

Invited Talks

Preference Elicitation and Preference Learning in Social Choice: New Foundations for Group Recommendation

Craig Boutilier

Social choice has been the subject of intense investigation within computer science, AI, and operations research, in part because of the ease with which preference data from user populations can now be elicited, assessed, or estimated in online settings. In many domains, the preferences of a group of individuals must be aggregated to form a single consensus recommendation, placing us squarely in the realm of social choice.

The application of social choice and voting schemes to domains like web search, product recommendation and social networks places new emphasis on issues such as: articulating suitable decision criteria; approximation; incremental preference elicitation; learning methods for population preferences; and more nuanced analysis of manipulation.

In this talk, I will provide an overview of some of these challenges and outline some of our recent work tackling them, including: learning probabilistic models of population preferences from choice data; robust optimization (winner determination) with incomplete user preferences; incremental preference elicitation for group decision making; and new analyses of manipulation. Each of these poses interesting modeling, knowledge representation and optimization challenges that are best tackled using a combination of techniques from AI, operations research, and statistics.

Craig Boutilier is a professor of computer science at the University of Toronto. He received his Ph.D from Toronto in 1992, and joined the faculty of University of British Columbia in 1991 (where he remains an adjunct professor). He returned to Toron-

to in 1999, and served as chair of the Department of Computer Science from 2004–2010. Boutilier has held visiting positions at Stanford, Brown, Carnegie Mellon and Paris-Dauphine, and served on the Technical Advisory Board of CombineNet for nine years.

Boutilier has published over 180 refereed articles covering topics ranging from knowledge representation, belief revision, default reasoning, and philosophical logic, to probabilistic reasoning, decision making under uncertainty, multiagent systems, and machine learning. His current research efforts focus on various aspects of decision making under uncertainty: preference elicitation, mechanism design, game theory and multiagent decision processes, economic models, social choice, computational advertising, Markov decision processes and reinforcement learning. Boutilier served as program chair for both UAI-2000 and IJCAI-09, and is currently associate editor-in-chief of the *Journal of Artificial Intelligence Research (JAIR)*. He is also a Fellow of the Association for the Advancement of Artificial Intelligence (AAAI).

Ontology-Based Data Management: Present and Future

Maurizio Lenzerini

The need of effectively managing the data sources of an organization, which are often autonomous, distributed, and heterogeneous, and devising tools for deriving useful information and knowledge from them is widely recognized as one of the challenging issues in modern information systems. Ontology-based data management aims at accessing, using, and maintaining data by means of an ontology, i.e., a conceptual representation of the domain of interest in the underlying information system. This new paradigm provides several interesting features,

many of which have been already proved effective in managing complex information systems. On the other hand, several important issues remain open, and constitute stimulating challenges for the knowledge representation and reasoning research community. In this talk we first provide an introduction to ontology-based data management, illustrating the main ideas and techniques for using an ontology to access the data layer of an information system, and then we discuss several important issues that are still the subject of extensive investigations, including the need of higher-order modeling capabilities, inconsistency tolerant query answering methods, and mechanisms supporting update operations expressed over the ontology.

Maurizio Lenzerini is a professor of computer science and engineering at the Università Roma “La Sapienza,” where he is currently leading the Artificial Intelligence and Databases research group. His main research interests are in knowledge representation and reasoning, ontology languages, semantic data integration, and service modeling. His recent work is mainly oriented towards the use of knowledge representation and automated reasoning principles and techniques in information system management, and in particular in information integration and service composition. He has authored over 250 papers published in leading international journals and conferences. He has served on the editorial boards of several international journals, and on the program committees of the most prestigious conferences in the areas of interest. He is currently the chair of the Executive Committee of the ACM Symposium of Principles of Database Systems, a Fellow of the European Coordinating Committee for Artificial Intelligence (ECCAI), a Fellow of the Association for Computing Machinery (ACM), and a member of The Academia Europaea — The Academy of Europe.

Great Moments in KR: Logic Programming Solution to the Frame Problem

Vladimir Lifschitz

Research on the frame problem included attempts to formalize the “commonsense law of inertia”: whatever we know about the state of affairs before executing an action can be presumed, by default, to hold after the action as well. A counterexample, known now as the Yale Shooting Scenario, showed that early attempts to make the idea of commonsense inertia precise were unsuccessful. A way to overcome that difficulty was suggested by experience with the use of Prolog and by work on the semantics of negation in logic programming.

Solving computational problems related to the commonsense law of inertia became possible when methods used in the design of SAT solvers were applied to nonmonotonic extensions of propositional logic. This technology has been applied, in particular, to verifying the reliability of the reaction control system of the Space Shuttle.

Vladimir Lifschitz is a professor of computer science at the University of Texas at Austin. He received a degree in mathematics from the Steklov Mathematical Institute (St. Petersburg, Russia) in 1971 and emigrated to the United States in 1976. His research interests are in the areas of computational logic and knowledge representation.

The Rise and Fall of Linear Temporal Logic

Moshe Y. Vardi

One of the surprising developments in the area of program verification in the late part of the 20th Century is the emergence of linear temporal logic (LTL), a logic that emerged in philosophical studies of free will, as the canonical language for describing temporal behavior of computer systems. LTL, however, is not expressive enough for industrial applications. The first decade of the 21st century saw the emergence of industrial temporal logics such as For-Spec, PSL, and SVA. These logics, however, are not clean enough to serve as objects of theoretical study. This talk will describe the rise and fall of LTL, and will propose a new canonical temporal logic: linear dynamic logic (LDL).

Moshe Y. Vardi is the George Distinguished Service Professor in Computational Engineering and director of the Ken Kennedy Institute for Information Technology Institute at Rice University. He is the co-recipient of three IBM Outstanding Innovation Awards, the ACM SIGACT Goedel Prize, the ACM Kanellakis Award, the ACM SIGMOD Codd Award, the Blaise Pascal Medal, and the IEEE Computer Society Goode Award. He is the author and coauthor of over 400 papers, as well as two books: *Reasoning about Knowledge* and *Finite Model Theory and Its Applications*. He is a Fellow of the Association for Computing Machinery, the Association for the Advancement of Artificial Intelligence, the American Association for the Advancement of Science, and the Institute for Electrical and Electronic Engineers. He is a member of the US National Academy of Engineering, the American Academy of Arts and Science, the European Academy of Science, and Academia Europea. He holds honorary doctorates from the Saarland University in Germany and Orleans University in France. He is the editor-in-chief of *Communications of the ACM*.