Toward User-Centric Recommender Systems

Craig Boutilier
Google, Inc.
Mountain View, CA, 94043
cboutilier@google.com

ABSTRACT

Artificial intelligence and machine learning technologies continue to broaden and influence our access to information, entertainment, products and services—and each other—through data-driven recommendations. While the increased access afforded by AI has undoubtedly improved certain aspects of social welfare, the ability of recommenders to generate genuinely personalized recommendations and engage users in meaningful ways remains limited. Furthermore, our understanding of how AI recommenders shape long-term user behavior is poorly understood. In this talk, I will discuss the role that various AI techniques have to play in next-generation, user-centric recommender systems. Among these are preference modeling and preference elicitation; reinforcement learning and latent state models; behavioral decision theory and economics; and modeling of user behavioral preferences. I will also highlight challenges that emerge when putting these methods into practice.

KEYWORDS

recommender systems; preferences; user behaviour

ACM Reference Format:
Craig Boutilier. 2018. Toward User-Centric Recommender Systems. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 1 pages.

Short Bio

Craig Boutilier is a Principal Scientist at Google. He received his Ph.D. in Computer Science from the University of Toronto in 1992, and worked as an Assistant and Associate Professor at the University of British Columbia from 1991–1999. He was a Professor of Computer Science at the University of Toronto from 1999–2017 and served as Chair of the Department of Computer Science at Toronto from 2004–2010. He was co-founder of Granata Decision Systems from 2012–2015, a consulting professor at Stanford University from 1998–2000, a visiting professor at Brown University in 1998, a visiting professor at Carnegie Mellon University in 2008–09, and at Université Paris-Dauphine in 2011. He also served on the Technical Advisory Board of CombineNet, Inc. from 2001 to 2010.

He has published over 200 refereed articles covering topics ranging from knowledge representation, belief revision, default reasoning, and philosophical logic, to probabilistic reasoning, decision making under uncertainty, multi-agent systems, and machine learning. His current research efforts focus on various aspects of decision making under uncertainty: preference elicitation, mechanism design, game theory and multi-agent decision processes, economic models, social choice, computational advertising, Markov decision processes and reinforcement learning.

Boutilier was Program Chair for both the 16th Conference on Uncertainty in Artificial Intelligence (UAI-2000) and the 21st International Joint Conference on Artificial Intelligence (IJCAI-09). He served as Editor-in-Chief of the Journal of Artificial Intelligence Research (JAIR). In addition, he serves or has served on nine editorial boards, over 60 international conference program committees, and 25 international workshop program committees.

Boutilier is a Fellow of the Royal Society of Canada (FRSC), the Association for Computing Machinery (ACM) and the Association for the Advancement of Artificial Intelligence (AAAI). He was awarded the Tier I Canada Research Chair in Adaptive Decision Making for Intelligent Systems, the Izaak Walton Killam Research Fellowship, an IBM Faculty Award, and the NSERC Discovery Accelerator Supplement (DAS). He also received the Killam Teaching Award from the University of British Columbia.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.