Artificial Intelligence Operations Research Workshop 2 (August 16-17, 2022)

Main Takeaways

Workshop Organized by

John Dickerson, University of Maryland
Bistra Dilkina, University of Southern California
Yu Ding, Texas A&M University
Swati Gupta, Georgia Institute of Technology
Pascal Van Hentenryck, Georgia Institute of Technology
Sven Koenig, University of Southern California
Ramayya Krishnan, Carnegie Mellon University
Radhika Kulkarni, SAS Institute, Inc. (retired)

With Support From

Catherine Gill, Computing Community Consortium (CCC)
Haley Griffin, Computing Community Consortium (CCC)
Maddy Hunter, Computing Community Consortium (CCC)
Ann Schwartz, Computing Community Consortium (CCC)

On August 16th, 2022, the Computing Community Consortium (CCC), INFORMS, and ACM SIGAI held the second Artificial Intelligence/Operations Research Workshop in Atlanta, GA. This workshop, which is the second of a three-part series, focused on exploiting the synergies of the AI and OR communities and envisioning brainstorming activities to engage both communities and solve outstanding societal problems.

The workshop included four sessions focused on the foundational elements of trustworthy AI: fairness, explainable AI/causality, robustness/privacy and human alignment/human-computer interaction. The speakers and participants were selected from computer science and O.R. communities to foster a healthy exchange of ideas between the two groups. The workshop concluded with discussing a few grand challenges that are worth pursuing jointly by AI and OR researchers working together to overcome current challenges, described on the back:

We discussed the disparities in fairness between users of ridesharing services, specifically in the city of New York. Transportation Network Companies (TNC's) such as Uber, Lyft, and taxi services have been observed to provide less reliable service in historically disadvantaged areas, because these areas usually have lower rider demand leading to longer wait times between trips for drivers. To rectify this inequity, financial incentives need to be put in place to encourage more drivers to travel to these areas. Who pays for these incentives, however, is the main point of contention, whether it be the rideshare riders or the TNC's themselves. To calculate who optimally ought to pay at any given time, and how to best allocate drivers to riders in this multi-agent setting and balance multiple objectives, we propose a competition be created with teams vying to develop the most efficient and equitable, and trustworthy algorithm combining AI and OR practices.

We also recommend creating a summer program, convening Artificial Intelligence and Operation Research experts to educate Ph.D. students on multidisciplinary training in AI and OR. This program would take place annually, similar to Machine Learning summer workshops which have taken place in the past. The summer program could focus on a number of topics on which research has been done in both AI and OR (such as decision making under uncertainty, local search, etc.) and invite both an AI and an OR speaker for each topic. This way, the Ph.D. students understand the techniques that have been developed in both AI and OR as well as their commonalities and differences. The Ph.D. students could work in interdisciplinary teams to solve a challenge problem that requires the integration of AI and OR techniques, to allow these students to overcome some of the challenges listed above.

We also put forth the idea of a joint AI and OR conference on decision making for intelligent robots. OR often uses its optimization techniques to support human decision makers (in business settings), while AI often focuses on autonomous decision making by agents. Robots, for example, must plan their motions and tasks, both individually and as a team, which are independently complex optimization problems but also need to be coordinated. To create robots which are as efficient, accurate, and coordinated as possible, a convergence of AI and OR expertise will be required.

Finally, we suggest another challenge problem which may demand a joint AI/OR approach. The setting is scheduling of nursing and physician staff in healthcare facilities. Healthcare delivery settings have seen significant fluctuations in demand patterns from pre-COVID times, to during COVID and post-COVID, resulting in increased variability of patterns for emergency care, out-of-hospital care, and clinic care. The classical OR version of this problem is an open-loop problem, that is, resources (staff) are assigned to rosters to meet (deterministic or stochastic) demand. However, in reality the process is often iterative, that is, staff often have changing priority over days or weeks of the roster, along with varying demand, which may require that the schedule be re-calculated dynamically to make it human-friendly. To better predict healthcare demands while creating optimal and fair schedules for healthcare staff is a challenge that could be addressed by AI and OR approaches.