Research Statement

Pradeep Varakantham
School of Computing and Information Systems, Singapore Management University
Tel: (65) 6828-0519; Email: pradeepv@smu.edu.sg
26 (Day) 12 (Month) 2023 (Year)

In an era where digital technology profoundly influences our decisions and behaviors, my aim is to harness the power of AI to guide organizations and individuals towards making effective, wiser, safer, and more informed choices. Thus, my long term goal is to build trustworthy intelligent agent systems (single and multi-agent cooperative/competitive systems) that either autonomously or in conjunction with a human make a sequential set of decisions to optimize quality of life metrics. My aim is to achieve this objective through a combination of computational and behavioural methods rooted in Artificial Intelligence and Machine Learning, Operations Research, Serious Gaming and Behavioural Economics.

Motivation 1:

The first motivation is to blend AI with the nuances of human behavior for greater good as outlined in a few examples below:

- **Scam Prevention through Better Awareness**: Amidst the digital landscape's ever-evolving threats, HumAIn endeavors to fortify the individual's ability to discern between genuine and deceptive communications. By delving into the depths of users' past interactions with digital content, the system is designed to craft and present simulated scenarios. These scenarios are meticulously tailored to train users in distinguishing between scam and legitimate emails/messages, thereby enhancing their cognitive flexibility and building a digital fortress of awareness and vigilance.
- **Customized Educational Experiences in Complex Domains**: In the realm of education, the aim is to first understand each user's current knowledge landscape, particularly in intricate subjects (e.g., Psychology, IRB concepts, or IP Law). The goal is to come up with personalized, scenario-based questions, allowing for interactive and engaging learning experiences. These tailored queries are more than mere questions; they are keys unlocking the doors to deeper understanding and retention of complex concepts.
- **Advancing Public Health through Personalized Nudge**: In health management, HumAIn acts as a compassionate companion, nudging patients along a path lined with doctor-approved health practices. By integrating data from doctor's reports and personal health records, the system thoughtfully curates personalized nudges. These nudges, aligned with medical guidance and the patient's unique health narrative, encourage adherence to beneficial behaviors such as medication compliance and exercise regimens. In doing so, HumAIn nurtures a garden of well-being and health.

Motivation 2:

The second motivation is from decision problems in urban environments, which can be characterised as requiring a match between limited resource supply and an unpredictable demand for resources. Given below are a few practical urban decision problems that motivate my research:

- Taxi fleets: Resource supply corresponds to the available taxis and demand corresponds to customers needing taxis. The goal in this problem is to increase revenues for taxis (or reduce wait times for customers) by continuously matching available taxis to customer demand.

- Emergency response: Resource supply corresponds to ambulances or fire trucks at base stations and demand corresponds to emergency events. The goal in this problem is to reduce response time for emergency events by dynamically moving the "right" ambulances to the "right" base stations.

- Traffic and security patrols: Resource supply corresponds to traffic or security personnel and demand corresponds to potential for traffic violations or security incidents at different nodes in a network. The goal in this problem is to prevent traffic violations and security incidents by reducing predictability in patrols of traffic/security personnel in the network without sacrificing on coverage of "important" locations.

- Theme parks: Resource supply corresponds to attractions and demand corresponds to patrons visiting the attractions. The goal in this problem is to reduce wait times by providing decision support to patrons on visiting the "right" attractions at the "right" times.

- Bike sharing systems: Resource supply corresponds to available bikes at base stations and demand corresponds to customers needing bikes. The goal in this problem is to reduce lost demand due to unavailability of bikes at base stations. We are focussed on lost demand, as it can lead to customers employing private vehicles, which in turn will lead to increased carbon emissions and traffic congestion. A similar problem is relevant to car sharing systems as well.

Given the scale, stochasticity, dynamism and objectives involved in real urban environments, there is a need to develop not only novel models and solution methods, but also new solution concepts that enable trustworthy collaboration with humans (to facilitate successful deployment of methods in real world).

Research Areas and Evolution of Research

Figure 1 provides the different research threads of interest and how my research has evolved over the years. The fundamental research problem has been on sequential decision making (planning/matching/scheduling/learning) under uncertainty. The different threads and the evolution are differentiated along multiple factors : (a) Scale of problems; (b) Cooperative/Competitive nature of decision makers. Cooperative includes environments where there is a single decision maker; (c) Offline/Online nature of decision making; and (d) Trustworthy collaboration with humans.

The major contributions and differentiating factors of my research so far have been in solving sequential decision making problems at a large scale and complexity (thousands of agents, city scale, stochasticity and dynamism). Specifically, we have developed methods and systems for single/multi-agent, offline and online settings. Our approaches have been demonstrated to work on real datasets and in some cases also have been successfully deployed by government agencies (Land Transport Authority, LTA; Singapore Civil Defense Force, SCDF; Singapore Police Coast Guard, PCG; and Singapore Police Force, SPF).

My research has evolved over the years and the different threads are highlighted in different colours.

(1) Single and Multi-Agent Planning/Matching under Uncertainty: Most of my earlier research focussed on small scale problems with stochasticity (worked actively during my PhD, post-doc years and a little bit during my SMU years) or determinism (early part of my PhD years and early part of my SMU years) and over the years has focussed on solving the large scale problems of interest in urban environments. Three of my graduated students (Supriyo Ghosh, Pritee Agrawal, Meghna Lowalekar) have done their thesis on this topic. **References**: [2P, 3P, 7P, 8P, 9P, 10P, 11P, 12P, 15P, 16P, 18P, 19P, 20P, 21P, 22P, 23P, 25P, 30P, 36P, 37P, 38P, 39P, 41P, 43P, 44P, 45P, 46P.]

(2) Single and Multi-Agent Scheduling under Uncertainty: The emphasis on scheduling problems under uncertainty started during my post-doc years and has evolved over time in advising Na Fu during her thesis and post-doc years. It started with single project scheduling and then moved to orienteering problems and multiple cooperative and competing agents. **References**: [4P, 5P, 42P, 16P, 52P, 31P,…]

(3) **Single and Multi-Agent Reinforcement Learning**: Due to the focus on online decision making, a natural framework of relevance is Reinforcement Learning. Due to the scale and special structure in urban environments, we have been able to make concrete contributions in this area. One of the graduated PhD students (Tanvi Verma) thesis was on this topic and currently Rajiv Ranjan KUMAR's thesis is also on this topic. One of the new PhD student will also be working on this topic. **References**: [6P, 24P, 26P, 27P, 28P, 34P, 35P,…]

(4) **Trustworthy and Collaborative Reinforcement Learning**: Through deploying our systems in the real world, we have realised the importance of considering feedback from and to humans while also satisfying robustness constraints. This is the major thread of research on which 2 of the new PhD students will be working on henceforth. **References**: [33P, 29P]

Solution Approaches

We have developed models and approaches that lie at the intersection of Artificial Intelligence, Operations Research, Game Theory and Behavioural Economics to solve scalability issues for these dynamic matching problems. **My solution methods for these urban decision problems are based on exploiting basic properties of the urban environments**:

(a) **Exploiting Homogeneity and Anonymity**: Typically in urban environments, there is homogeneity in supply (e.g., 90% of taxis in Singapore are identical and have same fare structure) and demand components (e.g., customers going from a source to a destination are identical from the perspective of taxis) and more over there is anonymity in interactions between supply and demand (e.g., assigning any one of the two taxis at a taxi stand to a near by customer typically have identical match value).

(b) **Exploiting Limited Influence of Individual Entities (Supply/Demand)**: While there are typically a large number of entities involved in urban environments, the impact of each of them on the overall outcome is typically very small.

(c) **Exploiting abstraction**: Urban decision problems where there is supply demand matching are amenable to abstraction. That is to say, we initially abstract a group of supply components (depending on specific domain properties) into an abstract supply component and create an abstract problem. The sequential decision strategy to match supply and demand can initially be computed for this abstracted problem and then incrementally improved by reducing the abstraction in the problem. We have successfully demonstrated this in the context of bike sharing systems as explained later.

(d) **Exploiting decomposability**: In many of the urban decision problems, it is easy to identify multiple parts of the overall problem that are nearly decomposable. For instance, in bike sharing systems, the problem of moving bikes between stations so as to reduce lost demand and the problem of finding routes for trucks that move bikes are nearly decomposable.

Results

These insights have helped generate concrete results that have been validated against current practice in the real world and against existing approaches. Many of the key results obtained through the use of intelligent systems for dynamic matching of demand and supply of resources are as follows:

- [32P] On ride sharing applications, we have been able to improve over existing best approach by around 15-16%.
- [36P, 38P, 39P] Our approaches for emergency response have been successfully deployed by Singapore Civil Defense Force and have been in use for over 3 years now. We have received the SCDF strategic partner award 2 years in a row. [2018 and 2019]
- [22P, 50P, 40P] In taxi fleets where taxi drivers employ applications (e.g., Uber, Ola etc.) to obtain customers, we have demonstrated the limitations of myopic reasoning adopted in those applications. More importantly, we show that online sequential decision making strategies that

anticipate future demand yield up to 90% of optimal solutions in comparison to 60% with myopic approaches.

- [1P] We were able to reduce the key performance indicator for emergency response systems, the α-quantile response time (α = 0.8) by at least 2 minutes on two real world datasets from Asian cities.

- [11P] In taxi fleets where taxi drivers operate individually and in their own selfish interest, we demonstrated an increase of both average and minimum revenue for taxi drivers, along with an increase in availability of taxis to customers by following equilibrium strategies. Concretely, we demonstrated a revenue increase of SGD 40 per day in expectation for each taxi driver. These results are based on a 2 year dataset of a major taxi company in Singapore.

- [48P] On bike sharing data sets, we demonstrated a reduction of 22% and 45% in lost demand on two real world bike sharing datasets over current practice (repositioning at the end of the day). We also demonstrated a reduction of 10% and 42% in lost demand over an online myopic approach.

- [19P] We improved scalability of decentralised power supply restoration by at least 30 fold.

- [53P] We were able to increase the energy savings in a commercial building by USD 17000 per year in expectation.

References

[1P] *Risk based Optimization for Improving Emergency Medical Response Systems* by Sandhya Saisubramanian, **Pradeep Varakantham** and Hoong Chuin Lau. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-15).

[2P] *Regret based Robust Solutions for Uncertain Markov Decision Processes* by Asrar Ahmed, **Pradeep Varakantham**, Yossiri Adulyasak and Patrick Jaillet. In Proceedings of Neural Information Processing Systems (NIPS-13).

[3P] *Solving Uncertain MDPs with Objectives that are Separable over Instantiations of Uncertainty* by Yossiri Adulyasak, **Pradeep Varakantham**, Asrar Ahmed and Patrick Jaillet. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-15).

[4P] *Dynamic Stochastic Orienteering Problems for Risk-Aware Applications* by William Yeoh, Lau Hoong Chuin, **Pradeep Varakantham**, Huaxing Chen and Duc Thien Nguyen. Twenty Eighth International Conference on Uncertainty in Artificial Intelligence (UAI-12).

[5P] *Optimization Approaches for Solving Chance Constrained Stochastic Orienteering Problems* by **Pradeep Varakantham** and Akshat Kumar. In Proceedings of Algorithmic Decision Theory (ADT-13).

[6P] *Learning and Controlling Network Diffusion in Dependent Cascade Models* by Jiali Du, **Pradeep Varakantham**, Akshat Kumar and Shih-Fen Cheng. In proceedings of Web Intelligence and Intelligent Agent Technology (WI-IAT,15)..

[7P] *Exploiting Coordination Locales in Distributed POMDPs via Social Model Shaping* by **Pradeep Varakantham**, Jun Young Kwak, Matthew Taylor, Janusz Marecki, Paul Scerri, and Milind Tambe. In Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS-09).

[8P] *Distributed Model Shaping for Scaling to Decentralized POMDPs with hundreds of agents* by Prasanna Velagapudi, **Pradeep Varakantham**, Paul Scerri and Katia Sycara. In proceedings of the Tenth International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS-11).

[9P] *Prioritized Shaping of Models for Solving DEC-POMDPs* by **Pradeep Varakantham**, William Yeoh, Prasanna Velagapudi, Katia Sycara and Paul Scerri. In Proceedings of Eleventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-12).

[10P] *Networked Distributed POMDPs: A Synthesis of Distributed Constraint Optimization and POMDPs* by Ranjit Nair, **Pradeep Varakantham**, Milind Tambe, and Makoto Yokoo. In Proceedings of AAAI-05.

[11P] *Decision Support for Agent Populations in Uncertain and Congested Environments* by **Pradeep Varakantham**, Shih-Fen Cheng, Geoff Gordon and Asrar Ahmed. In Proceedings of AAAI Conference on Artificial Intelligence (AAAI-12).

[12P] *Uncertain Congestion Games with Assorted Human Agent Populations* by Asrar Ahmed, **Pradeep Varakantham** and Shih-Fen Cheng. In Proceedings of Uncertainty in Artificial Intelligence (UAI-12).

[13P] *STREETS: Game Theoretic Traffic Patrolling with Exploration and Exploitation* by Matthew Brown, Sandhya Subramanian, **Pradeep Varakantham** and Milind Tambe. In Proceedings of Innovative Applications in Artificial Intelligence at AAAI Conference on Artificial Intelligence (AAAI-14).

[14P] *Scalable Randomized Patrolling for Securing Rapid Transit Networks* by **Pradeep Varakantham**, Hoong Chuin Lau and Zhi Yuan. In Proceedings of Innovative Applications in Artificial Intelligence (IAAI-13).

[15P] *Budgeted Personalized Incentive Approaches for Smoothing Congestion in Resource Networks* by **Pradeep Varakantham**, Fu Na, William Yeoh, Shih-Fen Cheng and Lau Hoong Chuin. In Proceedings of ADT-13.

[16P] *Unleasing Dec-MDPs in Security Games: Enabling Effective Defender Teamwork*, by Eric Shieh, Albert Jiang, Amulya Yadav, **Pradeep Varakantham** and Milind Tambe. In Proceedings of European Conference on Artificial Intelligence (ECAI-14).

[17P] *DIRECT: A Scalable Approach for Route Guidance in Selfish Orienteering Problems* by **Pradeep Varakantham**, Hala Mostafa, Na Fu and Hoong Chuin Lau. In proceedings of International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-15).

[18P] *Dynamic Redeployment to Counter Congestion or Starvation in Vehicle Sharing Systems* by Supriyo Ghosh, **Pradeep Varakantham**, Yossiri Adulyasak and Patrick Jaillet. In proceedings of International Conference on Automated Planning and Scheduling (ICAPS-15).

[19P] *Near-Optimal Decentralized Power Supply Restoration in Smart Grids* by Pritee Agrawal, Akshat Kumar and **Pradeep Varakantham**. In proceedings of International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-15).

[20P] *Decentralised Stochastic Planning with Anonymity in Interactions* by **Pradeep Varakantham**, Yossiri Adulyasak and Patrick Jaillet. In proceedings of AAAI Conference on Artificial Intelligence (AAAI-14).

[21P] *Robust Repositioning to Counter Unpredictable Demand in Bike Sharing Systems* by Supriyo Ghosh, Mike Trick and **Pradeep Varakantham**. In proceedings of International Joint Conference on Artificial Intelligence (IJCAI-16).

[22P] *Online Spatio-Temporal Matching in Stochastic and Dynamic Domains* by Meghna Lowalekar, **Pradeep Varakantham** and Patrick Jaillet. In proceedings of AAAI Conference on Artificial Intelligence (AAAI-16).

[23P] *Scalable Greedy Algorithms for Task/Resource Constrained Multi-Agent Stochastic Planning* by Pritee Agawal, **Pradeep Varakantham** and William Yeoh. In proceedings of International Joint Conference on Artificial Intelligence (IJCAI-16).

[24P] "Learn to Intervene: An Adaptive Learning Policy for Restless Bandits in Application to Preventive Healthcare" by Arpita Biswas, Gaurav Aggarwal, Pradeep Varakantham and Milind Tambe. International Joint Conference on Artificial Intelligence (IJCAI), 2021.

[25P] "Zone pAth Construction (ZAC) based approaches for Effective Real-Time Ridesharing" by Meghna Lowalekar, Pradeep Varakantham and Patrick Jaillet. International Joint Conference on Artificial Intelligence (IJCAI), 2021.

[26P] "CLAIM: Curriculum Learning Policy for Influence Maximization in Unknown Social Networks" by Dexun Li, Meghna Lowalekar and Pradeep Varakantham. Uncertainty in Artificial Intelligence (UAI), 2021.

[27P] "Adaptive Operating Hours for Improved Performance of Taxi Fleets" by Rajiv Ranjan Kumar, Pradeep Varakantham and Shih-fen Chang. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2021.

[28P] "Neural Approximate Dynamic Programming for On-Demand Ride-Pooling" by Sanket Shah, Meghna Lowalekar and Pradeep Varakantham. Published in AAAI 2020.

[29P] "Solving Online Threat Screening Games using Constrained Action Space Reinforcement Learning" by Sanket Shah, Arunesh Sinha, Pradeep Varakantham, Andrew Perrault and Milind Tambe. Published in AAAI 2020.

[30P] "Competitive Ratios for Online Multi-Capacity Ridesharing" by Meghna Lowalekar, Pradeep Varakantham and Patrick Jaillet. Published in AAMAS 2020.

[31P] "Online Traffic Signal Control through Sample-Based Constrained Optimization" by Srishti Dhamija, Alolika Gon, Pradeep Varakantham and William Yeoh. Published in ICAPS 2020.

[32P] " ZAC: A Zone Path Construction Approach for Effective Real-time Ridesharing " by Meghna Lowalekar, Pradeep Varakantham and Patrick Jaillet. Published in International Conference on Automated Planning and Scheduling, ICAPS 2019. Best Application Paper award.

[33P] " Resource Constrained Deep Reinforcement Learning " by Abhinav Bhatia, Pradeep Varakantham and Akshat Kumar. Published in International Conference on Automated Planning and Scheduling, ICAPS 2019.

[34P] " Correlated Learning for Aggregation Systems " by Tanvi Verma and Pradeep Varakantham. Published in International Conference on Uncertainty in Artificial Intelligence, UAI 2019.

[35P] " Entropy based Independent Learning in Anonymous Multi-Agent Settings " by Tanvi Verma, Pradeep Varakantham and Hoong Chuin Lau. Published in International Conference on Automated Planning and Scheduling, ICAPS 2019.

[36P] " Dispatch Guided Allocation Optimization for Effective Emergency Response " by Supriyo Ghosh and Pradeep Varakantham. Published at National Conference on Artificial Intelligence, AAAI-18

[37P] " Decentralized Planning for Non-Dedicated Agent Teams with Submodular Rewards in Uncertain Environments " by Pritee Agrawal, Pradeep Varakantham and William Yeoh. Published at the conference on Uncertainty in Artificial Intelligence, UAI-18

[38P] " Bounded Rank Optimization for Effective and Efficient Emergency Response " by Pallavi Manohar, Pradeep Varakantham and Hoong Chuin Lau. Accepted for publication at the International Conference on Automated Planning and Scheduling, ICAPS-18

[39P] " Reserved Optimization: Handling Incident Priorities in Emergency Response Systems " by Muralidhar Konda, Supriyo Ghosh and Pradeep Varakantham. Accepted for publication at the International Conference on Automated Planning and Scheduling, ICAPS-18

[40P] " Upping the game of taxi driving in the age of Uber " by Shashi Shekhar Jha, Shih-Fen Cheng, Meghna Lowalekar, Nicholas Wong Wai Hin, Rishikeshan Rajendram, Tran Trong Khiem, Pradeep Varakantham, Truong Trong Nghia and Firmansyah Rahman. Published at the conference on Innovative Applications of Artificial Intelligence (IAAI-18)

[41P] "Proactive and Reactive Coordination of Non-dedicated Agent Teams Operating in Uncertain Environments ", by Pritee Agrawal and Pradeep Varakantham. Accepted for Publication at International Joint Conference on Artificial Intelligence, IJCAI-17

[42P] "Mechanism Design for Strategic Project Scheduling ", by Pradeep Varakantham and Na Fu . Accepted for Publication at International Joint Conference on Artificial Intelligence, IJCAI-17.

[43P] "Decentralized Planning in Stochastic Environments with Submodular Rewards ", by Rajiv Ranjan Kumar, Pradeep Varakantham and Akshat Kumar. In Proceedings of the AAAI Conference on Artificial Intelligence, AAAI-17.

[44P] "Exploiting Anonymity and Homogeneity in Factored Dec-MDPs through Pre-computed Binomial Distributions ", by Rajiv Ranjan Kumar and Pradeep Varakantham. In Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS-17.

[45P] "Incentivizing the Use of Bike Trailers for Dynamic Repositioning in Bike Sharing Systems ", by Supriyo Ghosh and Pradeep Varakantham. In Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS-17.

[46P] "Augmenting Decisions of Taxi Drivers through Reinforcement Learning for Improving Revenues ", by Tanvi Verma, Pradeep Varakantham, Sarit Kraus and Hoong Chuin Lau. In Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS-17.

[47P] "Online Repositioning in Bike Sharing Systems ", by Meghna Lowalekar, Pradeep Varakantham, Supriyo Ghosh, Sanjay Dominik Jena and Patrick Jaillet. In Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS-17.

[48P] *Dynamic Repositioning to Reduce Lost Demand in Bike Sharing Systems* by Supriyo Ghosh, **Pradeep Varakantham**, Yossiri Adulyasak and Patrick Jaillet. Journal of Artificial Intelligence (JAIR).

[49P] "Zone pAth Construction (ZAC) based Approaches for Effective Real-Time Ridesharing" by Meghna Lowalekar, Pradeep Varakantham and Patrick Jaillet. Accepted for publication at Journal of Artificial Intelligence Research (JAIR), 2021. JAIR Award Track.

[50P] " Online Spatio-Temporal Matching in Stochastic and Dynamic Domains " by Meghna Lowalekar, Pradeep Varakantham and Patrick Jaillet. Published in Artificial Intelligence Journal (AIJ), 2018

[51P] " Sampling based Approaches for Minimizing Regret in Uncertain Markov Decision Problems (MDPs) ", by Asrar Ahmed, Pradeep Varakantham, Meghna Lowalekar, Yossiri Adulyasak and Patrick Jaillet. Accepted for Publication at Journal of Artificial Intelligence Research (JAIR)

[52P] "Risk-Sensitive Stochastic Orienteering Problems for Trip Optimization in Urban Environments ", by Pradeep Varakantham, Akshat Kumar, Hoong Chuin Lau, William Yeoh. Accepted for Publication at Transactions on Intelligent Systems and Technology (TIST)

[**53P**] "TESLA: An Extended Study of an Energy-saving Agent that Leverages Schedule Flexibility;", by Jun Young Kwak, Pradeep Varakantham, Rajiv Maheswaran, Yu-Han Chang, Milind Tambe, Burcin Becerik-Gerber and Wendy Wood. In the Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)