

Worldwide AI



■ IIIA is a public research center, belonging to the Spanish National Research Council (CSIC), dedicated to AI research. We focus our activities on a few well-defined subdomains of artificial intelligence, positively avoiding dispersion and keeping a good balance between basic research and applications, and paying particular attention to training Ph.D. students and technology transfer. In this article, I survey some of the most relevant results we have obtained during the last 12 years.

Ramon Lopez de Mantaras

A Survey of Artificial Intelligence Research at IIIA

The Artificial Intelligence Research Institute (IIIA) is a leading artificial intelligence research center in Spain that belongs to the Spanish National Research Council (CSIC). It was founded in 1991 and, since 1994, has been located on the campus of the Autonomous University of Barcelona. IIIA grew out of an AI research group at the Center for Advanced Studies in Blanes (Spain) that started AI research in 1985.

On average IIIA has had about 50 members per year during the last 12 years with a peak of almost 80 members in 2012. In total around 200 different people, including visiting researchers as well as master's and Ph.D. students, have been members of IIIA over the past 20 years. Seventy-seven students have completed their Ph.D. work at our Institute, 48 of them during the last 12 years.

The well over 1000 papers published by IIIA members over



Figure 1. The IIIA People in the Institute's Courtyard.

the last 12 years account for a significant share of the total Spanish publications in the main AI journals and conferences during this period, and many of them are highly cited. IIIA researchers have received 32 awards during these 12 years, including 9 best paper awards at international conferences, 6 outstanding Ph.D. thesis awards, the 2012 European Society for Fuzzy Logic and Technology (EUSFLAT) Best Dissertation Award, the 2011 IFAAMAS Victor Lesser Distinguished Dissertation Award, and the 2011 Association for the Advancement of Artificial Intelligence (AAAI) Robert S. Engelmore Award.

Intensive collaborations take place with academic institutions from numerous countries and particularly from France, UK, Italy, Australia, USA, Germany, Argentina, Czech Republic, Japan, Israel, Brazil, and Austria. As a result of these collaborations, about 50 percent of our publications have international coauthorship.

On average, about 40 percent of the institute's total budget (which is approximately 3 million euros per year, including salaries) comes from competitive external research grants as well as contracts with private companies. A total of 57 grants have been obtained during the last 12 years, including 15 European Community grants, 38 grants from the Spanish government, and 4 grants from the Catalan govern-

ment. Besides, we have signed 14 contracts with private companies. We have also received additional funding as a Consolidated Group from the Catalan government. Overall, more than 14 million euros have been raised during the last 12 years.

Our research has been and is always guided by concrete and challenging applications in fields such as health, e-commerce, automated negotiation, conflict mediation and resolution, music, tourism, logistics, supply-chain management, transport, energy, data privacy, and social networks, among others. Several of our systems, tools, and applications have been distributed outside the institute and in some cases have been commercialized. Among the many AI applications developed during the last 12 years the most recent ones are prediction of energy demand in intelligent buildings; early detection of potential failures in windmill turbines for electrical power generation; improving the customer shopping experience in supermarkets; managing safe personalized tourism for disabled persons; AI tools for social networks-based music education; online digital games that are worth playing by older people for active and positive aging; social networking using autonomic software agents to enrich, encourage, and enliven online cultural experiences in virtual visits to museums; recruitment intelligent matching system to improve online



Figure 2. The IIIA Building in the Autonomous University of Barcelona Campus.

job searching; and automatic generation of audiovisual narrative such as summaries of soccer matches or other types of TV broadcast events and news.

Our focus is on a few well-defined subdomains of artificial intelligence, avoiding dispersion and keeping a good balance between basic research and applications, and paying particular attention to training Ph.D. students and technology transfer.

The existence of the Technological Development Unit (UDT) provides technological support to our research activities and improves our technology transfer capabilities by channeling contacts with industry. In particular, we keep strong ties with our three spin-off companies: iSOCO,¹ Strands,² and CogniCor.³ Our first spin-off, iSOCO, was set up in 1999 and is dedicated to the design of intelligent software components for Internet-related applications. Today, iSOCO is a leading company within its sector in Spain and has received several awards including the 2012 Excellence Award in Technology from the

magazine *Dirigentes*. Strands was started in 2004 and is dedicated to recommendation systems particularly in the finances sector and now has labs in Corvallis (Oregon), San Francisco, Madrid, and Barcelona. CogniCor was founded in 2011 based on the results of a large (more than 5 million euros) project called Agreement Technologies. This company develops software products for the automatic resolution of customers' complaints using machine learning and case-based reasoning techniques. CogniCor has received several awards, including the prestigious 2012 European Union Tech All Stars Competition.

This article summarizes some contributions and main results of the Artificial Intelligence Research Institute over the last 12 years. For a survey of the contributions from 1996 to 2001, the reader should refer to the *Contributions to Science* article by Lopez de Mantaras (2001). For a full list of publications and further information regarding our projects and other research activities, please refer to our institute's website.⁴



Figure 3. Scene of a Virtual Auction House Electronic Institution.

Research Activities

Research activities are structured around three departments: Learning Systems, Multiagent Systems (MAS), and Logic, Reasoning, and Search. In what follows I describe some of the most relevant results we have obtained during the last 12 years. Most of the research has been and is being done in strong collaboration among the three departments. For this reason in what follows I describe our activities and results organized thematically instead of by department.

Negotiation Algorithms

Negotiation algorithms are fundamental for distributed computing. In particular, in areas like agent-based systems, service-oriented computing, cloud computing, or the grid. The IIIA scientists have proposed and tested several automatic negotiation algorithms and methods that have been highly influential (Faratin, Sierra, and Jennings 2004; Ramchurn et al. 2007). These algorithms are based on a mixture of utility and information theories and incorporate heuristics that allow for efficient search processes in large solution spaces (De Jonge and Sierra 2012).

Electronic Institutions

An electronic institution (EI) is a computational framework that regulates the behavior of autonomous entities (humans or software) in open systems, as conventional institutions do. They shift the traditional focus of autonomous and multiagent

systems research to the organizational and environmental aspects of collective action (d’Inverno et al. 2012). Electronic institutions originated at IIIA. Refinements, tools, and applications have been developed there and elsewhere and have become mainstream in the field giving rise to abundant theoretical and technological developments. Applications range from hotel information systems to virtual markets, water-management policy testing, and archaeological simulation. Research related to this topic was honored with different awards including the most cited paper (between 2005 and 2010) of the journal *Engineering Applications of Artificial Intelligence* (Arcos et al. 2005).

Norms in Multiagent Systems

Norms, as declarations of permissions, prohibitions, and obligations, may be used to achieve intended behavior of autonomous agents, with the advantage of using logic to reason about compliance and enforcement of conventions in multiagent systems, both at design and at run time. IIIA has contributed formalisms and tools for norm synthesis (Morales et al. 2013), norm and convention emergence (García-Camino et al. 2009), and norm-based programming (Villatoro, Sabater-Mir, and Sen 2011).

Computational Models of Trust and Reputation

Computational trust and reputation models have been recognized as one of the hot topics in the area

of multiagent systems. These social artifacts are necessary to maintain the social order in virtual communities as they have been used in real communities. Till now, however, the cognitive approach to these models has been somehow marginal even though it has been demonstrated to be the appropriate approach in socially complex environments where humans and artificial agents interact. Following the steps of cognitive scientists like Cristiano Castelfranchi and Rosaria Conte, IIIA has developed a state-of-the-art computational trust and reputation cognitive model that nowadays is a reference in the area (Sabater-Mir, Paolucci, and Conte 2006). Another area in which IIIA has been active is in the development of methods to evaluate trust and reputation models, first founding the Agent Reputation and Trust (ART) test-bed initiative (Fullam et al. 2005) and then with the development of the alpha test bed (Jelenc et al. 2013).

Optimization in Multiagent Systems

Optimization in multiagent systems has recently generated much attention because of its potential application to a wide number of domains such as e-commerce, disaster management, and information acquisition through embedded devices. The challenges posed by this type of optimization mainly take the form of very hard optimization problems that are substantially different from problems traditionally dealt with in other areas (for example, industrial processes or scheduling applications). First, we were able to provide the most general optimization model to support the automated formation of auction-based supply chains (Cerquides et al 2007; Parsons, Rodriguez-Aguilar, and Klein 2011) and to handle failures in markets (Ramchurn et al. 2009). Second, we have been the first to characterize how to provide quality guarantees for the max-sum algorithm (Vinyals et al. 2010), the state-of-the-art algorithm for distributed constraint optimization and one of the most popular algorithms in machine learning because of its wide applicability (for example, computer vision, error decoding). Both results have been achieved in collaboration with Ulle Endriss and Alessandro Farinelli. At present we are also actively investigating the application of our novel techniques to the coordination of unmanned air vehicles (UAVs) in dynamic environments (Pujol-Gonzalez et al. 2013) in collaboration with Milind Tambe.

Conceptual Systems in Computational Environments

Humans and computers interact more and more through environments that are highly distributed, dynamic, heterogeneous, and open. Managing and sharing conceptual systems under these conditions requires sound mathematical models as well as novel computational techniques that span from standard approaches to distributed knowledge represen-

tation to computational creativity techniques (Colton, Lopez de Mantaras, and Stock 2009) for concept invention. At IIIA we have studied the mathematical foundations of ontology-based semantic alignment and integration in distributed environments (Schorlemmer and Kalfoglou 2008); we have proposed a formal framework for reasoning about properties of knowledge transformations as they occur in decentralized interactions between large and complex computational artifacts (Schorlemmer and Robertson 2011); and we have developed semantic-alignment techniques that tackle the problem of semantic heterogeneity in the context of agent communication and interaction (Atencia and Schorlemmer 2012).

Autonomous Robots and Developmental Robotics

The real-time coordination of the behaviors of a team of robots in complex environments (uncertain, dynamic, adversarial) is a very challenging problem. We have been able to successfully deal with this problem by means of a case-based reasoning approach and we have tested it in robot soccer. We were able to achieve coordinated passing behaviors between teammates (Ros et al. 2009). This achievement required the invention of new case retrieval and case adaptation algorithms that are useful in other domains such as in the design of more sophisticated computer games. This result was obtained in close collaboration with Manuela Veloso. We have also contributed to the very challenging problem of autonomous robot localization based on visual detection of objects in the robot's natural environment by means of signatures of places constructed from constellations of visual features detected in panoramic images (Ramisa et al. 2009). More recently, we have started a new research line in developmental robotics in collaboration with Yiannis Demiris. We have obtained some initial results in sensorimotor anticipation based on optical flow information (Ribes et al. 2012).

Multiagent Learning from Communication and Argumentation

While machine learning often focuses on learning from data, we proposed a framework where several learners can, in a critic-based nonnaïve way, learn from communication among them as well as individual data. Communication is modeled as an argumentative process, so the systems do not "learn as told" but can criticize the information received, ask for elaborations, and even teach the opponent when wrong. Integrating autonomous learning with learning from communication allows a more general scheme for acquiring knowledge and information between humans and machines as well as among several autonomous systems. We are among the first to integrate machine learning and computational mod-



Figure 4. An IIIA Pet in the Courtyard.

els of argumentation with the A-MAIL approach. Moreover, we have developed a formal logic-based model to integrate learning and argumentation (Ontañón et al. 2012).

Case-Based Reasoning with Complex Cases

Case-based reasoning (CBR) is a powerful paradigm that goes beyond simple k-NN for domains where

cases are complex structures (Lopez de Mantaras et al. 2006). We have developed a methodology to model and implement CBR inference steps over complex structured cases. This methodology is based on refinement graphs that originated in ILP and that we have generalized to other representation languages. Specifically, we have developed two refinement-based measures of similarity over complex cases (Ontañón and Plaza 2012) and we have also developed a model for blending two or more solution cases by a formal operation that we call amalgam. An amalgam of two solutions is a third solution, which combines as much as possible from the original solutions. Amalgams allowed us to develop a declarative model of multiple case combinations in CBR as well as to model case-based transfer in CBR reuse as a form of analogical inference (Ontañón and Plaza 2010). Another source of complexity in cases, which is an important limitation for CBR systems, is when the cases in the case base and the problems to be solved are in different domains. We have addressed this problem, in collaboration with Simeon Simoff, by integrating analogical and commonsense reasoning components. The analogical reasoning component employs the structure mapping engine (Falkenhainer, Forbus, and Gentner 1989) that allows computing similarities between semantically different, but structurally similar, cases from different domains represented by means of semantic networks. We have successfully applied this approach to the automatic mediation of disputes in conflict resolution applications (Baydin et al. 2011).

Speeding Up Reinforcement Learning by Case-Based Reasoning

One of the main problems of reinforcement learning (RL) algorithms is that they typically suffer from very slow learning rates. One way to speed up RL algorithms is by making use of a conveniently chosen heuristic function, which is used for selecting appropriate actions to perform in order to guide exploration during the learning process. Several methods have been successfully applied for defining the heuristic function. We have been among the first to reuse previously learned policies using a case-based reasoning approach and have proposed a new algorithm called case-based heuristically accelerated reinforcement learning (CB-HARL) (Bianchi, Ros, and Lopez de Mantaras 2009; Bianchi and Lopez de Mantaras 2010) that uses cases in the case base as heuristics. This approach has shown its efficiency in speeding up RL. We have also investigated the use of the CB-HARL algorithm as a means to transfer learning knowledge acquired by one agent during its training in one problem to another agent that has to learn how to solve a similar, but more complex, problem. To do so, we have proposed a new algorithm, called L3 (Celiberto et al. 2011), which works in three stages: in the first stage, it uses Q-learning to learn

how to perform one task, and stores the optimal policy for this problem as a case base; in the second stage, it uses a neural network to map actions from one domain to actions in another, more complex domain; and in the third stage, it uses the cases in the case base learned in the first stage as heuristics in the CB-HARL algorithm, speeding up the learning process. These results have been achieved in collaboration with Reinaldo Bianchi.

Machine Learning with Probabilistic Graphical Models

Probabilistic graphical models (PGMs) are a marriage between probability theory and graph theory. They provide a natural tool for dealing with two problems that occur throughout applied mathematics and engineering — uncertainty and complexity — and in particular they are playing an increasingly important role in the design and analysis of machine-learning algorithms. On the theoretical side, we have provided mathematical results that allow the exact Bayesian learning of classifiers based on tree augmented naïve Bayes models (Cerquides and Lopez de Mantaras 2005a), and results for efficient averaging of ODE classifiers (Cerquides and Lopez de Mantaras 2005b) that have been proven to provide the highest accuracy (Yang et al. 2007). Furthermore, in collaboration with the research group of Stefan Douglas Webb, we have proposed (Zaidi et al. 2013) efficient alternatives that help alleviate the burden of the strong independence assumptions made by the naïve Bayes, the simplest probabilistic graphical model widely applied to classification problems. On the applied side, IIIA maintains research collaborations with Große on the application of PGMs to problems in bioinformatics, and in particular to phylogenetic models. Furthermore, we provide data-mining services for graph databases to SMEs in this area.

Learning for (and from) Time Series

Recently, we have been developing innovative tools and concepts for time series mining. Our work has been mainly focused on similarity measures for the classification of temporal signals (Serrà and Arcos 2012). However, we have also worked on other tasks such as segmentation, motif discovery, outlier detection, and clustering. In general, we are interested in the analysis of the information contained in real-world time series for understanding the systems behind them.

Computer Audition

Computer audition deals with algorithms and systems for audio understanding by machines. In this context, our group pioneered the modeling of music expressivity by learning from experience to automatically generate expressive performances by means of similarity-based reasoning to imitate examples of human performances (Lopez de Mantaras and Arcos




Figure 5. Gesture and Audio Analysis of Classical Guitar Performances.

2002), and to automatically change the tempo of a given performance preserving its expressivity (Grachten, Arcos, and Lopez de Mantaras 2006). More recently we have continued to work on this topic by studying intentional timing deviations in professional performances of classical guitar (Serrà, Ozaslan, and Arcos 2013). In addition, we have taken successful approaches to music structure segmentation (Serrà et al. in press) and on quantifying the evolution of popular music on a large scale (Serrà et al. 2012). Another work on the statistical nature of audio timbral codings (Haro et al. 2012) has been done in collaboration with the Musical Technology Group of Pompeu Fabra University and with the Mathematics Research Centre of the Autonomous University of Barcelona.

Embedding Knowledge into Data Privacy

We study different aspects related to the use of knowledge in data privacy. For example, we have

developed approaches so that data protection takes into account the constraints of databases (for example, relationships between variables) and the semantics of categorical terms. Knowledge has also an important role with respect to the transparency of data protection methods. To achieve transparency, published data has to be accompanied with a description of the data protection process. In this case, in order that disclosure risk measures be accurate enough, they have to use all the knowledge related to the data protection process (Torra 2011; Juárez and Torra 2013).

Boolean-Based Optimization.

Modeling challenging optimization problems as MaxSAT and MinSAT instances, using the language of propositional logic, and then deriving an optimal solution of the original problem with a MaxSAT/MinSAT solver is becoming a highly competitive approach in optimization. At IIIA we have

defined powerful solving techniques that have been decisive for the development of a new generation of exact MaxSAT/MinSAT solvers such as MaxSatz, MinSatz, and WPM2 (Li et al. 2012; Ansótegui, Bonet, and Levy 2013), have studied novel MaxSAT/MinSAT encodings of relevant combinatorial optimization problems, and have stimulated the research on Boolean-based optimization by organizing an international competition of MaxSAT solvers. Moreover, we have defined the first complete inference systems for MaxSAT and multiple-valued MaxSAT (Bonet, Levy, and Manyà 2007).

Study of the Structure of Industrial SAT Instances

SAT formulas provide a method to encode restrictions used in industrial processes. In fact, the SAT problem is a classic problem in computer science, and there exists an old tradition on the study of randomly generated SAT formulas. However, SAT formulas resulting from the encoding of industrial problems have a distinct nature. We have been pioneers in the study of such formulas and have proved that they have a scale-free structure (Ansótegui et al. 2008; Ansótegui, Bonet, and Levy 2009). This contribution will be crucial in the development of future SAT solvers. Currently, our solvers have been among the best in the industrial category in recent MaxSAT competitions.

Distributed Constraint Reasoning by Distributed Search

When constraints are distributed among several autonomous agents, a possible solution method is based on distributed search. This approach uses short messages but it may require an exponential number of them (worst case). We have worked on this approach, considering different ways to improve it: not adding new links among agents in the ABT algorithm (Bessiere et al. 2005), improving the privacy of constraints (Brito et al. 2009), connecting distributed search with arc consistency (both hard and soft cases) (Gutierrez et al. 2013), removing redundant messages (Gutierrez and Meseguer 2012) and generalizing algorithms in distributed constraint optimization (Brito and Meseguer 2008). We have advocated for the use of global constraints in this context (Bessiere et al. 2014). This work has been done in collaboration with Christian Bessiere and Jimmy H. M. Lee.

Distributed Constraint Reasoning by Distributed Inference

Distributed constraint reasoning can also be performed by distributed inference. This alternative method requires a polynomial number of messages, but of exponential size (in the induced width of the tree decomposition used). To improve their practical usage, we have explored different ways of approximating messages: function filtering (Brito and

Meseguer 2010) and zero-tracking decomposition (Pujol-Gonzalez et al. 2011). In the context of fast-changing scenarios, we have also explored the use of Max-Sum, an approximate algorithm performing inference in factor graphs, jointly with tractable high-order potentials with very promising results (Pujol-Gonzalez et al. 2013). This work has been done with Milind Tambe and Alessandro Farinelli.

General Methods and Results on Logics for AI

We have played a key international role in the definition and development of mathematical fuzzy logic and we have obtained important results in the following topics: (1) General and deep results for completeness of fuzzy logics either propositional or first order with respect to different semantics (real, hyper-real, rational, and finite) that cover and significantly extend previous results in the field. Our results have been possible as a consequence of a fruitful collaboration with researchers from different leading institutions on the topic; (2) Formalization of t-norm-based logics dealing with partial degrees of truth, with algebraic semantics, and axiomatization and completeness results, both for propositional and first-order languages (Esteva, Godo, and Noguera 2009; Cintula et al. 2009), which have high applicability in modeling graded notions (Casali, Godo, and Sierra 2011); (3) Development of different systems of fuzzy modal logic (Bou et al. 2011), with applications to reasoning under different forms to uncertainty on non-Boolean algebras of events (Flaminio, Godo, and Marchioni 2013); and (4) Development of a new hierarchy of fuzzy description logics, along with new complexity results based on results of mathematical fuzzy logic (Cerami, Garcia-Cerdaña, and Esteva 2014).

Defeasible Argumentation

We have significantly extended the argumentation framework of defeasible logic programming (DeLP) along two different axes: (1) by incorporating the treatment of possibilistic uncertainty at object level, allowing to stratify defeasible rules in a DeLP program according to their strength, and defining a new recursive semantics that avoids some undesired side effects of the original semantics based on dialectical trees (Alsinet et al. 2014); (2) by defining a discrete temporal extension of DeLP for defeasible causal reasoning, and its deployment for multiagent collaborative planning (Pardo and Godo 2013). All these results have been achieved in collaboration with leading international researchers in the area of computational argumentation such as G. Simari and C. Chesñevar, as well as with T. Alsinet and R. Bejar.

Conclusions

From this overview, which focuses on the last 12

years of our AI research contributions, one can see that our activities cover a wide spectrum from basic to applied research and technology transfer, including spin-off creation. We focus our research activities on a well-defined set of subdomains of artificial intelligence, positively avoiding dispersion. These domains are learning systems; multiagent systems; and logic, reasoning, and search. In all of them we believe we have played, and we are playing, a very significant role in their development. We pay particular attention to training Ph.D. students in these domains and we have intensive collaborations with prestigious academic institutions from numerous countries.

The existence of a technological development unit within IIIA provides technological support to our research activities and improves our technology transfer capabilities by channeling contacts with industry and keeping strong ties with our spin-off companies. Next time you are near Barcelona you are encouraged to visit us and find out more about our activities.

Acknowledgments

IIIA is funded by the Spanish government as one of the research centers belonging to the Spanish National Research Council, as well as through the Spanish R&D projects funding program. Our research is also partially supported by the Catalan government mainly through the “Consolidated Groups” schema for groups of excellence. A significant number of competitive research projects have been, and are, also funded by European Community’s R&D framework programs. Our technology transfer activities are mainly funded by means of contracts with the private sector as well as by creating spin-off companies.

Notes

1. www.isoco.com.
2. strands.com.
3. cognicor.com.
4. www.iiia.csic.es.

References

Alsinet, T.; Bejar, R.; Godo, L.; and Guitart, F. 2014. RP-DeLP: A Weighted Defeasible Argumentation Framework Based on a Recursive Semantics. *Journal of Logic and Computation*. dx.doi.org/10.1093/logcom/exu008

Ansótegui, C.; Bonet, M. L.; and Levy, J. 2013. SAT-Based MaxSAT Algorithms. *Artificial Intelligence* 19(1): 77–105. dx.doi.org/10.1016/j.artint.2013.01.002

Ansótegui, C.; Bonet, M. L.; and Levy, J. 2009. Towards Industrial-like Random SAT Instances. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, 387–392. Menlo Park, CA: AAAI Press.

Ansótegui, C.; Bonet, M. L.; Levy, J.; and Manyà, F. 2008. Measuring the Hardness of SAT Instances. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI 2008)*, 222–228. Menlo Park, CA: AAAI Press.

Arcos, J. L.; Esteva, M.; Noriega, P.; Rodriguez-Aguilar, J. A.

and Sierra, C. 2005. Engineering Open Environments with Electronic Institutions. *Engineering Applications of Artificial Intelligence* 18(2): 191–204. dx.doi.org/10.1016/j.engappai.2004.11.019

Atencia, M., and Schorlemmer, M. 2012. An Interaction-Based Approach to Semantic Alignment. *Journal of Web Semantics* 12/13: 131–147. dx.doi.org/10.1016/j.websem.2011.12.001

Baydin, A. G.; Lopez de Mantaras, R.; Simoff, S.; and Sierra, C. 2011. CBR with Commonsense Reasoning and Structure Mapping: An Application To Mediation. In *Proceedings of the 19th International Conference on Case Based Reasoning (ICCBR 2011) Lecture Notes in Artificial Intelligence 6880*: 378–392. Berlin: Springer.

Bessiere, C.; Brito, I.; Gutierrez, P.; and Meseguer, P. 2014. Global Constraints in Distributed Constraint Satisfaction and Optimization. *The Computer Journal* 57(6). dx.doi.org/10.1093/comjnl/bxt088

Bessiere, C.; Brito, I.; Maestre, A.; and Meseguer, P. 2005. Asynchronous Backtracking Without Adding Links: A New Member in the BT Family. *Artificial Intelligence* 161(1–2): 7–24. dx.doi.org/10.1016/j.artint.2004.10.002

Bessiere, C.; Meseguer, P.; Freuder, E. C.; and Larrosa, J. 2002. On Forward Checking for Non-Binary Constraint Satisfaction. *Artificial Intelligence* 141(1–2): 205–224. dx.doi.org/10.1016/S0004-3702(02)00263-1

Bianchi, R., and Lopez de Mantaras R. 2010. Case-Based Multiagent Reinforcement Learning: Cases as Heuristics for Selection of Actions. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, 355–360. Amsterdam, The Netherlands: IOS Press.

Bianchi, R.; Ros, R.; and Lopez de Mantaras R. 2009. Improving Reinforcement Learning by Using Case-Based Heuristics. In *Proceedings of the 8th International Conference on Case-Based Reasoning (ICCBR 2009)*, Lecture Notes in Artificial Intelligence 5650, 75–89. Berlin: Springer.

Bonet, M. L.; Levy, J.; and Manyà, F. 2007. Resolution for MAX-SAT. *Artificial Intelligence* 171(8–9): 606–618. dx.doi.org/10.1016/j.artint.2007.03.001

Bou, F.; Esteva, F.; Godo, L.; and Rodriguez, R. 2011. On the Minimum Many-Valued Modal Logic over a Finite Residuated Lattice. *Journal of Logic and Computation* 21(5): 739–790. dx.doi.org/10.1093/logcom/exp062

Brito, I., and Meseguer, P. 2010. Cluster Tree Elimination for Distributed Constraint Optimization with Quality Guarantees. *Fundamenta Informaticae* 102(3–4): 263–286.

Brito, I., and Meseguer, P. 2008. Connecting ABT with ARC Consistency. In *Proceedings of the 14th International Conference on Principles and Practice of Constraint Programming (CP 2008)*, Lecture Notes in Computer Science 5202, 387–401. Berlin: Springer.

Brito, I.; Meisels, A.; Meseguer, P.; and Zivan, R. 2009. Distributed Constraint Satisfaction with Partially Known Constraints. *Constraints* 14(2): 199–234. dx.doi.org/10.1007/s10601-008-9048-x

Casali, A.; Godo, L.; and Sierra, C. 2011. A Graded BDI Agent Model to Represent and Reason About Preferences. *Artificial Intelligence* 175(7–8): 1468–1478. dx.doi.org/10.1016/j.artint.2010.12.006

Celiberto, L. A.; Matsuura, J. P.; Bianchi, R.; and Lopez de Mantaras R. 2011. Using Cases as Heuristics in Reinforcement Learning: A Transfer Learning Application. In *Proceedings of the 22nd International Joint Conference on Artificial*

- Intelligence* (IJCAI 2011), 1211–1217. Menlo Park, CA: AAAI Press.
- Cerami, M.; García-Cerdaña, A.; and Esteva, F. 2014. On Finitely Valued Fuzzy Description Logics. *International Journal of Approximate Reasoning*.
- Cerquides, J., and Lopez de Mantaras, R. 2005a. Robust Bayesian Linear Classifier Ensembles. In *Proceedings of the 16th European Conference on Machine Learning (ECML 2005)*. Lecture Notes in Computer Science 3720, 72–83. Berlin: Springer.
- Cerquides, J., and Lopez de Mantaras, R. 2005b. TAN Classifiers Based on Decomposable Distributions. *Machine Learning Journal* 59(1): 1–32.
- Cerquides, J.; Endriss, U.; Giovannucci, A.; and Rodríguez-Aguilar, J. A. 2007. Bidding Languages and Winner Determination for Mixed Multi-Unit Combinatorial Auctions. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI 2007)*, 1221–1226. Menlo Park, CA: AAAI Press.
- Cintula, P.; Esteva, F.; Godo, L.; Gispert, J.; Montagna, F.; and Noguera, C. 2009. Distinguished Algebraic Semantics for T-Norm Based Fuzzy Logics: Methods and Algebraic Equivalentes. *Annals of Pure and Applied Logic* 160(1): 53–81. dx.doi.org/10.1016/j.apal.2009.01.012
- Colton, S.; Lopez de Mantaras, R.; and Stock, O. 2009. Computational Creativity: Coming of Age. *AI Magazine* 30(3): 11–14.
- De Jonge, D., and Sierra, C. 2012. Branch and Bound for Negotiations in Large Agreement Spaces. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2012)*, 1415–1416. Richland, SC: IFAAMAS.
- D’Inverno, M.; Luck, M.; Noriega, P.; Rodríguez-Aguilar, J. A.; and Sierra, C. 2012. Communicating Open Systems. *Artificial Intelligence* 186: 38–94. dx.doi.org/10.1016/j.artint.2012.03.004
- Esteva, F.; Godo, L.; and Noguera, C. 2009. First-Order T-Norm Based Fuzzy Logics with Truth-Constants: Distinguished Semantics and Completeness Properties. *Annals of Pure and Applied Logic* 161(2): 185–202. dx.doi.org/10.1016/j.apal.2009.05.014
- Falkenhainer, B.; Forbus, K. D.; and Gentner, D. 1989. The Structure Mapping Engine: Algorithm and Examples. *Artificial Intelligence* 41(1): 1–63. dx.doi.org/10.1016/0004-3702(89)90077-5
- Faratin, P.; Sierra, C.; and Jennings, N. 2002. Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiations. *Artificial Intelligence* 142(2): 205–237. dx.doi.org/10.1016/S0004-3702(02)00290-4
- Flaminio, T.; Godo, L.; and Marchioni, E. 2013. Logics for Belief Functions on MV-Algebras. *International Journal of Approximate Reasoning* 54(4): 491–512. dx.doi.org/10.1016/j.ijar.2012.08.006
- Fullam, K. K.; Klos, T. B.; Muller, G.; Sabater, J.; Schlosser, A.; Topol, Z.; Barber, K. S.; Rosenschein, J. S.; Vercouter, L.; and Voss, M. 2005. A Specification of the Agent Reputation and Trust (ART) Testbed: Experimentation and Competition for Trust in Agent Societies. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)*, 512–518. Richland, SC: IFAAMAS.
- García-Camino A.; Rodríguez-Aguilar J. A.; Sierra C.; and Vasconcelos W. 2009. Constraint Rule-Based Programming of Norms for Electronic Institutions. *Autonomous Agents and Multi-Agent Systems* 18(1): 186–217. dx.doi.org/10.1007/s10458-008-9059-4
- Grachten, M.; Arcos, J. L.; and Lopez de Mantaras, R. 2006. A Case-Based Approach To Expressivity-Aware Tempo Transformation. *Machine Learning Journal* 65(2–3): 411–437. dx.doi.org/10.1007/s10994-006-9025-9
- Gutierrez, P., and Meseguer, P. 2012. Removing Redundant Messages in N-Ary BNB-Adopt. *Journal of Artificial Intelligence Research* 45: 287–304.
- Gutierrez, P.; Lee, J. H. M.; Lei, K. M.; Mak, T. W. K.; and Meseguer, P. 2013. Maintaining Soft Arc Consistency in BNB-Adopt+ During Search. In *Proceedings of the 19th International Conference on Principles and Practice of Constraint Programming (CP 2013)*, Lecture Notes in Computer Science 8124: 365–380. Berlin: Springer.
- Haro, M.; Serra, J.; Herrera, P.; and Corral, A. 2012. Zipf’s Law in Short-Time Timbral Codings of Speech, Music, and Environmental Sound Signals. *Plos One* 7(3): E33993. dx.doi.org/10.1371/journal.pone.0033993
- Jelenc, D.; Hermoso, R.; Sabater-Mir, J.; and Trcek, D. 2013. Decision Making Matters: A Better Way to Evaluate Trust Models. *Knowledge-Based Systems* 52: 147–164. dx.doi.org/10.1016/j.knosys.2013.07.016
- Juárez, M., and Torra, V. 2013. Toward a Privacy Agent for Information Retrieval. *International Journal of Intelligent Systems* 28: 606–622. dx.doi.org/10.1002/int.21595
- Li, C. M.; Zhu, Z.; Manyà, F.; and Simon, L. 2012. Optimizing with Minimum Satisfiability. *Artificial Intelligence* 190: 32–44. dx.doi.org/10.1016/j.artint.2012.05.004
- Lopez de Mantaras, R. 2001. Contributions to Artificial Intelligence: The IIIA Perspective. *Contributions to Science* 2(1): 9–22.
- Lopez de Mantaras, R., and Arcos, J. L. 2002. AI and Music: From Composition to Expressive Performance *AI Magazine* 23(3): 43–57.
- Lopez de Mantaras, R.; McSherry, D.; Bridge, D.; Smyth, B.; Craw, S.; Faltings, B.; Maher, M. L.; Cox, M.; Forbus, K.; Keane, M.; Aamodt, A.; and Watson I. 2006. Retrieval, Reuse, Revision, and Retention in CBR. *Knowledge Engineering Review* 20(3): 215–240. dx.doi.org/10.1017/S0269888906000646
- Morales, J.; Lopez-Sanchez, M.; Rodríguez-Aguilar, J. A.; Wooldridge, M.; and Vasconcelos, W. 2013. Automated Synthesis of Normative Systems. In *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2013)*, 483–490. Richland, SC: IFAAMAS.
- Ontañón, S., and Plaza, E. 2012. Similarity Measures over Refinement Graphs. *Machine Learning Journal* 87(1): 57–92. dx.doi.org/10.1007/s10994-011-5274-3
- Ontañón, S., and Plaza, E. 2010. Amalgams: A Formal Approach for Combining Multiple Case Solutions. In *Proceedings of the 9th International Conference on Case-Based Reasoning (ICCB 2010)*, Lecture Notes in Artificial Intelligence 6176: 257–271. Berlin: Springer.
- Ontañón, S.; Dellunde, P.; Godo, L.; and Plaza, E. 2012. A Defeasible Reasoning Model of Inductive Concept Learning from Examples and Communication. *Artificial Intelligence* 193: 129–148. dx.doi.org/10.1016/j.artint.2012.08.006
- Pardo, P., and Godo, L. 2013. T-Delp: An Argumentation-Based Temporal Defeasible Logic Programming Framework. *Annals of Mathematics and Artificial Intelligence* 69(1): 3–35. dx.doi.org/10.1007/s10472-013-9334-1

- Parsons, S.; Rodriguez-Aguilar, J. A.; and Klein, M. 2011. Auctions and Bidding: A Guide for Computer Scientists. *ACM Computing Surveys* 43(2): Article 10. dx.doi.org/10.1145/1883612.1883617
- Pujol-Gonzalez, M.; Cerquides, J.; Meseguer, P.; and Rodriguez-Aguilar, J. A. 2011. Communication-Constrained DCOPS: Message Approximation in GDL with Function Filtering. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2011)*, 379–386. Richland, SC: IFAAMAS.
- Pujol-Gonzalez, M.; Cerquides, J.; Meseguer, P.; Rodriguez-Aguilar, J. A.; and Tambe, M. 2013. Engineering the Decentralized Coordination of UAVs with Limited Communication Range. In *Advances in Artificial Intelligence*, Lecture Notes in Artificial Intelligence 8109: 199–208.
- Ramchurn, S. D.; Mezzetti, C.; Giovannucci, A.; Rodriguez-Aguilar, J. A.; Dash, R. K.; and Jennings, N. R. 2009. Trust-Based Mechanisms for Robust and Efficient Task Allocation in the Presence of Execution Uncertainty. *Journal of Artificial Intelligence Research* 35: 119–159.
- Ramchurn, S. D.; Sierra, C.; Godo, L.; and Jennings, N. R. 2007. Negotiating Using Rewards. *Artificial Intelligence* 171(10–15): 805–837.
- Ramisa, A.; Tapus, A.; Aldavert, D.; Toledo, R.; and Lopez de Mantaras, R. 2009. Robust Vision-Based Localization Using Combinations of Local Feature Regions Detectors. *Autonomous Robots Journal* 27(4): 373–385. dx.doi.org/10.1007/s10514-009-9136-9
- Ribes, A.; Cerquides, J.; Demiris, Y.; and Lopez de Mantaras, R. 2012. Incremental Learning of an Optical Flow Model for Sensorimotor Anticipation in a Mobile Robot. In *Proceedings of the IEEE Conference on Development and Learning and Epigenetic Robotics (ICDL-EPIROB 2012)*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Ros, R.; Arcos, J. L.; Lopez de Mantaras, R.; and Veloso, M. 2009. A Case-Based Approach for Action Selection and Coordination in Robot Soccer Gameplays. *Artificial Intelligence* 173(9–10): 1014–1039. dx.doi.org/10.1016/j.artint.2009.02.004
- Sabater-Mir, J.; Paolucci, M.; and Conte, R. 2006. Repute: Reputation and Image Among Limited Autonomous Partners. *Journal of Artificial Societies and Social Simulation* 9(2): 3.
- Schorlemmer, M., and Kalfoglou, Y. 2008. Institutionalising Ontology-Based Semantic Integration. *Applied Ontology* 3(3): 131–150.
- Schorlemmer, M., and Robertson, D. 2011. Reasoning About Distributed Knowledge-Transforming Peer Interactions. *IEEE Transactions on Knowledge and Data Engineering* 23(9): 1419–1431. dx.doi.org/10.1109/TKDE.2010.265
- Serrà, J., and Arcos, J. L. 2012. A Competitive Measure to Assess the Similarity Between Two Time Series. In *Proceedings of the 11th International Conference on Case-Based Reasoning (ICCBR 2012)*, Lecture Notes in Artificial Intelligence 7466, 414–427. Berlin: Springer. dx.doi.org/10.1038/srep00521
- Serrà, J.; Corral, A.; Boguna, M.; Haro, M.; and Arcos, J. L. 2012. Measuring the Evolution of Contemporary Western Popular Music. *Scientific Reports* 2: 521.
- Serrà, J.; Muller, M.; Grosche, P.; and Arcos, J. L. 2014. Unsupervised Music Structure Annotation by Time Series Structure Features and Segment Similarity. *IEEE Transactions on Multimedia Special Issue on Music Data Mining* 47.
- Serrà, J.; Ozaslan, T. H.; and Arcos, J. L. 2013. Note Onset Deviations as Musical Piece Signatures. *Plos One* 8(7): E69268. dx.doi.org/10.1371/journal.pone.0069268
- Torra, V. 2011. Towards Knowledge Intensive Data Privacy. In *Data Privacy Management and Autonomous Spontaneous Security, Revised Selected Papers from the 5th DPM and 3rd SETOP Conferences*, Lecture Notes in Computer Science 6514: 1–7. Berlin: Springer.
- Villatoro, D.; Sabater-Mir, J.; and Sen, S. 2011. Social Instruments for Robust Convention Emergence. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, 420–425. Menlo Park, CA: AAAI Press.
- Vinyals, M.; Cerquides, J.; Farinelli, A.; and Rodriguez-Aguilar, J. A. 2010. Worst-Case Bounds on the Quality of Max-Product Fixed-Points. In *Proceedings of Neural Information Processing Systems (NIPS 2010)*, 2325–2333. La Jolla, CA: Neural Information Processing Systems Foundation.
- Yang, Y.; Webb, G. I.; Cerquides, J.; Korb, K. B.; Boughton, J.; and Ting, K. M. 2007. To Select or to Weigh: A Comparative Study of Linear Combination Schemes for Superparent-One-Dependence Estimators. *IEEE Transactions on Knowledge and Data Engineering* 19(12): 1652–1665. dx.doi.org/10.1109/TKDE.2007.190650
- Zaidi, N. A.; Cerquides, J.; Carman, M. J.; Webb, G. I. 2013. Alleviating Naive Bayes Attribute Independence Assumption by Attribute Weighting. *Journal of Machine Learning Research* 14: 1947–1988.

Ramon Lopez de Mantaras (www.iiia.csic.es/~mantaras) is a research professor of the Spanish National Research Council (CSIC) and director of the Artificial Intelligence Research Institute (IIIA). He earned his M.Sc. in computer science from the University of California Berkeley, his Ph.D. in physics from the University of Toulouse, and a Ph.D. in computer science from the Technical University of Barcelona. A pioneer of artificial intelligence in Spain, with contributions since 1976 in pattern classification, approximate reasoning, expert systems, machine learning, case-based reasoning, autonomous robots, and AI and music, Lopez de Mantaras has been an invited plenary speaker at numerous international conferences. He is the former editor-in-chief of *Artificial Intelligence Communications*, a current editorial board member of several international journals including *AI Magazine*, and an associate editor of the *Artificial Intelligence Journal*. Lopez de Mantaras has been a program committee member in numerous conferences, as well as chairing several AI conferences (UAI-94, ECML'00, ECAI-04, ECML-07, PKDD-07, and IJCAI-07). Lopez de Mantaras is an ECCAI Fellow, corecipient of six best paper awards at international conferences, a recipient, among other awards, of the City of Barcelona Research Prize from the Barcelona City Council, the 2011 AAAI Robert S. Engelmore Memorial Award, and the Spanish National Computer Science Award from the Spanish Computer Society. He was president of the Board of Trustees of IJCAI from 2007 to 2009. He serves on a variety of panels and advisory committees for public and private institutions based in the USA and Europe. Lopez de Mantaras is presently working on case-based reasoning, machine learning for autonomous robots, and AI and music.