

POOLE, David

Born: Port Lincoln, South Australia.

Specialties: Artificial intelligence, computational logic, logic and probability, logical representations for decision making under uncertainty, probabilistic inference.

Educated: B.Sc., Flinders Univ. of South Australia (Maths), 1979
Ph.D., Australian National University (Computer Science) 1984
Postdoctoral Fellow, University of Waterloo, 1983-1985

Dissertation: The Theory of CES: A Complete Expert System, Robin Stanton (supervisor)

Regular Academic or Research Appointments: PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, THE UNIVERSITY OF BRITISH COLUMBIA, 1998–; Associate Professor, Department of Computer Science, The University of British Columbia, 1993–1998; Assistant Professor, Department of Computer Science, The University of British Columbia, 1988–1993; Assistant Professor, Department of Computer Science, The University of Waterloo, 1985–1988

Research Profile: Poole is known for his work on assumption-based reasoning, combining logic and probability, and automatic inference algorithms.

In the 1980s, Poole worked on assumption-based reasoning, in particular investigating a simple logical framework with facts written in first order logic and explicit assumables that can be used as long as they are consistent with the facts. This framework can be used for default reasoning when the assumables are normality assumptions used for predictions, and can be used for abduction when the assumptions are used to explain

an observation and a design goal. This framework was implemented in the Theorist program. A common reasoning strategy was: when we make observations, we explain the observations then make predictions from these explanations. Different reasoning frameworks can be obtained by considering who chooses the assumptions: the agent, an adversary or nature.

In the early 1990s Poole developed probabilistic Horn abduction, a simple framework with independent probabilities on assumables and a logic program to give the consequences of the choices. Abduction from observations followed by prediction corresponded to reasoning in Bayesian networks. This showed that a Bayesian network can be interpreted as a deterministic system with (independent) stochastic inputs. It also showed how to extend Bayesian networks to a richer first-order language,

Probabilistic Horn abduction evolved into the independent choice logic which has a richer logic that includes negation as failure, multiple agents choosing assumptions, and various models of time, including the event calculus and the situation calculus.

When probabilistic Horn abduction was developed, the current inference algorithms were too complicated and needed to be extended to cover richer logic. With Nevin Zhang, we developed the variable elimination algorithm for probabilistic inference, and showed how it could be extended to implement causal independence and context-specific independence. I also developed various search algorithms for probabilistic inference.

In the early 2000s Poole developed a framework for lifted probabilistic inference. More recently he has worked on the probability of existence and identity and on combining probabilities with ontologies.

Main Publications:

1. David Poole, "Logical Generative Models for Probabilistic Reasoning about Existence, Roles and Identity", *Proc. Twenty Second AAAI Conference on AI (AAAI-07)*, July 2007.
2. David Poole and Alan Mackworth, "Dimensions of Complexity of Intelligent Agents", *International Symposium on Practical Cognitive Agents and Robots*, Perth, November 2006.
3. Rita Sharma and David Poole, "Probabilistic Reasoning with Hierarchically Structured Variables", *Proc. Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*, Edinburgh, August 2005.
4. David Poole and Clinton Smyth, "Type Uncertainty in Ontologically-Grounded Qualitative Probabilistic Matching", *Eighth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2005)*, Barcelona, July 2005.

5. Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger Hoos, and David Poole "CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements", *Journal of AI Research*, Volume 21, pages 135–191, February 2004.
 6. David Poole, First-order probabilistic inference, *Proc. Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acapulco, August 2003, 985–991.
 7. David Poole and Nevin Lianwen Zhang, "Exploiting contextual independence in probabilistic inference", *Journal of Artificial Intelligence Research*, 18, 263-313, 2003.
 8. D. Poole, "Abducting Through Negation as Failure: Stable models within the independent choice logic", *Journal of Logic Programming*, special issue on Abductive Logic Programming, 44, 5–35, 2000.
 9. D. Poole, "Learning, Bayesian Probability, Graphical Models, and Abduction", in P.A. Flach and A.C. Kakas (Eds.), *Abduction and Induction: Essays on their Relation and Integration*, Kluwer, 2000, 153–168.
 10. M. Horsch and D. Poole, "Estimating the Value of Computation", *Proc. Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI-99)*, Stockholm, Sweden, pages 297–304, July 1999.
 11. D. Poole, "Decision Theory, the Situation Calculus and Conditional Plans", *Linköping Electronic Articles in Computer and Information Science*, Vol 3 (1998):nr 8. <http://www.ep.liu.se/ea/cis/1998/008/>, June 15, 1998. *The Electronic Transactions on Artificial Intelligence*, Volume 2, 105–154, 1998.
 12. D. Poole, A. Mackworth, and R. Goebel, *Computational Intelligence: A Logical Approach*, Oxford University Press, January 1998 (556 pages).
 13. D. Poole, "Probabilistic Partial Evaluation: Exploiting rule structure in probabilistic inference", *Proc. Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, Nagoya, Japan, pp. 1284–1291, 1997.
 14. D. Poole, "The Independent Choice Logic for modelling multiple agents under uncertainty", *Artificial Intelligence*, special issue on Economic Principles of Multi-Agent Systems, 94, 7–56, 1997.
 15. D. Poole, "Probabilistic conflicts in a search algorithm for estimating posterior probabilities in Bayesian networks", *Artificial Intelligence*, 88, 69–100, 1996.
 16. N.L. Zhang and D. Poole, "Exploiting Causal Independence in Bayesian Network Inference", *Journal of Artificial Intelligence Research*, 5, 301–328, 1996.
 17. D. Poole, "Probabilistic Horn abduction and Bayesian networks", *Artificial Intelligence*, 64(1), 81–129, 1993.
 18. D. Poole, "The effect of knowledge on belief: conditioning, specificity and the lottery paradox in default reasoning", *Artificial Intelligence*, 49, 281-307, 1991. Republished in R. J. Brachman, H. J. Levesque and R. Reiter (Eds.), *Knowledge Representation*, MIT Press, 1991.
 19. D. Poole, "A methodology for using a default and abductive reasoning system", *International Journal of Intelligent Systems*, 1990, 5(5), 521–548, 1990.
 20. D. Poole, "Explanation and Prediction: An Architecture for Default and Abductive Reasoning", *Computational Intelligence* 5(2), 97-110, 1989.
 21. D. L. Poole, "A Logical Framework for Default Reasoning", *Artificial Intelligence*, 36(1), 27–47, 1988. *Work in Progress*
 22. He is currently working on semantic science: the idea that data and theories can be published referring to formal ontologies. (Probabilistic) theories can then be tested on all available data and used on prediction on new cases. This involves developing theories that can be specified without knowing the individuals, and where we have to reason about existence and identity.
- Service to the Profession:** Poole was a program chair of the Uncertainty in AI (UAI) conference in 1994, and the general chair in 1995. He has served as the secretary of the Association for uncertainty in AI from 2004–2009.
- Associate editor for the *Journal of Artificial Intelligence Research (JAIR)* 2000-2003, *Artificial Intelligence Journal (AIJ)* 2007-2011, and the *International Journal of Approximate Reasoning* 2005–. Member of the editorial boards of *AAAI press* (2006-2009), *JAIR* (1997-1999), *New Generation Computing* (1991-1997).
- Elected Member of the American Association for Artificial Intelligence Executive Council (2000–2003)
- He has been on program committees for *IJCAI*, *UAI*, *AAAI*, *KR*, *ECSQARU*, *CAI*, *ECAI*, and numerous workshops.
- Teaching:** Poole has graduated 12 Ph.D. students, 12 M.Sc students, and has supervised or co-supervised 9 postdoctoral fellows.
- He was the coauthor of an AI textbook, *Computational Intelligence: A Logical Approach*, published by Oxford University Press, 1998. He is a co-developer of "CISpace: tools for learning computational intelligence", a set of interactive tools designed to learn the fundamental of AI. He wrote *cilog*, a logic programming language designed for teaching that allows for interactive exploration of proofs and search strategies, and includes ask the user mechanisms and probabilistic reasoning.
- Vision Statement:** I believe that practical reasoning is decision making, and that decision making is best modelled in terms of probabilities and utilities. Following the decision theory and game theory traditions, reasoning them becomes finding actions that maximize expected utility (taking into account other agents' reasoning when there are multiple agents). In the past most decision-theoretic reasoning was carried out with simple, essentially propositional, languages.

Logic gives us much richer languages in terms of individuals and relations. Individuals can be referred to by naming them in constant, indirectly using function symbols, and they can be quantified. I believe that using richer first-order languages with probabilities and utilities can capture the the language of thought or mentales that logicians are searching for.

The main arguments for basing reasoning on probabilities and utilities are

- (1) probabilities are what you get from data (by observing and interacting with the world),
- (2) acting is gambling, and if an agent doesn't use probabilities and utilities in gambling it will lose to one that does, and
- (3) there is a well-defined principle for approximate reasoning: approximations can lead to a loss of utility, which lets us trade off thinking and acting.

Extending the language of probability to be richer is a major challenge (particularly when we need to be able to reason with the language and learn the representations).

Honours and Awards: Flinders University of South Australia University Medal 1978. Scholar Canadian Institute for Advanced Research 1992-1995. C. A. McDowell Medal for Excellence in Research U.B.C. 1994. Fellow of the Association for the Advancement of Artificial Intelligence (AAAI) 2000.

PRADE, Henri

Specialties: Artificial intelligence, applied logic, uncertainty modeling.

Born: 20 August 1953 in Mulhouse, France.

Educated: Engineer, and then Doctor-Engineer degrees from Ecole Nationale Supérieure de l'Aéronautique et de l'Espace (Toulouse, France) in 1975 and in 1977; "Doctorat d'Etat" (1982) and "Habilitation † Diriger des Recherches" (1986) both from Paul Sabatier University in Toulouse.

Dissertation: (Doctorat d'Etat): "Modèles Mathématiques de l'Imprécis et de l'Incertain en Vue d'Applications au Raisonnement Naturel".

Regular Academic or Research Appointments: "DIRECTEUR DE RECHERCHE" 1ST CLASS, AT C.N.R.S., IRIT ("INSTITUT DE RECHERCHE EN INFORMATIQUE DE TOULOUSE", TOULOUSE, FRANCE), 2002; "Directeur de Recherche", 2nd class, at C.N.R.S.,

IRIT, 1988-2002; "ChargÉ de Recherche" at C.N.R.S., Lab. LSI ("Langages et Systèmes Informatiques"), Toulouse, France, 1982-1988; "AttachÉ de Recherche" at C.N.R.S., Lab. LSI, 1979-1982.

Visiting Academic or Research Appointments: IRIA (Institut de Recherche en Informatique et Automatique) Post-doctoral Fellow, visiting scholar at Artificial Intelligence Lab., Stanford University, 1977-1978.

Research Profile: Henri Prade has been working on the handling of uncertainty in artificial intelligence for about twenty-five years, and more particularly on the modeling of different types of reasoning (reasoning under uncertainty, nonmonotonic reasoning and belief revision, similarity-based reasoning, reasoning under inconsistency). Jointly with Didier Dubois, he has contributed to the development of fuzzy sets and possibility theory (fuzzy interval analysis, typology of fuzzy if-then rules, study of possibilistic independence, relations between possibility and probability theories, ...), and their applications to approximate reasoning, since the late seventies. Possibility theory, initiated by Lotfi Zadeh in 1978, provides a framework for representing uncertainty in terms of a pair of dual measures of possibility and necessity (their duality expresses a graded version of the classical relationship between the modalities 'possibly' and 'necessarily'). Possibility measures are max-decomposable under disjunction, and are usually valued on the real unit interval, but any discrete linearly ordered scale may be used as valuation scale.

Regarding logic, Henri Prade (jointly with Didier Dubois) has mainly made three contributions: Developing possibilistic logic, relating logic of conditional objects with nonmonotonic reasoning, and introducing bipolarity in logic.

Possibilistic logic (whose first elements appear in Prade's French Doctorat d'Etat thesis in 1982, and whose name was coined in a 1987 paper) is a weighted logic that has been introduced for dealing with uncertain or prioritized information. Standard possibilistic logic expressions are classical logic formulas associated with weights, interpreted in the framework of possibility theory as lower bounds of necessity degrees. Possibilistic logic handles partial inconsistency by means of an inconsistency level that is associated with any possibilistic base (i.e. a set of possibilistic logic formulae). From a semantic point of view, a possibilistic base is understood as a possibility distribution representing the fuzzy set of models of the base. An interpretation is all the less possible as it falsifies formulae of higher degree. Soundness