

November, 2024

Subbarao Kambhampati

Department of Computer Science and Engineering
School of Computing & AI
Arizona State University, Tempe, AZ 85287
Office Phone: (480) 965-0113
Email: rao@asu.edu *WWW:* <http://rakaposhi.eas.asu.edu/rao.html>

Research Interests

- **Human-Aware AI Systems:** Explainable human-AI Interaction. Planning and Decision-making for human-in-the-loop AI systems. Human-robot teaming. Proactive decision support. Learnable planning models & Model-Lite Planning. Explicable behavior and Explanations. Human-factors evaluations.
- **Automated Planning (AI):** Planning & Reasoning capabilities of Large Language Models (LLMs). Plan Synthesis in metric, temporal, partially accessible and stochastic worlds. Heuristics for plan synthesis. Multi-objective optimization for planning. Reasoning with expressive actions. Scheduling. Speedup-learning to assist planners. Constraint Satisfaction and Operations Research Techniques. Applications of planning to automated manufacturing and space autonomy.
- **Social Media Analysis & Information Integration:** Analysis of human-behavior on social-media platforms. Adaptive techniques for query optimization and execution in information integration. Source discovery and source meta-data learning.

Education

Ph.D. in Computer Science (1989) *University of Maryland, College Park, USA*

M.S. in Computer Science (1985) *University of Maryland, College Park, USA*

B.TECH. in Electrical Engineering and Electronics (1983) *Indian Institute of Technology, Madras, India*

Professional Experience (Academic)

8/00 - present: Professor, Department of Computer Science and Engineering, School of Computing & AI, Arizona State University, Tempe

2/2006 -3/2006: Visiting Researcher, Australian National University (Sabbatical Visit).

2/98 -5/98: Visiting Professor, Dept. of Computer Science and Automation, Indian Institute of Science, Bangalore, India. (Sabbatical Visit)

11/97 -1/98: Visiting Professor, Dept. of Computer Science, Indian Institute of Technology, Kharagpur, India. (Sabbatical Visit)

8/96 -8/00: Associate Professor, Department of Computer Science and Engineering, Arizona State University, Tempe

8/91 - 8/96: Assistant Professor, Department of Computer Science and Engineering, Arizona State University, Tempe

10/89 - 8/91: Research Associate, Center for Design Research and Department of Computer Science, Stanford University, Stanford

6/84 - 8/89: Research Assistant, Center for Automation Research, Department of Computer Science, University of Maryland, College Park

6/88 - 8/88: Part-time Instructor, University College, University of Maryland, College Park

8/85 - 12/86: Part-time Teaching Assistant, University College, University of Maryland, College Park

1/84 - 6/84: Teaching Assistant, Department of Computer Science, University of Maryland, College Park

8/83 - 1/84: Research Assistant, Department of Electrical and Computer Engineering, Carnegie-Mellon University, Pittsburgh

Professional Experience (Industrial)

9/18 - present: Chief AI Officer, AI Foundation.

1/14 - present: Consultant, (incl. DiDi Research, SIFT LLC, Johnson Controls, Lockheed Martin)

Honors & Significant Service

- AAAI/EAAI Patrick Henry Winston Outstanding Educator Award, 2025.
- President, Association for the Advancement of Artificial Intelligence (AAAI), 2016-18. (President-Elect during 2014-16; Past-President 2018-20).
- Fellow, Association for the Advancement of Artificial Intelligence (AAAI), 2004.
- Fellow, American Association for the Advancement of Science (AAAS), 2017.
- Fellow, Association for Computing Machinery (ACM), 2019.
- Distinguished Alumnus, Indian Institute of Technology, Madras (Chennai). 2022
- Distinguished Alumnus, Dept. of Computer Science, University of Maryland, College Park. 2018.
- Chair, AAAS Section T (Information, Communication & computation), 2022-24.
- Privacy Fellow, Innovators Network Foundation. 2019-20; 2020-21; 2021-22.
- Program Chair, Intl. Joint Conf. on AI (IJCAI), 2016 (Trustee, IJCAI, 2013-18).
- Program Co-Chair, AAAI Conference on Artificial Intelligence, 2005.
- Member, Board of Directors, Partnership for Artificial Intelligence, 2017-2018.
- National Science Foundation Young Investigator (NYI) award, 1994-1999.
- Influential Paper Award (Honorable Mention), ICAPS 2010 (for a 2000 paper on viewing planning as Constraint Satisfaction with Minh B. Do).
- Best Researcher Award (Senior Faculty), ASU CIDSE, 2017

- Selected by ASU CSE students as the Teacher of the Year. 2012.
- Selected by ASU students for 16th Annual ASU Last Lecture Series, 2011 (lectured on "You can't do that HAL! Collateral Lessons from a Computational Quest to Design HAL").
- College of Engineering Teaching Excellence Award, 2002.
- Best Poster Award, World Wide Web Conference, 2010 (for the SourceRank work with Raju Balakrishnan).
- Conference Committee Chair, AAI 2012-2014.
- Program Co-Chair, Intl. Conference on Planning and Scheduling, 2013.
- J.P. Morgan AI Faculty Research Award, 2019; 2021.
- IBM Faculty Award, 2004.
- Google Research Grant. 2016; 2013; 2010; 2008.
- Amazon Research Award. 2024.
- National Science Foundation Research Initiation (RIA) award, 1992-1995.
- Track Co-Chair, AI and Web Track, AAI 2010.
- Elected to AAI Executive Council, 2009-2012.
- Distinguished Performance Award (Simple Preferences Track for Yochan-ps), International Planning Competition, 2006.
- Researcher of the year award. School of Computing and Informatics. Arizona State University. 2005.
- AAI Presidential Address (on Challenges of Human-Aware AI). AAI 2018. New Orleans. February 2018.
- AAMAS Keynote Talk (on Synthesizing Explainable Behavior for Human-AI Collaboration). AAMAS 2019. Montreal. May 2019.
- Presidential Distinguished Lecture (on Challenges of Human-Aware AI). Singapore Management University. April 2019.
- Invited Speaker, ACAI Summer School on Automated Planning, Freiburg, Germany. 2011.
- Invited Speaker, International Conference on Automated Planning and Scheduling, Trento, Italy. June 2003.
- Program Co-Chair, Intl. Conference on Planning and Scheduling (formerly known as AI Planning Systems Conference), 2000.
- Invited Speaker, Machine Learning summer school. (Canberra, Australia). February 2006.
- Invited Speaker, PLANET summer school on Planning. (Paphos, Cyprus). September 2000.
- Invited Speaker, National Conference on Artificial Intelligence, 1996 (Portland)
- Invited Speaker, 4th Intl. Conference on Knowledge-Based Computer Systems, Mumbai, India. December 2000.
- Nominated for "Outstanding Graduate Mentor Award"--an ASU program for recognizing Graduate Student Mentoring (2010).
- Arizona State University Nominee for NSF Presidential Faculty Fellowship, 1995.
- Nominated three times for ASU CEAS Teaching Excellence Award, 1995, 1998, (2002).

- Samuel N. Alexander A.C.M. Fellowship Grant, 1988 (Awarded by Washington D.C. Chapter of A.C.M. to a Computer Science doctoral student from A.C.M. Mid-Atlantic Region)
- Fellow, Systems Research Center, University of Maryland, 1987
- Institute Merit Certificate and Prize for securing First Rank in 2nd year of B.TECH., Indian Institute of Technology, Madras, India, 1980.
- State Special Merit Scholarship for Rank-holders, A.P., India, 1976-78, 1978-83
- Merit Scholarship, Indian Institute of Technology, Madras, India, 1978-83
- State Board of Education Merit Certificates for securing 11th Rank in state in class XII and 7th Rank in state in class X, A.P., India

Significant Student Honors

- **IEEE AI 10 to Watch** (2020) recognition for Tathagata Chakraborti (Ph.D. 2018).
- **Dean's Dissertation Award** (2024) recognition for Zahra Zahedi (Ph.D. 2023).
- **ICAPS Distinguished Dissertation. Runner-up** (2023) for Sarath Sreedharan (Ph.D. 2022).
- **Dean's Dissertation Award** (2022) recognition for Sarath Sreedharan (Ph.D. 2022).
- **ICAPS Distinguished Dissertation. Runner-up**, for Anagha Kulkarni. 2022.
- HRI Pioneer recognition for Zahra Zahedi. 2024.
- Best paper award at Bridging Planing & Reinforcement Learning workshop at ICAPS 2022 (with students Lin Guan and Sarath Sreedharan).
- ASU CIDSE Outstanding M.S. Student Award to Niharika Jain. 2021.
- People's Choice Award for the Best Demonstration, Intl. Conf. on Planning and Scheduling, 2020 (For Sarath Sreedharan and IBM co-authors).
- Outstanding PC Member (1 of 12) for AAAI 2020 for Sarath Sreedharan. 2020.
- **ICAPS Distinguished Dissertation, Runner-up**, for Tathagata Chakraborti. 2019.
- ASU CIDSE Outstanding Ph.D. Student Award to Tathagata Chakraborti. 2019.
- IBM Ph.D. Fellowship for Sailik Sengupta. 2018-19.
- IBM Ph.D. Fellowship for Tathagata Chakraborty. 2017-18.
- ASU CIDSE Outstanding Masters Student Award to Sarath Sridharan. 2017.
- Students reached Microsoft Imagine Cup National Finals with an entry on human-robot teaming. 2017.
- Nominted for AAMAS 2016 Best Student Paper Award, 2016 (for work with Xin Tian, Hankz Hankui Zhuo)
- IBM Ph.D. Fellowship for Tathagata Chakraborty. 2016-17.
- People's Choice Award for the Best Demonstration, Intl. Conf. on Planning and Scheduling, 2014 (for Lydia Manikonda, Tathagata Chakraborti and Kartik Talamadupula).
- School of Computing & Informatics Outstanding Masters Student Award to Manish Jha, 2012.
- Best Poster Award, World Wide Web Conference, 2010 (for the SourceRank work with Raju Balakrishnan).

- **Influential Paper Award (Honorable Mention)**, ICAPS 2010 (for a 2000 paper on viewing planning as Constraint Satisfaction with Minh B. Do).
- International Conference on Planning & Scheduling **Best Dissertation Award** for Daniel Bryce; 2009.
- International Conference on Planning & Scheduling **Best Dissertation Award honorable mention** for Menkes van den Briel; 2009.
- Computing Research Association Outstanding Undergraduate Student (honorable mention) to Kartik Talamadupula; 2008
- **Yahoo! Key Scientific Challenges Award to Raju Balakrishnan, 2009**
- International Planning Competition Distinguished Performance Award to J. Benton, 2006.
- School of Computing & Informatics Distinguished Senior Award to Kartik Talamadupula, 2008
- School of Computing & Informatics Outstanding Masters Student Award to Garrett Wolf, 2009.
- School of Computing & Informatics Outstanding Masters Student Award to Hemal Khatri, 2007.
- School of Computing & Informatics Outstanding Masters Student Award to Thomas Hernandez, 2005.
- School of Computing & Informatics Outstanding Ph.D. Student Award to Minh Binh Do, 2005.
- Shortlisted for the best paper award at ICAPS 2006 (Daniel Bryce)
- Shortlisted for the best student paper award at AAAI 2005 (William Cushing and Daniel Bryce)
- Two Science Foundation of Arizona Fellows (Tuan Nguyen, Kartik Talamadupula); two ARCS Foundation Fellows (Daniel Bryce, J. Benton); one Fulton Fellow (Will Cushing), one University Graduate Scholar (Will Cushing).

Full Publication List

Most of these papers are available online at <http://rakaposhi.eas.asu.edu/papers.html>. According to Google Scholar (<http://scholar.google.com/citations?user=y13L07sAAAAJ>), these papers have received over **18,696** citations to date, with an **h-index** of **68**, and **i-10 index** of **279**.

Note: In Computer Science, unlike many other fields, conference papers are rigorously reviewed, with top conferences having low acceptance rates and high citation impacts; publications in these conferences are considered archival and comparable to journal papers.

Journal articles:

- (1) Subbarao Kambhampati. Can Large Language Models Reason and Plan?. The Annals of New York Academy of Sciences. March 2024.
- (2) Sarath Sreedharan, Siddharth Srivastava and Subbarao Kambhampati. Explain It As Simple as Possible, but No Simpler â Explanation via Model Simplification for Addressing Inferential Gap. Artificial Intelligence. 2024.
- (3) Zahra Zahedi, Sailik Sengupta and Subbarao Kambhampati. A Game-Theoretic Model of Trust in Human Robot Teaming: Guiding Human Observation Strategy for Monitoring Robot Behavior. IEEE Transactions on Human-Machine Systems. 2024.
- (4) Sarath Sreedharan, Tathagata Chakraborty, Christian Muise and Subbarao Kambhampati. Planning with mental modelsâBalancing explanations and explicability. Artificial Intelligence. 2024.
- (5) Kebin Jin, Hankz Hankui Zhuo, Zhanhao Xiao, Hai Wan and Subbarao Kambhampati. Gradient-Based Mixed Planning with Symbolic and Numeric Action Parameters Artificial Intelligence. 2022. (DOI).
- (6) Niharika Jain, Alberto Olmo, Sailik Sengupta, Lydia Manikonda, Subbarao Kambhampati. Imperfect ImaGANation: Implications of GANs Exacerbating Biases on Facial Data Augmentation and Snapchat Face Lenses Artificial Intelligence Journal. 2022. (Accepted 12/21)
- (7) Sarath Sreedharan*, Tathagata Chakraborti*, Subbarao Kambhampati Foundations of Explanations as Model Reconciliation Artificial Intelligence Journal. 2021.
- (8) Sarath Sreedharan, Siddharth Srivastava, Subbarao Kambhampati. Using State Abstractions to Compute Personalized Contrastive Explanations for AI Agent Behavior. Artificial Intelligence Journal. 2021.
- (9) Subbarao Kambhampati. Challenges of Human-Aware AI Systems AI Magazine, Vol. 41 No. 3, Fall 2020.
- (10) Sachin Grover, Sailik Sengupta, Tathagata Chakraborti, Aditya Prasad Mishra and Subbarao Kambhampati. RADAR: Automated Task Planning for Proactive Decision Support. HCI Journal (Special Issue on Unifying Human Computer Interaction and Artificial Intelligence), 2020
- (11) Sailik Sengupta*, Ankur Chowdhary*, Abdulhakim Sabur, Dijiang Huang, Adel Alshamrani, Subbarao Kambhampati. A Survey of Moving Target Defenses for

Network Security IEEE Communications Surveys and Tutorials, March 2020

- (12) Hankz Hankui Zhao, Yantian Zha, Subbarao Kambhampati, Xin Tian. Discovering Underlying Plans Based on Shallow Models. *ACM TIST*, 2019.
- (13) Kartik Talamadupula, Grdon Briggs, Matthias Scheutz & Subbarao Kambhampati. Architectural Mechanisms for Handling Human Instructions for Open-World Mixed-Initiative Team Tasks and Goals **Advances in Cognitive Systems**. 2017.
- (14) Tuan Nguyen, Subbarao Kambhampati and Sarath Sreedharan. Robust planning with incomplete domain models. **Artificial Intelligence**. 2017.
- (15) Hankz Hankui Zhuo and Subbarao Kambhampati. Model-Lite Planning: Case-Based vs. Model-Based Approaches. **Artificial Intelligence** 2017.
- (16) Lydia Manikonda, Tathagata Chakraborty, Kartik Talamadupula & Subbarao Kambhampati. Herding the Crowd: Using Automated Planning for Better Crowdsourced Planning. *Human Computation Journal*. 2017
- (17) Sushovan De, Yuheng Hu, Venkata Vamsikrishna Meduri, Yi Chen, Subbarao Kambhampati: BayesWipe: A Scalable Probabilistic Framework for Improving Data Quality. *J. Data and Information Quality* 8(1): 5:1-5:30 (2016)
- (18) Raju Balakrishnan and Subbarao Kambhampati Click Efficiency: A Unified Optimal Ranking for Online Ads and Documents. *Journal of Intelligent Information Systems*. 2015.
- (19) Rohit Raghunathan, Sushovan De, Subbarao Kambhampati. Bayes Networks for Supporting Query Processing Over Incomplete Autonomous Databases *Journal of Intelligent Information Systems*. 42(3): 595-618 (2014)
- (20) Raju Balakrishnan, Subbarao Kambhampati, and Manishkumar Jha. Assessing Relevance and Trust of the Deep Web Sources and Results Based on Inter-Source Agreement *ACM Transactions on Web*. Volume 7, Issue 2, May 2013.
- (21) Nan Li, William Cushing, Subbarao Kambhampati and Sungwook Yoon. Learning Probabilistic Hierarchical Task Networks to Capture User Planning Preferences. *ACM Transactions on Intelligent Systems and Technology*. *ACM TIST* 5(2): 29 (2014)
- (22) Tuan Nguyen, Minh Do, Alfonso Gerevini, Ivan Serina, Biplav Srivastava and Subbarao Kambhampati. Planning with partial preference models *Artificial Intelligence*. Vol 190. Pages 1-31. Oct 2012.
- (23) Ravi Gummadi, Anupam Khulbe, Aravind Kalavagattu, Sanil Salvi, Subbarao Kambhampati. SmartInt: Using Mined Attribute Dependencies to Integrate Fragmented Web Databases *Journal of Intelligent Information Systems*. 2011 (Accepted)
- (24) Xiaoqin Zhan, Sungwook Yoon, and Subbarao Kambhampati and several other authors. An Ensemble Architecture for Learning Complex Problem-Solving Techniques From Demonstration. **ACM Transactions on Intelligent Systems and Technology**. (Accepted). 2011.
- (25) Daniel Bryce, William Cushing and Subbarao Kambhampati. State Agnostic Planning Graphs: Deterministic, Non-Deterministic, and Probabilistic Planning. **Artificial Intelligence**. Volume 175, Issues 3-4, March 2011.
- (26) Kartik Talamadupula, J. Benton, Subbarao Kambhampati, Paul Schermerhorn, and Matthias Scheutz. Planning for Human-Robot Teaming in Open Worlds. *ACM*

Transactions on Intelligent Systems and Technology. (Special Issue on Applications of Automated Planning). Vol 1. No. 2. 2010.

- (27) Garrett Wolf, Aravind Kalavagattu, Hemal Khatri, Raju Balakrishnan, Bhaumik Chokshi, Jianchun Fan, Yi Chen and Subbarao Kambhampati. Query Processing Over Incomplete Autonomous Databases: Query Rewriting Using Learned Data Dependencies. **VLDB Journal**. 18(5): 1167-1190 (2009). (Special issue on uncertain and probabilistic databases.)
- (28) J. Benton, Minh Do and Subbarao Kambhampati. Anytime Heuristic Search for Partial Satisfaction Planning **Artificial Intelligence Journal**. 173(5-6): 562-592 (2009)
- (29) Menkes van den Briel, Thomas Vossen and Subbarao Kambhampati. Loosely coupled formulations for Automated Planning: An Integer Programming Perspective **Journal of Artificial Intelligence Research**. Volume 31, pages 217-257.
- (30) Daniel Bryce, Subbarao Kambhampati and David E. Smith. Sequential Monte Carlo in Reachability Heuristics for Probabilistic Planning **Artificial Intelligence Journal**. Vol. 172. 2008.
- (31) How to skin a planning graph for fun and profit: A Tutorial on Planning Graph Based Reachability Heuristics. Daniel Bryce and Subbarao Kambhampati. ASU CSE TR-06-007, April 2006. To appear in **AI Magazine** Spring 2007.
- (32) Daniel Bryce, Subbarao Kambhampati and David E. Smith. Planning Graph Heuristics for Belief Space Search **Journal of Artificial Intelligence Research**. Volume 26, pages 35-99. May 2006.
- (33) Menkes van den Briel and Subbarao Kambhampati. Optiplan: A Planning System that Unifies Integerprogramming with Planning Graph. (Engineering Note) JAIR special track on IPC 2004. December 2005.
- (34) Zaiqing Nie, Subbarao Kambhampati and Ullas Nambiar. Effectively mining and using coverage and overlap statistics for data integration. **IEEE Transactions on Knowledge and Data Engineering**. 2005.
- (35) Terry Zimmerman and Subbarao Kambhampati. Using memory to transform search on the planning graph **Journal of Artificial Intelligence Research**. 2005.
- (36) Jianchun Fan and Subbarao Kambhampati. A Snapshot of Public Web Services SIGMOD Record, March 2005.
- (37) Thomas Hernandez and Subbarao Kambhampati. Integration of Bioinformatic Sources: Current Approaches and Systems. SIGMOD Record, Vol 33, No 3 September 2004.
- (38) S. Kambhampati, E. Lambrecht, U. Nambiar, Z. Nie and G. Senthil. Optimizing Recursive Information Gathering Plans in EMERAC. **Journal of Intelligent Information Systems**. Vol 22, No. 2. February 2004. pp. 119-153.
- (39) Minh B. Do and Subbarao Kambhampati. Sapa: A Multi-Objective Metric Temporal Planner. **Journal of AI Research**, Volume 20, pages 155-194. December 2003.
- (40) Sanchez, R. and Kambhampati, S. (2003) "AltAltp: Online Parallelization of Plans with Heuristic State Search." **Journal of AI Research**, Volume 19, pages 631-657.
- (41) Terry Zimmerman and Subbarao Kambhampati. Learning-assisted automated planning: Looking back, taking stock, going forward. **AI Magazine**. 24(2):73-96. Summer 2003

- (42) XuanLong Nguyen, Subbarao Kambhampati and Romeo Sanchez Nigenda. Planning Graph as the Basis for deriving Heuristics for Plan Synthesis by State Space and CSP Search. **Artificial Intelligence** 135(1-2): 73-123 (2002)
- (43) Minh B. Do and Subbarao Kambhampati. Planning as Constraint Satisfaction: Solving the planning graph by compiling it into CSP. **Artificial Intelligence** 132(2): 151-182 (2001)
- (44) Biplav Srivastava, Subbarao Kambhampati, Binh Minh Do. Planning the Project Management Way: Efficient Planning by Effective Integration of Causal and Resource Reasoning in RealPlan. **Artificial Intelligence** 132(2): 151-182 (2001)
- (45) S. Kambhampati. Planning Graph as a (Dynamic) CSP: Exploiting EBL, DDB and other CSP Search Strategies in Graphplan. **Journal of Artificial Intelligence**. JAIR 12: 1-34 (2000)
- (46) S. Kambhampati. On the relations between intelligent backtracking and explanation-based learning in planning and constraint satisfaction. **Artificial Intelligence**. Vol. 105, No. 1-2. October 1998.
- (47) B. Srivastava and S. Kambhampati. Synthesizing customized planners from specifications. **Journal of Artificial Intelligence Research**. Vol 8. pp 93-128. 1998. 1998.
- (48) S. Kambhampati. Refinement planning as a unifying framework for plan synthesis. **AI Magazine**. 18(2):67-97. Summer. 1997.
- (49) L. Ihrig and S. Kambhampati. Storing and Indexing Plan Derivations through Explanation-based Analysis of Retrieval Failures. **Journal of Artificial Intelligence**. Vol 7. pp 161-198. 1997.
- (50) S. Kambhampati, S. Katukam and Y. Qu. "Failure driven Dynamic Search Control for Partial Order Planners: An Explanation-based approach", **Artificial Intelligence**. 88(1-2):253-315. 1997.
- (51) S. Kambhampati and D.S. Nau, "On the nature and role of modal truth criteria in planning", **Artificial Intelligence**. 82(1-2):129-156. 1996.
- (52) S. Kambhampati, C. Knoblock and Q. Yang. Planning as Refinement Search: A Unified framework for evaluating design tradeoffs in partial order planning. **Artificial Intelligence**. Special issue on Planning and Scheduling. Vol. 76. No. 1-2, September 1995. pp. 167-238.
- (53) S. Kambhampati. Comparative analysis of Partial Order and HTN Planning. **SIGART Bulletin** Special section on Evaluation of Plans, Planners and Planning Agents. Vol. 6, No. 1, January, 1995. pp. 16-25.
- (54) S. Kambhampati, "Multi-Contributor Causal Structures for Planning: A Formalization and Evaluation," **Artificial Intelligence** Vol. 69, No. 1-2, pp. 235-278. 1994.
- (55) S. Kambhampati and S. Kedar, "A Unified Framework for Explanation-Based Generalization of Partially Ordered and Partially Instantiated Plans," ASU CSE-TR-92-008, **Artificial Intelligence**, Vol 67, No. 2, June 1994. pp. 29-70.
- (56) S. Kambhampati, "Exploiting Causal Structure to Control Retrieval and Refitting during Plan Reuse," **Computational Intelligence**, Vol. 10, No. 2, May 1994, pp 212-244.
- (57) S. Kambhampati, M.R. Cutkosky, J.M. Tenenbaum and S.H. Lee, "Integrating General Purpose Planners and Specialized Reasoners: Case Study of a Hybrid Planning

Architecture,” **IEEE Transactions on Systems, Man and Cybernetics (special issue on Planning, Scheduling and Control),** Vol 23, No. 6, November 1993. pp 1503-1517.

- (58) S. Kambhampati and J.A. Hendler, “A Validation Structure Based Theory of Plan Modification and Reuse,” **Artificial Intelligence**, 55 (2-3), pp. 193-258, June 1992.
- (59) S. Kambhampati and L.S. Davis, “Multiresolution Path Planning for Mobile Robots,” **IEEE Journal of Robotics and Automation RA-2(3)**, September 1986, pp. 135-145.

Commentaries, Reviews and Reports:

- (60) Subabroa Kambhampati. Changing the nature of AI research Communications of the ACM, September 2022. Vol. 65 No. 9. Pages 8-9
- (61) Subbarao Kambhampati. Polanyi’s Revenge & AI’s new romance with tacit knowledge. Communications of the ACM (CACM), February 2021, Vol. 64 No. 2, Pages 31-32.
- (62) Subbarao Kambhampati, Craig A. Knoblock: Guest Editors’ Introduction: Information Integration on the Web. **IEEE Intelligent Systems** 18(5): 14-15 (2003)
- (63) A. Naryek and S. Kambhampati. Research Issues at the Intersection of Planning and Constraint Programming. **Constraints Journal**. Special Issue on Planning. Vol 8. No. 4. pp. 335-338. October 2003.
- (64) S. Kambhampati and C. Knoblock. Information Integration on the Web: A view from AI and Databases (Report on IIWeb-03). **SIGMOD Record**. December 2003.
- (65) B.Srivastava, X. Long, M. Do, U. Nambiar, X. Nie, R. Nigenda, T. Zimmerman and S. Kambhampati. AltAlt: Combining Graphplan and Heuristic State Search B. Srivastava, X. Long, Minh B. Do, U. Nambiar, X. Nie, R. Nigenda, T. Zimmerman and S. Kambhampati **AI Magazine**, 2001.
- (66) S. Kambhampati. “Theoretical Contributions of Artificial Intelligence.” Side bar. **IEEE Computer Magazine**. 50th anniversary issue. Fall 1996.
- (67) S. Kambhampati. “Refinement planning: Status and Prospectus” *In Proc. National Conference on AI*, 1996. (Invited)
- (68) S. Kambhampati. “AI Planning: A prospectus on theory and applications” Position Statement. **ACM Computing Surveys**, Symposium on Artificial Intelligence, September 1995.
- (69) G. Olsen, M.R. Cutkosky and S. Kambhampati, “Real Physics for Real Engineers: A Response to “Prolegomena to Any Future Qualitative Physics,”” **Computational Intelligence**, Vol. 8, No. 2, pp. 286-288. September, 1992.
- (70) S. Kambhampati, "Report on the AAAI 1992 Spring Symposium on Computational Considerations in Supporting Incremental Modification and Reuse," **AI Magazine**, Vol. 13, No. 3, pp. 24-25.
- (71) A. Lansky, M. Drummond, S. Kambhampati, E. Pednault and Q. Yang. “Report on the AAAI 1993 Spring Symposium on Foundations of Automatic Planning, the classical approach and beyond,” **AI Magazine**, Vol 14, No. 3. pp. 34-35.

Book Publishing (Books, Edited Proceedings, Book chapters):

- (72) Sarath Sreedharan, Anagha Kulkarni and Subbarao Kambhampati. Explainable Human-AI Interaction: A Planning Perspective. Morgan & Claypool. January 2022. ISBN: 978-1-63639-289-9.
- (73) Subbarao Kambhampati (ed). Proceedings of the 25th International Joint Conference on Artificial Intelligence, 2016. (IJCAI 2016)
- (74) Daniel Borrajo, Subbarao Kambhampati, Angelo Oddi, Simone Fratini (Eds.): Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling, ICAPS 2013, Rome, Italy, June 10-14, 2013. AAAI 2013, ISBN 978-1-57735-609-7.
- (75) Subbarao Kambhampati and Sungwook Yoon. Explanation-Based Learning for Planning. Encyclopedia of Machine Learning. Springer-Verlag, New York, NY (Claude Sammut, ed.)
- (76) M. Veloso and S. Kambhampati (eds). Proceedings of National Conference on Artificial Intelligence, 2005. (AAAI 2005).
- (77) S. Kambhampati and C. Knoblock. (eds). Proceedings of IJCAI 2003 Workshop on Intelligent Information Integration. 2003.
- (78) S. Kambhampati, C. Knoblock and S. Chien. (eds) Proceedings of AIPS 2002 (International Conference on AI Planning Systems). 2000.
- (79) T. Dean and S. Kambhampati, "Planning and Scheduling," in "*The CRC Handbook of Computer Science and Engineering*," CRC Press. 1997.
- (80) S. Kambhampati, "Supporting Flexible Reuse of Plans, in "*Machine Learning Methods for Planning*," chapter 12, Editor: S. Minton, Morgan Kaufmann, Palo Alto, 1992 (ISBN 1-55860-248-8). pp. 397-434.
- (81) S. Kambhampati (ed), Working Notes of AAAI 1992 Spring Symposium on "Computational Considerations in Supporting Incremental Modification and Reuse," March 25-27, 1992, Stanford University.

Formally Refereed Conference Publications:

Conferences with very rigorous review criteria

- (82) Kaya Stechly, Karthik Valmeekam, Subbarao Kambhampati. Chain of Thoughtlessness: An Analysis of CoT in Planning NeurIPS 2024. (arXiv: <https://arxiv.org/abs/2405.04776>).
- (83) Lin Guan, Yifan Zhou, Denis Liu, Yantian Zha, Heni Ben Amor and Subbarao Kambhampati. Task Success is not Enough: Investigating the Use of Video-Language Models as Behavior Critics for Catching Undesirable Agent Behaviors. 1st Conference on Language Models (COLM). October 2024.
- (84) Subbarao Kambhampati, Karthik Valmeekam, Lin Guan, Kaya Stechly, Mudit Verma, Siddhant Bhambri, Lucas Saldyt, Anil Murthy. LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks. ICML 2024. (Spotlight)
- (85) Mudit Verma*, Siddhant Bhambri*, & Subbarao Kambhampati. Theory of Mind abilities of Large Language Models in Human-Robot Interaction : An Illusion? HRI 2024

- (86) Upasana Biswas, Lin Guan & Subbarao Kambhampati. On the Pitfalls of Learning to Cooperate with Self Play Agents Checkpointed to Capture Humans of Diverse Skill Levels. HRI 2024 (Late Breaking Paper)
- (87) Zahra Zahedi, Sailik Sengupta & Subbarao Kambhampati. Why didn't you allocate this task to them?' Negotiation-Aware Task Allocation and Contrastive Explanation Generation AAAI 2024.
- (88) Yantian Zha, Lin Guan & Subbarao Kambhampati. Learning from Ambiguous Demonstrations with Self-Explanation Guided Reinforcement Learning AAAI 2024.
- (89) Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, Subbarao Kambhampati. On the Planning Abilities of Large Language Models -- A Critical Investigation NeurIPS 2023. (Accepted as **Spotlight**)
- (90) Lin Guan*, Karthik Valmeekam*, Sarath Sreedharan, Subbarao Kambhampati. Leveraging Pre-trained Large Language Models to Construct and Utilize World Models for Model-based Task Planning NeurIPS 2023.
- (91) Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, Subbarao Kambhampati. PlanBench: An Extensible Benchmark for Evaluating Large Language Models on Planning and Reasoning about Change NeurIPS 2023. (Benchmark Track).
- (92) Sarath Sreedharan, Christian Muise, Subbarao Kambhampati. Generalizing Action Justification and Causal Links to Policies ICAPS 2023.
- (93) Lin Guan, Karthik Valmeekam & Subbarao Kambhampati. Relative Behavioral Attributes: Filling the Gap between Symbolic Goal Specification and Reward Learning from Human Preferences ICLR 2023.
- (94) Zahra Zahedi, Mudit Verma, Sarath Sreedharan & Subbarao Kambhampati. Trust-Aware Planning: Modeling Trust Evolution in Iterated Human-Robot Interaction HRI 2023.
- (95) Zahra Zahedi, Sailik Sengupta & Subbarao Kambhampati. Why didn't you allocate this task to them?' Negotiation-Aware Explicable Task Allocation and Contrastive Explanation Generation AAMAS 2023 (Extended Abstract)
- (96) Tung Thai, Utkarsh Soni, Mudit Verma, Sriram Gopalakrishnan, Ming Shen, Mayank Garg, Ayush Kalani, Nakul Vaidya, Subbarao Kambhampati, Neeraj Varshney, Chitta Baral, Jivko Sinapov, Matthias Scheutz. Methods and Mechanisms for Interactive Novelty Handling in Adversarial Environments AAMAS 2023. (Extended Abstract)
- (97) Lin Guan, Sarath Sreedharan, Subbarao Kambhampati. Leveraging Approximate Symbolic Models for Reinforcement Learning via Skill Diversity. ICML 2022. (Also in RLDM workshop 2022)
- (98) Sarath Sreedharan, Pascal Bercher & Subbarao Kambhampati. On the Computational Complexity of Model Reconciliation IJCAI 2022.
- (99) Siddhant Bhambri, Purv Chuhan, Frederico Arujo, Adam Doupe and Subbarao Kambhampati. Using Deception in Markov Game to Understand Adversarial Behaviors through a Capture-The-Flag Environment. GameSec 2022.
- (100) Karthik Valmeekam, Sarath Sreedharan, Sailik Sengupta & Subbarao Kambhampati. RADAR-X: An Interactive Mixed Initiative Planning Interface Pairing Contrastive Explanations and Revised Plan Suggestions ICAPS 2022.

- (101) Symbols as a Lingua Franca for Bridging Human-AI Chasm for Explainable and Advisable AI Systems Subbarao Kambhampati, Sarath Sreedharan, Mudit Verma, Yantian Zha, Lin Guan AAAI 2022 (Blue Sky Track)
- (102) Sriram Gopalakrishnan & Subbarao Kambhampati. Minimizing Robot Navigation Graph For Position-Based Predictability By Humans AAMAS 2022 (Extended Abstract)
- (103) Zahra Zahedi, Sarath Sreedharan, Mudit Verma and Subbarao Kambhampati. Modeling the interplay between human trust and monitoring HRI 2022 (Late breaking paper)(also at XAIP-ICAPS'20)
- (104) Lin Guan, Mudit Verma, Sihang Guo, Ruohan Zhang, Subbarao Kambhampati. Widening the Pipeline in Human-Guided Reinforcement Learning with Explanation and Context-Aware Data Augmentation. NeurIPS 2021 (**Spotlight Presentation**).
- (105) Utkarsh Soni, Sarath Sreedharan, Subbarao Kambhampati. Not all users are the same: Providing personalized explanations for sequential decision making problems IROS, 2021.
- (106) Sarath Sreedharan, Anagha Kulkarni, David E. Smith, Subbarao Kambhampati. A Unifying Bayesian Formulation of Measures of Interpretability in Human-AI. IJCAI 2021.
- (107) Valmeekam Karhik, Sarath Sreedharan, Sailik Sengupta, Subbarao Kambhampati RADAR-X: An Interactive Interface Pairing Contrastive Explanations with Revised Plan Suggestions AAAI'21 Demonstrations Program (Demo Paper)
- (108) Frederico Araujo, Sailik Sengupta, Jiyong Jang, Adam Doupa©, Kevin Hamlen and Subbarao Kambhampati. Software Deception Steering through Version Emulation. HICSS, 2021
- (109) Tathagata Chakraborti, Sarath Sreedharan and Subbarao Kambhampati. The Emerging Landscape of Explainable AI Planning and Decision Making. IJCAI, 2020.
- (110) Sailik Sengupta, Kaustav Basu, Arunabha Sen and Subbarao Kambhampati. Moving Target Defense for Robust Monitoring of Electric Grid Transformers in Adversarial Environments. GameSec 2020.
- (111) Anagha Kulkarni*, Sarath Sreedharan*, Sarah Keren, Tathagata. Chakraborti, David Smith and Subbarao Kambhampati. Designing Environments Conducive to Interpretable Robot Behavior IROS, 2020
- (112) Sarath Sreedharan, Siddharth Srivastava and Subbarao Kambhampati. TLdR: Policy Summarization for Factored SSP Problems Using Temporal Abstractions. ICAPS 2020.
- (113) Sarath Sreedharan, Tathagata Chakraborti, Christian Muise, Yasaman Khazaeni and Subbarao Kambhampati. D3WA+: A Case Study of XAIP in a Model Acquisition Task. ICAPS 2020.
- (114) Sarath Sreedharan, Tathagata Chakraborti, Christian Muise and Subbarao Kambhampati. Expectation-Aware Planning: A Unifying Framework for Synthesizing and Executing Self-Explaining Plans for Human-Aware Planning. AAAI 2020.
- (115) Subbarao Kambhampati Synthesizing Explainable Behavior for Human-AI Collaboration AAMAS. Invited Talk extended abstract. 2019.
- (116) Anagha Kulkarni, Siddharth Srivastava and Subbarao Kambhampati. Signaling Friends and Head-Faking Enemies Simultaneously: Balancing Goal Obfuscation and Goal

Legibility. AAMAS (Extended Abstract), 2020

- (117) Sarath Sreedharan, Alberto Olmo, Aditya Prasad Mishra, Subbarao Kambhampati. Model-Free Model Reconciliation. IJCAI, 2019.
- (118) Sarath Sreedharan, Siddharth Srivastava, David Smith, Subbarao Kambhampati Why Can't You do that, HAL? Explaining Unsolvability of Planning Tasks. IJCAI, 2019. Tathagata Chakraborti*, Sarath Sreedharan*, Subbarao Kambhampati. Balancing Explicability and Explanations for Human-Aware Planning. IJCAI, 2019.
- (119) Sailik Sengupta, Tathagata Chakraborti, Subbarao Kambhampati MTDeep: Moving Target Defense to Boost the Security of Deep Neural Nets Against Adversarial Attacks GameSec. 2019.
- (120) Sailik Sengupta, Ankur Chowdhary, Dijiang Huang, Subbarao Kambhampati General Sum Markov Games for Strategic Detection of Advanced Persistent Threats using Moving Target Defense in Cloud Networks GameSec. 2019.
- (121) Tathagata Chakraborti, Anagha Kulkarni, Sarath Sreedharan, David E. Smith, Subbarao Kambhampati. Explicability? Legibility? Predictability? Transparency? Privacy? Security? The Emerging Landscape of Interpretable Agent Behavior. ICAPS, 2019.
- (122) Anagha Kulkarni, Siddharth Srivastava and Subbarao Kambhampati. A Unified Framework for Planning in Adversarial and Cooperative Environments. AAI, 2019. (Also presented at ICAPS 2018 Workshop on Planning and Robotics.)
- (123) Tathagata Chakraborti*, Sarath Sreedharan*, Sachin Grover, Subbarao Kambhampati. Plan Explanations as Model Reconciliation--An Empirical Study. HRI, 2019.
- (124) Tathagata Chakraborti and Subbarao Kambhampati. (When) Can AI Bots Lie? AAAI/ACM Conference of Artificial Intelligence, Ethics, and Society (AIES), 2019.
- (125) MTDeep: Moving Target Defense to Boost the Security of Deep Neural Nets Against Adversarial Attacks. GameSec, 2019.
- (126) Sailik Sengupta, Ankur Chowdhary, Dijiang Huang, Subbarao Kambhampati. General Sum Markov Games for Strategic Detection of Advanced Persistent Threats using Moving Target Defense in Cloud Networks. GameSec, 2019.
- (127) Anagha Kulkarni, Yantian Zha, Tathagata Chakraborti, Satya Gautam Vadlamudi, Yu Zhang and Subbarao Kambhampati. Explicability as Minimizing Distance from Expected Behavior. AAMAS (Extended Abstract), 2019.
- (128) Sailik Sengupta, Ankur Chowdhary, Dijiang Huang, Subbarao Kambhampati. Moving Target Defense for the Placement of Intrusion Detection Systems in the Cloud GameSec 2018
- (129) Lydia Manikonda, Ghazaleh Beigi, Subbarao Kambhampati, Huan Liu. #metoo Through the Lens of Social Media International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation (SBP-BRIMS) 2018.
- (130) Sarath Sreedharan, Siddharth Srivastava, Subbarao Kambhampati. Hierarchical Expertise-Level Modeling for User Specific Robot-Behavior Explanations. IJCAI 2018.
- (131) Wenfeng Feng, Hankz Hankui Zhuo, Subbarao Kambhampati. Extracting Action Sequences from Texts Based on Deep Reinforcement Learning. IJCAI 2018.

- (132) Tathagata Chakraborti, Sarath Sreedharan, Anagha Kulkarni, Subbarao Kambhampati. Projection-Aware Task Planning and Execution for Human-in-the-Loop Operation of Robots in a Mixed-Reality Workspace IROS 2018.
- (133) Lydia Manikonda & Subbarao Kambhampati Tweeting AI: Perceptions of Lay vs. Expert Twitterati. ICWSM 2018.
- (134) Sarath Sreedharan, Tathagata Chakraborti & Subbarao Kambhampati. Handling Model Uncertainty and Multiplicity in Explanations via Model Reconciliation ICAPS 2018.
- (135) Tathagata Chakraborti*, Sarath Sreedharan* & Subbarao Kambhampati. Balancing Explicability and Explanations: Emergent Behaviors in Human-Aware Planning AAMAS 2018 (Extended Abstract).
- (136) Yantian Zha, Yikang Li, Sriram Gopalakrishnan, Baoxin Li, & Subbarao Kambhampati Recognizing Plans by Learning Embeddings from Observed Action Distributions AAMAS 2018 (Extended Abstract).
- (137) Lydia Manikonda, Aditya Deotale, Subbarao Kambhampati. What's up with Privacy?: User Preferences and Privacy Concerns in Intelligent Personal Assistants AAAI/ACM Conference on Artificial Intelligence, Ethics and Society (AIES) 2018 (To appear)
- (138) Tathagata Chakraborti*, Sarath Sreedharan*, Yu Zhang and Subbarao Kambhampati. Plan Explanations as Model Reconciliation: Moving Beyond Explanation as Soliloquy IJCAI, 2017.
- (139) Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, Hankz Hankui Zhuo, Subbarao Kambhampati. Plan Explicability and Predictability for Robot Task Planning. ICRA 2017.
- (140) Sailik Sengupta, Satya Gautam Vadlamudi, Subbarao Kambhampati, Adam Doupe, Marthony Taguinod, Ziming Zhao and Gail-Joon Ahn A Game Theoretic Approach in Strategy Generation for Moving Target Defense in Web Applications AAMAS 2017.
- (141) Tathagata Chakraborti, Kartik Talamadupula, Kshitij P. Fadnis, Murray Campbell, Subbarao Kambhampati UbuntuWorld 1.0 LTS - A Platform for Automated Problem Solving & Troubleshooting in the Ubuntu OS AAAI/IAAI, 2017
- (142) Xin Tian, Hankz Hankui Zhuo and Subbarao Kambhampati. Discovering Underlying Plans Based on Distributed Representations of Actions. AAMAS 2016 (**Nominated for AAMAS 2016 Best Student Paper Award**)
- (143) Tweeting the Mind and Instagramming the Heart: Exploring Differentiated Content Sharing on Social Media. Lydia Manikonda, Vamsi Meduri & Subbarao Kambhampati ICWSM 2016.
- (144) Tathagata Chakraborti, Sarath Sreedharan, Sailik Sengupta, T. K. Satish Kumar, and Subbarao Kambhampati. Compliant Conditions for Polynomial Time Approximation of Operator Counts Symposium on Combinatorial Search (SOCS), 2016.
- (145) Yu Zhang, Sarath Sreedharan & Subbarao Kambhampati A Formal Analysis of Required Cooperation in Multi-agent Planning ICAPS 2016.
- (146) Tathagata Chakraborti, Yu Zhang, Subbarao Kambhampati. Planning with Resource Conflicts in Human-Robot Cohabitation AAMAS 2016.
- (147) Satya Gautam Vadlamudi, Sailik Sengupta, Subbarao Kambhampati, Marthony Taguinod, Ziming Zhao, Adam Doupi, Gail-Joon Ahn Moving Target Defense For Web

Applications Using Bayesian Stackelberg Games AAMAS 2016

- (148) Satya Gautam Vadlamudi & Subbarao Kambhampati A Combinatorial Search Perspective on Diverse Solution Generation AAAI 2016
- (149) Tathagata Chakraborti, Gordon Briggs, Kartik Talamadupula, Yu Zhang, Matthias Scheutz, David Smith and Subbarao Kambhampati Planning for Serendipity. IROS 2015
- (150) Yu (Tony) Zhang, Vignesh Narayanan, Tathagata Chakraborty & Subbarao Kambhampati. A Human Factors Analysis of Proactive Assistance in Human-robot Teaming. IROS 2015.
- (151) Manikandan Vijayakumar, Tejas Mallapura Umamaheshwar, Kartik Talamadupula and Subbarao Kambhampati TweetSense: Context Recovery for Orphan Tweets by Exploiting Social Signals in Twitter WebScience 2015.
- (152) Acquiring Planning Knowledge via Crowdsourcing Jie Gao, Hankz Hankui Zhuo, Subbarao Kambhampati and Lei Li. HCOMP 2015
- (153) Vignesh Narayanan, Yu Zhang, Nathaniel Mendoza and Subbarao Kambhampati. Automated Planning for Peer-to-peer Teaming and its Evaluation in Remote Human-Robot Interaction 10th ACM/IEEE Intl. Conf on Human Robot Interaction (HRI), 2015.
- (154) Yilin Wang, Yuheng Hu, Subbarao Kambhampati and Baoxin Li Inferring Sentiment from Web Images with Joint Inference on Visual and Social Cues: A Regulated Matrix Factorization Approach ICWSM 2015
- (155) Yu Zhang, Sarath Sreedharan and Subbarao Kambhampati. Capability Models and their application in Multi-agent planning. AAMAS 2015
- (156) Sushovan De, Yuheng Hu, Yi Chen, and Subbarao Kambhampati. BayesWipe: A Multimodal System for Data Cleaning and Consistent Query Answering on Structured Data IEEE BigData 2014.
- (157) Tuan A. Nguyen, Subbarao Kambhampati: A Heuristic Approach to Planning with Incomplete STRIPS Action Models. ICAPS 2014
- (158) Kartik Talamadupula, Gordon Briggs, Tathagata Chakraborti, Matthias Scheutz, Subbarao Kambhampati: Coordination in human-robot teams using mental modeling and plan recognition. IROS 2014: 2957-2962
- (159) Yu Zhang, Lynne E. Parker, Subbarao Kambhampati: Coalition coordination for tightly coupled multirobot tasks with sensor constraints. ICRA 2014: 1090-1097.
- (160) Yuheng Hu, Lydia Manikonda, Subbarao Kambhampati: What We Instagram: A First Analysis of Instagram Photo Content and User Types. ICWSM 2014
- (161) Lydia Manikonda, Tathagata Chakraborti, Sushovan De, Kartik Talamadupula, Subbarao Kambhampati: AI-MIX: Using Automated Planning to Steer Human Workers Towards Better Crowdsourced Plans. IAAI/AAAI 2014: 3004-3009
- (162) Tuan Nguyen, Subbarao Kambhampati & Minh B. Do. Synthesizing Robust Plans Under Incomplete Domain Models. NIPS 2013.
- (163) Kartik Talamadupula, Gordon Briggs, Matthias Scheutz & Subbarao Kambhampati. Architectural Mechanisms for Handling Human Instructions in Open-World Mixed-Initiative Team Tasks Advances in Cognitive Systems, ACS, 2013.

- (164) Srijith Ravikumar, Kartik Talamadupula, Raju Balakrishnan, Subbarao Kambhampati: RARProp: ranking tweets by exploiting the tweet/user/web ecosystem and inter-tweet agreement. CIKM 2013: 2345-2350
- (165) Hankz Hankui Zhuo, Tuan Nguyen and Subbarao Kambhampati. Refining Incomplete Domain Models through Plan Traces. IJCAI 2013.
- (166) Yuheng Hu, Fei Wang & Subbarao Kambhampati. Listening to the Crowd: Automated Analysis of Live Events via Aggregated Twitter Sentiment. IJCAI 2013.
- (167) Yuheng Hu, Kartik Talamadupula and Subbarao Kambhampati. Dude, srsly?: The Surprisingly Formal Nature of Twitter ICWSM, 2013.
- (168) Hankz Hankui Zhuo, Subbarao Kambhampati. Action Model Acquisition from Noisy Plan Traces. IJCAI 2013.
- (169) Hankz Hankui Zhuo, Subbarao Kambhampati, Tuan Nguyen. Model-Lite Case-Based Planning AAAI 2013.
- (170) Hankz Hankui Zhuo, Qiang Yang & Subbarao Kambhampati. Action-Model Based Multi-agent Plan Recognition NIPS 2012.
- (171) Yuheng Hu, Ajita John, Fei Wang & Subbarao Kambhampati. ET-LDA: Joint Topic Modeling for Aligning Events and their Twitter Feedback AAAI 2012.
- (172) Rehj Cantrell, Kartik Talamadupula, Paul Schermerhorn, J. Benton, Subbarao Kambhampati & Matthias Scheutz. Tell me when and why to do it!: Run-time planner model updates via natural language instruction ACM/IEEE Intl. Conf. on Human Robot Interaction. 2012.
- (173) Manish Kumar Jha, Raju Balakrishnan and Subbarao Kambhampati. Agreement Based Source Selection for the Multi-Domain Deep Web Integration. COMAD 2011.
- (174) Raju Balakrishnan, Subbarao Kambhampati SourceRank: Relevance and Trust Assessment for Deep Web Sources Based on Inter-Source Agreement. Proceedings of WWW 2011, Hyderabad.
- (175) Ravi Gummadi, Anupam Khulbe, Aravind Kalavagattu, Sanil Salvi, Subbarao Kambhampati. SmartInt: Using Mined Attribute Dependencies to Integrate Fragmented Web Databases. Proceedings of WWW 2011, Hyderabad. (Poster from regular track; 16% acceptance rate)
- (176) Raju Balakrishnan, Subbarao Kambhampati Factal: Integrating Deep Web Based on Trust and Relevance. Proceedings of WWW 2011, Hyderabad. (Demo paper).
- (177) Kartik Talamadupula, J. Benton, Paul Schermerhorn, Matthias Scheutz and Subbarao Kambhampati. Integrating a closed-world planner with an open-world robot: A case study. AAAI 2010
- (178) William Cushing, J. Benton, Subbarao Kambhampati. Cost-based search considered harmful. Symposium on Combinatorial Search (SOCS) 2010.
- (179) J.Benton, Robert Mattmueller, Patrick Eyerich, S. Kambhampati G-Value Plateaus: A Challenge for Planning. ICAPS 2010. (Challenge Paper).
- (180) Raju Balakrishnan and Subbarao Kambhampati. SourceRank: Relevance and Trust Assessment for deep web sources based on inter-source agreement. WWW 2010. (**Recipient of the "Best Poster" award at WWW-2010.**)

- (181) Ravi Gummadi, Anupam Khulbe, Aravind Kalavagattu, Sanil Salvi, and Subbarao Kambhampati. SMARTINT: A System for Answering Queries over Web Databases Using Attribute Dependencies. ICDE 2010 (Demo).
- (182) Nan Li, Will Cushing, Subbarao Kambhampati & Sungwook Yoon. Learning user plan preferences obfuscated by feasibility constraints. ICAPS 2009.
- (183) Paul Schermerhorn*, J. Benton, Matthias Scheutz, Kartik Talamadupula, Subbarao Kambhampati. Finding and Exploiting Goal Opportunities in Real-time during Plan Execution. IROS 2009.
- (184) Nan Li, Subbarao Kambhampati & Sungwook Yoon. Learning Probabilistic Hierarchical Task Networks to Capture User Preferences. IJCAI 2009.
- (185) Tuan Nguyen, Minh Do, Biplav Srivastava and Subbarao Kambhampati. Planning with Partial Preference Models. IJCAI 2009.
- (186) Sungwook Yoon and Subbarao Kambhampati (along with many other authors). An Ensemble Learning and Problem Solving Architecture for Airspace management Proc. IAAI 2009.
- (187) Sungwook Yoon, Alan Fern, Robert Givan and Subbarao Kambhampati. Probabilistic Planning via Determinization in Hindsight. AAAI 2008. (Presented orally as well as as a poster. Acceptance Ratio: 5%)
- (188) Sungwook Yoon, J. Benton and S. Kambhampati. An online learning method for improving over-subscription planning. ICAPS 2008 (Acceptance ratio: 30%)
- (189) Garrett Wolf, Hemal Khatri, Bhaumik Chokshi, Jianchun Fan, Yi Chen & Subbarao Kambhampati. Query processing over Incomplete Autonomous Databases. VLDB 2007. (Acceptance ratio: 16%)
- (190) Raju Balakrishnan and Subbarao Kambhampati. Optimal Ad Ranking for Profit Maximization. WebDB 2008.
- (191) Subbarao Kambhampati. Model-lite Planning for the Web Age Masses: The Challenges of Planning with Incomplete and Evolving Domain Theories AAAI 2007.
- (192) Menkes van den Briel, J. Benton, Subbarao Kambhampati and Thomas Vossen. An LP-Based Heuristic for Optimal Planning. Constraint Programming Conference (CP-2007).
- (193) William Cushing, Daniel Weld, Subbarao Kambhampati, Mausam & Kartik Talamadupula. Evaluating Temporal Planning Domains. ICAPS 2007.
- (194) J. Benton, Menkes van den Briel & Subbarao Kambhampati. A Hybrid Linear Programming and Relaxed Plan Heuristic for Partial Satisfaction Planning Problems. ICAPS 2007.
- (195) William Cushing, Subbarao Kambhampati, Mausam, Daniel Weld. IJCAI 2007 When is Temporal Planning Really Temporal? (oral presentation, Acceptance rate < 15%)
- (196) Planning with Goal Utility Dependencies Minh Do, J. Benton, Subbarao Kambhampati, Menkes van den Briel. IJCAI 2007. (oral presentation, Acceptance rate < 15%)
- (197) Domain Independent Approaches for Finding Diverse Plans Biplav Srivastava, Subbarao Kambhampati, Tuan Nguyen, Minh Do, Alfonso Gerevini, Ivan Serina. IJCAI 2007. (oral presentation, Acceptance rate < 15%)

- (198) QUIC: Handling Query Imprecision & Data Incompleteness in Autonomous Databases. Subbarao Kambhampati, Garrett Wolf, Yi Chen, Hemal Khatri, Bhaumik Chokshi, Jianchun Fan, Ullas Nambiar. CIDR 2007.
- (199) Sequential Monte Carlo in Probabilistic Planning Reachability Heuristics Daniel Bryce, Subbarao Kambhampati, and David E. Smith, ICAPS, 2006.
- (200) Supporting Queries with Imprecise Constraints Ullas Nambiar and Subbarao Kambhampati. AAAI 2006 NECTAR program.
- (201) Answering Imprecise Queries over Autonomous Web Databases Ullas Nambiar and Subbarao Kambhampati. To appear in Proc. ICDE 2006.
- (202) Oversubscription planning with metric goals. J. Benton, Minh B. Do and S. Kambhampati. IJCAI 2005.
- (203) Cost Sensitive Reachability Heuristics for Handling State Uncertainty Daniel Bryce and Subbarao Kambhampati. Proc. UAI 2005.
- (204) Romeo Sanchez and Subbarao Kambhampati. Planning Graph Heuristics for Selecting Objectives in Over-subscription Planning Problems. ICAPS 2005.
- (205) Menkes van den Briel, Thomas Vossen and Subbarao Kambhampati. Reviving Integer Programming Approaches for AI Planning: A Branch-and-Cut Framework ICAPS 2005.
- (206) Effective approaches for Partial Satisfaction (over-subscription) Planning. Menkes van den Briel, Romeo Sanchez Nigenda, Minh B. Do and Subbarao Kambhampati. AAAI 2004 (Acceptance Ratio: 25%).
- (207) Mining Approximate Functional Dependencies and Concept Similarities to Answer Imprecise Queries Ullas Nambiar and Subbarao Kambhampati. WebDB 2004. (Acceptance Ratio: 20%)
- (208) A frequency-based approach for mining coverage statistics in Data Integration. Zaiqing Nie and Subbarao Kambhampati. International Conference on Data Engineering (ICDE). 2004 (To appear). (Acceptance ratio: 14%)
- (209) Daniel Bryce, Subbarao Kambhampati. Heuristic Guidance Measures for Conformant Planning. Proc. ICAPS. 2004. (Acceptance ratio: 30%).
- (210) Bibfinder/Statminer: Effectively mining and using Coverage and Overlap Statistics in Data Integration Zaiqing Nie, Subbarao Kambhampati and Thomas Hernandez. Proc. Very Large Databases (VLDB) 2003. (Demonstration Paper) (Acceptance ratio: 25%)
- (211) Minh B. Do and Subbarao Kambhampati. Improving the Temporal Flexibility of Position Constrained Metric Temporal Plans. In Proc. ICAPS. (Acceptance ratio: 30%). 2003.
- (212) Minh B. Do and Subbarao Kambhampati. Planning Graph-based heuristics for Cost-sensitive Temporal Planning AIPS 2002. (Acceptance ratio: 30%).
- (213) Z.Nie, U. Nambiar, S. Lakshmi and S. Kambhampati. Mining Coverage Statistics for Webservice Selection in a Mediator Proc. CIKM 2002.
- (214) Zaiqing Nie and Subbarao Kambhampati Joint Optimization of Cost and Coverage of Information Gathering Plans. In Proc. of ACM CIKM 2001. CIKM 2001: 223-230 (Acceptance ratio: 25%).

- (215) Minh B. Do and Subbarao Kambhampati. Sapa: A domain-independent heuristic metric temporal planner. In Proc. ECP 2001. Toledo, Spain. (Acceptance ratio: 30%).
- (216) XuanLong Nguyen, Subbarao Kambhampati. Reviving Partial Order Planning In Proc. IJCAI-2001. IJCAI 2001: 459-466. (Acceptance ratio: 24.7%).
- (217) Extracting effective and admissible heuristics from the planning graph. X. Nguyen and S. Kambhampati. AAAI/IAAI 2000: 798-805. (Acceptance ratio: 30%).
- (218) M. Do and S. Kambhampati. Solving Planning Graph by Compiling it into a CSP. Proc. 5th AIPS. AIPS 2000: 82-91. (Acceptance ratio: 30%) **(10 Year Influential Paper Award, Runner-up; 2010)**.
- (219) S. Kambhampati and R. Nigenda. Distance-based goal ordering heuristics for Graphplan. Proc. 5th AIPS 2000. AIPS 2000: 315-322. (Poster. Acceptance ratio: 40%)
- (220) M. Do, B. Srivastava and S. Kambhampati. Investigating the effect of Relevance and Reachability Constraints on SAT encodings of Planning. Proc. 5th AIPS 2000. AIPS 2000: 308-314. (Poster. Acceptance ratio: 40%)
- (221) S. Kambhampati. Improving Graphplan's search with EBL and DDB Techniques. Proc. Intl. Joint Conf. on Artificial Intelligence (IJCAI-99). Stockholm. 1999. IJCAI 1999: 982-987. (Acceptance ratio: 26%).
- (222) E. Lambrecht, S. Kambhampati and S. Gnanaprakasam. Optimizing recursive information gathering plans. Proc. Intl. Joint Conf on Artificial Intelligence. (IJCAI-99). Stockholm. 1999. IJCAI 1999: 1204-1211. (Acceptance ratio: 26%).
- (223) T. Zimmerman and S. Kambhampati. Exploiting symmetry in the planning graph via explanation-guided search. Proc. National Conference on AI (AAAI-99). 1999. AAAI/IAAI 1999: 605-611. (Acceptance ratio: 25%).
- (224) A. Mali and S. Kambhampati. On the utility of Plan-space (causal) encodings. In Proc. National Conference on AI (AAAI-99). 1999. AAAI/IAAI 1999: 557-563. (Acceptance ratio: 25%).
- (225) S. Kambhampati, A. Mali and B. Srivastava. Hybrid planning for partially hierarchical domains. Proc. National Conf. on AI (AAAI-98), 1998. (Acceptance ratio: 28%)
- (226) A. Mali and S. Kambhampati. HTN planning as propositional satisfiability. Proc. 4th AI Planning Systems Conference. 1998 AIPS 1998: 190-198. (Acceptance ratio: 30%)
- (227) B. Srivastava, A. Mali and S. Kambhampati. A structured approach for synthesizing planners from specifications. Proc. 12th IEEE Intl. conf. on Automated Software Engineering Conference, Lake Tahoe. 1997. pp. 18-26. (Acceptance ratio: 30%)
- (228) S. Kambhampati, E. Parker and E. Lambrecht. Understanding and Extending Graphplan. 4th European Conference on Planning. September, 1997. (Acceptance ratio: 30%). pp. 262-274. (Also appears as a poster at IJCAI-97)
- (229) S. Kambhampati. Challenges in bridging plan synthesis paradigms. Intl. Joint Conf. on Artificial Intelligence. IJCAI-97. 1997. pp. 44-49. (Acceptance ratio: 25%)
- (230) S. Kambhampati and X. Yang. "On the role of Disjunctive representations and Constraint Propagation in Refinement Planning" *Proc. of Knowledge Representation and Reasoning*, 1996. KR 1996: 135-146. (Acceptance ratio: 31%)
- (231) S. Kambhampati, Formalizing Dependency Directed Backtracking and Explanation-based Learning in Refinement Search. *Proc. National Conference on Artificial*

Intelligence (AAAI-96), 1996. (Acceptance ratio: 30%)

- (232) L. Ihrig and S. Kambhampati. "Design and Implementation of a Replay Framework based on a Partial order Planner." *Proc. National Conference on Artificial Intelligence (AAAI-96)*, 1996. (Acceptance ratio: 30%)
- (233) S. Kambhampati, L. Ihrig and B. Srivastava. "A Candidate Set based analysis of sub-goal interactions in conjunctive goal planning." In *Proceedings of 3rd Intl. Conf. on AI Planning Systems*, May 1996. (Acceptance ratio: 30%)
- (234) S. Kambhampati and B. Srivastava. "Universal Classical Planner: An Algorithm for Unifying State-space and Plan-space Planning," *Current trends in AI Planning*, IOS Press, 1995. (based on a paper presented at 3rd European Workshop on Planning, Setemeber 1995. Acceptance Ratio: 39%).
- (235) L. Ihrig and S. Kambhampati. "Integrating Replay with EBL to improve planning performance." *Current trends in AI Planning*, IOS Press, 1995. (based on a paper presented at Proc. 3rd European Workshop on Planning, Setemeber 1995; Acceptance Ratio: 39%).
- (236) Y. Qu and S. Kambhampati. "Learning control rules for expressive plan-space planners: Factors influencing the performance." *Current trends in AI Planning*, IOS Press, 1995. (based on a paper presented at Proc. 3rd European Workshop on Planning, Setemeber 1995; Acceptance Ratio: 39%).
- (237) S. Kambhampati. "Admissible pruning strategies based on plan-minimality for plan-space planning" *Proc. 14th Intl. Joint Confernce on Artificial Intelligence. August 1995. (Acceptance ratio: 23%)*
- (238) S. Kambhampati. "Refinement Search as a unifying framework for analyzing planning algorithms," *Proc. 4th Intl. Conf. on Principles of Knowledge Representation and Reasoning*, May 1994 (Acceptance ratio: 20%)
- (239) S. Kambhampati. "Design Tradeoffs in partial order (plan-space) planning," *Proc. 2nd Intl. Conf. on AI Planning Systems*, June 1994. (Acceptance ratio: 30%)
- (240) S. Kambhampati and D.S. Nau. "On the nature of modal truth criteria in planning," *Proc. 12th Natl. Conf. on Artificial Intelligence (AAAI-94)*, August 1994. (Acceptance ratio: 27%)
- (241) L. Ihrig and S. Kambhampati. "Derivation Replay for Partial-order Planning," *Proc. 12th Natl. Conf. on Artificial Intelligence (AAAI-94)*, August 1994. (Acceptance ratio: 27%)
- (242) S. Katukam and S. Kambhampati. "Learning Explanation-based Search control rules for partial-order planning," *Proc. 12th Natl. Conf. on Artificial Intelligence (AAAI-94)*, August 1994. (Acceptance ratio: 27%).
- (243) S. Kambhampati. "On the Utility of Systematicity: Understanding Tradeoffs between redundancy and commitment in partial-ordering planning," *Proceedings of the 13th Intl. Joint Conf. on Artificial Intelligence*, Chamberry, France. (Acceptance ratio: 25%)
- (244) S. Kambhampati and J. Chen. "Relative utility of EBG-based plan reuse in partial ordering vs. total ordering planning frameworks," *Proceedings of 11th National Conference on Artificial Intelligence*, Washington D.C. (Acceptance ratio: 24%)
- (245) S. Kambhampati, "Characterizing Multi-Contributor Causal Structures for Planning," *Proceedings of the 1st Intl. Conf. on AI Planning Systems (AIPS-92)*, College Park,

MD, June 1992, pp. 116-125. (Acceptance ratio: 27%)

- (246) S. Kambhampati, M.R. Cutkosky, J.M. Tenenbaum and S.H. Lee, "Combining Specialized Reasoners and General Purpose Planners: A Case Study," *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)*, Anaheim, CA, July 1991. (Acceptance ratio: 23%)
- (247) S. Kambhampati and S.T. Kedar, "Explanation Based Generalization of Partially Ordered Plans" *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)*, Anaheim, CA, July 1991. (Acceptance ratio: 23%)
- (248) S. Kambhampati, "A Theory of Plan Modification," *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI-90)*, Boston, MA, August 1990, pp. 176-182. (Acceptance ratio: 17%)
- (249) S. Kambhampati, "Mapping and Retrieval during Plan Reuse: A Validation Structure Based Approach," *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI-90)*, Boston, MA, August 1990, pp. 170-175. (Acceptance ratio: 17%)
- (250) S. Kambhampati and J.A. Hendler, "Control of Refitting during Plan Reuse," *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI-89)*, Detroit, MI, August 1989, pp. 943-948. (Acceptance ratio: 20%)
- (251) S. Kambhampati and J.A. Hendler, "Flexible Reuse of Plans via Annotation and Verification," *Proceedings of the 5th IEEE Conf. on Applications of Artificial Intelligence (CAIA-89)*, Miami, FL, March 1989, pp. 37-43. (Acceptance ratio: 22%)

National/International Conference Proceedings/Reviewed papers

- (252) AltAlt: Combining the advantages of Graphplan and Heuristic State Search. Romeo Sanchez Nigenda, XuanLong Nguyen and Subbarao Kambhampati. Proc. Intl. Conference on Knowledge-based Computer Systems (KBCS-2000). Mumbai. India.
- (253) X. Li, S. Kambhampati and J. Shah. "ASUPPA: A Framework for Interactive and Iterative Synthesis and Improvement of Process Plans." DETC2000/CIE-14628. Proc. of DETC'00, ASME 2000 Design Engineering Technical Conferences and Computer and Information Engineering Conference. September, 2000.
- (254) X. Li, S. Kambhampati, K. Hirode and J. Shah. "Process Planner's Assistant: An incremental and interactive approach to automating process planning." ASME Design Engineering Technical Conference, 1997
- (255) S.H. Lee, M.R. Cutkosky and S. Kambhampati, "Incremental and Interactive Geometric Reasoning for Fixture and Process Planning," *Proceedings of 1992 ASME Winter Annual Meeting on Computer based Approaches to Concurrent Engineering*, November 1991.
- (256) S. Kambhampati and J.M. Tenenbaum, "Planning in Concurrent Domains," *Proceedings of the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA, November 1990, pp. 93-99.
- (257) S. Kambhampati and M.R. Cutkosky, "An Approach Toward Incremental and Interactive Planning for Concurrent Product and Process Design," *Proceedings of 1990 ASME Winter Annual Meeting on Computer based Approaches to Concurrent Engineering*, Dallas, TX, November 1990.

- (258) S. Kambhampati and J.A. Hendler, "Adaptation of Plans via Annotation and Verification," *Proceedings of the 1st International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, Tullahoma TN, June 1988 pp. 164-170.
- (259) S. Kambhampati and L.S. Davis, "Multi Resolution Path Planning for Mobile Robots," *Proceedings of the DARPA Image Understanding Workshop*, pp. 421-432, Miami Beach, FL, December 1985, pp. 421-432.
- (260) F.P. Andresen, L.S. Davis and S. Kambhampati, "Visual Algorithms for Autonomous Navigation," *Proceedings of the IEEE International Conference on Robotics and Automation*, St. Louis, MO, March 1985, pp. 856-861.

Symposia and Workshop Publications (reviewed proceedings):

- (261) Karthik Valmeekam, Kaya Stechly, Subbarao Kambhampati LLMs Still Can't Plan; Can LRMs? A Preliminary Evaluation of OpenAI's o1 on PlanBench Preprint on Arxiv, Sept 2024. NeurIPS 2024 Workshop on Open-World Agents. 2024. (Accepted) (Coverage in The Decoder, Coverage in Sina.com, Top paper in Alpha Signal Newsletter)
- (262) Siddhant Bhambri, Amrita Bhattacharjee, Huan Liu and Subbarao Kambhampati. Efficient Reinforcement Learning via Large Language Model-based Search NeurIPS 2024 Workshop on Open World Agents. 2024.
- (263) Mudit Verma, Siddhant Bhambri and Subbarao Kambhampati. Do Think Tags Really Help LLMs Plan? A Critical Evaluation of ReACT Style Prompting NeurIPS 2024 Workshop on Adaptive Foundation Models. 2024 (accepted)
- (264) Lucas Saldyt and Subbarao Kambhampati. Algorithmic Language Models with Neurally Compiled Libraries. NeurIPS 2024 Workshop on System 2 Reasoning at Scale. 2024. (Accepted)
- (265) Kaya Stechky, Matthew Marquez and Subbarao Kambhampati. GPT-4 Doesn't Know It's Wrong: An Analysis of Iterative Prompting for Reasoning Problems NeurIPS 2023 workshop on Foundation Models for Decision Making. (Covered in 36Kr)
- (266) Karthik Valmeekam, Matthew Marquez and Subbarao Kambhampati. Can Large Language Models Really Improve by Self-critiquing Their Own Plans? NeurIPS 2023 workshop on Foundation Models for Decision Making. (Covered in 36Kr).
- (267) Mudit Verma*, Siddhant Bhambri*, Subbarao Kambhampati. Preference Proxies: Evaluating Large Language Models in capturing Human Preferences in Human-AI Tasks Workshop on Theory of Mind in Communicating Agents at ICML 2023 ,Ip Mudit Verma, Siddhant Bhambri, Subbarao Kambhampati. Exploiting Action Distances for Reward Learning from Human Preferences Many Facets of Preference Learning workshop at ICML 2023
- (268) Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, Subbarao Kambhampati. On the Planning Abilities of Large Language Models -- A Critical Investigation KLR workshop at ICML 2023
- (269) Lin Guan*, Karthik Valmeekam*, Sarath Sreedharan, Subbarao Kambhampati. Leveraging Pre-trained Large Language Models to Construct and Utilize World Models for Model-based Task Planning KLR workshop at ICML 2023
- (270) Daman Arora, Subbarao Kambhampati. Learning and Leveraging Verifiers to Improve Planning Capabilities of Pre-trained Language Models KLR workshop at ICML 2023 (also at PRL workshop at IJCAI 2023).
- (271) A Mental Model Based Theory of Trust Zahra Zahedi, Sarath Sreedharan, Subbarao Kambhampati. XAI workshop at IJCAI 2023
- (272) Siddhant Bhambri, Purv Chauhan, Frederico Araujo, Adam DoupÃ©, and Subbarao Kambhampati. Using Deception in Markov Game to Understand Adversarial Behaviors through a Capture-The-Flag Environment AICS workshop at AAAI 2023
- (273) Zahra Zahedi, Sarath Sreedharan and Subbarao Kambhampati. A Mental-Model Centric Landscape of Human-AI Symbiosis R2HCAI workshop at AAAI 2023

- (274) Mudit Verma, Siddhant Bhambri, Subbarao Kambhampati. Exploiting Unlabeled Data for Feedback Efficient Human Preference based Reinforcement Learning R2HCAI workshop at AAAI 2023
- (275) Utkarsh Soni, Sarath Sreedharan, Mudit Verma, Lin Guan, Matthew Marquez, Subbarao Kambhampati. Towards customizable reinforcement learning agents: Enabling preference specification through online vocabulary expansion Workshop on Human in the Loop Learning, NeurIPS 2022
- (276) Lin Guan, Karthik Valmeekam, Subbarao Kambhampati. Relative Behavioral Attributes: Filling the Gap between Symbolic Goal Specification and Reward Learning from Human Preferences Workshop on Human in the Loop Learning, NeurIPS 2022
- (277) Zahra Zahedi, Sarath Sreedharan, Karthik Valmeekam and Subbarao Kambhampati. Explicable or Optimal: Trust Aware Planning in Iterated Human Robot Interaction Accepted as a stand-alone video at ICRA 2023
- (278) Advice Conformance Verification by Reinforcement Learning agents for Human-in-the-Loop Mudit Verma, Ayush Kharkwal, Subbarao Kambhampati Workshop on Human in the Loop Learning, NeurIPS 2022 IROS RLCONFORM 2022 (Video)
- (279) Sriram Gopalakrishnan, Mudit Verma, Subbarao Kambhampati Computing Policies That Account for the Effects of Human Uncertainty During Execution in Markov Decision Processes. ICAPS Workshop on Explainable AI Planning (XAIP) 2022.
- (280) Sarath Sreedharan, Christian Muise, Subbarao Kambhampati. Why Did You Do That? Generalizing Causal Link Explanations to Fully Observable Non-Deterministic Planning Problems. ICAPS Workshop on Explainable AI Planning (XAIP) 2022
- (281) Yantian Zha, Lin Guan and Subbarao Kambhampati. Learning from Ambiguous Demonstrations with Self-Explanation Guided Reinforcement Learning. AAAI 2022 Workshop on Reinforcement Learning in Games.
- (282) Alberto Olmo, Sarath Sreedharan, Subbarao Kambhampati. GPT3-to-plan: Extracting plans from text using GPT-3 ICAPS 2021 Workshop on KEPS and ICAPS 2021 Workshop on Planning for Financial Services. 2021.
- (283) Sailik Sengupta*, Zahra Zahedi*, Subbarao Kambhampati. Game-theoretic Model of Trust to Infer Human's Observation Strategy of Robot Behavior. R4P workshop RSS, 2021
- (284) Planning for Proactive Assistance in Environments with Partial Observability Anagha Kulkarni, Siddharth Srivastava, Subbarao Kambhampati XAIP Workshop ICAPS, 2021.
- (285) Sarath Sreedharan, Subbarao Kambhampati. Leveraging PDDL to Make Inscrutable Agents Interpretable: A Case for Post Hoc Symbolic Explanations for Sequential-Decision Making Problems. XAIP Workshop ICAPS'21
- (286) Sriram Gopalakrishnan, Mudit Verma, Subbarao Kambhampati Synthesizing Policies That Account For Human Execution Errors Caused By State Aliasing In Markov Decision Processes XAIP Workshop ICAPS'21
- (287) Zahra Zahedi, Mudit Verma, Sarath Sreedharan, Subbarao Kambhampati. Trust-Aware Planning: Modeling Trust Evolution in Longitudinal Human-Robot Interaction. Plan-Rob Workshop, XAIP Workshop ICAPS'21
- (288) Niharika Jain, Alberto Olmo, Sailik Sengupta, Lydia Manikonda, Subbarao Kambhampati. Imperfect ImGANation: Implications of GANs Exacerbating Biases on Facial

Data Synthetic Data Generation Workshop at ICLR'21 (Oral presentation)

- (289) Zahra Zahedi*, Sailik Sengupta* and Subbarao Kambhampati 'Why didn't you allocate this task to them?' Negotiation-aware Task Allocation and Contrastive Explanation Generation NeurIPS Workshop on Cooperative AI, 2020
- (290) Sailik Sengupta, Subbarao Kambhampati Learning Movement Policies in Bayesian Stackelberg Markov Games for Adaptive Moving Target Defense NeurIPS Workshop on Real-world Reinforcement Learning, 2020
- (291) Karthik Valmeekam, Sarath Sreedharan, Sailik Sengupta and Subbarao Kambhampati RADAR-X: An Interactive Interface Pairing Contrastive Explanations with Revised Plan Suggestions ICAPS Workshop on Explainable AI Planning (XAIP), 2020
- (292) Sriram Gopalakrishnan and Subbarao Kambhampati Model Minimization For Online Predictability ICAPS Workshop on Explainable AI Planning (XAIP), 2020
- (293) Sachin Grover, David Smith and Subbarao Kambhampati Model Elicitation through Direct Questioning ICAPS Workshop on Explainable AI Planning (XAIP), 2020
- (294) Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, David Smith and Subbarao Kambhampati A Bayesian Account of Measures of Interpretability in Human-AI Interaction ICAPS Workshop on Explainable AI Planning (XAIP), 2020
- (295) Alberto Olmo*, Sailik Sengupta*, Subbarao Kambhampati Not all Failure Modes are Created Equal: Training Deep Neural Networks for Explicable (Mis)Classification ICML Workshop(s) on Uncertainty & Robustness in Deep Learning (UDL) and Human-in-the-loop Learning (HILL), 2020
- (296) Sarath Sreedharan, Utkarsh Soni, Mudit Verma, Siddharth Srivastava, Subbarao Kambhampati Bridging the Gap: Providing Post-Hoc Symbolic Explanations for Sequential Decision-Making Problems with Inscrutable Representations ICML Workshop on Human-in-the-loop Learning (HILL), 2020
- (297) Lin Guan*, Mudit Verma*, Subbarao Kambhampati Explanation Augmented Feedback in Human-in-the-Loop Reinforcement Learning ICML Workshop on Human-in-the-loop Learning (HILL), 2020
- (298) Sachin Grover, Sailik Sengupta, Tathagata Chakraborti, Aditya Prasad Mishra, Subbarao Kambhampati iPass: A Case Study of the Effectiveness of Automated Planning for Decision Support Naturalistic Decision Making, 2019.
- (299) Aditya Prasad Mishra, Sailik Sengupta, Sarath Sreedharan, Tathagata Chakraborti, Subbarao Kambhampati CAP: A Decision Support System for Crew Scheduling using Automated Planning Naturalistic Decision Making, 2019.
- (300) Tathagata Chakraborti and Subbarao Kambhampati (How) Can AI Bots Lie? ICAPS Workshop on Explainable Planning (XAIP), 2019.
- (301) Sarath Sreedharan, Tathagata Chakraborti, Christian Muise, Subbarao Kambhampati A General Framework for Synthesizing and Executing Self-Explaining Plans for Human-AI Interaction ICAPS Workshop on Explainable Planning (XAIP), 2019
- (302) Anagha Kulkarni, Sarath Sreedharan, Sarah Keren, Tathagata Chakraborti, David E. Smith, Subbarao Kambhampati Design for Interpretability. ICAPS Workshop on Explainable Planning (XAIP), 2019.

- (303) Sailik Sengupta*, Zahra Zahedi*, Subbarao Kambhampati. To Monitor or to Trust: Observing Robot's Behavior based on a Game-Theoretic Model of Trust 21st International Workshop on Trust in Agent Societies (co-located with AAMAS), 2019
- (304) Markov Game Modeling of Moving Target Defense for Strategic Detection of Threats in Cloud Networks Ankur Chowdhary*, Sailik Sengupta*, Dijiang Huang, Subbarao Kambhampati Artificial Intelligence for Cyber Security Workshop (AICS) at AAI, 2019.
- (305) Plan-Recognition-Driven Attention Modeling for Visual Recognition Yantian Zha, Yikang Li, Tianshu Yu, Subbarao Kambhampati, Baoxin Li Workshop on Plan, Activity, and Intent Recognition at AAI, 2019.
- (306) Towards Understanding User Preferences for Explanation Types in Model Reconciliation Zahra Zahedi, Alberto Olmo, Tathagata Chakraborti, Sarath Sreedharan, Subbarao Kambhampati HRI Late Breaking Report, 2019.
- (307) Sailik Sengupta, Andrew Dudley, Tathagata Chakraborti and Subbarao Kambhampati. Investigation of Bounded Misclassification for Operational Security of Deep Neural Networks AAI'18 Workshop on Engineering Dependable and Secure Machine Learning Systems
- (308) Sailik Sengupta, Tathagata Chakraborti and Subbarao Kambhampati. MTDeep - Boosting the Security of Deep Neural Nets Against Adversarial Attacks with Moving Target Defense AAI'18 Workshop on Engineering Dependable and Secure Machine Learning Systems
- (309) Sailik Sengupta, Tathagata Chakraborti, Sarath Sreedharan, Satya Gautam Vadlamudi and Subbarao Kambhampati RADAR - A Proactive Decision Support System for Human-in-the-Loop Planning AAI 2017 Fall Symposium on Human-Agent Groups: Studies, Algorithms and Challenges. (Also presented at the ICAPS Workshop on User Interfaces for Scheduling and Planning (UISP) 2017).
- (310) Sarath Sreedharan*, Tathagata Chakraborti* and Subbarao Kambhampati Explanations as Model Reconciliation - A Multi-Agent Perspective AAI 2017 Fall Symposium on Human-Agent Groups: Studies, Algorithms and Challenges.
- (311) Sarath Sreedharan*, Tathagata Chakraborti* and Subbarao Kambhampati Balancing Explicability and Explanation in Human-Aware Planning AAI 2017 Fall Symposium on Artificial Intelligence for Human-Robot Interaction (AI-for-HRI).
- (312) Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, Hankz Hankui Zhuo and Subbarao Kambhampati. Plan explicability for robot task planning. RSS Workshop on Planning for Human-Robot Interaction, 2016
- (313) Tathagata Chakraborti, Kartik Talamadupula, Yu Zhang & Subbarao Kambhampati A Formal Framework for Studying Interaction in Human-Robot Societies AAI 2016 Workshop on Symbiotic Cognitive Systems.
- (314) A Game Theoretic Approach to Ad-hoc Coalitions in Human-Robot Societies Tathagata Chakraborti, Venkata Vamsikrishna Meduri, Vivek Dondeti & Subbarao Kambhampati AAI 2016 Workshop on Multiagent Interaction without Prior Coordination.
- (315) Lydia Manikonda, Heather Pon-Barry, Subbarao Kambhampati¹, Eric Hekler & David W. McDonald Venting Weight: Analyzing the Discourse of an Online Weight Loss Forum AAI 2016 Workshop on World Wide Web and Population Health Intelligence

- (316) Tathagata Chakraborti; Yu Zhang; David Smith; and Subbarao Kambhampati. Planning with Stochastic Resource Profiles: An Application to Human-Robot Co-habitation. ICAPS Workshop on Planning and Robotics (PlanRob) 2015.
- (317) Tathagata Chakraborti, Gordon Briggs, Kartik Talamadupula, Matthias Scheutz, David Smith and Subbarao Kambhampati. Planning for Serendipity -- Altruism in Human-Robot Cohabitation. ICAPS Workshop on Planning and Robotics (PlanRob) 2015.
- (318) Lydia Manikonda, Tathagata Chakraborti, Sushovan De, Kartik Talamadupula, Subbarao Kambhampati. AI-MIX: Using Automated Planning to Steer Human Workers towards Better Crowdsourced Plans HCOMP 2014. Works in Progress.
- (319) Lydia Manikonda, Heather Pon-Barry, Subbarao Kambhampati, Eric Hekler & David W. McDonald. Discourse Analysis of User Forums in an Online Weight Loss Application ACL 2014 Joint Workshop on Social Dynamics and Personal Attributes.
- (320) Kartik Talamadupula, Subbarao Kambhampati, Yuheng Hu, Tuan Anh Nguyen, Hankz Hankui Zhuo: Herding the Crowd: Automated Planning for Crowdsourced Planning. HCOMP (Works in Progress / Demos) 2013.
- (321) Kartik Talamadupula & Subbarao Kambhampati. Easychair as a Pedagogical Tool Engaging Graduate Students in the Reviewing Process. EAAI 2014 (AAAI).
- (322) A Formal Analysis of Required Cooperation in Multi-agent Planning Yu Zhang & Subbarao Kambhampati ICAPS Workshop on Distributed Multi-Agent Planning (DMAP), 2014.
- (323) Srijith Ravikumar, Kartik Talamadupula, Raju Balakrishnan Subbarao Kambhampati. RARProp: Ranking Tweets by Exploiting the Tweet/User/Web Ecosystem and Inter-Tweet Agreement AAAI 2013 (Late-breaking paper)
- (324) Kartik Talamadupula, Matthias Scheutz, Gordon Briggs, and Subbarao Kambhampati. On the Many Interacting Flavors of Planning for Robotics ICAPS 2013 Workshop on Planning and Robotics. 2013
- (325) Kartik Talamadupula, David Smith, William Cushing and Subbarao Kambhampati. A Theory of Intra-Agent Replanning ICAPS 2013 Workshop on Distributed and Multi-Agent Planning. 2013
- (326) Srijith Ravikumar, Raju Balakrishnan and Subbarao Kambhampati. Ranking Tweets Considering Trust and Relevance. SIGMOD Workshop on Information Integration on the Web (IIWeb), 2012.
- (327) Kartik Talamadupula, Paul Schermerhorn, J. Benton, Subbarao Kambhampati, and Matthias Scheutz. Planning for Agents with Changing Goals ICAPS 2011 System Demonstration. Placed 3rd for the "Best Demo" award
- (328) Kartik Talamadupula, Subbarao Kambhampati, Paul Schermerhorn, J. Benton and Matthias Scheutz. Planning for Human-Robot Teaming. ICAPS 2011 Workshop on Scheduling and Planning Applications, 2011.
- (329) William Cushing, J. Benton, Subbarao Kambhampati. Cost Based Satisficing Search Considered Harmful Technical report arxiv:1103.3687; March 2011. ICAPS 2011 Workshop on Heuristics for Domain Independent Planning.
- (330) J. Benton, Patrick Eyerich and Subbarao Kambhampati Enhancing Search for Satisficing Temporal Planning with Objective-driven Decisions ICAPS 2011 Workshop on Heuristics for Domain Independent Planning.

- (331) Tuan Nguyen, Subbarao Kambhampati, Minh Do. Synthesizing Robust Plans under Incomplete Domain Models Technical report arxiv:1104.5069; April 2011. Presented at AAAI 2011 Workshop on Generalized Planning.
- (332) Tuan A. Nguyen and Subbarao Kambhampati and Minh B. Do. Assessing and Generating Robust Plans with Partial Domain Models ICAPS 2010 Workshop on Planning under Uncertainty. 2010.
- (333) Kartik Talamadupula, J. Benton, Paul Schermerhorn, Subbarao Kambhampati and Matthias Scheutz. Integrating a ClosedWorld Planner with an Open World Robot: A Case Study. ICAPS 2009 workshop on bridging the gap between action and motion planning. 2009
- (334) Fluent Merging: A General Technique to Improve Reachability Heuristics and Factored Planning Menkes van den Briel, Subbarao Kambhampati and Thomas Vossen. ICAPS 2007 Workshop on Heuristics for Domain Independent Planning: Progress, Ideas, Limitations and Challenges.
- (335) Towards Model-lite Planning: A Proposal For Learning & Planning with Incomplete Domain Models Sungwook Yoon and Subbarao Kambhampati ICAPS 2007 Workshop on AI Planning and Learning.
- (336) Hierarchical Strategy Learning with Hybrid Representations Sungwook Yoon and Subbarao Kambhampati AAAI 2007 Workshop on Acquiring Planning Knowledge via Demonstration. 2007.
- (337) Model-Lite Planning: Diverse Multi-option plans & Dynamic Objective Functions. Daniel Bryce, William Cushing & Subbarao Kambhampati ICAPS 2007 Workshop on Planning and Plan Execution for Real World Systems.
- (338) Towards Model-lite Planning: A Proposal For Learning & Planning with Incomplete Domain Models Sungwook Yoon and Subbarao Kambhampati ICAPS 2007 Workshop on AI Planning and Learning.
- (339) Finding Admissible Bounds for Over-subscription Planning Problems J. Benton, Menkes van den Briel & Subbarao Kambhampati ICAPS 2007 Workshop on Heuristics for Domain Independent Planning: Progress, Ideas, Limitations and Challenges.
- (340) QPIAD: Query processing over Incomplete Autonomous Databases Hemal Khatri, Jianchun Fan, Yi Chen & Subbarao Kambhampati. ICDE 2007 (poster).
- (341) Thomas Hernandez and Subbarao Kambhampati. Improving Text Collection Selection with Coverage and Overlap Statistics. Proc. of Word Wide Web Conference, 2005. (poster paper).
- (342) Providing ranked relevant results for Web database queries. Ullas Nambiar and Subbarao Kambhampati. Proc. of Word Wide Web Conference, 2004. (poster paper).
- (343) Over-subscription in Planning: A Partial Satisfaction Problem. Menkes van den Briel, Romeo Sanchez Nigenda and Subbarao Kambhampati. ICAPS 2004 Workshop on Integrating Planning into Scheduling.
- (344) Planning-Scheduling Connections through Exogenous Events. Minh B. Do, Subbarao Kambhampati and Terry Zimmerman. ICAPS 2004 Workshop on Integrating Planning into Scheduling.
- (345) Planning in Belief Space with a Labelled Uncertainty Graph Daniel Bryce, Subbarao Kambhampati and David E. Smith. AAAI 2004 workshop on Learning and Planning in

Markov Decision Processes.

- (346) Romeo Sanchez Nigenda, Subbarao Kambhampati. Parallelizing State Space Plans Online. IJCAI 2003. Poster.
- (347) Terry Zimmerman, Subbarao Kambhampati. Using available memory to transform Graphplan's search. IJCAI 2003. Poster.
- (348) Zaiqing Nie, Subbarao Kambhampati. Frequency-based Coverage Statistics Mining for Data Integration. IJCAI Workshop on Information Integration on the Web. 2003.
- (349) Ullas Nambiar and Subbarao Kambhampati. Answering Imprecise Database Queries. ACM WIDM 2003 (Workshop on Web Information and Data Management).
- (350) Subbarao Kambhampati. A Critique of Knowledge-based planning track at ICP Position paper presented at the ICAPS Workshop on the Competition: Impact, Organization, Evaluation and Benchmarks. 2003.
- (351) Dan Bryce and Subbarao Kambhampati. Heuristic guidance measures for Conformant Planning. ICAPS workshop on Planning under uncertainty and incomplete information. 2003
- (352) S. Kambhampati, Senthil, G. Source-call ordering techniques for information gathering. IJCAI Workshop on Intelligent Information Integration. 1999.
- (353) E. Lambrecht and S. Kambhampati. Minimizing recursive information gathering plans. *In Proc. AAI Workshop on Integrating Information Sources*. 1998.
- (354) X. Li, S. Kambhampati and J. Shah. An iterative and interactive approach for process planning. *In Proc. AAI/SIGMAN workshop on AI Approaches in Manufacturing*, 1998.
- (355) E. Parker and S. Kambhampati. Making Graphplan goal-directed. Working notes of AIPS-98 workshop on Planning as combinatorial search. 1998.
- (356) A. Mali and S. Kambhampati. Refinement Planning as Satisfiability. Working notes of AIPS-98 workshop on Planning as combinatorial search. 1998.
- (357) A. Mali and S. Kambhampati. Frugal propositional encodings for planning. Working notes of AIPS-98 workshop on Planning as combinatorial search. 1998.
- (358) L. Ihrig and S. Kambhampati. "Automatic storage and indexing of plan derivations based on replay failures." *In Proc. of IJCAI Workshop on formal methods for reuse of plans, proofs and programs.*, August 1995.
- (359) S. Kambhampati. "No Planner is an Island: Lessons from Nextcut process planning System." In working notes of AAI Spring Symposium on Integrated planning applications. Stanford, March 1995.
- (360) L. Ihrig and S. Kambhampati. "Evaluating the effectiveness of replay in partial order planning." In working notes of AAI workshop on Case-based Reasoning, Seattle, 1994.
- (361) S. Kambhampati et. al. "Integrated approaches to improving the effectiveness of plan reuse," Proceedings of ARPA Planning Initiative workshop, Tucson, February 1994.
- (362) J. Hendler et. al. "Massively parallel support for casebased planning," Proceedings of ARPA Planning Initiative workshop, Tucson, February 1994.

- (363) S. Kambhampati, "On the Utility of Systematicity: Understanding Tradeoffs between Redundancy and Commitment in Partial-Ordering Planning," *AAAI Spring Symposium on Foundations of Automatic Planning: The classical approach and beyond*, March 1993
- (364) S. Kambhampati, "Utility Tradeoffs in Incremental Plan Modification," *AAAI Spring Symposium on Computational Considerations in supporting incremental modification and reuse*, Stanford, CA, March 1992.
- (365) S. Kambhampati, J.M. Tenenbaum, S.H. Lee and M.R. Cutkosky, "Towards a Hybrid Planning Architecture," *AAAI Fall Symposium on Principles of Hybrid Reasoning*, Asilomar, CA, November 1991.
- (366) S. Kambhampati, J.M. Tenenbaum, S.H. Lee and M.R. Cutkosky, "Interactive and Incremental Planning for Concurrent Design: A Case Study," *Proceedings of the DARPA Workshop on Manufacturing*, Salt Lake City, Utah, February 1991.
- (367) S. Kambhampati, "A Framework for Replanning in Hierarchical Nonlinear Planning," *AAAI Spring Symposium on Planning in Uncertain, Unpredictable, or Changing Environments*, Stanford, CA, March 1990.
- (368) S. Kambhampati, "A Classification of Plan Modification Strategies based on Coverage and Information Requirements," *AAAI Spring Symposium on Case-Based Reasoning*, Stanford, CA, March 1990.
- (369) S. Kambhampati and J.M. Tenenbaum, "Towards a Paradigm for Planning in Interactive Domains with Multiple Specialized Domain Modules," *AAAI Workshop on Automated Planning for Complex Domains*, Boston, MA, August 1990.
- (370) S. Kambhampati, J.M. Tenenbaum and M.R. Cutkosky, "Towards an Incremental and Interactive Paradigm for Process Planning in Concurrent Design Environments," *AAAI Workshop on Concurrent Engineering Design*, Boston, MA, August 1990.
- (371) S. Kambhampati, "Integrating Planning and Reuse: A Framework for Flexible Plan Reuse," *Proceedings of the 2nd DARPA Workshop on Case-Based Reasoning*, Pensacola Beach, FL, May 1989.
- (372) S. Kambhampati, "Representational Requirements for Plan Reuse," *Proceedings of the 2nd DARPA Workshop on Case-Based Reasoning*, Pensacola Beach, FL, May 1989.
- (373) J.A. Hendler and S. Kambhampati, "Refitting Plans for Case-Based Reasoning," *Proceedings of the DARPA Workshop on Case-Based Reasoning*, Clearwater Beach, FL, May 1988.

Theses:

- (374) S. Kambhampati, "Flexible Reuse and Modification in Hierarchical Planning: A Validation Structure Based Approach," **Ph.D. Thesis**, University of Maryland, College Park, MD, 1989.
- (375) S. Kambhampati, "Multi Resolution Path Planning for Mobile Robots," **M.S. Thesis**, University of Maryland, College Park, MD, 1985.
- (376) S. Kambhampati, "Some Experiments on Isolated Word Speech Recognition for Confusable Vocabulary," **B.TECH. Project Report**, Indian Institute of Technology, Madras, India, 1983.

View Points, Op-eds, Public Outreach Pieces on AI

- Can LLMs Really Reason and Plan? Communications of ACM Magazine (Blog). (View Point). Sep 2023.
- AI as (an Ersatz) Natural Science?. Communications of the ACM (Blog). (Viewpoint) June 2022.
- Broad and Shallow AI: The Promise and Perils of Competence without Comprehension. The India Forum. Jan 2022. (Column)
- Language Imitation Games and the Arrival of Broad and Shallow AI Communications of the ACM Blog. October 2021. (Made it to the top-5 articles list at #1).
- Beauty, lies & ChatGPT: Welcome to the post-truth world. The Hill. Feb 2023 (Opinion Column).
- Will Artificial Intel get along with us? Only if we design it that way. The Hill. Feb 2021. (Opinion Column)
- Our Deepfake Future. Virginia Quarterly Review. Spring 2020. (Interview in comic form!).
- Why are Artificial Intelligence systems biased?. The Hill. July 2020. (Opinion Column)
- AI computing will enter the 'land of humans' in the 2020s: The promise and the peril. The Hill. January 2020. (Opinion Column)
- Perception won't be reality, once AI can manipulate what we see. The Hill. November 2019. (Opinion Column)
- What just happened? The rise of interest in Artificial Intelligence. The Hill. August 2019. (Opinion Column)

Invited Presentations (not including the conference presentations):

- **Challenges of Human-Aware AI Systems. AAAI 2018 Presidential Address. New Orleans. February, 2018**
- Can LLMs Reason and Plan? Keynote at 21st Intl. Conf. on Knowledge Representation & Reasoning. Hanoi, Vietnam. November 2024.
- Can LLMs Reason and Plan? Keynote at 62nd Conf. on Association for Computational Linguistics (ACL). Bangkok. August 2024.
- Can LLMs Reason and Plan? Invited Colloquium at Notre Dame. October 2024.
- LLM-Modulo Frameworks as Compound AI Architectures for Robust Planning (and Reasoning). Keynote @ 3rd Neuro-Symbolic AI Summer School. August 2024.
- Can LLMs Reason and Plan? Invited Speaker. UPenn CLUNCH Computational Linguistics Seminar. September 2024.
- Panelist for IBM Panel on AI Reasoning. September 2024.
- On the Role of Large Language Models in Planning. Tutorial (2hrs) delivered to ~1000 attendees. ICML 2024. Vienna. July 2024.
- LLMs are System 1s for Humanity(..and that ought to be enough..). Keynote @ ICML 2024 Workshop on LLMs & Cognition. Vienna. July 2024. (Also was a panelist).
- Panelist, ICML 2024 workshop on Theoretical Foundations of Foundation Models. Vienna. July 2024.
- Can LLMs Reason and Plan? Ellis Robust LLMs Workshop (Invited Talk). Oxford University. July 2024.
- Can LLMs Reason and Plan? ESSAI/ACAI Summer School (Invited Keynote). July 2024.
- Can LLMs Reason and Plan? IBM Research Distinguished Lecture. Yorktown Heights. June 2024.
- Can LLMs Reason and Plan? Talk at Apple NLU Workshop. Cupertino. June 2024.
- Can LLMs Reason and Plan? Keynote at 4th Tailor Conf on Trustworthy AI from Lab to Market. Lisbon, Portugal. June 2024.
- LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks. Invited talk at ICRA 2024 Workshop on Vision-Language Models for Navigation and Manipulation (VLMNM). May 2024.
- Demystifying LLMs. Webinar for PanIT Frontiers of AI Forum. May 2024.
- AI, Law and The Courts. Primer on AI for Maricopa County Superior Court. May 2024.
- Can LLMs Reason and Plan? Google/DeepMind LLMs & Reasoning Seminar. April 2024.
- Can LLMs Reason and Plan? Distinguished Lecture at Megagon Labs. April 2024.
- Can LLMs Reason and Plan? Invited Talk at Qualcomm Research. March 2024.
- Primer on AI for AZ Supreme Court's Steering Committee on AI. Arizona Courts Complex. March 2024.
- Implications of LLMs. AAAI 2024 Plenary Panel. Vancouver. Feb 2024.
- On the role of large language models in Planning. AAAI 2024 half-day Tutorial. Feb 2024.

- Can LLMs Really Reason and Plan? Keynote at SCAI AI Day. Dec 2023. Seminar at IIT Kharagpur Center of Excellence in AI.
- AI, Law and The Courts. Keynote at Arizona Summit on AI, Law and The Courts. December 2023.
- Can LLMs Really Reason and Plan? Keynote at JP Morgan AI Summit. Nov 2023.
- Can LLMs Really Reason and Plan? Keynote at SCAI AI Day. Nov 2023.
- Leveraging Mental Models for Cooperation & Competition in Human-AI Teams. Keynote @ AAAI 2023 Fall Symposium on Agent Teaming in Mixed Motive Situations. October 2023.
- AI for Scientific Discovery. Briefing and Panel Remarks at National Academies workshop on AI for Scientific Discovery. October 2023.
- Avenging Polanyi's Revenge (Exploiting the Approximate Omniscience of LLMs in Planning..). Invited Symposium talk at Naval Research Labs' Naval Center for Applied Research in Artificial Intelligence. Sep 2023.
- Generative #AI Revolution: What Happens to our Reality When AI Can Spin Many Synthetic Realities? Invited keynote at the Triservices AI Workshop at INS Valsura. Aug 2023.
- Avenging Polanyi's Revenge (Exploiting the Approximate Omniscience of LLMs in Planning..). Invited keynote at KDD 2023 Finance Day. Aug 2023.
- Avenging Polanyi's Revenge (Exploiting the Approximate Omniscience of LLMs in Planning..). Invited keynote at ICML 2023 workshop on Knowledge, Logic and Reasoning in the era of Data Driven Learning. July 2023.
- On the role of large language models in Planning. ICAPS 2023 half-day Tutorial. July 2023.
- LLMs can't plan (but they can help you in planning). Invited Keynote at ICAPS 2023 workshop on AI in Finance. July 2023.
- Symbols as a Lingua Franca for Supporting Human-AI Interaction For Explainable and Advisable AI Systems. Keynote talk. 21st Intl. Conf. of Italian Association of Artificial Intelligence. Nov 2022.
- Symbolic Mental Models as Lingua Franca for Supporting Explainable and Advisable Human-AI Teaming. Keynote talk. IROS 2022 Workshop on Theory of Mind for Human-Agent Teams. Nov 2022.
- Generative #AI Revolution: What Happens to our Reality When AI Can Spin Many Synthetic Realities?. Public Lecture. Manthan India (Hyderabad). March 2023.
- Beyond Nodding & Pointing: Widening Explainable and Advisable Human-AI Interaction. Keynote Lecture at the Kohli Day. IIIT-Hyderabad. March 2023. (Also given at Neuro Science Dept, Indian Institute of Science, Bangalore, and Google Research India).
- Beauty, Lies & ChatGPT (On the Promise, Perils & Limits of the Generative AI Revolution). Inaugural SCAI Faculty Colloquium. SCAI. ASU. April 2023.
- Symbols as a Lingua Franca for Bridging Human-AI Chasm For Explainable and Advisable AI Systems. Colloquium at IIT Madras (in conjunction with their Distinguished Alumnus Award). September 2022.
- Planning to Advise and Explain Reinforcement Learning. Invited talk. ICAPS 2022 Workshop on Bridging Planning and RL. June 2022.

- Tamig Inscrutable AI Systems with Explicit Knowledge and a Symbolic Lingua Franca. Intel Workshop on Vectors of Cognitive AI. May 2022.
- Getting AI Agents to Interact and Collaborate with us on our terms. Strachey Lecture. Oxford. March 2021.
- Human-Aware AI. Invited Talk. AI/Machine Learning Working Group. Institute for research in sensing. 2022.
- Getting AI Agents to Interact and Collaborate with us on our terms. AI in Automation Lecture Series. IBM Research. May 2021.
- Getting AI Agents to Interact and Collaborate with us on our terms. Keynote Address. PAKDD 2021. May 2021.
- Taming Broad/Shallow AI wit Explicit Knowledge & Bridging Human-AI Chasm with Symbolic Lingua Franca. Keynote at CIKM 2021 Workshop on Knowledge Injection in Neural Networks. November 2021.
- Rise of AI and the Challenges of Human-Aware AI Systems. Talk at IIT Madras Forum Director's Circle Event. October 2021.
- Explainable Plans and Decisions. CRA/Informs/ACM SIGAI workshop on AI & OR. Invited presentation. September 2021.
- Blind Men and an Explanation: Towards a Unifying Perspective on Explainable AI. Invited talk at Explainable AI Conference. Robert Bosch Center on Data Science & AI. IIT Madras. June 2021.
- AI & Ethics. Presentation at TTI/Vanguard Forum on Future of AI. June 2021.
- Conversation on Societal Impacts of AI. Jawaharlal Nehru Center for the Study of Law and Governance. New Delhi. April 2021.
- AI & The courts. Remarks at the National Academies Workshop on Emerging Areas of Science, Engineering and Medicine for the Courts. Feb 2021.
- Helping our Robot Overlords Help Us. Invited talk at AAAI 2021 Plan, Activity and Intent Recognition Workshop. Feb 2021.
- Synthesizing Explainable Behavior for Human-AI Collaboration. Invited Keynote at IEEE Cognitive Machine Intelligence Conferences, December 2020.
- Synthesizing Explainable Behavior for Human-AI Collaboration. Invited Keynote at Samsung AI Forum. November 2020.
- Polanyi vs. Planning (Planning around AI's New Romance with Tacit Knowledge). Invited talk at ICAPS DC, 2020.
- Defending our Reality in the era of Deep Fakes. Talk at Big Data & AI, Toronto. 2020.
- Has the terminator arrived? AI becomes smarter than humans? India Science Festival. Pune, India. January 2020. (Public lecture)
- Synthesizing Explainable Behavior for Human-AI Collaboration. Invited Keynote at Taiwan Association of AI (TAAI) Conference. Kaohsiung. November 2019.
- Synthesizing Explainable Behavior for Human-AI Collaboration. Invited keynote talk at AAMAS 2019. Montreal. May 2019.
- Rise of AI and Challenges of Human-Aware AI Systems. Presidential Distinguished Lecture. Singapore Management University. April 2019.

- Synthesizing Explainable Behavior for Human-AI Collaboration. Invited Keynote talk at Symposium on Combinatorial Search. SOCS 2019. Napa, CA. July 2019.
- Challenges of Human-Aware AI Systems. Invited Keynote talk at The ACM India Joint Intl. Conf. on Data Science & Management of Data (CODS-COMAD). Kolkata, India. January 2019.
- Explainable Planning for Human-Robot Collaboration. Invited Talk at ICAPS Workshop on Planning and Robotics (PlanRob). Berkeley, CA. July 2019.
- Hopes and Fears In a World with Artificial Intelligence. Public Lecture. Manthan India. Hyderabad. July 2019.
- Rise of AI and Challenges of Human-Aware AI Systems. Distinguished Lecture at Purdue Integrative Data Science Initiative. Purdue University. April 2019.
- Challenges of Human-Aware AI Systems. Jonsson School Distinguished Lecture, University of Texas, Dallas. February 2018.
- AI: Current and Emerging Capabilities and Research. National Academies of Science, Engineering and Medicine. March 2018.
- Rise of AI: Status, Thresholds, Impacts. Invited Presentation at AAAS 2019. Washington D.C.
- Rise of AI: Status, Thresholds, Attack Surfaces. Keynote at FCC Forum on AI & ML. November 2018.
- Challenges of Human-Aware AI Systems. AI Lecture Series. Computer Science Department. Utah University. November 2018.
- The Rise of AI: Status, Thresholds, Attack Surfaces. XIV Seminario Internacional Inteligencia Artificial. Fundacion Copec-UC. Santiago, Chile. November 2018.
- Rise of AI: Status, Challenges, Impacts. ICAPS/CPAIOR 2018 Plenary Public Lecture, TU Delft. June 2018.
- Rise of AI: Status, Prospects and Thresholds. IEEE Aerospace Conference. Big Sky, Montana. Plenary. March 7, 2018.
- Rise of AI: Status, Prospects & Thresholds. Materials Research Society Spring Meeting. Phoenix. April 4, 2018.
- Our Relationship with AI. Panelist. Computer History Museum. September 2017.
- Explicability and Explanations in Human-Aware AI Agents. Invited talk at IJCAI 2017 Workshop on Explainable Artificial Intelligence (XAI). August 2017.
- Rise of AI: Status, Thresholds, Attack Surfaces. National Academies of Science. August 2017.
- The Rise of AI and The Challenges of Human-Aware AI Systems. Keynote Talk. CCF-GAIR Meeting. Shenzhen. June 2017.
- The Rise of AI and The Challenges of Human-Aware AI Systems. IIIT Hyderabad. June 2017.
- Challenges of Human-Aware AI Systems; Microsoft Faculty Summit. July 2017.
- Societal Impacts of Artificial Intelligence. Talk at Spirit of the Senses Salon. May 2017.
- Machines Take Over the World. Panelist. **Bill Nye Saves the World. Season 1. Episode 3. Netflix.** April 2017.

- Testimony on Impacts of Artificial Intelligence to Canadian Senate Standing Committee on Social Affairs, Science & Technology, Senate of Canada. March 2017.
- Origins Great Debate: Future of Artificial Intelligence: Who is in control?. Panelist. Tempe. February 2017.
- Planning Challenges in Human-Aware AI Systems. Invited talk at AAAI 2017 Workshop on Human-Aware Artificial Intelligence. San Francisco. February 2017.
- Planning Challenges in Human-Machine Collaboration (In Praise of Human-Aware AI). Invited talk at Johnson Controls. Milwaukee. January 2017.
- Innovation Talks: Advancing the Scientific Frontiers of Cognitive Systems. IBM World of Watson. Panelist. Las Vegas. October 2016.
- Planning Challenges in Human-Machine Collaboration (In Praise of Human-Aware AI). Key-note speech at Chinese Conference on Artificial Intelligence. Beijing. August 2016
- Aspen Institute AI Round Table, August 2016. (Invited participant). Aspen.
- The Path to General AI goes through Human-Machine Collaboration. An invited briefing to the DOD JASON Committee. San Diego. June 2016.
- Existential Threat, Moi? Panel presentation at the Governance of Emerging Technologies Conference. Tempe. May 2016,
- Symbols-Neurons, Logic-Probability, Replace-Augment, Disappointment-Doomsday: Where will the AI pendulum swing next? Invited talk at AAAI-2016 Open house. Phoenix. February 2016.
- Future of Artificial Intelligence. Part of a discussion series held by the School of Future of Innovation. ASU. April 2016.
- AI Thresholds. A presentation at Hastings Center workshop on Control and Responsible Innovation in Artificial Intelligence. April 2016.
- Planning challenges in human-machine collaboration. Lockheed Martin Distinguished Colloquium. Cherry Hill, PA. April 2016.
- Planning challenges in human-machine collaboration. AI Seminar. University of Washington. February 2016.
- Human-Aware and Human-in-the-Loop Planning and Decision Support. Invited short "What's Hot" talk at AAAI 2014. July 2014.
- Challenges of Human-in-the-loop Planning & Decision Support. Seminar at University of Connecticut 2014. October 2014.
- Human-in-the-Loop/Human-Aware Planning and Decision Support. Talk at ARL Invitational Workshop. October 2014.
- Wittgenstein's papers and Faraday's talks: Maxims for a Milk-fed Researcher. Invited Mentoring Talk at IJCAI 2013. August 2013.
- "You can't do that, Dave! Collateral Lessons from a Computational Quest to Design HAL", 16th ASU Last Lecture Series, April 2011.
- "Back to the Future of Planning" Invited lectures at ACAI Summer School on Automated Planning, Freiburg, Germany, 2011.
- "Incomplete Domain Models, Uncertain Users, Unending Planning Open Worlds" ONR/IPAM Machine Reasoning Workshops 3 & 4. November 2010.

- "Incomplete domain models, uncertain users and open worlds: Foundations of Model-lite Planning" Invited talk at the CMU Robotics Institute Seminar. April 2010.
- "How to write a good research paper?" Invited talk at the Doctoral Consortium. International Conference and Planning and Scheduling. Sep 2009.
- "Do Robots need a Bill of Rights?" A Science Cafe Discussion at AZ Science Center. 11/2007.
- "Human-Aware AI" Talk given at University of Washington Seminar on Future of AI. 11/2007.
- "Real World Planning: Soft Constraints & Incomplete Models" Colloquium at Dept of CSE, Washington University, St. Louis. 10/2007
- "Real World Planning: Soft Constraints & Incomplete Models" Colloquium at Institute for Human-Machine Cognition. Pensacola, FL. 9/2007.
- "Learning for Planning" A 4-hour Tutorial given at Intl. Conf. on Planning and Scheduling 2007 (joint with Sungwook Yoon).
- "Information Integration on the Web" A 4-hour Tutorial given at the National Conference on Artificial Intelligence, 2007 (joint with Craig Knoblock)
- "Planning & Scheduling with Oversubscribed Resources, Preferences and Soft Constraints" A 4-hour Tutorial given at the National Conference on Artificial Intelligence, 2007 (Joint with Minh Binh Do and Terry Zimmerman).
- "Planning Graph Based Reachability Heuristics" A 4-hour Tutorial given at Intl. Joint Conference on Artificial Intelligence, 2007. (Joint with Daniel Bryce).
- "Planning and Learning" Lectures at the 2006 Machine Learning Summer School. Canberra. Australia.
- "Scalability Revolution in Planning" Lecture at Royal Melbourne Institute of Technology. 2006.
- "Adaptive Information Integration." Invited Seminar at USC/Information Sciences Institute. November, 2004.
- "1001 ways to skin a planning graph for heuristic fun and profit." Invited talk at 13th International Conference on Planning and Scheduling (ICAPS). Subbarao Kambhampati. ICAPS 2003.
- "Information Integration on the Web," a 4-hour tutorial at National Conference on AI, Edmonton, Alberta. Canada. August 2002 (Co-presented with Craig Knoblock). <http://rakaposhi.eas.asu.edu/i3-tut.html>
- "Integrating Planning and Scheduling: Status and Prospects," Invited talk. Intl. Conference on Knowledge Based Computer Systems. Mumbai, India. December, 2000.
- "A Unifying and Brand-name free introduction to Planning," Invited lecture. PLANET Summer School on Planning. Cyprus. September 2000.
- "Recent Advances in AI Planning: A unified view", a 4-hour tutorial at National Conference on AI, Austin, Texas. July 2000. <http://rakaposhi.eas.asu.edu/planning-tutorial.html>
- "Recent Advances in AI Planning: A unified view", a 4-hour tutorial at International Joint Conference on AI, Stockholm, Sweden. July 1999.
- "Refinement Planning: Status and Prospectus", Invited Talk at National Conference on AI. 1996.

- “Winning by being lazy: Ideas of abstraction, hierarchy and least-commitment in Planning”, Invited Talk at NIPS-98 workshop on Reinforcement Learning. 1998.
- “Refinement planning”. Departmental Colloquium at Indian Institute of Technology, Kharagpur, India. 1998.
- “Refinement planning”. Departmental Colloquium at Indian Institute of Science, Bangalore, India. 1998.
- “Disjunctive refinement planning as a unifying framework for scaling up plan synthesis”, Invited Talk at ARPA Planning Initiative PI meeting. June 1997. Boston.
- “A Universal Classical planning algorithm based on refinement search”, University of Kaiserslautern, Germany, September 1995.
- “Universal classical Planning” Seminar at Naval Research Laboratory, Artificial Intelligence Center, April 1995.
- “Refinement Planning as a unifying framework for classical planning”, Seminar at University of Texas, El Paso, March 1995.
- “Planning and Learning,” Panelist at the 2nd Intl. Conference on AI Planning Systems. Chicago, June 1994.
- “Refinement Search as a Unifying Framework for Analyzing Planning Algorithms,” Colloquium talk at Edinburgh University, Department of Artificial Intelligence. Edinburgh. May 1994.
- “Computational Considerations in Supporting Incremental Modification and Reuse,” Fifth Annual Spring Symposium Series Public Forum on *Research Advances of Artificial Intelligence*, American Association for Artificial Intelligence, Stanford University, CA, March 26, 1992.
- “Improving Generality and Efficiency of Classical Planning,” A series of seminar talks at Indian Institutes of Technology (Madras, Delhi, Kanpur, Kharagpur, Bombay), and Indian Institute of Science, Bangalore, India. June 1992.
- “Supporting Flexible Plan Reuse,” Symposium on Learning Methods for Planning and Scheduling, Palo Alto, CA, January 1991.
- “A Theory of Incremental Plan Modification and its applications,” SIGLUNCH, Knowledge Systems Laboratory, Stanford University, February 1991.

Graduate Student Dissertations, Theses and Projects supervision

Ph.D. Dissertations

- (1) **Lin Guan.** Ph.D. 2024. Worked on sample efficient frameworks for Reinforcement Learning with Human Feedback. Research Scientist at Meta GenAI.
- (2) **Mudit Verma.** Ph.D. 2024. Worked on incorporating sample-efficient active and passive human feedback in Reinforcement learning. Research Scientist at Google AI Gemini Group.
- (3) **Utkarsh Soni.** Ph.D. 2024. Worked on Explainable Human-AI Interaction in Sequential Decision Making. Data Scientist at ManuLife.
- (4) **Zahra Zahedi.** Ph.D. 2023. Worked on trust in human-AI interaction.
- (5) **Alberto Olmo Hernandez.** Ph.D. 2022. Worked on investigating failure modes of large learned models. ML Ops Engineer at Vianai.
- (6) **Sriram Gopalakrishnan.** Ph.D. 2022. Worked on incorporating human limitations into plans suggested for them. AI Research Scientist at JP Morgan AI Research.
- (7) **Sachin Grover.** Ph.D. 2022. Worked on active learning of human mental models. Research Scientist at PARC.
- (8) **Sarath Sreedharan.** Ph.D. 2022. Worked on explainable human-AI interaction. Faculty member at Colorado State University.
- (9) **Yantian Zha.** Ph.D. 2022. Worked on combining vision and high-level planning. Post-doc at Maryland Robotics. University of Maryland.
- (10) **Anagha Kulkarni,** Ph.D. 2021. Worked on Human-Aware AI systems. Research Scientist at Invitae.
- (11) **Sailik Sengupta,** Ph.D. 2020. Worked on AI and Cyber Security. Research Scientist at Amazon AI.
- (12) **Lydia Manikonda,** Ph.D. 2019. Worked on analysis and decision-making with social media. Faculty member at Rensselaer Polytechnic Institute.
- (13) **Tathagata Chakraborti,** Ph.D. 2018. Worked on foundations of human-aware planning. Currently at IBM AI Research. *ICAPS Distinguished Dissertation Awar (Runner-Up). 2019. ASU CIDSE Outstanding Ph.D. Student. 2019. IEEE AI 10 to Watch. 2020*
- (14) **Yuheng Hu,** Ph.D. 2014. Worked on event analytics for social media. Currently on the faculty at University of Illinois, Chicago.
- (15) **Kartik Talamadupula,** Ph.D. 2014. Worked on planning for human-robot teaming. Currently Research Staff Member at IBM Watson Labs.
- (16) **Tuan Nguyen,** Ph.D. 2014. Worked on planning with partial preferences and domain models. Currently at Mathworks.
- (17) **Sushovan De,** Ph.D. 2014. Worked on Data Rectification for Big Data. Currently at Google.
- (18) **William Cushing,** Ph.D. 2012. Worked on the foundations of temporal planning. Currently a post-doctoral fellow at UC Berkeley (with Prof. Stuart Russell).

- (19) **J. Benton**, Ph.D. 2012. Worked on partial satisfaction planning. Currently a research scientist at SIFT LLC.
- (20) **Raju Balakrishnan**, Ph.D. 2012. Worked on deep web source selection. Currently at Groupon.
- (21) **Menkes van den Briel**, Ph.D. 2008. Worked on AI & OR Techniques for scaling up Automated Planning. *Honorable Mention for ICAPS Best Dissertation Award*, 2009. First Employment: Research Faculty Member at University of Colorado (Leeds Business School).
- (22) **Daniel Bryce**, Ph.D. 2007. Worked on scaling up planning under uncertainty. *Winner of ICAPS Best Dissertation Award, 2009*. Currently Assistant Professor at Utah State University (Computer Science Department).
- (23) **Ullas Nambiar**, Ph.D. 2005. Worked on supporting imprecise queries over autonomous databases. Currently research scientist at IBM India Research Labs.
- (24) **Romeo Sanchez**, Ph.D. 2005. Worked on reachability heuristics for planning. Currently research scientist at USC/Information Sciences Institute.
- (25) **Binhminh Do**, Ph.D. August 2004. Worked on metric Temporal Planning--the SAPA project. *Winner of the Department Outstanding Ph.D. Student Award, 2005 (also ASU nominee for ACM Distinguished Dissertation Award)*. Also received an honorable mention for 2010 ICAPS 10-year Influential Paper Award. Currently a Research Scientist at PARC (formerly called Xerox Palo Alto Research Labs).
- (26) **Zaiqing Nie**, Ph.D. Spring 2004. Worked on multi-objective query optimization and similarity queries for Havasu Data Integration system. Currently with Microsoft Research Asia.
- (27) **Terry Zimmerman**, Ph.D. 2003. Dissertation on Effective use of memory in Graphplan. Currently Post Doctoral Fellow at CMU Robotics Institute.
- (28) **Biplav Srivastava**, Ph.D. Dissertation: “*Efficient Planning by Effective Resource Reasoning*”. March 2000. M.S. Thesis: “*Using refinement search to unify and synthesize classical planners*” (1996). Faculty member at University of South Carolina.
- (29) **Amol D. Mali**, Ph.D. Dissertation: “*Hierarchical task network planning as Satisfiability*,” Ph.D. 1999. Current position: Faculty member at University of Wisconsin, Milwaukee.
- (30) **Xiaomin Li**. Ph.D. Dissertation: “*ASUPPA: Interactive and Iterative Framework for Process Planning*” Ph.D. 2001. Quest Communication, Denver.
- (31) **Laurie Ihrig**, Ph.D. Dissertation: “*Improving planning performance through derivational replay*.” (1996). Boeing Corp., Phoenix.

M.S. Theses

- (1) **Karthik Valmeekam**, M.S. 2021. Worked on RADAR-X decision support system for human-in-the-loop mission planning.
- (2) **Niharika Jain**, M.S. 2020. Worked on biases in GAN-based data augmentation. (**Recipient of CIDSE Outstanding M.S. Student Award, 2021; CIDSE outstanding Undergraduate, Computer Science, 2019**)

- (3) **Aditya Prasad Mishra**. M.S. 2019. Worked on human-in-the-loop support for mission-planning.
- (4) **Sarath Sridharan**. M.S. 2016. Worked on multi-agent planning. Department Outstanding Masters Student Award, 2017. Continuing Ph.D. (**Recipient of CIDSE Outstanding M.S. Student Award, 2017**)
- (5) **Vignesh Narayanan**. M.S. 2015. Worked on human-factors analysis of automated planning for human-robot teaming. First Employment: Amazon.
- (6) **Anirudh Acharya**. M.S. 2015. Worked on joint topic modeling for tweet/event alignment. First Employment: Yahoo!
- (7) **Tejas Mallapura Umamaheshwar**. M.S. 2015. Worked on Hashtag Recovery for Twitter Messages. First Employment: Mathworks.
- (8) **Manikandan Vijaykumar**, M.S. 2014. Worked on Hashtag Recovery for Twitter Messages. First Employment: American Express.
- (9) **Preet Inder Singh Rihan**, M.S. 2013. Worked on query processing for probabilistic data. First Employment: Blomberg LP R&D.
- (10) **Srijith Ravikumar**, M.S. 2013. Worked on ranking tweets. First Employment: Amazon.
- (11) **Manish Jha**, M.S. 2011. Worked on topic-sensitive source rank for deep web sources. Department Outstanding Masters Student Award. 2012. First Employment: Amazon.
- (12) **Rohit Raghunathan**, M.S. 2011. Worked on principled probabilistic frameworks for query processing over autonomous databases in the presence of incompleteness. First Employment: Amazon.
- (13) **Chris White**, (Integrated) M.S. 2011. Worked on distributed approaches to Taxi World. First Employment: Lockheed Martin. .
- (14) **Sanil Salvi**, M.CS 2009. Worked on integrating databases on the web in the absence missing primary-foreign key relations.
- (15) **Ravi Gummadi**, M.S. 2009. Worked on integrating databases on the web in the absence missing primary-foreign key relations. Currently at Facebook.
- (16) **Anupam Khulbe**, M.S. 2009. Worked on integrating databases on the web in the absence missing primary-foreign key relations. Currently at Amazon.
- (17) **Garrett Wolf**, MCS 2008. Worked on handling incompleteness and imprecision in autonomous databases. Department Outstanding Masters Student Award. 2008.
- (18) **Aravind Kalavagattu**, M.S. 2008. Worked on learning approximate functional dependencies from data. Currently at Yahoo.
- (19) **Bhaumik Chokshi**, M.S. 2007. Worked on overlap-aware techniques for collection selection. Currently at MSN Search Live.
- (20) **Hemal Khatri**, M.S. 2006. Worked on handling incompleteness in autonomous databases. Currently at MSN Search Live. **Winner of the Department Outstanding M.S. Student Award, 2007..**
- (21) **Jianchun Fan**, M.S. 2006. Worked on multi-objective query optimization for data aggregation. Currently at Amazon.

- (22) **Thomas Hernandez**, M.S. 2004. Worked on collection selection with overlapping text collections. Currently at Amazon. **Winner of the Department Outstanding M.S. Student Award, 2005.**
- (23) **Sreelakshmi Vaddi**, M.C.S. 2003. Worked on execution issues in data integration.
- (24) **XuanLong Nguyen**, M.S. 2001 Thesis: “*Heuristic Search Control for Plan Synthesis Algorithms and Dynamic Constraint Satisfaction Problems*”
- (25) **Senthil Gnanaprakasam**, M.S. 2001. Thesis: “*A System-R style Join Order Optimization for Internet Information Gathering.*” Currently at Intel, Seattle.
- (26) **Eric Lambrecht**, M.S. Student. RA since Fall 1996. M.S. thesis: “*Optimizing Recursive Information Gathering Plans.*” Founded *shareyourworld.com.*, San Francisco.
- (27) **Yong Qu**, M.S. Thesis: “*Extending EBL framework to planning with expressive action representations*” (1996). Currently at Microsoft Corp., Seattle.
- (28) **Suresh Katukam**, M.S. Thesis: “*Explanation based learning for search control rules for partial order planning*” (1995). Currently at Hypercom, Phoenix.
- (29) **Bulusu Gopi Kumar**, M.C.S. Project: “*Systematic and complete dependency directed backtracking* (1994)
- (30) **Eric Cohen**. M.C.S. Project: “*Minimality based pruning techniques for Partial Order Planning*” (1993)
- (31) **Dennis Chen**, M.C.S. Project: “*Comparisons beteen Opaque and Interleavable macros.*” (1993)
- (32) **Jengchin Chen**, M.C.S. Project: “*Relative utility of EBG based plan reuse in total order vs. partial order planning.*” (1992)

Undergraduate Research Students

- (1) Daniel D’Souza (Spring 2015-). Honors Thesis. Spring 2017. URAP, FURI.
- (2) Gabriel Saba (Spring 2016-) Honors thesis, Spring 2017.
- (3) Cameron Dudley (Spring 2016-)
- (4) Jake Merdich (Summer 2015; URAP)
- (5) Nathaniel Mendoza (Summer 2014-Spring 2015)(URAP; FURI)
- (6) Sumbhav Sethia (Summer 2013-Fall 2014)(URAP; FURI)
- (7) Wyatt Tyree (Summer 2014-Fall 2014)(URAP)
- (8) Paul Reesman (Summer 2013)(URAP)
- (9) Matthew Mellott (Fall 2012--Spring 2013)(NSF REU, FURI Scholar)
- (10) Nicole Ang (Fall 2012)(NSF REU)
- (11) James Fiacco (Fall 2011) (FURI Scholar)
- (12) Alex Wallace(Summer 2009-Spring 2010) (B.S. honors thesis. Thesis) (Currently at U. Arizona Med school)
- (13) Chris White(Fall 2009-Spring 2010) (B.S. honors thesis)(Continued for Integrated M.S.)

- (14) Alan Hogan (Spring 2008) (Undergraduate FURI Scholar)
- (15) Kartik Talamadupula(Spring 2007-) (Bachelors; Dec 2007)(Continued for PhD)CSE Distinguished Senior award, 2008. CRA Outstanding Undergraduate Student (honorable mention), 2008 (local press) Receptient of FURI and CS UG Research grants) Science Foundation of Arizona (SFAZ) Fellow (2008).
- (16) Wes Dyer (Spring 2004-) Undergraduate (honors thesis) (Thesis) (Currently at Microsoft).

Professional Activities and Service to the Research Community

- **President, Association for the Advancement of Artificial Intelligence. 2016-18.**
- Chair, Section T (Information, Computation and Communication), American Association for the Advancement of Science (AAAS). 2022-2024.
- Trustee, Partnership for Artificial Intelligence. 2017-2018.
- **Program Chair, Intl. Joint Conference on Artificial Intelligence, 2016. Trustee 2013-18.**
- **Program Co-Chair; National Conference on Artificial Intelligence, 2005 (Pittsburgh, PA).**
- **Program Co-Chair; International Conference on Planning and Scheduling, 2013 (Rome, Italy)**
- High level Government service including
 - (1) Testimony on Impacts of Artificial Intelligence to Canadian Senate Standing Committee on Social Affairs, Science & Technology, Senate of Canada. March 2017.
 - (2) Panel on AI Capacity Building in India (Organized by Hon. Sushma Swaraj; External Affairs Minister of India), June 2018.
 - (3) Keynote at the Federal Communications Commission's forum on AI & ML. 2018
 - (4) Whitehouse Roosevelt Room discussion on National AI strategy (eventually lead to the National AI research priorities, with human-AI collaboration a top priority). March 2019.
 - (5) (Multiple briefings at the National Academies of Science and Engineering)
- Founding Steering Committee Member, AAAI/ACM Conference on AI, Ethics and Society, 2018--
- Steering Committee Member, IBM AI X-Prize. 2017-
- **Track Co-Chair; AI and Web Track, National Conference on Artificial Intelligence, 2010.**
- **Program Co-Chair; AI Planning Systems Conference, 2000 (Breckenridge, CO)**
- **Executive Council Member (Elected), American Association for Advancement of Artificial Intelligence, 2009-2012.**
- **Executive Council Member, International Conference on Planning and Scheduling. 2002-2008.**
- **Conference Committee Chair, AAAI. 2012-14.**
- Advisory Board Member, Journal of Artificial Intelligence Research (JAIR), 2008-present.
- Editorial Board Member, AI Magazine. 2007-2014
- **Advisor**, Intl. Conf. on Automated Planning and Scheduling. 2021.
Co-Chair, Special track on Human-Aware Planning and Scheduling. ICAPS 2022.
- **Area Chair**, AAAI Conference on Artificial Intelligence, 2020, 2021, 2022, 2024, 2025
- **Area Chair**, NeurIPS (Benchmark track), 2023, 2024.
- **Senior Area Chair**, International Joint Conference on Artificial Intelligence, 2021,2022

- **Area Chair**, International Joint Conference on Artificial Intelligence, 2007, 2013, 2015, 2017, 2019, 2020
- **Area Chair**, Autonomous Agents and Multi-Agent Systems Conference, 2021
- **Senior Program Committee Member**, National Conference on Artificial Intelligence, 2007, 2008, 2011, 2012, 2014.
- Senior Program Committee Member, International Joint Conference on Artificial Intelligence, 2009,2014.
- Program Committee Member, Intl. Conference on Planning and Scheduling, 2007, 2008, 2009, 2010, 2011, 2012, 2017, 2019.
- Program Committee Member, AI, Ethics & Society Conference. 2018, 2019.
- Editorial Board Member, IEEE Intelligent Systems. 2003-present
- Editorial Board Member, ACM Trans. on Intelligent Systems and Technology. 2012-present.
- Senior Program Committee Member, National Conference on Artificial Intelligence, 2006.
- Program Committee Member, Intl. Conference on Planning and Scheduling, 2006.
- Program Committee Member, European Semantic Web Conference, 2007.
- Program Committee Member, Intl. Conference on Planning and Scheduling, 2005. (Also organized the first-ever *Festivus*). • Associate Editor, Journal of Artificial Intelligence Research (JAIR), 2003-2007.
- Executive Council Member, AI Planning and Scheduling Conference. 2000-2002
- Maintainer of Planning List Digest, a mailing list for AI Planning researchers, 1996-.
- Member, Editorial Board, Journal of Artificial Intelligence, 1995-97
- Co-Chair, IJCAI Workshop on Intelligent Information Integration on the Web. 2003.
- Co-Editor, Special Issue on Intelligent Information Integration. IEEE Intelligent Systems. 2003.
- Editor, Special Issue on Planning and Scheduling of the journal Constraints. 2001.
- Program Committee member, National Conference on AI, 2003.
- Program Committee member, AAMAS, 2003.
- Program Committee member, 13th International Conference on Planning and Scheduling. 2003.
- Program Committee member, Intl. Conference on Autonomous and Multi-Agent Systems, 2003.
- Program Committee member, Intl. Conf. on Knowledge Based Computer Systems.2002.
- Program Committee member. European Conference on Planning. Toledo, Spain. September, 2001.
- Reviewer, IJCAI 2001.
- Program Committee Member and Session Chair, ACM SIGMOD Conference (ACM conference on Databases). Dallas, TX in May 2000.
- Program committee member. National Conference on AI. July 2000.
- Organizing Committee member. AAAI workshop on CSP approaches to planning. July 2000.

- Program committee member, 5th European Conference on Planning, Durham, UK, September 2000.
 - Organizing Committee Member, AAAI Spring Symposium on Search in uncertain, incomplete and Dynamic environments. Held at Stanford, CA, March 1999.
- Reviewer, and Challenge paper coordinator, International Joint Conference on Artificial Intelligence, held at Stockholm, Sweden. August 1999.
- Program committee member and session chair, 15th National Conference on AI. Held at Madison, Wisconsin, July 1998.
 - Program committee member and session chair, 4th AI Planning Systems Conference. Held at Pittsburgh, PA, May, 1998.
 - Co-organizer, Workshop on planning as combinatorial search, 4th AI Planning Systems Conference, 1998.
 - Program committee member, ASME Conference on Design for Manufacturing, 1998.
 - Program committee member, 13th National Conference on AI, 1996.
 - Workshop Chair, 13th National Conference on Artificial Intelligence, 1996
 - Program Committee Member, 3rd European Planning Workshop, Assisi, Italy, 1995.
 - Program Committee Member, 2nd International Conf. on AI Planning Systems, Chicago, 1994.
 - Organizing Committee Member, AAAI Fall Symposium on “Adaptation of Knowledge for Reuse,” 1995.
 - *Chair*, 1992 AAAI Spring Symposium on Computational Considerations in Supporting Modification and Reuse (March 1992)
 - *Program Committee Member and Session Chair*, 10th National Conference on Artificial Intelligence (AAAI-92)
 - *Program Committee Member and Session Chair*, 1st International Conferences on AI Planning Systems, College Park, MD, 1992.
 - Organizing Committee Member, AAAI 1993 Spring Symposium on Foundations of Automatic Planning: The Classical Approach and Beyond
 - Proposal Reviewer, National Science Foundation (1992)
 - *Member*, American Association for Artificial Intelligence
 - *Panel Member*, NSF KMCS proposal review panel, 1993; 1995; 1997; 1999; 2002.
 - **Referee:**
Journal of ACM, Artificial Intelligence Journal, Machine Learning, Journal of AI Research, Journal of Logic and Computation, IEE Expert, IEEE Trans. on Systems, Man and Cybernetics, ASME Transactions on Computer-aided Manufacturing, IEEE Intl. Conference on Robotics and Automation, IEEE Intl. Conference on Applications of Artificial Intelligence.

Media Coverage

(Quoted/profiled/wrote in multiple outlets, including Netflix, The New York Times, Washington Post, National Public Radio, The New Yorker, The Guardian, The Atlantic, Washington Post, Wired, The Hindu, etc. Hyperlinked list available at <http://rakaposhi.eas.asu.edu/media.html>)

- Why multi-agent AI tackles complexities LLMs can't. Venture Beat. November 2024. (Quoted)
- OpenAI is So Doomed if Inference Time Scaling for o1 Fails. Analytics India Magazine. 2024. (Quoted)
- Apple Proves OpenAI o1 is Actually Good at Reasoning. Analytics India Magazine. October 2024. (Quoted)
- AIM 100: The Most Influential Global Leaders in AI. Analytics India Magazine. October 2024. (Listed as an AI Creator in their AI 100)
- Researchers put OpenAI's o1 through its paces, exposing both breakthroughs and limitations. The Decoder. September 2024. (Coverage of our recent research).
- OpenAI o1 Likely Uses RL over Chains of Thought to Build System 2 LLMs. Analytics India Magazine. September 2024. (Cites my speculations on OpenAI o1 model)
- Why Google Will Make a Better Model Than OpenAI's o1. Analytics India Magazine. September 2024. (Quoted)
- After Software Engineers, LLMs Are Coming After AI Researchers. Analytics India Magazine. September 2024. (Quoted)
- OpenAI Thinks ChatGPT Thinks. Analytics India Magazine. September 2024. (Quoted).
- Do you think that ChatGPT can reason?. Machine Learning Street Talk. August 2024. (~2hr podcast interview) (Over 67K views)
- What is Stopping Devs from Building an LLM? Analytics India Magazine. Aug 2024. (Research Covered)
- Video showing human-like tiger shows dangers in social media posts. Channel 12 News. August 2024. (Interview).
- Wait! Did OpenAI Just Solve 'Jagged Intelligence'? Analytics India Magazine. August 2024. (Coverage)
- LLMs Have Stalled the Progress of AGI. Analytics India Magazine. August 2024. (Quoted).
- Ray Kurzweil Still Says He Will Merge With A.I. **New York Times**. July 2024.(Quoted)
- The New ChatGPT Offers a Lesson in A.I. Hype. **New York Times**. May 2024. (Covers our research)
- Can Artificial Intelligence Make the PC Cool Again?. **New York Times**. May 2024. (Quoted)
- How AI is Changing Political Landscape. ABC Channel 15. May 2024. (On air TV interview).
- Arizona House votes to regulate deepfakes. Courthouse News. May 2024. (Quoted)
- Don't Be Fooled by A.I. Katy Perry Didn't Attend the Met. **New York Times**. May 2024. (Quoted)

- Even Katy Perry’s mom was fooled by what appeared to be AI-generated Met Gala pics. **NBC News (National)**. May 2024. (Quoted)
- A Mom Was Convinced Her Daughter Was Kidnapped After Hearing Her Voice. Romper. May 2024. (Quoted)
- AI Chatbots Will Never Stop Hallucinating. **Scientific American**. April 2024. (Multiple quotes)
- Is OpenAI’s Sam Altman’s future worth 1M kidnapping scam: “I’ve got your daughter”. New York Post. April 2023.
- A mom thought her daughter had been kidnapped—it was just AI mimicking her voice Popular Science. April 2023. (Quoted)
- ASU research expands artificial intelligence knowledge. ASU News. April 2023. (Research Profiled)
- That panicky call from a relative? It could be a thief using a voice clone, FTC warns. NPR (National Public Radio). March 2023. (Quoted) OpenAI’s GPT-4 Is Closed Source and Shrouded in Secrecy. Vice. March 2023. (Quoted)
- Beauty, lies & ChatGPT: Welcome to the post-truth world. The Hill. Feb 2023 (Opinion Column).
- EntÃ£o tÃ¡, vamos falar sobre o ChatGPT (So, let’s talk about ChatGPT). Communica.ufu.br (Brazilian News Portal). (Quotes The Hill piece). March 2023.
- The State of Artificial Creativity. ASU News. February 2023. (Interview) ChatGPT Worries Professors, Excites Them For Future of AI. State Press. Jan 2023. (Interviewed).
- IIT Jodhpur concludes NAIBS-2023 “Next-gen AI: Inspiration from Brain Science” international conference. NE India Broadcast. Jan 2023. (News Item)
- Subbarao Kambhampati: Avamanamtho.. Kalla neellu thirigaayi. Eenadu (Telugu). 25th December 2022.
Profile in Sunday Magazine.
- How Is Everyone Making Those A.I. Selfies?. **New York Times**. December 2022. (Quoted and research linked)
- Diver Rahita Kaarlu Ippude Kaadu (Self-driving cars are not here yet). Eenadu (Telugu). December 2022.
- What is Augmented Intelligence?. BuiltIn (Explainer article). October 2022.
- Ready or not, mass video deepfakes are coming. **Washington Post**. August 2022. (Quoted)
- The ethical and privacy concerns over deep fakes and AI and our democracy. Cincinnati Edition (NPR/WVXU). August 2022 (Interview)
- Large language models can’t plan, even if they write fancy essays. The Next Web. July 2022. (Coverage of our research). (First appeared in Tech Talks)
- Meet the LA Startup That Lets People Talk to the Dead Dot.LA. August 2022. (Quoted)
- Amazon’s Alexa could soon speak in a dead relative’s voice, making some feel uneasy. NPR (**National Public Radio**). Quoted. June 2022.
- DALL-E Mini and AI Sentience Good Morning Arizona. Channel 3/Phoenix. (Interview Segment) June 2022.

- Amazon's Alexa Can Mimic Any Voice. Morning Wave Busan. Special Interview. June 2022.
- Explainability, Human Aware AI & sentience in large language models. Podcast interview with Jay Shah. June 2022.
- AI as (an Ersatz) Natural Science?. Communications of the ACM (Blog). (Viewpoint) June 2022.
- Dall-e 2: What is the AI image generator creating strange artworks out of nothing and how do you use it?. Independent (UK Newspaper). Quoted. June 2022.
- Meet DALL-E, the A.I. That Draws Anything at Your Command. **New York Times**. April 2022. (Quoted)
- Distinguished career earns Kambhampati Alumnus Award. In the Loop (ASU Fulton). April 2022. (Profile)
- Deepfake, come funzionano e che ruolo hanno nella guerra in Ucraina. . Domani. March 2022. (Quoted)
- Broad and Shallow AI: The Promise and Perils of Competence without Comprehension. The India Forum. Jan 2022. (Column)
- Access for all: the democratisation of AI E&T Engineering and Technology. November 2021. (Quoted).
- Language Imitation Games and the Arrival of Broad and Shallow AI Communications of the ACM Blog. October 2021. (Made it to the top-5 articles list at #1).
- After Covid A.C.10 JTBC TV Channel. South Korea. October 2021. (Interviewed as part of a 3 part documentary on future after Covid--with special attention to AI and--privacy).
- A Stanford Proposal Over AI's 'Foundations' Ignites Debate. Wired. September 2021. (Quoted)
- S2E10: Subbarao Kambhampati with Dhruv Batra on Humans of AI: Stories, Not Stats. Podcast. An hour long interview. September 2021. (Interview)
- Covid-19 May Have Rewired People's Relationships with Artificial Intelligence Forever. Observer. April 2021. (Quoted)
- What's next for iris-recognition systems? ASU Now. March 2021. (Interview)
- Will Artificial Intel get along with us? Only if we design it that way. The Hill. Feb 2021. (Opinion Column)
- Who should stop unethical AI?. **The New Yorker**. Feb 2021. (Quoted)
- Dancing, vacuuming, learning: What's next for robots and their creators?. ASU News. Jan 2021. (Profiles our research)
- Innovators Network Foundation Announces 2020-21 Privacy Fellows. December 2020.
- Samsung AI Forum 2020: Humanity Takes Center Stage in Discussing the Future of AI. Samsung Newsroom. Nov 2020.
- Samsung AI Forum Day 2: Putting People at the Center of AI Development. Samsung Newsroom. Nov 2020. (Covers my talk on Human-Aware AI).
- Why are Artificial Intelligence systems biased?. The Hill. July 2020. (Opinion Column)
- Deepfakes Are Becoming the Hot New Corporate Training Tool. Wired. July 2020. (Quoted)

- Artificial Intelligence Can Predict Our Behavior, But It Had To Adjust To COVID-19. Word. NPR/KJZZ. June 2020. (On Air Interview)
- Enlisting AI in our war on coronavirus: Potential and pitfalls The Hill. March 2020. (Opinion Column)
- Our Deepfake Future. Virginia Quarterly Review. Spring 2020. (Interview in comic form!).
- Trump Proposes a Cut in Research Spending, but a Boost for AI. Wired. February 2020. (Quoted)
- WOULD YOU LIKE TO BE OPERATED [ON] BY ROBOTS? Analytics India Magazine. February 2020 (Tweet quoted).
- AI computing will enter the 'land of humans' in the 2020s: The promise and the peril. The Hill. January 2020. (Opinion Column)
- Facebook says it will ban 'Deepfakes'. **New York Times**. January 2020. (Quoted)
- Is seeing still believing? The deepfake challenge to truth in politics. Brookings Institution. January 2020. (Quoted)
- Fake Trump video? How to spot deepfakes on Facebook and YouTube ahead of the presidential election. USA Today. January 2020. (Quoted)
- Searching for information on Iran's missile strike? Google, YouTube lean on trusted sources; Facebook doesn't. USA Today. January 2020. (Quoted)
- ASU engineering professor elected as ACM Fellow. ASU Now. December 2019. (Profile/News)
- Internet Companies Prepare to Fight the 'Deepfake' Future. **New York Times**. November 2019. (Quoted)
- Perception won't be reality, once AI can manipulate what we see. The Hill. November 2019. (Opinion Column)
- What just happened? The rise of interest in Artificial Intelligence. The Hill. August 2019. (Opinion Column)
- AI Stats News: Data is eating the world quote of the week Forbes. August 2019. (Quoted)
- # Wahl19 #Image: How pictures make politicians Der Standard (Austrian). September 2019. (Quotes our ICWSM 2016 paper).
- 'Deepfake face swap' app ZAO is making people afraid of the future. BoingBoing(Blog). Septemeber, 2019. (Tweet-quoted)
- The Age of Artificial Intelligence . Channel News Asia. August 2019. (Interviewed on the news story)
- Data biases can skew outcomes of AI-based systems. **The Hindu** Business Line. July 2019. (Coverage of my lecture in Hyderabad, India)
- When Atta Halilintar, the YouTube Child (is) Rejected on Twitter. Tirto (Indonesian). August 2019. (Quotes our ICWSM 2016 paper).
- When it comes to #MeToo, is Instagram the new Twitter?". Footwear News(!). July 2019. (Quotes our ICWSM 2016 paper).

- Comments on AI-based Wimbledon highlights; Starts 13th min at the link Details with Andrey Derkach. Voice of America Russian Edition. July 2019.
- AI Conference AAMAS: With warm-heartedness and irony, the robot dialogue succeeds. Heise Online, May 2019 (Coverage of my AAMAS invited talk)
- Renowned AI Expert Subbarao Kambhampati Speaks at PDLs. SMU News. (April 2019) (Coverage of my talk)
- The Rise of AI and Its Challenges Interview on Channel News Asia. April, 2019.
- When is it OK for AI to Lie? ASU Now. January 2019. (Coverage of our research)
- Leading experts in Philosophy, Economics, Law and More Convene at AIES 2019 to Assess Impact of AI on Society. ACM News Release. January 2019. (Our paper is one of the 7 highlighted).
- Pai Advises Real Regulatory Humility on Artificial Intelligence B+C November 2018. (My remarks at FCC covered).
- FCC Chair Calls for Regulatory Restraint for Developing AI. Cablefax. November 2018. (Quoted).
- Google Tweaks Email Program That Assumed An Investor Was Male. NPR. November 2018. (Quoted).
- Que le falta mejorar a la inteligencia artificial Las Ultimas Noticias. (Chile). November 2018. (Coverage of a speech).
- Inteligencia Artificial: Las maquinas que sorprendieron en encuentro internacional realizado santiago. Etstrategia. November 2018. (Coverage).
- Das Silicon Valley will Algorithmen Moral beibringen e Konzerne schlieen sich an Handelsblatt (Germany). October 2018. (Quoted)
- Infra to collect data is the tough part: Subbarao Kambhampati, professor, computer science, Arizona State University . The Economic Times (India). Profile/Interview. June 2018.
- Voices in AI: A conversation with Subbarao Kambhampati. Voices in AI Podcast (audio + transcript). Jun 2018.
- How India is carving out a niche for itself in the field of Artificial. Intelligence. The Economic Times (India). Quoted. June 2018.
- The price of panic: ASU experts ground us in what AI really is and can be . ASU State Press Magazine. (Featured). June 2018. Fear and hope in the age of AI. Fountain Ink. (Long form magazine; India). May 2018. The Trump Administration Finally Looks at Artificial Intelligence Le Temps (French). May 2018.
- Inside the Transformation: Our Relationship with AI (Episode 107H). KQED Television (Bay Area). April 2018.
- Researchers study effects of photo filters on social media The State Press. April 2018. (Coverage of our research)
- Uber Self-Driving Car Fatality Reveals the Technology's Blind Spots. Scientific American. March 2018. (Quoted).
- What Uber's fatal accident could mean for the autonomous-car industry
- Business leaders discuss the future of a digital workforce at ASU event. ASU Now. March 2018. (Coverage)

- The pros and cons of AI ASU Now. March 2018. (Interview)
- Global Governance of AI Roundtable: World Government Summit 2018, The Future Society. March, 2018. (Interview)
- Should AI bots lie? Hard truths about Artificial Intelligence ZDNet.com. (Feb 2018). (Coverage of our research)
- China's AI Imperative Science Magazine. (Feb 2018). (Quoted).
- Questioning AI: What are the key research challenges?" **The Guardian** Science Weekly Podcast. (Jan 2018). (Participant)
- #1: The Landscape of Artificial Intelligence with Rao Kambhampati Vocies of AI. Jan 2017. (Podcast Interview). Sophia: This is downright scam Leiphone. (Jan 2018). (Quoted) [Also on 36KR.com]
- Artificial Intelligence in 2017 still can't understand humans Mashable. December 2017 (Interviewed). Exemplary leadership, significant contributions earn ASU professors AAAS Fellow status ASU Now. December 2017. (Profile about AAAS Fellowship).
- Over a Dozen Outstanding Indian American Scientists Named 2017 AAAS Fellows India West. December 2017. (Mention)
- 19 Indian Americans recognized by American Association for the Advancement of Science News India Times. December 2017. (Mention)
- How Artificial Intelligence is Transforming local business Phonenix Business Journal. (Quoted). November 2017.
- The future of autonomous weapon regulation relies on public awareness Venture Beat. November, 2017. (Quoted)
- AI taught itself to beat us at our own game--what does it mean? ASU Now. (Q&A) November 2017.
- I'm a pacifist, so why don't I support the Campaign to Stop Killer Robots? **The Guardian**. UK. November, 2017. (Op-Ed)
- Artificial intelligence risks GM-style public backlash, experts warn. **The Guardian**. UK. November, 2017. (Quoted)
- Our Relationship with AI KQED Radio. Radio Special. November, 2017. (Radio broadcast of CHM panel)
- How Closer Are We to Robots Replacing Us? Seeker Media Video Segment. How Close Are We? series. Oct 2017. (Featured)
- Elon Musk predicts that A.I. is 'Most Likely Cause of World War III'. Inverse. September 2017 (Reprinted in Yahoo).
- Solar Reflections: Have Machines Eclipsed Humans? Alley Watch, August 2017. (Quoted).
- AI Scientists to Elon Musk: Stop Saying Robots Will Kill Us All. Inverse (interview), July 2017.
- AI Researchers Disagree With Elon Musk's Warnings About Artificial Intelligence Futurism (quoted). July 2017.
- Elon Musk SLAMMED by top researchers for scaremongering over dangers of AI, Daily Express (UK), July 2017.

- Artificial Intelligence Experts Rebut Elon Musk's Warning and Call for Regulation, Christian Post, July 2017.
- Elon Musk Called An Alarmist By AI Experts, Machine Overlords A Far-Fetched Idea EconoTimes, July 2017.
- Elon Musk und KI: Widerspruch aus der Forschung Trends der Zukunft. July 2017. (German)
- Des experts en IA reprochent Elon Musk de se montrer trop alarmiste Numerama. July 2017. (French)
- Yapay zeka araştırmacıları Elon Musk ile aynı fikirdedir Webrazzi, July 2017 (Turkish)
- Le gouvernement américain néglige-t-il l'intelligence artificielle? Le Temps (French). July 2017.
- Voice AI might still have a long way to go TV interview. CGTN Live. TV Segment, Shenzhen, China. July 2017.
- Interview with Leifeng Net. July 2017.
- China press coverage of my plenary presentation at CCF-GAIR, 2017, Shenzhen. Leiphone, July 2017. (Coverage of the talk and Interview.)
- AI Revolution Reshaping Chinese Society. Today (Singapore). June 2017. (Quoted)
- Dewey Decimal Drama. Every Little Thing (Gimlet Media). Interview about biases in Dewey Decimal System and AI. June 2017.
- Artificial intelligence's potential impacts raise promising possibilities, societal challenges. Phys.Org. June, 2017. (Also on ASU Now).
- Coverage of the inauguration of AAAI India Chapter, and invited lecture at IIIT-Hyderabad. Various Indian news media. June 2017.
- Is China Outsmarting America in A.I.?. **New York Times**. May 2017. (Quoted on the proposed U.S. budget and preparing for the future with A.I.)
- Tech majority disagrees with AI warnings from Hawking, Musk and Gates. Network World. May 2017. (Coverage of paper)
- ASU Team Taking Concept for Closer Human-Robot Connection to U.S. Imagine Cup Finals. Full Circle, April 2017 (Coverage of Work) (Version in ASU Now)
- Human-Robot Communication. Cronkite News. April 2017. (Coverage of Work)
- Episode 3: Machines Take Over the World. From robots that paint Rembrandts to Twitterbots gone rogue, Bill and his guests explore the potential and pitfalls of artificial intelligence Panelist. Bill Nye Saves the World. **Netflix Original Series**. April 2017.
- Eminent Astrophysicist Issues a Dire Warning on AI and Alien Life (Quoted). Futurism April, 2017.
- How Trump's 'extreme vetting' hurts science and tech ComputerWorld (Quoted). March, 2017.
- Fighting AI with AI Off Air, Science Friday. March 2017. (Quoted)
- Science Friday segment with Ira Flatow about AI & Security. February, 2017.
- Our Bots, Ourselves. **The Atlantic**. March 2017. (Print edition) (Quoted).
- China's Artificial Intelligence Boom. **The Atlantic**. Feb 2017. (Online) (Quoted).

- AAI and internationalization: An interview Leiphone; Feb 2017. (Interview)
- Asm pretenden Google, Apple y el resto de gigantes educar a la inteligencia artificial; Article on Partnership for AI; El Diario. Feb 2017. (Quoted).
- Coverage about the inaugural board of trustees for the Partnership for AI. (Jan 2017) (TechCrunch, Wallstreet Journal, Quartz, Inverse, MS Power User, The Tech Portal Forbes, Kabir News). Tech Giants Team Up To Tackle The Ethics Of Artificial Intelligence. Quoted in an article on Partnership on AI, NPR All Tech Considered. September, 2016. Why tech giants are forming an AI coalition. Christian Science Monitor. September, 2016. (Quoted) Amazon, Facebook, Google, IBM og Microsoft skal samarbeide om kunstig intelligens, Digi.No (In Norwegian) (Quoted) (September 2016) Gemeinsame Standards ORF.at (Austrian site in German) (quoted) (September 2016)
- How 'human-aware' AI could save us from the robocalypse A write-up on IJCAI 2016 theme in PC World, ComputerWorld etc. July, 2016 The Difference Between Robot's Artificial Intelligence And Humans An interview on KJZZ Public Radio Station; 3/17/2016.
- Experts explain the biggest obstacles to creating human-like robots. Tech Insider. 3/6/2016
- Artificial Intelligence Just Mastered Go, But One Game Still Gives AI Trouble Discover Magazine (quoted in), 1/27/2016.
- Can AI solve information overload? CIO Magazine. (Quoted). 12/29/2015. Old School Photography Techniques are making a comeback: On Instagram (Phoenix New Times, 12/15)
- 18 AI researchers reveal the most impressive thing they've ever seen (TechInsider, 11/15)
- This phenomenon explains what everyone gets wrong about AI (TechInsider, 10/15)
- The future may be full of intentionally clumsy and apologetic robots (TechInsider, 10/15)
- I binge-watched 7 movies about artificial intelligence and the most accurate one was a cartoon (Business Insider, Australia; 8/15)
- 19 A.I. experts reveal the biggest myths about robots (Tech Insider; 10/15).
- HUMAN BRAINS ARE TOO PUNY and 15 more reasons why top researchers are obsessed with AI (Business Insider, India; 10/15)
- Los mejores filtros para que tus fotos sean un ixito en redes sociales, segzn la ciencia BBC Mundo (in Spanish) 6/15.
- Turns Out There Are a Lot of Academics Studying Photo Filters (Wired.com; 5/15)
- Which Instagram filter should you use? Science has the answer (Sydney Morning Herald; 5/15)
- Instagram Is Vast But Its Favorite Themes Are Surprisingly Few, Study Shows (ABC News Online; 10/14)
- Is Twitter Bad For Language? Statistical Analysis Says No (Huffington Post excerpts from the book Dataclysm which profiles our work on Twitter Language (9/14)
- Twitter Analysis of Public Opinion (PBS Horizon, 10/13). Tapping Twitter data for in-depth analysis of public opinion (ASU News; 9/13)
- Professors using blogs, podcasts as teaching tools (The Arizona Republic and other venues; 2/07)

Research Grants (Reverse Chronological Order; Total \$16.3M)

- Unrestricted gift to support research on LLMs and Planning. Qualcomm Research. \$250,000. (5/2024)
- Understanding and Leveraging Planning, Reasoning & Self-Critiquing Capabilities of Large Language Models. AWS AI Amazon Research Awards. \$70,000 cash and \$50,000 AWS credits. (3/2024).
- Engendering & Leveraging Trust in Longitudinal Human-AI Interactions. Office of Naval Research (ONR). PI: Kambhampati. \$626,883. (8/2023-7/2026).
- Automated Extraction and Execution Support for Cognitive Workflows in Finance. J.P. Morgan AI Research Faculty Award. (Unrestricted Gift). \$90,000 (9/1/2021-8/31/2022).
- ACT-NOW: Autonomous Cognitive Technologies for Novelty in Open Worlds. DARPA SAIL-ON Initiative. \$2,349,749 (Co-PI; share 50%) (11/15/2019-6/30/2023)
- Human-Aware AI Assistants for Interactive Decision Support in Finance. J.P. Morgan AI Research Faculty Award. (Unrestricted Gift). \$150,000. (3/1/2019-2/29/2020)
- RAISE: C-Accel Pilot - Track B2 (National Talent Ecosystem): Safe Skill-Aligned On-The-Job Training with Autonomous Systems. National Science Foundation (NSF). 9/1/2019-5/31/2020. \$998,588 (Co-PI. Share 10%).
- Hardened Orchestrated Response for Uncertain Settings (HORUS); DARPA CHASE Program. \$674,000 (10/1/2019-9/30/2022) (Co-PI. Share 50%)
- Human-Aware Planning & Decision Support for Collaborative Complex Decision-Making. Office of Naval Research (ONR). PI: Kambhampati (9/15/2018-9/14/2024). \$1,931,087.
- Proactive Decision Support for Human-in-the-Loop Mission Planning. Office of Naval Research (ONR). PI: Kambhampati (7/1/2018--6/30/2023). \$1,153,076.
- Plan-recognition-Driven Attention Models for Supporting Active Perception in Decentralized Multi-Agent System. Office of Naval Research (ONR). PI: Kambhampati (with Baoxin Li). (3/25/2019-3/24/2023). \$1,193,410. (Share: 40%).
- EAGER: Hierarchical Contrastive Explanations for Robot-Human Communication. National Science Foundation (NSF). PI: Siddharth Srivastava; Co-PI: Subbarao Kambhampati. (9/1/2018--8/30/2020). \$274,581. (Share: 40%)
- Improving Situation Awareness in Distributed Human-Robot Teams. Airforce Office of Scientific Research (AFOSR). PI: Nancy Cook. Co-PI: Subbarao Kambhampati. (3/1/2018--1/31/2022) \$2,014,294.
- Explicable planning and replanning for human-in-the-loop decision support. (Early stage innovations (ESI)). NASA. PI: Kambhampati. (1/15/2017-1/14/2020). \$500,000.
- Planning Challenges in Human-Robot Teaming: An Integrated Exploration of Representations, Algorithms and Human Factors. ONR. PI: Kambhampati Co-PI: Nancy Cooke (8/1/2016-7/31/2019) \$774,376 (Kambhampati share: \$550,782).
- Instrumentation for Research on Planning for Peer-to-Peer Human-Robot Teaming. ONR. PI: Kambhampati. (8/15/2015--8/14/2016), \$282,500.
- Leveraging Automated Planning Technology to Develop Mixed-Initiative Decision Support Tools for the Web. \$51,500. Google Research Award (unrestricted gift). 2016.

- RADAR: A Framework for Human-in-the-loop Planning and Data-based decision support. ONR. PI: Kambhampati (1/1/2015--12/31/2017) \$596,800.
- Video-based Activity Recognition through Tight Integration of Visual Reasoning and Plan Recognition. ONR. PI: LI (Co-PI: Kambhampati) (5/1/2015--4/30/2018) \$753,553.
- Instrumentation for Research on Planning for Human-Robot Teaming. ARO. PI: Kambhampati. (8/1/14-7/31/15) \$112,864.
- Exploratory Research on RADAR: Proactive Decision and Data Support. ONR. PI: Kambhampati. (6/1/14-12/31/15) \$20,000.
- Listening To the Crowd: Automated Alignment, Aggregation and Analysis of Social Media Responses to Public Events \$55,000. Google Research Award (unrestricted gift). 2013.
- Robust Planning in the Presence of Partially Specified Domain Models. ONR. PI: Kambhampati. (1/1/13-12/31/15) \$450,000.
- Long-Term Continual Planning for Remote Human-Robot Teaming in Open Worlds. ONR. PI: Kambhampati. (1/1/13-12/31/15). \$450,000 (Kambhampati share \$349,675).
- Planning for Peer-to-Peer Human Robot Teaming in Open Worlds. ONR. PI: Kambhampati. Co-PI: Nancy Cooke. (5/1/13-4/30/16). \$450,000.
- S3: Suggestions for Suitable Scenarios. ASU subcontract on DARPA Oh By the Way program (subcontract through SIFT). \$341,270, 12/2010-2/1014.
- Relevance and Trust Aware Ranking for the Deep Web. \$50,000. Google Research Award (unrestricted gift). 2011.
- Metaplanning Framework to Evaluate and Select Decision Making Paradigms for UAV Missions. Northrup Grumman Corporation. \$174,524. 1/1/2010--12/31/2010. (PI: Sandeep Gupta).
- Solving Stochastic Planning Problems through Principled National Science Foundation. (Robust Intelligence Medium Collaborative Grant with Purdue and Oregon State) \$328,821. 7/2009--6/2012.
- Foundations of Model-lite Planning. Office of Naval Research. \$342,854. 10/2008--8/2011.
- Source and User Adaptive Information Integration. Office of Naval Research. \$472,589. 10/2008--8/2011.
- Google Research Grant. Unrestricted cash gift. \$50,000. 1/2008.
- Effective human-robot interaction through natural language dialog and dynamic autonomy. ONR Multi-University Research Initiative Award. ASU Share: \$1,000,000. 7/2007-7/2012. (One of three ASU PIs) (Other Institutions: Indiana University, Notre Dame, Stanford).
- Computational Approaches to Creativity Through Goal-Directed Cross-Domain Analogy. National Science Foundation. \$200,000. 8/2007--8/2009.
- ASU subcontract of DARPA Integrated learning program. DARPA (subcontract through Lockheed Martin), \$420,000. 6/2007-12/2008. (Phase 2 of a 4-phase program)
- ASU subcontract of DARPA Integrated learning program. DARPA (subcontract through Lockheed Martin), \$420,000. 6/2006-6/2007. (Phase 1 of a 4-phase program)

- AOC: Archaeological Data Integration for the Study of Long-Term Human and Social Dynamics. National Science Foundation. \$750,000. 1/2007-1/2010. (One of 6 PIs).
- Supporting Partial Satisfaction Planning and Replanning in Expressive and Mixed Initiative Domains. Office of Naval Research. \$300,000. 10/2005-9/2008
- Enabling the Study of Long-Term Human and Social Dynamics: A Cyberinfrastructure for Archaeology. National Science Foundation. Planning grant. \$100,000. (One of 10 PIs). 7/2004-6/2005.
- IBM Faculty Award. Unrestricted cash gift. \$40,000. 7/2004.
- "Scalable Multi-Objective Planning for Metric Temporal Domains: Heuristics, Algorithms and Tradeoffs." National Science Foundation. Sole PI. \$459,042. (1/7/2003-30/6/2006).
- Enabling Technologies for Intelligent Information Integration. Prop 301 Initiative. State of Arizona. \$350,000. (1/1/2003-6/1/2005). (Lead PI; with 4 others).
- "Heuristic control of Plan Synthesis and Execution." NASA Intelligent Systems Initiative. \$162,000. (3/1/2001-2/28/2004).
- Supporting Continual Planning and Replanning in Metric-Temporal Domains" NASA Cross-enterprise Technology. \$483,000. (1/29/2001-1/28/2004). (Open competition among 1000 proposals; 10% award rate).
- Disjunctive planning as a unifying framework for scaling up plan-synthesis. National Science Foundation. 1998. IRI-9801676 \$273,473. (9/1/98-8/31/02).
- Exploring Tradeoffs in Disjunctive Planning as a means to scale-up plan synthesis, Air Force Research Lab, F20602-98-1-0182, \$100,000 (4/3/98-4/3/99). Extended to 7/3/00.
- NYI: Improving the generality and efficiency of AI planning Techniques," **NSF Young Investigator Award**, \$500,000 (\$125,000 base level + Matching funds) IRI-9457634 (8/94-8/2001)
- Exploring the role of constraint propagation in refinement planning. ARPA AASERT Award. DAAH-04096-1-0247. (my share \$227,571, 8/96-8/99).
- Automated Synthesis of Planners and Schedulers: Generic synthesis technology for crisis management. ARPA/ROME Laboratory Planning Initiative grant (with Doug Smith, Kestrel Institute), (my share \$273,000, 9/95-9/98).
- "Automated Annotation for the Analysis and Modification of Plans," ARPA/ROME Laboratory Planning Initiative grant F30602-93-c-0039 (with J.A. Hendler and A. Agrawala, University of Maryland), (my share \$225,000) (3/93-3/96).
- "*Exploring Fundamental Utility Tradeoffs in Plan Reuse*," Research Initiation Award, National Science Foundation, \$90,000, IRI-9210997, (7/92-12/95)