

Uncertainty in AI

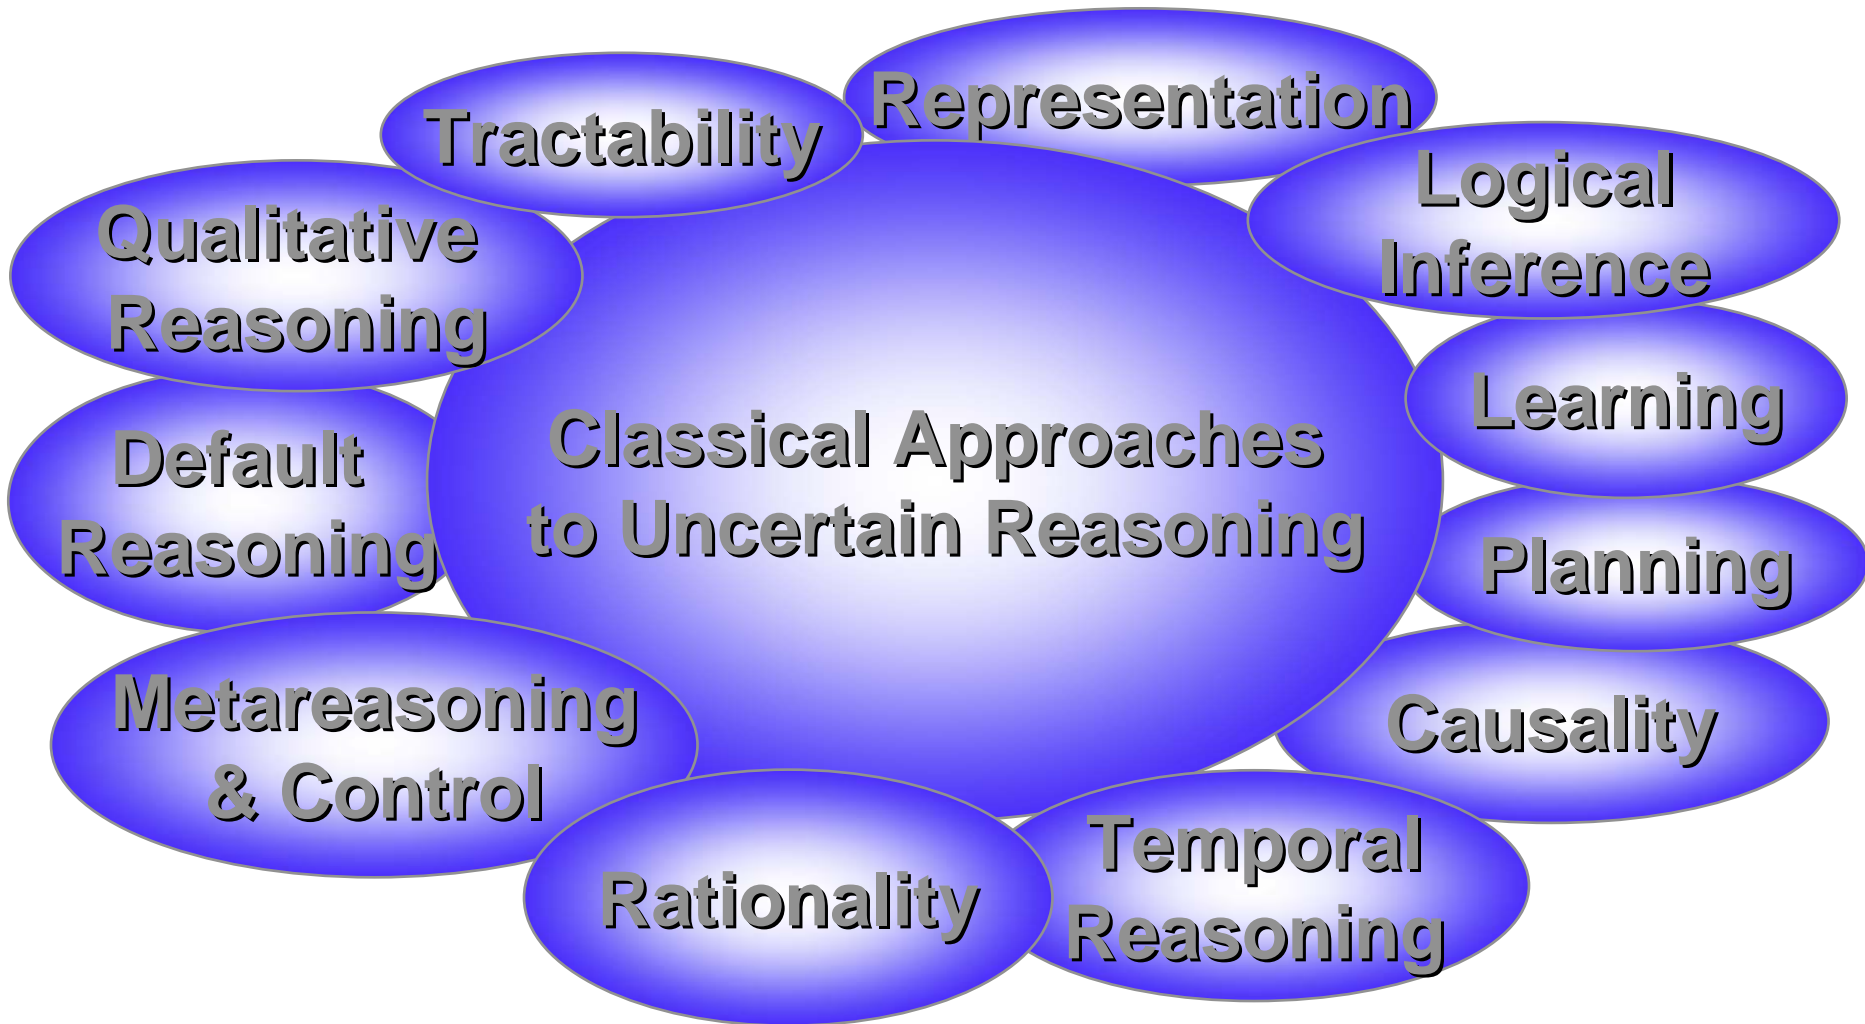
(UAI)

Joe Halpern
Cornell


A Little History

- ◆ The UAI Conference started in 1985
 - Meeting annually since then
 - Premier conference on *reasoning and decision making under uncertainty*
- ◆ Goal: To bring together researchers working on various ways of representing uncertainty
 - In the 1980's, more papers on alternate formalisms for belief and action (Bayesian methods, Dempster-Shafer, certainty factors, possibility measures, fuzzy-set theory, etc.)
 - In recent years, Bayesian paradigms have been dominant
 - ◆ Work on Bayesian networks has had major impact on AI and beyond
 - Have gone from ad hoc analyses to technical sophistication
- ◆ Most conference papers available online at <http://www.auai.org/>
- ◆ Typically colocated with either AAAI or IJCAI
 - But colocated with ICML and COLT in 2004

UAI: A Rich Synthesis




Submissions by Year




253 Submissions

Acceptance By Year



UAI 2004: 76 Papers Accepted

UAI is Selective!



UAI 2004: 30% Acceptance Rate

UAI 2004: Organization

- ◆ Program Committee Co-Chairs
 - Max Chickering
 - Joe Halpern
- ◆ Conference Chair: Chris Meek
- ◆ Local Arrangements: Rob Holte
- ◆ Dates: July 7-11, 2004
- ◆ Location: Banff, Alberta
 - Colocated with COLT and ICML

UAI 2004: Program

◆ 5 Invited Talks

- **Jon Kleinberg**, Cornell University
Cascading Behavior and Bursty Dynamics in Computational Models of Social Networks
- **Lillian Lee**, Cornell University
What is the matter? Explorations in text categorization
- **Alon Orlitsky**, University of California at San Diego
Good-Turing estimation and its applications
- **Moshe Y. Vardi**, Rice University
Markov Processes and Markov Decision Processes: The Verification Perspective
- **Ed George**, University of Pennsylvania
Bayesian or Frequentist? How to Double Your Chances for a Date

UAI 2004: Program

- ◆ Four tutorials
 - **Ronen Brafman**, Ben-Gurion University
Eliciting, modeling, and reasoning about preference using CP nets
 - **Rina Dechter**, UC Irvine
Constraint processing: a graphic models perspective
 - **Nir Friedman**, Hebrew University
Graphical models in computational molecular biology
 - **Martin Wainwright**, UC Berkeley
Graphical models, exponential families, and variational inference
- ◆ 27 Plenary Papers
- ◆ 49 Posters in 2 sessions

UAI 2004: Topics

	#papers
◆ Foundations	4
• L. Schubert, <i>A new characterization of probabilities in Bayesian networks</i>	
◆ Probabilistic modeling	5
• Y. Mao, F. Kschischang, B. Frey, <i>Convolutional factor graphs as probabilistic models</i>	
◆ Multiple agents (inc. game theory, collaborative filtering)	9
• M. Tennenholtz, <i>Reputational systems: an axiomatic approach</i>	
◆ Temporal/dynamic models	5
• A. Howard, T. Jebara, <i>Dynamical systems trees</i>	
◆ Nonprobabilistic/qualitative reasoning	7
• L. Amgoud, H. Prade, <i>Using argument for making decisions: A possibilistic logic approach</i>	
◆ Learning	11
• O. Mamdani, D. Lizotte, R. Greiner, <i>Active model selection</i>	

UAI 2004: Topics

	#papers
◆ Inference	6
• C. Yuan, T. Lu, M. Druzdzel, <i>Annealed MAP</i>	
◆ Decision-theoretic planning (MDPs, etc.)	7
• T. Smith, R. Simmons, <i>Heuristic search value iterations for POMDPs</i>	
◆ Causality	2
• J. Tian, <i>Identifying conditional causal effects</i>	
◆ Computation in Graphical Models	9
• V. Goghate, R. Dechter, <i>A complete anytime algorithm for treewidth</i>	
◆ Applications	9
• L. Kuncheva, C. Whitaker, P. Cockcroft, Z. Hoare <i>Pre-Selection of Independent Binary Features: An Application to Diagnosing Scrapie in Sheep</i>	
◆ Miscellaneous	2

UAI 2004

BEST PAPER AWARD

Case-Factor Diagrams for Structured Probabilistic Modeling

David McAllester

Michael Collins

Fernando Pereira

Max Chickering

Joseph Halpern

UAI 2004 Program Chairs

Case-Factored Diagrams

- ◆ The problem: finding efficient representation of structured probabilistic models.
 - Structured model defines a distribution on structured object, like a sequence or parse tree
- ◆ There are lots of structured statistical models:
 - Markov random fields (MRFs)
 - Probabilistic context-free grammars (PCFGs)
 - Hidden Markov models (HMMs)
 - Conditional random fields (CRFs)
 - Dynamic Bayes nets (DBNs)

Why Another Representation?

- ◆ We are typically interested in $P(y \mid x)$
 - x is the observation; y is the object of interest
 - E.g. x is a word string, y is a parse tree
- ◆ Want a way of representing $P(y \mid x)$ such that inferring the most likely y for a given x is tractable
- ◆ Case-factored diagrams (CFDs) are a new representation for which both the representation and inference problems are often tractable
 - Similar to *binary decision diagrams* (BDDs)
 - Allow a concise representation of problems that factor into independent subproblems
 - Can express the probability distribution over parse trees of sentence of length n with $O(n^3)$ nodes
 - ◆ Better than other formalisms (like PCFGs)

UAI 2005

- ◆ UAI 2005 will be held in Edinburgh, July 26-29
 - Just before IJCAI

See you there!