

Pradeep Reddy VARAKANTHAM

School of Computing and Information Systems
Singapore Management University (SMU)
80 Stamford Road
Singapore 178902

Email: pradeepv@smu.edu.sg
Office Phone: (+65) 68280519



Education

PhD, University of Southern California, United States of America, 2007
Master of Science, University of Southern California, United States of America, 2005
Bachelor of Information Systems, International Institute of Information Technology, India, 2002

Academic Appointments

Professor of Computer Science, School of Computing and Information Systems, SMU, Jul 2022 - Present
Associate Professor of Computer Science, School of Computing and Information Systems, SMU, Apr 2021 - Jun 2022
Associate Professor of Information Systems, School of Computing and Information Systems, SMU, Jan 2017 - Mar 2021
Assistant Professor of Information Systems, School of Computing and Information Systems, SMU, Jul 2009 - Dec 2016
Lecturer of Information Systems, School of Computing and Information Systems, SMU, Jun 2009 - Jun 2009

Academic Administrative Positions

Director, CARE.AI Lab, School of Computing and Information Systems, SMU, Jul 2021 - Present
Principal Investigator/Co-PI, AISG Project, School of Computing and Information Systems, SMU, Apr 2021 - Mar 2025
Coordinator, BSc (CS) Artificial Intelligence Track, School of Computing and Information Systems, SMU, Jul 2018 - Present

RESEARCH

Research Interests

Planning under Uncertainty: Markov Decision Problems (MDPs), Partially Observable MDPs, Decentralized MDPs and POMDPs, Influence Maximization, Online Matching

Scheduling under uncertainty: Resource Constrained Project Scheduling Problems (RCPSP) with durational uncertainty and Selfish RCPSPs

Game Theory: Congestion Games, Stochastic Games, Selfish Routing, Behavioral Models

Learning: Reinforcement Learning (RL), Deep Learning and Inverse RL

Publications

Journal Articles [Refereed]

Creating talking points for client advisers at banks to promote sustainable investing, by VARAKANTHAM, Pradeep; MEGARGEL, Alan. (2025). *Journal of AI, Robotics and Workplace Automation*, 3 (4), 350-361. <https://doi.org/10.69554/CMZB1679> (Published)

Zone pAth Construction (ZAC) based approaches for effective real-time ridesharing, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; JAILLET, Patrick. (2021). *Journal of Artificial Intelligence Research*, 70 119-167. <https://jair.org/index.php/jair/article/view/11998/26644> (Published)

Online spatio-temporal matching in stochastic and dynamic domains, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; JAILLET, Patrick. (2018). *Artificial Intelligence*, 261 71-112. <https://doi.org/10.1016/j.artint.2018.04.005> (Published)

Risk-Sensitive Stochastic Orienteering Problems for Trip Optimization in Urban Environments, by VARAKANTHAM, Pradeep; KUMAR, Akshat; LAU, Hoong Chuin. (2018). *ACM Transactions on Intelligent Systems and Technology*, 9 (3), 1-25. <https://doi.org/10.1145/3080575> (Published)

Sampling Based Approaches for Minimizing Regret in Uncertain Markov Decision Processes (MDPs), by AHMED, Asrar; VARAKANTHAM, Pradeep; LOWALEKAR, Meghna; ADULYASAK, Yossiri; JAILLET, Patrick. (2017). *Journal of Artificial Intelligence Research*, 59 229-264. <https://doi.org/10.1613/jair.5242> (Published)

Dynamic Repositioning to Reduce Lost Demand in Bike Sharing Systems, by GHOSH, Supriyo; VARAKANTHAM, Pradeep; ADULYASAK, Yossiri; JAILLET, Patrick. (2017). *Journal of Artificial Intelligence Research*, 58 387-430. <https://doi.org/10.1613/jair.5308> (Published)

An extended study on addressing defender teamwork while accounting for uncertainty in attacker defender games using iterative Dec-MDPs, by SHIEH, Eric; JIANG, Albert Xin; YADAV, Amulya; VARAKANTHAM, Pradeep Reddy; TAMBE, Milind. (2016). *Multiagent and Grid Systems*, 11 (4), 189-226. <http://doi.org/10.3233/MGS-150236> (Published)

Robust execution strategies for project scheduling with unreliable resources and stochastic durations, by FU, Na; LAU, Hoong Chuin; VARAKANTHAM, Pradeep. (2015). *Journal of Scheduling*, 18 (6), 607-622. <http://doi.org/10.1007/s10951-015-0425-1> (Published)

TESLA: an extended study of an energy-saving agent that leverages schedule flexibility, by KWAK, Jun Young; VARAKANTHAM, Pradeep Reddy; MAHESWARAN, Rajiv; TAMBE, Milind; BECERİK-GERBER, Burcin. (2014). *Autonomous Agents and Multi-Agent Systems*, 28 (4), 605-636. <https://doi.org/10.1007/s10458-013-9234-0> (Published)

Coordinating occupant behavior for building energy and comfort management using multi-agent systems, by Klein, Laura; Kwak, Jun Young; Kavulya, Geoffrey; Jazizadeh, Farrokh; Becerik-Gerber, Burcin; VARAKANTHAM, Pradeep; Tambe, Milind. (2012). *Automation in Construction*, 22 525-536. <http://doi.org/10.1016/j.autcon.2011.11.012> (Published)

Robust Local Search for Solving RCPSP/max with Durational Uncertainty, by FU, Na; LAU, Hoong Chuin; VARAKANTHAM, Pradeep; XIAO, Fei. (2012). *Journal of Artificial Intelligence Research*, 43 43-86. <http://www.jair.org/papers/paper3424.html> (Published)

Introducing communication in Dis-POMDPs with locality of interaction, by Tasaki, Makoto; Yabu, Yuichi; Iwanari, Yuki; Yokoo, Makoto; Marecki, Janusz; VARAKANTHAM, Pradeep Reddy; Tambe, Milind. (2010).

Web Intelligence and Agent Systems, 8 (3), 303-311. <http://dx.doi.org/10.3233/WIA-2010-0193> (Published)

Privacy loss in distributed constraint reasoning: A quantitative framework for analysis and its applications, by Maheswaran, Rajiv; Pearce, Jonathan; Bowring, Emma; Varakantham, Pradeep Reddy; Tambe, Milind. (2006). *Autonomous Agents and Multi-Agent Systems*, 13 (1), 27-60. <http://dx.doi.org/10.1007/s10458-006-5951-y> (Published)

Book Chapters

Adjustable Autonomy Challenges in Personal Assistant Agents: A Position Paper, by Maheswaran, Rajiv; Tambe, Milind; Varakantham, Pradeep; Myers, Karen. (2006). In Nickles, Matthias; Rovatsos, Michael; Weiss, Gerhard (Ed.), *Agents and Computational Autonomy: Potential, Risks, and Solutions* (pp. 187-194) Berlin: Springer Verlag. http://dx.doi.org/10.1007/978-3-540-25928-2_15 (Published)

Conference Proceedings

No experts, no problem: Avoidance learning from bad demonstrations, by HOANG, Minh Huy; MAI, Anh Tien; VARAKANTHAM, Pradeep. (2025.0). *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems, San Diego, CA, December 2-7*, (pp. 1-43) United States: Advances in Neural Information Processing Systems. (Advance Online)

On learning informative trajectory embeddings for imitation, classification and regression, by GE, Zichang; CHEN, Changyu; SINHA, Arunesh; VARAKANTHAM, Pradeep. (2025.0). *AAMAS '25: Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, Detroit, Michigan, USA, May 19-23*, (pp. 858-866) New York : ACM. <https://doi.org/10.5555/3709347.3743604> (Published)

EduQate: Generating adaptive curricula through RMABs in education settings, by TIO, Sidney; LI, Dexun; VARAKANTHAM, Pradeep. (2025.0). *AAMAS '25: Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, Detroit, Michigan, USA, May 19-23*, (pp. 2042-2050) New York : ACM. <https://doi.org/10.5555/3709347.3743842> (Published)

Unlocking the planning capabilities of LLMs through maximum diversity fine-tuning, by LI, Wenjun; CHEN, Changyu; VARAKANTHAM, Pradeep. (2025.0). *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), Albuquerque, New Mexico, April 29 - May 4*, (pp. 3318-3340) New Mexico: ACL. <https://doi.org/10.18653/v1/2025.findings-naacl.183> (Published)

On minimizing adversarial counterfactual error in adversarial reinforcement learning, by BELAIRE, Roman; SINHA, Arunesh; VARAKANTHAM, Pradeep. (2025.0). *Proceedings of the Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28*, (pp. 1-25) Singapore: ICLR. <https://openreview.net/forum?id=eUEMjwh5wK> (Published)

Semantic loss-guided data-efficient supervised fine-tuning for safe responses in LLMs, by LU, Yuxiao; SINHA, Arunesh; VARAKANTHAM, Pradeep. (2025.0). *Proceedings of the Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28*, (pp. 52123-52135) Singapore: ICLR. <https://openreview.net/forum?id=kO0DgO07hW> (Published)

Bootstrapping language models with DPO implicit rewards, by CHEN, Changyu; LIU, Zichen; DU, Chao; PANG, Tianyu; LIU, Qian; SINHA, Arunesh; VARAKANTHAM, Pradeep; LIN, Min. (2025.0). *Proceedings of the Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28*, (pp. 1-13) Singapore: ICLR. <https://openreview.net/forum?id=POXfn3OH9G> (Published)

On generalization across environments in multi-objective reinforcement learning, by TEOH, Jayden; VARAKANTHAM, Pradeep; VAMPLEW, Peter. (2025.0). *ICLR 2025: 13th International Conference on Learning Representations, Singapore, April 24-28: Proceedings*, (pp. 83582-83609) Singapore: International Conference on Learning Representations, ICLR. (Published)

Offline safe reinforcement learning using trajectory classification, by GONG, Ze; KUMAR, Akshat; VARAKANTHAM, Pradeep. (2025.0). *Proceedings of the 39th Annual AAAI Conference on Artificial Intelligence, Philadelphia, Pennsylvania, 2025 February 25 - March 4*, (pp. 16880-16887) Philadelphia, Pennsylvania: <https://doi.org/10.1609/aaai.v39i16.33855> (Published)

Marginal benefit driven RL teacher for unsupervised environment design, by LI, Dexun; LI, Wenjun; VARAKANTHAM, Pradeep. (2025.0). *Proceedings of the Thirty-Ninth AAAI Conference on Artificial*

Intelligence (AAAI-25), Philadelphia, Pennsylvania, February 25 - March 4, (pp. 18253-18261) Philadelphia: AAAI. <https://doi.org/10.1609/aaai.v39i17.34008> (Published)

Safety through feedback in constrained RL, by CHIRRA, Shashank Reddy; VARAKANTHAM, Pradeep; PARUCHURI, Praveen. (2024.0). *Proceedings of 38th Annual Conference on Neural Information Processing Systems (NeurIPS 2024) : Vancouver, Canada, December 10-15*, Vancouver, Canada: NeurIPS. (Published)

Improving environment novelty quantification for effective unsupervised environment design, by TEOH, Jayden; LI, Wenjun; VARAKANTHAM, Pradeep. (2024.0). *Proceedings of 38th Annual Conference on Neural Information Processing Systems (NeurIPS 2024): Vancouver, Canada, December 10-15*, (pp. 1-28) Vancouver, Canada: NeurIPS. <https://neurips.cc/virtual/2024/poster/94954> (Published)

SPRINQL : Sub-optimal demonstrations driven offline imitation learning, by HOANG, Minh Huy; MAI, Tien; VARAKANTHAM, Pradeep. (2024.0). *Proceedings of 38th Annual Conference on Neural Information Processing Systems (NeurIPS 2024) : Vancouver, Canada, December 10-15*, Vancouver, Canada: NeurIPS. (Accepted)

IRL for restless multi-armed bandits with applications in maternal and child health, by JAIN, Gauri; VARAKANTHAM, Pradeep; XU, Haifeng; TANEJA, Aparna; DOSHI, Prashant; TAMBE, Milind. (2024.0). *PRICAI 2024: Trends in artificial intelligence: 21st Pacific Rim International Conference on Artificial Intelligence, Kyoto, November 18-24*, (pp. 165-178) Cham: Springer. https://doi.org/10.1007/978-981-96-0128-8_15 (Published)

Imitating cost-constrained behaviors in reinforcement learning, by SHAO, Qian; VARAKANTHAM, Pradeep; CHENG, Shih-Fen. (2024.0). *Proceedings of the 34th International Conference on Automated Planning and Scheduling, ICAPS 2024: Banff, June 1-6*, (pp. 514-522) Banff: Association for the Advancement of Artificial Intelligence. <https://doi.org/10.1609/icaps.v34i1.31512> (Published)

Reward penalties on augmented states for solving richly constrained RL effectively, by HAO, Jiang; MAI, Tien; VARAKANTHAM, Pradeep; HOANG, Minh Huy. (2024.0). *Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence : Vancouver, Canada, February 20-27*, (pp. 19867-19875) Vancouver, Canada: Association for the Advancement of Artificial Intelligence. <https://doi.org/10.1609/aaai.v38i18.29962> (Published)

Imitate the good and avoid the bad: An incremental approach to safe reinforcement learning, by HOANG, Minh Huy; TIEN, Mai Anh; VARAKANTHAM, Pradeep. (2024.0). *Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence, Vancouver, Canada, 2024 February 20-27*, Washington: AAAI. <https://doi.org/10.48550/arXiv.2312.10385> (Published)

Handling long and richly constrained tasks through constrained hierarchical reinforcement learning, by LU, Yuxiao; SINHA, Arunesh; VARAKANTHAM, Pradeep. (2024.0). *Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence, Vancouver, Canada, 2024 February 20-27*, Washington: AAAI. (Accepted)

Unsupervised training sequence design: Efficient and generalizable agent training, by LI, Wenjun; VARAKANTHAM, Pradeep. (2024.0). *Proceedings of the AAAI Conference on Artificial Intelligence 2024: Vancouver, February 20-27: Proceedings*, (pp. 13637-13645) Briarcliff Manor, NY: AAAI Press. <https://doi.org/10.1609/aaai.v38i12.29268> (Published)

Generative modelling of stochastic actions with arbitrary constraints in reinforcement learning, by CHEN, Changyu; KARUNASENA, Ramesha; NGUYEN, Thanh Hong; SINHA, Arunesh; VARAKANTHAM, Pradeep. (2023.0). *Proceedings of the 37th Annual Conference on Neural Information Processing Systems, New Orleans, USA, 2023 December 10-16*, New Orleans: <http://doi.org/10.48550/arXiv.2311.15341> (Published)

Generalization through diversity: Improving unsupervised environment design, by LI, Wenjun; VARAKANTHAM, Pradeep; LI, Dexun. (2023.0). *Proceedings of the 32nd International Joint Conference on Artificial Intelligence, IJCAI 2023: Macao, August 19-25*, (pp. 5411-5419) Washington, DC: AAAI Press. <https://doi.org/10.24963/ijcai.2023/601> (Published)

Transferable curricula through difficulty conditioned generators, by TIO, Sidney; VARAKANTHAM, Pradeep. (2023.0). *Proceedings of the 32nd International Joint Conference on Artificial Intelligence, IJCAI 2023: Macao, August 19-25*, (pp. 4883-4891) Washington, DC: AAAI Press. <https://doi.org/10.24963/ijcai.2023/543> (Published)

Avoiding starvation of arms in restless multi-armed bandit, by LI, Dexun; VARAKANTHAM, Pradeep. (2023.0). *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent*

Systems, London, May 29-June 02, (pp. 1303-1311) Taipei: International Foundation for Autonomous Agents and Multiagent Systems . (Published)

Strategic planning for flexible agent availability in large taxi fleets, by KUMAR, Rajiv; VARAKANTHAM, Pradeep; CHENG, Shih-Fen. (2023.0). *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, London, 2023 May 29 - June 2, (pp. 552-560) New York: ACM.* <https://doi.org/10.5555/3545946.3598683> (Published)

Knowledge compilation for constrained combinatorial action spaces in reinforcement learning, by LING, Jiajing; SCHULER, Moritz Lukas; KUMAR, Akshat; VARAKANTHAM, Pradeep. (2023.0). *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, London, Great Britain, 2023 May 29-June 2, (pp. 860 - 86-868) London: International Foundation for Autonomous Agents and Multiagent Systems. (Published)*

Constrained reinforcement learning in hard exploration problems, by PATHMANATHAN, Pankayaraj; VARAKANTHAM, Pradeep. (2023.0). *Proceedings of the AAAI Conference on Artificial Intelligence, (pp. 15055-15063) Washington DC: Association for the Advancement of Artificial Intelligence.* <https://doi.org/10.1609/aaai.v37i12.26757> (Published)

Future aware pricing and matching for sustainable on-demand ride pooling, by ZHANG, Xianjie; VARAKANTHAM, Pradeep; JIANG, Hao. (2023.0). *Proceedings of the 37th AAAI Conference on Artificial Intelligence, Washington, USA, 2023 February 7-14, (pp. 14628-14636) Washington: AAAI.* <https://doi.org/10.1609/aaai.v37i12.26710> (Published)

Efficient resource allocation with fairness constraints in restless multi-armed bandits, by LI, Dexun; VARAKANTHAM, Pradeep. (2022.0). *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence, (UAI 2022), Eindhoven, Netherlands, 2022 August 1-5 , (pp. 1158-1167) USA: AUAI Press. (Published)*

Joint pricing and matching for city-scale ride pooling, by SHAH, Sanket; LOWALEKAR, Meghna; VARAKANTHAM, Pradeep. (2022.0). *Proceedings of the 32nd International Conference on Automated Planning and Scheduling, Singapore, Singapore, Virtually, 2022 June 13-24, (pp. 499-507) California: AAAI.* <https://doi.org/10.1609/icaps.v32i1.19836> (Published)

Hierarchical value decomposition for effective on-demand ride pooling, by HAO, Jiang; VARAKANTHAM, Pradeep. (2022.0). *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Virtual Conference, May 9-13, (pp. 1-8) Virtual: IFAAMAS. (Forthcoming)*

Field study in deploying restless multi-armed bandits: assisting non-profits in improving maternal and child health, by MATE, Aditya; MADAN, Lovish; TANEJA, Aparna; MADHIWALLA, Neha; VERMA, Shresth; SINGH, Gargi; HEGDE, Aparna; VARAKANTHAM, Pradeep; TAMBE, Milind. (2022.0). *Proceedings of the 36th AAAI Conference on Artificial Intelligence 2022, Virtual, February 22 - March 1, (pp. 12017-12025) Washington, DC: AAAI Press.* <https://doi.org/10.1609/aaai.v36i1.21460> (Published)

Learn to intervene: An adaptive learning policy for restless bandits in application to preventive healthcare, by BISWAS, Arpita; AGGARWAL, Gaurav; VARAKANTHAM, Pradeep; TAMBE, Milind . (2021.0). *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI-21), Virtual Conference, 2021 August 19-26, (pp. 4039-4046) Virtual Conference: IJCAI. (Published)*

CLAIM: Curriculum learning policy for influence maximization in unknown social networks , by LI, Dexun; LOWALEKAR, Meghna; VARAKANTHAM, Pradeep. (2021.0). *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence (UAI 2021), Virtual Conference, July 27-30, (pp. 1-11) Virtual Conference: AAMAS . (Published)*

Adaptive operating hours for improved performance of taxi fleets, by KUMAR, Rajiv Ranjan; VARAKANTHAM, Pradeep; CHENG, Shih-Fen . (2021.0). *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Virtual Conference, May 3-7, (pp. 728-736) Richland, SC: IFAAMAS.* <https://doi.org/10.5555/3463952.3464040> (Published)

Learning index policies for restless bandits with application to maternal healthcare, by BISWAS, Arpita; AGGARWAL, Gaurav; VARAKANTHAM, Pradeep; TAMBE, Milind. (2021.0). *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, (pp. 1467-1468) London: (Published)*

Online traffic signal control through sample-based constrained optimization, by DHAMIJA, Srishti; GON,

Alolika; VARAKANTHAM, Pradeep; YEOH, William. (2020.0). *Proceedings of the 30th International Conference on Automated Planning and Scheduling, ICAPS 2020: Nancy, France, October 26-30*, (pp. 366-374) Menlo Park, CA: AAAI Press. (Published)

Competitive ratios for online multi-capacity ridesharing, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; JAILLET, Patrick. (2020.0). *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, (pp. 771-779) UK: IFAAMAS. (Published)

Solving online threat screening games using constrained action space reinforcement learning, by SHAH, Sanket; SINHA, Arunesh; VARAKANTHAM, Pradeep; PERRAULT, Andrew; TAMBE, Millind. (2020.0). *Proceedings of 34rd AAAI Conference on Artificial Intelligence (AAAI), New York, 2020 February 7-12*, (pp. 1-10) Palo Alto, CA: AAAI Press. <https://doi.org/10.1609/aaai.v34i02.5599> (Published)

Neural approximate dynamic programming for on-demand ride-pooling, by SHAH, Sanket; LOWALEKAR, Megha; VARAKANTHAM, Pradeep. (2020.0). *Proceedings of the AAAI Conference on Artificial Intelligence: 34th AAAI Conference on Artificial Intelligence, February 7-12*, (pp. 507-515) Menlo Park, CA: AAAI press. <https://doi.org/10.1609/aaai.v34i01.5388> (Published)

Correlated learning for aggregation systems, by VERMA, Tanvi, VARAKANTHAM, Pradeep. (2019.0). *Uncertainty in Artificial Intelligence UAI 2019: Tel Aviv, Israel, July 22-25: Proceedings*, (pp. 1-11) Tel Aviv: UAI. (Published)

ZAC: A Zone pAth Construction approach for effective real-time ridesharing, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; JAILLET, Patrick. (2019.0). *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling: Berkeley, CA, July 11-15*, (pp. 528-538) Menlo Park, CA: AAAI Press. (Published)

Entropy based independent learning in anonymous multi-agent settings, by VERMA, Tanvi; VARAKANTHAM, Pradeep; LAU, Hoong Chuin. (2019.0). *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019): Berkeley, CA, July 11-15*, (pp. 655-663) Menlo Park, CA: AAAI Press. (Published)

A homophily-free community detection framework for trajectories with delayed responses, by HAN, Chung-Kyun; CHENG, Shih-Fen; VARAKANTHAM, Pradeep. (2019.0). *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-19), Montreal, Canada, May 13-17*, (pp. 2003-2005) Montreal, Canada: ACM. <https://doi.org/10.5555/3306127.3331991> (Published)

RE-ORG: An online repositioning guidance agent, by KONDA, Muralidhar; VARAKANTHAM, Pradeep; SAXENA, Aayush; LOWALEKAR, Meghna. (2019.0). *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems AAMAS 2019, Montreal, May 13-17*, (pp. 2369-2371) Richland, SC: IFAAMAS. (Published)

A driver guidance system for taxis in Singapore, by JHA, Shashi Shekhar; CHENG, Shih-Fen; LOWALEKAR, Meghna; WONG, Nicholas; RAJENDRAM, Rishikeshan; VARAKANTHAM, Pradeep; TROUNG, Nghia Troung; BIN ABD RAHMAN, Firmansyah. (2018.0). *AAMAS '18: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems, Stockholm, July 10-15*, (pp. 1820-1822) Richland, SC: IFAAMAS. (Published)

Decentralized planning for non-dedicated agent teams with submodular rewards in uncertain environments, by AGRAWAL, Pritee; VARAKANTHAM, Pradeep; YEOH, William. (2018.0). *Uncertainty in Artificial Intelligence, UAI - 18, Monterey, California, US, 2018 July 7-9*, Monterey, California, US: (Published)

Reserved optimistion: Handling incident priorities in emergency response systems, by KONDA, Muralidhar Konda; GHOSH, Supriyo; VARAKANTHAM, Pradeep. (2018.0). *Proceedings of International Conference on Automated Planning and Scheduling, ICAPS-18, Netherlands, Europe, UK, 2018 June 24-29*, (pp. 330-338) Palo Alto: AAAI Press. <https://doi.org/10.1609/icaps.v28i1.13915> (Published)

Bounded rank optimization for effective and efficient emergency response, by MANOHAR, Pallavi; VARAKANTHAM, Pradeep; LAU, Hoong Chuin. (2018.0). *Proceedings International Conference on Automated Planning and Scheduling ICAPS 2018: Delft, Netherlands, June 24-29*, (pp. 375-382) Palo Alto, CA: AAAI Press. (Published)

Upping the game of taxi driving in the age of Uber, by JHA, Shashi Shekhar; CHENG, Shih-Fen; LOWALEKAR, Meghna; WONG, Wai Hin; RAJENDRAM, Rishikeshan; TRAN, Trong Khiem; VARAKANTHAM, Pradeep; TRUONG TRONG, Nghia ; ABD RAHMAN, Firmansyah; (2018.0). *Proceedings of the Thirtieth AAAI*

Conference on Innovative Applications of Artificial Intelligence, New Orleans, 2018, February 2-7, (pp. 7779-7785) Menlo Park, CA: AAAI Press. (Published)

Dispatch guided allocation optimization for effective emergency response, by GHOSH, Supriyo; VARAKANTHAM, Pradeep. (2018.0). *Proceedings of the 32nd AAAI Conference on Artificial Intelligence 2018, February 2-7, New Orleans, (pp. 775-783) Menlo Park, CA: AAAI Press. (Published)*

Mechanism design for strategic project scheduling, by VARAKANTHAM, Pradeep; FU, Na. (2017.0). *Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia, 2017 August 19-25, Melbourne, Australia: (Published)*

Proactive and reactive coordination of non-dedicated agent teams operating in uncertain environments, by AGRAWAL, Pritee; VARAKANTHAM, Pradeep. (2017.0). *Proceedings of the 26th International Joint Conference on Artificial Intelligence IJCAI-17, Melbourne, Australia, August 19-25, (pp. 28-34) Vienna: IJCAI. https://doi.org/10.24963/ijcai.2017/5 (Published)*

Online repositioning in bike sharing systems, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; GHOSH, Supriyo; JENA, Sanjay Dominik; JAILLET, Patrick. (2017.0). *Proceedings of the 27th International Conference on Automated Planning and Scheduling 27th ICAPS 2017, Pittsburgh, PA, June 18-23, (pp. 200-208) Palo Alto, CA: AAAI. (Published)*

Augmenting decisions of taxi drivers through reinforcement learning for improving revenues, by VERMA, Tanvi; VARAKANTHAM, Pradeep; KRAUS, Sarit; LAU, Hoong Chuin. (2017.0). *Proceedings of the 27th International Conference on Automated Planning and Scheduling ICAPS 2017: Pittsburgh, June 18-23, (pp. 409-417) Palo Alto, CA: AAAI Press. (Published)*

Incentivizing the use of bike trailers for dynamic repositioning in bike sharing systems, by GHOSH, Supriyo; VARAKANTHAM, Pradeep. (2017.0). *Proceedings of the 27th International Conference on Automated Planning and Scheduling ICAPS 2017, Pittsburgh, PA, June 18-23, (pp. 373-381) Palo Alto, CA: AAAI Press. (Published)*

Exploiting anonymity and homogeneity in factored Dec-MDPs through pre-computed binomial distributions, by KUMAR, Rajiv Ranjan; VARAKANTHAM, Pradeep. (2017.0). *16th International Conference on Autonomous Agents and Multiagent Systems, AAMAS, (pp. 732-740) Sao Paolo, Brazil: International Foundation for Autonomous Agents and Multiagent Systems.*
<http://www.ifaamas.org/Proceedings/aamas2017/pdfs/p732.pdf> (Published)

Decentralized planning in stochastic environments with submodular rewards, by KUMAR, Rajiv Ranjan; VARAKANTHAM, Pradeep; KUMAR, Akshat. (2017.0). *Proceedings of the 31st AAAI Conference on Artificial Intelligence 2017: San Francisco, February 4-10, (pp. 3021-3028) Menlo Park, CA: AAAI Press. (Published)*

An intelligent system for personalized conference event recommendation and scheduling, by GUNAWAN, Aldy; LAU, Hoong Chuin; VARAKANTHAM, Pradeep; WANG, Wenjie. (2016.0). *ECAI 2016: 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, (pp. 1797-1802) Amsterdam: IOS Press. https://doi.org/10.3233/978-1-61499-672-9-1797 (Published)*

Detecting communities using coordination games: A short paper, by ARAVA, Radhika; VARAKANTHAM, Pradeep. (2016.0). *Proceedings on the 22nd European Conference on Artificial Intelligence: ECAI 2016, The Hague, Netherlands, 2016 August 29 - September 2, (pp. 1752-1753) Amsterdam: IOS Press. http://doi.org/10.3233/978-1-61499-672-9-1752 (Published)*

Robust repositioning to counter unpredictable demand in bike sharing systems, by GHOSH, Supriyo; TRICK, Michael; VARAKANTHAM, Pradeep. (2016.0). *Proceedings of the 25th International Joint Conference on Artificial Intelligence IJCAI 2016: New York, July 9-15, (pp. 3096-3102) Palo Alto, CA: AAAI Press. (Published)*

Scalable greedy algorithms for task/resource constrained multi-agent stochastic planning, by AGRAWAL, Pritee; VARAKANTHAM, Pradeep; YEOH, William. (2016.0). *Proceedings of the 25th International Joint Conference on Artificial Intelligence IJCAI 2016: New York, July 9-15, (pp. 10-16) Palo Alto, CA: AAAI Press. (Published)*

Sequential decision making for improving efficiency in urban environments, by VARAKANTHAM, Pradeep. (2016.0). *Proceeding of the 25th International Joint Conference on Artificial Intelligence: IJCAI 2016, New York; United States; 2016 July 9-15, New York: AAAI Press. (Published)*

Strategic planning for setting up base stations in emergency medical systems, by GHOSH, Supriyo;

VARAKANTHAM, Pradeep. (2016.0). *Proceedings of the 26th International Conference on Automated Planning and Scheduling (ICAPS 2016): London, June 12-17*, (pp. 385-393) Palo Alto, CA: AAAI Press. (Published)

Robust partial order schedules for RCPSP/max with durational uncertainty, by FU, Na; VARAKANTHAM, Pradeep; LAU, Hoong Chuin. (2016.0). *Proceedings of the 26th International Conference on Automated Planning and Scheduling: London, June 12-17, 2016*, (pp. 124-130) Menlo Park, CA: AAAI. (Published)

Robust influence maximization, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; KUMAR, Akshat. (2016.0). *AAMAS '16: Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems, Singapore, May 9-13*, (pp. 1395-1396) Richland, NC: IFAAMAS. (Published)

Robust decision making for stochastic network design, by KUMAR, Akshat; SINGH, Arambam James; VARAKANTHAM, Pradeep; SHELDON, Daniel. (2016.0). *Proceedings of the 30th AAAI Conference on Artificial Intelligence 2016: Phoenix, Arizona, February 12-17*, (pp. 3857-3863) Palo Alto, CA: AAAI Press. (Published)

Online spatio-temporal matching in stochastic and dynamic domains, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; JAILLET, Patrick. (2016.0). *Proceedings of the 30th AAAI Conference on Artificial Intelligence 2016: Phoenix, Arizona, February 12-17*, (pp. 3271-3277) Palo Alto, CA: AAAI Press. (Published)

Solving risk-sensitive POMDPs with and without cost observations, by HOU, Ping; YEOH, William; VARAKANTHAM, Pradeep Reddy. (2016.0). *Proceedings of the AAAI Conference on Artificial Intelligence 2016: Phoenix Arizona, February 12-17, 2016*, Menlo Park, CA: AAAI Press. (Published)

A proactive sampling approach to project scheduling under uncertainty, by VARAKANTHAM, Pradeep; FU, Na; LAU, Hoong Chuin. (2016.0). *Proceedings of the AAAI Conference on Artificial Intelligence 2016: Phoenix, Arizona, February 12-17, 2016*, (pp. 3195-3201) Menlo Park, CA: AAAI Press. (Published)

NLU Framework for voice enabling non-native applications on smart devices, by LANKA, Soujanya; PANTHANIA, Deepika; KUSHALAPPA, Pooja; VARAKANTHAM, Pradeep Reddy. (2016.0). *Proceedings of the 30th AAAI Conference on Artificial Intelligence 2016: Phoenix, Arizona, February 12-17*, (pp. 4359-4360) Menlo Park, CA: AAAI Press. (Published)

PRESS: Personalized event scheduling recommender system (demonstration), by LAU, Hoong Chuin; GUNAWAN, Aldy; VARAKANTHAM, Pradeep; WANG, Wenjie. (2016.0). *15th International Conference on Autonomous Agents and Multiagent Systems: AAMAS 2016, Singapore, 2016 May 9-13*, (pp. 1513-1514) Singapore: International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS). (Published)

Learning and controlling network diffusion in dependent cascade models, by DU, Jiali; VARAKANTHAM, Pradeep; KUMAR Akshat; CHENG, Shih-Fen. (2015.0). *2015 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology WI-IAT 2015: 6-9 December, Singapore*, Piscataway, NJ: IEEE. <https://doi.org/10.1109/WI-IAT.2015.126> (Published)

Incremental DCOP Search Algorithms for Solving Dynamic DCOP Problems, by YEOH, William; VARAKANTHAM, Pradeep Reddy; SUN, Xiaoxun; KOENIG, Sven. (2015.0). *Web Intelligent and Intelligent Agent Technologies, WI-IAT, 2015*, (pp. 1-8) Singapore: IEEE. <http://idm-lab.org/bib/abstracts/Koen15k.html> (Published)

Probabilistic inference based message-passing for resource constrained DCOPs, by GHOSH, Supriyo; KUMAR, Akshat; VARAKANTHAM, Pradeep. (2015.0). *Proceedings of the 24th International Joint Conference on Artificial Intelligence IJCAI 2015: Buenos Aires, Argentina, 25-31 July*, (pp. 411-417) Palo Alto, CA: AAAI Press. (Published)

Dynamic Redeployment to Counter Congestion or Starvation in Vehicle Sharing Systems, by SUPRIYO GHOSH; VARAKANTHAM, Pradeep Reddy; ADULYASAK, Yossiri; JAILLET, Patrick. (2015.0). *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling: Jerusalem, Israel, June 7-11, 2015*, (pp. 79-87) Menlo Park, CA: AAAI Press. <http://www.aaai.org/ocs/index.php/ICAPS/ICAPS15/paper/view/10602> (Published)

DIRECT: A scalable approach for route guidance in Selfish Orienteering Problems, by VARAKANTHAM, Pradeep; MOSTAFA, Hala; FU, Na; LAU, Hoong Chuin. (2015.0). *AAMAS '15: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems: 4-8 May, Istanbul*, (pp. 483-491) Richland, SC: AAMAS. <http://dl.acm.org/citation.cfm?id=2772942> (Published)

Near-optimal decentralized power supply restoration in smart grids, by AGRAWAL, Pritee; KUMAR, Akshat; VARAKANTHAM, Pradeep. (2015.0). *AAMAS '15: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems: Istanbul, Turkey, May 4-8*, (pp. 1275-1283) Richland, SC: IFAAMAS. (Published)

Solving uncertain MDPs with objectives that are separable over instantiations of model uncertainty, by ADULYASAK, Yossiri; VARAKANTHAM, Pradeep; AHMED, Asrar; JAILLET, Patrick. (2015.0). *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence: 25-30 January 2015, Austin, TX*, (pp. 3454-3460) Palo Alto, CA: AAAI Press. (Published)

Risk based optimization for improving emergency medical systems, by SAISUBRAMANIAN, Sandhya; VARAKANTHAM, Pradeep; LAU, Hoong Chuin. (2015.0). *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence: 25-30 January 2015, Austin, Texas*, (pp. 702-708) Palo Alto, CA: AAAI Press. <https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/viewFile/9427/9316> (Published)

Unleashing Dec-MDPs in Security Games: Enabling Effective Defender Teamwork, by Shieh, Eric; Jiang, Albert; Yadav, Amulya; VARAKANTHAM, Pradeep Reddy; Tambe, Milind. (2014.0). *ECAI 2014: 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic: Proceedings*, (pp. 819-824) Amsterdam: IOS Press. <http://teamcore.usc.edu/pubDetails.aspx?id=739> (Published)

Decentralized stochastic planning with anonymity in interactions, by VARAKANTHAM, Pradeep; ADULYASAK, Yossiri; JAILLET, Patrick. (2014.0). *Proceedings of Twenty Eighth AAAI Conference of Artificial Intelligence: 27-31 July 2014, Québec City*, (pp. 2505-2511) Palo Alto, CA: AAAI Press. (Published)

STREETS: Game-theoretic traffic patrolling with exploration and exploitation, by BROWN, Matthew; SAISUBRAMANIAN, Sandhya; VARAKANTHAM, Pradeep; TAMBE, Milind. (2014.0). *Proceedings of Twenty-Sixth Innovative Applications in Artificial Intelligence Conference: 29-31 July 2014, Québec City*, (pp. 2966-2971) Palo Alto, CA: AAAI Press. (Published)

On understanding diffusion dynamics of patrons at a theme park, by DU, Jiali; KUMAR, Akshat; VARAKANTHAM, Pradeep Reddy. (2014.0). *AAMAS '14: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems:Paris, May 5-9*, (pp. 1501-1502) Richland, SC: AAMAS. (Published)

Building THINC: User Incentivization and Meeting Rescheduling for Energy Savings, by Kwak, Jun Young; Kar, Debarun; Haskell, William; VARAKANTHAM, Pradeep Reddy; Tambe, Milind. (2014.0). *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May, 5-9, 2014, Paris*, (pp. 925-932) Richland, SC: AAMAS. <http://aamas2014.lip6.fr/proceedings/aamas/p925.pdf> (Published)

Revisiting Risk-Sensitive MDPs: New Algorithms and Results, by Hou, Ping; Yeoh, William; VARAKANTHAM, Pradeep Reddy. (2014.0). *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling (ICAPS-14): 21-26 June 2014, Portsmouth, New Hampshire*, (pp. 136-144) Palo Alto, CA: AAAI Press. <http://aaai.org/ocs/index.php/ICAPS/ICAPS14/paper/view/7798> (Published)

An agent-based simulation approach to experience management in theme parks, by CHENG, Shih-Fen; LIN, Larry Junjie; DU, Jiali; LAU, Hoong Chuin; VARAKANTHAM, Pradeep. (2013.0). *WSC '13: Proceedings of the 2013 Winter Simulation Conference: December 8-11, 2013, Washington DC*, (pp. 1527-1538) Piscataway, NJ: IEEE. <http://doi.org/10.1109/WSC.2013.6721536> (Published)

Optimization approaches for solving chance constrained stochastic orienteering problems, by VARAKANTHAM, Pradeep; KUMAR, Akshat. (2013.0). *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, Proceedings*, (pp. 387-398) Cham: Springer Verlag. https://doi.org/10.1007/978-3-642-41575-3_30 (Published)

Budgeted personalized incentive approaches for smoothing congestion in resource networks, by VARANTHAM, Pradeep; FU, Na; YEOH, William; CHENG, Shih-Fen; LAU, Hoong Chuin. (2013.0). *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, 2013, Proceedings*, (pp. 375-386) Cham: Springer Verlag. http://doi.org/10.1007/978-3-642-41575-3_29 (Published)

Scalable randomized patrolling for securing rapid transit networks, by VARAKANTHAM, Pradeep; LAU, Hoong Chuin; YUAN, Zhi. (2013.0). *Proceedings of the Twenty-Fifth Innovative Applications of Artificial Intelligence Conference IAAI 2013: Bellevue, WA, July 14-18, 2013*, (pp. 1563-1568) Palo Alto, CA: AAAI Press. <http://www.aaai.org/ocs/index.php/IAAI/IAAI13/paper/download/6203/6436> (Published)

TESLA: An Energy-saving Agent that Leverages Schedule Flexibility, by Kwak, Jun Young; Varakantham, Pradeep Reddy; Maheswaran, Rajiv; Becerik-Gerber, Burcin; Tambe, Milind. (2013.0). *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems, AAMAS'13, May 6-10, 2013, St. Paul, MN, USA.*, (pp. 965-972) Minnesota, USA: (Published)

Lagrangian relaxation for large-scale multi-agent planning, by GORDON, Geoffrey J.; VARAKANTHAM, Pradeep; YEOH, William; LAU, Hoong Chuin; ARAVAMUDHAN, Ajay S.; CHENG, Shih-Fen. (2012.0). *2012 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT 2012: Proceedings, Macau, December 4-7*, (pp. 494-501) Piscataway, NJ: IEEE. <https://doi.org/10.1109/WI-IAT.2012.252> (Published)

Dynamic stochastic orienteering problems for risk-aware applications, by LAU, Hoong Chuin; YEOH, William; VARAKANTHAM, Pradeep; NGUYEN, Duc Thien. (2012.0). *Proceedings of the Twenty-Eighth Conference Conference on Uncertainty in Artificial Intelligence 2012: August 14-18, Catalina Island, CA*, (pp. 448-458) Corvallis, OR: AUAI Press. (Published)

Uncertain congestion games with assorted human agent populations, by AHMED, Asrar; VARAKANTHAM, Pradeep; CHENG, Shih-Fen. (2012.0). *Uncertainty in Artificial Intelligence: Proceedings of the Twenty-eighth Conference: August 15-17, 2012, Catalina Island, United States*, (pp. 44-53) Corvallis, OR: AUAI Press. (Published)

Decision Support for Agent Populations in Uncertain and Congested Environments, by VARAKANTHAM, Pradeep Reddy; CHENG, Shih-Fen; GORDON, Geoff; AHMED, Asrar. (2012.0). *Proceedings of the Twenty-sixth AAAI Conference on Artificial Intelligence, 22-26 July 2010, Toronto, Ontario*, (pp. 1471-1477) Menlo Park, CA: AAAI Press. <http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/5030> (Published)

Lagrangian relaxation for large-scale multi-agent planning, by GORDON, Geoffrey J.; VARAKANTHAM, Pradeep; YEOH, William; LAU, Hoong Chuin; ARAVAMUDHAN, Ajay S.; CHENG, Shih-Fen. (2012.0). *AAMAS '12: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems: June 4-8, Valencia, Spain*, (pp. 1227-1228) Richland, SC: IFAAMAS. (Published)

Prioritized Shaping of Models for Solving DEC-POMDPs, by Varakantham, Pradeep Reddy; Yeoh, William; Velagapudi, Prasanna; Scerri, Paul. (2012.0). *AAMAS '12: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, June 2012, Valencia Spain*, (pp. 1269-1270) Valencia, Spain: ACM. <http://dl.acm.org/citation.cfm?id=2343896.2343957> (Published)

Delayed Observation Planning in Partially Observable Domains, by Varakantham, Pradeep Reddy; Marecki, Janusz. (2012.0). *Proceedings of the Eleventh International Conference on Autonomous Agents and Multi-Agent Systems, published by ACM*, Valencia, Spain: (Published)

Active Malware Analysis using Stochastic Games, by WILLIAMSON, Simon; VARAKANTHAM, Pradeep; GAO, Debin; ONG, Chen Hui. (2012.0). *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012): Valencia, Spain, June 4-8*, (pp. 29-36) Richland, SC: IFAAMAS. <http://dl.acm.org/citation.cfm?id=2343580> (Published)

Distributed model shaping for scaling to decentralized POMDPs with hundreds of agents, by VELAGAPUDI, Prasanna; VARAKANTHAM, Pradeep; SYCARA, Katia; SCERRI, Paul. (2011.0). *AAMAS '11: Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems: May 2-6, Taipei, Taiwan*, (pp. 955-962) Richland, SC: IFAAMAS. (Published)

Incremental DCOP Search Algorithms for Solving Dynamic DCOP Problems, by YEOH, William; VARAKANTHAM, Pradeep; SUN, Xiaoxun; KOENIG, Sven. (2011.0). *AAMAS 2011: Proceedings of 10th International Conference on Autonomous Agents and Multiagent Systems: May 2-6, 2011, Taipei, Taiwan*, (pp. 1069-1070) Richland, SC: IFAAMAS. <http://dl.acm.org/citation.cfm?id=2034422> (Published)

Adaptive decision support for structured organizations: A case for OrgPOMDPs, by VARAKANTHAM, Pradeep; SCHURR, Nathan; CARLIN, Alan; AMATO, Christopher. (2011.0). *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems AAMAS 2011: Taipei, May 2-6*, (pp. 1149-1150) Richland, SC: IFAAMAS. <http://dl.acm.org/citation.cfm?id=2034396.2034461> (Published)

Decentralized Decision Support for an Agent Population in Dynamic and Uncertain Domains, by VARAKANTHAM, Pradeep Reddy; CHENG, Shih-Fen; NGUYEN, Thi Duong. (2011.0). *AAMAS-11: Proceedings of the Tenth International Joint Conference on Autonomous Agents and MultiAgent Systems: Taipei, Taiwan, May 2-6, 2011*, (pp. 1147-1148) Richland, SC: IFAAMAS. <http://dl.acm.org/citation.cfm?id=2034460> (Published)

Effect of human biases on human-agent teams, by PARUCHURI, Praveen; VARAKANTHAM, Pradeep; SYARA, Katia; SCERRI, Paul. (2010.0). *2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology: Toronto, Canada, August 31 - September 3*, (pp. 327-334) Piscataway, NJ: IEEE. <https://doi.org/10.1109/WI-IAT.2010.104> (Published)

A decision theoretic approach to data leakage prevention, by MARECKI, Janusz; SRIVASTAVA, Mudhakar; VARAKANTHAM, Pradeep. (2010.0). *SocialCom 2010: Proceedings of the 2nd IEEE International Conference on Social Computing: PASSAT 2010, 2nd IEEE International Conference on Privacy, Security, Risk, and Trust: Minneapolis, August 20-22*, (pp. 776-784) Los Alamitos, CA: IEEE. <https://doi.org/10.1109/SocialCom.2010.119> (Published)

Towards finding robust execution strategies for RCPSP/max with durational uncertainty, by FU, Na; VARAKANTHAM, Pradeep; LAU, Hoong Chuin. (2010.0). *ICAPS 2010: Proceedings of the Twentieth International Conference on Automated Planning and Scheduling: Toronto, Canada, May 12-16, 2010*, (pp. 73-80) Menlo Park, CA: AAAI Press. <http://www.aaai.org/ocs/index.php/ICAPS/ICAPS10/paper/view/1425> (Published)

Exploiting Coordination Locales in Distributed POMDPs via Social Model, by Varakantham, Pradeep Reddy; Kwak, Jun Young; Taylor, Matthew; Marecki, Janusz; Scerri, Paul; Tambe, Milind. (2009.0). *International Conference on Automated Planning and Scheduling (ICAPS)*, (pp. 51-58) Thessaloniki, Greece: AAAI. <http://aaai.org/ocs/index.php/ICAPS/ICAPS09/paper/view/733> (Published)

Caching Schemes for DCOP Search Algorithms, by Yeoh, William; Varakantham, Pradeep Reddy; Koenig, Sven. (2009.0). *Proceeding: AAMAS '09 Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, (pp. 609-616) Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. <http://portal.acm.org/citation.cfm?id=1558098> (Published)

Linear relaxation techniques for task management in uncertain settings, by VARAKANTHAM, Pradeep; SMITH, Stephen F.. (2008.0). *Proceedings of the 18th International Conference on Automated Planning and Scheduling ICAPS 2008: Sydney, Australia, September 14-18*, (pp. 363-371) Palo Alto, CA: AAAI Press. (Published)

Not All Agents are Equal: Scaling up Distributed POMDPs for Agent Networks, by Marecki, Janusz; GUPTA, Tapana; Varakantham, Pradeep Reddy; Tambe, Milind; Yokoo, Makoto. (2008.0). *Proceedings of the 7th international joint conference on autonomous agents and multiagent systems, AAMAS 2008, Estoril, Portugal, 12-16.05.2008.*, (pp. 485-492) New York (NY): ACM. <http://dl.acm.org/citation.cfm?id=1402453> (Published)

Letting loose a SPIDER on a network of POMDPs: Generating quality guaranteed policies, by Varakantham, Pradeep Reddy; Marecki, Janusz; Yokoo, Makoto; Tambe, Milind. (2007.0). *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi Agent Systems, AAMAS*, ACM. <http://dx.doi.org/10.1145/1329125.1329388> (Published)

Towards efficient computation of quality bounded solutions in POMDPs, by Varakantham, Pradeep Reddy; Maheswaran, Rajiv; GUPTA, Tapana; Tambe, Milind. (2007.0). *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence, IJCAI*, (pp. 2683-2643) <http://portal.acm.org/citation.cfm?id=1625700> (Published)

Winning back the CUP for Distributed POMDPs: Planning over continuous belief spaces, by Varakantham, Pradeep Reddy; Nair, Ranjit; Tambe, Milind; Yokoo, Makoto. (2006.0). *AAMAS'06 : proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems Hakodate, Japan, May 8-12, 2006*, New York: ACM. <http://portal.acm.org/citation.cfm?id=1160683> (Published)

Implementation Techniques for Solving POMDPs in Personal Assistant Domains, by Varakantham, Pradeep Reddy; Maheswaran, Rajiv; Tambe, Milind. (2006.0). *Programming Multiagent Systems: Third International Workshop, ProMAS 2005, Utrecht, Netherlands, July 26, 2005: Revised and Invited Papers*, (pp. 76-89) Berlin: Springer Verlag. http://dx.doi.org/10.1007/11678823_5 (Published)

Conflicts in teamwork: Hybrids to the rescue, by Tambe, Milind; Bowring, Emma; Jung, H.; Kaminka, G.; Maheswaran, Rajiv; Marecki, Janusz; Modi, P. J.; Nair, Ranjit; Pearce, Jonathan; VARAKANTHAM, Pradeep Reddy. (2005.0). *AAMAS 2005 : the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems : Utrecht University, the Netherlands, July 25-29, 2005*, New York: ACM. <http://dx.doi.org/10.1145/1082473.1082474> (Published)

Valuations of Possible States (VPS): A unifying quantitative framework for evaluating privacy in collaboration, by MAHESWARAN, Rajiv T.; PEARCE, Jonathan; VARAKANTHAM, Pradeep; Bowring, Emma; <http://dx.doi.org/10.1145/1082473.1082474> (Published)

Tambe, Milind. (2005.0). *Proceedings of the 4th International Conference on Autonomous Agents and Multi Agent Systems, AAMAS 2005, July 25-29, (pp. 1030-1037) Palo Alto, CA: AAAI Press.* <https://doi.org/10.1145/1082473.1082629> (Published)

Exploiting Belief Bounds: Practical POMDPs for Personal Assistant Agents, by Varakantham, Pradeep Reddy; Maheswaran, Rajiv; Tambe, Milind. (2005.0). *AAMAS 2005 : the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems : Utrecht University, the Netherlands, July 25-29, 2005, (pp. 978-985) New York: ACM.* <http://dx.doi.org/10.1145/1082473.1082621> (Published)

Networked Distributed POMDPs: A Synthesis of Distributed Constraint Optimization and POMDPs, by Nair, Ranjit; Varakantham, Pradeep Reddy; Tambe, Milind; Yokoo, Makoto. (2005.0). *Proceedings of the Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference : July 9-13, 2005, Pittsburgh, Pennsylvania, (pp. 133-139) Menlo Park: AAAI Press.* <http://portal.acm.org/citation.cfm?id=1619356> (Published)

Taking DCOP to the Real World: Efficient Complete Solutions for Distributed Event Scheduling, by Maheswaran, Rajiv; Tambe, Milind; Bowring, Emma; Pearce, Jonathan; VARAKANTHAM, Pradeep Reddy. (2004.0). *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems : AAMAS 2004 : New York City, New York, USA, 19-23 July 2004, New York, N.Y.: IEEE.* <http://dx.doi.org/10.1109/AAMAS.2004.257> (Published)

Conference Papers

Interacting Knapsack Problem in Designing Resource Bundles, by NGUYEN, Truong Huy D.; VARAKANTHAM, Pradeep; LAU, Hoong Chuin; CHENG, Shih-Fen. (2013.0). *Metaheuristics International Conference 10th MIC 2013, August 5-8, Singapore.* (Published)

Regret based Robust Solutions for Uncertain Markov Decision Processes, by Ahmed, Asrar; Varakantham, Pradeep Reddy; Adulyasak, Yossiri; Jaillet, Patrick. (2013.0). *Advances in Neural Information Processing Systems (NIPS), Lake Tahoe, San Francisco.* (Presented)

Analyzing the Impact of Human Bias on Human-Agent Teams in Resource Allocation, by Paruchuri, Praveen; Varakantham, Pradeep Reddy; Sycara, Katia; Scerri, Paul. (2010.0). *International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS), Toronto, Canada.* <http://www.cs.cmu.edu/~pscserri/papers/ParuchuriAAMAS10.pdf> (Published)

Risk-Sensitive Planning in Partially Observable, by Marecki, Janusz; Varakantham, Pradeep Reddy. (2010.0). *International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS), Toronto, Canada.* (Published)

Hybrids in Multiagent Teamwork, by Paruchuri, Praveen; Bowring, Emma; Nair, Ranjit; Pearce, Jonathan; Schurr, Nathan; Tambe, Milind; VARAKANTHAM, Pradeep Reddy; Maheswaran, Rajiv. (2006.0). *Communications of the Computer Society of India,* <http://teamcore.usc.edu/papers/2006/tambe.pdf> (Published)

Magazine Articles

Reimagining education: Keeping the human in the loop, by VARAKANTHAM, Pradeep; TIO, Sidney. (2025, March). *Asian Management Insights, 12(1), 12-19.* (Published)

Closing the data-decisions loop: Deploying artificial intelligence for dynamic resource management, by VARAKANTHAM, Pradeep. (2020, June). *Asian Management Insights, 7(1), 28-33.* (Published)

Artificial Intelligence research in Singapore: Assisting the development of a smart nation, by VARAKANTHAM, Pradeep; AN, Bo; LOW, Bryan; ZHANG, Jie. (2017, October). *AI Magazine, 38(3), 102-105.* <https://doi.org/10.1609/aimag.v38i3.2749> (Published)

Electric Elves: What Went Wrong and Why, by TAMBE, Milind; BOWRING, Emma; PEARCE, Jonathan; VARAKANTHAM, Pradeep; SCERRI, Paul; PYNADATH, David V.. (2008, June). *AI Magazine, 29(2), 23-31.* <https://doi.org/10.1609/aimag.v29i2.2123> (Published)

Working Papers

Zone pAth Construction (ZAC) based approaches for effective real-time ridesharing, by LOWALEKAR, Meghna; VARAKANTHAM, Pradeep; JAILLET, Patrick. (2020). (Published)

Solving Selfish Routing Problems with Involuntary Movements and Bounded Rationality, by VARAKANTHAM, Pradeep Reddy; Ahmed, Asrar; Cheng, Shih-Fen. (2014). Submitted to Artificial Intelligence Journal: (Published)

Uncertain congestion games with assorted human agent populations, by VARAKANTHAM, Pradeep; AHMED, Asrar; CHENG, Shih-Fen. (2013). (Published)

Decision Support for Assorted Populations in Uncertain and Congested Environments, by Varakantham, Pradeep Reddy; Ahmed, Asrar; CHENG, Shih-Fen. (2013). Will be submitting to Journal of Artificial Intelligence Research (JAIR). (Published)

Other Outputs and Contributions

Reports

Towards efficient planning for real world partially observable domains by Varakantham, Pradeep Reddy. (2007). <http://portal.acm.org/citation.cfm?id=1329821> (Published)

Research Grants

Singapore Management University

Trust to Train and Train to Trust: Agent Training Programs for Safety-Critical Environments, AI Singapore Research Programme, AI Singapore , PI (Project Level): Pradeep Reddy VARAKANTHAM , Co-PI (Project Level): Akshat KUMAR, Arunesh SINHA, David LO, 2021, S\$6,086,963.76

The “Other Me” : Human-Centred AI Assistance In Situ, AI Singapore Research Programme, AI Singapore , Co-PI (Project Level): Pradeep Reddy VARAKANTHAM, Akshat KUMAR, 2021, S\$9,551,520

Autonomous Onboarding and Periodic KYC review (PKR), AI Singapore 100 Experiments, AI Singapore , PI (Project Level): Alan MEGARGEL , Co-PI (Project Level): Pradeep Reddy VARAKANTHAM, 2020, S\$625,607.6

Autonomous Prospecting and Product Recommendations, AI Singapore 100 Experiments, AI Singapore , PI (Project Level): Alan MEGARGEL , Co-PI (Project Level): Pradeep Reddy VARAKANTHAM, 2020, S\$625,607.6

MHA-Merlion Initiative – SCDF 02, MHA Merlion Initiative, Ministry of Home Affairs , PI (Project Level): Pradeep Reddy VARAKANTHAM , Co-PI (Project Level): LAU Hoong Chuin, 2019

MHA-Merlion Initiative – SPF 02, MHA Merlion Initiative, Ministry of Home Affairs , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2018

MHA-MERLION INITIATIVE – SPF 02 (Extension), MHA Merlion Initiative, Ministry of Home Affairs , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2017

MHA-MERLION INITIATIVE – SCDF, MHA Merlion Initiative, Ministry of Home Affairs , PI (Project Level): Pradeep Reddy VARAKANTHAM , Co-PI (Project Level): LAU Hoong Chuin, 2016

Moving Beyond Data Insights: Optimizing Dynamics in Safety and Security Networks, Academic Research Fund (AcRF) Tier 2, Ministry of Education (MOE) , PI (Project Level): Pradeep Reddy VARAKANTHAM , Co-PI (Project Level): Akshat KUMAR, 2016, S\$674,046

MHA-MERLION INITIATIVE – SPF 02, MHA Merlion Initiative, Ministry of Home Affairs , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2016

Sequential Decision Making for Matching of Supply and Demand in Mobility Systems, Future Urban Mobility (FM) IRG, Singapore-MIT Alliance for Research and Technology (SMART) , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2016, S\$49,710.7

Urban Computing and Engineering Centre of Excellence, Corporate Laboratory@University, National

Research Foundation (NRF) & Fujitsu (Multiple Funding Sources) , PI (Programme Level): LAU Hoong Chuin , PI (Project Level): CHENG Shih-Fen, Pradeep Reddy VARAKANTHAM, Akshat KUMAR, 2014, S\$18,500,000

Robust Decision Support in Stochastic Selfish Routing (With Applications to Social Transportation), Future Urban Mobility (FM) IRG, Singapore-MIT Alliance for Research and Technology (SMART) , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2013, S\$199,085.17

Decision Support for Taxi Fleet Optimization, Future Urban Mobility (FM) IRG, Singapore-MIT Alliance for Research and Technology (SMART) , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2012, S\$71,500

Decentralised Decision Support to Improve Performance of a Taxi Fleet, Future Urban Mobility (FM) IRG, Singapore-MIT Alliance for Research and Technology (SMART) , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2011, S\$66,000

Planning Against "Intelligent" Adversaries, DSO National Laboratories , PI (Project Level): Pradeep Reddy VARAKANTHAM, 2010, S\$202,200

Other Institutions

Autonomous Onboarding and Periodic KYC Review (PKR), AI Singapore 100 Experiments Programme, AI Singapore PI (Project Level): Alan MEGARGEL, Co-PI (Project Level): Pradeep Reddy VARAKANTHAM, 2020, SGD325,000

Autonomous Prospecting and Product Recommendations, AI Singapore 100 Experiments Programme, AI Singapore PI (Project Level): Alan MEGARGEL, Co-PI (Project Level): Pradeep Reddy VARAKANTHAM, 2020, SGD325,000

TEACHING

Courses Taught

Singapore Management University

Undergraduate Programmes :

Design and Analysis of Algorithms

Introduction to Artificial Intelligence

IS Project Experience (Applications)

IS Project Experience (Research)

Reasoning, Planning and Learning under Uncertainty

Postgraduate Professional Programmes :

Capstone Project

Capstone Project - Data Science and Engineering

Introduction to Artificial Intelligence

Project

Postgraduate Research Programmes :

Dissertation

Empirical Research Project 1
Empirical Research Project 2
Empirical Research Project 3
Empirical Research Project 4
Empirical Research Project I
Empirical Research Project II
Introductory Research Project
Learning and Planning in Intelligent Systems

THESES AND DISSERTATIONS

Theses and Dissertations Supervised

Singapore Management University

Supervisor, "Proactive and Reactive Resource/Task Allocation for Agent Teams in Uncertain Environments", Dissertation by AGRAWAL PRITEE, PhD in Information Systems, Singapore Management University, 2018

Supervisor, "Proactive Resource (Re)distribution for Improving Efficiency in Urban Environments", Dissertation by SUPRIYO GHOSH, PhD in Information Systems, Singapore Management University, 2018

Theses and Dissertations Assessed

Singapore Management University

Committee Member, "Recommending Personalized Schedules in Urban Environments", Dissertation by CHEN CEN, PhD in Information Systems, Singapore Management University, 2017

Committee Member, "Robust Execution Strategy for Scheduling Under Uncertainty", Dissertation by FU NA, PhD in Information Systems, Singapore Management University, 2012