Experiences - A forgotten component of epistemic states*

Leila Amgoud[1], Mark d'Inverno[2], Nardine Osman[3], Henri Prade[1], and Carles Sierra[3]

[1] IRIT, UPS-CNRS, 118 route de Narbonne, 31062 Toulouse Cedex 09, France
[2] Computing Department Goldsmiths, University of London, London SE14 6NW, UK
[3] IIIA-CSIC, Campus UAB, 08193 Cerdanyola, Catalonia, Spain

email: {amgoud, prade}@irit.fr, dinverno@gold.ac.uk, {nardine, sierra}@iiia.csic.es

Abstract. This paper investigates a specific type of information that should have an important place in the epistemic state of individual agents, namely their experiences. Just as beliefs, desires, or intentions, experiences should be properly represented, and their specific role in reasoning and decision processes clearly identified. After formally defining what an experience is, the paper explains in what respect experiences differ from and complement beliefs, and are not just ordinary cases. The added value of experiences in agent reasoning and decision making is then discussed in detail.

Keywords: belief; experience; expectation; case.

1 Introduction

The idea of *experience* is central in different human activities such as scientific inquiry, learning processes for education, or even in the understanding of art, as well as the role of shared experience in communication. While it has received some attention in human sciences [10,11,2,16,20], its specificity has not been entirely recognized in artificial intelligence. Certainly, the idea of experience as being the result of the accumulation of cases of interest is at work in case-based reasoning [1], and some authors, e.g. [3,17,9,18,19], have used the word 'experience' for referring to valuable cases (then stored in memory). But then, a set of experiences, viewed as a repertory of cases, is not considered as an original component of agents' epistemic states that interacts with generic beliefs, and which should be handled specifically, together with beliefs, in reasoning and in decision making tasks. It is this latter view that we shall advocate in the following. It indeed contrasts with case-based reasoning, since cases are then the only kind of information that is considered. Moreover, an ordinary case is not usually enough

* AT2012, 15-16 October 2012, Dubrovnik, Croatia.

to constitute an experience, as explained in the following.

According to Dewey [10], there is experience each time an agent performs an action and associates consequences to it. This author also emphasized that neither an action alone nor a consequence alone does make an experience. This leads us to the following informal definition: An experience is an action performed by an agent x who perceives consequences of that action. An experience can also be made from watching another agent who performs an action. Agent x may also evaluate (positively or negatively) the perceived consequences of the action. For instance, in a case like "Mary experienced a traffic jam when going to airport last Saturday morning, got angry and missed her flight", the action is "experiencing a traffic jam", the retained consequences are "getting angry" and "missing a flight", and the evaluation is (implicitly) negative here.

Thus, an experience has two main features that distinguish it from ordinary actions. Indeed, at least some of its consequences: i) are unexpected for the agent, and ii) emotionally significant for the agent. Then, the experience is worth being kept in mind by the agent. Strictly speaking, the surprise created by the unexpected consequence of the experience is part of the emotion, which also includes other components that make the consequence relevant for the agent. Thus, while the agent is emotionally affected by the experience, the experience is associated with some consequences that are unexpected with respect to the agent's beliefs. In the following we concentrate on this latter aspect. As such, an experience constitutes a piece of knowledge which differs from generic beliefs.

In the following, after having introduced the representation setting used in the paper in Section 2, we formally define the notion of experience, and show how it interacts with the notion of belief in Section 3. We then discuss the added values of experiences in reasoning and decision making in Section 4, before indicating lines for further research in the conclusion.

2 Logical language

We assume a logic $(\mathcal{L}, \mathtt{CN})$ where $\mathcal{L}$ is a set of all well-formed formulas, $\mathtt{CN}$ is a consequence operator, that is a function from $2^{\mathcal{L}}$ to $2^{\mathcal{L}}$. Indeed, for $X \subseteq \mathcal{L}$, $\mathtt{CN}(X)$ returns all the consequences (under $\mathtt{CN}$) that follow from X. We do not require here any particular property from $\mathtt{CN}$ (beyond the assumption that $\mathtt{CN}(\mathtt{CN}(X)) = \mathtt{CN}(X)$). A notion of *consistency* is associated with the logic $(\mathcal{L}, \mathtt{CN})$:

Definition 1 *A set* $X \subseteq \mathcal{L}$ *is* consistent *in logic* $(\mathcal{L}, \mathtt{CN})$ *iff* $\mathtt{CN}(X) \neq \mathcal{L}$. X *is* inconsistent *otherwise.*

Regarding the logical language, we assume that it contains two kinds of formulas ($\mathcal{L} = \mathcal{F} \cup \mathcal{G}$ with $\mathcal{F} \cap \mathcal{G} = \emptyset$):

i) $\mathcal{F} \subseteq \mathcal{L}$ which represent *factual* beliefs about what is true in the environment. For instance, 'My laptop is heavy' is a factual belief since it concerns a particular instance.

ii) $\mathcal{G} \subseteq \mathcal{L}$ contains the *generic* beliefs about the (*default*) laws that hold in the environment. An example of generic beliefs is 'The first generation laptops are heavy'. Unlike the previous example, it concerns a class of laptops and not one particular instance.

Moreover, we shall use a list of atoms $\mathcal{V}$ that do not appear in $\mathcal{L}$, and which play the role of *names*. Finally, we let `Ags` denote a set of agents.

Action is a key notion in the formulation of experience. Indeed, as we shall see in the next section, an experience amounts to performing an action and observing and evaluating its consequences. An action has a name (e.g., "going to the USA"), it has a set of pre-conditions that should be fulfilled before performing the action, and a set of post-conditions which represent the consequences of the action. Finally, an action is always performed by one or more agents.

Definition 2 *An* action *is a tuple* $a = (N, P, C, A)$ *such that: i)* $N \in \mathcal{V}$ *is the* name *of the action; ii)* $P \subseteq \mathcal{F}$ *are the* pre-conditions *of applicability of the action; iii)* $C \subseteq \mathcal{F}$ *are the* consequences *of the action; and iv)* $A \subseteq$ `Ags` *is the set of* agents *who perform the action.*

It is worth mentioning that the same action name may be associated with different pre-conditions and/or with different post-conditions. Let us consider the action "going to the USA". From one year to another, the visa requirement may change and thus, the pre-conditions are not the same. Similarly, for the same action, the consequences may change.

3 The notion of experience

In this section, we focus on one particular agent $ag \in$ `Ags` who is equipped with a consistent knowledge base $\Sigma \subseteq \mathcal{L}$. An experience of this agent (called *actor* of the experience) is an action performed by the agent itself. The consequences attached to this action are also assigned and evaluated (positively or negatively) by this agent. Thus, the agent *establishes a causality relation* between the action and its consequences. Note that in contrast with the use of action in planning, the consequences are not given a priori (as part of the definition of the action), but identified by the actor. In some sense, experiences might feed the libraries of actions that we start with when planning.

An experience is personal to the actor, since the same action made by two different agents may not manifest the same consequences for both. Although an agent may report some of their own experiences to another agent, the experience reported does not in general become an experience of the latter agent. Moreover, what makes an action a genuine experience for an agent is the fact that it *emotionally affects* the actor positively or negatively, and its consequences are *not expected* by the actor (the agent either expects exactly the contrary of the perceived consequences, or the agent had no particular expectations a priori).

Thus for instance, our weekly shoppings are certainly not experiences. Assuming that amounts to say that there is nothing peculiar or different from a normal visit. However, if one meets the president in the supermarket, then shopping gives birth to an experience since it is unexpected to see the president there.

Besides, experiences can both be made in the past or envisaged in the future. All these remarks lead us to the following definition.

Definition 3 *An* experience *is a pair* $e = (y, z)$ *such that:*

- y *is an action* (N, P, C, A) *such that:*
  a) $A = \{ag\}$*;*
  b) $P \subseteq \mathtt{CN}(\Sigma) \cap \mathcal{F}$*;*
  c) $C \subseteq \mathcal{F}$ *s.t.* $C \not\subseteq \mathtt{CN}(\Sigma)$*.*
- $z = \mathtt{eval}(C)$ *s.t.* $\mathtt{eval} : \mathcal{F} \to \mathcal{S}$ *where* $\mathcal{S}$ *is a scale.*

If $C = \emptyset$*, then* e *is a* future experience*. Otherwise, it is a* past *one.*

It should be emphasized that the action may be a *simple observation* (e.g., looking at a sunset), or a *physical action* (e.g., going to a restaurant). So, you hearing somebody who calls you an idiot counts as an experience. Moreover, the preconditions of the action may be misrepresented or misunderstood by the actor. For example, Peter thought he did not need a business visa for China.

Similarly, the consequences may be wrongly associated to the action. For instance, one may wrongly think that a train is delayed because of the snow whilst in reality the train in question has some engine problem. What matters here is perceived causality [6]. The perceived consequences may also be a strict *subset* of the real consequences.

Note that the condition $C \not\subseteq \mathtt{CN}(\Sigma)$ captures the idea that consequences are unexpected for the actor. The idea is that at least some consequences in C do not follow from the knowledge base of the agent. This covers two cases:

- either $\exists c \in C$ such that $\neg c \in \mathtt{CN}(\Sigma)$
- or $\exists c \in C$, $c \notin \mathtt{CN}(\Sigma)$ and $\neg c \notin \mathtt{CN}(\Sigma)$.

This means that the generic beliefs of the agent were not enough to forecast all the consequences of the action in the circumstances were the action took place. In the last of the two above cases, C is consistent with $\mathtt{CN}(\Sigma)$, but not deducible from $\mathtt{CN}(\Sigma)$. In the first case, the consequences obtained are in contradiction with what was expected. In both cases, the experience brings a supplementary piece of knowledge w.r.t. the generic knowledge, and in case of contradiction with what was expected, the agent learns that some (new) *exception* do exist with respect to its generic beliefs.

Thus, past experiences are not just ordinary reported cases. They have been singled out for their unexpected consequences that did not leave the agent feeling indifferent when it perceived them. As a consequence, the epistemic state of an agent facing a current situation about which it is trying to draw conclusions, is made of three parts, i) the factual beliefs about the current situation, ii) the generic beliefs applicable to this situation, and iii) the experienced cases that supplement this generic knowledge. Note that the experienced cases if they are sufficiently repeated may give rise to new pieces of generic beliefs, or may lead to revise these beliefs. However, experienced cases, as such, are not generic beliefs, and should be handled in a specific manner.

As just said, experiences supplement generic knowledge, but we have to distinguish between two quite different cases: i) the consistency case, and ii) the abnormality case. In both cases, the epistemic state of the agent will be represented by a pair of subsets of possible worlds, say K and E, where K is supposed to account for the worlds that are *not impossible* according to the agent's generic knowledge (and the factual beliefs about the current situation), or if we prefer the set of interpretations that are consistent with this knowledge, while E gathers the worlds that are guaranteed to be possible in the sense that they have been already encountered by the agent in past experiences.

In the *consistency case*, we have $E \subseteq K$. This corresponds to a situation of *bipolar* information [5], where one can distinguish between a positive part of the information, here represented by E, which represents what is possible for sure (because it was experienced), and a negative part of the information which corresponds to all the possible worlds that are ruled out for sure (due to beliefs), namely here $\Omega \setminus K$, where Ω denotes the set of all the possible worlds. More generally, K and E may be replaced by a pair of possibility distributions with intermediary grades of possibility [5].

At this step of the discussion, we will view K and E as binary possibility distributions, i.e., $\omega \in K$ (resp. $\omega \notin K$) means that ω is not impossible (resp. impossible), while $\omega \in E$ means that ω is possible for sure since it was experienced. But $\omega \notin E$ does not mean at all that ω is impossible, but just that it has not been experienced yet. As can be seen, the "classical" knowledge representation situation corresponds to the case $K = E$, where no distinction is made between what is "not impossible" and what is "possible for sure". Making such a distinction is important for acknowledging the specific nature of experiences with respect to generic knowledge, since agents tend to favor the former when reasoning and making decisions.

The *abnormality case* is more tricky, since then $E \cap (\Omega \setminus K) \neq \emptyset$, i.e. there is some consequence associated with some experience that is in contradiction with the agent's generic knowledge. As long as the generic knowledge is not revised (in order to recover the consistency condition between E and K), the only way to handle this inconsistency is to consider that in fact K represents the set of all the possible worlds that are *normally* possible, while there exist some abnormal (exceptional) worlds outside K. This is similar to nonmonotonic reasoning where one has to reason under incomplete factual knowledge with rules having potential exceptions. Here some exceptions factually exist since they were experienced.

4 Reasoning with and about experiences

Different forms of reasoning may involve experiences. In particular, one may reason about experiences for making a kind of compilation of a set of experiences. One may also reason about a factual situation on the basis of a set of generic beliefs and of a set of experiences. One may also wonder how to take advantages of both beliefs and experiences in a decision, or in an argumentation process. We briefly examine these three classes of problems.

4.1 Percolation of experiences

There are two situations where it makes sense to aggregate elementary experiences into a *compound* experience.

The first one is when the agent performs the same action in very similar circumstances, and perceives the same consequences, then the description of circumstances may be *minimally enlarged* to cover all the slightly different circumstances of the elementary experiences. This includes the case of getting rid of details that are finally found irrelevant for the description of the experience. For instance, an agent gets sick each time after eating a particular dish in different restaurants, and finally retains as a "global" experience that eating this dish she gets sick, without no longer associating any particular restaurant to this state of facts.

The second case occurs when the same action that is repeated in the same circumstances has led to differently perceived consequences. This does not mean that there is any inconsistency: experiences never contradict each other, they just accumulate. Still it is important to keep track of the variability of the consequences by taking as a new consequence the *disjunction* of the consequences of the elementary experiences. Such situation is encountered in case-based reasoning, where, e.g., two second hand cars that are quite similar in all respects may have been sold to significantly different prices.

Generally speaking, it is a distinctive feature of positive information to be accumulated *disjunctively*: the fact that a possible world ω is possible for sure because it is observed does not prevent that another world ω' to be also possible for sure because it is observed later. This contrasts with negative information, which corresponds to the classical view of logical knowledge bases, where the more pieces of the information the agent has, the more restricted the set of worlds remaining possible is (since they are not impossible), as a result of the *conjunctive* combination of these pieces of information.

4.2 Reasoning from experiences and beliefs

Let us first observe that from beliefs on may infer new beliefs, and even from generic beliefs, one may infer new pieces of generic belief, see, e.g., [4,13], but from experiences, or from experiences and beliefs, one does not infer experiences. So, reasoning from experiences and beliefs makes only sense when applying them to a factual situation under consideration.

Generally speaking, when reasoning about a factual situation s, experiences may either complement or challenge beliefs, as already explained. They complement beliefs when the conclusions obtained from beliefs that apply to s are imprecise or even amount to complete ignorance. However, s may appear in the preconditions of some past experiences. Then, one may take advantage of the consequences of those experiences in order to draw plausible conclusions about s.

Still the problem of reasoning with *both* beliefs and cases, which requires a bipolar setting, has received little attention up to very few exceptions [21]. For illustrating the issue, let us consider the following example which corresponds to a simple situation where both beliefs and experiences bring some valuable and consistent information.

Example. We have two pieces of generic beliefs (coming from general regulations), namely the two following rules:

- "if an employee is in category 1, his monthly salary (in euros) is necessarily in the interval [1000, 2000]",
- "if an employee is in category 2, his monthly salary (in euros) is necessarily in the interval [1500, 2500]".

Suppose we also know by experience that

- "an employee in category 1 has a monthly salary typically in the interval [1500,1800]",
- "an employee in category 2 has a monthly salary typically in the interval [1700, 2000]".

Strictly speaking, all the values in [1500,1800] or in [1700, 2000] have not been observed for employees of categories 1 and 2 respectively, but one may consider that these intervals are the convex hulls of set of values that have been experienced for such employees. Then suppose we have the imprecise piece of factual information that "Peter is in category 1 or 2". What can be concluded about Peter's salary? Using the two generic rules we get that his salary is in $[1000, 2000] \cup [1500, 2500] = [1000, 2500]$. Indeed, from "if p_1 then q_1", "if p_2 then q_2", "p_1 or p_2", one concludes "q_1 or q_2" (or if prefer in propositional logic terms $\neg p_1 \vee q_1, \neg p_2 \vee q_2, p_1 \vee p_2 \vdash q_1 \vee q_2$).

Besides, the values in $[1500, 1800] \cap [1700, 2000] = [1700, 1800]$ are for sure possible for Peter's salary. The use of the intersection $\cap$ here may be justified intuitively by the fact that the more imprecise the situation, the smaller the associated set of values that are possible for sure. Indeed, it is only those values that are common to the two contexts 'category 1' and 'category 2' that can be retained as values that are possible for sure in the imprecise context 'category 1 or category 2'. A formal logical mechanism that can handle inferences with positive information can be found in [5]; see [14] for the case of fuzzy rules (i.e., with *graded* possibility distributions).

When an experience leads to consequences that are not consistent with the conclusions obtained from beliefs about s, we have to give *priority to experiences over beliefs* (at least if there is no uncertainty on the consequences associated with experiences). This also means that we need a case-based reasoning mechanisms working together with a non-monotonic mechanism for handling the default rules expressing generic beliefs. Something that has still to be done.

Besides, the fact that agents possess both beliefs and experiences should also considerably enrich argumentation, since experiences may provide a natural basis for counter-argumention, by offering counter-examples to generic rules that have been used in a previous argument for supporting a conclusion, which then becomes challenged.

4.3 Making decision from experiences and beliefs

Decision under uncertainty takes advantage of the available knowledge (possibly pervaded with uncertainty) about what can be deduced about the state of the world in the current situation, and of how good or bad are the consequences of actions performed in precisely known states of the world. Then, depending on what are the rationality postulates, different decision criteria, such as, for example, the expected utility, can be justified ; see, e.g. [7].

In contrast with such approaches, Gilboa and Schmeidler have already advocated the need for making decisions on the basis of experiences, called *cases* in [15]. Indeed, it is common in everyday life that people make decisions on the basis of their own experiences, or maybe those of others. In [15], they propose a formal model for case-based decision making, and justify the use of a similarity-based decision criterion. It is worth noting that in their model, beliefs are not involved at all. However, it seems clear that the application of generic beliefs to the current situation for determining the possible states of the world where the action will take place, should enrich the case-based decision process. Thus, it suggests the need for a new decision model which satisfactorily combines beliefs and experiences.

A formal parallel between decision under uncertainty and case-based decision has been already made for pessimistic and optimistic qualitative decision criteria [12]. Intuitively speaking, in decision under uncertainty, according to the pessimistic criterion, a good decision in an ill-known situation, is a decision that in any precise state compatible with the ill-known situation leads to a good consequence, while in case-based decision, the pessimistic criterion privileges the decisions (if any) that in all cases similar to the current situation have led to a good consequence according to past experiences. The optimistic criteria in decision under uncertainty (resp. in case-based decision) considers that a decision is good as soon as there exists at least one state compatible with the ill-known situation that leads to a good consequence (resp. at least one case similar to the current situation where the decision has led to a good consequence according to past experience).

Such decision criteria could be adapted when we have both generic beliefs and experiences about candidate decisions. Consider the pessimistic point of view, in the case where experiences are consistent with beliefs (i.e., $E \subseteq K$ in the notations of Section 3). Then asking for having all cases similar to the current situation leading to a good consequence is less requiring than asking that all non impossible states compatible with the ill-known situation lead to a good consequence. Note that with the optimistic point of view, this is the converse: it is more requiring to have at least one case similar to the current situation leading to a good consequence rather than having one state compatible with the ill-known situation that leads to a good consequence (since the set of worlds similar to E remains strictly included in K).

When some experiences contradict generic beliefs (i.e. in the abnormality case of Section 3), then the agent has to assume that the generic knowledge covers *normal* situations only, and that some abnormal cases may be encountered as well by experience. Then a more difficult situation is when all *normal* situations lead to good consequences for a candidate decision, while it has been observed that this decision in some abnormal cases led to bad consequences. If the agent is moderately pessimistic, it may be already satisfied if in all *normal* states compatible with the ill-known situation, the decision leads to a good consequence, even if some (but not all) experiences where the decision was experimented led to bad consequences. But, if the agent is very pessimistic, the agent will be afraid to make a choice that once led to a bad consequence, even it was abnormal. This suggests the need for designing sophisticated decision criteria able to cover all these situations.

5 Concluding remarks

This paper has presented a first manifesto in favor of a forgotten component in the modeling of the epistemic states of an agent in Artificial Intelligence: the experiences held by the agent. Indeed, we believe that experiences are a component which is fundamentally different from beliefs, desires and intentions, and which deserves to be formalized and clearly distinguished from the others. We have explained in detail why beliefs and experiences are not of the same nature, and we have indicated in general terms how they may impact both reasoning and decision making in general. This should lead to a new generation of agent architecture, that we may call *XBDI agents* (where X stands for experience), which enlarges the original BDI architecture [8]. Then, the addition of this component raises many basic issues, such that how to compute intentions on the basis of beliefs, experiences, and desires.

The new architecture will make it possible to define platforms that allow several agents to share experiences. For instance, one may imagine an application in which agents bring their past experiences to others about a particular topic. We can also think of another type of application in which a group of agents want to have a joint future experience such as visiting a place of interest together, for instance. In this case, the future experience is a common desire for the agents. Both these examples are only possible if the notion of experience for agents is clearly defined.

6 Acknowledgements

The authors acknowledge the support of the ACE project (Autonomic Software Engineering for online cultural experiences) of the European Chist-Era programme.

References

1. A. Aamodt, E. Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. AI Commun. 7(1): 39–59, 1994.
2. L. Andresen, D. Boud and R. Cohen. Experience-based learning. In: Understanding Adult Education and Training, (G. Foley, ed.), 2nd Edition, Allen & Unwin, Sydney, 225–239, 1995.
3. S. Arai, K. Sycara, T. R. Payne. Experience-based reinforcement learning to acquire effective behavior in a multi-agent domain. Proc. 6th Pacific Rim inter. conf. on Artificial intelligence (PRICAI'00), Springer-Verlag, Berlin, 125-135, 2000.
4. S. Benferhat, D. Dubois, H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. Artificial Intelligence, 92, 259–276,1997.
5. S. Benferhat, D. Dubois, S. Kaci, H. Prade. Modeling positive and negative information in possibility theory. Int. J. Intell. Syst., 23, 1094–1118, 2008.
6. J.-F. Bonnefon, R. Da Silva Neves, D. Dubois, H. Prade. Predicting causality ascriptions from background knowledge: Model and experimental validation. Int. J. Approx. Reasoning 48, 752–765, 2008.
7. D. Bouyssou, D. Dubois, M. Pirlot, H. Prade (eds.) Decision-Making Process - Concepts and Methods, Wiley, 2009.
8. M. E. Bratman. Faces of Intention. Cambridge University Press, 1999.

9. L. Brito, P. Novais, J. Neves. Temporality, priorities and delegation in an E-commerce environment. Proc. 14th Bled Electronic Commerce Conf., Bled, Slovenia, June 25 - 26, 2001, 53–68.
10. J. Dewey. Reality as experience. In: Collected Works of John Dewey. The Middle Works 1899-1924, vol. III. South. Illin. Univ. Press, 1977.
11. J. Dewey. Art as experience. Minton, Balch & Company, New York, 1934.
12. D. Dubois, L. Godo, H. Prade, A. Zapico. On the possibilistic decision model: From decision under uncertainty to case-based decision. Inter. J. of Uncertainty, Fuzziness and Knowledge-Based Systems 7, 631–670, 1999.
13. D. Dubois, H. Prade: Focusing vs. belief revision: A fundamental distinction when dealing with generic knowledge. Proc. 1st Inter. Joint Conf. on Qualitative and Quantitative Practical Reasoning (ECSQARU-FAPR'97), (D. M. Gabbay, R. Kruse, A. Nonnengart, H. Jrgen Ohlbach, eds.), Bad Honnef, Germany, June 9-12, LNCS 1244, 96–107,1997.
14. D. Dubois, H. Prade, L. Ughetto. A new perspective on reasoning with fuzzy rules. Int. J. Intell. Syst., 18, 541–567, 2003.
15. I. Gilboa, D. Schmeidler. Case-based decision theory. The Quarterly Journal of Economics, 110 (3), 605-639, 1995.
16. D. Middleton, S. D. Brown. The Social Psychology of Experience. Studies in Remembering and Forgetting. Sage Publ., London, 2005.
17. P. Novais, L. Brito, J. Neves. Experience-based mediator agents as the basis of an electronic commerce system. Workshop 2000 - Agent-Based Simulation, Passau, Germany, 2000.
18. Z. H. Sun, G. Finnie. Experience based reasoning for recognising fraud and deception. Proc. 4th Inter. Conf. on Hybrid Intelligent Systems (HIS '04), IEEE Computer Society, 80–85, 2004.
19. Z. H. Sun, G. Finnie. A fuzzy logic approach to experience-based reasoning. Int. J. Intell. Syst., 22, 867–889, 2007.
20. K. Thorgersen, C. F. Thorgersen. Impetus, imagination and shared experience in aesthetic communication. Inter. Network of Philosophers of Education, INPE, Biennal Conf., Kyoto, Japan, 11 p., Aug. 9-12, 2008.
21. L. Ughetto, D. Dubois, H. Prade. Implicative and conjunctive fuzzy rules - A tool for reasoning from knowledge and examples. Proc. 16th National Conf. on AI (AAAI'99), Orlando, 214-219, 1999.

Preface

This volume contains the papers presented at AT 2012: Agreement Technologies 2012, held on October 15-16, 2012, in Dubrovnik, Croatia. Agreement Technologies refer to computer systems in which autonomous software agents negotiate with one another, typically on behalf of humans, in order to come to mutually acceptable agreements. An agent may choose whether to fulfill an agreement or not, and it should fulfill it when there is an obligation to do so derived from the standing agreements. Autonomy, interaction, mobility and openness are key concepts studied within the AT approach. Semantic alignment, negotiation, argumentation, virtual organisations, learning, real time, and several other technologies are part of the sandbox to define, specify and verify such systems.

The first International Conference on Agreement Technologies, AT 2012, is an interdisciplinary forum that brings together researchers and practitioners working on the various topics comprising this emergent and vibrant field. It provides an avenue to discuss and exchange new ideas and techniques for the design, implementation and verification of next generation open distributed systems centered on the notion of agreement among computational agents. The AT 2012 conference focuses on the following topics: Argumentation, negotiation, Trust and reputation, Coordination and distributed decision making, Computational Social Choice, Semantic alignment, Inter-theory Relations, Decision and game theoretic foundations for agreement, Agent Commitments, Semantic Service Coordination, Normative Systems, Individual reasoning about norm adoption, Collective deliberation about norm adoption, Autonomic Electronic Institutions, Group planning agreements, Deliberative Agreement: social choice and collective judgment, Evolution of organisational structures, Social intelligence, Logics for Agreements, Realtime agreements, Agreement patterns, Agreement technologies architectures, environments and methodologies, Applications of agreement technologies (e.g. web service composition, contract automation, supply chain automation, sensor networks, etc.)

There were 71 submissions. Each submission was reviewed by at least 2 program committee members or reviewers. The committee decided to accept 27 full papers reporting on original and previously unpublished work that is currently not under review in any conference or journal and 25 position papers that represent either a summary of original work that has already been published in a conference or journal, or a summary of original results obtained as a product of STSMs funded by the Agreement Technologies Cost Action, or ongoing research that can lead to important contributions to or of Agreement Technologies. The program also includes one invited talk and a panel, sponsored by The European Network for Social Intelligence (SINTELNET) .

The conference is supported by COST Action IC0801 on Agreement Technologies, the European Science Foundation (ESF), The European Network for Social Intelligence (SINTELNET), CETINIA, the Ministry of Science, Education and Sports of the Republic of Croatia, the Faculty of Electrical Engineering and Computing of the University of Zagreb and the Spanish Project on Agreement Technologies. We wish to extend our warm thanks to the AT 2012 Organizing

Committee, the Program Committee members, the reviewers and all authors of submitted papers for making this conference so rewarding.

October 15-16, 2012

Sascha Ossowski (Conference chair)
Francesca Toni (Programme co-chair)
George Vouros (Programme co-chair)

Program Committee

Leila Amgoud	IRIT - CNRS
Cristiano Castelfranchi	Institute of Cognitive Sciences and Technologies
Carlos Chesñevar	UAB, Universidad Nacional del Sur
Marco Colombetti	University of Lugano
Paul Davidsson	Malmö University
Ioanna Dionysiou	University of Nicosia
Juergen Dix	Clausthal University of Technology
Maria Ganzha	University of Gdańsk, Poland
Mirjana Ivanovic	Faculty of Science, Department of Mathematics and Informatics
Gordan Jezic	University of Zagreb
Antonis Kakkas	University of Cyprus
Matthias Klusch	DFKI
Ryszard Kowalczyk	Swinburne University of Technology
Mario Kusek	University of Zagreb, FER
Lea Kutvonen	University of Helsinki
Christian Lemaitre	Universidad Autonoma Metropolitana, UAM
Merik Meriste	University of Tartu, Tallinn University of Technology
John-Jules Meyer	Utrecht University
Pavlos Moraitis	Paris Descartes University
Viorel Negru	West University of Timisoara
Eugénio Oliveira	Faculdade de Engenharia Universidade do Porto - LIACC
Sascha Ossowski	University Rey Juan Carlos
Marcin Paprzycki	IBS PAN and WSM
Marco Schorlemmer	Artificial Intelligence Research Institute, IIIA, CSIC
Michael Ignaz Schumacher	University of Applied Sciences Western Switzerland (HES-SO)
Carles Sierra	IIIA
Kuldar Teveter	Tallinn University of Technology
Francesca Toni	Imperial College
Denis Trcek	Faculty of Computer and Information Science, University of Ljubljana
László Zsolt Varga	MTA SZTAKI
Harko Verhagen	Dept. of Computer and Systems Sciences, Stockholm University/KTH
George Vouros	University of Piraeus
Emil Weydert	CSC, University of Luxembourg

Additional Reviewers

B

Bikakis, Antonis
Bonzon, Elise
G

Gao,Yang
Gjermundrod, Harald
K

Koster, Andrew
Kusek, Mario
L

Lopes Cardoso, Henrique
M

Mihalas, Stelian
O

Osman, Nardine
P

Peppas, Pavlos
Podobnik, Vedran
U

Urbano, Maria Joana