the size of such a case base could be impracticable. There is, hence, an implicit compromise between the amount of cases stored and the values of similarity achieved.

In the third phase, the solution is tested or simulated in an attempt to determine the result of its application. It may be possible that the results are not as good as expected, which should lead to the revision of the action taken. In the last phase, the solution adopted may be stored in the case memory, along with the description of the new case, contributing to the enrichment of the knowledge about the domain.

CBR is obviously suitable to be used in the legal domain, once the ability to predict or estimate an outcome is an important component of legal advice. A legal practitioner frequently examines past similar cases and their outcomes to try to predict the outcome of a new case. There is even a similar legal concept: the legal precedent (Landes and Posner, 1976). The notion of legal precedence defines a case that establishes a rule or principle that could or should be used by practitioners when deciding on subsequent similar cases.

CBR models are, in principle, particularly useful in common law systems, in which The Law is interpreted and applied by judges. Nevertheless, civil law systems (in which The Law is written by a legislature's enactment (Zweigert and Koetz, 1998)) can also be approached from a case-based perspective, namely through the development of systems that target the retrieval of information with the objective of informing the users instead of producing outcomes. Considering this topic, Ashley poses the question: "Should researchers in a civil law jurisdiction pursue work on implementing AI & The Law models of case-based legal reasoning in a civil law context?" He answers with a conclusive "the answer may well be, "Yes"!" (Ashley, 2004).

Although being object of research, CBR is already one of the most commonly used approaches in the development of intelligent and learning systems, for the most varied purposes. (Watson, 1997) gives some examples of big companies operating in fields such as air and fraud management in which the use CBR is routine. Likewise, the legal field has some implementations of CBR that address specific problems. HYPO models the way attorneys argue when confronted with a case, either real or hypothetical (Ashley, 1991). CATO is an intelligent learning environment, designed to help law students to learn the basic skills of argument building when leading with a case (Aleven, 1997).

Multi-agent Systems

Multi-agent Systems (MAS) (Wooldridge, 2002) emerged from the combination of AI-based methodologies and techniques for problem solving with distributed computational models, generating a new area of research: Distributed Artificial Intelligence. Many different definitions for a MAS exist. In this thesis they are seen from a legal point of view. A MAS is a group of entities (software or hardware) that make intelligent decisions in order to achieve some common goal (like proposing a solution for the parties in dispute) based on information that is shared between every agent in the system.

(Parunak, 1997) proposes a detailed definition, based on the presupposition that a MAS is not only defined by the agents or their properties. A MAS is rather defined by a triple: a set of agents, an environment and a pairing between them. It is easy to agree with Parunak since an agent is genuinely associated with the environment, as its actions depend on the state of its peers. As an analogy, we humans commonly look to ourselves in function of our social or geographical positioning, i.e., our environment and our social relations make us who we are.

Agents materialize an appealing computational tool as they allow for a wide range of behaviours and/or functionalities to be analysed, specified and implemented. In particular, there is a set of assets proposed by (Wooldridge and Jennings, 1995) that make what the authors call the weak notion of agent: autonomy, social ability, reactivity and pro-activeness. This vision entails that the most basic agent should at least: (1) operate without the direct intervention of humans and formulate its own decisions in an autonomous way; (2) be able of interacting with other agents (independently of their nature); (3) perceiving the environment and responding on time to the stimuli and; (4) take the initiative of pursuing its goals. The same authors also proposed a stronger notion of agent that may include properties such as mobility, veracity, benevolence or rationality. This means that additionally, an agent may move between locations by means of a network, will not give false information on purpose, will not have conflicting goals and always will try to do what is asked for (an agent will always act in order to achieve its goals).

The main objective of the present approach to computing is to address the complexity of intelligent behaviour intra and inter communities of simple entities or agents (Olson et al., 2001), i.e., agents must be able to autonomously make their undemanding assessments that, once combined, may lead the communities of agents to perceive intelligent behaviour. This approach has nowadays a major role in the design of intelligent systems. Especially interesting for the legal domain is the research on argumentation theory. In argumentation, agents debate, defend their beliefs and try to convince their peers of the rationality of their causes (Rahwan, 2009). There is here an evident parallelism with the argumentation procedures that take place during dispute resolution processes. Agents may also implement negotiation techniques (Beer et al., 1999).

In the context of a MAS, negotiation refers to the modelling of human conciliation techniques so that they can be used for solving conflicts between agents. The main field of application of this modus operandi is in conflicts that arise from auctions and e-Commerce. In this specific sub-field of dispute resolution, agents may represent the parties in a negotiated settlement and try by themselves to get to an end, then suggesting it to the parts in dispute. An important analogy may also be done with negotiation procedures that take place in the legal arena, between parties that are trying to achieve a common agreement. A different kind of added value that comes with the use of MAS, from which the legal domain may profit, is distributed problem solving. Significant virtues in the legal domain (e.g. veracity, benevolence) can also be instilled into agents, namely in the so-called emotion-based ones (Velasquez, 1997). In the legal field, this kind of work may lead to the implementation of the second generation of ODR systems, with the ability to understand the feelings of the parties according to each topic of the dispute.

Legal Ontologies

In philosophy, an ontology represents the study of the nature of existence in general. In that sense, ontologies deal with the questions that concern the definition of a given entity, its existence, and how that entity relates with others. In computer science, ontologies are a way of formally representing knowledge in terms of concepts within a domain and the relationships between those concepts. According to (Gruber, 1993), an ontology is a "formal, explicit specification of a shared conceptualisation". In order for the ontology to be understood, a shared vocabulary must be provided. This vocabulary must contain all the concepts that can be used to model the domain of discourse, i.e., the ontology must define the type of each concept as well as their properties and relations. Therefore, in ontology specification, one defines classes and subclasses of individuals as well as the properties of each individual in a class or subclass. If, on top of that, there are also defined relationships between individuals, it will be possible to infer properties, namely by inheritance.

In computer science, ontologies are nowadays paramount, mainly because they are the enablers of the so-called Semantic Web. The Semantic Web describes a group of methods and technologies that allow machines to understand the meaning of information on the Web, rather than simply accessing it. That is indeed the main innovation that ontologies brought along, i.e., allow machines to read, interpret and understand information. Logically, such technologies can also be used in other domains than the Web, ranging from software engineering or biomedical informatics to library science or information architecture.

Indeed, for complex domains such as the legal one, the advantages are considerable. For instance, by systematizing knowledge, it becomes readily available. It allows not only for the extraction of rich patterns of information that otherwise would not be perceptible, but also to draw inferences. Indeed, computer models that can efficiently deal with huge amounts of structured information are being object of research, namely by (Corcho and Fernández-lópez, 2005), that presents a methodology to build an ontology in the legal domain following the development method METHONTOLOGY, and using the ontology engineering workbench WebODE; (Visser and Bench-Capon, 1998) that presents, compares and critiques four different legal ontologies. For more interesting examples on legal ontologies, the book "Law and the Semantic Web" presents a selection of revised papers drawn from two meetings devoted to the Semantic Web and the legal domain: The International Workshop on Legal Ontologies and Web-Based Legal Information Management (Benjamins et al., 2005).

Rule-based Systems

Rule-based Systems (RbS) are, in general, the straightforward way of implementing an intelligent behaviour, i.e., they stand for the simplest form of building Artificial Intelligent systems. Using a RbS, it is possible to encode the knowledge and skill of a human expert in a given domain in the form of IF-THEN rules, in which each rule denotes a small piece of the expert's knowledge. Rules have a left and a right hand side. On the left side there is information about facts that must be true in order for the rule to be enforced. On the right hand side, the rule contains the actions that should be carried out whenever the rule is fired. The model of execution of a RbS consists, therefore, in analysing on the fly the left hand side of all rules. The rules whose left hand side is evaluated to true are placed on an execution agenda. Then, rules in the agenda will be executed, without any explicit order, and then removed from the agenda. One singularity of RbS, contrasting with standard object-oriented programming, is that there is not an effecting order that can be predicted beforehand.

RbS are a way to store, interpret and manipulate knowledge about a given domain. In fact, if appropriate design strategies are followed, these systems allow for an ease access to an expert knowledge, i.e., whenever knowledge about the domain changes, only specific rules need to be transformed. This can be even more straightforward if a proper and intuitive rule editor is used. In order to have a fully operational RbS, a rule engine is necessary.

If we consider specifically the use of RbS on the legal domain, several advantages can be put forward, namely when representing the legal corpus and other legal concepts as RbS. The most obvious one is that when capturing the expertise of an expert in a given field, that expertise will become available to all. However, when representing legal rules in a RbS, some issues must be kept in mind in order to avoid some possible problems. Indeed, if one tries to encode considerable amounts of knowledge into a single RbS, the system may become inefficient, since it must search through a very large number of rules. Another possible disadvantage is that rules may not exactly implement the reasoning process used by an expert as no specific execution sequence can be dictated. Last but not the least, the open textured nature of The Law and reasoning mechanisms being used must be considered (Popple, 1990). In fact, when a judge decides on verdicts, he/she does not look only at the rules that apply in that situation. There is more information that influences the outcome, such as recurrence or intention of the defendant, information that is hard to model in such systems.

Nevertheless, RbS are broadly used in different fields of application, namely in insurance companies, banks, fraud detection, e-commerce and evidently in The Law. In conclusion, these systems implement a fairly simple and efficient way of modelling knowledge and expertise of a human practitioner in a well-defined field. Such systems can be particularly useful in the legal domain, once this field is rather rule-based, i.e., legal practitioners are usually comfortable about using such systems as they reflect their way of reasoning.

2.3 PROJECTS IN THE INTERSECTION OF AI AND THE LAW

In this section it is considered the role that Artificial Intelligence based techniques may play in the development of better ODR tools, focusing on practices that emulate the behaviour of human experts in a given field. A set of projects of reference are depicted and analysed to determine to which extent Artificial Intelligence tools can improve ODR.

Rule-based Legal Decision-making Systems (LDS)

This work dates from the eighties and was one of the first decision support systems to be developed (Waterman and Peterson, 1980). It's domain is liability law, which holds responsible product distributors and manufacturers for the injuries their products cause. The system embodies the skills and knowledge of a human expert in the shape of antecedent-consequent rules. The project has as objective the capture of the decision-making processes of attorneys and claim adjusters involved in product liability litigation in the shape of rule-based systems, so that the effects that changes in legal doctrine have in settlement strategies and practices can be studied. The authors formalized the strict-liability concept on ROSIE language, so that the defendant could or could not be considered liable.

The knowledge embodied in the system is divided into five layers: (1) the formal doctrine (given in terms of rules that emerge from the legally responsible and common law); (2) the informal principles (depicted in terms of rules that are not explicitly expressed in The Law, but are generally agreed upon by legal practitioners); (3) the strategies (where the authors code the methods used by legal practitioners to accomplish a given goal); (4) the subjective considerations (set in terms of rules that anticipate the subjective responses of people involved in the process); and (5) the secondary effects layer (set in terms of meta-rules that describe the rules interactions at the object level). In this project the authors concluded that despite the number of rules needed for shaping a given scenario, the rule-based model was feasible and suited for the legal domain.

EXPERTIUS

EXPERTIUS is a decision-support system that advices Mexican judges and clerks upon the determination of whether the plaintiff is or not eligible for granting him/her a pension (on the basis of the "feeding obligation"), and, if so, upon the determination of the amount of that pension (Caceres, 2008). The system comprises three main modules: a tutorial, an inferential, and a financial one. 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 presuppositions are defeated and which prevail. Finally, the financial module assists the user on the calculus of the pension values.

Expert knowledge is stated in terms of three interrelated layers. A base one that contains a representation of the expert knowledge. An intermediate one that denotes the decisions regulated by the law procedures. And a the third layer that keeps up a correspondence between the dialogues (written as conversation and measured in terms of a confrontation pattern) and the cases that arise as a result of the decisions taken at the intermediate level.

SmartSettle

Thiessen's SmartSettle stands for a decision support system that intends to find the middle ground among parties to settle disputes based on satisfaction functions acknowledged by them. Initially the parties declare to the system their tenure to each item under dispute, either using mathematical tools and/or by sketching it. The assigned preferences are however not static as they may change during the negotiation process in order to shape eventual changes in the objectives of the parties. Besides assigning their preferences, the parties must decide on what would be a constructive outcome for each one and try to join it on a single text.

On the one hand, during the negotiation process one may simply exchange messages. On the other hand SmartSettle may produce suggestions according to the current state of the case, which the parties may or may not accept. When parties reach an impasse, they may ask SmartSettle for an equal distribution of the items in dispute. It is important in this phase that the preferences are well defined once the allocation of items depends on that. A final document is then produced. All these steps are supported by a web site on which the parties log in, access their personal data and perform all the actions related with the negotiation process. SmartSettle is based on the doctoral thesis of Ernest Thiessen (Thiessen, 1993). This work resulted in a commercial ODR provider, whose president & CEO is Thiessen himself.

Family_Winner

The Family_Winner project is being developed by Zeleznikow and Bellucci and provides support on the Australian family law domain (Bellucci and Zeleznikow, 2001). In order to attain this goal, the system uses game theory and heuristics (Zeleznikow and Bellucci, 2003, 2004), relying on an algorithm that is an adapted version of the AdjustedWinner (Brams and Taylor, 1996) one. Working in a similar way as to SmartSettle, parties must provide as input to the system their tenure to each item under dispute, a value that denotes how much they hope for for each specific item. The system, according to these values, will try to assign the items to the parties, having in mind that each allocated matter may change the preferences of the parties on the other issues.

Once this is done, the parties are asked whether they agree or not with the distribution; if the answer is no, the system embarks on a negotiation of item-by-item, starting with the piece considered less controversial. The users are asked to rank the piece, so that it may be better distributed (Guasco and Robinson, 2007). This can be seen as a classical "divide the pie" problem. These problems are characterized by a fixed number of items with an associated value that must be divided by the parts. This process of decomposition and division goes on until there are no more items under dispute.

ALIS

The Automated Legal Intelligent System (ALIS) stands for a decision support system that will provide European citizens and private companies with a transparent, fast, secure and reliable access to legal data in the field of intellectual property rights, in Europe. The motivation in what concern this work is that the daily observation of legal systems in democratic countries or supranational institutions reveals severe problems of understanding, application and adhesion. There are different reasons for this.

Firstly, there is a considerable increase in the number of laws and regulations that make it more difficult to comply with the applicable legal and regulatory framework. This complexity often generates repetitions, lacks, and contradictions within the system itself. Furthermore, legal professionals may be overwhelmed by simple cases requiring time and effort that could be better allocated to solve more complex issues that cannot be, still, dealt with only by technology. The ALIS system aims at solving these problems by providing tools for regulatory compliance, alternative dispute resolution, conflict prevention, support in law making and scientific and technologic developments.

PERSUADER

PERSUADER (Sycara, 1993) is a framework for intelligent computer-supported conflict resolution through negotiation/mediation. The model integrates Artificial Intelligence and decision theoretic techniques to provide enhanced conflict resolution and negotiation support in group problem solving settings. This model has been implemented in PERSUADER, a computer program which operates in the domain of labour management disputes. The main objective of PERSUADER is to act as a mediator, facilitating the disputants' problem solving so that a mutually agreed upon settlement can be achieved. The PERSUADER embodies a general negotiation model that handles multi-agent, multi-issue, singe or repeated encounters based on an integration of Case-Base Reasoning and Multi-Attribute Utility Theory.

2.4 PRESENT AND FUTURE OF AI AND THE LAW

At the outset of this work, a thorough review of the state of the art in commercial ODR providers and research projects has been performed (Appendix B). Based on this, in this section the current state of ODR is critically analysed. If there is a first general conclusion that can be drawn it's that currently technology is not being exploited as it could be. Indeed, there is a limited use of IT by legal practitioners that use it mostly for text processing, office automation, case management (still at a rudimentary level), management of client and case databases or for electronic document interchange.

Technological tools are sill used mostly for text processing, office automation, case management, management of client and case databases or for electronic document interchange

Going into more detail, first of all most of the existing ODR implementations rely on traditional web forms for acquiring information, providing little to no assistance at all. Moreover, the visualization of the information is performed at a very low level, i.e., users see information in a way that is very close to how it is stored. This can indeed be pointed out as a major disadvantage as the lack of intelligent and intuitive interfaces can constitute a barrier for a wide acceptance and use of these systems.

There is also little automation of processes, even the simplest ones. This automation could boost the throughput of legal institutions and practitioners by automating simple tasks that do not explicitly need human intervention. Finally, another major drawback can be pointed out: very few systems use IT for knowledge management and goal achievement, i.e., the use of technologies able to handle complex models of legal information would improve information structuring and retrieval, improving the work of legal practitioners.

Another drawback identified in current systems is that there is no structured support for communication. In fact, current ODR tools rely on common communication technologies. It would be useful that specific tools for communication would be developed that included functionalities such as the exchange of structured information (e.g. ideas, solutions, proposals) or that even supported the exchange of additional context information (e.g. when parties communicate over character-based tools they have no access to modalities such as body language).

Technology still plays a secondary role in the ODR arena. In fact, the technologies with the higher degree of penetration are the ones required to implement the communication channels. Consequently, current ODR systems have little to no autonomy at all and are barely automated. In a few words it can be argued that, despite some innovative research projects, first generation ODR systems are still the rule. A research effort must thus take place in order to achieve the so-called second generation ODR. The path to follow relies in the use of intelligent techniques that can enhance ODR systems with notions as autonomy or proactivity.

However, progress in the field of Artificial Intelligence and The Law has been slower than expected. In fact, in the excitement of the early years, it was expected that computers would soon have the skills and the computational power to take over the role of judges and attorneys. This is still far from happening. Moreover, nowadays this is not even any more the main objective of research in this area.

The main reason against sitting computers in the chairs of judges and attorneys, mainly uttered by lawyers, is that doing so is morally undesirable. However, that alone would not hold back the research being done in the field; it would, at most, delay its implementation but not its development. One of the main reasons is that computers act as simple executors of rules while the legal field requires interpretation. Until computers are able to actually interpret norms and their context, they will not be sufficient to make judicial systems. John Searle formalized this restriction on the well-known thought test of the Chinese Room (Searle, 1980):

> *Suppose I am in a closed room and that people are passing in to me a series of cards written in Chinese, a language of which I have no knowledge; but I do possess rules for correlating one set of squiggles with another set of squiggles so that when I pass the appropriate card back out of the room it will look to a Chinese observer as if I am a genuine user of the Chinese language. But I am not; I simply do not understand Chinese; those squiggles remain just squiggles to me.*
>
> — John R. Searle, American philosopher (1932-)

Moreover, law is not straightforward and ambiguous. That is, the interpretation of norms frequently raises doubts among legal practitioners, leading to different and clashing interpretations and, consequently, different outcomes. Thus, at a first glance, one would conclude that we need a more specific definition of the norms, one that would lead to unambiguous interpretations. The problem is that the society is complex, with many conflicting values and norms of conduct. This task seems thus rather utopic.

Nonetheless, let us admit that such achievement is possible, that we can define norms to the point that their interpretation is straightforward. It is evident that this would only be achieved with a larger number of much more specific norms. Would it be efficient to handle such a complex legal system? Would it be feasible to develop computer systems to handle such complexity?

Another challenge that future research will face is related with the changing nature of the laws. Indeed, in civil law systems, the frequency of legislation changes is increasingly higher. Moreover, as the number of cases solved by courts in common law systems increases, more and more different cases can be considered when solving a new one. Thus, another major challenge will be to deal with increasing and ever changing amount of information.

From the technological point of view, for ODR systems that work in civil law domains (tendentiously rule-based), this means that whenever a legal norm changes someone will have to search the system for the rules or ontologies that implemented that norm and change them accordingly. Thus, there will be a growing effort to manage such systems and keeping them up to date without creating ambiguities. The same happens in common law domains, in which systems tend to be case-based. In these systems, the question is about whether a past case should or should not be considered after a clear trend of change in more recent cases. Here, there is also a growing effort to maintain a database of relevant information.

There are evidently many challenges to be addressed in the development of AI and The Law research. It is not clear if the development of fully autonomous software agents that can take the role of judges and attorneys will or not take place in a near future. Nevertheless, by aiming at this ambitious objective, researchers will continue to develop useful tools that will slowly but steadily improve the legal systems, making them more efficient and, ultimately, more accessible to people.

It is not clear if the development of fully autonomous software agents that can take the role of judges and attorneys will take place in a near future

This should be the goal of future AI and The Law research and is also the motivation behind this work: not to develop highly advanced and specialized systems that barely no one will use but to develop systems that can actually be used by individuals that have little to no knowledge at all about the legal field, essentially as support decision tools.

2.5 Summary

In this section the current state of the research on AI and The Law has been analysed. It can be concluded that current tools lack autonomy and still rely largely on the intervention of human experts. On the one hand there are delicate issues and scenarios in which this is understandable and acceptable. However, on the other hand, there are many scenarios in which simple tasks could be automated, which would release practitioners for higher-level ones and increase the throughput of legal institutions.

The main drawbacks of current approaches can thus be summarized:

- acquisition of information is performed in rudimentary ways, without support;
- the visualization of the information is performed at a very low level;
- there are no advanced tools for the structuring and retrieval of complex legal information;
- current ODR systems have no autonomy and still rely largely on the human;
- there is no structured and domain-dependent support for communication.

In order to achieve this goal several sub-fields of AI should be regarded as particularly interesting. Decision Support Systems and Expert Systems can be an interesting approach to generate possible solutions and strategies for problem-solving. Rule and case based approaches could also be interesting, specifically for information retrieval. Finally, in what concerns the representation and access to information, Intelligent Interfaces and Ontologies can play an important role towards the design of more intuitive and user-friendly platforms.

Despite an early view on this field that pointed to fully autonomous conflict resolution systems, the future now points more to autonomous decision support systems. In fact, it is broadly accepted that human practitioners should have the commanding role, although they could be supported in this as much as possible by technology in order to improve the efficiency and the quality of their work. Some research projects have been pointed out that look at the automation of several specific aspects, mostly focusing the division of assets while trying to maximize the satisfaction of the parties.

This section established the current state of the art, identified its main drawbacks and pointed out potential future lines. The following one builds on these conclusions to state the motivation of this thesis and realistically define a research plan and objectives that can contribute to the improvement of some of the negative aspects pointed out.

3 CONCEPTUALIZATION

Imagination is more important than knowledge.
— Albert Einstein

This section is dedicated to the conceptualization and formalization of the research plan. It starts by pointing out the motivation behind the choice of this research field, by defining the problem being addressed and by setting the main goals. Based on this and on the analysis of the state of the art performed, the section continues by detailing the main research hypothesis and objectives. Finally, the strategy and research plan that will guide the work are defined.

3.1 PROBLEM STATEMENT, MOTIVATION AND GOALS

Motivation is the feature that pushes an organism into taking an action toward the achievement of a desired goal. While the goal is not achieved, the organism defines, controls, and maintains a set of goal-directed behaviours. With this thought in mind, three key issues for this thesis are described in this section: the nature of the problem being addressed; the set of issues that constituted the spark for the beginning of this research; and the goals that directed it.

In a broad sense, this work is concerned with improving the current paradigms for conflict resolution, namely the ones based on technology. The problems identified in Chapter 1 can be organized in two main groups: (1) problems inherited from the current state of judicial systems and (2) problems related with the current approaches on alternative conflict resolution methods. For the problems pointed out in each group the main causes will be identified. Based on this, possible solutions will be pointed out which will constitute the guiding line for the remaining of the research work.

Concerning the first group, there is at this moment one major problem with the legal systems: they are very slow and very expensive. A list of causes and their relationships can be put forward, depicted graphically in Figure 3.

There is at this moment one major problem with the legal systems: they are very slow and very expensive

1. Uncooperative environments - Courtrooms are highly uncooperative environments in which parties have as main objective to gain as much as possible, at the expenses of the other side. These classical win-loose scenarios difficult the achievement of a mutually agreeable outcome, with consequences on cause 2);

2. Number of appeals - Defendants have the right to several appeals which is a manifestation of increasing rights but one of the major causes for the low throughput of courtrooms as each appeal adds to the amount of time needed and costs associated;

3. Increasing new cases - Currently every minor action is liable to end up in court, independently of the costs involved. Presently, a significant amount of processes involves small-value claims. Unfortunately, the time to conclude the process is not proportional to the value of the claim;

4. Increasing rights - Defendants have increasing rights which is the reflection of a positive evolution on the legal system but, on the other hand, also contributes to the aggravation of the problem being analysed, mainly through points 5) and 6);

5. Complexity of legal systems - Legal systems are becoming increasingly complex due to the increase in the number and complexity of rules, the existence of several ways to do similar things, the use of specialist terms and specialized language, the rapid appearance of new concepts or the increasing amount of information in each case, just to name a few;

6. Complexity of cases - Each legal process is becoming increasingly complex due to a need for a more precise specification of each word, decision or rule used. The emergence of new (legally challenging) concepts and the easiness with which proof is managed and presented electronically also adds to this;

7. Costs of legal practitioners - The costs of a legal process may be divided by the legal system and some of the parties but tend to be significant. They include the costs of the whole court staff (e.g. clerk, administration, security, legal staff, support staff, court reporter), solicitor fees, disbursements, lawyer services (generally paid hourly), among others;

Concerning problems related with the current approaches on alternative conflict resolution methods, six main problems have been identified:

1. Drawbacks of online communication - Depending on the communication mean, communicating online poses new challenges as the lack of contextual information such as body language or physiological responses threatens the efficiency of the communication process. Misunderstandings are frequent and are particularly threatening in the legal domain;

2. Info-exclusion - The lack of training in the generic use of technological tools poses an obstacle to the use of ODR mechanisms.

Complex or unfriendly user interfaces are other common deterrents;

3. Costs of ODR - The use of ODR is generally cheaper than litigation but it is not necessarily cheap;

4. Security, privacy, data protection, identity - These issues are not specifically related to ODR but rather result from the use of online communication methods and online tools. Nevertheless, they reasonably constitute one important concern in a field such as the legal one;

5. Rudimentary access to information - ODR tools frequently have data access methods that are tightly coupled with the data representation layer, providing no abstraction, making it difficult for the parties to efficiently understand and manage the data;

6. Rudimentary conciliation methodologies - Conciliation methodologies in ODR frequently place emphasis on human factors and depend significantly on the parties' decisions and judgement. This lack of a formal structure results in a sense of disorganization as well as in longer processes and less successful processes;

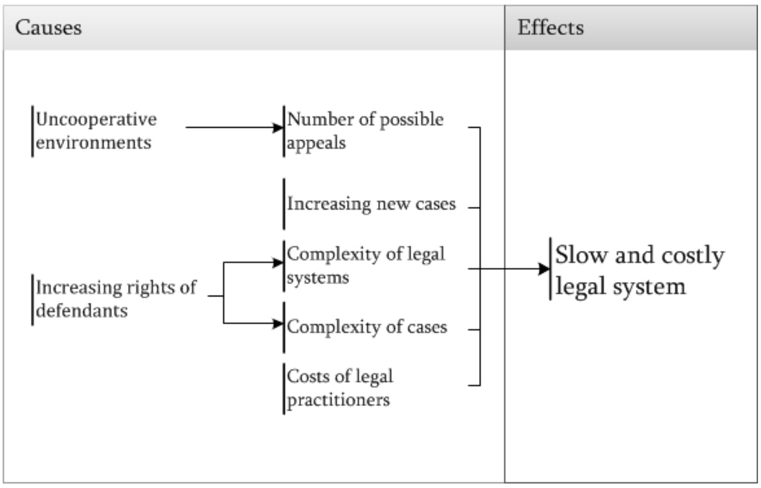

Figure 3: The main problem with litigation is currently related to its costs and inefficiency. The main causes identified are pointed out.

A listing describing the main causes for the problems pointed out is given below. Figure 4 depicts the relationship between the causes and the problems pointed out.

1. Lack of body language - Body language is one of the most important modalities of communication. Our gestures, our posture, our attitudes or our facial expressions provide the necessary framework for our interlocutors to correctly understand

the words we say. The lack of this information significantly hampers communication;

2. Lack of contextual factors - There are other modalities involved in communication besides the ones mentioned in 1). These modalities include the physiological response of our body, the speaking velocity, the tone of voice or the accentuation and often give away more information than words themselves. The lack of such modalities in communication negatively affects it;

3. Lack of training - Still today there are people who are ill at ease with the use of the typical technological tools that support ODR. This factor, as well as the lack of specific training or support in using a given tool, may make it difficult for parties to efficiently use an ODR solution;

4. Non-intuitive interfaces - Many ODR tools available nowadays still rely largely on traditional web forms and static pages as an interface for information. This may pose an obstacle, mainly when they are not adapted to the specific context or to the needs/characteristics of the users;

5. Cost of access to technology - There are costs associated to the use of technological solutions that, although not dependent on the ODR approach, are inherent to their use. These costs may include the costs of the hardware necessary as well as the costs of using the necessary telecommunication means (e.g. internet or telephone service provider);

6. Fee for using tools - Although ODR tools tend to be less expensive than traditional litigation, their use may be subject to a fee (fixed or per unit of time), which may amount to a significant sum. Moreover, the use of conciliation services mediated by a neutral also tends to have an associated fee;

7. Security issues related to online environments - The sheer fact of communicating online rises issues related to online identity (how am I sure that I am talking to whom I think I am), privacy (how do I know that my information won't be accessed by someone who should not do it), data protection (how to I know that my personal data will not be available to anyone else), among others. Although these issues are transversal, they are particularly worrying in the legal domain;

8. Low level access to data - Data representation models tend to be very close to the data layer. Representing information as it is stored, without abstraction methods, makes it difficult for parties to efficiently understand and manage it;

9. Lack of structure/formalization - Many ODR processes, despite taking place in online environments, are still largely unstructured. This informal approach, largely human-based, may result in difficulties for taking decisions, for managing information and for following some desirable line of attack or strategy;

10. Highly human-dependent interaction methods - Interaction methods in ODR tools still rely largely on natural language and on the rhythm or pace that the participants impose. Moreover, they focus too much on subjective issues such as the parties' personal views on the problem. This hampers the use of tools to support conciliation;

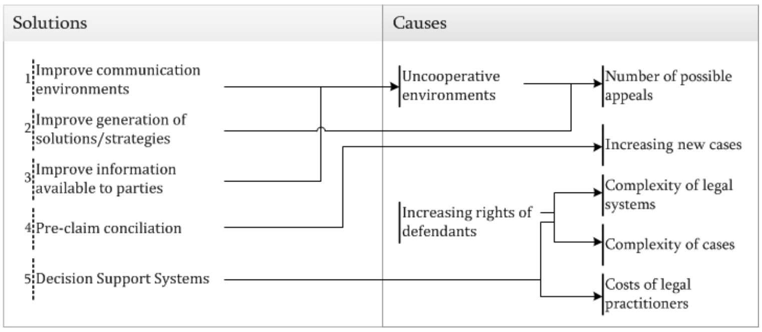

Figure 4: The main causes for the problems identified in the current ODR approaches.

The current main problems in litigation and ODR as well as their causes have been put forward. The issues mentioned were compiled from a revision of the current state of the art in the legal systems (depicted in Chapter 1) as well as from an analysis of several commercial ODR providers and research projects (detailed in Appendix B).

Based on this study of the state of the art, several solutions that would contribute to its improvement are pointed out. These solutions constitute the foundation for the definition of the goals of this PhD work. A brief description of each one is given below. Each solution proposed tackles one or more of the causes described, thus tackling the problems associated to such causes, as depicted in Figure 5.

1. Development of autonomous solutions - The development of technological solutions that can, to some extent, alleviate the work of human practitioners could have positive effects on the efficiency of the legal systems. On the one hand, it could allow for legal practitioners to work more efficiently and with increased quality by releasing them from monotonous and repetitive tasks. This would have an impact on the throughput of conflict resolution approaches, making them more efficient. Consequently, operating and specialist-related costs would be reduced. On the other hand, this would also allow to reduce the

apparent complexity of legal problem-solving, potentially making it more intuitive and supportive;

2. Improvement of Communication Environments - A positive and cooperative communication environment is paramount for the implementation of efficient conciliation approaches. The development of communication environments that look at contextual information such as body language as an actual communication modality may allow to improve the efficiency of the simple "message passing" approaches used nowadays. Moreover, such approach would allow neutrals working behind a computer to better understand the state of the parties, allowing a better management of the process (e.g. making a pause when a party shows signs of stress or fatigue). It would have a positive effect on the success rate of conflict resolution processes, under the assumption that people that communicate better will understand each other's fears and objectives better, will cooperate more willingly and will consequently be more likely to work together towards a solution. A particularly important subject here is the level of stress as an indicator of the inclination of a given party towards undesired behaviours such as hasty decisions, lost of interest, rudeness in communication or giving up on the process. Mediators could take profit of the access to this information in order to better manage the process;

3. Improve generation of solutions/strategies - Not infrequently, the main obstacles to conflict resolution reside within parties themselves. The generation of solutions, which is often the responsibility of the parties (especially in alternative conflict resolution approaches), is an example of a potential problem since most of us are not familiar with the action of devising a solution for a given problem (or we are simply not willing to), namely when it includes complex issues with complex relationships. The definition of a potentially successful strategy may be even more challenging. Nevertheless, these two features are central in the conflict resolution process. Tools to support them are thus necessary that could improve the quality of the solutions/strategies used so that parties can achieve better outcomes faster;

4. Improve information available to parties - In order for an individual to take good decisions he must be able to analyse different different courses of actions, weight their possible outcomes, decide on one over the others and learn from the consequences of that choice. This can nevertheless be challenging without the right amount of information with a minimum level of quality. The lack of these condition implies that individuals take decisions based on incomplete or poor information, being most

likely far from optimum. Tools for compiling information are thus needed that can provide concise and useful data for parties involved in a conflict resolution process to take realistic and weighted decisions, that will be more likely to be accepted by other parties;

5. Pre-claim conciliation - Currently many of the processes pending in courtrooms involve small-value claims, that can nevertheless make their way through different courts by means of consecutive appeals, making judicial processes inefficient. Pre-claim conciliation should be seen by the parties as a potential way to a faster and cheaper resolution of their differences. Successful alternative conflict resolution methods could thus not only improve the satisfaction of the parties by allowing them to solve their problems more efficiently but also contribute to the decrease the main problem of litigation currently;

6. Decision Support Systems - Decisions in the legal domain are frequently multi-issue, multi-value and multi-party with complex relationships between these variables, generally hardly to understand at first sight. Parties thus need support in analysing their possible choices in real time and understand the relationship between them and their most likely consequences. Tools to support parties in this could not only contribute to more satisfying decision processes but also to induce a more structured and formal way of reasoning and taking decisions in the legal arena.

As Figure 5 depicts, some causes are not addressed by the solutions proposed. In fact, the solutions described here include only the ones that will be tackled during this research work. Despite the need for additional solutions for many other problems, this selection was performed in order to focus the research work. The objective of defining a fully-functional architecture for an online dispute resolution platform also contributed to this choice.

Crucial to the definition of this specific line of attack was the group in which the candidate is inserted and his background on Intelligent Systems and Ambient Intelligence. This advances a hint about the strategy to be followed: the definition of "intelligent" solutions that can alleviate the problems identified. Moreover, the problems to attack were selected in order to address all the stages of conflict: definition, management and resolution. Based on this, in Section 3.2 the research hypothesis and objectives are defined.

3.2 RESEARCH HYPOTHESIS AND OBJECTIVES

The last figure depicted in the previous section starts to unveil both the aim and scope of this PhD work. In fact, most of the problems se-

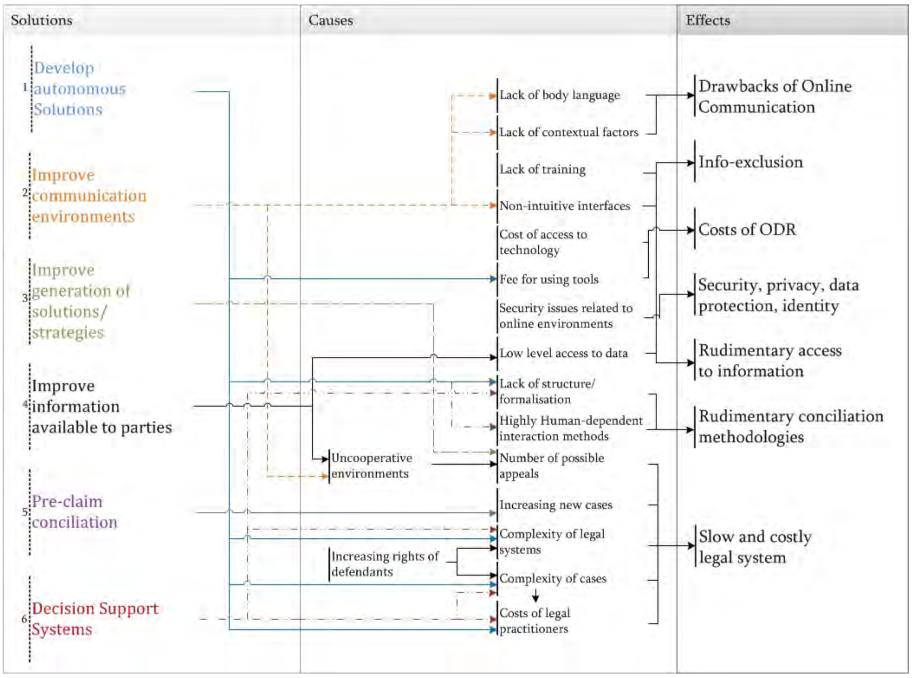

Figure 5: General view of the main problems identified in conflict resolution, the identified causes and the proposed solutions. Different colors and line shapes were added merely to increase the readability of the diagram.

lected to be tackled refer directly or indirectly to the decision-making processes of the parties in the different phases of the conflict resolution (e.g. initial building of knowledge, devising of possible solutions, negotiation).

This lays the foundations for the idea that the development of an integrated decision-support system that could assist the parties throughout all the process could result advantageous. Indeed, at this point and, given the understandable technological and ethical constraints (DeMars et al., 2010), it is more realistic to focus on decision-support rather than on autonomous decision-making, i.e., it is more feasible and more promising to support parties in taking better decisions than to devise solutions for automatic conflict resolution.

The notion that parties need access to valuable knowledge in order to take good decisions is fundamental

In this view of the problem, the notion that parties need access to valuable knowledge in order to take good decisions is fundamental. In that sense, such a decision-support system should be centred on providing the right knowledge at the right time.

Moreover, the decision-support system should be more than a traditional tool that responds to user's requests. Much in line with the Ambient Intelligence view, it should be "watching" the user, intervening if and when necessary, without an explicit request of the user. This leads to the last concept needed in order to clearly define the research hypothesis: the one of an environment for supporting conflict resolution. Here, the *environment* should be seen as a set of devices, software solutions and methodologies that result in a seamless group of services that assist the user in tasks related to conflict resolution.

Thus being, the main research hypothesis that will steer the work performed during this PhD is put forward: *Research hypothesis*

RESEARCH HYPOTHESIS

The unfolding of a conflict resolution process, from its definition to its outcome, under a knowledge-driven decision support environment, will enable richer and optimized decisions by the parties.

This research hypothesis touches different topics. In that sense, it is broken down into more specific and better defined research questions, to be addressed during the PhD work. Addressing and validating each of the different research questions will allow to determine the eventual validity of the research hypothesis. Each of the research questions, together with their objectives and validation methodologies is put forward.

RESEARCH QUESTION NO.1: *Multi-agent systems are a suitable way to implement an Online Dispute Resolution platform*

Multi-agent systems emerged as more than an approach to solve problems in a distributed way. In fact, software agents are seen as autonomous entities with the ability to take decisions towards the achievement of some individual or shared goal. In this work the candidate argues that an agent-based architecture, in which software agents contribute with their particular expertise towards the achievement of a common goal of the community, is a not only a suitable approach but a recommended one indeed.

OBJECTIVES: With the assessment of this research question, an actual agent-based architecture will be developed with a basic set of services that implement functionalities identified as central for a successful conflict resolution process. Focus will be placed on exploiting especial features of Multi-agent Systems such as autonomy or communication.

VALIDATION: This question will be considered valid if an agent-based architecture is formally defined that is validated by the scientific community. Besides, an actual prototype of the architecture must be deployed at the end of the research plan.

RESEARCH QUESTION NO.2: *Artificial Intelligence techniques can be used to develop information retrieval methods suited for the legal domain*

Information retrieval is a central problem in many knowledge-based domains, being the legal no exception. In the work developed in the scope of this research topic, Artificial Intelligence techniques will be used to develop information retrieval methods that can be used in the legal domain. Particular emphasis will be placed on making these methods loosely coupled with particular legal domains so that they can easily be adapted to others. In that sense, the information retrieved will revolve around concepts that are useful for the parties and that are more or less universal in the The Law, rather than very specific concepts with a narrow potential field of application. To accomplish this, well-known information retrieval methods will be considered, namely the ones based on rules and on cases as these share many concepts with the practice of The Law.

OBJECTIVES: The main objective of this research question is to determine an appropriate or the most appropriate method for information retrieval in the legal domain, given the requirements provided above. A working prototype will be developed.

VALIDATION: The work implemented in the scope of this question will result in a prototype whose performance will be assessed statistically and will define its validity. Moreover, the approach as well as the prototype will be documented in renowned scientific publications and thus be validated by scientific peers in the field.

RESEARCH QUESTION NO.3: *Artificial Intelligence techniques can be used to generate solutions for legal problems*

The generation of solutions for legal problems may be a very complex topic. As more complex problems need to be considered, both Human and machine-based techniques increase in complexity and make it harder to understand the relationship between all the variables and the solution.

In this research question it will be assessed the fitness of Artificial Intelligence techniques for generating solutions for legal problems. Once again, emphasis will be placed on domain-independence so that the range of potential applications in the legal field can be wide. Particularly, nature-inspired techniques will be considered since they have been used recently to generate solutions for complex problems in other domains, given that they can cover virtually the whole search space.

OBJECTIVES: It is the aim of this research question to develop a prototype for generating solutions for legal problems. The approach should be independent of the legal domain.

VALIDATION: The research question will be validated if a prototype can be successfully deployed that generates solutions for legal problems. These solutions must be valid within the particular legal domain addressed. The scientific community will also validate this approach by accepting scientific publications describing it.

RESEARCH QUESTION NO.4: *The incorporation of contextual factors can improve the efficiency of legal practitioners working online*

Our complex communication mechanisms involve, by nature, many more modalities than our mere words. When we communicate face-to-face we rely on our body language, our gestures, our tone of voice, our physiological responses or our facial expressions. These modalities are important not only for passing our message but also for perceiving feedback from our interlocutors. The problem is that this information is absent when we communicate online and must rely on words only. This may lead to misunderstandings in communication, being worrisome in conflict resolution processes.

Particularly, legal practitioners may find it difficult to accurately manage a process without interpreting the state of the parties. In this research question it is argued that the incorporation of such modalities of communication would increase the information available to legal practitioners, allowing them to take better decisions regarding the management of the conflict resolution process.

OBJECTIVES: When studying this research question, the main objective is to determine which contextual factors may be important for the legal practitioners and are currently absent from online communication means.

VALIDATION: This research question will be put forward for the scientific community in the legal field, properly grounded by justifying the utility of each modality and by giving plausible paths to the implementation of mechanisms for acquiring and providing such information. If the legal community accepts the conclusions of such enunciation, the topic will be considered validated.

Given the importance of the contextual factors in communication and, particularly, in conflict resolution processes (as already described in Chapter 1), a last research question is put forward concerning their assessment in real-time, using non-invasive methods. Given its complexness, it is divided into better defined ones:

RESEARCH QUESTION NO.5.1: *The conflict handling style of the parties can be assessed in a non-invasive way*

The conflict handling style represents each individual's way of reacting before a conflict. While some tend to be cooperative, others tend to be more competitive. Moreover, individuals may change their style depending on the parties they are in conflict with, on the nature of the issues being disputed or on other factors like past experiences or level of stress. Knowing the personal conflict style of an individual may point out to the mediator how he is likely to behave during the conflict resolution process, allowing a forehand planning. However, until now, only self-reporting instruments have been used for such purpose.

OBJECTIVES: The main objective of this research question is to develop a formal approach for evaluating the conflict resolution style, without using traditional self-reporting instruments such as questionnaires. The goal is that this analysis is performed in real-time and in a non-invasive way, so as to allow to detect not only the personal style but also changes in this style, that are frequent during the resolution process.

VALIDATION: The validation of this research question will be given in terms of the development or not of a functional prototype and by its validation by peers in scientific events.

RESEARCH QUESTION NO.5.2: *Stress influences in a significant manner our interaction patterns with technological devices*

Stress is a manifestation of the effects of external demands on our body and our mind. Its effects influence virtually all of one's actions. Knowing the level of stress of individuals may be important in many domains, namely when managing teams of collaborators working in sensitive domains such as healthcare, finance, insurance or air traffic control. This knowledge may allow a better management of the team, resulting in an increase in the productivity and quality of work. In the legal context it would also be interesting to have a measure of the stress of the individuals participating in the conflict resolution process so as to better manage it (e.g. making a pause when stress escalates).

OBJECTIVES: In the pursuit of non-invasive ways to measure stress, the main objective of this research question is to find out if stress has or has not a significant influence on the way one individual interacts with technological devices. An eventual positive conclusion on this will lay the path to topic 6.3.

VALIDATION: In order to validate this approach an experiment will be set in which the effect of stress on interaction patterns will be studied by comparing statistically the behaviour of individuals under and without the effect of stressors. Statistical hypothesis tests will be performed that will account for the validity of the research question.

RESEARCH TOPIC NO.5.3: *It is possible to accurately measure the influence of stress in a non-invasive and non-intrusive way, in real time*

The ability to assess the level of stress of an individual in real-time is very important to allow a conflict-manager to plan ahead and avoid the worsening of the interpersonal relationships. While this may be fairly intuitive to perform face-to-face, it results very challenging or nearly impossible using current online communication means. Here, it is considered the hypothesis of measuring stress in a non-invasive and non-intrusive way, by analysing features of the interaction patterns of individuals with technological devices.

OBJECTIVES: The main objective of this research question is to define a personalized and multi-modal model for assessing the level of stress of individuals from the analysis of their interaction patterns. This approach must be non-invasive, non-intrusive and work in real-time.

VALIDATION: This question will be validated through the implementation of a prototype for stress assessment using a model validated by the scientific community.

The research questions that will be addressed during this work have been described, together with the respective objectives and validation methodology. In the following sections the research methodology, strategy and plan are detailed.

3.3 STRATEGY AND RESEARCH PLAN

To accomplish the objectives enumerated, the Action-Research methodology is followed (Somekh, 2005). This methodology starts by identifying the problem so that a hypothesis can be formulated on which the development is based. Subsequently, the information is recompiled, organized and analysed, continuously building a proposal for solving the identified problem. Finally, one can make its conclusions based on the results obtained during the investigation. For this research model to be followed, six complementary stages are defined to achieve the objectives stipulated. The stages defined are described:

- Specification of the problem and its characteristics;
- Constant and incremental update and review of the state of the art;
- Idealization and gradual and interactive development of the proposed model;

- Experimentation and implementation of the solution through the development of a prototype;
- Result analysis and formulation of conclusions;
- Constant diffusion of knowledge, results obtained and experiences in the scientific community.

This research methodology guides the strategy to follow during the implementation of the research plan. A strategy can be seen as the group of activities implemented in order to achieve a given objective.

For each research question described in Section 3.2, potential approaches/techniques from Artificial Intelligence are analysed. The fitness of each approach analysed is determined based on the objectives and validation method of the topic as well as on the characteristics and scope of the approach. Based on this, a first selection is performed in order to obtain the set of approaches that may apply for the specific problem being deal with.

If there are indeed approaches that may be used in the context of the topic, a specific and reduced set of these approaches will be implemented in order to determine the best one, by evaluating the results of their application. If the results are satisfactory, conclusions can be drawn about the validity of the topic. Otherwise, potential approaches may be re-analysed with different constraints. This strategy admits the possibility of non-existence of a suitable approach based on Artificial Intelligence.

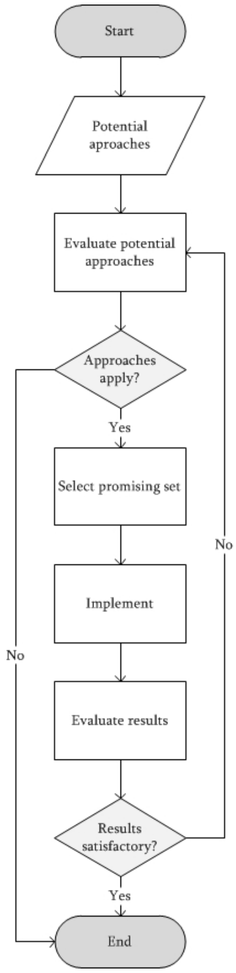

Figure 6

Figure 6 depicts a flowchart describing the research strategy followed. This strategy is followed for each of the main tasks identified that involve some kind of development or experimentation. In total, nine high-level tasks are identified in the research plan. Of these, six involve development/experimentation or analysis of results. The tasks planned for this research work are described:

1 - STUDY OF THE STATE OF THE ART

TIME-FRAME: 01-04-2010 – 30-09-2010

OBJECTIVES: With the execution of this task, the candidate aims to acquire a thorough knowledge about the current state of conflict resolution, concerning both traditional litigation as well as alternative methods. This includes the recognition of the main advantages and disadvantages of current approaches and the identification of possible future paths. This will constitute the base for a solid research.

RESULTS: By the end of this task the candidate should be able to identify the main issues in the intersection of AI and The Law and argue about them.

DELIVERABLES: At lest one scientific document will be published describing the current state of the art of the intersection of AI and the Law.

2 - DESIGN OF THE ARCHITECTURE

TIME-FRAME: 01-10-2010 – 31-03-2011

OBJECTIVES: This task aims to evaluate the main requirements for an agent-based architecture for conflict resolution, given a set of target functionalities. Based on this evaluation, a specification for a group of agents will be defined that will later implement the intended functionalities.

RESULTS: The main result of this task will be the definition of an agent-based architecture for supporting a conflict resolution platform.

DELIVERABLES: By the end of this task, the specification of an agent-based architecture for conflict resolution will be provided and made available in scientific publications.

3 - INFORMATION RETRIEVAL

TIME-FRAME: 01-04-2011 – 31-07-2011

OBJECTIVES: During this task it will be developed part of the agent-based architecture, namely the modules that concern the retrieval of information in an initial stage of the conflict, in order to inform the parties of the most important issues in their conflict.

RESULTS: As a result of this task a set of software agents will be developed that will implement information retrieval algorithms, incorporated in the architecture of the conflict resolution platform.

DELIVERABLES: By the end of this task several software agents will be added to the architecture of the conflict resolution platform, constituting the prototype of the information retrieval module. Moreover, this module will be detailed in scientific publications.

4 - GENERATION OF SOLUTIONS

TIME-FRAME: 01-08-2011 – 30-11-2011

OBJECTIVES: In this task the main objective is to develop mechanisms for generating possible and valid solutions for a given case, with specific constraints and requirements.

RESULTS: This task will result in the development of an agent-based module for generating solutions for a given case.

DELIVERABLES: By the end of this task a software prototype will be delivered that generates solutions for conflict scenarios. This will be detailed in scientific publications.

5 - GENERATION OF STRATEGIES

TIME-FRAME: 01-12-2011 – 01-04-2012

OBJECTIVES: This task aims at the development of mechanisms that can generate strategies for conflict resolution. Particular emphasis will be placed on generating strategies for the cases whose parties are unable or unwilling to do so and thus constitute an obstacle to the resolution themselves.

RESULTS: This task will result in an agent-based module that generates strategies for solving conflicts.

DELIVERABLES: By the end of this task, a prototype of a module for generating solutions for conflict scenarios will be delivered and included in the architecture. Moreover, this will be accompanied by publication(s) on the scientific community.

6 - STUDY OF CONTEXT INFORMATION

TIME-FRAME: 01-08-2011 – 30-03-2012

OBJECTIVES: During this task specific research work will be conducted in order to identify the main contextual issues that may have an important role in alternative conflict resolution but are, at present time, being ignored. Particular emphasis will be placed on the contextual issues in communication processes.

RESULTS: This task will result in a thorough knowledge about the major contextual factors in a conflict resolution processes.

DELIVERABLES: When this task is concluded one model will be delivered detailing the role of contextual factors in conflict resolution. This will be detailed in at least on scientific publication.

7 - BEHAVIOURAL ANALYSIS

TIME-FRAME: 01-04-2012 – 30-09-2012

OBJECTIVES: This task aims at the development of a software module for acquiring context information in a non-invasive and non-intrusive way in the context of a conflict resolution process. Given the complexity of the domain, only some types of information will be considered for this module, selected according to its potential advantages for the process.

RESULTS: The main result of this task will be a software module for acquiring context information that can improve some aspects in a conflict resolution process.

DELIVERABLES: By the end of this task, a prototype will be delivered that can acquire important context information in a non-invasive way and make it available for the remaining modules of the conflict resolution platform.

8 - ANALYSIS OF RESULTS

TIME-FRAME: 01-10-2012 – 30-03-2013

OBJECTIVES: The main objective of this task is to perform a critical analysis of the results achieved during the implementation of the work plan, namely in what concerns the efficiency of the methods for information retrieval, generation of solutions or acquisition of context information. Whenever possible, this analysis will build on statistical methods that can validate it.

RESULTS: This task will result in an understanding of the fitness of each of the solutions implemented for solving the problems that they were designed for.

DELIVERABLES: The conclusions achieved will be made available for the scientific community through their publication in scientific venues.

9 - DISSEMINATION OF RESULTS

TIME-FRAME: 01-10-2010 – 30-03-2013

OBJECTIVES: This task aims to make the conclusions of the work developed during the PhD available for the scientific community. In order to achieve this, the work will be documented in scientific publications such as conference proceedings, book chapters and international journals. Moreover, it will also be presented in scientific venues and discussed with scientific peers. This also has the objective of being an important feedback for the research work.

RESULTS: The main result of this task will be the dissemination of the conclusions of this research at an international level.

DELIVERABLES: A list of the scientific publications in the context of this work will be made available.

3.4 SUMMARY

In this section, the process of conceptualizing the research plan was detailed. This included the statement of the problem and the motivation behind the will to do research on these topics. It continued by clearly pointing out the main research hypothesis and objectives. Based on this, the research plan was defined in terms of the main tasks identified, their interdependence and their approximate duration. The strategy to guide the research plan was also set out.

This section closes the preamble of this thesis. So far, the state of the art on conflict resolution has been described, with a particular focus on the main drawbacks identified. On the other hand, fields of Artificial Intelligence research have been analysed in search for promising approaches that can deal with these drawbacks. Finally, the plan that proposes to use such approaches to deal with the identified problems has been set out. In the following chapter focus is moved towards the description of the actual development and implementation work, detailing the accomplishment of most of the tasks defined.

Part II

DESIGN AND IMPLEMENTATION

The second part of this thesis is devoted to the actual implementation of the necessary steps to validate the research hypothesis postulated in the first part. It is divided into four chapters. The first one is dedicated to the description of the agent-based architecture built. In the second chapter, the knowledge retrieval methods analysed and developed are described. The third chapter is dedicated to supporting the negotiation process, particularly in what concerns the generation of valid solutions. Finally, the fourth chapter details the use of behavioural analysis to build a stress-aware conflict resolution environment. In its essence, this second part of the thesis defines the main layers of the proposed conflict resolution platform.

4 ARCHITECTURE OVERVIEW

> There are two ways of constructing a software design;
> one way is to make it so simple
> that there are obviously no deficiencies,
> and the other way is to make it so complicated
> that there are no obvious deficiencies.
> The first method is far more difficult.
>
> — Charles Antony Richard Hoare

In the present time, many different ODR methods may be considered, ranging from negotiation and mediation to modified arbitration or modified jury proceedings (Goodman, 2003). In this context, it must be considered the existence of legal knowledge-based systems, appearing as tools that provide legal advice to the disputant parties and also systems that (help) settle disputes in an online environment (Vries et al., 2005).

A gradual tendency to foster the intervention of autonomous software agents is clearly assumed in this work

In devising such systems, the Katsch/Rifkin vision of the four parties in an ODR process must be taken into consideration: the two opposing parties, the third neutral party (e.g. mediator, arbitrator) and the technology that works with the mediator or arbitrator (Katsh and Rifkin, 2001). A gradual tendency to foster the intervention of autonomous software agents is clearly assumed in this research work, with these acting either as decision support systems or as "real" electronic mediators. This approach is wittingly close to the second generation ODR, proposed by Peruginelli and Chiti (Peruginelli and Chiti, 2002), as it is guided by three main lines:

- The aim of such systems does not end by putting the parties in contact, but it consists of actually proposing solutions for solving the disputes;
- The human intervention is reduced and the software intervention enhanced;
- These systems act through the use of autonomous software agents.

The consideration of this wider role for software agents is based in the use of artificial intelligence techniques and information and knowledge representation tools. Yet, merely representing facts and events is not enough for dispute resolution. In order to have useful actions performed by software agents it is required that they not only know the terms of the dispute, but also the rights and wrongs of the parties and foresee the legal consequences of facts and events (Peruginelli and Chiti, 2002).

The complexness and width of the legal domain, as well as the particularities of each sub-field of The Law, must also be considered. Thus, the development of fully functional technology-based tools that can assist parties may be a challenge. In fact, until now, only a very limited number of ODR tools with some real degree of autonomy have been developed (some have been detailed in Chapter 2). Moreover, these tools are very domain-dependent, focusing on very specific and well defined issues. There is thus the need to develop architectures that are based on high-level concepts, common to every sub-field of law, so as to apply them to a broader range of problems. Such architectures should be abstract enough to cover multiple legal domains but at the same time possess the specific knowledge of each specific domain.

In this chapter a proposal of such an architecture is put forward, henceforth designated as UMCourt . It is based on the Jade agent platform. This agent platform is aimed at the simplification of the development of multi-agent systems by providing support for agent communication and registry (Bellifemine et al., 2007). The platform runs on a Java Virtual Machine, thus agents are also programmed using the Java language.

The objective of UMCourt is to enable a range of services targeted at assisting the disputant parties, independently of the domain of their dispute. The architecture is abstract in the sense that it encompasses many concepts that are common to virtually all the legal domains. However, each domain has its particularities.

In that sense, the architecture presented here covers a broad range of domains but also contains the specificities that allow it to operate correctly in each of those domains, in a transparent manner. This is the main goal of the work detailed in this chapter.

To accomplish it, abstract entities that are present in conflict resolution regardless of the domain of the dispute will be identified. As an example, to a certain point, a negotiation process will always be a sequence of rounds in which, in each round, each party states his opinion about the current proposal on the table. That is, this process goes on like this independently of the domain of the negotiation. The same happens with certain concepts: independently of the domain of the dispute, the concept of best or worst possible scenarios apply.

The first step in the development of this abstract architecture is thus to identify which concepts and processes apply in different legal domains and model them. It must however be kept in mind that although these concepts and processes have the same meaning for legal domains, they must be implemented differently. However, this should be transparent to the user. The final step is the definition of the minimum set of software agents developed to support the architecture.

4.1 IDENTIFYING USEFUL ABSTRACT CONCEPTS

While identifying the abstract concepts for a multi-domain ODR tool aimed at assisting the parties there is the need to determine, in the first place, which information would actually be useful for the parties. Only after this should work move on to formalize the concepts. In a first instance, it would be interesting for a party to determine to which extent it is reasonable to engage in an alternative dispute resolution process. That is, the party should ask himself "will I reach a better outcome using an alternative dispute resolution process instead of litigation?"

In a first instance, it would be interesting for a party to determine to which extent it is reasonable to engage in an alternative dispute resolution process.

The concept can be defined in terms of the BATNA - Best Alternative to a Negotiated Agreement, or the possible best outcome "along a particular path if I try to get my interests satisfied in a way that does not require negotiation with the other party" (Notini, 2005). This concept is abstract as it is useful for any dispute and it is useful for parties as it will, at least, contribute to the acknowledgement that an agreement may be disadvantageous (Klaming et al., 2008). In fact, by knowing their BATNA, parties would on the one hand become better protected against agreements that should be rejected and, on the other hand, in a better condition to reach an agreement that better satisfy their interests (Vries et al., 2005).

On the opposite side one should also consider the notion of a WATNA, or the Worst Alternative to a Negotiated Agreement (van Steenbergen, 2005). The WATNA aims to estimate the worst possible outcome along a litigation path. It can be quite relevant in the calculation of the real risks that parties will face in a judicially determined litigation, imagining the worst possible outcome for the party.

With these two abstract concepts, the party is aware of the best and worst scenario if the dispute is to be solved in a court.

It is also interesting for a party to analyse the space between the BATNA and the WATNA as a useful element to be taken into account for taking decisions. Of course, the less space there is between BATNA and WATNA, the less dangerous it becomes for the party not to accept the agreement (unless, of course, their BATNA is really disadvantageous). A wider space between BATNA and WATNA would usually mean that it can become rather dangerous for the party not to accept the ODR agreement (except in situations when the WATNA is really not inconvenient at all for the party).

It can thus be argued that knowledge about the space between the BATNA and the WATNA is very important and is also an abstract concept, independent of the domain. This concept is close to the Zone of Possible Agreement (ZOPA) proposed by Raiffa in 1982 (Raiffa, 1985). It is the zone where an agreement can be met that is acceptable for both parties.

It would also be interesting for a party to understand the region of this space in which an outcome is more likely. That is, if the parties are to solve the dispute through litigation, what is the most likely outcome? In fact, sticking only with the BATNA and the WATNA may not be realistic as these are usually not the most likely outcomes but merely informative values that establish boundaries. Thus, an informed party should also consider the concept of MLATNA – Most Likely Alternative to a Negotiated Agreement (van Steenbergen, 2005).

Using the same arguments, it can also be argued that the existence of metrics that measure the probability of each possible outcome could also be extremely useful for a party. Thus, it is also considered the concept of probable case: a possible outcome with an associated value of likeliness.

Concluding, several abstract concepts that are important for parties can be considered in the development of an ODR tool: (1) the BATNA; (2) the WATNA; (3) the ZOPA; (4) the MLATNA and (5) the probable cases. These concepts establish what is identified in this work as the minimum set of information that a party should consider prior to getting involved in a litigation or alternative conflict resolution process.

4.2 IDENTIFYING ABSTRACT PROCESSES

Besides abstract concepts, abstract processes in conflict resolution should also be considered. The main aim is to point out processes that follow the same process model, independently of the legal domain. Several abstract processes have been identified and implemented during the research work. They are detailed in the following sub-sections.

4.2.1 *Information Retrieval*

The retrieval of information generally involves a given number of predetermined tasks such as querying a database, analysing the results or possibly filtering or sorting them (Korfhage, 1997). In this subsection the proposed abstract model for information retrieval in the legal domain is presented Figure 7. It is further detailed in Chapter 5. The main aim of the process is to retrieve a number of past cases that are, to some extent, related to the current case of the parties. It then computes the similarity of each past case to the current case (which tells a party how likely the outcome of a past case is on their case) as well as the utility (which tells the party how much they would win/loose if an outcome such as the one of the past case were to occur).

The Information Retrieval process starts with the reception of a request. This request is analysed as it may be accepted or rejected (e.g. unavailability of service, requester not authorized). If it is rejected the process ends. If it is accepted, a query to the database is generated. This query can be seen as a pre-selection of potentially interesting cases, i.e., cases that are expected to have some degree of similarity. Given that there is no control on the number of cases that are retrieved, the pre-selection returned from the database is analysed.

If no cases are retrieved, the process ends unsuccessfully as it cannot currently provide a solution for the problem. If, on the other hand, there are too many or too few cases, the database query is rebuilt so as to make the selection rules more restricted or more relaxed, respectively. The main objective of this step is to obtain the optimum amount of information that parties can analyse efficiently. This has an effect on the similarity of the cases retrieved, as detailed in Chapter 5.

When a suitable selection of cases is achieved, a value of similarity and of utility is computed for each one. Similarity is obtained by comparing key characteristics of the cases, whose resemblance will contribute, with a given weight, to the overall measure of similarity. In order to compute the utility of a past case the structure of the solution is changed with the values of the current case, allowing to compute a value of the solution for the characteristics of the current case. This constitutes the adaptation phase.

Finally, when these steps are completed, a result is returned to the user, containing past cases that are potentially similar, a measure of this similarity and a value of utility. This will allow each party to understand to which extent a given outcome would satisfy them, as well as to understand how likely that outcome is.

4.2.2 *Generation of Solutions*

The process of generating solutions can be implemented in an abstract way, given that the specific rules for the domain are provided that allow for valid and meaningful solutions to be created. In order to develop this abstract model for generating solutions for conflict resolution, a nature-inspired approach was followed, based on the concept of Genetic Algorithms (Davis, 1991). Although a much more detailed definition of the process is provided in Section 6.1, the process model is depicted in this section (Figure 8).

The process starts with the reception of a request for generating solutions for a given problem, with associated data describing the problem and the domain. At this point, the request may be rejected if it is malformed (e.g. necessary information missing) or refers to an invalid/unknown domain. If the request is valid, the model proceeds to reading the necessary settings from the database, specific to the domain of the problem. With these settings, the model generates the

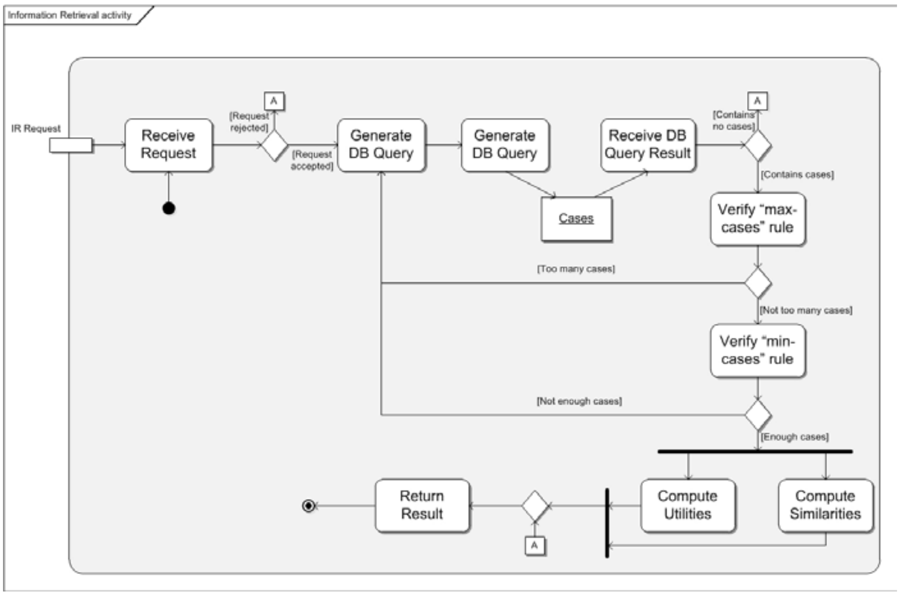

Figure 7: UML 2 activity diagram depicting the sequence of main activities needed to implement the process of retrieving information. The processes of computing utilities and computing similarities will be detailed in Chapter 5

validity rules (the rules that state whether a given solution is valid in the domain) and proceeds to the initialization of its main part.

During the initialization, the model defines the settings that will control the remaining of its execution (e.g. the type and weight of the genetic operands, terminating conditions, size of population). Moreover, it also generates a random population of the pre-determined size.

Having done this, the model moves on to measuring the fitness of each of the individuals. If there is a number of solutions with a minimum level of fitness, the model terminates successfully by returning the solutions generated. Otherwise, the model goes on to a process of selection.

In this process, the main objective is to select the most promising solutions: the ones that are likely to give birth to even better ones. Having concluded it, the next step is to apply the genetic operators on the selected solutions. Each of these operators will take some solutions to generate new ones, which will be validated prior to continuing. When a new set of valid solutions is obtained, the terminating condition is verified. If a maximum number of generations has been reached, the process terminates by returning the best solutions obtained so far. Otherwise, the model returns to the computation of the fitness, repeating the whole process.

When the model reaches the terminating condition, it is expected to have generated a number of valid satisfactory solutions. The solutions may however have low values of utility, which is generally due to bad terminating conditions not allowing the model to evolve for the necessary time.

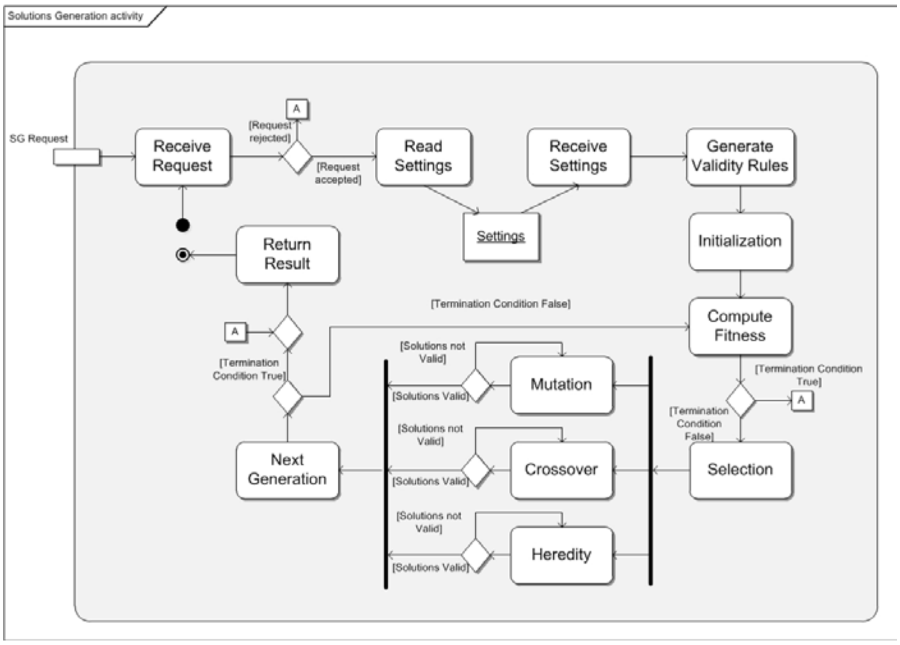

Figure 8: UML 2 activity diagram depicting the sequence of main activities needed to implement the process of generating solutions, using a nature-inspired approach. All the steps are further detailed in Section 6.1

4.2.3 *Negotiation*

Negotiation assumes a central role in any conflict resolution process. It is through negotiation that parties cooperate through the sharing of successive proposals and counter-proposals, in order to reach a mutually satisfactory outcome. A typical negotiation process can be seen as a sequence of rounds in which parties exchange proposals, with a potential intervention of a negotiator that may control the process and contribute with proposals as well.

In that sense, an abstract negotiation model can be developed to be used in the legal domain. Figure 10 depicts such a model. It starts by receiving a request and verifying its validity. Hence it can be accepted or rejected. Being accepted, the process moves on to allow all the parties to register. By doing so, each party agrees into engaging in a negotiation process with the given characteristics and provides the necessary information (e.g. personal information, personal objectives, personal appraisal of each issue).

Once all the parties are registered, the model receives a list of possible solutions (obtained from any of the two methods described in Section 4.2.2 and Section 4.2.1) and informs them about the issues being negotiated in the following round: number of issues, type (e.g.

monetary, real estate, divisible) and description. Then, the model, on the role of an electronic mediator, proposes what is considered the best solution, i.e., the solution that maximizes the sum of the gains of the parties. After publishing the proposal the model waits for all the parties to respond. Parties can do it in several ways: (1) ignore/reject ; (2) accept and (3) exit the process.

When all the answer are received, the first step is to verify if some party abandoned the process, in which case the negotiation ends unsuccessfully. Otherwise, the model checks is all the parties agreed on the proposal, in which case the process ends successfully. If that is no so, the model asks the parties for counter-proposals, which may be a minor modification of the original proposal or an entirely new one. After receiving all the proposals (if any) the models tries to find intersection points in which the parties may have agreed (e.g. two parties suggesting the attribution of the same item to the same party).

If there are indeed intersecting points, these will be marked as "agreed upon" and removed from the list of pending issues. This being the case, parties are informed of the new list of issues and the process repeats. However, if there are no intersecting points, the model itself will propose the following solution of the list of possible solutions compiled. If there are none, the process ends unsuccessfully. If there are, the model publishes it and waits for the replies of the parties, repeating the process. The model will fail when parties do not agree on a solution before one of the final states and will succeed if they do.

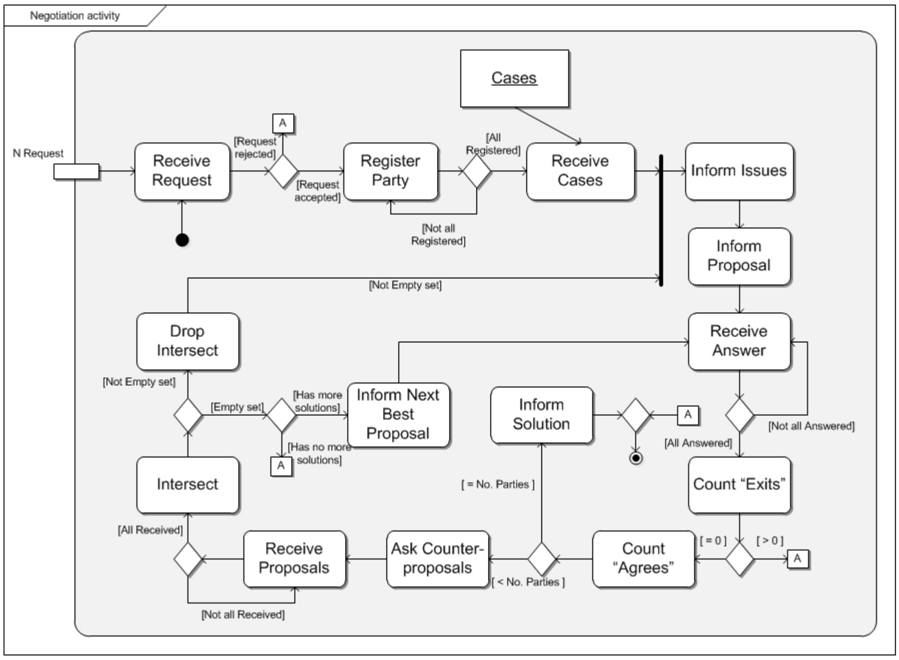

Figure 9: UML 2 activity diagram depicting the sequence of main activities needed to implement the negotiation process. All the steps are further detailed in Chapter 6

4.2.4 *Generation of Strategies*

One common challenge in ODR concerns the creation of effective strategies for problem-solving (generally the definition of a line of attack in a negotiated procedure). Frequently parties are unable (because they lack the training or are unaware of all the details of the dispute) or not motivated (because they are avoiding dealing with some sensitive issues) to think on possible solutions or strategies for solving the conflict. When this happens, the conflict resolution process risks stalling.

The main aim of this abstract model is to provide the framework for an assisted negotiation process which supports parties in going through a negotiated process. It is similar to de model described in Section 4.2.3 in the sense that it is a negotiation model. However, while the previous one only suggested solutions, this one tries to guide the parties through a path (strategy) that has already worked in the past for similar disputes.

The model starts similarly, by analysing the validity of the request and registering all the parties. Having concluded the registration it will receive a list of past successful negotiations that concerned a similar case. These negotiations (here denominated "strategies") are sorted according to the degree of similarity of their case. The model hence starts by informing the parties about the issues in dispute and proposing the first solution of the most similar strategy.

When all the parties have answered, the model checks for waivers. If no party gave up, the model counts how many of the parties have agreed on the proposal. Three scenarios are possible. (1) If all agreed, the current proposal is the solution for the negotiation and the process ends successfully. (2) If the majority of parties agreed, this is considered as a positive sign and the model suggests the following proposal on the current strategy, repeating the process. If no more proposals exist on the strategy, the model tries to move to the next best strategy. (3) If less than half of the parties agreed, the system will also move to the next best strategy.

Whenever the model moves to the next strategy, it verifies if there is actually another strategy on the list. If there is none, the process ends unsuccessfully. If there is, it will move on to suggest the first proposal of the new strategy, repeating the process.

This model will succeed if all the parties agree on one solution proposed. It will fail if some party abandons the process or the last proposal of the last strategy is not accepted by all the parties.

4.2.5 *Generic Conflict Resolution Model*

The high level abstract processes that were described in the previous sub-sections can now be composed in order to design a generic con-

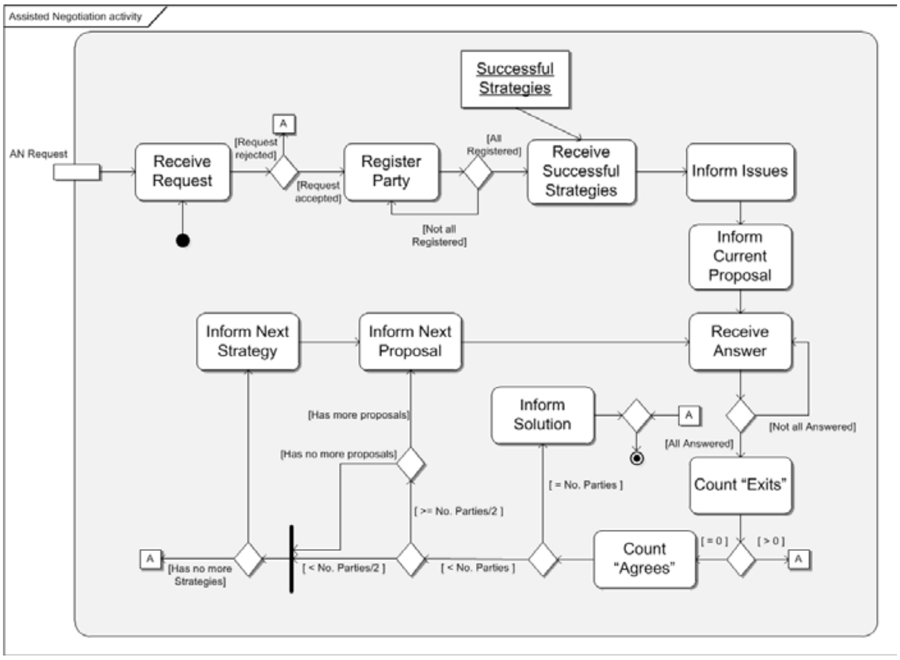

Figure 10: UML 2 activity diagram depicting the sequence of main activities needed to implement the assisted negotiation process. This model is especially suited for scenarios in which parties are unable or unwilling to generate solutions.

flict resolution model (Figure 11). This model aims to accompany the conflict resolution process since its beginning to its end. When the process starts, parties register providing some personal information (e.g. personal objectives, background information). After this, parties provide and agree on the details of the case (e.g. the number of issues being disputed and their attributes).

Based on this, the model computes a set of useful information that will allow the parties to fully understand the most important aspects of their dispute. This information includes the abstract concepts identified in Section 4.1 or past similar cases. At the same time, the model generates a set of possible solutions (with the associated values of utility) and a set of strategies that may result in a successful outcome.

The information compiled is provided to the parties so that they can be informed before the actual negotiation starts. Having done it, parties must choose whether they intend to follow a regular negotiation process (in which the responsibility of cooperating to reach a solution is theirs) or an assisted one (in which the model tries to guide the parties into a successful outcome). In the first case, if the process fails, parties can still try an assisted negotiation process. At the end of any of these processes, independently of the result of the negotiation, the model ends, with the success being the success of any of the negotiation processes used.

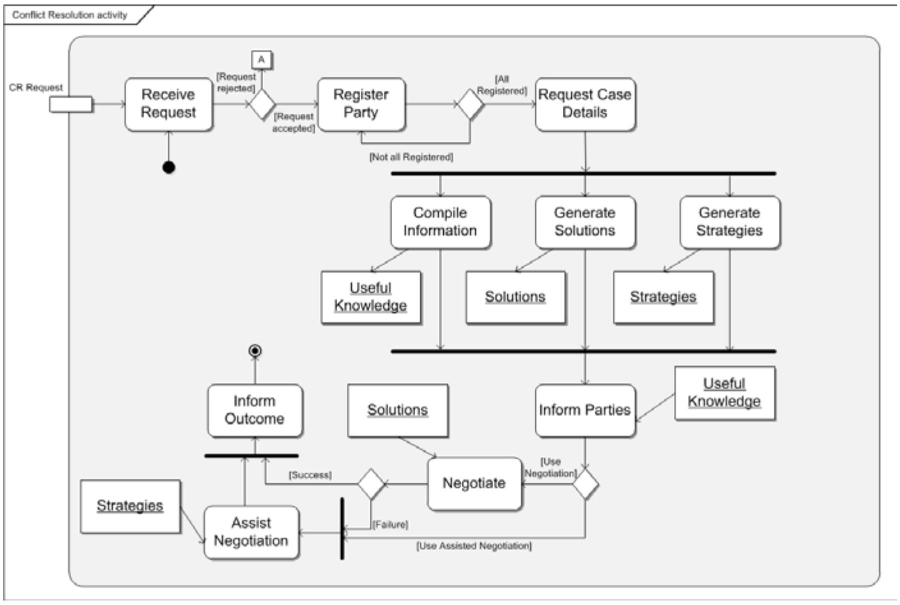

Figure 11: UML 2 activity diagram depicting the sequence of main activities needed to implement a generic conflict resolution model. It can be seen as a high level composition of the previously depicted models.

4.3 SOFTWARE AGENTS

The activities depicted in the previous section are implemented by a number of different software agents through the composition of their services. Agents are classified as *Main Agents* or *Secondary Agents* according to their degree of autonomy.

The activities depicted in the previous section are implemented by a number of different software agents through the composition of their services.

Main agents have an increased degree of autonomy and knowledge about the process models that implement abstract activities. In that sense, these agents can control secondary agents and compose their services in order to implement high level activities. Moreover, they can take decisions in run-time targeting the optimization of the processes implemented.

Secondary agents are simpler and have little or no autonomy at all. Their main aim is to implement specific and relatively simpler processes. Main agents can then request services from these agents.

This organization results in a highly modular architecture. Moreover, it allows to implement abstract processes independently of the domain, but containing the necessary particularities to use them to solve specific problems. In the following sub-section the agents' functionalities are described briefly. Afterwards the approach followed to use this abstract architecture to solve specific problems is detailed.

4.3.1 *Functionalities*

The following listings briefly describe the main agent types that make up the proposed architecture and the functionalities they implement. It must be kept in mind that the actual architecture may encompass more than one instance of some of the agents enumerated, so as to address specific legal domains (this is detailed in Section 4.3.2). Figure 12 depicts the layered nature of the proposed architecture. Given what is depicted in Chapter 7 and the nature of the topmost layers, a special emphasis was placed while devising this architecture on making it dynamic and modular, so as to seamlessly integrate users, devices, sensors and environment (Preuveneers and Novais, 2012).

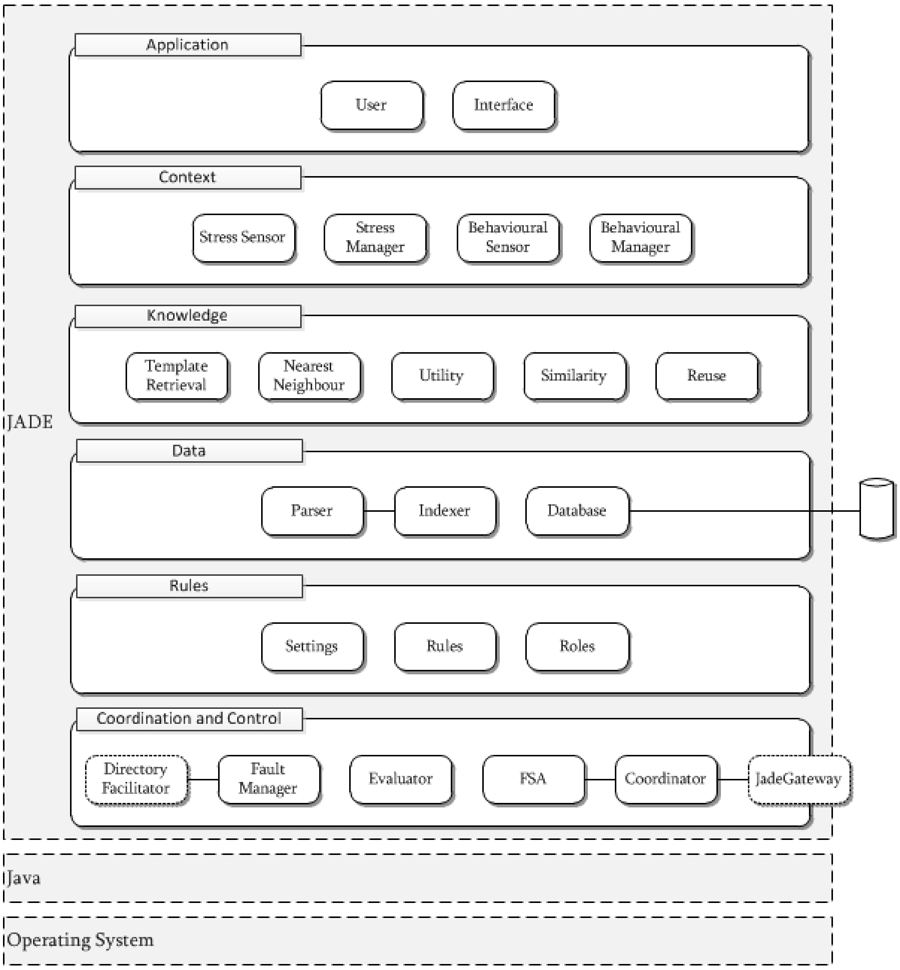

Figure 12: Diagram depicting the layered nature of the architecture, with the organization of the software agents according to their functionalities. Lines between agents depict important communication paths. More than one instance of some agents may exist simultaneously.

Coordination and Control Layer

The agents belonging to this layer have as main purposes to control the life-cycle of the platform and the correct execution of the complex processes implemented.

- Coordinator - Receives task requests from other agents (e.g. external agents, interface agents) and takes the necessary steps (requesting tasks to other agents) in order to implement them. This agent maintains a list of active tasks and has access to a list of finite state automata that completely define each task or activity, provided by the FSA agent.
- JadeGateway - The JadeGateway agent is a special agent that allows non-Jade and Jade-based code to communicate. It is used to allow the multi-agent system to communicate with external entities.
- FaultManager - Contains the necessary mechanisms to detect malfunctions in software agents. It is able to stop, restart or start the agents that make up the architecture. Its main purpose is to start the agents on startup and monitor their life-cycles.
- Evaluator - Receives messages from the software agents detailing the timestamps of certain events, in order to compute performance indicators and point out possible paths to improve the platform (e.g. most common causes for case retrieval failure, average time of each activity).
- FSA - Contains a list of Jade FSM behaviours that describe the guidelines or steps necessary for an agent to perform given actions.

Rules Layer

The agents that make up this layer are intended to encode the rules than ensure the validity of the operations of the platforms, as well as legal norms.

- Settings - Defines several pre-determined settings, mostly for cases-related operations (e.g. retrieval, computation of similarity, computation of utility) according to which retrieve parameters can be changed.
- Roles - Each software agent has associated a set of roles, describing the services that it can request within the architecture. This agent encodes the access schemes.
- Rules - This agent maintains a set of rules for different purposes (e.g. legal norms, validity of a solution).

Data Layer

The agents in the Data Layer have as main objective to facilitate access to data by providing an abstraction layer to files and database.

- Parser - This agent receives a list of XML files describing cases and parses them, converting them into software objects that can be then used by the remaining agents.
- Indexer - Given a list of cases, this agent indexes them in the database so that they can be easily retrieved.
- Database - All the interactions with the database go through this agent. It provides several methods for data retrieval and storage that facilitate the interaction with the database.

Knowledge Layer

The agents in this layer aim at compiling meaningful knowledge from the lower layer so that parties can be better informed and improve their decision-making processes.

- TemplateRetrieval - Encodes a pre-selection algorithm that selects a broad range of cases that match given pre-selection rules. This agent has the autonomy to change the search settings, the similarity parameters and parameters of the retrieve algorithm in order to perform a better selection of cases.
- NearestNeighbour - Encodes a pre-selection algorithm to retrieve cases from memory based on the Nearest Neighbour algorithm. It has the autonomy to iteratively change the search settings so that better search results can be achieved.
- Utility - Encodes the necessary mechanisms for computing the utility of an outcome for a given party.
- Reuse - Contains the necessary knowledge to adapt past known cases to new scenarios by replacing key parameters.
- Similarity - Is able to compute a value of similarity between two cases, based on a weighted sum of their differences in key characteristics.

Context Layer

The agents that make up the context layer aim at providing, in a seamless way, the tools to assess the context of the user, namely their level of stress and their behaviour.

- StressSensor - This agent represents one or several sources of information that are used to assess the level of stress.

- StressManager - Contains the necessary mechanisms for assessing the level of stress of users, in real-time, based on the data provided by one or several StressSensors.
- BehaviouralSensor - Represents one or more sources of information that are used to evaluate the behaviour of the user in what concerns the actions during the conflict resolution process.
- BehaviouralManager - Encodes the necessary mechanisms for drawing conclusions about the behaviour of the users based on the information provided by the BehaviouralSensors.

Application Layer

The main aim of the agents in this layer is to isolate the user from the complexness of the remaining platform. In that sense, these agents implement abstract interfaces to the functionalities provided by the platform.

- User - This agent represents one party, while he is using the platform, encoding characteristics such as objectives, personal information, context, among others.
- Interface - This agent is responsible for setting up the several interfaces with which the user can interact.

4.3.2 *Using an Abstract Architecture to solve Specific Problems*

One of the main objectives of this research work is to define an agent-based architecture that can be used in different legal domains, so as to increase the fields of application. In fact, many of the current ODR tools are very domain-specific, resulting in a low use/acceptance.

Some abstract processes and concepts have already been identified that are valid in any legal domain. However, the way they are implemented varies according to the domain. As an example, let us consider the computation of the BATNA and WATNA in two different legal domains, namely: Commercial Law and Labour Law. These two abstract conceptions denote, as described above, the best and worst possible scenario in a litigation process. Nevertheless, depending on that domain, it is computed in different ways.

In Labour Law it will have to be considered issues such as of worker antiquity, monthly salary, seniority, a just cause for dismissal, just to name a few (Listing 1). The following listing details a simplification of the rules that allow the computation of the BATNA and WATNA for the Portuguese Labour Law, as it is given in Decree of Law (DL) 7/2009 (Portuguese laws). This simplified rule considers only the case in which a worker ends the contract with a just cause.

Listing 1: Computing the BATNA and the WATNA in the Labour Law domain

```
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
```

If rule 394 is true (i.e. if there is a just cause for dismissal), in the worst case the worker is entitled to the value of three months' salary. It must be kept in mind that the BATNA of one party is often the WATNA of the other and vice versa (an indemnity of three months' salary is the BATNA of the employer and the WATNA of the employee). In the case of a temporary contract, the employee is entitled to an indemnity that amounts to the number of months still in the contract times salary and seniority. On the other hand, if the contract is not temporary, in the worst case the worker is entitled to 15 days of wage plus seniority per each year of antiquity, at a minimum of three months. In the best case, the employee receives 45 days of salary per each year of antiquity. This value may be higher if there has been some damage done to the employee.

On the other hand, in Commercial Law key issues are the date in which the product was bought, the type of the product, the type of the warranty, the state of the product, among others (Listing 2). Thus, although the concept means the same it is implemented differently. The following listing depicts a simplification of the rules that allow the computation of the BATNA for the Commercial Law, as it is given in DL 67/2003. In this example rule it will be considered only numbers 1 to 4 of Article 5th.

Basically this rules states that if, (a) the item is a mobile device, (b) the complaint has been filed until 60 days from the date of detection of the defect and (c) less than two years have passed since the date of bought, the consumer is entitled to a reparation in 30 days or to the replacement of the product for an equivalent one. In other case there is no indemnity due. On the other hand, if the item is not mobile, the dates change respectively to one and five years and the user is entitled to the reparation of the product.

Thus, two concepts that have the same meaning in different legal domains, are implemented differently. The approach followed to address this issue focuses on two key points:

Listing 2: Computing the BATNA and the WATNA in the Commercial Law domain

```
Def_Rule 5
if IS_MOBILE then
  if DEFECT_COMPLAINT_DELAY < 60 then
    if WARRANTY_DELAY < 730 then
      BATNA := {"product repair in 30 days";
                "product replacement"}
    else BATNA := {"no indemnity due"}
  else BATNA := {"no indemnity due"}
else
  if DEFECT_COMPLAINT_DELAY < 365 then
    if WARRANTY_DELAY < 1810 then
      BATNA := {"product repair in reasonable time"}
    else BATNA := {"no indemnity due"}
  else BATNA := {"no indemnity due"}
```

- Each secondary agent is broken into several similar agents (with the same service interface) and contains the knowledge necessary to implement its methods in a specific legal domain;
- Each coordinator agent forwards a request to specific secondary agents according to the domain of the request.

There are some advantages in following such approach:

- The complexity is hidden from the other agents as the process of requesting a service for a particular domain is completely transparent;
- It is easy to extend the architecture to other domains by adding the new secondary agents;
- When new domains are included, no change is necessary to the existing agents;
- The architecture is highly modular and easily reconfigurable.

A simplified view of this hierarchical nature of the architecture is depicted in Figure 13. It details the agents involved in computing the value of the BATNA for the Labour law domain. When a request is received by the "Coordinator" agent, it reads its domain and forwards it to the corresponding "Domain Coordinator", which possesses the knowledge necessary to coordinate tasks in the specific domain. In this case, only one service is needed from the BATNA agent, thus the request is forwarded to the BATNA agent. However, this agent is in turn a high-level agent that will, given the domain of the request, forward it to the "BATNAL" agent, which possesses the knowledge

New legal domains can be added without any changes being made to the existing agents

about how to compute the BATNA in the domain of the Labour Law. The result of the service request is returned via the same agents.

When forwarding messages, the following approach is followed. The name of each agent is composed of a first part detailing the service provided and a second part detailing the domain of specialization. As an example the "BATNAL" agent provides the computation the BATNA in the Labour law domain. The following domains are considered (for mode detail on the case studies please consult Section 9.7): Labour Law (L), Consumer Law (C), Property Division (P) and Virtual Organizations (VO).

In that sense, when a coordinator agent receives a request for a given service S in the domain D, it will search the agent registry for an agent identified as SD. If it exists the message is forwarded, otherwise the request is rejected. The main implication of this approach is that new legal domains can be added by programming new agents, without any changes being made to the existing agents.

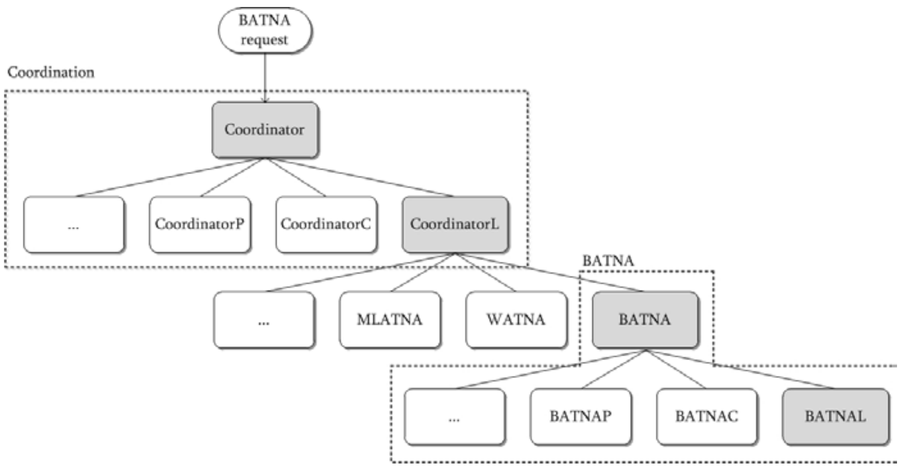

Figure 13: Hierarchical nature of the architecture in which the complexity of providing specific services is hidden. Each agent forwards the request to more specific agents. Domain coordinators (e.g. "CoordinatorL") are used to implement more complex activities involving more than one services (the computation of the BATNA requires the use of a single service from a single agent).

4.4 SUMMARY

One of the factors holding back a faster development and use of conflict resolution platforms is, without doubt, their highly domain-dependent nature. In fact, current approaches that result in an actual working prototype or service are generally target at a single domain. This significantly narrows the target audience, reducing eventual revenues and slowing the natural development of these platforms.

From this starting point, in this chapter it was analysed the possibility of developing conflict resolution services or platforms that could be, insofar as possible, independent of the domain. The main objective is clearly that the same services can be used in different fields, thus widening the possible domains of application.

In order to accomplish this vision, the chapter started by identifying some abstract concepts and processes, i.e., elements that mean the same independently of the legal domain, although the way of their computation may differ slightly. This would allow to implement a conflict platform based on such concepts and processes, that would provide the same services in different domains.

The final part of the chapter detailed the actual development of an agent-based architecture for the UMCourt platform, which provides a set of services detailed in the following chapters, in a domain-independent way. The approach followed is a highly modular one, being also transparent: the processes are the same, what changes are their domain-related specificities, which are hidden from the remaining platform.

Conflict resolution platforms that follow such an approach will increase the reuse of already defined functionalities and processes, resulting in more efficient and solid development processes.

5 KNOWLEDGE RETRIEVAL

The secret of success
is to know something nobody else knows.
— Aristotle Onassis

In Chapter 4 it has been discussed the most important abstract concepts and processes in the context of an Online Dispute Resolution platform. The software agents that make up the supporting framework for this platform have also been briefly described.

In this chapter a step forward is taken towards a more precise specification and implementation of such notions. Specifically, here it is described how meaningful knowledge is retrieved and compiled in order to be made available for the parties.

This chapter describes how meaningful knowledge is retrieved and compiled in order to be made available for the parties.

Firstly, one must bear in mind that the ideal conflict resolution process is the one in which the two parties change for the better at the end of the course of action. Unfortunately, not all conflicts come to this end. In order to improve such odds, it is of the highest importance to provide the parties with more knowledge about the dispute so as to enhance their possibilities throughout all the process.

In fact, parties that have poor access to important information generally end making bad choices or, at least, they hardly take the best one. Moreover, parties usually have a reduced role on the process, resulting in suspicion about the outcome, mostly because they do not fully understand how or why it was achieved (Deutsch, 1958).

Given the growing amount of information with which parties must deal with, efficient information retrieval tools are nowadays essential. In fact, there is a significant amount of information to be considered, including legal norms, legal texts, past cases, rights and obligations, standard procedures and standards of conduct, among others. Thus being, information retrieval tools can be useful for both law practitioners and disputant parties. Actually, such tools are of use in any knowledge-based domain as a way to improve decision making mechanisms.

Specifically, in UMCourt , information retrieval techniques are used for compiling many different types of knowledge. Rule-based techniques are used to compute the values of the BATNA and WATNA. These are computed according to the legal norms that define the best and worst possible outcomes of the specific legal domain. Two examples of computing the BATNA and the WATNA in different legal domains are given in Section 4.3.2. Their implementation consists in implementing the corresponding rules in the software agents. Computing the BATNAs and WATNAs is thus a fairly simple process once the important legal norms are identified. Based on these two boundary values the value of the ZOPA can be easily estimated by computing the distance between them.

Significantly more complex but equally or more important is however the retrieval of past cases. Past similar cases are important because parties can look at them and their respective outcomes to have a feeling about their odds in the process, by assuming that past similar cases in similar contexts will have similar outcomes. Moreover, software agents can make recommendations based on the observation of a number of retrieved similar cases. Legal practitioners, namely mediators and arbitrators, can also look at past cases in order to better decide on how to guide the dispute resolution process.

Given the importance of past cases in the context of conflict resolution, in this chapter three different information retrieval methods are described whose objective is to select and retrieve past cases that may be of use to solve a given problem. Although their objective is the same, their implementations differ. These methods are implemented and analysed in order to determine the ones more fit for the problem domain.

The first method is based on a similarity function that selects cases above a given threshold. Similarity is determined by analysing the differences in key parameters, with given importances or weights. The second one categorizes cases according to their main features and uses association rules to decide to which category a case belongs to. When retrieving, cases that belong to the same category are considered potentially similar. Finally, the third method is an hybrid one in which the two former ones are combined.

5.1 THREE KNOWLEDGE RETRIEVAL METHODS

As in other case-based knowledge domains, information retrieval methods in UMCourt are used to select among a group of past cases the ones that may be relevant to solve a particular problem. In what concerns this issue, several issues must be considered.

The first one is related with the amount of information retrieved. The existence of some mechanism to control the amount of cases retrieved is desirable, namely because: (1) there is no utility in retrieving a large number of cases if only a few have a significant value of similarity, and (2) there should be a minimum/maximum amount of cases that can provide the necessary information for the parties to take rational and informed decisions. Few information will not be enough for good decisions and too much information will make the decision-making process inefficient.

Another issue is related with the nature of the domain: such methods must be autonomous, i.e., they cannot depend significantly on human experts at running-time. In that sense, information retrieval that can operate autonomously, with minimum interaction required, are valued. This implies the ability to interpret the party's request and request's context (e.g. legal domain, party's objectives) and to set or reset the search parameters in order to return optimal results.

The third important issue is the information that makes up a case. A case is, under this setting, the basic unit of information. It denotes a past experience that took place in a context that is also described in the case. It is therefore a contextualized piece of information and allows to estimate the outcome of an experience that is now happening by looking at a past similar one (that took part in an equivalent context) and its respective outcome. Generally, the information contained in a case may be organized into three distinct categories (Leake, 1996):

The information contained in a case may be organized into: problem. solution and outcome.

- The Problem – the description of the problem, including the context of applicability, the state of the world and all other information that may 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 autonomously by the machine and abstract enough to be adapted to new problems.
- The Outcome – what happened when the described solution was applied to solve the problem, i.e., the final state of the world.

Although these categories depend on the problem domain, some applications may only consider the problem and the solution (in which case it is possible to derive solutions to new problems), while others may consider the problem and the outcome (making it possible to estimate outcomes to new problems).

In the particular case of this work, the information detailed in Table 1 is considered.

CAT	INFO 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 towards the dispute, the intended outcome.

	Legal	Legal information such as the norms used by the parties and witnesses to support their claims or the guilty statement.
	Dates	All the important dates of the case.
Solution	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 actions that describe the outcome.
	Value	A value denoting the percentage of successful applications of this case in the dispute resolution process.

Table 1: This table describes the main components of a case as considered in this research work.

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 (Watson and Marir, 1994). Three different approaches that were implemented and tested are presented in the three following sub-sections.

5.1.1 *Method 1: Using Similarity Functions*

In this approach, the Nearest Neighbour Algorithm is used that is able to compute a value of similarity between two cases by comparing some key characteristics. Cases are then selected according to their value of similarity with the new case: if they are above a given threshold, they are selected.

$$sim = \frac{\sum_{i=1}^{n} w_i * fsim_i(Arg_i^N, Arg_i^R)}{\sum_{i=1}^{n} W_i} \tag{1}$$

In Equation 1, the equation for the computation of the similarity is depicted. In this equation we have:

- n – number of issues to consider to compute the similarity;
- W_i – weight of issue i in the overall similarity;
- $fsim$ – a specific similarity function for issue i;

- Arg – arguments for the similarity function representing the values of the issue i for the new case and the retrieved case, respectively N and R.

In a broad sense, for each issue, this similarity function looks at the values in each case being compared and computes a value of similarity. The way the similarity is computed depends on the nature of the issue and is described further ahead. Each of the issues has a given significance for the overall computation of similarity (e.g. the legal norms used by the parties to support their points of view is far more important than the years in which the cases took place). These weights are, at this moment, determined by a law expert, based on the substance that, according to his experience, each of the components of the similarity measure has.

The function looks at the values in each case being compared and computes a value of similarity.

Let us now detail the issues that make up the description of the problem as well as the way the similarity is computed. According to application domain, and as stated above, three types of information are considered to define the problem: (1) the objectives stated by each party in the beginning of the dispute, (2) the norms addressed by each party and by the eventual witnesses, and (3) important dates of the dispute. Both the norms addressed and the objectives of the parties are lists of elements, thus the similarity function consists in comparing two lists, as depicted in Equation 2. The similarity is higher when the two lists have a higher percentage of common members. As for the date, the similarity function verifies if the two dates are within a given time range, having a higher similarity when the two dates are closer.

$$fsim_{list} = \frac{|L_N \bigcap L_R|}{n}, n = \begin{Bmatrix} |L_N|, & |L_N| \geqslant |L_R| \\ |L_R| & |L_N| < |L_R| \end{Bmatrix} \tag{2}$$

Once all the values of the several similarity functions are summed in accordance to their weights, a value of similarity is obtained that describes to which extent a past known case is similar to the new one. By applying this algorithm to each known case, it is possible to select the most similar cases.

The main disadvantage of this approach is that the algorithm requires significant computational power and may take some time to perform, depending on the size of the case-base. Another disadvantage is that, for each new case or problem under consideration, all the values of similarity must be computed again as each new case is different from the previous ones.

On the opposite side, the main advantage is that once all the similarity values are computed for a given new case, it is easy to determine which cases to select: this is done by changing the similarity threshold. This is especially useful for controlling the number of cases that are retrieved.

Let us consider, for example, a case-base that is fairly big. It is likely that under this scenario, a relatively large amount of cases will be retrieved. That is not desirable as too much information may impede an efficient analysis and decision-making process. However, using this approach, retrieving the ideal number of cases is as easy as changing the similarity threshold.

That is, if too many cases were retrieved, one can increase the similarity threshold, ensuring that fewer cases will be presented to the user, with a higher average value of similarity. On the other hand, if too few cases are selected, the disputant party may not be able to get the whole picture. In that sense, it might be useful to decrease the similarity threshold in order to retrieve more cases. Once the similarity values are computed, this process becomes straightforward and the results space can easily be changed.

5.1.2 *Method 2: Using Association Rules*

The aim of this second method is to identify relationships between the values of given variables that make up a case. This is a fairly common task in data-mining, with a wide range of applications. The main objective is to find hidden patterns that may help to explain or determine some behaviour. The most traditional example is the use of association rules in a shopping environment to determine the behaviour of the customers. Generally, shopping lists are analysed to determine which products are bought together, with the aim of better placing the products in the store or better directing marketing campaigns.

The aim of this second method is to identify relationships between the values of given variables that make up a case.

From such analyses, rules can be defined that describe conducts such as "seventy percent of the people who buy beer also buy appetizers". Alternatively, if we think in the legal domain, we can consider rules such as "sixty percent of the cases in which norms A and B are properly used by one party to support his argument, that party has a successful outcome". Such rules are stated in the form ifXthenY, or $X \Rightarrow Y$, in which X is the antecedent of the rule ("use of norms A and B" in the first example) and Y is the consequent ("party wins").

In order to support the generation of the rules and select the ones that are essential, statistical aspects can be considered. These include the notion of *Support*(the ratio of transactions supporting the association rule and the total number of transactions within the database of transactions) and *Confidence* (the percentage of transactions supporting the rule out of all transactions supporting the rule body). Generally, only rules that have a confidence factor above a given threshold are considered.

There are currently several software tools that can be used to explore large sets of data and find such rules. One of the most widely used is Weka (Holmes et al., 1994). Weka contains tools for data preprocessing, classification, regression, clustering, association rules, and visualization. In this research work, Weka is mainly used to find association rules in values of a database using different algorithms. Figure 14 shows the use of the Apriori algorithm in Weka, to find rules in a database that contains a list of past cases. Weka can also be called from a Java application. This was also an important reason to choose this tool since it can be easily integrated with the conflict resolution platform.

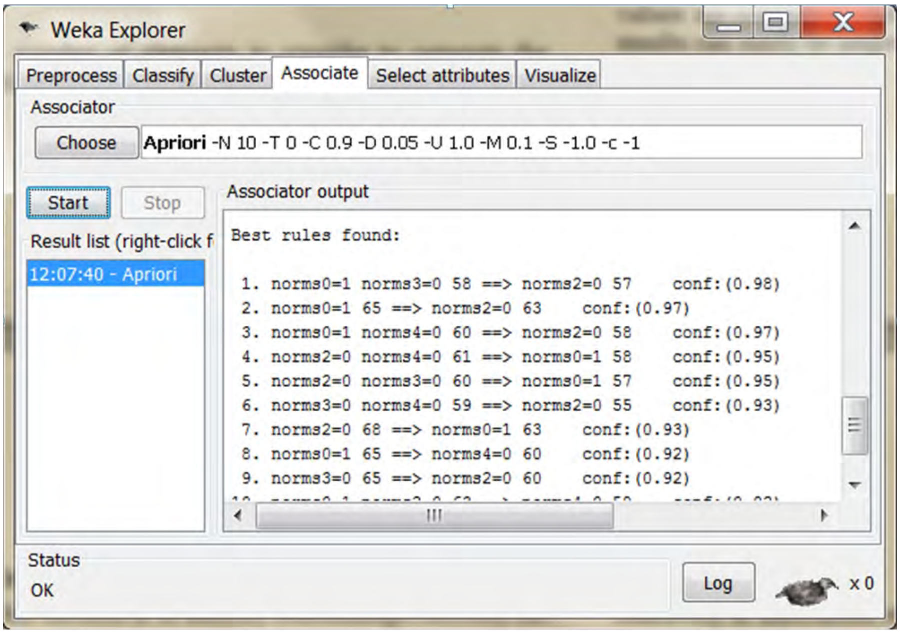

Figure 14: In this work, Weka is used to find association rules in the database. In this example it is used for searching for legal norms that are frequently used together in a conflict resolution process.

However, there is still the need to determine which of the identified rules make sense given the domain of application. Let us consider, for example, the best rule found in the example depicted in Figure 14. According to this rule, with a confidence of 0.98, in the cases in which norm 0 is used and norm 3 is not used, norm 2 is also not used. At first sight this rule does not contain any useful information. However, we could be interested in a rule stating that cases in which norm 1 and norm 2 are used and the objective of the party is to solve the dispute at all cost, that party wins.

The work of identifying and selecting relevant and meaningful rules is done by a legal practitioner. Although this might be an extensive work, it must only be done once for a database of cases. Once it is done, these rules can be used to create categories or classes of cases. Under this approach, cases are assigned to categories according to the rules they comply with: the cases for which the same rules are true belong to a same category.

In order to follow this approach, the information contained in the database about the cases is represented according to the vector space model (Salton et al., 1975). 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.

The vector space model allows to represent textual documents as vectors of values.

Specifically, in this work, a case is seen as a vector V of binary entries, in which each entry $i \in [0..n-1]$ corresponds to a fixed descriptor from the descriptor vector $\vec{D}$ of size n. The 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, outcome 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. Considering the example depicted in Equation 3, a case represented by vector $\vec{V}$ would address $Art.128.°n.°1, c)$, $Art.128.°n.°1, d)$, and so on, according to the n-dimension descriptors vector.

$$\vec{V} = \begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 0 & \cdots & 1 \end{bmatrix} \qquad \left.\begin{matrix} Art.128.°n.°1\ a) \\ Art.128.°n.°1\ c) \\ Art.128.°n.°1\ d) \\ Art.128.°n.°2 \\ \cdots \\ Art.1114.°n.°1 \\ Art.1114.°n.°3 \end{matrix}\right\} D \tag{3}$$

This representation of a case allows to encode the norms that each party addressed, their objectives, the outcome, among other issues. It results in a very concise manner of representing all this information (one bit for each issue being considered), demanding very few resources to handle and to store it. 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 and a descriptors vector (Equation 4).

$$database = \begin{bmatrix} 0 & 1 & 0 & \cdots & 1 \\ 1 & 1 & 0 & \cdots & 1 \\ 1 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 1 & \cdots & 1 \end{bmatrix} \qquad \left.\begin{matrix} Art.128.°n.°1\ a) \\ Art.128.°n.°1\ c) \\ Art.128.°n.°1\ d) \\ Art.128.°n.°2 \\ \cdots \\ Art.1114.°n.°1 \\ Art.1114.°n.°3 \end{matrix}\right\} D \tag{4}$$

Given this data representation, it is possible, as stated before, to apply association procedures to determine relationships between the data. The objective is to create groups of documents, or cases, in which the same rules return a truth-value true. Then, the retrieval process becomes relatively simple: whenever cases have to be retrieved for a given problem, the system first determines which rules return a truth-value true for the new case. This will allow determining to which category the case belongs to. Then, all the cases of that group can be retrieved, as they are potentially similar and appealing to find a solution to the new problem. Thus being, this approach consists in classifying cases using association procedures. The purpose is thus to group the cases in such a way that retrieval will be faster.

The main advantage of this method is its computational and spacial effectiveness.

The main advantage of this method is indeed its effectiveness. Once all the cases of the database are classified, it becomes very easy to retrieve the cases from a given group. However, on the down-side, there is no control on the number of cases that are retrieved. That is, this method cannot actively control the amount of cases that are retrieved to be presented to the users. In fact, if the new case is classified as belonging to a group that contains hundreds of cases, all those cases will be retrieved, rendering the analysis of the results nearly impracticable for a human user.

5.1.3 *Method 3: A Hybrid Method*

The third method combines the advantages of the previous two in order to make it dynamic, efficient and autonomous.

The third method implemented is intended to combine the advantages of the two previously presented methods in order to make it dynamic, efficient and autonomous. Recalling, the first method presented had as main advantage the ability to control with precision the amount of cases retrieved while the second presented had as main advantage a fast retrieval by means of case classification. These are the two advantages that are merged in this hybrid approach.

Specifically, this method has a preparation phase and two running phases: pre-select and evaluation. In the pre-select phase, association rules are used to pre-select cases in an efficient way. In the evaluation phase, the system assesses the amount of cases retrieved and may refine the pre-selection by means of similarity functions. Figure 15 depicts this process.

In the preparation phase, association rules are discovered in the database. Then, a human expert determines which of these rules are to be considered and which ones are to be discarded. Once this task is finished, the system goes through all the cases in the database and classifies them according to the rules that they comply with. This will organize the cases into groups or categories: each case in a category is more similar to any case in the same category than it is to any case in other categories. This preparation phase must only take place

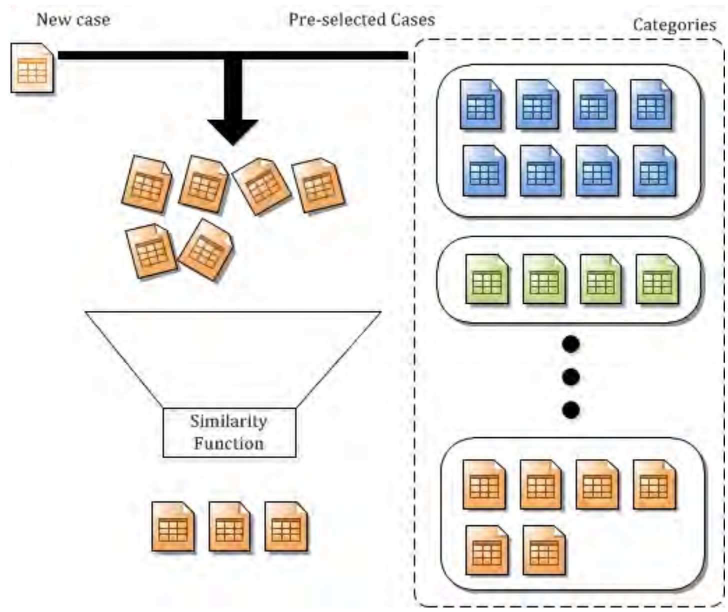

Figure 15: This hybrid method uses classification rules to make a pre-selection of cases that have the potential to be similar (because they share some characteristics). It then uses a similarity function to decide, among the cases inside a category, which ones are more relevant.

once for a given database of cases. When it is finished, the database is ready to be used by the system to retrieve cases.

A common retrieval process starts with the pre-select phase. This phase makes use of the previously mentioned method that relies on association rules to select cases. In that sense, it starts by determining which of the known rules return a truth-value true for the new case. Having done that, the system determines to which category or group the new would belong to, if it were to be indexed in the database. Then, it returns all the cases that belong to that category since they are potentially similar. As an example, if rules 1, 3 and 4 are true for the new case, the process will return all the cases for which the same rules are true, if any. Such cases belong to a category identified as C-1;3;4. Up to this point, this process is quite efficient as the cases were already indexed in the respective categories from the preparation phase.

In the second phase, the evaluation one, the system analyses the results of the pre-selection and determines if further actions are needed. If the amount of the retrieved cases is inside the desirable range (this depends on the request properties or target user), the process ends and all the cases pre-selected are returned. On the other hand, the amount of cases pre-selected may not be the desirable one. This being the case, two scenarios are possible.

In the first one, there are few cases selected. In such a scenario, the system will try to relax the pre-selection rules. Still considering the previous example in which rules 1, 3 and 4 were true for a new case, let us assume that the category C-1;3;4 does not exist (this is possible if, in the preparation phase, no case existed in the database for which rules 1, 3 and 4 were simultaneously true). Under this assumption, no case would be pre-selected as no group matched the same rules of the new case. In that sense, the system would relax the pre-selection rules by looking at the rules individually. Thus, under this assumption, if categories C-4 or C-1;3 existed, their cases would be pre-selected as they represent groups of cases that have something in common. This relaxing of the rules is performed to ensure the needed minimum amount of cases, although their average value of similarity is lower.

In the second scenario, too many cases are selected. In this event, a fine-tuning of the pre-selection must be performed in order to obtain fewer cases. However, this fine-tuning must be performed under the requirement that the less similar cases are discarded in favour of the most similar ones. In this scenario, this method makes use of a similarity function to decide on which cases to discard and which ones to consider. This will pair each pre-selected case with a similarity value. From this point on, the system only has to change the similarity threshold in order to change the amount of cases retrieved. In that sense, one of two different approaches may be selected.

On the one hand, the system can make use of the similarity function presented before. This method is straightforward and consists on the application of the algorithm described. Nevertheless, it involves retrieving the cases from the database in their original form in order for the algorithm to be applied. However, as the indexes of cases are already known from the pre-selection phase, the process is relatively fast. In this algorithm, output values range from 0 to 1, with 0 meaning that there is no similarity at all and 1 denoting an exact match.

A faster approach, however, is to use the vectorial representation of the cases, which is available from the pre-selection phase, to compute the similarity. Similarity in the vectorial representation model can be computed by means of the cosine similarity (Tan et al., 2005). This is based on the principle of computing the similarity between two vectors 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 Equation 5. Given the modelling of the data as vectors of binary entries described previously, the cosine similarity of two cases will range from 0 to 1, i.e., the angle between two vectors cannot be greater than $90°$.

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 now allow to assign weights to the several components of the case. This may or may not be a disadvantage, depending on the scope of the application.

$$sim = \frac{A.B}{\|A\| \, \|B\|} = \frac{\sum_{i=1}^{n} A_i * B_i}{\sqrt{\sum_{i=1}^{n} A_i^2} * \sqrt{\sum_{i=1}^{n} B_i^2}} \quad (5)$$

5.2 PROMOTING BETTER DECISIONS BY PROVIDING KNOWLEDGE

Given its characteristics, the third method is being used as the privileged knowledge retrieval method. Specifically, the knowledge retrieved can be presented to the parties, in order for them to take more informed decisions. Moreover, it can be presented to a mediator or arbitrator, which will analyse the past cases in order to take better decisions, supported by previous occurrences. Alternatively, the information can be forwarded to a software agent in order to implement higher level functionalities such as the ones presented in the following chapters.

The knowledge build can be depicted graffically, to efficiently support decision making.

An example of what can be done with the knowledge retrieved is depicted in Figure 16.In this example, a software agent uses the information provided by the system to build a graphical representation with some added value, to show to a party. Several knowledge components are included in this graphical representation:

- Past similar cases - Each past similar case is represented as a coloured circle. Each past case has a value of similarity to the current case and a value of utility for a party (quantifying how much that party would win/lose if the outcome of that case would happen). In that sense, cases are plotted in a Cartesian plane, in which the x-axis represents the utility of the case and the y-axis represents the similarity. Cases can be clicked in order to access additional information (e.g. date, how was the utility/similarity computed). This allows a party to quickly understand how likely and how desirable a given past case is;

- Linear regression - A linear regression is plotted that allows a party to understand how the values of utility and similarity vary in the space of solutions. It is particularly useful to provide information about areas of the space where no past cases were retrieved;

- BATNA, WATNA and MLATNA - The values of the BATNA and WATNA are represented by the extreme values of the utility range, i.e., the WATNA is the leftmost value of utility and the BATNA is the rightmost. As for the MLATNA region, it is represented by the green portion of the line of linear regression. With this data the party becomes aware of what are the worst and best possible outcomes as well as the region of utility in which an outcome is more likely;

- Clusters - Clusters are computed that try to identify past cases that share similar values of similarity/utility. For each cluster, the average value of similarity and utility is computed. The main aim is that a user can get a general view of the knowledge retrieved by looking at the clusters and their values and then, if necessary, analyse particular cases;

- Visualization controls - The party can use several controls that allow to fine-tune the amount of knowledge displayed. The similarity and utility threshold sliders allow to adjust the minimum value of utility/similarity, respectively, of the displayed past cases. These conditions may be set in union (pressing the "OR" button) or in conjunction. The user can also use the "Complexity threshold" slider. This allows to hide cases that are very close in terms of similarity and utility, appearing almost overlapping in the Cartesian plane. Adjusting this slider removes cases within a given distance of already plotted cases;

Therefore, this information retrieval method results in an efficient and multifaceted one, enabling the implementation of very different functionalities that can considerably improve the experience of people and systems involved in dispute resolution processes.

Figure 16: A detail of the graphical representation of the knowledge retrieved. Past similar cases are represented as coloured dots, accompanied by their value of utility for a specify party and similarity to the new case. The utility ranges from the BATNA to the WATNA. The MLATNA is the region in which the curve of the linear regression is green. Some clusters are also computed, with a respective mean similarity and utility.

5.3 SUMMARY

This chapter focused on describing several viable approaches for retrieving meaningful knowledge in the context of Online Dispute Resolution. The main functionality implemented by these approaches is to retrieve past cases that are significant for the solution of a given problem, typically a conflict resolution one. Past cases are important as they provide a good basis for comparison and for understanding the possible outcomes: an outcome of a similar case that took place in similar circumstances is likely to happen in the current case. Particularly interesting due to its simplicity and performance is the hybrid method, the one currently at use by the UMCourt platform.

The use of these functionalities allows a party to be aware of their best and worst possible alternatives, as well as the region of the most likely outcome. Parties can look at the knowledge from a high point of view, by considering clusters of cases, or specifically, by considering particular cases. Each case has a value of similarity and utility, that allows its plot on a Cartesian plane. This allows a party to rapidly look at their "odds", i.e., many cases with high values of utility and similarity are a good auspice. Some controls are also provided by the interface that allow a party to manage the amount of knowledge that is accessible at each time. The same knowledge and controls is also available to other software agents, which significantly increases the possibilities. In the following chapters it will be depicted how these base functionalities are used in UMCourt to build higher-level ones such as the suggestion of solutions and the planning of strategies.

6 NEGOTIATION SUPPORT

Diplomacy is the art of letting someone else have your way.

— Sir David Frost

The problem of generating solutions for conflict resolution is generally addressed by devising Expert Systems. Two main trends can be identified in the literature: rule-based and case-based solution generation. However, both rule-based and case-based approaches face criticism, given their prospective disadvantages, one of the worst being that most of the solutions developed so far are restricted to specific domains of The Law (e.g. the Peruvian Cibertribunal[1], focused on electronic commerce; WIPO[2], focused on Domain Name Dispute Resolution). This makes it hard to reproduce the results of a given project in other domains.

Both rule-based and case-based approaches to negotiation face criticism .

Considering RBSs, the main disadvantages of this approach concerning the generation of solutions are linked with the human experts that formulate the rules. In fact, a big challenge is effectively to determine the Quality-of-Information (QoI) coded by the rules. In that sense, there must be a big concern in involving experts from both the Legal and/or the Computer Science arenas in the task of defining the rules. Moreover, a RBS is not an optimal solution generator for all problems. Consequently, considerable knowledge is needed not to misapply these systems. Finally, it is a fact that the ease of rule creation and edition can be an advantage. However, this can be also seen as a potential disadvantage as the system can easily be sabotaged by a non-knowledgeable user. Typical reasons for the failure of RBSs include the negligence to employ simple tools for system audit that can detect conflicts in rules (Popple, 1990).

Considering CBR, its main disadvantage stems from its potential complexness. In fact, in a multifaceted knowledge-based domain, the simple definition of the content of a case may be an overwhelming task, involving many different experts from different fields of knowledge. Moreover, this also implies the use of significant amount of resources to store and deal with such cases. This means that case-based approaches are generally more intricate and resource demanding. Additionally, most of the systems analysed in the scope of this PhD work are static rather than adaptive. This means that once a strategy is defined, generally at the outset of the process, it will be

1 http://www.cibertribunalperuano.org/
2 http://www.wipo.int

followed disregarding probable changes in the environment that sets the context of interaction.

We can thus succinctly enumerate the main drawbacks of both RBS and CBR approaches in what concerns the generation of solutions for a conflict resolution problem:

- Laws change constantly thus implying updates to the rules that establish how solutions are generated, in RBS. This may result in inconsistencies and/or redundancy. Moreover, this might be quite a complex task (depending on the complexity of the legal domain) that must be performed manually, despite the use of some supporting tool;
- The quality of a RBS ODR tool is directly dependent on the quality of the work of the humans translating the legal norms into rules. The quality of information of the rules may be hard to determine;
- RBS are static and will not shape changes in the legal domain, unless these are coded manually by a human expert;
- The quality of a CBR ODR tool is directly dependent on the quality and amount of past cases known;
- The fact that legal norms change frequently also has a negative impact in CBR approaches, rendering past cases potentially useless under the light of the new norms;
- Both CBR and RBS approaches are domain dependent. This implies that rules are defined independently for each legal domain and that cases from a specific domain can hardly be reused to another.

Recently, soft computing techniques have been used to address similar complex problems, although in different domains of application. The works described in (Sedano et al., 2010; Corchado et al., 2010), in which soft computing techniques are used to detect the lifetime of building thermal insulation failures and to identify typical meteorological days, respectively, can be seen as good examples. More specifically, evolutionary computation has been successfully used to deal with problems that involve a significant amount of complex variables in which traditional approaches would not be suited (Abraham, 2009; Corchado and Herrero, 2011).

In line with this new trend, this section describes the use of genetic algorithms to generate potential solutions for the resolution of a conflict. The main objective is to assist parties and mediators in generating consensual solutions in a conflict resolution process, generally a negotiated one.

A nature-inspired approach is used to create solutions for conflicts.

Under this approach, each chromosome will represent a possible solution for a conflict, i.e., a possible distribution of the items in dispute. The population evolves from generation to generation by means of genetic operators that change the items' distribution and thus have an effect on the fitness of the solutions. This fitness is computed from the point of view of each party, i.e., a solution that is good for one party may not be so good for another party, given that they have potentially conflicting objectives. With the successive generations, some chromosomes and their offspring tend to be more fit to one party, i.e., there are lines of evolution that evolve naturally towards the maximization of the fitness of a given party.

In this work, a *species* is defined as the group of chromosomes that encode good solutions for a given party. Thus, there will be a species for each party. Moreover, a chromosome may belong to more than one species if it denotes a solution that is good for more than one party. Evidently, chromosomes that belong to more than one species are more attractive for the problem resolution, since they correspond to solutions that will be more easily accepted by the parties.

6.1 GENERATION OF SOLUTIONS

6.1.1 *Definition of Chromosome*

A population P of size s is defined by a set of chromosomes Ch (Figure 17). Each chromosome $Ch_i, i \in 1, 2, ..., s$ represents a possible solution for the problem, i.e., a chromosome represents a distribution of the items in dispute or, in other words, who gets how much of what. For a dispute involving n parties and m issues, a chromosome Ch can be represented as an m-by-n matrix such as the one depicted in Equation 6.

$$Ch = \begin{bmatrix} V_{1,1} & \cdots & V_{1,n} \\ \vdots & \ddots & \vdots \\ Vm, 1 & \cdots & V_{m,n} \end{bmatrix} \quad (6)$$

$$P = \boxed{Ch_1 = \begin{bmatrix} V_{1,1} & \cdots & V_{1,n} \\ \vdots & \ddots & \vdots \\ V_{m,1} & \cdots & V_{m,n} \end{bmatrix}} \boxed{Ch_2 = \begin{bmatrix} V_{1,1} & \cdots & V_{1,n} \\ \vdots & \ddots & \vdots \\ V_{m,1} & \cdots & V_{m,n} \end{bmatrix}} \boxed{\cdots} \boxed{Ch_s = \begin{bmatrix} V_{1,1} & \cdots & V_{1,n} \\ \vdots & \ddots & \vdots \\ V_{m,1} & \cdots & V_{m,n} \end{bmatrix}}$$

Figure 17: Under this model a population of solutions of size s is represented as a set of chromosomes with a cardinality of s.

Under this representation, the value $V_{m,n}$ denotes the amount of issue m that party n receives in terms of the solution Ch. The correctness of the solutions is ensured by specific rules, that are generally independent of the legal domain. Let us take as an example the general model of distributive negotiation, in which a group of resources must be divided. Example scenarios of this model include divorces without children, division of inherited property or the winding up of a company. Under this model each entry stands for a value between 0 and 1 (Equation 7), and the sum of the values of each line must be 1 (Equation 8). The total amount of resources received by party n, R_n, is defined as the sum of the values of column n (Equation 9).

$$V_{m,n} \in A, A = \{x \in \mathbb{R} | 0 \leqslant x \leqslant 1\} \tag{7}$$

$$\sum_{i=1}^{n} V_{m,n} = 1, \forall m \in \{1, 2, ..., m\} \tag{8}$$

$$R_n = \sum_{i=1}^{n} V_{m,n} \tag{9}$$

This model also considers the existence of indivisible items, including items that due to their nature or to party constraints cannot be divided or sold to divide their monetary value. This is common in scenarios such as divorces, in which there are assets (e.g. car, house) that parties are not willing to sell in order to split the money, and must be assigned to one party only. Thus, for every indivisible issue m, Equation 10 must hold true.

$$\forall V_{m,i} = 1 \Rightarrow V_{m,x} = 0, \forall x \in 1, 2, ..., n, x \neq i \tag{10}$$

6.1.2 *Initialization*

Initially, some information must be provided to the system in order for it to be initialized. Figure 18 depicts the prototype of the interface used to do this. Concerning the conflict resolution process in itself, the name, value and type of each issue under negotiation must be provided, as well as a name and a color for each party. Moreover, each party must state its preferences concerning each item in dispute by distributing 100 points among the several items. This will allow the system to determine how much each party values each item, from a personal perspective (e.g. despite the monetary value of these two items, a car may have a more significant sentimental value than a house).

In terms of the genetic algorithm, some information is also needed for the initialization process. (1) A predefined number of cycles sets the termination condition. (2) The size of the population which represents the number of chromosomes at each running cycle. (3) The verbose option which, when checked, will make the algorithm detail all its steps. The algorithm can also be configured in terms of (4) how many individuals are considered from each species to create the next generation, being that a higher number of individuals will result in a higher diversity. Concerning the genetic operators, it is also possible to (5) define the weight of each one in the making of new generations, i.e., it is possible to decide on how many individuals of the new population will be generated by each operator. Finally, it is also possible to (6) configure the fitness function in terms of the weights of the personal and monetary values in the computation of the overall fitness of each solution.

Once this information is provided, the algorithm can be initialized. In this process, a population of the specified size is generated, where each chromosome encodes a random distribution of the items, i.e., a valid solution generated randomly.

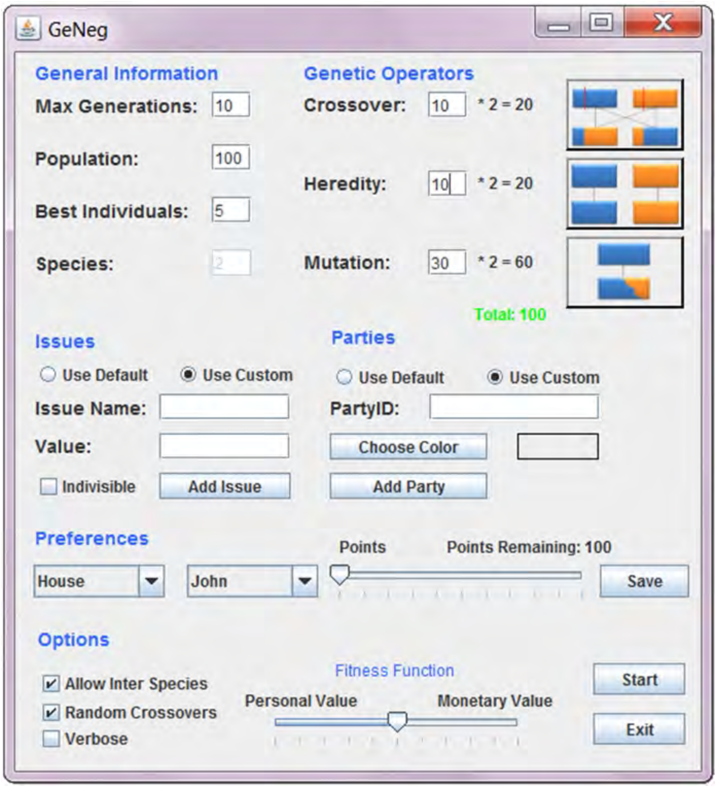

Figure 18: The prototype of the interface used to configure the genetic algorithm, including information about the parties, the issues and the weight of each genetic operator.

6.1.3 *Selection*

At each running cycle there is a part of the population that is selected from each species to give birth to a new one. Given that only the best individuals are selected to breed the new population, a fitness-based process must be defined to set which of the fitter solutions of each species may be pointed out. As each solution has a different fitness for each party, the fitness of each solution for each party must be computed. This means that for a conflict involving n parties and for a population of size s, $n * s$ values of fitness will be computed at each running cycle.

The fitness of a solution is based on the distribution of the items and on their monetary and personal values.

The fitness function takes each solution and assigns it a value that depends on three issues: (1) the amount of items that each party receives, (2) the economic value of the items and (3) the personal value. Two fitness functions were initially devised, as defined in Equation 11 and Equation 12, where.

- I defines the number of issues;
- mv_i stands for the monetary value of issue i;
- tmv (total monetary value) denotes the economic value of the case, i.e., the total amount of money that the issues in dispute are worth, being defined as $tmv = \sum_{i=1}^{I} mv_i$;
- $fit_{j,p}$ represents the fitness of chromosome j for party p;
- W_m denotes the weight of the monetary value while W_p stands for the weight of the personal value;
- $prefs_i$ denotes the preferences of a given party regarding issue i;

The use of Equation 11 will tend to result in solutions in which each party receives approximately what he valued the most, i.e., it minimizes the difference between the individual preferences for each item and the personal value assigned to it.

$$fit_{j,p} = W_m * \frac{\sum_{i=1}^{I} Ch_{j,p} * mv_i}{tmv} + W_p * \left(1 - \sum_{i=1}^{I} \frac{|Ch_{j,p} - prefs_i|}{I}\right) \tag{11}$$

On the other hand, Equation 12 focuses on maximizing the individual gain of each party, i.e., Equation 11 will tend to generate populations that may be described as fair, while the ones generated by Equation 12 will be more greedy.

$$fit_{j,p} = W_m * \frac{\sum_{i=1}^{I} Ch_{j,p} * mv_i}{tmv} + W_p * \sum_{i=1}^{I} \frac{|Ch_{j,p} - prefs_i|}{I} \quad (12)$$

The system is currently using Equation 11 as it generates solutions that, by being more balanced, are more likely to be accepted by the parties. At each cycle, the fitness of the population is computed and the best individuals from each species are selected to generate a new population.

6.1.4 *Reproduction*

In genetic algorithms, reproduction concerns the creation of new populations, making the heuristic search move towards the maximization of the fitness function. In this work three different genetic operators are used: mutation, crossover and heredity. They are applied to each generation of the population according to what was defined during the initialization process.

6.1.4.1 *Mutation*

In genetics, a mutation is defined as a spontaneous and random change in a genomic sequence. In the particular case of this work, a mutation is a random change at the chromosome level, i.e., in terms of the issues' distribution. The issues are mutated according to a mutation threshold, here designated as μ.

To implement the mutation, a random issue is selected as well as two random parties. The allocation of the issue is then changed for the selected parties, according to the mutation threshold. If the issue is divisible, the amount of the issue is subtracted from one party according to μ and added to the other party. If the issue is indivisible, there is a probability that a change of ownership may occur, according to the mutation threshold. Once a new chromosome is generated its validity is checked in order to determine if all the invariants hold. If the chromosome is not valid (e.g. contains a negative value) it must be generated again. This process repeats until the chromosome generated is valid, which eventually happens. Let us consider the scenario in which two parties, with a given μ, are disputing four issues. Assuming that issue 2 is divisible and is selected to be changed between party 1 and party 2, the result is shown in Equation 13, where Ch

Listing 3: Description of the Mutation algorithm

```
Algorithm Mutation is
Input: List of parties, L
       List of issues, I
       Parent chromosome, C
Output: A new chromosome, C'
Do
   i  := select random issue from I
   p1 := select random party from L
   p2 := select random party from L such that p1 != p2
   C' := C
   if (i is divisible)
      C'i,p1 := C'i,p1 + μ * C'i,p1
      C'i,p2 := C'i,p2 - μ * C'i,p2
   else if (randomNumber > μ)
           temp := C'i,p1
           C'i,p1 := C'i,p2
           C'i,p2 := temp
While (C' is invalid solution)
Return C'
```

and Ch' represent a chromosome, respectively, before and after the operation.

$$Ch = \begin{bmatrix} V_{1,1} & V_{1,2} & V_{1,3} \\ V_{2,1} & V_{2,2} & V_{2,3} \\ V_{3,1} & V_{3,2} & V_{3,3} \\ V_{4,1} & V_{4,2} & V_{4,3} \end{bmatrix} Ch' = \begin{bmatrix} V_{1,1} & V_{1,2} & V_{1,3} \\ V_{2,1+\mu} & V_{2,2} - \mu & V_{2,3} \\ V_{3,1} & V_{3,2} & V_{3,3} \\ V_{4,1} & V_{4,2} & V_{4,3} \end{bmatrix} \quad (13)$$

The mutation, as defined here, has an effect on the fitness of the solution for each party. That is, the new solution will be more favourable to party 1 and less favourable to party 2. Below the mutation algorithm is described.

6.1.4.2 *Crossover*

Crossover is a binary genetic operator by means of which two offspring are generated from two parent chromosomes. In this work, a two-point crossover technique is used. Concisely, the two points that the name of the technique refers to define the beginning and the end of an issue for both parents, meaning that everything between the beginning and the end of the issue will be swapped to create two offspring. Crossover consists in swapping two distributions of the same issue, generating two new solutions.

Two different approaches have been implemented and can be selected that will influence the variety of the new offspring: *inter species* and *random parents*. If the *inter species* option is set to true, chromosomes from different species can be crossed, generating a more diverse population. If the option is not selected, only chromosomes from the same species will be crossed. The *random parents* option will allow one to decide on which parents should be used. If it is set to true, random parents will be selected. On the other hand, if it is set to false, only the best parents of each species will be selected for the crossover operation. Equation 14 depicts an example of the use of crossover in two parent chromosomes Ch1 and Ch2 to generate two offspring Ch1′ and Ch2′. In this example the distribution of issue 2 was swapped.

$$\begin{gathered} Ch1 = \begin{bmatrix} A & B & C \\ D & E & F \\ G & H & J \\ I & K & L \end{bmatrix} Ch2 = \begin{bmatrix} M & N & O \\ P & Q & R \\ S & T & U \\ V & W & X \end{bmatrix} \\ Ch1' = \begin{bmatrix} A & B & C \\ P & Q & R \\ G & H & J \\ I & K & L \end{bmatrix} Ch2' = \begin{bmatrix} M & N & O \\ D & E & F \\ S & T & U \\ V & W & X \end{bmatrix} \end{gathered} \tag{14}$$

Given that this technique changes the distribution of each solution, it will have effect on the fitness function. Below it is given a description of the algorithm that implements the crossover technique being used.

6.1.4.3 *Heredity*

In genetics, heredity can be defined as the passing of traits from parent to offspring, either positive or negative. Once this process is finished, the offspring acquires characteristics that may be compared to the ones of the parent. The evolution of the species is thus achieved by accumulating variations exhibited by different individuals. In this work, heredity is the simplest genetic operator to be considered, in the sense that the new offspring have exactly the same distribution that its parents, i.e., this operator generates no diversity at all and should be used when one wants to avoid "losing" the best individuals of a population, ensuring that the best traits will be passed to the next generation.

6.1.5 *Termination*

The process of selection and reproduction is repeated until the maximum number of generations stated in the initialization phase or a

Listing 4: Description of the Crossover algorithm

```
Algorithm Crossover is
Input: List of parties, L
       List of issues, I
Output: New chromosomes, C1', C2'
i  := select random issue from I
if (interspecies)
   s1 = select random species
   s2 = select random species such that s1 != s2
   if (randomparents)
      C1 := select random ch from s1
      C2 := select random ch from s2
   else
      C1 := select best ch from s1
      C2 := select best ch from s2
else
   s1 = select random species
   if (randomparents)
      C1 := select random ch from s1
      C2 := select random ch from s1 such that C1 != C2
   else
      C1 := select best ch from s1
      C2 := select second best ch from s1
swap issues and generate C1', C2'
return C1', C2'
```

minimum level of fitness for at least one solution for each party is reached. At this point, the system may give a picture of the state of the evolutionary process, in terms of the solutions so far attained, and their lines of evolution, as depicted in Figure 19.

Each solution is represented by a rectangle in one or more colors. A solution with a given color means that it belongs to the species (party) of that color, i.e., it makes part of the group of the best solutions for that party. This will allow one to see the natural emergence of the lines of evolution that tend towards the maximization of the fitness value for a given party. Colourless rectangles denote solutions that are not among the most fit for any particular party but generate offspring that are.

It is possible to select chromosomes to have access to further information concerning the solution it encodes, including its fitness for each party, the average value of fitness or the parent solutions. In the interface the lines between individuals represent the parent-offspring relationships. A unary genetic operator (such as mutation or heredity) generates individuals that have a single line connecting to the previous generation, while an individual that has two lines was generated by crossover.

The tree generated can then be used by a mediator, either electronic or human, to easily access potential alternative solutions for the prob-

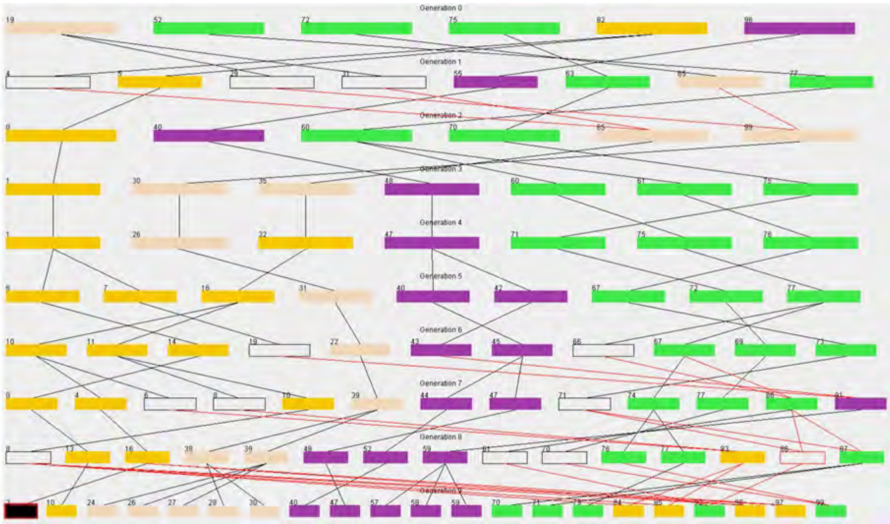

Figure 19: The lines of evolution of the genetic course and their outcome. Each coloured rectangle represents a solution, being the color representative of the party for whom the solution is better. Different species of solutions are visible.

lem. This will support a dynamic conflict resolution model, in which the strategy devised at the outset of the process can be easily adapted, in real time, to respond to significant changes in the state of the parties.

To support this, the system is able to perform two main operations on the populations produced, namely *jumps* and *walks*. In a jump, a solution is selected from a different species (Figure 20). This may happen in a scenario in which the mediator realizes that a given party is being penalized or when he is getting close to abandoning the process. In that sense, the process will jump to the species that has more constructive solutions to that party.

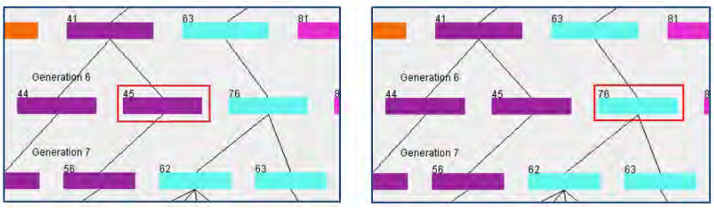

Figure 20: The mediator (electronic or human) may change from a solution that is better for one party to a solution that is better for another one.

Walks, on the other hand, are used to move within the same species (Figure 21). This usually takes place when the state of one or more parties changes slightly. Walks down the tree are performed in order to search for solutions that are better for the current party (given that the fitness of solutions improves as the generations increase) and worse for the remaining ones. Walks up the tree are performed with the opposite objective.

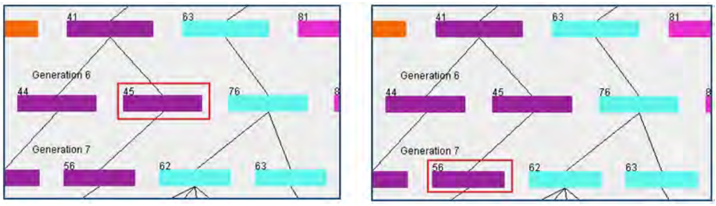

Figure 21: The mediator (electronic or human) may move inside a species' tree in order to reach a solution that is better (or worse) for a given party, thus worse (or better) for the remaining ones.

6.2 SUMMARY

At the outset of this PhD, a wide range of existing ODR tools were analysed. One of the objectives of this analysis was to assess current approaches on the generation of solutions for the disputes. This may be a somewhat complex problem, involving several issues (with specific traits), several strategies and several parties, each one with his own perspective, point of view and personal evaluation of the issues.

The approaches analysed could be generally organized in two groups: case-based and rule-based. While case-based solutions try to suggest solutions for a case based on the observation of past similar cases, rule-based ones do it by applying rules that shape the domain of the problem. Either approach has advantages but also disadvantages: they are closely related to the legal domain, will fail if not enough/variate cases of the domain exist and do not cope well with changes in the legal rules.

In order to tackle the problems identified, an approach based on genetic algorithms was devised. In this approach, each chromosome encodes a possible and valid solution for the problem. A population of chromosomes is evolved in each cycle, through the application of genetic operators that change the value of fitness of the population. This approach has three main advantages:

- It is a more complete solution when compared with case-based approaches, as it may cover virtually the entire search space. Case-based approaches often fail when the cases describing a domain are not variate enough;

- The number of generated solutions depends only on the initialization settings;

- It is domain-independent, contrary to case and rule based approaches in which cases and rules can, generally, only be reused in the same legal domain.

This algorithm results in the natural emergence of lines of evolution that tend to generate solutions that are better for a given party. One of these lines exists for each party, in each problem. This allows a mediator to move in and/or between branches of the evolution tree in search for a solution he might find more suited in an instant, given the *current* characteristics of the problem.

A solution that may appear to be suited in the beginning of the process may not be so through all its lifecycle.

In the last paragraph, the word *current* was emphasized specifically to point out that the characteristics of a given problem may not be static: they change as the process develops. This means that a solution that may appear to be suited in the beginning of the process may not be so through all its lifecycle. This is a very important notion, that many of the current approaches fail to acknowledge. The approach presented in this section supports this by generating a wide range of solutions, ensuring that even when the characteristics of the problem change, alternative solutions will be available. In the next section the other part of the problem is addressed, i.e., how does one know when the characteristics of a given problem have changed so significantly that it calls for a rethinking of the strategies?

7

BEHAVIOURAL ANALYSIS IN ODR

Behaviour is a mirror in which every one displays his own image.
— Johann Wolfgang von Goethe

Until now, several different approaches have been followed by researchers to implement Online Dispute Resolution (ODR) tools, some of them trying to take advantage of powerful mechanisms made available by the newest information and communication technologies (Katsh and Rifkin, 2001). Specifically, interest has grown in the use of Artificial Intelligence models and techniques that include but are not limited to Argumentation, Game Theory, Heuristics, Intelligent Agents and Group Decision Systems, as described by (Peruginelli and Chiti, 2002; Lodder and Thiessen, 2003; Bellucci et al., 2004; Andrade et al., 2007; Marreiros et al., 2010), mainly aiming at supporting the parties' decision making process.

In this section a new line of research is put forward, based on a relatively recent sub-field of Artificial Intelligence: Ambient Intelligence (Aarts and Grotenhuis, 2011; Augusto, 2009). Under this paradigm, computational power is seamlessly embedded into the environment, ultimately creating computational environments that perform their actions in an ideally transparent way for the user (Cook et al., 2009). It represents a major shift, placing the user in the center of the interaction model, with intelligence surrounding him. It is thus a step forward in Artificial Intelligence, making "intelligence" seamlessly available through actions on the environment rather than through complex software interfaces (Ramos et al., 2008).

The main innovation of this PhD work is thus the concept of a Conflict Resolution Environment (CRE), rather than the traditional Conflict Resolution Platform. This environment supports the already traditional conflict resolution tools by providing important context information derived from the environment. This context information is very important for disputant parties and mediators alike to take better decisions. Specifically, two interesting contextual aspects are addressed in this work: the level of stress of the participants and their personal conflict handling style. The aim is to develop methodologies that allow to acquire such information in a non-invasive and transparent way.

The main innovation of this PhD is the concept of a Conflict Resolution Environment rather than the traditional Conflict Resolution Platform

In this approach, from the point of view of the user, he still interacts with a regular conflict resolution platform. However, he does it from within an intelligent environment that is transparently supporting the conflict resolution tool with additional information that may be of importance for the evolution of the conflict resolution. This information, that is transparently provided to the conflict resolution platform, potentially without the conscious interference of the user, will allow developing more complete conflict resolution models that, by being richer in information, will more realistically shape the ones undertaken by human experts.

7.1 THE IMPORTANCE OF ANALYSING BEHAVIOUR

In a pursuit of better remote communication frameworks, our society relies more and more in Virtual Environments (VEs). (Jim Blascovich, 2002) define VEs as synthetic sensory information that lead to perceptions of environments and their contents as if they were not synthetic. In other words, VEs may be seen as simulated environments that, in some way, try to look like the real environments being simulated, with the aim of implementing some kind of interaction scenario. Typical fields of application of VEs include teaching in classrooms, informal learning, distance learning, business, e-commerce, gaming, real-life simulation or conflict resolution. However VEs, in their current form, are still not a suitable replacement for traditional face-to-face communication.

In fact, a VE is frequently regarded as "cold", with emotions and other traces of our complex interaction modalities playing little to no role at all. One of the most important aspects here is that of body language. In their day-to-day interactions people unconsciously rely on body language to express themselves in a richer way. (Mehrabian, 1980b) concludes that in face-to-face communication there are three key elements: the words, the tone of voice and the non-verbal behaviour. The author also concludes that the non-verbal elements are particularly important for communicating feelings and attitudes, stating that they account for the majority of the information transmitted. That is to say: the way words are said is more important than words themselves.

The way *words are said is often more important than words themselves*

In a related line of research, (Dodds et al., 2011) conclude that the lack of gestural information from both speaker and listener limits successful communication in VEs. The authors experimentally prove not only that body language is very important for transmitting information but it is also important to perceive feedback from that transmission, i.e., to perceive if the communication is being successful or a different approach should be followed. Both the lack of feedback from the environment and meaningful content are also pointed out as a drawback by other researchers (Campbell, 1997; Marucci et al., 2001).

To deal with this issue, several approaches can be found. Recently, (Alsina-Jurnet and Gutiérrez-Maldonado, 2010) analysed the influence of five user characteristics - test anxiety, spatial intelligence, verbal intelligence, personality and computer experience - on the sense of presence. Also, (Rehm et al., 2008) deal with the idea of the analysis of the user's behaviour and interpretations regarding the cultural background, using accelerometers to uncover the user's cultural background by analysing his/her patterns of gestural expressivity in a model based on cultural dimensions. (Jaimes and Sebe, 2007) describe the concept of multimodal interaction as a way to communicate between humans and computers using more than one modality or communication channel (e.g., speech, gesture or writing).

Also important is the affective aspect of communication (Beale and Creed, 2009; Hudlicka, 2003). Emotions appear in almost all models of human communication, measurable in facial expressions, gestures, voice tone, respiration rate, skin temperature, and so on. Moreover, depending on the emotions, the message changes: once again, the most important is not what is said, but how it is said. As noted by (Picard, 2000) affect recognition is most likely to be accurate when it combines multiple modalities, information about the user's context, situation, goal, and preferences.

The importance of stress must also be considered. Evidently, stress is a very important factor in interpersonal communication, as it is in virtually any other aspect in our lives. However, current approaches on VEs lack stress models that can support it. This constitutes an obstacle to effective communication between the participants. In fact, research on stress applied to VEs does not exist. Bearing in mind that this PhD work focuses in the conflict resolution domain, the importance of the personal Conflict Handling Style should also be considered. As described in Section 1.2, each individual has a particular way of behaving before a conflict. However, these behaviours can be classified according to their degree of assertiveness and cooperativeness. Until now, to understand one's conflict handling style, a mediator would use a written questionnaire that would define how a party would react in specific situations, thus defining their conflict handling style. This, however, is not suitable for the virtual domain thus, this important information, is also not considered in present ODR tools.

The loss of all this context information in virtual settings makes it hard for the intervening parties to understand the state of each other in each moment. This constitutes the motivation for this research. In fact, when communicating online, people tend to forget that there is another person behind the screen on the other side. In that sense, there is a disinhibition effect and people tend to forget about social restrictions and the other's feelings and simply do not worry that much about the consequences of the words they utter and the actions they commit. Thus, it is often easier to offend people online. However, the advantages of this approach go further. Take as an example a party that is very stressed or feeling extremely angry at a given time in the dispute resolution. He should not take important or binding decisions as he might regret them later. Instead, he should be advised to make a pause and think about it later on. The development of ODR systems that are indeed able to understand the state of the different parties is thus interesting from more than one perspective.

The main aim of this section is thus to describe the techniques developed that allow such context information to be acquired from the real environment and be passed onto the VE, enabling a more accurate description of the real individual in the virtual setting, supporting better decisions. Specifically, emphasis is placed on the interpretation of the parties' personal conflict handling style and on the assessment of their stress level, given that these are two of the most important factors in conflict resolution. This section thus constitutes the higher level one and closes the cycle that started with the compilation of important knowledge and now ends with the acquisition of information that allows to characterize the state of the parties, in real-time, in a non-invasive way. It is the most innovative part of this PhD thesis as it opens a new line of research in which contextual (and often subjective) information is regarded as important, contrary to what was postulated until now by researchers in the field.

This section puts forward an controversial idea in which contextual information is regarded as important, contrary to what was postulated until now by researchers.

7.2 SYSTEMATIC BEHAVIOURAL ANALYSIS

The behaviour of a living organism, system or artificial entity, includes all the activities that the organism would not perform if it were not "living", with living being interpreted in the classical sense but also in the sense of, for example, a software lifecycle.

These activities are always the response of the system to some stimuli, or lack of them. Stimuli are said to be internal, when they are originated inside the "body" of the system (e.g. thought, pain, change of state). They are said to be external when originated outside of the "body" of the system and perceived by some receptor cell (e.g. change in the temperature, visual change in the environment, reception of a given message).

The behaviour of an organism includes all the activities that the organism would not perform if it were not living

The system may be conscious or unconscious of the perceived stimuli. Despite this, the system may respond. This response is said to be voluntary when the system undergoes some reasoning process before acting or involuntary when the system reacts in pre-determined ways, without reasoning about it.

Nothing characterizes a system better than its behaviour. Knowing how the system reacts to each stimulus allows the prediction of future states. Controlling stimuli allows to, indirectly, control the system. This is the base notion of behaviourism. Thus, in psychology, behaviour is analysed and influenced in order to address some behavioural issue of the individual, ranging from psychological disorders to other issues such as smoking habits or eating disorders, just to name a few.

In the specific context of this PhD work, the interest is on knowing how a given party responds to specific scenarios: how does a party behave when under stress, how does a party behave during a negotiation, among others. If the mediator has access to this information, he will be able to take better decisions. As an example, if a mediator knows that a given party generally assumes a highly competitive style during a negotiation, he may try to show that party that such a style might be an obstacle for a successful outcome.

The approach described in the following sections focuses on acquiring context information that allows to characterize the behaviour of the human users of the ODR tool. Moreover, it does it in an absolutely transparent and non-invasive way, i.e., rather than relying on traditional self-reporting mechanisms such as questionnaires in order to infer behaviours, it analyses the actions of the parties.

In order to implement such processes, procedures from the social science were analysed that would allow to study and, if needs be, influence the behaviour, in a systematic and scientifically valid way. Specifically, a procedure defined by (Cooper et al., 1987) was followed. It provides a complete description of the procedures and principles necessary to (1) correctly identify the sources for a behaviour, (2) understand the relationship between sources and behaviours and (3) how to change them in order to influence the behaviour as desired.

According to the authors, all behavioural studies should include the following:

1. At least one participant;
2. At least one behaviour (which is the dependent variable);
3. At least one setting or environment;
4. A system for measuring the behaviour and ongoing visual analysis of data;
5. At least one treatment or intervention condition;
6. Manipulations of the independent variable so that its effects on the dependent variable may be quantitatively or qualitatively analysed;
7. An intervention that will benefit the participant in some way.

These issues thus served as the guiding lines that directed the work here described, as the following subsections detail. The actions implemented in order to address each of the points enumerated are summarized at the end of the chapter.

Figure 22: Devices selected to build the test environment. These specific devices were selected because they allow to acquire some kind of information detailing the behaviour of the user.

7.3 AN ENVIRONMENT FOR ACQUIRING BEHAVIOURAL FEATURES

To implement the research plan a test environment was set up in the Intelligent Systems Lab[1] of the University of Minho. In this environment, the user is isolated from external stimuli. Data about his behavioural patterns is collected while he performs the requested tasks. The environment can be configured to induce stress in different levels.

This environment allows to analyse user behaviour in a transparent and non-invasive way.

The final aim is that this environment extends the UMCourt conflict resolution platform as an intermediate layer that receives sensory information and converts it into meaningful knowledge about the context of interaction. This knowledge is then used by the mediator to take better decisions, and by the parties to receive feedback from the environment.

All this process is completely transparent to its users. That is, the data about the users' context is collected, processed and transparently added to the operational data. Thus, when parties participate in the technology-supported conflict resolution they sign in a usual virtual mediation or negotiation room, or they may even just be using traditional web forms. Anyway, the conflict resolution platform and the mediator have a way to assess their state when taking their decisions that allows for a better management of the whole process.

Figure 22 depicts the devices that are installed in the user-area network and make up the environment. This specific devices were selected because they are available in the market and can provide some kind of feedback about the user's behaviour, when appropriate software is developed to handle them. Another requirement was that the devices were common, i.e., that their use or their presence in the environment would not intimidate or awe the user. Table 2 briefly describes each device and the main characteristics of interest.

1 The website of the Intelligent System Lab is available at http://islab.di.uminho.pt *(accessed in June, 2013)*

DEVICE	BRIEF DESCRIPTION	MAIN FEATURES
HP Touchsmart	All-in-one PC	touchscreen, web cam, large screen
Samsung Galaxy Tab	Tablet PC	touchscreen, web cam, accelerometer, relatively large screen, mobile, Android OS
HTC PDAs	Smartphones	touchscreen, camera, accelerometer, mobile, Android OS
Sony FCB-EX780BP	25x Super HAD PAL Color Block Camera with External Sync	25x Optical Zoom, Image stabilizer, Day/Night Mode, Privacy Zone Masking

Table 2: Brief description of the functionalities of the devices that constitute the environment.

The following features, acquired from the respective devices, are considered:

TOUCH PATTERN the touch pattern represents the way in which a user touches the device and constitutes a variation of intensity over a period of time. This information is acquired from touch screens with support for touch intensity.

TOUCH ACCURACY a comparison between touches in active controls versus touches in passive areas (e.g. without controls, empty areas) in which there is no sense in touching. This information is acquired from touch screens.

TOUCH INTENSITY the intensity of the touch represents the amount of force that the individual is putting into the touch. It is analysed in terms of the maximum, minimum and mean intensity of each touch. This information is acquired from touch screens.

TOUCH DURATION this represents the time span between the beginning and the end of the touch event. This data is acquired from devices with touch screens.

AMOUNT OF MOVEMENT the amount of movement represents how and how much the user is moving inside the environment. An estimation of the amount of movement from the video camera is built, provided by the INT3-horus framework Figure 23. A communication module was implemented to integrate this framework with UMCourt . The image processing stack uses the principles established by (Castillo et al., 2011) and uses image difference techniques to calculate the amount of movement between two consecutive frames (Fernández-Caballero et al., 2010).

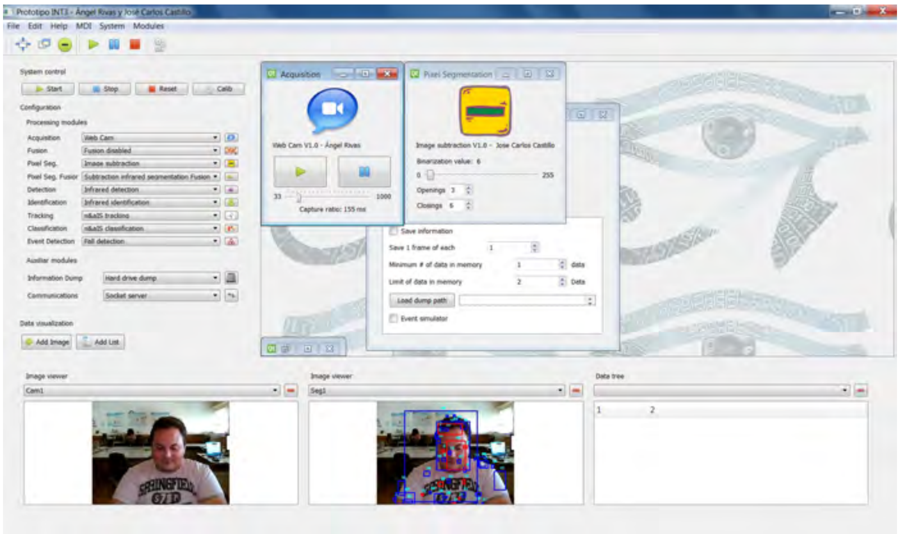

Figure 23: Interface of the INT3-horus: a multilevel framework for intelligent multisensor monitoring and activity interpretation.

ACCELERATION the acceleration is measured from accelerometers in mobile devices. It is useful for building an estimation of how much an individual is moving, how he is doing it (e.g. is the user having sudden movements?) and how he holds the hand-held device. Moreover, information from the accelerometer is used to support the estimation of the intensity of touch. This device is collected from smartphones or other hand-held devices.

SCORE this feature quantifies how well the individual fares in the tasks that are given to him.

STRESSED TOUCHES this feature describes which touches are classified as stressed and which are classified as calm, according to their intensity curve over time.

In this PhD work, these features are analysed in search for significant changes due to the influence of stress. When it is known how an individual reacts to stress in terms of these features, it is possible to build a real-time stress detection and assessment module that relies on the analysis of changes in these features to compute a value of stress.

7.4 ANALYSIS OF CONFLICT HANDLING STYLE

Conflict and personal Conflict Handling Styles have already been addressed in detail in Section 1.2. They encode the way each individual reacts before a conflict, in terms of their tendency to be more cooperative or competitive. Information about how each party reacts before a conflict is particularly interesting for a mediator to better manage the conflict resolution process. In that sense, this section explores the hypothesis of automatically classifying the personal conflict handling style of parties.

Conflict handling styles encode the way each individual reacts before a conflict.

In fact, a conflict resolution tool that has the ability to make a prediction about the evolution of the parties' personal conflict handling style may bring along several advantages, namely from the point of view of the mediator. Generally, this can be done by following two different approaches: either questioning the parties or analysing their behaviour.

The first approach will provide information before the start of the process, being possible to plan ahead. However, the main disadvantage is that it is easy to lie and fake a behaviour in an attempt to undermine the process. Moreover, once the (potentially stressful) conflict resolution starts, parties are likely to change their behaviour.

The second approach will gradually provide information as the process evolves. Although it may be a slower way of building knowledge about the personal conflict style, it will more reliably reflect the conflict resolution style and, more important than that, will reveal eventual changes in real time.

In this work, the potential of the second approach is explored. In that sense, the actions of the parties are analysed in each stage of the conflict resolution process. Moreover, the nature of the solutions proposed is also taken into consideration (e.g. is a party being too greedy?, is a party being realistic?), framed within the values of the BATNA, WATNA and ZOPA. Specifically, a negotiation game was developed to be played with the parties inside the environment. This allows not only to assess the personal conflict handling style of each individual but also to correlate that information with stress, which results interesting. The experiment performed is described in Section 7.4.2 and the results are described in Chapter 8.

7.4.1 *From the Computation of Utility to the Classification of the Personal Conflict Handling Style*

In this approach the potential relation between the personal conflict style and the utility of the proposals is explored. The utility quantifies how good a given outcome is for a party. In that sense, it is acceptable to argue that a competing party will generally propose solutions that maximize its own utility in expense of that of the other party, while for example a compromising party will most likely search for solutions in an intermediary region. Essentially, the aim is to classify the personal conflict style of parties by constantly analysing the utility of the proposals they create. The relation between the utility of the proposals and the conflict style is introduced in Figure 24.

In this approach the potential relation between the personal conflict style and the utility of the proposals is explored.

During the negotiated process, parties make successive proposals and counterproposals in order to achieve a mutually agreeable solution. It is thus possible to make an analysis of the proposals of

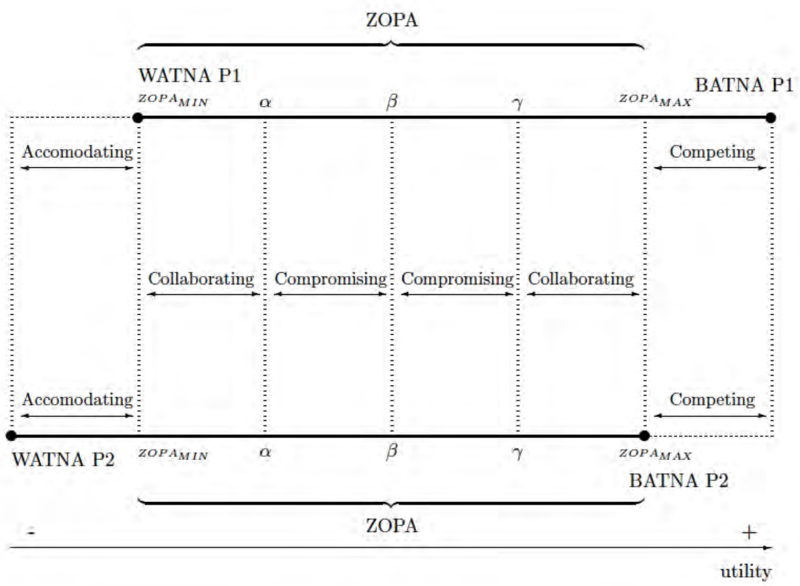

Figure 24: The mapping of a value of utility of a proposal into the corresponding personal conflict handling style of each party.

each party and, together with additional information (e.g. the space defined by the BATNA and WATNA of each party), classify the behaviour of the party in terms of its personal conflict style.

Each action that a party performs contributes to the overall characterization of his personal conflict handling style. In that sense, the style that is computed for each party in each round is a result of all the previous interactions, although the weight of older interactions decreases exponentially. Two main scenarios are possible: the party ignores the proposal or the party answers to the proposal. If a party, upon receiving a proposal for a solution, simply ignores it, he is not satisfying his interests nor the ones of the other party. In such a scenario, the conflict style evidenced is the *Avoiding* one.

If the party makes a proposal or a counterproposal, he is cooperating on the process. However, the nature of the proposal must be analysed, namely in what concerns its utility for each party. When the utility of the proposal is higher than the BATNA of the other party, he is clearly showing a *Competing* style as he is trying to maximize his own gain, in a way that is potentially unrealistic and disregards the other party. On the other hand, if the utility of the proposal is lower than the WATNA of the other party, he is neglecting his own gain or even maximizing the gain of the other party. In such a limit scenario, it is reasonable to state that the party is evidencing an *Accommodating* behaviour.

The utility of the proposal falling within the range of the ZOPA indicates that the party is being reasonable and realistic and is trying to propose a settlement in which both parties will not win everything they could but will not lose everything either. In such a scenario, the conflict style is determined according to the distance to the mean point of the ZOPA, as defined in Equation 15.

$$\beta = \left(\frac{\mathrm{ZOPA}_{\mathrm{MIN}} + \mathrm{ZOPA}_{\mathrm{MAX}}}{2}\right) \tag{15}$$

Two additional points are defined that allow the classification of the remaining conflict styles, as depicted in Equation 16 and Equation 17, by defining additional intervals.

$$\begin{aligned} \alpha &= \left(\mathrm{ZOPA}_{\mathrm{MIN}} + \frac{\beta - \mathrm{ZOPA}_{\mathrm{MIN}}}{2}\right) \\ &= \left(\frac{\mathrm{ZOPA}_{\mathrm{MIN}} + \beta}{2}\right) \end{aligned} \tag{16}$$

$$\begin{aligned} \gamma &= \left(\mathrm{ZOPA}_{\mathrm{MAX}} + \frac{\mathrm{ZOPA}_{\mathrm{MAX}} - \beta}{2}\right) \\ &= \left(\frac{\mathrm{ZOPA}_{\mathrm{MAX}} + \beta}{2}\right) \end{aligned} \tag{17}$$

When the utility of a proposal falls within the range $[\alpha, \gamma]$, the party is trying to negotiate in intermediary points of the ZOPA, i.e., the party is trying to compromise, which implies a loss from both parts. In such a scenario, the behaviour of the party is classified as *Compromising*. On the other hand, if the value of the utility belongs to the range defined by $[\mathrm{ZOPA}_{\mathrm{MIN}}, \alpha[\cup]\gamma, \mathrm{ZOPA}_{\mathrm{MAX}}]$, the party is proposing a solution that is closer to the limits of the ZOPA. This is interpreted as the party trying to work out a mutually agreeable solution, although he may be trying to explore the weaknesses of the opposing party by trying to force him to accept a given solution. Behaviours such as this one are classified as *Collaborating*.

Nevertheless, as depicted in the literature (Munduate et al., 1999) and as evidenced by our own daily interactions, we do not make use of a single conflict handling style throughout a conflict resolution process. More likely, we evidence a combination or sequence of several styles. In that sense, in order to more accurately define the boundaries, we propose an approach in which a main conflict style is inferred, accompanied by a potential trend style. This means that a party shows a given style with a possible tendency towards another one. The notation used to denote a main conflict style with a trend to a secondary one is as follows: $\mathtt{Main}_{\rightarrow \mathtt{secondary}}$. Let φ be the value of the utility of a proposal. The following additional styles are defined:

$$\text{Compromising}_{\to\text{Accomm}}, \text{if}\varphi \in \left[\text{ZOPA}_{\text{MIN}}, \frac{\text{ZOPA}_{\text{MIN}} + \alpha}{2}\right[$$
$$\text{Collaborating}_{\to\text{Comprom}}, \text{if}\varphi \in \left[\frac{\text{ZOPA}_{\text{MIN}} + \alpha}{2}, \alpha\right[$$
$$\text{Compromising}_{\to\text{Collab}-\text{Accomm}}, \text{if}\varphi \in \left[\frac{\text{ZOPA}_{\text{MIN}} + \alpha}{2}, \alpha\right[$$
$$\text{Compromising}_{\to\text{Collab}-\text{Compet}}, \text{if}\varphi \in [\alpha, \beta[$$
$$\text{Collaborating}_{\to\text{Comprom}}, \text{if}\varphi \in \left[\gamma, \frac{\text{ZOPA}_{\text{MAX}} + \gamma}{2}\right[$$
$$\text{Collaborating}_{\to\text{Collab}-\text{Compet}}, \text{if}\varphi \in \left[\frac{\text{ZOPA}_{\text{MAX}} + \gamma}{2}, \text{ZOPA}_{\text{MAX}}\right] \quad (18)$$

By determining the personal conflict style of each party in each round, it is possible to analyse its temporal evolution throughout the conflict resolution process Figure 25. When the conflict resolution platform or the mediator have a temporal representation of the evolution of the conflict styles, it is possible to not only see this evolution but also predict future instances. This constitutes the principle for adaptive, dynamic and context-aware conflict resolution models.

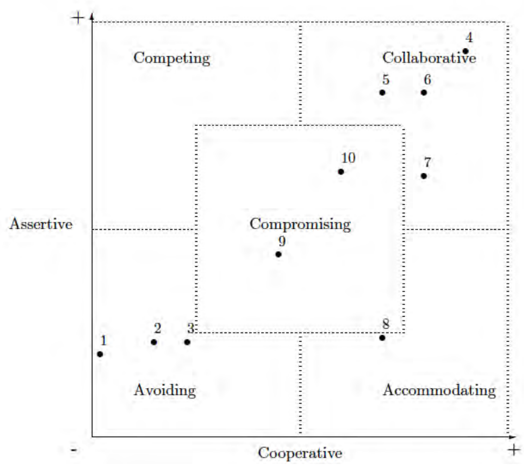

Figure 25: A temporal evolution of the conflict handling style allows a mediator or a conflict resolution system to not only analyse the evolution of the parties' behaviour towards the conflict resolution process but also to predict potential future attitudes. Each point in the Cartesian plane and its corresponding number represents the style computed and the respective round.

7.4.2 *Methodology and Data Collection*

This experiment was implemented with two main aims in mind. On the one hand it aims at validating the approach described in the previous sub-section, showing that it is possible to assess the personal conflict handling style of an individual without the use of questionnaires. On the other hand, given that it takes place inside the context acquisition environment developed, it allows to acquire information about the behavioural features of the participants and later correlate it with the conflict handling style.

In that sense, the collection of the data was organized into two phases. In a first phase, participants played the negotiation game in a stress-free environment. In a second phase, they played it under the effect of stressors such as the vibration of the devices, loud and annoying sounds, unexpected behaviours of the devices, among others. This allows to later assess the effect of stress not only on the conflict handling style but also on the outcome of the negotiation.

The developed negotiation game (Figure 26) simulates a business scenario where each party must achieve a desired result in the negotiation or go bankrupt. The desired result was a win/win situation for both parties. The game starts with the application randomly giving one of the predetermined roles to each party. The instructions to win the game were to negotiate a successful deal and make sure that the party in question didn't go bankrupt. Each party's instructions were clearly presented, visible to them through the application interfaces. The objectives and the *persona* for each party were given as:

role a party A is a light bulb manufacturer who specializes in specific light bulbs. He is not the only supplier of this light bulb. In order to stay in business, he needs to sell 6,000 light bulbs at 1 euro or more per light bulb. If he does not achieve this, he will go bankrupt. He is also given the information that Party B needs to make this deal.

role b party B is a designated retailer of light bulbs. He has recently received a contract to supply a hotel chain 6,000 of these specific light bulbs. The hotel is prepared to pay 2 euros per light bulb. If Party B does not manage to negotiate with Party A to buy the light bulbs at 1,20 or less, he will go bankrupt. Party B is told that party A is in a little financial trouble and needs to make the deal to survive.

The game lasts at most ten rounds (in the training phase) or five minutes (in the stressful phase). If a successful outcome is not achieved in this time, both parties go bankrupt. The ZOPA is bounded by the values of the BATNA (1 euro) and the WATNA (1,20 euro). The range of possible agreement is 0.20, but this information is not explicitly provided to the parties.

While the parties play the game, their messages as well as their proposals are recorded, together with the behavioural features of interest. This allows to later determine how each party behaves and negotiates under stress. The data collected in this experiment is described in Table 3. It was analysed statistically, in a process that shapes the one

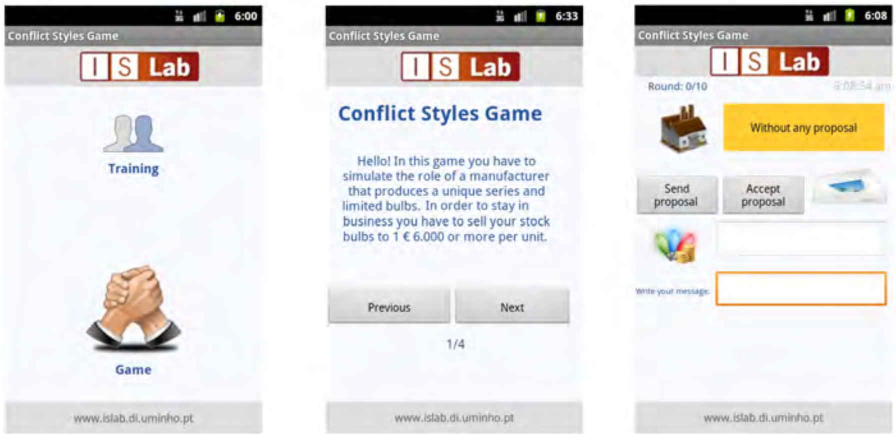

Figure 26: Some screenshots of interfaces from the Negotiation game implemented to analyse the conflict handling style together with the level of stress.

described in Section 7.5.3. This allowed to determine how each party behaves during the negotiation, under the existence or absence of stressors. The results of this study are described in Chapter 8.

7.5 ANALYSIS OF STRESS

Stress is another important factor in conflict resolution, much like in any other activity in our lives. In fact, even in an unconscious way, stress (or the lack of it) influences all our decisions and activities. It is a multi-modal phenomenon: it influences peoples lives in many different dimensions, from the physical to the psychological. Moreover, there is no one-size-fits-all solution: each individual is affected in different ways and each one has a natural level of stress that brings balance to his life. More stress will make the individual take hasty decisions, be weary, strained or less tolerable to stressors. Less stress will make the individual's life dull, will decrease productivity and ultimately lead depressive states.

Stress, or the lack of it, influences virtually al our decisions and actions.

In the particular domain of conflict resolution, stress is also one of the key factors conditioning decisions. Specifically during the negotiation phase, the levels of stress may make people take wrong decisions that they may regret later. Mediators have a major role in avoiding this: they should possess the necessary skills to interpret the state of the parties and intervene when necessary in order to prevent a escalation of the emotions.

DATA	BRIEF DESCRIPTION	SIZE
Acceleration	Data concerning the acceleration felt on the handheld device while playing the game	33366

Movement	A dataset containing information about the amount of movement during the tests	9137
Touches	SThis dataset contains information about the touches	590
Proposals	Data concerning the proposals made by the parties in all rounds	60

Table 3: Dataset generated from the Negotiation game, describing the actions taken during the game and the behavioural features of the participants.

When negotiations take place with all the parties in the physical presence of each other, this is fairly easy for a skilled mediator (Herrman et al., 2001). However, when negotiations take part in virtual environments, in which parties share generally text-based messages only, it becomes very difficult for a mediator to accurately evaluate the state of the parties. The work described in this section aims at tackling this challenge, namely by providing a context layer that allows mediators to take better informed decisions.

7.5.1 *A Multi-modal view on Stress*

One of the first definitions of stress was proposed by (Selye, 1978). According do Selye, stress can be seen as a non-specific response of the body to external demands. These demands (the load or stimulus that triggered a response) are denominated *stressors* while the internal body changes that they produce constitute the actual *stress response*. Selye was also the first to document the chemical and hormonal changes that occur in the body due to stress.

Nevertheless, the definition of stress is still not consensual in the scientific community, remaining as an open topic of discussion. In fact, stress involves a multiplicity of factors, many of them subjective, leading to multiple interpretations that make it difficult to be objectively defined. Thus being, some researchers argue that such a concept is elusive because it is poorly defined (Cox, 1985) while others prefer not to provide an actual definition of the concept until a more accurate and consensual view of the phenomenon is achieved.

In an attempt to address this issue, researchers started dealing with stress from an empirical point of view. In this sense, a strong focus was put on its cognitive and behavioural effects and it started to be viewed as a mind-body, psychosomatic or psycho-physiologic phenomenon. A more up-to-date view of stress thus looks at it as a physico-physiologic arousal response occurring in the body as result of stimuli. It should also be added that these stimuli only become stressors by virtue of the cognitive interpretation of the individual, i.e., the effects of stressors depend on the individual. This is the interpretation of stress considered in this PhD and is the starting point for the definition of a stress model.

Given the multiplicity of factors that influence stress and the different modalities of the behaviour and cognition that are affected, a single-modality approach for measuring the effects of stress would not be suited.

Given the multiplicity of factors that influence stress and the different modalities of the behaviour and cognition that are affected, a single-modality approach for measuring the effects of stress would not be suited, as some experimental results demonstrate (Wenhui Liao A, 2006). In fact, for a sufficiently precise and accurate measurement of stress, a multi-modal approach should be considered. The diagram depicted in Figure 27 represents the multi-modal approach followed to measure stress. This diagram has two main parts. The upper part concerns the predictive aspects of stress, i.e., aspects influencing stress that can be estimated from the background or context of the user. The lower part concerns the diagnostic aspects of stress, i.e., it deals with the effects of stress on several spheres of the individual.

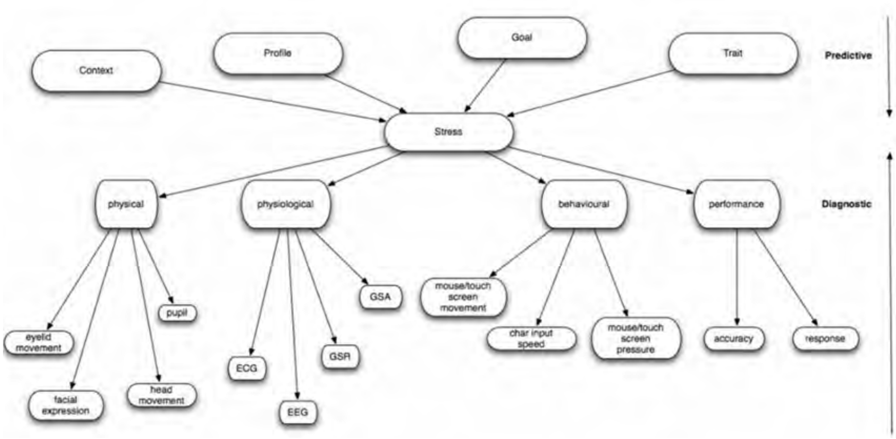

Figure 27: The stress recognition model developed during this PhD work. It includes two main groups of aspects: predictive and diagnostic. Predictive aspects are the ones that can be estimated from the background or context of the individual. Diagnostic aspects are the ones that can be observed and measured and that have a relationship with the level of stress.

- Predictive

    CONTEXT Includes meaningful information to describe the different contexts of the individual, including the historical, economical, social or geographical contexts. Numerous studies exist that map such information to a base level of stress. In (Soylu et al., 2000) the authors evaluate the effect of socioeconomic status on blood pressure of children living in areas with different degrees of development. In (Evans et al., 1998) the context in which people live is studied as a source of stress, namely in what concerns environmental stressors such as noise. The important role of the economical situation of the individual on his level of stress is studied in (Elder and Caspi, 1988). In (Meyer, 2003), the authors evaluate the weight of the social stress, i.e., the stress that society makes each individual feel for their choices, limitations or conditions. As a last example, in (Stockdale et al., 2007) the authors analyse the importance of the social context for stress, namely in what concerns neighbourhood stressors. As examples, commonly known notions include: living in a metropolitan area is more stressful than living in a rural area; a poor bank account or a debt to the bank is more stressful than a rich account or no debts; an active social life helps relief stress.

PROFILE The profile of the individual includes more personal information such as the age, gender, marital status, number of children, type (or lack) of employment, job category, among others. The relationship between these factors and stress has been thoroughly studied by researchers of different fields. In (Fagin, 1985), the authors evaluate the role of unemployment on stress. Parental stress is thoroughly analysed by (Berry and Jones, 1995), with the definition of a stress scale. A good review of empirical research on job stress is also performed on (Beehr and Newman, 1978).

GOAL While the previous aspects were general, this one is more intrinsically related with the scope of this PhD work. The goal of the individual in the conflict resolution process is related to his level of stress: an individual that aims at maximizing his gain at all cost is generally under a higher level of stress than an individual that wants only to reach an agreement, despite of the losses (Tidd and Friedman, 2002).

TRAIT The trait is related with the personality of each individual, i.e., habitual patterns of behaviour, thought or emotion. Some specific traits can be related to stress. As an example, an impulsive individual is generally a stressed one, with stress driving his hasty decisions. Specifically interesting here are the personal conflict handling styles described before, that can be seen as traits of individuals before a conflict and can also be related to the level of stress (Tidd and Friedman, 2002).

- Diagnostic

PHYSICAL The physical aspects include, in a general way, body movements or postures that have some particular meaning that can be related to stress. Specifically interesting are aspects such as the eyelid movement, facial expressions, body movements (e.g. specific gestures, head movements, repetitive movement patterns) or pupil movement and dilatation.

PHYSIOLOGICAL Physiological aspects are the ones that provide an easier diagnose of stress. In fact, several approaches exist nowadays that can evaluate the level of stress of an individual from the Galvanic Skin Response (the electrical conductivity of the skin, influenced by sweating, which is caused by stress), from Electroencephalography (a measure of one's brain electrical activity) or from Electrocardiography (a measure of the electrical activity of the heart) (Picard, 2000). Other physiological indicators include the respiratory rate, heart rate or body temperature. All these aspects are reliable indicators of the level of stress.

BEHAVIOURAL As addressed before in this chapter, the behaviour of an individual is the visible end of his inner self. In that sense, besides other things, behaviours (and especially changes on the behaviours) may also be a good indicator of the level of stress. Given the scope of this PhD work, particularly interesting are the behaviours related to the interaction patterns with technological devices (e.g. how an individual writes, how a individual interacts with a smartphone, how a individual moves when in front of a computer).

PERFORMANCE The performance of an individual is significantly affected by stress. The optimum level of stress will maximize performance. A higher level of stress may increase performance temporarily but will soon tire the individual. A lower level of stress will decrease productivity and lead the individual into increasing lethargy. Thus, tests that evaluate the performance of the individual in given tasks for which standard performance values are known can be a good indicator of the effects of stress on the individual.

From an high level point of view, different types of stress can also be identified, namely acute and chronic stress. Acute stress comes from recent demands and pressures and from anticipated demands in the near future. On the other hand, chronic stress is a long-term one, due to social or health conditions, dysfunctional families, among many other issues. This type of stress will have nefarious effects on the body and mind of the individual, slowly wearing him away day after day. On the other hand, acute stress, because it is short term, won't do the extensive damage associated with chronic stress. Nevertheless, it will instantaneously influence the performance of the actions being performed. The interest here is clearly in the analysis of acute stress and its effects, given that they may be more determinant for the outcome of a conflict resolution process.

7.5.2 *Methodology and Data Collection*

The main objective of this research line is to identify which factors of the human physiology influenced by stress can be pointed out using standard devices available on the market.

The following research questions and hypothesis are analysed: (a) Does stress actually influences, in a significant manner, our interaction patterns while using common technological devices? (b) Is it possible to accurately measure this influence of stress in a non-invasive and non-intrusive way?

With regard to the first question, it is hypothesized that stress does influence interaction patterns in a significant manner. The work of other researchers, who proved that stress does influence people's behaviours and physical responses in a significant way, serves both as a motivation and as a pre-validation of the approach (some examples can be found in (Healey and Picard, 2005), (Vizer et al., 2009) and (Rehm et al., 2008)).

Concerning the second question, it is hypothesized that it is possible to measure, using non-invasive and non-intrusive methods, the level of stress of the users by analysing key features in their interaction with technological devices.

There are also some sub-questions that, although not compulsory, will contribute to the achievement of better results. Namely: (1) Determine how the interaction patterns of each user are affected by stress and develop personalized stress models that will maximize the accuracy of the output. (2) Find groups of people that are affected in similar ways (e.g. same parameters show similar tendencies when subject to stress) in order to develop more accurate generic stress models, to be applied when a personalized model is not available.

In order to test these hypothesis, an experiment was set up in the environment described in Section 7.3 aimed at collecting data about how people interact with technological devices, when calm and when under stress. The collection of the data was organized into two phases. In a first phase, individuals were required to perform these tasks in a stress-free environment. In a second phase, the the same individuals performed the same tasks subject to stressors such as the vibration of the devices, loud and annoying sounds, unexpected behaviours of the devices, time constraints, among others.

The empirical data gathered in both phases about the user interaction patterns and physical response is described in Table 4. This data was synchronized and transformed/normalized to allow its joint analysis. The participants of the proposed experiment were volunteer students and professors from the University of Minho. 19 male and female individuals participated in the experiment aged between 20 and 57. All these individuals are familiar with the technological devices used and the interaction with them was not an obstacle.

The data gathered was analysed in order to determine statistically significant differences between phase 1 and phase 2 of the data collection. Measures of central tendency and variability were calculated for all variables of interest. The Mann-Whitney-Wilcoxon Statistical test was used to test whether there are actual differences in the distributions of the data. A 0.05 level of significance was considered. The data analysis was performed using Wolfram Mathematica, Version 8.0.

This analysis of the data allowed to determine which parameters, for each individual, were effectively affected by stress and to which extent. This information allows to develop personalized models for stress estimation in real time. Moreover, a more generic model can also be developed taking into consideration the data of several or all the users. This generic model can be applied in the cases in which a personalized one is not available. The models developed are used to develop a real-time stress estimation software layer to be used in a VE or in other domains.

In order to identify which factors vary with stress and in what magnitude, a game with the following features was developed: a mentally challenging objective and stressors. The main objective of the game is that the user performs mental calculations using the four basic arithmetic operations and a group of numbers given randomly in order to get as close as possible to a target number. The score is given in function of the distance of the result to the target number: the closer to a given target number, the higher the score, up to a maximum of 100 points when the result is equal to the target number. There are however some rules. In each round, it is only possible to use each number and each operator once. As an example, let us say that the target number is 198 and you have four random numbers such as 50, 45, 3 and 8. A fairly good solution would be 3*50+45.

In each consecutive round of the game the user has to perform one of such calculations. There are however stressors that make the effects of stress visible, including a time limit, vibrations and sounds on the handheld devices. The time limit decreases as the number of rounds increases. Moreover, the longer the user takes to answer in a round, the smaller the time limit in the following round. Thus, there is a pressure on the user to answer quickly. At the same time, there is also the pressure to make the best possible score. Vibration and sounds are used to increase the level of stress when the time is almost over at each round. When only a few seconds remain, the handheld device starts vibrating and making a disturbing beeping sound whose frequencies are higher when the time limit is smaller. All this increases the physiological effects of stress on the users, as the results show. Two screenshots of the game are shown in Figure 28.

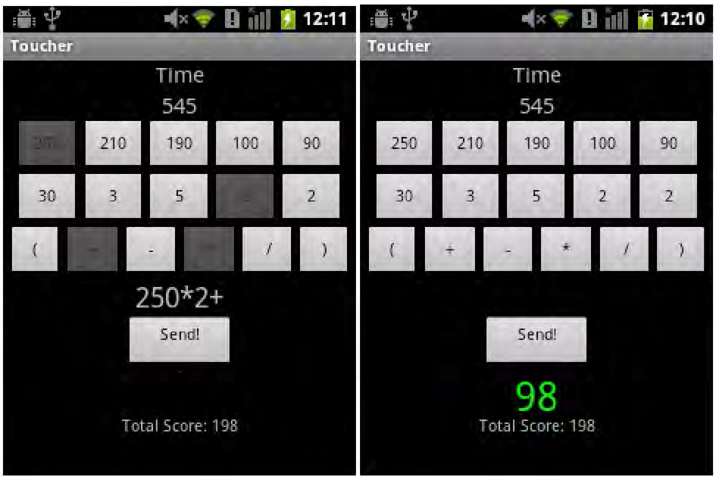

Figure 28: Screenshots of the game's interface in which the user must perform mental calculations while under stress.

As stated, the game is developed to induce stress in the individual and to determine in which way each one is affected. To study this, the 8 parameters mentioned in Section 7.3 are analysed for each user.

7.5.3 *Statistical Data Analysis*

Each individual was requested to play the game for some rounds without any source of stress. In this version, the game has no time limit and no vibration or annoying beeping sounds. In that sense, the user calmly plays the game, with enough time to think on the different possibilities. This phase allows collecting some data about how the individual normally behaves when he/she is not under stress (e.g. how he touches the screen, how he holds the smartphone, how well he fares in the game). This enables establishing a baseline for comparison. Afterwards, the same data is collected when the individual is playing the game with the stressors. The data gathered without stress is from now on designated as *baseline data*, whereas the data gathered with the influence of the stressors is from now on designated as *stressed data*.

Given that the data about the state of the user comes from different sources (e.g. handheld device, video camera) and is synchronized by a timestamp, the clocks of the devices are previously synchronized. The collected data is organized into five groups. The first one contains the baseline data. The second one contains the stressed data. The other three are subgroups of the stressed data created according to the time that was left for the end of the round at the instant of the event. Three groups were created: one for events that took place when there were less than 8 seconds left for the end of the round, another for events that took place when there were less than 5 seconds and the last one for events that took place when there were less than 3

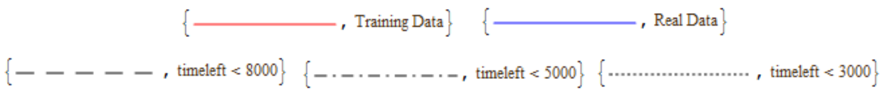

Figure 29: Different styles of the lines used to depict the different groups of the data collected. These styles will be used in the remaining of this document.

seconds remaining. The notation depicted in Figure 29 is used for the graphical representation of the data from this point on in the thesis, in order to distinguish between the different groups of data.

To determine to which extent each feature considered is or is not influenced by stress, the baseline data is compared with each of the remaining four groups. Provided that most of the distributions are not normal, the Mann-Whitney test is used to perform the analysis. This test is a non-parametric statistical hypothesis test for assessing whether one of two samples of independent observations tends to have larger values than the other. The null hypothesis is thus: H_0 = *The medians of the two distributions are equal*. For each two distributions compared, the test returns a p-value, with a small p-value suggesting that it is unlikely that H_0 is true. For each parameter, the training data is compared with the remaining four groups. In all the tests, a value of $\alpha = 0.05$ is used. Thus, for every Mann-Whitney test whose $p - value < \alpha$, the difference is considered to be statistically significant, i.e., H_0 is rejected.

A significant difference between the baseline data and the stressed data means that the parameter is effectively influenced by stress for this specific user. This is the most desirable result as it indicates a higher level of confidence. This type of analysis is from now on designated as *first order*. If this is not the case, the baseline data is compared with each of the three subgroups of the stressed data (subgroups describing increased levels of stress), in search for significant differences. If statistically significant differences between one of these groups and the baseline data are found, it may still be concluded that the individual is affected by stress in this parameter, although not in such an explicit manner. From now on, this kind of analysis is designated as *second order*.

Table 4 depicts the sizes and a brief description of each dataset generated for each feature studied. Each of these datasets consists of a list of values describing the respective feature as well as a time stamp and a user name. A detailed analysis of the results achieved is provided in Chapter 8.

DATA	BRIEF DESCRIPTION	SIZE

Acceleration	Data concerning the acceleration felt on the handheld device while playing the game	27291
Maximum intensity of touch	Data about the maximum intensity of each touch in a touchscreen	1825
Mean intensity of touch	This dataset contains data about the mean intensity of each touch event in a touchscreen	1825
Amount of movement	A dataset containing information about the amount of movement during tests	25416
Touches on target	This dataset contains information about the accuracy of the touches	1825
Stressed touches	A dataset containing information that allows to classify each touch as stressed or not stressed	1825
Score	A dataset describing the performance of the user playing the game, during the tests	321
Touch duration	A dataset containing the duration of each touch event	1825

Table 4: Dataset generated from the experiment, describing the behavioural features of the participants.

7.5.4 *Real-time Assessment of Stress*

Until now, it has been seen how a concrete model was defined for shaping the effects of stress on individuals, after which the process of collecting the necessary data was described. Essentially, in this work, data describing how each user behaves in a calm state and in a stressed state was collected and analysed. The final step is now to develop a solution that allows to assess the level of stress of each individual, in real-time, so that the mediator can take real-time decisions as well. Two approaches are described in this section. The first one is based on the analysis of the data for each feature individually. The second one looks at the touch patterns of the users, i.e., it looks at the shape of the curve of intensity of touches.

Only a real-time analysis of stress allows a mediator to take insightful decisions.

From a high level point of view, the first approach consists in collecting data in real-time about the features studied, and compare it to the data that was already collected for the individual. In a few words, if the data being collected is very similar to the data collected when the user was calm, it is concluded that user is in a calm state. Alternatively, if the data is similar to the data collected when the user was stressed, it is concluded that the user is stressed. The main problem here is thus the one of finding the similarity between distributions of data. Moreover, if it is possible to evaluate the degree of similarity, it will be possible to provide a quantitative value of *how stressed* or *how calm* an individual is, by determining *how similar* the data being collected is to each of the known distributions.

Two approaches have been implemented to address this problem and are described in this subsection. The results of their use are described in Chapter 8. The first one is based on the notion of Confidence Intervals while the second is based on a statistical classifier.

The Confidence Interval is a term used in statistics that measures the probability that a population parameter will fall between two set values. As an example: given a 95% confidence interval, an individual's maximum touch intensity falls between 0.8 and 0.14 when the individual is stressed. What this means is that we are 95% confident that when the individual is stressed, the values for the maximum intensity of his touches on a touch screen will fall between 0.8 and 0.14.

The first approach thus compares two distributions of data by comparing their confidence intervals. Figure 30 depicts how it is done. The lower part depicts a distribution of data, with the 90% confidence interval highlighted by a blue rectangle. This represents a known distribution. The upper part of the figure (the orange rectangles) depicts several confidence intervals of another distribution (the one being generated in real-time), under different scenarios. Specifically, four examples are provided. In scenario a), the new distribution is not similar to the new one as their confidence intervals do not overlap at all. Under scenario b) there is some significant degree of similarity as their confidence intervals overlap considerably. Scenario c) presents a similar situation. Finally, scenario d) shows some similarity as well, although smaller, given that the overlap area is smaller.

Hence, under this approach, the similarity between two distributions is given in terms of the percentage of overlapping area of the confidence intervals of the two distributions. Distributions are constructed in real-time for each feature and compared with the known stressed and calm distributions, at regular time intervals. Thus, in each time interval, the system outputs a measure of how similar the behaviour of an individual is to the recorded stressed and calm behaviours, for each feature studied.

Under the second approach, standard and well known pattern recognition tool is used: the *k*-nearest neighbour algorithm, specifically, the `weka.classifiers.lazy.IBk` (Aha et al., 1991) implementation for `java`, provided by the Weka workbench (version 3.6.3) (Holmes et al., 1994). It is a method for classifying objects based on closest training examples in the feature space: an object is classified

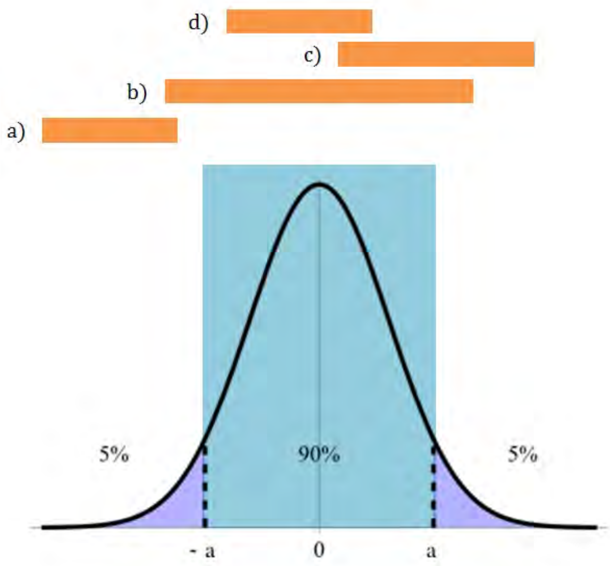

Figure 30: Quantifying the similarity between two distributions of data by measuring the overlap of their confidence intervals.

by a majority vote of its neighbours, with the object being assigned to the class most common amongst its k nearest neighbours (k is a positive integer, typically small). If k = 1, then the object is simply assigned to the class of its nearest neighbour.

Figure 31 depicts the working of this algorithm with a real example. The data detailed describes several instances of the maximum intensity of touch: the red dots correspond to the data collected under stress while the blue dot to the data collected without stress. This data shows a tendency that will be analysed in detail in Chapter 8: under stress the intensity of touch is higher. Given the new instance to be classified represented by the green dot (a new touch), and k = 4, the algorithm would classify it as stressed given that there are more neighbours from the class "stressed" (3) than from the class "not stressed" (1). The working of the algorithm is the same for each of the other features and for each new instance that must be classified. A classifier was trained for each feature and for each individual that participated in the experiment, using the data collected.

The analysis of the classifiers' performances for each feature is performed in detail in Chapter 8. This analysis is validated by the Cohen's kappa coefficient, which is a statistical measure of inter-rater agreement that tells how much of the accuracy is due to chance. Afterwards, for each feature the average performance of all the classifiers in terms of their average sample accuracy is also considered. Finally, box-and-whisker diagrams are plotted with the values of the sample accuracies of all the classifiers for each feature and look at the median and upper and lower quantiles. Finally, the average values of

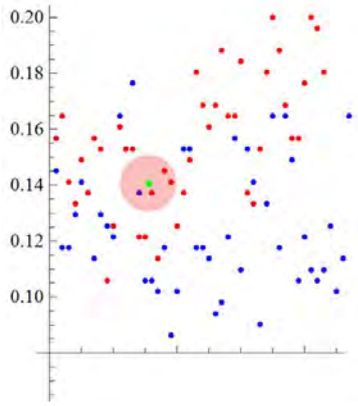

Figure 31: Example of the working of the nearest neighbour algorithm by a majority vote: in this case the green dot (new instance) would be classified as "stressed" (the class of the red dots).

the kappa coefficient of all the classifiers trained for each feature are also considered.

Independently of the approach used, the system will compute a value of stress for each feature, in each time instant. These values are normalized and the final value of stress is a weighted sum of the level of stress corresponding to each feature. The weight is given in terms of the results of the significance tests (e.g. the touch intensity has a bigger weight than the acceleration if the individual is more significantly affected by stress on the touch intensity than on the acceleration). More details on this process are given on the results' analysis (Chapter 8).

The shape of the intensity curve of the touch can also provide information about the user's level of stres.

Concerning the second approach mentioned at the beginning of this subsection, it looks at the shape of the touch patterns of the users in order to provide an additional feature for assessing stress. The hypothesis being assessed is that stressed users will have different touch patterns (e.g. longer, shorter, steeper). This approach is clearly targeted at data collected from devices with touch screens, particularly the Android smartphones used to build the environment Section 7.3.

Thus, to assess the level of stress of each touch, this approach relies on the event listeners provided by the Android framework. An event listener is an interface in the `View` class that contains a single callback method that will be called by the Android framework when the View to which the listener has been registered is triggered by user interaction with the item in the UI. For this purpose, the `onTouch()` callback method is used, which is called when the user performs an action qualified as a touch event, including a press, a release, or any movement gesture on the screen (within the bounds of the item). Thus, in each touch of the user on an item of the UI several touch events are fired: one when the finger of the user first touches the screen (identified by the action event `ACTION_DOWN`), several while the user is touching (depending on the duration of the touch) and one when the finger releases the screen (identified by the action event `ACTION_UP`).

Each of these events has information about the intensity of the touch (via the `getPressure()` method) and about the position of the event. Moreover,, when each event is fired, it is registered together with a time stamp. This allows to visualize the evolution of a touch in terms of its intensity over time. From this information it is also possible to extract the duration and intensity features.

DURATION FEATURE is defined as the difference between the timestamps of the action events `ACTION_UP` and `ACTION_DOWN`. One of the hypothesis being tested is that the stress of a user will have an influence on the duration of the touch, hence our interest. The duration of the touch can however be influenced by factors other than the stress. Namely, the type of item of the UI being touched. In that sense, events fired by items such as sliders or by scrolling pages are not considered. For the purpose of the experiment implemented, the interest lies in the standard touches used to interact with buttons, inputting text and similar actions.

INTENSITY FEATURE The intensity of a touch event depicts the force exerted by the finger of the user while touching the device. Given that each touch event includes a pressure and that each touch fires several touch events (as described above), it is possible to analyse the variation of the intensity throughout all the touch, from the moment the finger touches the screen to the moment it releases it.

The main goal of this approach is thus to investigate if it is possible to build a classifier for touch patterns that can be used in real-life applications to provide some information about the level of stress of the user. Two standard and well known machine learning tools are used for this purpose: a decision tree constructor and a support vector machine. As the decision tree constructor the J48 algorithm is used - the `java` implementation of the C4.5 (Quinlan, 1993). As support vector machine, the SMO function was selected, which implements John Platt's sequential minimal optimization algorithm for training a support vector classifier (Platt, 1999). These experiments were also performed using the Weka workbench (version 3.6.3).

The results of the two classifiers will be compared in Chapter 8 by looking at some performance measures such as the percentage of correctly classified instances, the Kappa statistic (which is a chance-corrected measure of agreement between the classifications and the true classes) and the ROC area.

As previously described, each touch in the screen results in several touch events that are fired during the time of the touch. This number varies according to the duration of the touch. In that sense, the data for each touch, as it is, cannot be used to build a classifier (each touch would have a potentially different list of values of intensity, one for each touch event). Figure 32 (a) highlights this by depicting different types of touches.

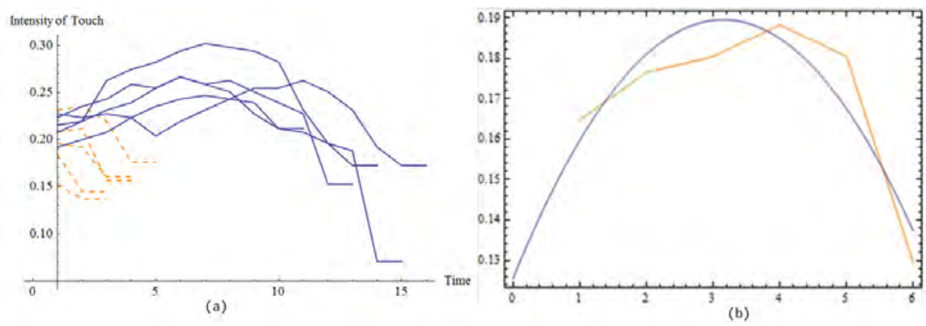

Figure 32: (a) 10 different touch patterns from users: touches can be composed of a different number of touch events. The orange lines depict touches classified as "calm" whereas the blue lines belong to touches classified as "stressed". (b) Fitting a polynomial curve (blue curve) to a given touch (orange line).

To tackle this problem it was observed that the intensity of touches follows a similar shape: a convex curve that grows to a maximum point and then decreases. Thus, the approach was to fit a second polynomial degree curve to each touch pattern. To perform this fitting in real-time the proposed approach uses `J/Link`, the Mathematica's Java interface that allows for controlling Mathematica Kernels from Java programs. Specifically, the `Fit[`*`data, funs, vars`*`]` function is used which finds a least-squares fit to a list of data as a linear combination of the functions *funs* of variables *vars*. To implement this, Mathematica v8.0 was used. An example of this approach is depicted in Figure 32 (b). Given that the second degree polynomial curves are of the type $y = ax^2 + bx + c$ it is possible to compare the parameters of the curve of each touch pattern: similar values of `a,b` and `c` indicate similar curves, thus similar touch patterns. Hence, the input for the classifier are three numeric attributes `a,b` and `c` (the independent variables) and a nominal attribute that describes the state of the user at the time of the touch as "stressed" or "not stressed" (the dependent variable).

In a few words, this last approach characterizes the behaviour of an individual based on the way he touches the screen of the hand held device. Each touch is classified as "stressed" or "not stressed", allowing a temporal analysis of the type "number of touches classified as stressed per unit of time".

These two approaches (analysis of each feature individually and analysis of touch patterns) allow the development of a real-time solution for estimating the level of stress of an individual in a non-invasive way. All the results concerning this subject are provided in Chapter 8.

7.6 SUMMARY

In this section, the focus has been placed on the process of devising and implementing an approach for acquiring context information from a physical environment. Specifically, two different types of context information were studied in detail, given their importance in the conflict resolution process: the personal conflict handling style and the level of stress. While the first defines the way that each individual reacts before a conflict, the last defines how much external pressure each one is able to successfully handle. A particular emphasis has been placed on acquiring this information in a non-invasive, automated and transparent way, clearly breaking with the traditional self-reporting approaches in use nowadays, inherited from the Social Science.

The approach implemented for the estimation of the personal conflict handling style was inspired by an economical view on the problem and is fairly straightforward. The estimation of the level of stress however, by being such a complex and subjective issue, was more laboured. Several challenges can be put forward in the time of developing a practical human stress monitoring system, namely:

1. the expression and the measurements of human stress are very much person-dependent and even time or context dependent for the same person;

2. the sensory observations are often ambiguous, uncertain, and incomplete;

3. the user stress is dynamic and evolves over time;

4. the lack of a clear criterion for feasible stress states greatly increases the difficulty of validating stress recognition systems.

In order to tackle these challenges, the approach developed is multimodal and personalized. It's multi-modal nature allows it to consider inputs from the different spheres of an individual that are affected by stress, making it exhaustive. Its personalized nature stems from the models developed, in which data from each user is analysed individually. Moreover, a systematic approach has been implemented, inspired by the applied behaviour analysis methodologies of the Social Science.

Thus, each of the seven points enumerated in Section 7.2 were addressed in the following way:

1. At least one participant is needed to study a given behaviour. The experiments performed in the scope of this work involved 19 participants from the faculty;

2. There must be at least one behaviour being studied. In this study, eight features were considered that depict how individuals behave using technological devices in face of stress;

3. An environment was developed to implement the experiment. This environment allowed to control stressors that induced stress on the participants and to collect information about behavioural features being studied;

4. A system for measuring the behaviour of the participants was developed, based on the statistical analysis of the data gathered. Specifically, a measure of the similarity between distributions of data was developed. This included real-time and visual analysis of data;

5. A small difference from the classical approaches exists at the level of the intervention condition. It does not exist during the experiment but only posteriorly, during the actual conflict resolution process, i.e., the intervention takes place by the mediator or the conflict resolution platform, when the level of stress escalate, and not during the collection of the data;

6. Stressors (the independent variable) may be manipulated in order to study their effect on the behavioural features (the dependent variables). This is done by changing their intensity (e.g. changing the time limit, increasing the volume/frequency of the sounds, increasing the rate of vibrations on the handheld device);

7. The participant will benefit from the experiment in the sense that the mediator or the conflict resolution platform will be able to take better-framed decisions, ultimately leading to better and more mutually satisfactory outcomes;

The final goal of this approach inspired on applied behavioural analysis research is the development of conflict resolution environments Figure 33. Under this new view, each individual participates in the conflict resolution from his/her own *real* environment, which is equipped with sensors and devices that acquire different kinds of information in a non-intrusive way. While the individual conscientiously interacts with the system and takes his/her decisions and actions, a parallel and transparent process takes place in which this information is sent in a synchronized way to the conflict resolution platform. The platform, upon converting the sensory information into useful data, allows for a contextualized analysis of the operational

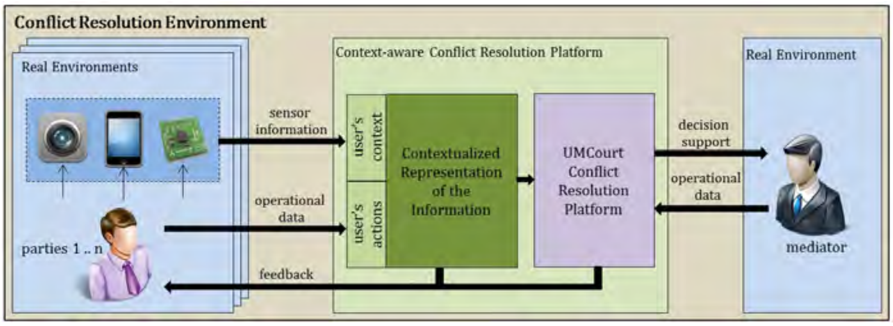

Figure 33: The proposed concept of Conflict Resolution Environment is central in this PhD Work: the user's context is acquired from sensors placed in the environment of the user. This allows to build a contextualized representation of his actions, supporting the conflict resolution and the mediator in taking better actions.

data of the individuals. This contextualized analysis is performed by the platform itself (e.g. for performing decision-support related tasks) and by the mediator. Then, the parties receive feedback from the platform (e.g. a new proposal, information updates, notifications), which may also include some kind of feedback from the state of the parties (e.g. an avatar showing the state of the parties).

Part III

RESULTS, CONCLUSIONS AND FUTURE LINES

This is the last of the three parts of this thesis, in which the results and conclusions are presented. It's first chapter details the practical results achieved with the implementation of prototypes or experiments. They will allow to understand the extent of the findings and their importance. The second chapter provides some quantitative results, namely concerning the publications and the prototypes and case-studies implemented. Finally, the third chapter concludes this part by making a critical analysis of the accomplishments, detailing the main conclusions achieved and providing some insights into the future work.

8 ANALYSIS OF PRACTICAL RESULTS

You cannot teach a man anything;
you can only help him discover it in himself.

— Galileo Galilei

This last part of the thesis starts with a description of the main results achieved during the implementation of the work plan. It details the performance of some of the methods developed and fully describes the results of the practical experiments performed, namely in what concerns stress and conflict handling styles. It is still not a critical analysis of the results, which is performed in Chapter 10, but rather an objective and statistical view on the data collected.

8.1 CONCERNING INFORMATION RETRIEVAL

In this section the performance of the information retrieval methods described is analysed. To do it, a test case was implemented in the field of Portuguese Labour Law. However, given the complexity of the Portuguese Labour law, a restrict set of norms was used. In that sense, a group of 36 norms of the Decree of Law 7/2009, from February 12th, 2009 were considered. These norms were selected because they are generally present in most of the disputes in Labour Law 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 Figure 34. The representation of these cases in vectors of binary entries was generated previously and was made available, in the form of a file.

8.1.1 *Efficiency*

During this performance assessment process, the focused has been 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, and their times of execution measured. This is useful to compare them in terms of their execution times. These tests thus evaluate the efficiency of each algorithm presented. One of the

```
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
```

Figure 34: A description of the database detailing which aspects of the Portuguese labour 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.

factors that influences the performance of the algorithms in a more significant way is the size of the database. The more cases there are the higher the efficacy of the algorithms. However, the efficiency will necessarily decrease as more cases need to be analysed and compared. The database used contains 127 cases. Given that the algorithms presented here have linear complexity, the efficiency of these algorithms will decrease proportionally to the size of the database.

Pre-Select (Classification)

The data depicted in Figure 35 describes 100 iterations of the preselection 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 analyse the case, determine which rules were true 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. Analysing Figure 35, it is possible to conclude that this is a considerably efficient 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 is 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.

Pre-select (Template Retrieval)

To test this algorithm, several pre-selection tasks were requested, using a random case. To perform the pre-selection, the algorithm had to analyse the case and then interact with the database, requesting all the cases that matched a given criteria. This algorithm was tested

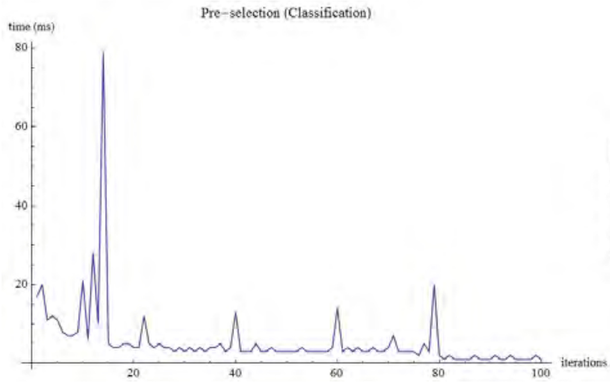

Pre-Selection (Classification)	
Average	5.53
Standard Deviation	8.78757
Maximum	79.
Minimum	1.
Size	100.

Figure 35: Summary of the performance of the pre-selection task using association rules for 100 requests generated randomly.

using a local instance of a database and a remote one. Looking at Figure 36, 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, the system depends on external factors like the speed of the internet connection and (3) in order to satisfy the pre-selection rules, the algorithm may need to adapt strategies and make several iterations (requests to the database). This can be improved by choosing the best pre-selection rules, minimizing the number of times that the algorithm must adapt the search queries in search for more/less/better cases. When compared with the previous algorithm, this shows significantly higher times of execution, mainly due to the interaction with a remote database but also to the potential need for multiple iterations.

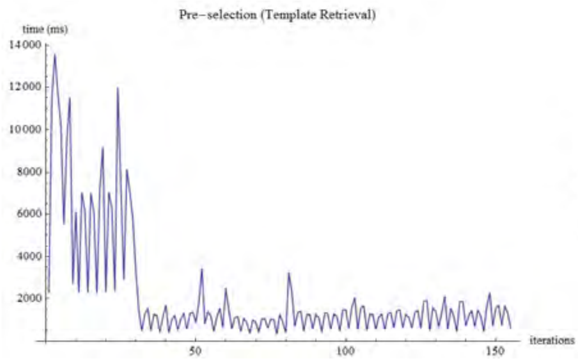

Pre-Selection (Template Retrieval)	
Average	2203.85
Standard Deviation	2674.8
Maximum	13607.
Minimum	313.
Size	155.

Figure 36: Summary of the performance of the pre-selection task using the Template Retrieval technique for 155 requests generated randomly.

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 with an isolated case and a list of cases, with the objective of determining the similarity between the isolated case and each one of the cases in the list. The results depicted in Figure 37 show 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.

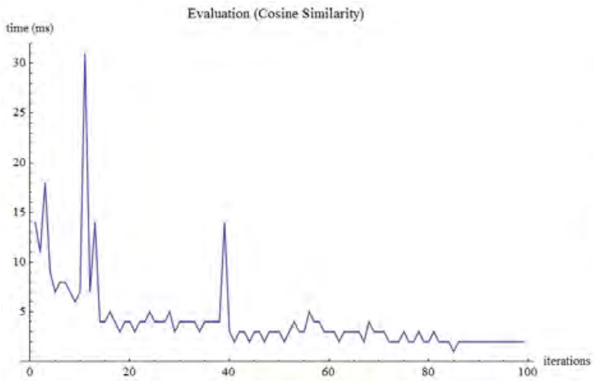

Evaluation (Cosine Similarity)	
Average	4.18182
Standard Deviation	3.94959
Maximum	31.
Minimum	1.
Size	99.

Figure 37: Summary of the performance of the evaluation task using the Cosine Similarity technique for 99 requests generated randomly.

Evaluation (Nearest Neighbour)

To test this algorithm, an approach similar to the previous one was followed. When analysing the collected data depicted in Figure 38, it is possible to conclude that the Nearest Neighbour 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 that are more complex to handle (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.

Get Complete Info

In this test, the objective was to determine the efficiency of a request of all the information regarding a case, i.e., given 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.

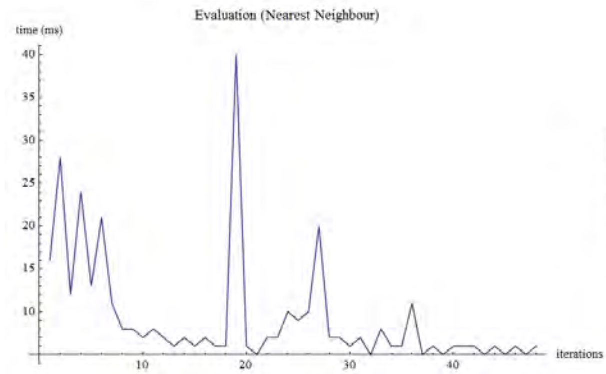

Evaluation (Nearest Neighbour)	
Average	9.27083
Standard Deviation	6.80969
Maximum	40.
Minimum	5.
Size	48.

Figure 38: Summary of the performance of the evaluation task using the Nearest Neighbour algorithm for 48 requests generated randomly.

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, similarly to the previously presented results. Thus, the execution time of a complete info request depends mostly on which algorithms are selected. The resemblance of the graph depicted in Figure 39 with the one depicted in Figure 36 is also not a coincidence as, given the potentially high values of its execution times, it has a considerable influence on the overall performance.

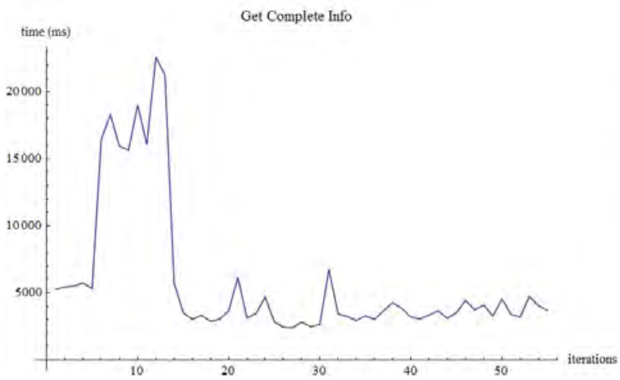

Get Complete Info	
Average	5891.65
Standard Deviation	5284.57
Maximum	22580.
Minimum	2394.
Size	55.

Figure 39: Summary of the performance of the "Get Complete Info" task for 55 requests generated randomly.

8.1.2 *Efficacy*

More than the efficiency of the algorithm, its efficacy must also be considered, i.e., it is not enough to perform a given task quickly, it must also be performed as correctly as possible. In this sense, the platform keeps record of some key actions as well 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. The objective is that, in the long term, the platform is able to apply these recommendations autonomously. However, this still requires further validation as it is a very sensitive topic.

A typical output of a self-assessment request is shown in Figure 40, in which the platform is assessing the performance of the preselect 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. Succeeding means that the pre-selection respects all the rules. In the example, this action is failing in 66% of the times, 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.

```
*****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%)
```

Figure 40: The output of a self-assessment request for the pre-select action.

8.2 CONCERNING CONFLICT HANDLING STYLES

In this section an analysis of the effect of stressors on the conflict handling style of the parties is performed. In order to do it the data collected in the two phases of the game (without and with stressors) was compared for the same pairs of players in search for statistically significant differences due to the action of the stressors. The main aim of the study is to assess the influence of stress on the behaviour of the parties and on the outcome of the negotiated process. This experiment involved 14 users playing the game, in a total of 60 negotiation rounds. The data gathered included behavioural features provided by the environment, which were used to estimate the level of stress of each user using the personalized stress models developed, described in the next section. These models were used to determine how the level of stress relates with the behaviour of the participants in a negotiation.

In order to statistically deal with this data, a numeric scale was used to describe the conflict handling styles. Table 5 depicts the conflict handling styles considered, the number of times that each style was evidenced by each participant and the ordinal rank attributed in order to be used by data-mining algorithms. The exact numeric quantity of a particular value has no purpose beyond its ability to establish a ranking over a set of data points. Therefore, rank-ordering was used which describes an order but does not establish relative size or degree of difference between the items measured. This was a mandatory step to make the data suitable for statistical and machine-learning techniques.

STYLE	TIMES USED	ORDINAL RANK
Competing	30	1
Collaborating	12	2
Compromising	8	3
Accommodating	10	4
Avoiding	0	5

Table 5: Summary of the conflict handling styles, the number of times that each style was evidenced and the rank-ordering of each style.

One of the first conclusions achieved when analysing the data is that parties show a competitive style of negotiation most of the times, both in stressed and calm settings. However, when calm, the use of more cooperative style is slightly larger. The histograms depicted in figure Figure 41 depict this: the green curve refers to the distribution of the conflict handling styles in the calm phase while the red line refers to the distribution in the stressed phase. However, the differences are not statistically significant (MannWhitneyTest = 0.33). On average, the style of a party when stressed is 1.525 (closer to competing) and when calm is 1.74 (closer to collaborating).

The evaluation of the progress of the conflict styles during the negotiation process was centred on the average slope of its numeric values. In other words, the object of study was the variation of the conflict

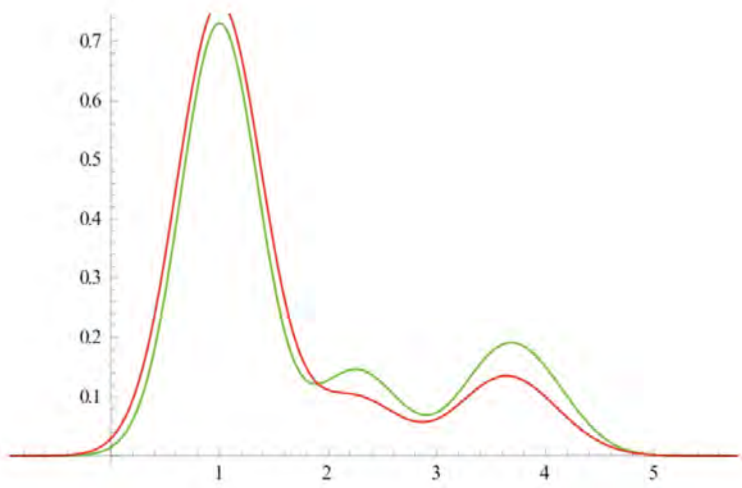

Figure 41: The distribution of the styles used by the parties: the red line represents data from the stressed phase while the green line represents data from the calm phase. The X axis represents the ordinal rank of the conflict handling styles as defined in Table 5. In a calm state the users evidence more cooperative styles.

styles used by each party within the negotiation game. It was concluded that in a stressful state the parties tend to vary their conflict handling style more (on average 0.71 points between the beginning and the end of the game) than when they are calm (on average 0.61 points). This is in line with other results that point out to more sudden and less weighted decisions under stress. Besides that, it was also concluded that the 'manufacturer' role presents a higher average slope (faster change in the conflict style), being on average 0.83 than the 'retailer' (on average 0.51 points).

Concerning the values of the proposals exchanged by the parties during the negotiation, it can be concluded that both parties change more the values of the proposals (on average 0.19) when under stress than during the calm phase (on average 0.14). Moreover, the 'manufacturers' present a more dynamic proposal evolution (changing 0.19 in average) than the 'retailers' (average slope is 0.11).

In a stressful situation it is more likely that the parties propose more uncooperative solutions

The euclidean distance to the optimum value was also analysed, i.e., it was studied the deviation given the most desirable negotiation outcome (the value that prevented both parties from entering into bankruptcy). Under a stressful situation both parties were at a distance of, in average, 0.154 euros from the optimum value while in stress-free situation the distance decreases to 0.071 euros, in average. Therefore, it can be concluded that in a stressful situation it is more likely that the parties propose more uncooperative values. This can be explained as a consequence of acting too quickly or relying too much on coercion. When parties are under pressure they can commit strategic mistakes or give in unwanted concessions. It may also lead to bad agreements.

Acting too quickly is also a known response to external and internal stressors. Indeed, considering the duration of the rounds, one can state that 90% of the negotiation rounds had a shorter duration under a stressful environment than under a stress-free one. However, only in 30% of these cases was the different statistically significant (at a level of 0.05).

Concerning the evolution of the conflict handling style in each game played, it is possible to conclude that 80% of the participants used a competitive conflict style, which is assertive and uncooperative, in the early rounds. During the game 55% of the players improve their styles (shifting towards more cooperative solutions), 35% remain on the same style and 10% become more competitive. It is known that 'competitors' often use power as the primary tool for handling conflict, and work to prove the importance of one side of the argument in order to win. This can be one explanation. Moreover, they are usually more concerned with winning the game than finding the best solution. Taking into consideration the pre-conditions of the game, the second hypothesis is more plausible.

In order to provide a more specific view of the results, one actual case is highlighted. It shows the evolution of the values proposed during the negotiation with stress (Figure 42 (a)) and without stress (Figure 42 (b)). The blue line represents the values proposed by the 'retailer' and the red one by the 'manufacturer'. It is possible to conclude that the 'manufacturer' is more flexible (changes more often). This is a recurrent behaviour and can be explained by the fact that the seller (in this case, the 'manufacturer' role), in a buyer's market, needs to be more flexible and expect more negotiation about contingencies. Comparing the lines, in a calm state the 'manufacturer's' average slope is 0.31 and the 'retailer's' is 0.056. When under stress, the values rise to 0.5 and 0.1, respectively. Similar results are also observed in other pairs of players and are in line with the previously described conclusions: stressed participants take hastier and less weighted decisions.

8.3 CONCERNING STRESS

In this section, the results that accrued from the stress experiments are detailed. Particularly, a preliminary analysis of the data is performed in which the datasets from stressed and calm states for each feature studied are analysed in search of measurable significant differences. This is done using the Mann-Whitney test. Afterwards, the performance of the classifiers developed with such datasets is assessed.

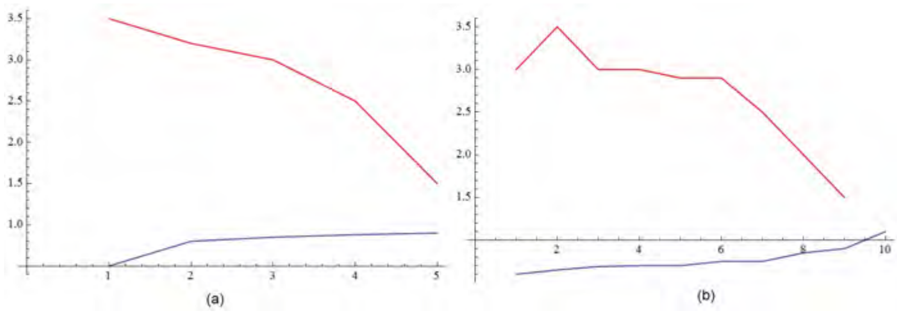

Figure 42: Evolution of the values proposed during the negotiation when under stress (a) and without stress (b). The red line depicts the values proposed by the manufacturer while the blue line depicts the values proposed by the retailer. It is possible the see that under stress the values proposed vary faster.

8.3.1 *Preliminary data analysis*

Acceleration

The information about the acceleration is provided by the handheld device. The interest in this parameter lies in a potential relation between Human movements and the level of stress, specifically in what concerns the way that the user holds and interacts with the handheld device. In fact, a stressed user generally exhibits sudden hand gestures and movements and may also touch the device in a more brusque way.

In fact, considering the first order analysis (as defined in Section 7.5.3), 80% of the users show a significant difference between the baseline and stressed data. Moreover, when comparing the baseline data with the data from higher levels of stress (second order analysis), the results are even more expressive: 93.3(3)% of users show a statistically significant difference between the histograms. In that sense, the acceleration felt in the handheld devices is effectively different between calm and stressed users.

Moreover, it should also be concluded that, for most users, the amount of acceleration measured tends to increase. Figure 43 shows two examples in which the baseline values of the acceleration are more centred in a given value (less deviation) while the stressed data is slightly shifted to the right. The three subgroups of the real data are even more shifted, although the standard deviation also increases. For these two examples, the p-value returned by the Mann-Whitney test for the first order analysis is $5.9975 * 10^{-7}$ (for the data in Figure 43(a)) and $2.75591 * 10^{-14}$ (for the data in Figure 43(b)).

Maximum intensity

From the touchscreen of the handheld device it is possible to acquire data about how the user touches it. In this case, the maximum inten-

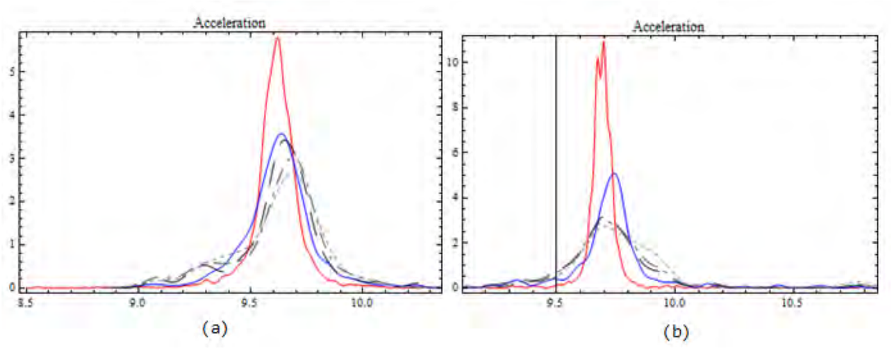

Figure 43: Histograms of data from the module of the acceleration concerning two different users. The difference between the baseline data and the stressed data (and its subgroups) is clearly visible: the data from stressed users has more variability, i.e., stressed users move their hands more or in more sudden ways.

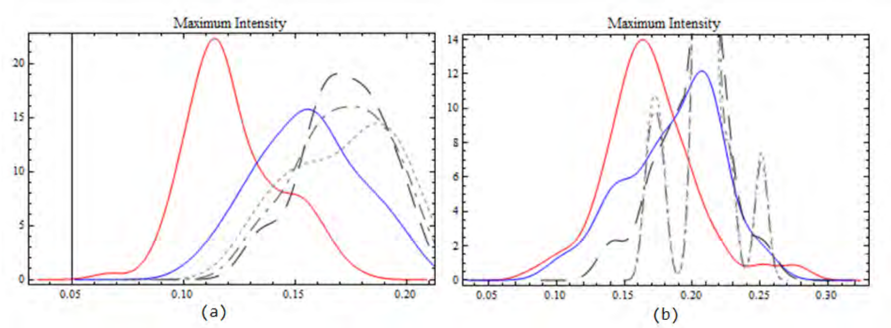

Figure 44: Histograms of two different individuals concerning the maximum intensity of the touch. These two histograms show the tendency observed in most of the data: stressed individuals touch the screen with more intensity.

sity of the touch is analysed. The initial assumption is that a more stressed user touches the screen with more intensity. The results obtained prove the assumption (see Figure 44(a) and Figure 44(b) that provide two example histograms). In both cases there is a clear shift in the values of the intensity towards higher values. Moreover, the differences observed between the distributions of the baseline data and the stressed data are statistically significant: p-value = $1.94289 * 10^{-11}$ for Figure 44 (a) and p-value = 0.00169036 for Figure 44 (b). And, if taking into consideration the second order analysis, we notice that the higher the level of stress, the more the distributions are shifted to the right (black lines in Figure 44).

This is, in fact, the general tendency: for all the participants who show a statistically significant difference, higher levels of stress are associated to an increased touch intensity. From a global point of view, considering a first order analysis, 53.3(3)% of the users under stress exhibit important differences in their touch intensities. If considering the second order analysis, this value rises to 60% of the users. Thus

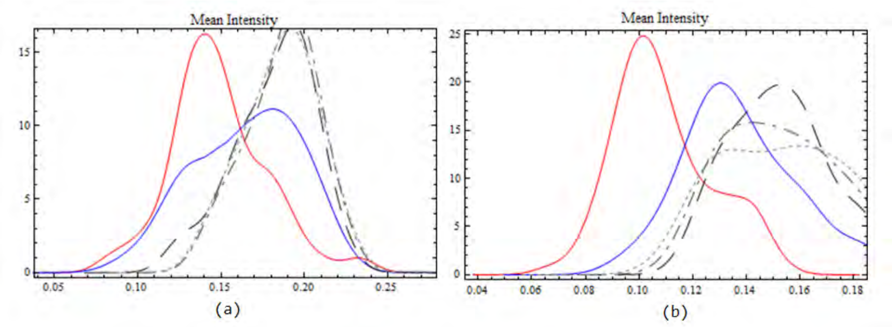

Figure 45: Histograms of two participants concerning the average intensity of the touch. As in almost every participant, the average intensity of the touch increases with increased levels of stress.

being, it is concluded that for approximately half of the users, the maximum intensity of the touch is significantly related to stress.

Mean intensity

A conclusion similar to the one of the previous section is achieved when the object of the analysis is the average value of the intensity during the touches. As depicted in Figure 45, the mean value of the intensity tends to increase as greater levels of stress are considered. These two particular examples have a p-value = 0.00265927 for Figure 45(a) and p-value = $6.5901 * 10^{-11}$ for Figure 45(b), which means that the differences observed are statistically significant. Moreover, when analysing data from all the users, the results seem to be slightly better than the ones of the maximum intensity. In fact, in a first order analysis 60% of the users show a considerable difference, while in a second order analysis this number rises to nearly 73.3%. In this sense, the average intensity of the touch can also be used to detect the effects of stress.

Amount of movement

The amount of movement represents a measure of the movement in front of the camera. The process is carried out from the information captured by cameras placed in the users' environment and using computer vision techniques to extract features regarding the users' states. The initial hypothesis is that a stressed user moves more and in more sudden ways when under an increased level of stress. However, the results obtained point the other way around: when users are under increased levels of stress they tend to move less. In fact, it was observed that users become more tense, rigid and highly focused on the tasks they are doing, decreasing the normal movements that people have when they are calmer.

Qualitative examples are depicted in Figure 46 and Figure 47. The differences between the movement detected in the two scenarios are

Figure 46: Results of movement detection on a non-stressed user. Row a) shows input images, row b) shows binarized and filtered movement and row c) shows the amount of movement detected.

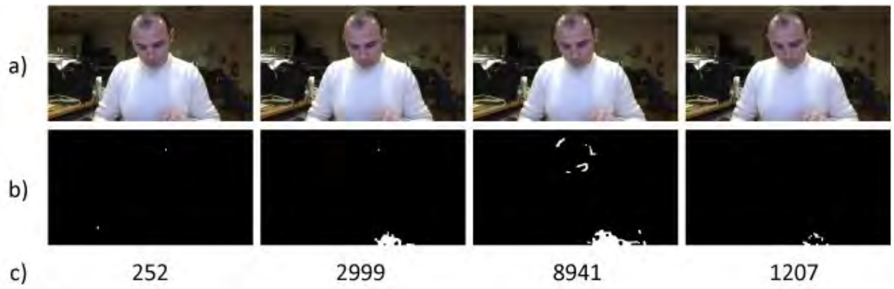

Figure 47: Results of movement detection on a stressed user. Row a) shows input images, row b) shows binarized and filtered movement and row c) shows the amount of movement detected.

noticeable. The first excerpt belongs to a non-stressed user while the second one belongs to the same user in a more stressed state. Furthermore, a quantitative example is offered in Figure 48. Given the nature of the data concerning the amount of movement, histograms are not the best way to depict them graphically. Instead, a Box-and-Whisker plot shows how the values are distributed for each of the five different analysis is shown. The results show how the user moves more during the calm phase. On the contrary, when the level of stress increases, the user moves less (Figure 48 (a)). More accurate values for this specific example are shown in Figure 48(b). It is possible to see that, not only the average value of the amount of movement decreases but also the values of the standard deviation.

From a general point of view, the analysis of the amount of movement shows statistically significant differences for around 47% of the users in a first order analysis. In a second order one, this value increases to 60%. Although the results achieved with this parameter point out to a different conclusion than the expected one, the parameter can be considered for the estimation of the levels of stress.

Duration of touch

The duration of the touch defines the time between the beginning and the end of each touch. The initial assumption is that stressed

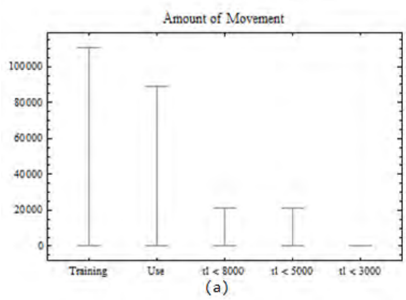

Amount of Movement – User: catalao		
	Mean	Standard Deviation
Training Data	2852.34	12140.
All Live Data	2980.33	12617.
timeleft < 8000	282.193	1991.41
timeleft < 5000	410.697	2479.05
timeleft < 3000	22.7143	105.158

(b)

Figure 48: An example of how the amount of movement generally varies with the amount of stress: higher stress is related to a smaller amount of movement. (a) Range of the values of movement for each dataset. (b) Average and standard deviation for the different datasets for a specific user.

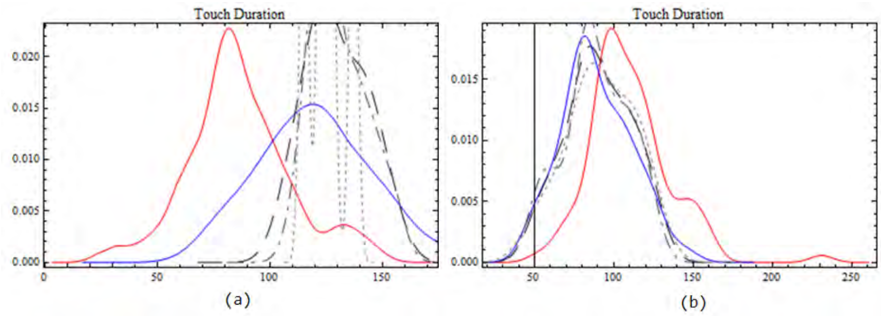

Figure 49: Histograms of two different participants concerning the duration of their touches. This feature does not have a homogeneous behaviour: some stressed participants have shorter touches while others have longer ones.

users have longer touches (also backed up by the assumption that stressed users have more intense touches). In fact, according to the results obtained, it is not possible to state that there is a marked tendency towards this: there are participants that have longer touches when stressed while others have shorter ones. Figure 49 depicts two examples of these opposing behaviours: (a) shows data from a participant whose touch duration increases with stress, while (b) shows data from a participant whose touch duration decreases. The fact that there is a time limit in each round may also have an influence on the results, as participants must finish the calculation rapidly, making use of shorter touches.

Nevertheless, both examples depict statistically important differences: p-value $= 2.70933 * 10^{-8}$ for (a) and p-value $= 9.54313 * 10^{-6}$ for (b). In fact, when analysing data from all the participants, nearly 47% show significant differences in a first order analysis and around 60% show significant differences in a second order analysis. Thus being, although different users react differently for this parameter, it still provides valuable data about how the individual user is affected.

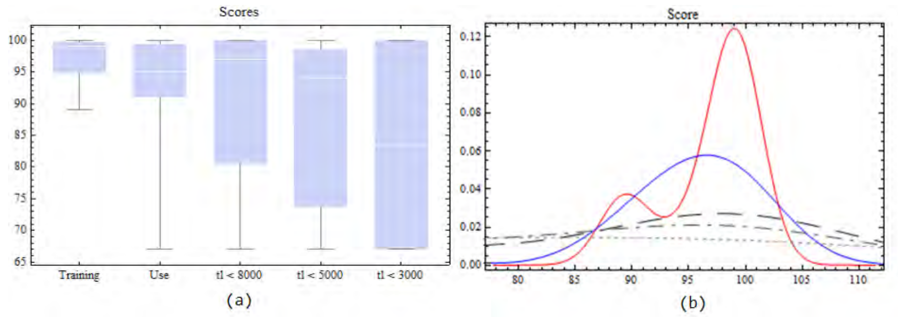

Figure 50: Box-and-Whisker plot denoting the distribution of the scores for a given participant, in different levels of stress (a). The same data is shown in (b). Although the several distributions are visually different, these differences are not statistically significant. However, the general tendency is that participants under stress tend to have worse performances.

Scores

In analysing the scores of the users, the main objective is to determine to which extent the cognition of the users is affected by stress. As it would be expected, the scores achieved by the users tend to decrease as the level of stress increases. Figure 50(a) shows this tendency for a specific user: without stressors, it is fairly easy and common to achieve good scores, with the average scores decreasing when under stress. The histogram depicted in Figure 50(b) depicts the distribution of the same data. For this specific user, although a visual analysis may point otherwise, the differences are not statistically significant: the Mann-Whitney test returns a p-value = 0.398694.

In fact, when analysing all the users, the score does not appear as a parameter that can be said to be related to stress. In a first order analysis, only around 7% of the users show statistically significant differences between the scores in the baseline data and in the stressed data. Moreover, in a second order analysis this number only increases to 13%. In that sense, the score does not seem to be very preponderant in analysing the effects of stress in an individual. Nevertheless, score is solely used in these experiments as one way to induce an objective "worth fighting for", thus making the participants commit to the experiment. In any case, the concept of score is most likely not present in a real conflict resolution scenario or, if it is, it is certainly estimated differently. Thus being, the fact that a low number of significant differences in analysing score was found constitutes no drawback.

Accuracy

Here, the accuracy of the touches of the participants is analysed. Accuracy defines the relation between the touches on active versus passive areas. The type of data from this parameter is different from the

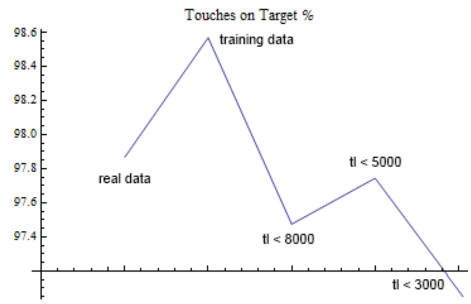

Figure 51: This plot shows the percentage of touches on target in the five different analysis for all the participants.

previous ones. In fact, instead of having a list of continuous values, there is a list of *true* or *false* values, indicating whether touches did or did not occur in an active area. Thus being, this parameter is not analysed with the Mann-Whitney test but by looking at plots such as the one depicted in Figure 51, that shows how the accuracy varies when considering the data collected from all the users.

It can be concluded that the accuracy of the touch remains relatively high, even for increased levels of stress (above 95%). However it is possible to identify the expected decreasing tendency.

8.3.2 *Classifiers*

The main goal of this section is to analyse the datasets built in order to determine if it is possible to train classifiers for the purpose of analysing the participants' behaviours. A feature-by-feature analysis is performed, in which the performance of the classifiers is analysed according to indicators such as the correctly classified instances and kappa statistics.

Acceleration

For each participant, a classifier was trained in order to classify each instance of the acceleration measured on the handheld device as "stressed" or "not stressed". In the best case, the classifier was able to correctly classify 99.85% of the instances ($\kappa = 0.995$). The worst classifier trained for a user was successful in 95.36% of the instances ($\kappa = 0.866$). When analysing all the classifiers built for all the participants, the average value of correctly classified instances is 98.1%. The median is 98.01%, the lower median is 96.92% and the upper median 99.85% (Figure 52). For all this data, $\overline{\kappa} = 0.94$. Given these results, it is possible to conclude that this classifier performs very well when used to classify the acceleration on the handheld device. This is in line with the very satisfying results of the significance tests on this feature.

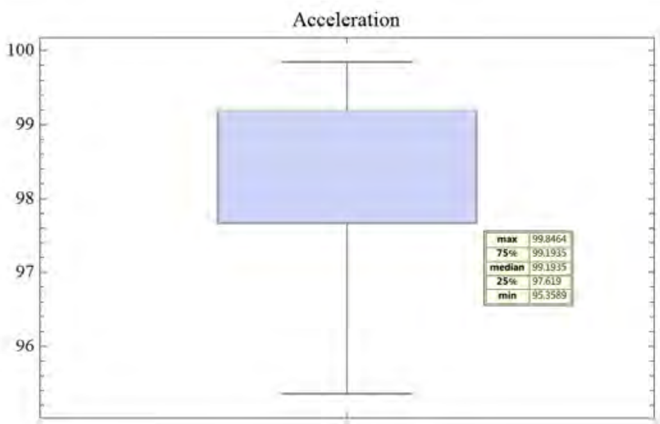

Figure 52: Box-and-whisker diagram detailing the sample accuracies of the classifiers trained for the acceleration feature.

Amount of Movement

Another classifier was trained to classify the amount of movement of each user. The best classifier trained was able to correctly classify 97.41% of the instances ($\kappa = 0.86$). On the other hand, the worst classifier trained from the data gathered resulted in a sample accuracy of 56.15% ($\kappa = 0.03$). When a global analysis is performed of all the classifiers built for this feature, it is possible can conclude that the average sample accuracy is of 78.84%. Looking at the distribution of the accuracy results, the median is 78.89%, the lower median 75.53% and the upper median 82.72% (Figure 53). The average value of the Cohen's kappa coefficient for this data is $\overline{\kappa} = 0.23$.

The results point out that it is possible to train such classifiers that look at the amount of movement of the participants to categorize it as "stressed" or "not stressed". However, this data has not shown as good results as the acceleration on the handheld device. Especially negative is the lower value of the Cohen's kappa coefficient. Moreover, unlike all the remaining features, the sample accuracy of these classifiers shows a large variation. Nevertheless, the positive results of the sample accuracy are encouraging and validate the approach.

Maximum touch intensity

The objective of analysing this feature is to train a classifier able to distinguish between "stressed" and "not stressed" touches by looking at the maximum intensity of each touch. The best classifier that resulted from this exercise was able to correctly classify 86.46% ($\kappa = 0.71$) of the instances. On the other hand, the worst classifier classified 71.82% ($\kappa = 0.21$) of the instances. When a broad analysis of all the classifiers trained for this feature is made, the average sample accuracy is 77.56%. The median value of the accuracy is 76.487% while the lower median is 75.21% and the upper median is 86.46% (Figure 54). The

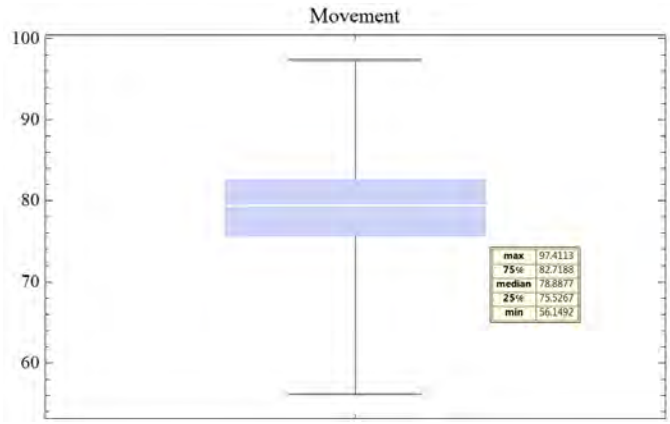

Figure 53: Box-and-whisker diagram detailing the sample accuracies of the classifiers trained for the movement feature.

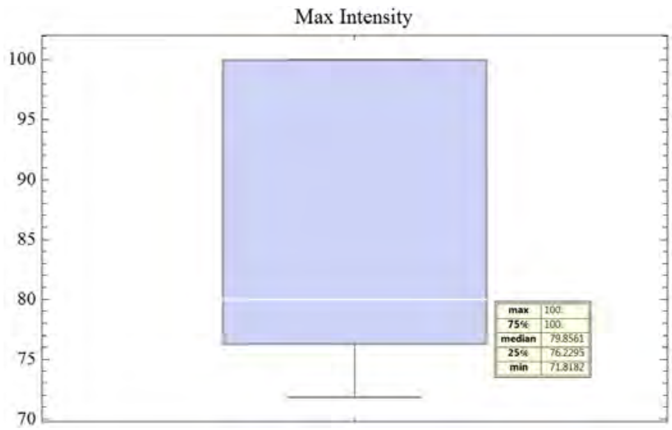

Figure 54: Box-and-whisker diagram detailing the sample accuracies of the classifiers trained for the maximum touch intensity feature.

average value of the Kappa coefficient is $\overline{\kappa} = 0.43$. The classifiers trained for this parameter have shown the worst results of this study. It is argued that this is due to the fact that this parameter shows a significant variability, sometimes due to outlayer values that do not reflect the values of intensity measured during the touch. Not only are the values of the performance relatively lower than the others, but also the values of the Cohen's Kappa coefficient.

Average touch intensity

The aim of analysing this feature is similar to the previous one. The only difference is that instead of considering the maximum pressure exerted on the screen during the touch, the average value during the touch is considered. The best classifier trained with this data was able to correctly classify 100% of the touches ($\kappa = 1.0$). On the opposite side, the classifier with the worst performance for this parameter classified 87.79% of the instances correctly ($\kappa = 0.69$). The average value of the sample accuracy when analysing all the classifiers trained

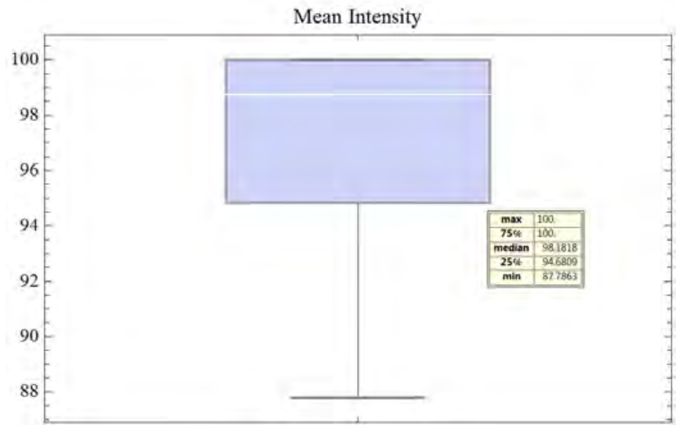

Figure 55: Box-and-whisker diagram detailing the sample accuracies of the classifiers trained for the average touch intensity feature.

was of 95.13%. The median value of the accuracy of the classifiers is 95.87%, the lower median is 92.14% and the upper median is 98.18% (Figure 55). The average value of the Cohen's Kappa is $\overline{\kappa} = 0.89$. This classifier also shows some very interesting results, significantly better than the ones of the maximum touch intensity. This may be related to the fact that the maximum intensity of the touch often does not reflect the value of the pressure during the touch and may be an outlayer value. On the other hand, the average value shapes in a better way the value of the pressure exerted during the touch. This results in a classifier with a fair better accuracy.

Duration of touch

With the analysis of this feature we want to train a classifier that looks at the duration of the touches in order to assign a label to each value, marking them as "stressed" or "not stressed". The classifier with the best performance was able to correctly classify 93.92% of the touches ($\kappa = 0.86$). On the other hand, the classifier with the worst performance classified 80.37% of the instances ($\kappa = 0.56$). In average, the sample accuracy was of 87.32%. Looking at the Box-and-whisker diagram for the performance of the classifiers, we conclude that the median value is 86.26%, while the lower median is 85.12% and the upper median 89.34% (Figure 56). The average value of the Cohen's coefficient is $\overline{\kappa} = 0.71$. On the overall the results of this classifier are satisfactory, given the high sample accuracy verified and the generally high value of the measure of inter-rater agreement chosen.

Intensity curve

The last feature considered in this study is the nature of the touches on the touchscreen. The objective is to classify the variation of intensity and time during a touch as belonging or not to a stressed user.

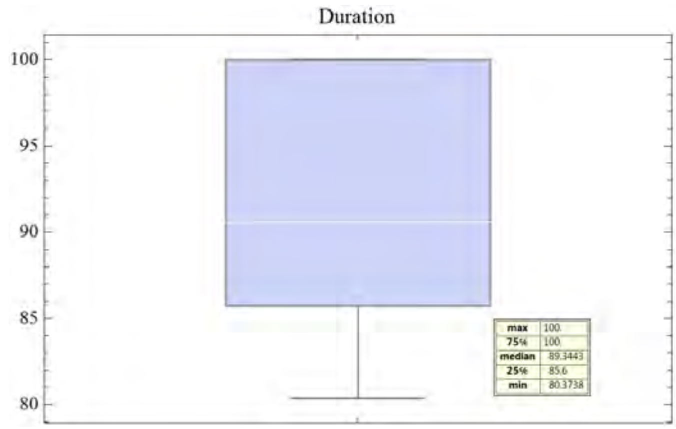

Figure 56: Box-and-whisker diagram detailing the sample accuracies of the classifiers trained for the touch duration feature.

However, each touch has an arbitrary duration and this results in more or less intensity points being generated during the touch. Thus, given that the number of values is arbitrary for each touch, it is not possibly to simply make a classifier using this data as it is. The approach used to deal with this issue was described in Section 7.5.4.

In the case of this feature, the meta-classifier `weka.classifiers.meta.CVParameterSelection` provided by weka was used that allows to optimize a given base-classifier. After finding the best possible configuration of parameters, the meta-classifier trains an instance of the base classifier with these parameters and uses it for subsequent predictions. The meta-classifier was used with lower bound 0.01, upper bound 0.5 and 10 optimization steps.

When using the J48 classification tree as the base classifier for the meta-classifier, the model is able to correctly classify 271 out of the 349 instances, which amounts to 77.6504%. The Kappa statistic for this model is 0.5434 and the value of the ROC area is 0.796. The constructed tree has a size of 15 nodes and a total of 8 leaves (Figure 57 (a)). In this tree, attributes `x0,x1` and `x2` correspond to the values of `a, b` and `c` of the polynomial curve, respectively. Given this, it is possible to use the rules of this tree to build a classifier for distinguishing between stressed and calm touch curves.

When the SMO function is used to build a classifier, the results achieved are similar. In fact, the correctly classified instances amount to 79.9427% (279 out of 349), the value of the Kappa statistics is 0.5809 and the value of the ROC area is 0.781. These results also show that a classifier can be trained with this data to distinguish between stressed and calm touches. Given that the results of both classifiers are similar, the decision was on using the J48 tree since the rules it generates can easily be used to classify touches in real time.

The classifier was assessed with data from 16 of the participants. Concerning this data, 13 participants show an increase in the touches classified as stressed by this classifier, when comparing the data from the baseline with the data from the highest stress. The minimum value of increase detected was of 6%, the maximum value of increase was of 60% and the average increase of touches classified as stressed, for all users, was of 32.3077%. The three participants for which the classifier reported a decreasing percentage of stressed touches for increased levels of stress have shown relatively low values of decrease (-2.5%, -5% and -1%). This means that the results of the classifier are consistent with the ones previously achieved in 81.25% of the cases. Figure Figure 57 (b) depicts the average increase of the touches classified as "stressed" in each of the five levels of stress of the experiment.

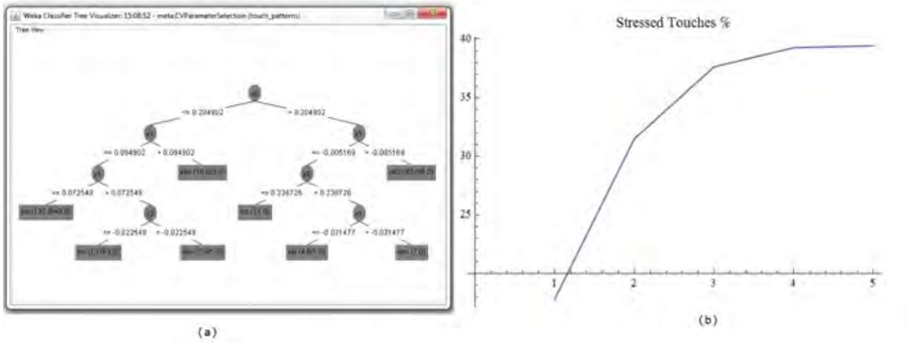

Figure 57: (a) J48 prunned tree generated by the algorithm. This tree can be used to classify touches in real time as stressed ("yes" leaves) or not stressed ("no" leaves). (b) Mean increase in the percentage of touches classified as "stressed" in each of the five levels of stress concerning all the users.

8.4 SUMMARY

When making an overview of the results achieved, it is possible to conclude that the feature that is most significantly affected by stress is the acceleration, with 80% of the participants exhibiting significant differences between baseline and stressed data. On the other hand, the parameter that presents the worst results is the score. However, as stated before, the score feature is merely indicative and is not considered in the real application. It is used solely with the purpose of inducing an objective on the user: without the objective of maximizing the score, the user would not *feel* stress.

It was also concluded that, as expected, all the parameters show a better performance in a second order analysis. However, as already stated, the first order analysis is preferred as it considers all the data (and not only data from more stressed states), which means that a statistically significant difference (when it exists) is more solid. Analysing all the data, each participant has in average 2.93(3) features with considerable differences (out of 6). Moreover, in the worst case there are 3 participants with only 1 parameter with significant

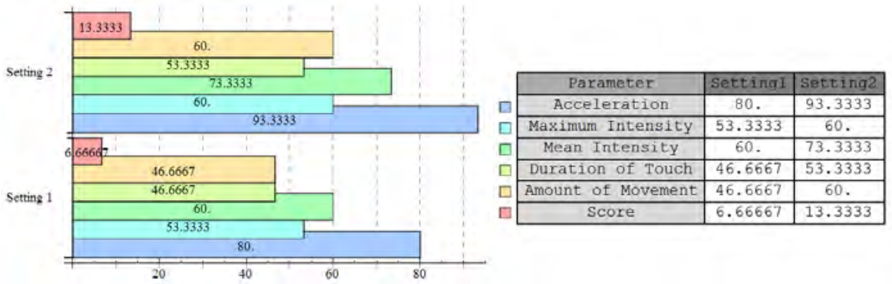

Figure 58: Percentage of users that reveal significant differences between baseline and stressed data considering a first order analysis (Setting 1) and second order analysis (Setting 2).

Table 6: Summary of the results of the analysis of the performance of the classifiers

Dataset	Best		Worst		Average	
	%	κ	%	κ	%	$\overline{\kappa}$
Acc.	99.85	0.995	95.36	0.866	98.1	0.94
Mov.	97.41	0.86	56.15	0.03	78.84	0.23
Max.	86.46	0.71	71.82	0.21	77.56	0.43
Avg.	100	1.0	87.79	0.69	95.13	0.89
Dur.	93.92	0.86	80.37	0.56	87.32	0.71

differences, and in the best case there are 4 participants with 5 features showing significant differences. Given this, it can be stated that a generic model can be applied to users for whom no baseline data exists. However, personalized models are certainly more accurate and that is the aim.

From the results, it is also concluded that the most affected features are the ones that people do not conscientiously control, such as the acceleration of our hand gestures (i.e. how much they move) or the intensity of the touch. On the other hand, the parameters that are more rational (in this case the score) are not influenced in such a significant manner. This can be seen as positive, in the sense that unconscious behaviours and reactions are more difficult to forge and are usually true reactions of the human body. Thus, the results that stem from their analysis are more solid.

Concerning the performance of the classifiers developed with this data, the results are briefly summarized in Table 6. In each row the table describes the name of the parameter as well as the the best and worst classifier trained (in terms of the percentage of correctly classified instances) and the respective kappa coefficients. Finally, for each parameter it also shows the average performance of the classifiers.

Using the classifiers described it is possible to build a solution for estimating the level of stress of users, in a personalized way, since

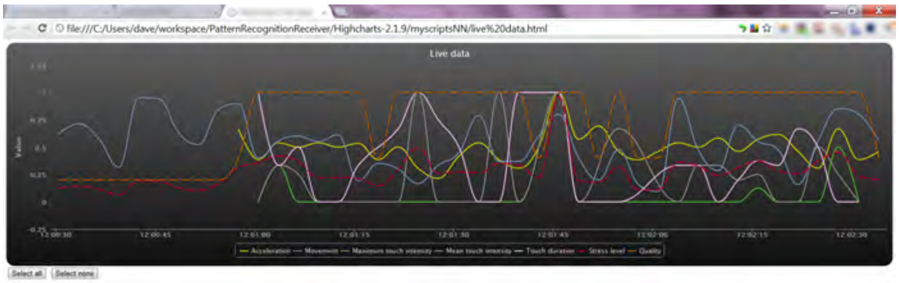

Figure 59: The evolution of the level of stress for a given user. The red dashed line represents the level of stress, the orange dashed line represents the quality of the information and the remaining lines represent the contribution of each parameter for the level of stress computed.

each user has a number of classifiers that were trained using his personal interaction patterns. These can thus be used in real time to classify each instance of data being generated by the stress sensor agents. The information generated is then depicted graphically. Figure Figure 59 shows a prototype of a web interface developed for such purpose. The dashed red line depicts the level of stress computed while the dashed orange line depicts the quality of that information.

The quality of the information generated is computed based on the type of inputs available and their significance, i.e., the assessment of stress is more reliable (higher quality of the information) if it is based on features that have shown very significant differences. On the other hand, if it is based on a low number of inputs or on inputs that have shown low levels of significance, the information compiled is not very reliable. This measure allows the decision-makers to have a more realistic view on the data provided. The remaining lines represent the contribution of each stress source for the overall computation of stress. They allow to determine in which feature the user is being more or less affected by stress.

9

QUANTITATIVE RESULTS

Measure what can be measured,
and make measurable what cannot be measured.

— Galileo Galilei

introduçao

9.1 PUBLICATIONS

The work described in this manuscript was disseminated through the participation in diverse scientific meetings in the fields of Computer Science and The Law, as well as through the publication in international journals and conference proceedings.

9.1.1 *International Journals*

Knowledge and Information Systems

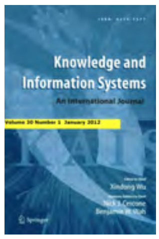

In this publication a decision-support system for conflict resolution is described. This system uses Case-based Reasoning techniques in order to retrieve information that can be useful during the conflict resolution process. Then, using that information, if provides support for decision making in the context of negotiated procedures. The system is validated through performance measures.

Carneiro D., Novais P., Andrade F., Zeleznikow J., Neves J., Using Case Based Reasoning and Principled Negotiation to provide Decision Support for Dispute Resolution. Knowledge and Information Systems. (Accepted to appear).

Impact Factor: 2.225
DOI: http://dx.doi.org/10.1007/s10115-012-0563-0
Indexed in: DBLP, ISI Web of Knowledge (Document Type: Article)

Expert Systems With Applications

This publication details the experiment performed in order to collect data about how users react to stress. Specifically, issues like movement and behavioural patterns were analysed in order to determine how people move and how people interact with technological devices, when under stress. The models developed allow to estimate the stress of users in real-time, in a non-invasive way.

Carneiro D., Carlos Castillo J., Novais P., Fernández-Caballero A., Neves J., Multimodal Behavioural Analysis for Non-invasive Stress Detection, Expert Systems with Applications, Volume 39, Issue 18, 15 December 2012, Pages 13376—13389, Elsevier, 2012.

Impact Factor: 2.203
DOI: http://dx.doi.org/10.1016/j.eswa.2012.05.065.
Indexed in: DBLP, ISI Web of Knowledge (Document Type: Article)

Neurocomputing

This paper describes a biologically-inspired method for generating valid solutions for a conflict resolution scenario with specific characteristics. Specifically, genetic algorithms are used to search virtually the whole solution space, generating solution trees that can then be followed during the negotiation process.

Carneiro D., Novais P., Neves J., Using Genetic Algorithms to Create Solutions for Conflict Resolution, Neurocomputing. Elsevier, ISSN: 0925-2312 (In Press).

Impact Factor: 1.580
DOI: http://dx.doi.org/10.1016/j.neucom.2012.03.024
Indexed in: DBLP, ISI Web of Knowledge (Document Type: Article)

Artificial Intelligence Review

In this journal an overview of the current state of the art of Online Dispute Resolution if performed, from the point of view of Artificial Intelligence. Specifically, current approaches are analysed critically, both in the commercial and academic fields. Some possible future lines of research to address the main challenges identified are put forward.

Carneiro D., Novais P., Andrade F., Zeleznikow J., Neves J., Online Dispute Resolution: an Artificial Intelligence Perspective, Artificial Intelligence Review, Springer, ISSN: 0269-2821 (2012)

Impact Factor: 1.213
DOI: http://dx.doi.org/10.1007/s10462-011-9305-z.
Indexed in: DBLP, ISI Web of Knowledge (Document Type: Article)

Expert Systems

This paper focused on the (sometimes ignored) importance of contextual information in the traditional conflict resolution process. Particularly, it analysis how similarly rich communication processes can be implemented online, shaping the ones undertaken face-to-face.

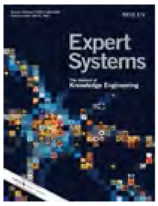

• Carneiro D., Gomes M., Costa A., Novais P., Neves J., Enriching Conflict Resolution Environments with the Provision of Context Information, Expert Systems, Wiley-Blackwell, ISSN: 0266-4720, (Accepted to appear in 2013).

Impact Factor: 0.769
To be indexed in: DBLP, ISI Web of Knowledge (Document Type: Article)

9.1.2 *Springer Book Chapters*

1. Gomes M., Oliveira T., Carneiro D., Novais P., Neves J.: Extracting Behavioural Patterns from a Negotiation Game. In Highlights on Practical Applications of Agents and Multi-Agent Systems Communications in Computer and Information Science Volume 365, 2013, pp 304-315.

   http://link.springer.com/chapter/10.1007%2F978-3-642-38061-7_29

2. Carneiro D., Novais P., Catalão F., Marques J., Pimenta A., Neves J.: Dynamically Improving Collective Environments through Mood Induction Procedures. Ambient Intelligence - Software and Applications - Advances in Intelligent Systems and Computing Volume 219, 2013, pp 33-40.

http://link.springer.com/chapter/10.1007%2F978-3-319-00566-9_5

3. Rodrigues M., Gonçalves S., Carneiro D., Novais P., Fdez-Riverola F.: Keystrokes and Clicks: Measuring Stress on E-learning Students. Management Intelligent Systems - Advances in Intelligent Systems and Computing Volume 220, 2013, pp 119-126.

   http://link.springer.com/chapter/10.1007%2F978-3-319-00569-0_15

4. Novais P., Carneiro D., Gomes M., Neves J., The Relationship between Stress and Conflict Handling Style in an ODR Environment, New Frontiers in Artificial Intelligence - JSAI-isAI 2012 Workshops, LENLS, JURISIN, AMBN, ISS, Revised Selected Papers. Lecture Notes in Computer Science, Springer (Accepted to appear 2013).

5. Carneiro D., Gonçalves S., Novais P., and Neves J.: Studying Stress on e-Learning Users. Accepter do appear, EPIA14 – Encontro Português de Inteligência Artificial.

6. Carneiro D., Novais P., Gomes M., Oliveira P., Neves J.: A Statistical Classifier for Assessing the Level of Stress from the Analysis of Interaction Patterns in a Touch Screen. In: Advances in Intelligent Systems and Computing, Volume 188, 257–266, DOI: 10.1007/978-3-642-32922-7_27, 2012.

   http://www.springerlink.com/content/v2vt355pu5663617/

7. Novais P., Carneiro D., Gomes M., Neves J.: Non-invasive Estimation of Stress in Conflict Resolution Environments. In: Advances on Practical Applications of Agents and Multi-Agent Systems - 10^{th} International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2012), Yves Demazeau, Jörg P. Müller, Juan M. Corchado Rodríguez, and Javier Bajo Pérez (Eds.), Springer - Series Advances in Intelligent and Soft Computing, vol. 155, pp 153–160, ISBN 978-3-642-28785-5, 2012.

   http://rd.springer.com/chapter/10.1007/978-3-642-28786-2_17

8. Novais P., Carneiro D., Neves J.: Incorporating Stress Estimation into User-Centred Agent-Based Platforms. In: Advances on Practical Applications of Agents and Multi-Agent Systems - 10^{th} International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2012) (Demo paper), Yves Demazeau, Jörg P. Müller, Juan M. Corchado Rodríguez, and Javier Bajo Pérez (Eds.), Springer - Series Advances in Intelligent and Soft Computing, vol. 155, pp 277–282, ISBN 978-3-642-28785-5, 2012.

   http://www.springerlink.com/content/164052v05723528w/

9. Carneiro D., Analide C., Novais P., Neves J.: An Evolutionary Approach to Generate Solutions for Conflict Scenarios. Emilio Corchado, Václav Snásel, Ajith Abraham, Michal Wozniak, Manuel Graña, Sung-Bae Cho (Eds.): Hybrid Artificial Intelligent Systems - 7^{th} International Conference, HAIS 2012, Springer - Lecture Notes in Computer Science 7208, ISBN 978-3-642-28941-5, pp 440–451, 2012.

   http://www.springerlink.com/content/148u3481n5393309/

10. Gomes M., Carneiro D., Novais P., Neves J.: Modelling Stress Recognition in Conflict Resolution Scenarios. Emilio Corchado, Václav Snásel, Ajith Abraham, Michal Wozniak, Manuel Graña, Sung-Bae Cho (Eds.): Hybrid Artificial Intelligent Systems - 7^{th} International Conference, HAIS 2012, Springer - Lecture Notes in Computer Science 7208, ISBN 978-3-642-28941-5, pp 533–544, 2012.

    http://rd.springer.com/chapter/10.1007/978-3-642-28942-2_48

11. Carneiro D., Carlos Castillo J., Novais P., Fernández-Caballero A., Neves J., López M.: Stress Monitoring in Conflict Resolution Situations. In Ambient Intelligence - Software and Applications - 3^{rd} International Symposium on Ambient Intelligence (ISAmI 2012), Paulo Novais, Kasper Hallenborg, Dante I. Tapia, and Juan M. Corchado Rodríguez (Eds.), Springer - Series Advances in Intelligent and Soft Computing, vol. 153, pp 137–144, ISBN 978-3-642-28785-5, 2012.

    http://www.springerlink.com/content/dr887t24q883657l/

12. Carneiro D., Lemos F., Novais P., Neves J.: Issues on Conflict Resolution in Collaborative Networks. PRO-VE'11, Adaptation and Value Creating Collaborative Networks, IFIP Advances in Information and Communication Technology, Volume

362/2011, 271–278, DOI: 10.1007/978-3-642-23330-2_30, 2011.

http://link.springer.com/chapter/10.1007/978-3-642-23330-2_30

13. Carneiro D., Gomes M., Novais P., Neves J.: Developing Dynamic Conflict Resolution Models Based on the Interpretation of Personal Conflict Styles. In: Progress in Artificial Intelligence, Antunes L., Sofia Pinto H. (eds), Lecture Notes in Artificial Intelligence 7026, Springer-Verlag, ISBN 978-3-540-24768-2, pp 44–58, 2011.

    http://rd.springer.com/chapter/10.1007/978-3-642-24769-9_4

14. Carneiro D., Novais P., Neves J.: An Agent-based Architecture for Multifaceted Online Dispute Resolution Tools. In: Developing Concepts in Applied Intelligence, Mehrotra et al. (eds), Springer – Studies in Computational Intelligence, ISBN: 978-3-642-21331-1, pp. 89-—94, The Twenty-fourth International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE 2011), 2011.

    http://rd.springer.com/chapter/10.1007/978-3-642-21332-8_13

15. Carneiro D., Novais P., Neves J.: Toward Seamless Environments for Dispute Prevention and Resolution, International Symposium on Ambient Intelligence, Advances in Intelligent and Soft Computing, Springer, Volume 92/2011, 25–32, DOI: 10.1007/978-3-642-19937-0_4, 2011.

    http://rd.springer.com/chapter/10.1007/978-3-642-19937-0_4

16. Carneiro D., Novais P., Machado L., Analide C., Costa N., Neves J.: Role Playing Games and Emotions in Dispute Resolution Environments. In: Advances in Intelligent and Soft Computing, Springer, Volume 87/2011, 155–162, DOI: 10.1007/978-3-642-19644-7_17, 2011.

    http://www.springerlink.com/content/nq201r36x6417072/

17. Carneiro D., Novais P., Andrade F., Neves J.: Improving Mediation Processes with Avoiding Parties. New Frontiers in Artificial Intelligence - JSAI-isAI 2010 Workshops, Revised Selected Papers, T. Onoda, D. Bekki, E. McCready (Eds) Springer – LNAI, 2011.

http://www.springerlink.com/content/xg4n3095l1620132/

18. Costa N., Carneiro D., Novais P., Barbieri D., Andrade F.: An Advice System for Consumer's Law Disputes. In: Enterprise Information Systems (12th International Conference, ICEIS 2010, Revised Selected Papers), Joaquim Filipe and José Cordeiro (Eds), Springer: Lecture Notes in Business Information Processing, Volume 73, ISBN: 978-3-642-19802-1, pp 237–248, 2011.

    http://rd.springer.com/chapter/10.1007/978-3-642-19802-1_17

19. Carneiro D., Novais P., Costa R., Neves J.: Developing Intelligent Environments with OSGi and JADE. Artificial Intelligence in Theory and Practice III, Bramer M. (Ed), Springer-Verlag, Series: IFIP International Federation for Information Processing, ISBN 978-3-642-15285-6, pp 174–183, The Third IFIP International Conference on Artificial Intelligence in Theory and Practice (IFIP AI 2010), 2010.

    http://www.springerlink.com/content/v4p1216mm1t84464/

20. Carneiro D., Novais P., Andrade F., Zeleznikow J., Neves J.: Using Case-based Reasoning to Support Alternative Dispute Resolution. In: Distributed Computing and Artificial Intelligence, Andre Carvalho, Sara Rod-González, Juan Paz and Juan M. Corchado (Eds), 7th International Symposium on Distributed Computing and Artificial Intelligence (DCAI 2010), Springer - Series Advances in Intelligent and Soft Computing, vol. 79, ISBN: 978-3-642-14882-8, pp. 123–130, 2010.

    http://www.springerlink.com/content/e3030372442r4304/

21. Andrade F., Novais P., Carneiro D., Zeleznikow J., Neves J.: Using BATNAs and WATNAs in Online Dispute Resolution. In: New Frontiers in Artificial Intelligence, Kumiyo Nakakoji, Yohei Murakami and Eric McCready (Eds), (JSAI-isAI 2009 Workshops, LENLS, JURISIN, KCSD, LLLL, Tokyo, Japan, 2009, Revised Selected Papers), Springer - LNAI 6284, ISBN 978-3-642-14887-3, pp. 5–18, 2010.

    http://www.springerlink.com/content/e424237450682x76/

22. Carneiro D., Novais P., Costa R., Neves J.: Enhancing the Role of Multi-agent Systems in the Development of Intelligent Environments. In Trends in Practical Applications of Agents and Multiagent Systems, Yves Demazeau et al. (eds), in Advances in Intelligent and Soft Computing, Vol. 71, Springer- Verlag,

ISBN 978-3-642-12432-7, pp. 123–130, (PAAMS 10 - The 8th International Conference on Practical Applications of Agents and Multi-Agent Systems), 2010.

http://www.springerlink.com/content/2m12x5n5084t1652/

9.1.3 *Book Chapters*

1. Andrade F., Carneiro D., Novais P., Neves J.: Resolução de conflitos em linha na contratação pública electrónica? A Arbitragem Administrativa e Tributária: Problemas e Desafios, Isabel Celeste M. Fonseca (Ed.), Edições Almedina, S.A., ISBN 978-972-40-4659-4, 2012. *in portuguese*

2. Andrade F., Novais P., Carneiro D., Neves J.: Conflict Resolution in Virtual Locations. In: Information Communication Technology Law, Protection and Access Rights: Global Approaches and Issues, Portela I., Cunha M. (Eds), IGI Global, ISBN: 978-1615209750, pp 33–50, 2010.

   http://www.igi-global.com/chapter/conflict-resolution-virtual-locations/43486

3. Andrade F., Carneiro D., Novais P.: A Inteligência Artificial na Resolução de Conflitos em Linha, Scientia Ivridica – Tomo LIX, n.º 321, pp 137-164, ISSN: 0870-8185, 2010. *in Portuguese*

4. Barbieri D., Carneiro D., Andrade F., Novais P.: Resolução de Conflitos em Linha - Uma aplicação ao direito do consumo, Scientia Ivridica, Tomo LIX - Número 323, pp 581-607, ISSN: 0870-8185, 2010. *in Portuguese*

9.1.4 *Peer-Reviewed Conference Proceedings*

1. Novais P., Carneiro D., Gomes M., Neves J.: Behavioural and Context Analysis in an Online Dispute Resolution Environment. In: Proceedings of the Sixth International Workshop on Juris-Informatics (JURISIN 2012), hosted by the Fourth JSAI International Symposia on AI, ISBN 978-4-915905-52-0.

2. Carneiro D., Novais P., Andrade F., Zeleznikow J., Neves J.: Context-aware Environments for Online Dispute Resolution. In: GDN 2012 - The 12th international annual meeting of the Group Decision and Negotiation conference, Recife, Brasil, 2012 (to appear).

3. - Carneiro D., Novais P., Andrade F., Neves J.: Retrieving Information in Online Dispute Resolution Platforms: A Hybrid

Method. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Law, University of Pittsburgh School of Law, ACM ISBN: 978-1-4503-0755-0, 2011.

4. Carneiro D., Gomes M., Novais P., Andrade F., Neves J.: Automatic Classification of Personal Conflict Styles in Conflict Resolution. In: Legal Knowledge and Information Systems - JURIX 2011: The Twenty-Fourth Annual Conference, Katie M. Atkinson (Ed), IOS Press - Frontiers in Artificial Intelligence and Applications, Vol. 235, ISBN 978-1-60750-980-6, pp 43–52, 2011.

5. Carneiro D., Novais P., Neves J.: Towards Domain-Independent Conflict Resolution Tools, The 2011 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2011), 2011.

   http://dl.acm.org/citation.cfm?id=2052174

6. Carneiro D., Costa A., Novais P., Andrade F., Neves J.: Providing Relevant Knowledge in Disputes. In Proceedings of the 6th International Workshop on Online Dispute Resolution, pp. 63–78, JURIX 2010, Liverpool, UK, 2010.

7. Costa N., Carneiro D., Novais P., Barbieri D., Andrade F.: An Agent-Based Approach To Consumer's Law Dispute Resolution. In: Proceedings of the ICEIS 2010 - 12th International Conference on Enterprise Information Systems, (ICEIS 2010, Madeira, Portugal, 8th-12th June 2010), ISBN: 978-989-8425-05-8, pp 103–110, 2010.

8. Andrade F., Barbieri D., Carneiro D., Novais P.: Artificial Intelligence Applications in ODR – Online Dispute Resolution: The UMCourt Project, Proceedings of the 17th Annual Northwest Dispute Resolution Conference, University of Washington, USA, 2010.

9. Carneiro D., Novais P., Andrade F., Neves J.: Using Mediation to Solve Disputes with Avoiding Parties. In: Proceedings of the JURISIN 2010 – Fourth International Workshop on Jurisinformatics, Tokyo, Japan, ISBN 978-4-915905-42-1, pp 17-28, 2010.

10. Machado L., Carneiro D., Analide C., Novais P.: Jogos de Papéis e Emoções em Ambientes Assistidos, INFORUM - 2º Simpósio de Informática, Luís Barbosa and Miguel Correia (Eds), Braga, Portugal, ISBN 978-989-96863-0-4, pp 745-756, 2010. *in portuguese*

11. Café A., Carneiro D., Andrade F., Novais P.: Sistema de Resolução Online de Conflito para Partilhas de bens – Divórcios

e Heranças, INFORUM - 2º Simpósio de Informática, Luís Barbosa and Miguel Correia (Eds), Braga, Portugal, ISBN 978-989-96863-0-4, pp 779-790, 2010. *in portuguese*

12. Carneiro D., Novais P., Andrade F., Zeleznikow J., Neves J.: The Legal Precedent in Online Dispute Resolution. In Legal Knowledge and Information Systems, ed. Guido Governatori (proceedings of the Jurix 2009 - the 22nd International Conference on Legal Knowledge and Information Systems, Rotterdam, The Netherlands), IOS press, ISBN 978-1-60750-082-7, pp. 47–52, 2009.

13. Andrade F., Novais P., Carneiro D., Zeleznikow J., Neves J.: Using BATNAs and WATNAs in Online Dispute Resolution. In: Proceedings of the JURISIN 2009 - Third International Workshop on Juris-informatics, Tokyo, Japan, ISBN 4-915905-38-1, pp. 15–26,2009.

9.2 PARTICIPATIONS IN SCIENTIFIC EVENTS

9.2.1 *Summer Schools*

During the period of this PhD, the candidate participated in one the summer school of the 2011 IEEE/WIC/ACM International Joint Conference on Web Intelligence (WI 2011) and Intelligent Agent Technology (IAT 2011), which took place in Lyon France.

The following topics were addressed:

- Orland Hoeber (Memorial University of Newfoundland, Canada): "Human-Centred Web Search: Searchers, Visual Interfaces, and Artificial Intelligence"
- Fabien Gandon (INRIA Sophia-Antipolis, France): "Social Web and Semantic web: knowledge graphs meet sociograms to capture social semantics"
- François Bancilhon (Chief Executive Officer of Data Publica): "Open Data and its impact in terms of information access on the Web"
- Devis Bianchini (Università degli Studi di Brescia, Italy): "Data Management and the Web of Data: interleaving two worlds"
- Katia Vila (University of Matanzas, Cuba): "Adapting Question Answering Systems for Restricted Domains"
- Maxime Morge (University Lille 1, France): "Argumentation technologies for agents and multiagent systems"

- Marco Calabrese (Polytechnic of Bari, Italy & University of Milan, Italy): Hierarchical-Granularity Holonic-Based Engineering: a Computational Intelligence Perspective
- Humbert Fiorino (Laboratory of Informatics of Grenoble, France) & Damien Pellier (Paris Descartes, France): "Automated Planning, a Key Component for Agent Rationality"
- Longbing Cao (University of Technology Sydney, Australia): "Complex Behavior Modeling, Analysis and Mining"

9.2.2 *Conferences and Invited Talks*

In the context of the PhD, the candidate participated in several scientific events in which preliminary results and findings of the work performed were presented. These meetings were important in order to collect insights and opinions from experts in the fields. The candidate has delivered the following oral presentations:

1. "Acquisition of context information for online dispute resolution", invited talk presented at the Universidad de Castilla-La Mancha, Albacete, Spain, 2012.
2. "An Evolutionary Approach to Generate Solutions for Conflict Scenarios", presented at HAIS 2012, Salamanca, Spain.
3. "Stress Monitoring in Conflict Resolution Situations", presented at ISAmI 2012, Salamanca, Spain.
4. "Non-invasive Estimation of Stress in Conflict Resolution Environments", presented at PAAMS 2012, Salamanca, Spain.
5. "Automatic Classification of Personal Conflict Styles in Conflict Resolution", presented at Jurix 2011, Vienna, Austria.
6. "Towards Domain-Independent Conflict Resolution Tools", presented at the International Conference on Web Intelligence and Intelligent Agent Technology, Lyon, France, 2011.
7. "Role Playing Games and Emotions in Dispute Resolution Environments", presented at the International Conference on Soft Computing Models in Industrial and Environmental Applications, Salamanca, Spain, April, 2011.
8. "Toward Seamless Environments for Dispute Prevention and Resolution", presented at the International Symposium on Ambient Intelligence, Salamanca, Spain, April, 2011.
9. "Using Case-based Reasoning to Support Alternative Dispute Resolution" presented at the International Symposium on Distributed Computing and Artificial Intelligence 2010 (DCAI'10), Valencia, Spain, September 2010

10. "Resolução de Disputas fora dos Tribunais" invited talk presented at the Master Course on Informatics, School of Engineering, University of Minho, Braga, Portugal, October, 2010.
11. "Sistema de Suporte para Resolução de Conflitos fora dos Tribunais" presented at the Week of the School of Engineering, University of Minho, Guimarães, Portugal, October, 2010.
12. "Providing Relevant Knowledge in Disputes: UMCourt Project", presented at the 6^{th} International Workshop on Online Dispute Resolution (ODR Workshop'10), Liverpool, United Kingdom, December 2010.
13. "The Legal Precedent in Online Dispute Resolution" presented at the 22^{nd} International Conference on Legal Knowledge and Information Systems (Jurix'09), Rotterdam, the Netherlands, December 2009.

9.3 ORGANIZATION OF SCIENTIFIC EVENTS

During the execution of his PhD plan, the candidate participated in the organization of the following scientific events:

1. Scientific Committee - International Symposium on Distributed Computing and Artificial Intelligence, Salamanca, Spain (DCAI 2012)
2. Program Committee - 6^{th} International Conference on Hybrid Artificial Intelligence Systems, Salamanca, Spain, 2011 (HAIS 2011)
3. Scientific Committee - International Symposium on Distributed Computing and Artificial Intelligence, Salamanca, Spain, 2011 (DCAI 2011)
4. Publicity and Web Chair - Ambient Intelligence Environments, Lisbon, Portugal, 2011 (AmI Environments 2011)
5. Local Organization Committee - International Symposium on Ambient Intelligence 2010, Guimarães, Portugal (ISAmI 2010)
6. Local Organization Committee - International Workshop on Soft Computing Models in Industrial Applications 2010, Guimarães, Portugal (SOCO 2010)

9.4 LECTURES

Lectures in the Master Course in Law and Informatics, in the "Online Dispute Resolution" Curricular Unit, the academic year of 2011/2012.

9.5 SUPERVISION OF STUDENTS

The candidate has successfully co-supervised several MSc students that developed their work in the context of this PhD thesis. The students and a summary of each work plan is described in this section.

FÁBIO ALEXANDRE MARQUES CATALÃO

Analysis of the Influence of Stress on the Interaction with Computers.
In co-supervision with prof. Paulo Novais.

While performing this work, Fábio is implementing an experiment similar to the ones described in this thesis. He is collecting data about how people interact with the computer (e.g. velocity of the mouse, keyboard typing rate, acceleration on the mouse, precision of the touch), in a total of more than 20 features. He is collecting data in two moments for each user, with and without stress. This work will result in a stress model similar to the one developed in this thesis that, by describing the interaction with the computer rather than with the smartphone, will complement it.

SÉRGIO MANUEL DE CARVALHO GONÇALVES

Evaluation as a Learning Moment – Stress Analysis in an e-learning Environment.
In co-supervision with prof. Paulo Novais.

In this work, Sérgio is applying the non-invasive approach on stress analysis detailed in this thesis in the domain of e-learning. Particularly, the interest is in analysing stress in students when taking exams on the Moodle platform. This work is important as it will allow professors to better understand which subjects stress students the most, thus allowing to learn from the moment of the evaluation.

MARCO VIEIRA GOMES

Behavioural and Context Analysis in an Online Dispute Resolution Enviroment.
In co-supervision with prof. Paulo Novais.

Some of the aspects detailed in thesis were implemented with the collaboration of Marco Gomes. Specifically, during the work that lead to his dissertation, Marco implemented the Negotiation Game that, in conjunction with the stress analysis approach described in this thesis, allowed to study how people behave during a negotiation, under the influence of stress.

FLÁVIO LEMOS
CONFLICT RESOLUTION IN VIRTUAL ORGANIZATIONS.
In co-supervision with prof. Paulo Novais.

Flávio was able to successfully take the agent-based architecture defined in this thesis and, with minimal changes, adapt it to the context of Virtual Organizations. His work consisted in studying the nature of conflicts in Virtual Organizations, identifying the particularities, and transforming that into system requirements. This resulted in a software solution that was able to detect conflicts and accompany the organizations from the inception of the conflict to its resolution, supporting them in the several phases.

9.6 RESEARCH STAY

In the context of the research work performed during the thesis, the candidate spent a research period in the University of Castilla-La Mancha. During this stay he was under the supervision of Prof. Dr. Antonio Fernandez-Caballero. The purpose of this partnership was to develop a link between the Universities of Minho and Castilla-La Mancha in the form of knowledge exchange and possible collaboration on current and future projects. It also increased the scientific content of the PhD work by including additional technologies and approaches to detect stress in a non-invasive way, using video sources.

The collaboration took place between the 2nd of July and the 5th of October, in 2012. The following tasks were implemented:

- Adaptation to the work methodology of the host group
- Analysis of the State of the Art in current projects
- Implementation of a software module to acquire high-level information from video sources to assess stress from user interaction patterns
- Analysis of the results achieved
- Validation by the supervisors

In practical terms, this work resulted in the development of a module for the interoperability between the agent-based architecture defined in this PhD work and the Int^3-Horus framework (a multilevel framework for intelligent multisensor monitoring and activity interpretation). This allowed for valuable data to be acquired concerning the behaviour of the individuals, namely describing their movement patterns.

9.7 PROTOTYPES AND CASE STUDIES

In the work described in this thesis several prototypes were implemented, using different legal domains as specific case-studies. All the functionalities described in the previous chapter were assessed through the implementation of prototypes in the domain of Labour Law. The prototypes described in this section were implemented by undergraduate students that collaborated with this work by implementing specific functionalities of interest in other domains. On the one hand, this allowed to validate the domain-independent nature of the architecture depicted in Chapter 4. On the other hand, it allowed to assess the aptitude of specific algorithms/approaches in more than one domain. Given that the prototypes on Labour Law domain were already presented throughout this thesis, in this section only the prototypes developed in collaboration with undergraduate students are described.

UMCourt Consumer Law

Nuno Costa, an undergraduate student of the host institution, collaborated with this PhD work by implementing a prototype of a part of the defined architecture in the Consumer Law domain. This domain has a wide set of rules, defining the relationship between concepts such as producer, seller, consumer, products or services, warranty contract, defect, among others. In that sense, it is particularly useful for testing rule-based approaches. In that sense, the main objective in the implementation of this prototype was the assessment of the rule-based approaches presented in this PhD thesis, namely in what concerns the retrieval of information and the estimation of the possible outcomes. Figure 60 depicts one of the interfaces of the prototype, that allows to input all the information necessary to describe a conflict between a consumer and a producer or seller. Based on the information provided and on the rules transcribed from the Portuguese Consumer Law domain, the prototype will provide a likely outcome. In that sense, more than conflict resolution, this is a prototype of a decision support system, intended for an individual to be aware of his/her rights as a consumer, should something happen to the product bought, at a given time.

Additionally, a mobile interface for the prototype was also developed (Figure 61). This had as main objective to allow the prototype to be used while the consumer is shopping for a given product, allowing the consumer to be informed of his/her rights (for example by simulating scenarios). This is particularly useful when we are shop-

Figure 60: Screenshot from the interface developed for the UMCourt Consumer Law prototype.

Figure 61: Screenshot from the some of the mobile interfaces developed for the UMCourt Consumer Law prototype.

ping abroad and are unaware of our rights as consumers in foreign countries.

Online Dispute Resolution for Property Division - Divorces and Heritage share

Under this work a prototype was developed in collaboration with the undergraduate student Ana Café to test some of the notions presented in this PhD work, namely in the field of property division. Particularly, this prototype denominated *UMCourt Divider* focused on dividing purely monetary assets over a group of individuals, including notions of fairness and justice. Its implementation was the first test bench for the computation of concepts such as the BATNA, WATNA and MLATNA from legal norms, in an autonomous way.

Figure 62: Two screenshots from the interfaces developed for the UMCourt Divider prototype.

Conflict Resolution in Virtual Organizations

Virtual Organizations (VOs) have very specific characteristics, namely a temporary nature, flexibility and distribution. In that sense, conflicts and their resolution also present specific issues. Traditional conflict resolution mechanisms have their known disadvantages, which clash with the VO's main advantages. Thus being, in a first phase, the undergraduate student Flávio Lemos focused on determining why aren't current conflict resolution methods (alternative or traditional) not suited for the context of conflicts in a VO. Then, based on the UMCourt architecture, Flávio proposed a hybrid solution mixing a case-base and a rule-based approach. This solution accompanied the VO's lifecycle from the beginning to the end, watching the definition of the contract and later intervening when and if the norms on the contract were broken. Figure 63 depicts its flowchart.

9.8 SUMMARY

The process of obtaining a PhD degree is often regarded as one that is undertaken in solitude, placing the emphasis on individual achievement. After obtaining a doctorate degree, the individual is supposed to be a world-level expert in a very specific subject. Matt Might, assistant Professor at the University of Utah, depicts this view graphically as in Figure 64. The whole human knowledge can be represented as a circle. When you finish elementary school (represented in blue) you know a little, after high school (green circle) you know a little bit more, with a bachelor's degree (pink circle) you gain a speciality. In a master's degree that speciality is deepened. Finally, during the attaining of a PhD degree, the candidate reaches the boundary of human knowledge in a very specific topic and pushes that boundary a little bit forward.

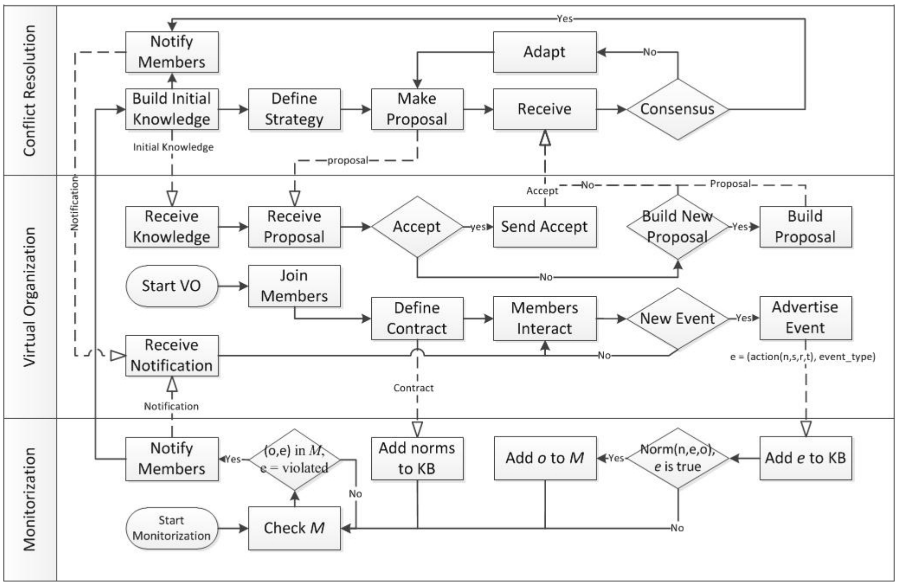

Figure 63: Flowchart denoting how the lifecycle of the organization integrates with the its monitoring and the eventual conflict resolution.

This is the traditional view on a PhD and is also what many of the PhD projects are still about. Not infrequently, candidates work in such a specific and advanced topic and do it in such an isolated manner that the results they achieve, however relevant they are, result in few practical advantages for this circle of knowledge that we all share.

In this PhD, a different approach was followed. As depicted in this chapter, emphasis was placed, obviously, on the individual achievements of the candidate, but it was also placed on establishing connecting points with the other elements of the human knowledge. The main objective in doing this is to maximize the utility that may accrue from the PhD work. In that sense, the candidate, with the support of the supervisors also focused significantly on:

- Participating on conferences and other scientific meetings so that the work could be disseminated, discussed and validated by scientific peers;
- Making part of organizing committees of scientific events, thus giving back to the scientific community;
- Giving lectures about the topic to undergraduate students of different fields, sharing the lessons learned;
- Supervising students, allowing them to go further by sharing the experience already accumulated;
- Cooperating with other researchers, in fields of intersection in order to mutually share knowledge;

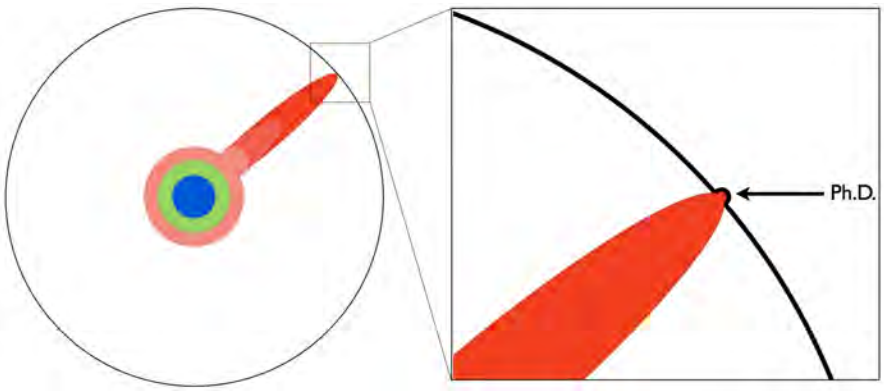

Figure 64: The traditional view on a PhD: the candidate works individually to push forward knowledge in a very specific subject. Source: http://matt.might.net/articles/phd-school-in-pictures/

- Working closely with other undergraduate students not necessarily being supervised.

Focusing on these items evidently takes its toll on the available time and, consequently, on the quantity and quality of the work that the candidate can perform. Thus, in a first view, the circle of Human knowledge will most likely not stretch as it could, given that the candidate is being hold back by meetings, classes, talks and supervisions. From an individualist point of view, this is evidently negative.

However, one should be able to look at the whole picture. The time spent in supporting students, in helping them into the right direction, in making them avoid your past mistakes, in sharing your knowledge and practical work, will result in a faster growth of their own circles of knowledge. This will prepare them better for their own endeavours. When, one day, they decide to enrol in their own personal adventure, they will be in a better start position to stretch their own circles and push the boundary of human knowledge, hence will do it faster and better.

And this should be the second goal of a PhD degree. The first one is, evidently, the individual achievements of the candidate. But the second one should be the maximization of the utility of the work done so that it is not inglorious or one more exercise of futility. The pursuit of this goal would result in a different view on the PhD, one followed by this candidate, and depicted in Figure 65: the candidate stretches his own circle of knowledge while fomenting and supporting the growth of his fellows' circles.

Given this, in this chapter the description of the quantitative results focused, in a first part, on the different scientific publications through which the work was successfully disseminated. In the remaining of the chapter, results were also evaluated in terms of the lectures and talks given and in terms of the students officially supervised or simply supported on their academic work.

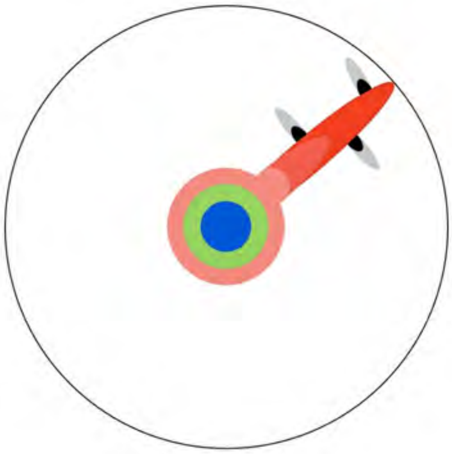

Figure 65: The candidate followed a different view on his PhD, focusing on sharing the lessons learned with the surrounding community, supporting his fellows in stretching their own circles of knowledge.

10 OUTTRO

Somewhere, something incredible is waiting to be known.

— Carl Sagan

Up until now, this manuscript described the motivation for the PhD work, its implementation details and the qualitative and qualitative results. By way of conclusion, in this last chapter an analysis of the accomplishments is performed, remembering the research hypothesis put forward in Section 3.2 and assessing their successful or unsuccessful validation. Then, the contribution to the current state of the art is briefly summarized. Finally, the conclusions achieved are summarized and some future lines are pointed out.

10.1 ANALYSIS OF ACCOMPLISHMENTS

In order to analyse the accomplishment of the research hypothesis of this thesis, the several research questions defined are first recalled:

1. Multi-agent systems are a suitable way to implement an Online Dispute Resolution platform;
2. Artificial Intelligence techniques can be used to develop information retrieval methods suited for the legal domain;
3. Artificial Intelligence techniques can be used to generate solutions for legal problems;
4. Artificial Intelligence techniques can be used to define strategies for conflict resolution;
5. The incorporation of contextual factors can improve the efficiency of legal practitioners working online;
6. Context information can improve the conflict resolution process in several ways:
   - 6.1 The conflict handling style of the parties can be assessed in a non-invasive way;
   - 6.2 Stress influences in a significant manner our interaction patterns with common technological devices;
   - 6.3 It is possible to accurately measure the influence of stress in a non-invasive and non-intrusive way, in real time.

Given this, a systematic question-by-question analysis is performed in which each of the research question's objectives is put forward and its validation is assessed by considering the validation conditions established in Section 3.2.

RESEARCH QUESTION 1: Concerning the first research question, the main objective was to determine if multi-agent systems were a suitable way to implement a domain-independent conflict resolution platform. It was stated that this hypothesis would be considered validated if an agent-based architecture could be defined after the mentioned requisites and implemented in the form of a prototype. Moreover, validation by scientific peers was also a requisite.

Concerning the first condition, the agent-based architecture defined in Chapter 4 was successfully implemented as a prototype in the Labour Law domain, implementing all the functionalities described. Moreover, parts of this architecture were implemented in other domains by undergraduate students, allowing to validate the domain-independent nature of the architecture. Concerning the validation by scientific peers, the architecture or specific aspects of its definition/implementation were described specifically in five publications, and partially in other publications that also touched other aspects of the PhD work. In general, the architecture was well received by the scientific community and its advantages understood and acknowledged. Given all this, this hypothesis is considered successfully validated.

RESEARCH QUESTION 2: Retrieving valid and useful information is a central and complex problem in conflict resolution, given the multi-value, multi-issue and multi-party nature of its problems. In that sense, the second research question aimed to identify Artificial Intelligence techniques that could be used to implement information retrieval methods in the framework of a domain-independent architecture. As a measure of validity, a prototype and validation by scientific peers was required.

Concerning the implementation of a prototype, several important functionalities related to information retrieval in the legal domain were implemented, namely to retrieve information such as the BATNA or the WATNA, past similar cases or possible valid solutions for given scenarios. In the process, both case-based, rule-based and hybrid approaches were assessed and implemented. Chapter Chapter 5 was devoted to de description of the prototypes implemented. Concerning the validation by scientific peers, 8 publications were specifically dedicated to describing the information retrieval methods implemented, one of them on the International Journal of Knowledge and Information Systems. Other publications focused on other issues of this work have also grasped this topic. In that sense, the approaches implemented and the reception by the scientific community validate this second hypothesis.

RESEARCH QUESTION 3: Devising a valid, fair and appealing solution for a conflict resolution process may be a challenge, mostly when the concept of fairness may be subjective. In that sense, there is the need to not only generate a solution that is valid but also to detail why such solution is good, fair, or even the best one possible given the circumstances. This was the aim of the third research question: to implement a prototype for generating solutions for conflict resolution problems using Artificial Intelligence techniques.

In light of the several AI-based approaches analysed, described in the introductory section, the choice was on the use of a nature-inspired approach, based on Genetic Algorithms. The main advantage of the prototype implemented is that it is independent of domain, in line with the remaining of the framework, and that it is not bounded by a number of known cases. It results in an approach that can search virtually the whole search space, limited only by the halting conditions. The execution of the algorithm results in a tree of solutions of growing specificity (thus utility for each party) that can then be used intuitively by the mediator to suggest solutions according to the circumstances of the case. In what concerns the validation of this approach, besides the successful implementation of the prototype, it must also be noted its publication in one conference proceeding and in the Neurocomputing journal. In that sense, the hypothesis is considered validated.

RESEARCH QUESTION 4: Online communication has emerged in the last decades as a thrilling way of bringing people closer together. However, such radical changes must always be analysed with care. In fact, on the one hand, technologies such as chats, telephones, forums, among many others allow us to communicate synchronously or asynchronously with virtually anyone in the world, hence the increased sense of proximity. However, in the other hand, such approach to communication is simplistic, leaving aside the contextual factors that are so important for successful communication processes, including our body language or the tone of our voice. This must be considered with even more care when such communication means are used in sensitive domains such as Online Dispute Resolution, in which an communication between the parties may jeopardize the whole process, ultimately putting personal relationship at risk.

In this research question the main contextual factors that have an influence in the success of the communication processes were identified, particularly the ones that are important in the context of conflict resolution. Then, a possible technological solution to address this problem was devised. This question would be considered valid if the legal community could acknowledge the anticipated advantages and accept that such approach would result, in the overall, in better decisions being taken. These ideas were put forward in three renown conferences/workshops in the field of Juris-Informatics (Jurix, Jurisin and GDN) and were received with great interest. In the overall, despite the controversy of the approach and the conservativeness of the legal field and its practitioners, the community accepted that there could be advantages in having access to contextual factors and that, in doing so, better communication and decision-making processes could be implemented.

RESEARCH QUESTION 5.1: The personal conflict handling style is seen as one of the major issues to consider in the evaluation of the factors that influence or constrain the conflict resolution process. Up until now, the only way of evaluating one's conflict handling style was to use questionnaires. This has several disadvantages, the most significant being that one individual frequently changes his conflict handling style upon changes in the circumstances, which cannot be foreseen when the questionnaire is first filled.

The main aim of this hypothesis was to determine if a different approach could be devised that would not rely on questionnaires. In that sense, in this work the estimation of the conflict handling style was implemented through the analysis of the interactions between the parties, during the negotiation process, particularly focusing on the evolution of the utility of the proposals' issues. This results in a dynamic approach, that provides in real-time an estimation of the conflict handling style of each party, thus detecting significant changes. A prototype was successfully implemented that transparently monitors a negotiation process and, with the information compiled initially by the UMCOURT platform, is able to classify the conflict handling styles. Moreover, this approach was very well received, particularly by the legal community, that acknowledge the valuable information that is compiled: it was published and presented in four conferences, two in the legal field and two in the computer sciences field, which validated both its technical implementation as well as its utility for the conflict resolution process.

RESEARCH QUESTION 5.2: In this research question it was investigated the influence that stress has on our interaction patterns with technological devices. In fact, stress influences virtually all our actions. The challenge here was to devise a non-invasive way of measuring it by analysing the behaviour of the individuals.

It was concluded that stress influences many interaction features that can be measured just by analysing how one individual interacts with the devices, namely the acceleration on the handheld devices, the amount of movement, the touch patterns, among others. In order to validate this research question, an experiment was conducted that collected data about the interaction patterns of the users. The data was collected in two different moments, with and without stressors. Statistical methods were used to analyse the significance of the differences between the two distributions of the data, for each feature studied and for each user. It was concluded that there is an actual influence of stress over these interaction patterns, although it naturally varies from user to user. This idea and its implementation has also been validated by the scientific community through the publication of the experiment and its results in the journal of Expert Systems with Applications and through its publication and discussion in several conferences.

RESEARCH QUESTION 5.3: More than only determining if stress has an influence on one's interaction patterns, in this research question the possibility of devising a real-time and non-invasive method for measuring this influence has been studied. In fact, the implementation of such idea would open the door for many interesting application scenarios, going much further than Online Dispute Resolution.

With the objective of validating this research question, an approach to estimate stress in real-time was thus devised, based on the personalized models built during the assessment of the previous research question. The approach consisted in analysing the interaction patterns in real-time, and comparing them with the already known patterns for each user. A measure of the similarity between the pattern being built and the two known patterns is translated into a level of "stress" or "calm". This approach was introduced in a publication in the international journal of Expert Systems with Applications and was published and presented later in more detail in three conferences, which welcomed the idea with interest, understanding its potential applications and its potential socio-economical impact. In that sense, this research question is considered validated.

Given the analysis of the research questions performed, it can be concluded that, in a general sense, all of them have been validated, some through statistical methods, others through the development of working prototypes and all by the scientific peers. On the one hand, this allowed to collect feedback in order to better steer the pursuit of the goals of the PhD work plan. On the other hand, it also served to testify both of the usefulness and applicability of the approaches as well as their scientific validity.

In that sense, the research hypothesis of this thesis is considered validated. Hence, the conclusion is that the notion of a conflict-resolution environment, focused on providing the right knowledge in the right moment, will improve decision-making, communication and, in general, the outcomes of the actions taken.

10.2 CONTRIBUTION TO THE STATE OF THE ART

This PhD work resulted in several contributions to the State of the Art, in both the fields of technology and the law. Several Artificial Intelligence methods were used to solve the most significant problems in conflict resolution identified, including the compilation of important information and the generation of solutions. This, when incorporated into conflict resolution tools, will result in better conflict resolution processes, with better informed parties taking better decisions.

From the technological point off view a step forward was also taken concerning the architectures of Online Dispute Resolution platforms. In fact, most of the existing platforms focus on very specific domains and functionalities. This hinders their wide use and the development of further solutions for other domains. In this thesis, an abstract architecture that can be specified to particular legal domains has been put forward. It is based on the identification of processes and concepts that have the same meaning independently of the domain, but are obviously implemented differently. The architecture has been implemented with several prototypes on different legal fields. Such approach is expected to improve the development process of conflict resolution platforms by increasing functionality reuse, widening the domains of application, and thus, the target users.

However, the most significant innovation of this PhD is a new vision on the Online Dispute Resolution problem, in which aspects linked to the emotional sphere of the individual are regarded as important. In fact, up until now, emotions and other subjective responses of ours were regarded as an obstacle to good decision, that, traditionally, should be taken in a cold and objective way, with emotions playing no role at all. However, as António Damásio puts it, emotions are indeed essential in our decision-making processes: they constitute shortcuts for rational thinking, providing meaningful information Damasio (2005).

In fact, emotions and other aspects such as stress or body language are fundamental in our daily social interactions, whether we want it or not. Hence, their lack in online communication processes constitutes a serious drawback, worsened when such processes are conducted to negotiate or mediate a conflict.

The idea that emotions, stress or body language is very important in conflict resolution, particularly the one taking place online, was thus put forward. More than that, a scientific approach was devised to measure the effects of stress on behavioural indicators, in a non-invasive way. This resulted in the concept of context-aware conflict resolution environments. Under this view, the parties still interact with a conflict resolution tool. However, they do it from within an environment that is transparently monitoring their state and providing such information to the conflict resolution tool, that can then take and support better decisions.

Such view results, sometimes, controversial, mostly to the legal community that is traditionally conservative and resistant to significant and sudden changes. However, this same conservative community also acknowledge high potential in these ideas and continues interested in the developments, although they not foresee their use in a near future. Such environments could, however, be implemented in more receptive domains, such as learning environments or workplaces, with benefits in aspects such as productivity, quality of work and quality of human relationships.

10.3 CONCLUSIONS

Given the different aspects addressed by this work, conclusions are organized in subsections, one for each of the main topics covered.

10.3.1 *On the use of Case-based Reasoning in the Legal Context*

The use of case-based reasoning for conflict resolution is a natural approach in common law contexts, although it's use in civil law systems also makes sense, as described previously. 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 the decision to study case-based approaches.

Traditionally, two main drawbacks are associated to case-based reasoning (Aamodt and Plaza, 1994).

On the one hand, case-based approaches may suffer from inefficiency, mostly when the size of the database grows. Concerning this subject, the decision was to develop different approaches (from pure case-based to hybrid ones) so that their performances could be tested in order to determine the more efficient ones. Moreover, it must also be acknowledged that conflict resolution is generally an asynchronous process, which takes place over several hours or even days. Thus being, given that the methods presented have relatively low times of execution and that the context of application does not mandatorily demand for higher performance, it can be concluded that they can be used in this context.

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 must be acknowledged as a 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, i.e., cases take place in a given legal framework and are not easily adapted to other domains as the norms are different. In the case-study implemented a database whose cases focused mostly on employee's rights and obligations was used (articles 129 and 128 of the Portuguese labour law, respectively), although some cases also addressed other articles. The case-study consisted of conflict resolution scenarios in the labour 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. It was concluded that disputes involving employee's rights and obligations had many more useful solutions proposed by the platform than disputes involving other issues, diminishing the success rate of the resolution of the second ones.

In that sense, the main conclusion drawn 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 case-based reasoning is supported by other tools that can generate solutions when a case-based approach alone is not enough. This possibility was tackled in this work by devising an approach using genetic algorithms. The main advantage of such approach is that it depends only on the rules of each specific domain, which are needed to ensure the validity of the solutions. Alternatively, it is also possible to rely on parties themselves to generate solutions, although this approach tends to fail when parties are unable or unwilling to do it.

10.3.2 *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 the experiments implemented it can be learned that these processes, by themselves, may not suffice. One of the first issues that emerged was a consequence of the fact that most of the participants from the informatics field had very little to no knowledge at all about Portuguese labour 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 participants generally had no idea about the possible outcomes for each side and could not realistically evaluate how good or bad a given solution was.

The 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, it was found to be 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 realism.

An important conclusion was thus that methods for compiling information should be carried out before the actual conflict resolution process starts, so that the parties can gain a realistic view about their conflict. Following this line of attack, parties converge more quickly to a solution that is realistic and in line with the solutions retrieved from the past similar cases. This increases the success rate of conflict resolution processes.

Another issue that was addressed was that the success rate of the process depended significantly on the attitude of the parties, i.e., when participants 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.

Evidently, the challenge here is still the one of accurately determining the conflict handling style of each party. In that sense, an automated and non-invasive method that is able to classify the conflict style of a party based on their behaviour during the process was developed (e.g. is the party proposing solutions?, are the solutions proposed selfish?, is the party simply answering positively or negatively?). Based on that, UMCourt is able to determine the conflict style in real time and to adapt or point out possible adaptations to any of the parties involved.

10.3.3 *On the relationship between Personal Conflict Handling Styles and Stress*

Concerning the interpretation of the conflict handling styles, the results achieved were consistent for the majority of the participants. They can be summarized as follows:

- Stressed participants take hastier decisions, taking less time to think them through;
- Stressed participants are more prone to change their behaviour and do it in more significant ways (we focused on the values of the proposals exchanged and on the conflict resolution style evidenced);
- Under a stressful environment, outcomes tend to be farther away from the optimum result;
- Under stress, participants tend to be more competitive.

This points out the need for seeking calm and harmonious environments for conflict resolution. As a consequence, it can be concluded that courtrooms are definitely not the ideal conflict resolution environment as these are highly competitive milieus in which parties forget each other's natural ambitions and focus on the maximization of the own gain. Alternative environments, focused on cooperative strategies, should be preferred.

Moreover, in order for conflict managers to improve their action, access to the context information mentioned should be provided. This would allow them to detect, in due time, an escalation on the level of conflict and prevent a degradation on the relationships. This context information, that is available in face-to-face settings, must also be considered in virtual settings so that conflict managers can increase the efficiency of their decisions by considering more complete information.

However, despite its apparent advantages, this approach may also encompass risks. One of the main concerns raised is related with the risk of people trying to control the system if they know how it works. On the one hand it is known that people have tried to cheat systems as much as they try to make them cheat-proof (e.g. parties in court will also try to manipulate decision-makers leading them into believing what they want). It is argued that following this approach, this kind of behaviours may be hindered. On the one hand, the process is transparent, i.e., parties will not be aware of how the process of compiling this information is implemented. On the other hand, even if parties try to cheat the system, it is difficult to fake expressions, gestures or other behavioural features since they are more reflexes than conscious behaviours.

Given this, it is concluded that this approach may encompass several interesting advantages for mediators, specifically for the ones operating in online environments, allowing them to take more informed decisions.

10.3.4 *On the assessment of Stress*

In this work some features and their relation with stress have been studied. One of the main objectives was to do it in a non-invasive way, so that the experiments performed would not influence the results. In that sense, the focus has been set on features whose data could be transparently acquired from the environment, without the user being explicitly aware of it. Specifically, it has been studied how stress influences the acceleration of the handheld device, the maximum and mean intensity of touch, the duration of the touch, the amount of movement and the cognitive performance. Essentially, this work allowed to define a way to measure how each user is affected by stress.

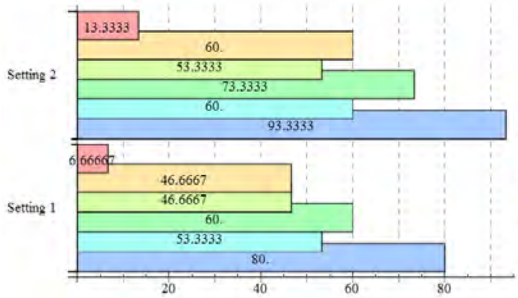

Parameter	Setting1	Setting2
Acceleration	80.	93.3333
Maximum Intensity	53.3333	60.
Mean Intensity	60.	73.3333
Duration of Touch	46.6667	53.3333
Amount of Movement	46.6667	60.
Score	6.66667	13.3333

Figure 66: Percentage of users that reveal significant differences between baseline and stressed data considering a first order analysis (Setting 1) and a second order analysis (Setting 2).

A general overview of its effects on the participants can be performed by looking at the data collected from all the users (Figure 66). It can be concluded that the feature most significantly affected by stress is the acceleration, with 80% of the users exhibiting statistically significant differences between baseline and stressed data.

On the other hand, the parameter that presents the less satisfying results is the score. However, it must be stressed that this last parameter is merely indicative and shall not be considered in eventual final applications, since the concept of a score may not even exist. It was however used in this study with the purpose of inducing an objective on the user: without the objective of maximizing the score, the user would not feel any stress.

It can also be concluded that, as expected, all the parameters show a better performance in a second order analysis. However, as already mentioned before, the first order analysis is preferred as it considers all the data (and not only data from more stressed states), which means that a statistically significant difference (when it exists) is more solid.

When considering all the data, each participant has in average 2.93(3) parameters with significant differences (out of 6). Moreover, in the worst case there are 3 participants with only 1 parameter with significant differences, and in the best case there are 4 users with 5 parameters with significant differences. Given this, it can be stated that a generic model can be applied to users for whom no training data exists. However, personalized models are certainly more accurate and that is the final aim.

The features most affected by stress are the ones that people do not conscientiously control

Moreover, from the results, it can also be concluded that the most affected features are the ones that people do not conscientiously control. Examples are the acceleration of our hands (i.e. how much they move) or the intensity of the touch. On the other hand, the parameters that are more rational (in this case the score) are not influenced in such a significant manner.

This is definitely positive in the sense that unconscious behaviours and reactions are more difficult to forge and usually represent true reactions of Human body. In that sense, the results that stem from their analysis are more solid and reliable.

10.4 FUTURE WORK

It was clear from the beginning of this PhD work that it would be a multidisciplinary one, drawing mostly from Computer Science but also largely from Legal Science. From the technological point of view, it also merged many different fields, including Multi-agent Systems, Behavioural Analysis, Service-oriented Architectures, Natural Computing, among others. In that sense, it is only natural that at the closure of this endeavour, a multitude of possible future paths exists, each one following a different line of research under the different fields addressed.

Several especially interesting research problems that could be followed after the conclusion of this thesis can be pointed out, each one worth at least another PhD plan:

- Pursuit of case-based models that can minimize certain drawbacks of this model that have been known for years, specifically the ones that matter in the legal field such as its high domain-dependence;
- Development of mechanisms that allow for an easier interpretation of norms and arguments written in natural language, its translation into rules and the verification of their correction;
- Pursuit of computational models for legal reasoning, with an ethical background, that can not only take decisions but also argue about them. This would be a step forward in accurately informing the parties;
- The inclusion of the notion of fairness into conflict resolution platforms. Although algorithms for fair division have been deemed since the 1940's, there is the need to not only use them but also to expose thoroughly how decisions are undertaken and how the notion of fairness is treated computationally;

Additional research lines could eventually be put forward. The ones listed here have in common the fact that they are already known topics, although still open research problems. This may make them less innovative research lines, although still valid ones with potential for success.

There is however a particularly innovative research line that stemmed from this PhD work: the non-invasive multi-modal acquisition of stress. The approach implemented opens the door to define personalized models to measure the influence of stress in the users of technological devices. It could lead to further developments with high social and economical impact.

Particularly, future work will undergo the development of full stress-aware environments that can model the level of stress of a group of human users that are interrelated in some way (e.g. team-mates, work colleagues). This will allow human resources specialists to better manage teams, with potential beneficial impact in the quality of the working environment and on the productivity.

The validity of this line of research, that will constitute the future work of the candidate, was recently recognized with the attribution of a funded research project by FCT, the Portuguese Science and Technology Foundation, with the reference PTDC/EEI-SII/1386/2012.

Part IV

BIBLIOGRAPHY

BIBLIOGRAPHY

Aamodt, A. and Plaza, E. (1994). Case-based reasoning; Foundational issues, methodological variations, and system approaches. *AI COMMUNICATIONS*, 7(1):39–59.

Aarts, E. and Grotenhuis, F. (2011). Ambient Intelligence 2.0: Towards Synergetic Prosperity. *Journal of Ambient Intelligence and Smart Environments*, 3(1):3–11.

Abraham, A. (2009). Editorial - Hybrid Soft Computing and Applications. *International Journal of Computational Intelligence and Applications*, 8(1).

Abrahams, B. and Zeleznikow, J. (2010). Including notions of fairness in development of an integrated multi-agent online dispute resolution environment. pages 102–111.

Aha, D. W., Kibler, D., and Albert, M. K. (1991). Instance-Based Learning Algorithms. *Machine Learning*, 6(1):37–66.

Aleven, V. (1997). *Teaching case-based argumentation through a model and examples*. PhD thesis, University of Pittsburgh.

Alsina-Jurnet, I. and Gutiérrez-Maldonado, J. (2010). Influence of personality and individual abilities on the sense of presence experienced in anxiety triggering virtual environments. *International Journal of Human-Computer Studies*, 68(10):788–801.

Andrade, F., Novais, P., Machado, J., and Neves, J. (2007). Contracting agents: legal personality and representation. *Artificial Intelligence and Law*, 15(4):357–373.

Ashley, K. D. (1991). *Modeling Legal Arguments: Reasoning with Cases and Hypotheticals (Artificial Intelligence and Legal Reasoning)*. The MIT Press.

Ashley, K. D. (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*.

Ashley, K. D. and Aleven, V. (1991). Toward an intelligent tutoring system for teaching law students to argue with cases. In *Proceedings of the third international conference on Artificial intelligence and law - ICAIL '91*, pages 42–52, New York, New York, USA. ACM Press.

Augusto, J. C. (2009). Past, Present and Future of Ambient Intelligence and Smart Environments. In Filipe, J., Fred, A., and Sharp, B., editors, *Proceedings of the International Conference on Agents and Artificial Intelligence*, pages 11–18. INSTICC Press.

Beale, R. and Creed, C. (2009). Affective interaction: How emotional agents affect users. *International Journal of Human-Computer Studies*, 67(9):755–776.

Beehr, T. A. and Newman, J. E. (1978). Job Stress, Employee Health, and Organizational Effectiveness: a Facet Analysis, Model, and Literature Review. *Personnel Psychology*, 31(4):665–699.

Beer, M., D'Inverno, M., Luck, M., Jennings, N., Preist, C., and Schroeder, M. (1999). Negotiation in multi-agent systems. *The Knowledge Engineering Review*, 14(03):285–289.

Bellifemine, F. L., Caire, G., and Greenwood, D. (2007). *Developing Multi-Agent Systems with JADE (Wiley Series in Agent Technology)*. Wiley.

Bellucci, E., Lodder, A., and Zeleznikow, J. (2004). Integrating artificial intelligence, argumentation and game theory to develop an online dispute resolution environment. In *16th IEEE International Conference on Tools with Artificial Intelligence*, number Ictai, pages 749–754. IEEE Comput. Soc.

Bellucci, E. and Zeleznikow, J. (2001). Representations of decision-making support in negotiation. *Journal of decision systems*, 10(3-4):449–479.

Benjamins, R., Selic, B., and Gangemi, A. (2005). *Law and the Semantic Web: Legal Ontologies, Methodologies, Legal Information Retrieval, and Applications (Lecture Notes in Computer Science / Lecture Notes in Artificial Intelligence)*. Springer.

Bennet, S. (2002). *Arbitration: Essential Concepts*. Incisive Media, LLC.

Berry, J. O. and Jones, W. H. (1995). The Parental Stress Scale: Initial Psychometric Evidence. *Journal of Social and Personal Relationships*, 12(3):463–472.

Bonczek, R. H., Holsapple, C., and Whinston, A. (1981). *Foundations of Decision Support Systems*. Academic Press Inc, 1st edition.

Brachman, R. and Levesque, H. (2004). *Knowledge Representation and Reasoning (The Morgan Kaufmann Series in Artificial Intelligence)*. Morgan Kaufmann.

Brams, S. J. and Taylor, A. D. (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

Brown, H. and Marriott, A. (2012). *ADR: Principles and Practice*. Sweet & Maxwell, 3rd edition.

Buchanan, B. and Headrick, T. (1970). *Some Speculation About Artificial Intelligence and Legal Reasoning*. Defense Technical Information Center.

Burton, J. (1990). *Conflict: Resolution and Provention*. St Martin's Press.

Caceres, E. (2008). EXPERTIUS: A Mexican Judicial Decision-Support System in the Field of Family Law. In *JURIX'08*, pages 78–87. IOS Press.

Campbell, D. (1997). Explorations into virtual architecture: a HIT Lab gallery. *IEEE Multimedia*, 4(1):74–76.

Castillo, J. C., Rivas-Casado, A., Fernández-Caballero, A., López, M. T., and Martínez-Tomás, R. (2011). A multisensory monitoring and interpretation framework based on the model-view-controller paradigm. pages 441–450.

Cook, D. J., Augusto, J. C., and Jakkula, V. R. (2009). Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and Mobile Computing*, 5(4):277–298.

Cooper, J. O., Heron, T. E., and Heward, W. L. (1987). *Applied Behavior Analysis*. Prentice Hall.

Corchado, E., Arroyo, A., and Tricio, V. (2010). Soft computing models to identify typical meteorological days. *Logic Journal of IGPL*, 19(2):373–383.

Corchado, E. and Herrero, A. (2011). Neural visualization of network traffic data for intrusion detection. *Applied Soft Computing*, 11(2):2042–2056.

Corcho, O. and Fernández-lópez, M. (2005). Building legal ontologies with methontology and webode.

Cox, T. (1985). The nature and measurement of stress. *Ergonomics*, 28(8):1155–63.

Damasio, A. (2005). *Descartes' Error: Emotion, Reason, and the Human Brain*. Penguin Books.

Davis, L. D. (1991). *Handbook Of Genetic Algorithms*. Van Nostrand Reinhold.

de L'Europe, C. (2010). CEPEJ rapport 2010 (donnees 2008).

DeMars, J., Exon, S. N., Kovach, K. K., and Rule, C. (2010). Ethical Considerations for an Online Dispute Resolution (ODR) Practice. *Dispute Resolution Magazine*, pages 6–10.

Deutsch, M. (1958). Trust and suspicion. *Journal of Conflict Resolution*, 2(4):265–279.

Dodds, T. J., Mohler, B. J., and Bülthoff, H. H. (2011). Talk to the Virtual Hands: Self-Animated Avatars Improve Communication in Head-Mounted Display Virtual Environments. *PLoS ONE*, 6(10):e25759.

Elder, G. H. and Caspi, A. (1988). Economic Stress in Lives: Developmental Perspectives. *Journal of Social Issues*, 44(4):25–45.

Evans, G. W., Bullinger, M., and Hygge, S. (1998). Chronic Noise Exposure and Physiological Response: A Prospective Study of Children Living Under Environmental Stress. *Psychological Science*, 9(1):75–77.

Fagin, L. (1985). Stress and unemployment. *Stress Medicine*, 1(1):27–36.

Fernández-Caballero, A., Castillo, J. C., Martínez-Cantos, J., and Martínez-Tomás, R. (2010). Optical flow or image subtraction in human detection from infrared camera on mobile robot. *Robotics and Autonomous Systems*, 58(12):1273–1281.

Forsyth, R. (1986). The Anatomy of Expert Systems. In *Artificial Intelligence: Principles and Applications*, chapter 8, pages 186–187. Chapman & Hall, London.

Goodman, J. (2003). The Pros and Cons of Online Dispute Resolution: an Assessment of Cyber-Mediation Websites. *ReviewDuke Law and Technology*, 4.

Greinke, A. (1994). Legal Expert Systems - A Humanistic Critique of Mechanical Legal Inference. *Murdoch University Electronic Journal of Law*, 1(4).

Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220.

Guasco, M. and Robinson, P. (2007). *Principles of Negotiation: Strategies, Tactics, Techniques to Reach Agreement (Entrepreneur Magazine's Legal Guide)*. Entrepreneur Press.

Haig, D. (2010). Colloquium papers: Transfers and transitions: parent-offspring conflict, genomic imprinting, and the evolution of human life history. *Proceedings of the National Academy of Sciences of the United States of America*, 107 Suppl:1731–5.

Harmon, P. and King, D. (1985). *Expert Systems: Artificial Intelligence in Business (General Trade)*. John Wiley & Sons.

Hayes-Roth, F., Waterman, D. A., and Lenat, D. B. (1983). Building expert systems.

Healey, J. and Picard, R. (2005). Detecting Stress During Real-World Driving Tasks Using Physiological Sensors. *IEEE Transactions on Intelligent Transportation Systems*, 6(2):156–166.

Herrman, M. S., Jerry Gale, N. H., and Foster, M. (2001). In Practice: Defining Mediator Knowledge and Skills. *Negotiation Journal*, 17(2):139–153.

Hogan, J. P. (1998). *Mind Matters*. Del Rey.

Holmes, G., Donkin, A., and Witten, I. (1994). WEKA: a machine learning workbench. In *Proceedings of ANZIIS '94 - Australian New Zealnd Intelligent Information Systems Conference*, pages 357–361. IEEE.

Hudlicka, E. (2003). To feel or not to feel: The role of affect in human-computer interaction. *International Journal of Human-Computer Studies*, 59(1-2):1–32.

Jackson, P. (1990). Introduction to Expert Systems.

Jaimes, A. and Sebe, N. (2007). Multimodal human-computer interaction: A survey. *Computer Vision and Image Understanding*, 108(1-2):116–134.

Jim Blascovich, J. L. (2002). Immersive Virtual Environment Technology as a Methodological Tool for Social Psychology. *Psychological Inquiry*, 13(2):103–124.

Katsh, E. and Rifkin, J. (2001). *Online Dispute Resolution: Resolving Conflicts in Cyberspace*. Jossey-Bass.

Kilmann, R. and Thomas, K. (1977). *Educational and Psychological Measurement*. XICOM.

Klaming, L., van Veenen, J., and 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 Poblet, M., editor, *ODR*. CEUR-WS.org.

Kolodner, J. (1992). An introduction to case-based reasoning. *Artificial Intelligence Review*, 6(3):3–34.

Kolodner, J. (1993). *Case-Based Reasoning (Morgan Kaufmann Series in Representation & Reasoning)*. Morgan Kaufmann.

Korfhage, R. (1997). *Information Storage and Retrieval*. John Wiley & Sons, 1st edition.

Kurzweil, R. (2000). *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. Penguin Books.

Landes, W. and Posner, R. (1976). Legal Precedent: A Theoretical and Empirical Analysis. *Journal of Law and Economics*, 19(2):309–313.

Leake, D. B. (1996). Case-Based Reasoning: Experiences, Lessons and Future Directions.

Lodder, A. (2006). The Third Party and Beyond. An analysis of the different parties, in particular The Fifth, involved in online dispute resolution. *Information & Communications Technology Law*, 15(2):143–155.

Lodder, A. and Thiessen, E. M. (2003). The role of artificial intelligence in online dispute resolution. In *Workshop on Online Dispute Resolution at the International Conference on Artificial Intelligence and Law*.

Lodder, A. R. and Zeleznikow, J. (2010). *Enhanced Dispute Resolution Through the Use of Information Technology*. Cambridge University Press.

Maria, B. A. D. and Silva, V. T. D. (2005). VisualAgent: A Software Development Environment for Multi-Agent Systems.

Marreiros, G., Santos, R., Ramos, C., and Neves, J. (2010). Context-Aware Emotion-Based Model for Group Decision Making. *IEEE Intelligent Systems*, 25(2):31–39.

Marucci, L., Mori, G., Paterno, F., and Costalli, F. (2001). Design Criteria for Usable Web-Accessible Virtual Environments.

Matthijssen, L. (1995). An intelligent interface for legal databases. *Proceedings of the 5th international conference on ...*, pages 71–80.

Matthijssen, L. (1999). *Interfacing Between Lawyers and Computers:An Architecture for Knowledge-Based Interfaces to Legal Databases (Law and Electronic Commerce)*. Kluwer Law International.

Mehrabian, A. (1980a). *Basic dimensions for a general psychological theory: Implications for personality, social, environmental, and developmental studies*. Oelgeschlager, Gunn & Hain (Cambridge).

Mehrabian, A. (1980b). *Silent Messages: Implicit Communication of Emotions and Attitudes*. Wadsworth Publishing Company.

Meyer, I. H. (2003). Prejudice, social stress, and mental health in lesbian, gay, and bisexual populations: Conceptual issues and research evidence. *Psychological Bulletin*, 129(5):674–697.

Munduate, L., Ganaza, J., Peiró, J. M., and Euwema, M. (1999). PATTERNS OF STYLES IN CONFLICT MANAGEMENT AND EFFECTIVENESS. *International Journal of Conflict Management*, 10(1):5–24.

Notini, J. (2005). Effective Alternatives Analysis In Mediation: "BATNA/WATNA" Analysis Demystified.

Olson, G., Malone, T. W., and Smith, J. B., editors (2001). *Coordination Theory and Collaboration Technology (Volume in the Computers, Cognition, and Work Series)*. Psychology Press.

Parunak, V. D. (1997). Go to the Ant: Engineering Principles from Natural Multi-Agent Systems.

Peruginelli, G. (2002). Artificial Intelligence in alternative dispute resolution. In *Proceedings of the Workshop on the law of electronic agents – LEA*.

Peruginelli, G. and Chiti, G. (2002). Artificial Intelligence in Online Dispute Resolution. In *Proceedings of the Workshop on the law of electronic agents – LEA*. Springer-Verlag.

Picard, R. W. (2000). *Affective Computing*. The MIT Press.

Platt, J. C. (1999). Fast training of support vector machines using sequential minimal optimization. pages 185–208.

Popple, J. (1990). Legal Expert Systems: The Inadequacy of a Rule-Based Approach.

Popple, J. (1996). *A Pragmatic Legal Expert System*. Dartmouth.

Preuveneers, D. and Novais, P. (2012). A survey of software engineering best practices for the development of smart applications in Ambient Intelligence. *Journal of Ambient Intelligence and Smart Environments*, 4(3):149–162.

Quinlan, J. R. (1993). C4.5: programs for machine learning.

Rahwan, I., editor (2009). *Argumentation in Artificial Intelligence*. Springer.

Raiffa, H. (1985). *The Art and Science of Negotiation*. Belknap Press of Harvard University Press.

Ramos, C., Augusto, J. C., and Shapiro, D. (2008). Ambient Intelligence—the Next Step for Artificial Intelligence. *IEEE Intelligent Systems*, 23(2):15–18.

Rehm, M., Bee, N., and André, E. (2008). Wave like an Egyptian: accelerometer based gesture recognition for culture specific interactions. pages 13–22.

Russell, S. and Norvig, P. (2009). *Artificial Intelligence: A Modern Approach (3rd Edition)*. Prentice Hall.

Salton, G., Wong, A., and Yang, C. S. (1975). A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620.

Searle, J. (1980). Minds, Brains and Programs. *Behavioral and Brain Sciences*, 3(3):417–457.

Sedano, J., Curiel, L., Corchado, E., de la Cal, E., and Villar, J. R. (2010). A soft computing method for detecting lifetime building thermal insulation failures. *Integrated Computer-Aided Engineering*, 17(2):103–115.

Selye, H. (1978). *The Stress of Life*. McGraw-Hill.

Silva, V. and Garcia, A. (2003). Taming Agents and Objects in Software Engineering.

Somekh, B. (2005). *Action Research: a methodology for change and development (Doing Qualitative Research in Educational Settings)*. Open University Press.

Sowa, J. F. (1999). *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks / Cole.

Soylu, A., Kavukcu, S., Turkmen, M., Cabuk, N., and Duman, M. (2000). Effect of socioeconomic status on the blood pressure in children living in a developing country. *Pediatrics International*, 42(1):37–42.

Span, G. (1993). LITES, an intelligent tutoring system for legal problem solving in the domain of Dutch Civil law. In *Proceedings of the fourth international conference on Artificial intelligence and law - ICAIL '93*, pages 76–81, New York, New York, USA. ACM Press.

Stockdale, S. E., Wells, K. B., Tang, L., Belin, T. R., Zhang, L., and Sherbourne, C. D. (2007). The importance of social context: neighborhood stressors, stress-buffering mechanisms, and alcohol, drug, and mental health disorders. *Social science & medicine (1982)*, 65(9):1867–81.

Susskind, R. E. (1988). *Expert Systems in Law: A Jurisprudential Inquiry*. Oxford University Press, USA.

Sycara, K. P. (1993). Machine learning for intelligent support of conflict resolution. *Decision Support Systems*, 10(2):121–136.

Tan, P.-N., Steinbach, M., and Kumar, V. (2005). *Introduction to Data Mining*. Addison-Wesley.

Thiessen, E. M. (1993). *ICANS: An Interactive Computer-Assisted Multiparty Negotiation Support System*. PhD thesis, Cornell University.

Tidd, S. T. and Friedman, R. A. (2002). Conflict Style and Coping with Role Conflict: an Extension of the Uncertainty Model of Work Stress. *International Journal of Conflict Management*, 13(3):236–257.

Turban, E. (1993). Decision Support and Expert Systems: Management Support Systems.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, 59:433–460.

Tyler, C. (2005). 115 and Counting: The State of ODR 2004. In Tyler, C., Katsh, E., and Daewon, C., editors, *Proceedings of the Third Annual Forum on Online Dispute Resolution*.

van Steenbergen, W. (2005). Rationalizing Dispute Resolution: From best alternative to the most likely one. In *Proceedings of the 3rd ODR workshop*, Brussels.

Velasquez, J. D. (1997). Modeling Emotions and Other Motivations in Synthetic Agents.

Visser, P. and Bench-Capon, T. (1998). A Comparison of Four Ontologies for the Design of Legal Knowledge Systems. *Artificial Intelligence and Law*, 6(1):27–57.

Vizer, L. M., Zhou, L., and Sears, A. (2009). Automated stress detection using keystroke and linguistic features: An exploratory study. *International Journal of Human-Computer Studies*, 67(10):870–886.

Vries, B. R. D., Leenes, R. E., and 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.

Walton, R. E. (1991). *A Behavioral Theory of Labor Negotiations: An Analysis of a Social Interaction System*. ILR Press.

Waterman, D. and Peterson, M. (1980). Rule-based models of legal expertise. In *Proceedings of the First National Conference on Artificial Intelligence*. Stanford University.

Watson, I. (1997). *Applying Case-Based Reasoning: Techniques for Enterprise Systems (The Morgan Kaufmann Series in Artificial Intelligence)*. Morgan Kaufmann.

Watson, I. and Marir, F. (1994). Case-Based Reasoning: A Review. *The Knowledge Engineering Review*, 9(4):355–381.

Wenhui Liao A, W. Z. A. (2006). Toward a decision-theoretic framework for affect recognition and user assistance. *International Journal of Human-Computer Studies*, 64(9):847–873.

Wooldridge, M. (2002). *An Introduction to MultiAgent Systems*. John Wiley & Sons.

Wooldridge, M. and Jennings, N. R. (1995). Intelligent Agents: Theory and Practice.

Zeleznikow, J. and Bellucci, E. (2003). Family Winner: Integrating Game Theory and Heuristics to Provide Negotiation Support.

Zeleznikow, J. and Bellucci, E. (2004). Building negotiation decision support systems by integrating game theory and heuristics.

Zeleznikow, J. and Hunter, D. (1994). *Building intelligent legal information systems: representation and reasoning in law*. Kluwer Law and Taxation Publishers, Deventer Boston.

Zeleznikow, J. and Stranieri, A. (1995). The split-up system. In *Proceedings of the fifth international conference on Artificial intelligence and law - ICAIL '95*, pages 185–194, New York, New York, USA. ACM Press.

Zweigert, K. and Koetz, H. (1998). *An Introduction to Comparative Law*. Oxford University Press, USA.

Part V

APPENDIX

A

APPENDIX A

A.1 LONG TABLES

	LITIGATION		ALTERNATIVE METHODS
√	Can/must be used when all else fails		Often parties do not commit since they can always use litigation
√	Allows appeals		Some methods, such as arbitration, may be binding, without possibility of appeal
	Is a public process	√	Alternative methods are private
	Takes many time, generally months or years; depends on the amount of pending processes on the court	√	Depending on the willingness of the parties, may take only a few days
	Parties and witnesses are involved emotionally and are led to take opposing places	√	Parties are led to cooperate in a constructive environment
	Is an expensive process	√	Is generally cheaper
	Is mentally and emotionally wearing	√	Tries to maintain a positive environment, focused on improving and maintaining relationships

Table 7: A summary of the main differences between litigation and alternative methods.

	RES. MODES								
	I	S	T	C	N	M	A	CORE	USABILITY
ANDRC	⋆	⋆	⋆	⋆	⋆	⋆	⋆	HE	-
ARyME	⋆					⋆	⋆	WP	-
BBB	⋆		⋆	⋆		⋆	⋆	HE	-
Peruvian Cybertribunal		⋆		⋆			⋆	WP, HE	Web forms
CPR	⋆		⋆			⋆	⋆	WP, HE	-
CyberLaws								WP	-
Electronic Courthouse		⋆	⋆			⋆	⋆	WP, information management tools, virtual workrooms, settlement conferences	Web forms, Automatic update of information, notifications
Eucon	⋆	⋆				⋆		HE	-
ICC		⋆				⋆		HE	-
JAMS		⋆				⋆	⋆	HE	-
Mediation Now	⋆	⋆				⋆		WP	-
Mediation Room		⋆	⋆			⋆		WP, HE, single platform for unified information management	Web forms, privileged access to information according to roles, tailoring of sites to specific clients
Forrest Mosten	⋆	⋆	⋆			⋆		HE	-

Continued on next page

Table 8 – *Continued from previous page*

	I	S	T	C	N	M	A	CORE	USABILITY
ODRWorld		⋆			⋆	⋆	⋆	WP, HE	Web forms, Message board, online chat
Private Judge		⋆				⋆	⋆	HE	-
Resolution F.		⋆	⋆			⋆	⋆	WP, HE, virtual break-out rooms	Simple text-based conferencing system, session transcript
Settle the Case		⋆				⋆	⋆	WP, HE	-
SmartSettle	⋆	⋆	⋆		⋆			Flash WP, optimization algorithms, unified case manager	Intuitive interface, quantitative and qualitative analysis tools
TRUSTe	⋆	⋆				⋆		WP	Web forms
WIPO		⋆				⋆	⋆	HE	E-mail filing
Camera Milano	⋆	⋆				⋆	⋆	HE	E-mail filing
OAN		⋆					⋆	HE	-
V. Courthouse		⋆				⋆	⋆	HE, Flash WP	Intuitive interface, web forms
Arbitrare		⋆					⋆	HE, Flash WP	Intuitive interface, web forms

Table 8: A summary of some major ODR service providers and resources.

Continued on next page

I: A web-site that provides valuable information about ODR.
S: A web-site that directly or indirectly (e.g. by containing a list of mediators or arbitrators) provides ODR services.
T: A web-site that provides training for either parties or neutrals on the field of ODR.
Resolution Modes: The several resolution modes that a web site deals with, either in the form of the actual provision of services or simply the categories that the available information refers to.
C: Conciliation - the neutral meets with the parties separately, hears them, and tries to solve their differences.
N: Negotiation - the parties dialogue and expose their interests in order to achieve an agreement.
M: Mediation - the neutral tries to guide the discussion in a way that optimizes the parties needs.
A: Arbitration - the neutral meets the parties and produces an Award that may or may not be enforceable.
HE: systems that rely significantly on Human Experts including Solicitors, Arbitrators, Attorneys, Judges, among others.
WP: Systems in which the web page plays a central role.
Core: the elements under this category describe the core of the system, i.e., the key resource or resources that support the system.
Usability: this category describes the main points of each system regarding the usability, i.e., what are the tools or approaches available for the users to interact with the system.

TECHNOLOGY	MAJOR FEATURES
DSS	Compile and provide useful information Provide support for decision processes Propose actions based on the analysis of facts
ES	Model human knowledge and inference mechanisms Reason similarly to human experts Automation of "simple" tasks by applying an inference engine to knowledge
KBS	Model complex knowledge Represent norms and judgement under uncertainty
II	Build a layer of abstraction for complex systems Faster, intuitive and more efficient access to information
CBR	Reasoning processes similar to the legal ones Contextualized retrieval of information Information is organized according to meaningful attributes
MAS	Distributed problem solving Implement negotiation protocols Support for argumentation
Ontologies	Representation of legal knowledge Inference Pattern extraction
RbS	Encode knowledge, expertise and processes of human experts Fairly simple way of interpreting and reasoning with rules

Table 9: The most interesting features of several sub-fields of Artificial Intelligence from the point of view of conflict resolution.

A.2 IMAGES

Gantt Chart of the research plan

ID	Task Name	Start	Finish	Duration
1	Study of the SoA	01-04-2010	30-09-2010	26,14w
2	Design of the Architecture	01-10-2010	31-03-2011	26w
3	Information Retrieval	01-04-2011	31-07-2011	17,43w
4	Generation of Solutions	01-08-2011	30-11-2011	17,43w
5	Generation of Strategies	01-12-2011	01-04-2012	17,43w
6	Study of Context Information	01-08-2011	30-03-2012	34,71w
7	Behavioural Analysis	01-04-2012	30-09-2012	26,14w
8	Analysis of Results	01-10-2012	30-03-2013	25,86w
9	Dissemination of Results	01-10-2010	30-03-2013	130,29w

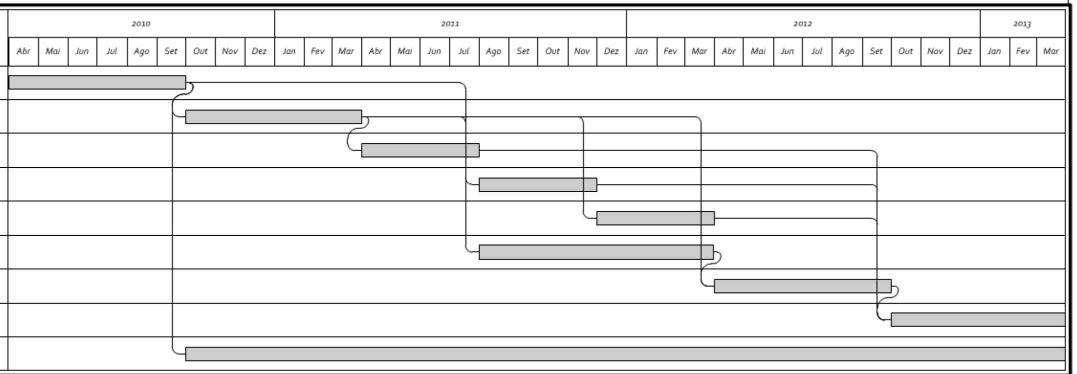

Figure 67: Gantt chart of the research plan detailing the main tasks that were identified as well as their relationships.

B

RESEARCH AND COMMERCIAL ODR PROJECTS

B.1 RESEARCH PROJECTS

In this section a few projects being developed by the research community are described. Most of the projects are maintained by universities and other institutions and focus on particular questions and not on the whole dispute resolution problem.

GENEVA LAW SCHOOL Geneva University has a Research Team on Online Dispute Resolution, Arbitration, and Information Technology. One of the projects developed in this group conducted research on the fields of Online Dispute Resolution on a first stage and on IT on a second phase. This project was financed by the Swiss National Science Foundation and ended in December 2005. This research group continues to conduct research on these topics mainly driven by the executive director Dr Thomas Schultz who did his academic career in the field of ODR. The results can be analyzed in many published articles and books.
http://www.mids.ch/school/university.html

LAUTERPACHT CENTRE The Centre is part of the Faculty of Law in the University of Cambridge and one of the Faculty's specialist law centres. The Centre is the scholarly home of international law at Cambridge University. International law is a major aspect of the Law Faculty's teaching programme at undergraduate and research level. The Centre's objectives are to promote the development of international law through research and publication, to serve as a forum for the discussion of current events and issues in international law and to provide an intellectual home in Cambridge for scholars of international law from around the world to pursue their research in a stimulating and congenial atmosphere.
http://www.lcil.cam.ac.uk

MERC The mission of the McMaster eBusiness Research Centre is to provide leadership and infrastructure support for eBusiness research to academic and industry partners. To accomplish this, MeRC focuses on three main activities: research, education and outreach. Research is looked as a way to identify and undertake important areas of cross-disciplinary eBusiness research, and secure the resources to support them. By fomenting education it provides eBusiness training and knowledge to current and future business and academic leaders.

Finally, MeRC aims to provide an interface and enhance the dialogue between the business community and academia. For achieving these goals, MeRC's activities are supported by a number of sources: private donations, e-business scholarships and bursaries, government grants, donations from the private sector, research contracts, commercialization of intellectual property and memberships in the MeRC research consortium. There is currently a large number of research projects being held at MeRC that cover a wide range of disciplines, among them mobile commerce, identity theft, online trust, online negotiation, among others. In these projects better and more secure ways of doing online transactions are investigated as well as the dispute resolution methods to apply when these fail. This research project is located in Ontario, Canada at MacMaster University.
http://www.merc-mcmaster.ca

VICTORIA UNIVERSITY The Faculty of Law of the Victoria University conducts research in the areas of business and law with special focus on ODR procedures, under the supervision of Prof. John Zeleznikow, who is responsible for the Family_Winner project. Some key areas of research include developing software tools to support negotiated decision making, building ethical standards into corporate governance, assessing the economic impacts of climate change, studying the impact of WorkChoices legislation on Victorian workers or analysing the social impact of tourism. Research is undertaken by academic staff members, research students and the faculty's three Research Centres: Centre for International Corporate Governance Research (CICGR), Centre for Strategic Economic Studies (CSES) and Centre for Tourism and Services Research (CTSR).
http://www.businessandlaw.vu.edu.au/index.asp

NATIONAL CENTER FOR TECHNOLOGY AND DISPUTE RESOLUTION This center is located at the University of Massachusetts at Amherst and provides an interdisciplinary approach to the study of law and society. Research efforts of faculty have included the impact of new information technologies on law, alternative dispute resolution, law and multinational corporations, law and popular culture, law and education, law and indigenous peoples, the legal profession, and law and education. The Department organizes a clinical project in conjunction with the Massachusetts Fair Housing Center (MFHC) and it also sponsors the National Center for Technology and Dispute Resolution. One of the main figures behind ODR investigation in this university is Professor Ethan Katsh who besides being professor of legal studies in this university is also Director of the National Center for Technology and Dispute Resolution. Among all its work in this field he has also published three books all related with the new role that technology has in law, namely in the dispute resolution topic.

http://www.odr.info

TIARAC TIARAC - Telematics and Artificial Intelligence in Alternative Conflict Resolution constitutes the first Portuguese experience in the field of Online Dispute Resolution. This project, located at the University of Minho, intends to develop tools to help parties involved in legal disputes. It currently addresses the specific domains of the Portuguese labor, consumer and property division law. The main objective of this project is to study the role that artifacts from the Artificial Intelligence field can play on the enhancement of the capabilities of ODR platforms. Specifically, this project aims at the development of tools that are proactive and autonomous and can effectively generate useful strategies and solutions for current problems. The main objective is that these tools are used as a path to avoid litigation in court.
http://tiaracserver.di.uminho.pt/tiarac/

B.2 COMERCIAL PROJECTS

In this section a number of commercially available ODR providers in the most different fields of law is described.

ADNDRC the Asian Domain Name Dispute Resolution Centre is a joint undertaking by the China International Economic and Trade Arbitration Commission (CIETAC), the Hong Kong International Arbitration Centre (HKIAC) and the Korean Internet Address Dispute Resolution Committee (KIDRC) that is credited by the ICANN as a dispute resolution provider. Disputes may be solved by offices located in Hong Kong, Beijing and Seoul.
http://www.adndrc.org

ICANN The Internet Corporation for Assigned Names and Numbers (ICANN) is responsible for managing and coordinating the Domain Name System (DNS), a distributed hierarchical naming system for computers, services, or any resource connected to the Internet or a private network. ICANN is also responsible for accrediting the domain name registrars. In that sense, ICANN's dispute resolution focus is on subjects such as domain transfer, unsolicited renewals, accreditation, trademark infringement, among others.
http://www.icann.org/en/dispute-resolution/

ARYME ARbitration & MEdiation was established in 1996 and is a private institution dedicated solely to providing specialized information on international ADR, and to promoting a better understanding and use of alternative dispute resolution mechanisms. It is not a case administrator as it only acts as a source of information and a meeting

point for specialists, international business and legal community and general public that can share ODR services throughout the world.
http://www.adrresources.com

BETTER BUSINESS BUREAU BBB is a service that certifies e-Commerce companies. Companies that are members of the Bureau pledge to meet the BBBOnLine Seal standards for ethical online business practices and have agreed to solve complaints using the BBB's dispute resolution program or a similar one. Currently this company certifies more than 50000 web sites and offers a search engine for consumers that helps to find certified companies.
http://www.bbb.org

PERUVIAN CIBERTRIBUNAL The Peruvian Cibertribunal is a non-profit, non-governmental organization created on 1999. It is an on-line and offline alternative resolution centre of the Minister of Justice which encourages conciliation among the parties and arbitration as an alternative to litigation in courts. The disputes that can be solved using this online tool comprise electronic commerce or agreements, intellectual property, publicity and marketing on the internet, consumer protection or intimacy protection.
http://www.cibertribunalperuano.org/ingles_prin.htm

CPR the international institute for Conflict Prevention and Resolution is a membership-based non-profit organization located in New York City whose mission is to spearhead innovation and promote excellence in public and private dispute resolution. To fulfil this mission, CPR is engaged in an integrated agenda of research and development, education and advocacy. CPR defends that companies and law firms should consider alternatives to litigation when disputes arise. It deals with more than 4,000 corporations - representing more than two-thirds of the U.S gross national product – and more than 1,500 law firms, including 400 of the nation's 500 largest firms. CPR does not create judicially enforceable rights or obligations, nor does it constitute a waiver of any substantive or procedural right or obligation. It rather aims at encouraging greater use of flexible, creative and constructive dispute resolution approaches.
http://www.cpradr.org

CYBERLAWS.NET this company provides consultancy services related to internet law. Through its services, consumers can get assistance on topics like E-commerce, domain names, online banking, privacy, telemedicine, spamming, hacking, intellectual property rights, among others. This site does not provide services for solving disputes. It is a site that provides information that may be used by parties in

dispute or by the neutrals, in this specific case to help solve a dispute related with internet law.
http://www.cyberlaws.net/new/index.php

ELECTRONIC COURTHOUSE this is an American dispute resolution service provider created in 2000 that offers services of mediation, arbitration and evaluation of cases. It counts with the participation of major law firms in North America comprising over 2500 corporate, commercial, trade, and business lawyers. Evaluation is used when the parties are not aware of the laws and are simply looking for some legal framework, being therefore a non-binding and merely informative process. The company also offers a mixed service which starts with a mediation process that may evolve into a binding arbitration process if the first one fails. Customers of this service include employees, unions, professional associations, individuals, sole proprietors, governments, public agencies, businesses and enterprises.
http://www.electroniccourthouse.com

EUCON Europäisches Institut für Conflict Management e.V., (European Institute for Conflict Management) headquartered in Munich is a non-profit organisation and was formed in October 2006 with the objective of advising and assisting parties in solving their disputes. This company provides mediation services that may be used by customers such as institutions, professional organizations, law firms and corporate business. The eucon site also offers a comprehensive explanation about mediation and its advantages, including a film entitled "Die Schöne in der Ostsee" ("The Beautiful Woman in the Baltic") intended to provide a practical case study showing the advantages of mediation and illustrating how a case proceeds.
http://www.eucon-institute.com

HKIAC Hong Kong International Arbitration Centre was established in 1985 to assist disputing parties to solve their disputes by arbitration and by other means of dispute resolution. It was established by a group of the leading business and professional people in Hong Kong to be the focus of dispute resolution in Asia. Besides arbitration, parties can chose alternative ways of solving their disputes such as negotiation, conciliation, mediation and finally litigation. The company also has a free service for providing information about the alternative dispute resolution methods. The Centre maintains a growing information services centre of books and publications which are available for reference to interested members of the public.
http://www.hkiac.org

ICC the International Chamber of Commerce is a global business organization whose activities cover a broad spectrum, ranging from

arbitration and dispute resolution to making the case for open trade and the market economy system, business self-regulation, fighting corruption or commercial crime. The arbitration services provided are on increasing use, having received cases at a rate of more than 500 a year since 1999.
http://www.iccwbo.org

JAMS JAMS provides a wide range of ODR services that include litigation, direct negotiation and several different types of mediation such as facilitative, evaluative, mini-trial or evaluation to customers all over the world. It has a service for mediator and arbitrator search by name, by location or by fields of expertise. The processes are all mediated by neutrals that are retired judges or attorneys and are experts in their areas. JAMS has been providing these services for almost 30 years and has strategic alliances with other similar companies around the world such as HKIAC.
http://www.jamsadr.com

MARS Mediation Arbitration Resolution Services has the conviction that every dispute, no matter how trivial or insignificant may first appear, has the right to be discussed and solved with the help of specialists. This company is specialized in e-commerce and virtual disputes and looks at new technologies for providing tools that can be used for designing new ODR mechanisms. MARS uses the concept of Virtual ADR and provides mediators, arbitrators, attorneys and other legal practitioners the opportunity to carry on mediation and arbitration conferences using the latest state-of-the-art videoconferencing to create a unique environment, very similar to traditional conferences, saving time, travel expenses and reducing costs. MARS also offers, through its Online Dispute Resolution program, the "Shop with Confidence" Seal (Trustmark) Program which provides merchants and consumers the opportunity to solve disputes which may be caused by misunderstandings from online transactions. This program offers the parties, access to a third party dispute resolution process, with an easy to use, inexpensive, automated, online web-based interface.
http://www.resolvemydispute.com

MEDIATIONNOW.COM the MediationNow.com site contains a list of mediators spread all over US and overseas that offer services of mediation. Through this online resource it is possible to search for a mediator that is close to a given location and access its description and contact information so that the mediation process may begin. The site also features a repository of useful information on alternative dispute resolution, including links to other providers.
http://www.mediationnow.com

THE MEDIATION ROOM this company provides a virtual mediation space for parties trying to solve their disputes. Users of TheMediationRoom.com system include The Ministry of Justice (UK), The National Institutes of Health (USA), The Law Council of Australia (Australia), The National Mediation Board (USA), The Commonwealth Telecommunications Organisation, The European Consumer Centres and eBay/PayPal (Europe and Australia). Each of these entities has very different disputes but the software, by being very dynamic, may be used by all of them. When a client authenticates in the web site, he has access to all the information gathered so far, including previous conversations, the current state of the process, information about the mediator, and in a later phase, about the outcome. The web site can also be used by mediators which have access to privileged information, organized according to their perspective.
http://mediation.orcawebsites.com

FORREST WOODY MOSTEN this is the website of a certified American family law specialist who provides mediation services mainly for divorce cases including the ones with parental disputes. His web site does not provide any tool for solving disputes but it does provide a lot of resources on family law such as articles, books and videos that disputants in this scenario may consult. Besides family law, this mediator also offers its services for partnerships, employment, probate, real estate, and commercial disputes. This attorney also offers training and coaching services.
http://www.mostenmediation.com

ODRWORLD the lemma of this company is that "justice is paramount and should be available to all". It therefore provides ODR services for solving any type of dispute, ranging from the simple and trivial to higher value ones. Disputes can be solved online and offline through assisted negotiation, mediation and arbitration. ODRWorld focuses on disputes that emerge from the global online activities and defends that these disputes should be solved online. The site includes tools such as a case message board and an online chat.
http://www.odrworld.com

PRIVATE JUDGE Private Judge is a site run by Judge Sullivan who has retired and now offers ADR services. In this site one will not find any tool for solving a dispute but will instead find the contact information of this ADR provider as well as extensive information about different types of ADR and their advantages and some example cases. This judge offers arbitration and mediation services.
http://www.privatejudge.com

RESOLUTION FORUM, INC. The Resolution Forum is a non-profit educational organization founded in close association with the Center for Legal Responsibility at South Texas College of Law. Its mission is to improve the quality and efficiency of dispute resolution services by making On-Line ADR more accessible and affordable to the general public, providing businesses and institutions with training and guidance about the effective use of ADR processes, encouraging the use of ADR to resolve business-related disputes, offering corporate volunteers the opportunity to provide On-Line ADR to the public, particularly to persons and organizations of limited means. Resolution Forum has developed the CAN-WINSM system which includes a conferencing system for conducting ADR over the Internet. This conferencing system is tailored for the Mediation or Arbitration process and allows communication among parties located anywhere in the world using standard browser software. Currently the system is text-based but currently a video-conferencing system is being implemented. The system saves transcripts of the conversations that are later available for the parties and ensures the confidentiality of the data of each party. This company also provides training services for individuals intending to perform the tasks of a neutral party.
http://www.resolutionforum.org

SETTLETHECASE.COM this site helps design, refine and rehearse the best strategies for ADR. It is used for arbitration, mediation and summary jury trials. It has as objective to accomplish this important goal in a cost effective and time efficient manner. With SettleTheCase it is possible to test ADR strategies in an attempt to assure the most favorable outcome for disputes. It is used by Attorneys, Mediators, Corporations, Government and Individual Clients. It provides therefore services for the parties or their legal representatives.
http://www.settlethecase.com

SMARTSETTLE Smartsettle is an online negotiation system (eNegotiation) that can be described as a generic tool for decision-makers with conflicting objectives that wish to reach a formal agreement. This platform can be used to solve problems relating to family, insurance, real estate, labor-management, contract negotiations, among others. In the site it is possible to find a few simulations that explain and show how the platform works. The Smartsettle suite is organized into SmartsettleOne which deals with simple and single-issue disputes to Smartsettle Infinity which deals with complex and multivariate cases.
http://www.smartsettle.com

TRUSTE TRUSTe provides privacy services that benefit both companies and individuals. Trough the TRUSTe Seals the company certifies more than 3500 web sites and this way provides privacy protection

and confidentiality to millions of ecommerce clients. If a client feels that a TRUSTe seal holder has violated a privacy agreement, a complaint can be issued and the company will activate its ODR mechanism to mediate the dispute in search for a fair outcome. The issues dealt with may comprise spam, identity theft or any other type of privacy violation.
http://www.truste.org

WIPO the World Intellectual Property Organization Arbitration and Mediation Center was established in 1994 to offer Alternative Dispute Resolution (ADR) options, in particular arbitration and mediation, for the resolution of international commercial disputes between private parties. Developed by leading experts in cross-border dispute settlement, the procedures offered by the Center are widely recognized as particularly appropriate for technology, entertainment and other disputes involving intellectual property. Among the services offered are arbitration, mediation and expert determination (submitting a dispute to the non-binding evaluation of a neutral).
http://www.wipo.int

CAMERA ARBITRALE DE MILANO the Chamber of Arbitration of Milan is a special branch of the Chamber of Commerce of Milan and specializes in commercial dispute resolution. The Chamber provides an array of services and tools which allow for a resolution of disputes within time-limits and through methods that are different and specific to the type of dispute and resolution that is needed. These include arbitration or mediation, which can be handled online as well. The Research Centre for ADR also provides material and publications for those who wish to develop their knowledge in the field of alternative justice.
http://www.camera-arbitrale.it

ONLINE ARBITRATION NETWORK OAN offers a protocol in an electronic format on the internet. The OAN protocol is claimed to be more accessible, economical, and faster than the traditional method of arbitration. It is also described as being considerably less formal than going to Court and completely private and confidential. The arbitration process is governed by Federal and State laws and is conducted by attorneys that are subject to Professional Ethics and Standards of Conduct described in the web site. Being an arbitration process, the decision is as binding on the parties as a Court Judgment and can be enforced by the Courts, if necessary.
http://oanlive.com/

VIRTUAL COURTHOUSE VirtualCourthouse.co (VCH) is an Internet-based service that enables parties to submit disputed claims,

responses and supporting material in digital form for resolution by a neutral provider. It is a combination of multimedia technologies and business processes aimed at replicating the processes of dispute resolution in online environments, removing constraints of time, expense and distance. It also acts as a platform in which neutrals can publish their services.
http://www.virtualcourthouse.com

ARBITRARE ARBITRARE is an institutionalized arbitration center that focuses on disputes over industrial property (trademarks and patents), Portuguese domain names, trade names and corporate names. These legal disputes can occur between private individual sor between private individuals and the bodies legally able to grant or refuse registrations. These include the National Institute of Industrial Property, the Foundation for National Scientific Computing and the Institute of Registration and Notary Affairs.
https://www.arbitrare.pt

C MULTI-AGENT SYSTEMS

—TODO : Given the scope of the thesis, this section details MAS—

The development of ODR tools that might act "as an autonomous agent" (Peruginelli, 2002) is indeed an appealing way for solving disputes. Such tools imply that agents are able of reading their environment (which comprises the parties, the problem domain and characteristics, the norms and other parameters). Agents also need to have enhanced communication skills that allow them to exchange complex knowledge with both parties. Thus, agents need a knowledge representation mechanism able of storing the data gathered during all the phases of the process (which may include data about norms addressed, problem domain, items in dispute, among others). Agents also need advanced cognitive skills for dealing with this information and eventually infer conclusions and propose strategies and advice for the parties.

Additionally, agents are a tool suited for addressing some of the new challenges that the legal field is facing. Namely there is a need for distributed web-based solutions that allow for parties and legal practitioners to access information everywhere seamlessly. Moreover, legal systems also have an urgent need for tools that can automate processes and replicate human cognitive processes in order to fasten the processes. As Multi-agent Systems can be used to address all of these challenges, they are paid special attention in this section. Specifically, we will analyse the main agent specification languages, development tools and methodologies, as well as the most used agent platforms.

C.1 AGENT SPECIFICATION LANGUAGES

Agent-based technologies are increasingly occupying a major position as viable solutions for deploying large-scale applications. In that sense, it is necessary to ensure that these systems that often play major and key roles are robust and reliable. Until recently, agent-based development was done in an ad-hoc fashion, following no particular specification and design methodologies. This lack of rigour has in fact delayed the adoption of agent-based solutions for wide-scale reliable applications. However, numerous methodologies and design principles now exist that are changing this situation. Generally, research conducted in this field can be divided into two main trends: one in which logic is used and another in which tools from software and knowledge engineering are adapted. However, we can also look at the languages from the point of view of the systems that they define, i.e., according to the types of agents and their relations. This is done in this section.

C.1.0.1 *BDI Languages*

2APL 2APL is an agent-oriented programming language aimed at that facilitating the implementation of multi-agent systems. Specifically, it provides programming constructs to define a multi-agent system in terms of a group of individual agents, a set of environments in which they can operate, and the relationship between the agents and the environments. At the individual agent level, the language supports the creation of cognitive agents based on the BDI architecture. It therefore provides programming constructs to implement agent's beliefs, goals, plans, actions, events, and a set of rules through which the agent can decide which actions to perform. With 2APL both reactive and pro-active agents can be implemented. 2APL also concerns a agent platform in which agents can be developed and tested, with the assistance of a graphical interface. This platform is built on top of the Jade platform and makes use of the Jade communication facilities.

METATEM Concurrent METATEM is an agent programming language based upon the direct execution of temporal logical formulae. It is made out of two key aspects: an execution mechanism for temporal formulae and an operational model that treats single executable temporal logic programs as asynchronously executing objects in a concurrent object-based system. The logic used is a discrete, linear temporal one. Each object asynchronously executes its own set of temporal formulae and this generates an infinite sequence of states. Moreover, a communication mechanism between separate objects is provided that consists of a broadcast message-passing mechanism.

DMARS The distributed Multi-Agent Reasoning System is the successor of the Procedural Reasoning System (PRS) which has its conceptual roots in the belief-desire-intention model (BDI). The BDI model is materialized in dMARS agents by plans. Each agent has a plan library, which is a list of plans specifying courses of action that may be activated by an agent in order to achieve its intentions. Each plan is defined by several components. A trigger condition establishes the scenarios under which the plan should be considered. The context of a plan specifies the circumstances or pre-conditions that must be met so that the execution of the plan may commence. Maintenance conditions specify the factors that must remain true while the plan is executed. Finally, a plan has a body which defines a potentially complex course of action, which may consist of both goals and primitive actions.

AGENTSPEAK The AgentSpeak programming language is based on logic programming and follows the BDI model. An AgentSpeak agent is defined by a set of beliefs and a set of plans. Beliefs are ground (first-order) atomic formulae. Plans contain a head which consists of a triggering event and a context, which is a logical consequence of that agent's current beliefs if the plan is to be considered. Plans also contain a body which consists of sequences of basic actions that represent atomic operations that can be performed by the agent and that affect the environment. Furthermore, two types of goals are considered in AgentSpeak: achievement goals (formed by an atomic formulae prefixed with the '!' operator) and test goals (prefixed with the '?' operator). An achievement goal states that the agent wants to achieve a given state of the world in which the associated atomic formulae is true. In the other hand, a test goal means that the agent wants to test whether the associated atomic formulae can be unified with one of its beliefs. An AgentSpeak agent can be seen as a reactive planning system in which plans are triggered by the addition or deletion of beliefs. These operations have origin either in the perception of the environment or in the addition or deletion of goals as a result of the execution of plans triggered by previous events.

SLABS is a formal specification language for agent-based systems aimed at facilitating the development of large-scale complicated multi-agent systems. The specification of MAS in SLABS consists of a set of specifications of agents and castes, which is a novel concept in agent-based development. Indeed, the concept of caste is a natural evolution of the concepts of classes in object-orientated development. In this approach, castes can play a significant role in the phases of requirements analysis and specification as well as in the ones of design and implementation. In that sense, a caste description contains a description of the structure of its states and actions, a description of its behavior, and a description of its environment. It can be equivalently represented in a text form or in a graphic form. Furthermore, the language can also be used to specify the environment of an agent as a subset of the agents in the system that may affect its behavior. This constitutes another major difference from the object-oriented approach. It is also possible to define the state and action spaces as a set of variables and behaviors as a set of transition rules. Each of these rules consists of a description of a scenario of the environment, the action to be taken by the agent in the scenario and a condition of the agent's internal state and previous behavior.

C.1.0.2 *Environment Description Languages*

PDDL Planning Domain Definition Language defines a standard encoding language for planning tasks. A PDDL planning task is made of five components. Objects are entities that exist in the world and are of interest. Predicates represent properties of the objects of interest and have an associated logical value. The initial state represents the initial state of the world, before the execution of the plan. The goal specification consists of a list of predicates that the agent wants to execute, i.e., predicates that are to have a true value. Finally, a plan has actions or operators, which are ways to interact with the world that influence the final state of the world. In a PDDL plan this information is stored in two separate files: a domain file contains predicates and actions and a problem file contains the description of objects, initial state and goals specification.

ELMS ELMS is a language used for the specification of simulated multi-agent environments. This language follows an approach based on specific agent technologies for cognitive agent programming and high-level agent communication. ELMS allows the description of the environments of the agents during simulations. With this language is it possible to define the perceptible properties of the agents as well as the kinds of interactions through actions (which influence the world and its objects) and perceptions (which read information from the world and its objects). ELMS can be used to characterize several classes of environments, namely environments that are (from the point of view of the agents): inaccessible, non-deterministic, non-episodic, and dynamic. ELMS can be used to define rather complex environments, supporting a wide range of agent-based applications, particularly in the fields of social simulation.

OWL The Web Ontology Language (OWL) is a family of knowledge representation languages characterized by formal semantics. It is especially target at web information processing and was designed to be interpreted by computers and not for being read by people. This language was developed essentially for defining and instantiating Web Ontologies which describe entities and their relations. Essentially, OWL is a set of XML elements and attributes, whose understanding is given by means of a standard, that are used to define terms and their relationships. An OWL ontology may include descriptions of classes, properties and their instances.

XML The Extensible Markup Language consists of a set of rules for encoding electronic documents defined by the W3C. XML aims at emphasizing simplicity, generality, and usability of information over the Internet. As it is a textual data format it has a strong support via Unicode for the languages of the world with numerous programming interfaces that are used to implement applications that interact with XML documents. Although the main focus of this language is on documents, it is widely used for the representation of arbitrary data structures.

C.1.0.3 *Other Specification Languages*

AGENT UML AUML is an extension to UML. The need for this extension comes from the fact that UML provides an insufficient basis for modeling agents and multi-agent systems. In fact, when compared to objects, agents are active because they show initiative, autonomy and proactiveness, and have control over whether and how they process external requests. Moreover, agents are not isolated software processes, they rather make up societies of interdependent agents. As a consequence, system designers are often troubled when they have to capture unique features of multi-agent systems using general-purpose object-oriented tools. In AUML, agents are represented in terms of their name, role and the organization they belong to. Moreover, the capabilities of agents also define them. Each capability may be constituted by a number of properties. The input denotes what the agent must receive in the form of an input to achieve his task. The output denotes what the capability generates as a result of the work. Input constraints represent constraints that are expected to be true in order for the capability to be performed. Output constraints define constraints that must be true after the capability has been performed. Input-output constraints represent constraints that must be true during the duration between input and output situations. Finally the capability may contain a description in natural language. Moreover, AUML also defines the notion of services in terms of name, description, type, protocol, ontology, content language and properties. Interactions among agents are represented using standard UML sequence diagrams.

AGENT OWL the motivation behind Agent OWL is related with the fact that multi-agent systems lack the interconnection with semantic web standards such as OWL. The aim of this language is to present a semantic knowledge model of an agent suitable for discrete environments. It makes use of the Jena semantic web library and the JADE agent system and allows interconnection of Agent and Semantic Web technologies. As the language addresses knowledge management domains, it is largely based on CommonKADS which is divided into two main parts: a knowledge model based on three other models (organizational model, agent model and task model) and the design of the system, which in this case is based on the MAS. When modeling a knowledge model for an application the first three CommonKADS models are followed, namely: (1) The Organizational; (2) The Task Model and (3) The Agent or Actor Model.

DESIRE DESIRE (framework for DEsign and Specification of Interacting REasoning components) is intended to be used to specify an operational multi-agent system. It allows the system designer to precisely define the intra-agent functionality and the inter-agent functionality. In that sense, it provides mechanism to specify the expertise required to perform the domain tasks for which the agent is responsible in terms of the knowledge requirements and the reasoning capabilities. Moreover, it also allows specifying the expertise necessary to define the mechanisms that guide coordination, cooperation and other forms of social interaction in terms of the knowledge requirements and the reasoning capabilities. DESIRE looks at the agents and the system they build as a compositional architecture. Therefore, all functionality is designed as a series of interacting, task-based, hierarchically structured components. Tasks are defined by their inputs, their outputs and their relationship to other tasks. On the other hand, interaction and coordination between components, the external world and users is defined by the way that information is exchanged, sequenced and controlled. These components can be of any complexity, ranging from simple functions and procedures to whole knowledge-based systems, and can perform any domain-specific function.

C.2 AGENT DEVELOPMENT TOOLS AND METHODOLOGIES

PROMETHEUS Prometheus is a methodology for building agent-based software specifically especially oriented to the development of BDI agents. This methodology consists of three phases. In the system specification phase the basic functionalities of the system are identified, along with inputs (percepts), outputs (actions) and important shared data sources. The next phase, architectural design, uses the outputs from the system specification to determine which agents will shape the system and how their interactions will take place. The last phase, called detailed design, looks at the internals of each agent and how it will accomplish its tasks within the overall system.

VISUALAGENT VisualAgent (Maria and Silva, 2005) is a software development environment aimed at supporting the development of multi-agent systems throughout the development lifecycle, from design to implementation. This tool relies on the Model Driven Architecture (MDA) and uses the MAS-ML modeling language. MAS-ML is an extension of UML and constitutes a modeling language specifically oriented for multi-agent systems specification. Using this modeling language, it is possible to model the objects, agents, environments, organizations and roles.

In order to do this, three different structural diagrams can be used: extended UML class diagram, organization diagram and role diagram. Class diagrams represent agents, organizations, environments and relationships between these entities and classes. Organization diagrams are used to model the agent's organization, their properties, the roles that they define and the agents that play these roles. The role diagrams model the relationship between the roles identified in the organization diagrams and between the roles and the resources available in the organization entities.

In order to make an intuitive use of this tool, this development environment provides a graphical interface in which agents and their roles and organization can be defined. It is organized in three components. The graphical tool supports the visual modeling using MAS-ML. The transformation tool gets the elements generated by the graphical tool and transforms them into UML XMI. Finally, the code generation tool takes this input and generates the corresponding java code based on the Agent Society Framework (ASF) (Silva and Garcia, 2003).

SADAAM SADAAM is an agile methodology, based loosely on Agent UML, that supports the development of multi-agent systems. It utilizes techniques derived from a variety of Agile Development methods. The agent development process is organized in four key phases: Design, Test-Driven Implementation, Release and Review, and Refactor and Enhancement. These phases are applied iteratively until a finished state is reached.

The main objective of the Design phase is to translate system requirements into design decisions. Specifically, it identifies system behaviors and the roles performed by the agents whilst engaging in these behaviors. Moreover, in this phase the agents and their relationships, interactions and activities are defined. The test-driven implementation phase is aimed at help managing and controlling agent development via agile techniques. It is organized into 5 steps, from the creation of a test agent to the review, refactor and restart of the phase, as necessary. In the Release and Review phase, working code is delivered to the costumer for review. The main objective in this phase is to test the behavior of the application while on its target environment and context and, ultimately, end the development process. In the last phase, continuous improvements and enhancements are applied to the finished code, as needed.

TROPOS Tropos is a software development methodology that relies on notions such as agent, goal, task and (social) dependency. These are used to model and analyze early and late software requirements, architectural and detailed design, and to implement the final system. It therefore covers the whole software development process. Tropos has two baseline ideas. First, agent-related notions (e.g. goals, plans) are used in all phases of software development. Second, Tropos covers also the very early phases of requirements analysis. This allows for a deeper understanding of the environment where the software will operate and of the kind of interactions that will occur. The proposed methodology concerns four phases.

In the Early Requirements phase, the problem is understood and an organizational setting is studied. This phase produces as output an organizational model which considers the relevant actors, their goals and their inter-dependencies. The second phase is called Late Requirements. In this phase, the system-to-be is already described within is operational environment, properly contextualized by its relevant functions and qualities. In the Architectural Design phase, the global architecture of the system is defined in terms of subsystems. The bridges that make the interconnections between these subsystems are also defined by means of data exchange, control and other dependencies. The last phase is called the Detailed Design. In this phase, the behavior of each component of the architecture is defined in a more refined fashion. Specifically, each agent is describe in terms of its goals, beliefs and capabilities, along with the interaction between each one.

MASE Multiagent Systems Engineering attempts to define a methodology for designing and developing multi-agent systems. The key idea is to extend the concept of object-orientation to multi-agent systems. While similar to objects, agents are typically defined to have traits such as autonomy, cooperation, perception, and pro-activeness that imply characteristics objects generally do not have. The authors of the project highlight two basic differences: (1) objects are passive as they react to external stimuli, but do not exhibit goal directed behavior; (2) agents typically use a common messaging language whereas object messages are usually class dependent. Agents are thus modeled as "active objects", or, in other words, as objects with goals. This approach has the advantage that it can build on existing object-oriented analysis and design techniques such as OMT and UML. Moreover, it adds additional features such as goals, sensors, and effectors to take on agent specific semantics.

The main focus of MaSE is to help a designer take an initial set of requirements and analyze, design, and implement a working multi-agent system. Moreover, it does so on a platform-independent fashion, not targeted at a specific agent architecture, programming language or communication framework. This means that a multi-agent system designed in MaSE can be implemented in different ways. This methodology is defined by the following phases. The Analysis phase, which consists of three steps: Capturing Goals, Applying Use Cases, and Refining Roles. The Design phase has four steps: Creating Agent Classes, Constructing Conversations, Assembling Agent Classes, and System Design. This methodology is the foundation for the agentTool development platform.

GAIA Gaia is a methodology for agent-oriented analysis and design. The Gaia methodology is on the one hand general in the sense that it can be applied to a wide range of multi-agent systems. On the other hand, it is comprehensive as it deals with both the macro-level (a society of agents) and the micro-level (each individual agent) aspects of systems. Essentially, Gaia is founded on the view of a multi-agent system as a computational organization of various interacting roles. In that sense, in this approach one moves away from the agents towards the actual roles. Similarly to MaSE, Gaia borrows some terminology and notation from object-oriented analysis and design. It provides an agent-specific set of concepts through which a software engineer can understand and model a complex system as a process of organizational design.

This process is organized into two phases. The first one, is the analysis phase. In this phase, the objective is to develop an understanding of the system and its structure without reference to any implementation detail. In Gaia, this understanding is captured in the system's organization which is defined in terms of the roles of the agents, which systematically interact with other roles. The design phase of Gaia is the second one and derives slightly from the "classical" design phase. In fact, usually the aim of a "classical" design process is to transform the abstract models derived during the analysis stage into models at a sufficiently low level of abstraction so that they can be easily implemented. In the other hand, the aim of the design phase of Gaia is to transform the analysis models into a sufficiently low level of abstraction that traditional design techniques (including object-oriented techniques) may be applied in order to implement agents. To put is simple, Gaia is concerned with how a society of agents cooperate to execute the high level goals of the system, and what is required of each individual agent in order to achieve this. Concretely how an agent executes its services is beyond the scope of the methodology and is dependent on the particular application domain.

C.3 AGENT PLATFORMS

Agent platforms are the environments in which agents actually live their lifecycle. These platforms aim to concentrate a set of tools that can make the development and execution of agents easier. Generally, agent platforms provide support for agent development, at the level of specific programming languages and agent models. They also provide tools for agent management and communication, namely at the level of message transport. It is also common to find support for mobility and remote access, allowing to agents to migrate between different instances of the platform in run-time and allowing the platform to be accessed remotely. Finally, these platforms also provide support at the service access level, with tools that allow for searching, announcing and registering services. In the next sub-section an analysis of several platforms that meet these requirements is presented.

JADE JADE (Java Agent DEvelopment Framework) (Bellifemine et al., 2007) is a software Framework completely implemented in Java. It results in a lightweight platform which minimal requirements are the existence of a Java 1.4 runtime environment. This means that it can run in a wide range of devices, which include mobile ones. It was developed with the aim of simplifying the implementation of multi-agent systems by means of a middleware in compliance with FIPA standards. FIPA promotes standards which are intended to improve the interoperation of heterogeneous agents and the services that they can represent. These standards are organized into categories: agent communication, agent transport, agent management, abstract architecture and applications. Agent communication is the core category of the FIPA multi-agent system model.

Jade has also several graphical tools that provide support for debugging and deployment. Moreover, the agent platform can be distributed across machines that can have differences even at the Operating System level. As a way of improving its highly distributed nature, Jade can be controlled via a remotely accessible GUI. Jade is also highly configurable and dynamic as agents can be moved, at run-time, between machines. The configuration can be even changed at run-time by moving agents from one machine to another one, as and when required.

JADEX BDI AGENT SYSTEM Jadex is a Java based agent environment which allows developing goal oriented agents following the so-called BDI software model – Belief, Desire, Intention. In essence, this software model separates the activity of selecting a plan from a plan library, from the actual execution of currently active plans. Plans are constituted by a number of steps, some of which may include invocation to other plans. BDI agents are thus able to balance the time spent on deliberating about plans, i.e., choosing what to do, and actually executing those plans. This separation is base on the work of Michael Bratman on the theory of human practical reasoning. This agent environment is a FIPA compliant one and provides a framework and a set of development tools to simplify the creation and testing of agents. The Jadex framework consists of an API, an execution model, and predefined reusable generic functionality. The API provides access to the Jadex concepts when programming plans which are plain Java classes, extending a specific abstract class, which provides useful methods.

As Jadex is built on top of the Jade agent platform, several tools are readily available. This is not only true for the tools included in JADE, such as the Sniffer or the DummyAgent, but also regards third party tools like the beangenerator plug-in for the ontology design tool Protégé. Moreover, several new tools that are related to BDI are also provided. Specifically, the BDI viewer allows to view the internal state of a Jadex agent, that is, its current beliefs, goals, and plans.

The Jadex Introspector is similar to the JADE Introspector, allowing monitoring and influencing the execution of an agent, by observing and influencing how incoming events are handled.

DIET AGENTS DIET Agents is an Open Source multi-agent platform developed in Java. With the objective of achieving a lightweight, scalable, robust, adaptive and extensible platform, a bottom-up design was used. It is especially targeted at the rapid development of peer-to-peer prototype applications and adaptive and distributed applications that use a bottom-up, nature-inspired techniques.

Among the key features of the platform, the authors highlight a clean layered architecture, with a kernel that is lightweight, simple and general. This kernel is fail-fast which means that it constrains and minimizes the use of threads, sockets and memory. The developed agents are autonomous yet lightweight, making it possible to run 100,000s of agents in a single VM. The platform has a model-event infrastructure which provides sophisticated visualization support. Another feature that is highlighted is the thread-safe agent execution model which makes programming new agent behaviors a straightforward process.

The platform uses a layered architecture organized into three layers. The application layer contains code specific to particular applications, along with specific debugging and visualization components. The ARC layer provides Application Reusable Components which provide optional functionality that is useful to many applications. Finally, the core layer provides the minimal software needed to implement multi-agent functionality in the framework.

MADKIT MadKit is an agent platform written in Java intended to be highly modular and scalable. It is built upon the AGR organizational model: Agent – Group – Role. The key point in this model is that agents are situated in groups and play roles. MadKit allows high heterogeneity in agent architectures and communication languages, and several customizations. At the communication level, the platform uses a peer to peer mechanism, and allows developers to quickly develop distributed applications based on multi-agent concepts.

As this is a very versatile platform, agents may be programmed in Java, Scheme (Kawa), Jess (rule based engine) or BeanShell. Moreover, other script languages may be easily added.

COUGAAR Cougaar (Cognitive Agent Architecture) is a Java-based platform that supports the construction of highly scalable distributed agent-based applications. This platform is the result of a DARPA research project that had as main objective to develop an open-source agent-based architecture able to support applications ranging from small-scale systems to large-scale highly-survivable distributed systems.

The design of Cougaar is based on nodes (Java Virtual Machines) which can contain one or more agents. Agents in turn contain one or more plugins which, through their interactions, define the behavior of each agent. There are also services which are provided at the node level, such as the inter-agent message transport, which are advertised by pluggable infrastructure components. Moreover, Cougaar can extend, be extended or interoperate with other agent platforms, such as Jade.

AJANTA Ajanta is a system for programming agent based applications over the Internet. In that sense, this platform specifically focuses on mobile agents: agents that are able to migrate autonomously from node to node. The main advantage of this approach is that agents can move with the users they represent or, moreover, they can move to remote server resources in case of need. A typical use of this type of agents is as personal assistants of their owners, moving between networks in order to follow them.

Therefore, Ajanta particularly focuses on mechanisms for secure and robust executions of mobile agents in open systems. In this platform, the mobile agent paradigm is based on the generic concept of a network mobile object. Agents in this system are active and mobile objects, which encapsulate code and execution context along with data. Ajanta is implemented in the Java language and its security mechanisms are designed according to Java's security model. It also builds on top of several other facilities of Java, such as object serialization, reflection, and remote method invocation.

MAGE MAGE (Multi-AGent Environment) is a software environment for designing, integrating and running agent-oriented software fully implemented in Java language. Specifically, this platform provides support for agent mental state representation, reasoning, negotiation, planning, cooperation and communication. Moreover, it provides system designing support, description and assembling of knowledge and capability, negotiation and cooperation designing for agent-based computing on the Internet.

MAGE simplifies the development of multi-agent system in several ways. First, it provides a middleware layer in compliance with FIPA specifications as well as a set of tools that supports the debugging and deployment phase. Second, it simplifies integration of applications through multiple schemes of software reuses and an agent-oriented software design with a graphic user interface. Moreover, it simplifies the running management through a powerful user interface with significant run-time support.

The agent platform can be distributed across machines (which not even need to share the same OS) and the configuration can be controlled via a remote GUI. The configuration can be even changed at run-time by moving agents from one machine to another one when required. The system needs only Java Runtime version 1.2 or later version as a running environment.

AGENTSCAPE AgentScape is a middleware layer especially targeted at the support of large-scale agent systems. The main objectives of this platform are to provide a platform for large-scale agent systems, to support multiple code bases and operating systems and to ensure the interoperability with other agent platforms. In that sense, the middleware provides a minimal but sufficient support for agent applications with the conviction that the middleware should be adaptive or reconfigurable in such a way that it can be tailored to very specific application scenarios.

There are three basic concepts in AgentScape: agents, objects and locations. Agents are autonomous entities that interact with each through exchange of messages. Agent migration in the form of weak mobility is also supported, which means that only the data state moves. Objects are passive entities that are only engaged into computations reactively on an agent's initiative. The location is the virtual place in which agents and objects reside. Besides agents, objects, and locations, the AgentScape model also defines services which provide information or activities on behalf of agents or the AgentScape middleware.

The AgentScape Architecture is organized into two components: the AgentScape Operating System kernel and the AgentScape Operating System middleware. This second component includes a number of system services, among which agent servers, object servers, service access providers, location managers, and host managers.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both LaTeX and LyX:

http://code.google.com/p/classicthesis/