Artificial Intelligence Operations Research Workshop 1 Report Out

Workshop

September 23-24, 2021

Workshop Organized by

Sanmay Das, George Mason University
John Dickerson, University of Maryland
Pascal Van Hentenryck, Georgia Tech
Sven Koenig, University of Southern California
Ramayya Krishnan, Carnegie Mellon University
Radhika Kulkarni, SAS Institute, Inc.
Phebe Vayanos, University of Southern California

With Support From

Ann Schwartz, Computing Community Consortium (CCC)
Daniel Lopresti, Leigh University and CCC
Madeline Hunter, Computing Community Consortium (CCC)

Background and Introduction

Artificial intelligence (AI) received significant attention in recent years, primarily due to breakthroughs in game playing, computer vision, and natural language processing that captured the imagination of the scientific community and the public at large. Many businesses, industries, and academic disciplines are now contemplating the application of AI to their own challenges. INFORMS and ACM SIGAI joined together with the Computing Community Consortium (CCC) to organize a series of three workshops. The objective is to explore ways to exploit the synergies of the AI and Operations Research (OR) communities to transform decision making. The aim of the workshops is to establish a joint strategic research vision for AI/OR that will maximize the societal impact of AI and OR in a world being transformed by technological change and a heightened desire to tackle important societal challenges such as growing racial and social inequity, climate change, and sustainable solutions to our food-water and energy needs.

The two communities have complementary strengths, though synergies have started to emerge. In its quest to build intelligent agents, the AI community often incorporates methods and technologies from other disciplines. Machine learning has adopted, refined, and expanded optimization algorithms that originated from the OR community, including stochastic gradient descent and convex optimization. The hybridization of constraint programming, mathematical programming, and satisfiability has been an active research topic in both communities for more than two decades, resulting in massive improvements in optimization solvers. AI also has much to offer to the OR community. Machine learning is bound to transform the next generation of optimization solvers and create a new generation of optimization algorithms. The cross-fertilization between AI and OR may also create breakthroughs in reinforcement learning and multi-agent coordination and optimization. These methodological advances combined with the specific needs of societal problems and contributions from allied social sciences (e.g., causal inference, economics) are bound to transform the art and science of modeling, solving and deploying solutions to business and societal problems.

However, there are barriers and difficulties in realizing this vision due to cultural differences between AI and OR. Some of the challenges are obvious. The two fields use different vocabulary for the same concepts, e.g., reinforcement learning in AI and (approximate) dynamic programming in OR. The missions of the two fields are also quite distinct: whereas AI focuses on building intelligent agents, OR traditionally focuses on process improvement to enhance quality of service or efficiency or to develop deep causal understanding of the state of a problem prior to designing and evaluating interventions. The two fields also focus on different application areas, which makes communication and knowledge transfer more difficult.

The workshop series aims to overcome these difficulties and to provide a stepping stone for a strong and sustained collaboration between the two fields. It is organized around a series of three workshops whose respective goals are: (1) articulating a strategic vision; (2) implementing a strategic vision: and (3) progressing the strategic

vision. The first of these workshops was held virtually on September 23-24, 2021, with the following goals:

1. to review the state of AI/OR along several axes;
2. to articulate a number of high-level research and education opportunities for each of these axes;
3. to identify potential grand challenges of mutual interest that would bring research and industry together with advances on these challenges powered by benchmark data sets and competitions (e.g., the ARPA-E GO competition);
4. to select a few topics for summer schools and research programs to foster AI/OR collaborations.

Grand challenges may be fundamental societal problems for which an AI/OR approach has significant promise and/or scientific advances that would fundamentally transform both fields. The grand societal challenges may be in areas such as resilient supply chains, sustainable energy, health care and crisis management, equitable transportation, and the modeling of human behavior on digital and physical platforms. The workshop invited thought leaders from both fields and balanced a variety of considerations to create a stimulating, diverse, and inclusive forum for creating this joint vision.

This report out should be considered an interim report with a final report to be published after the end of the third workshop. We expect the next two workshops to be held in the first half of 2022. Note that the workshops are held on behalf of the community, by invitation only. As with the first workshop, the intention is to involve as diverse and representative a group of participants as possible within the constraints of attendance.

Opening Remarks

Prof. Ramayya Krishnan
Opening Remarks Video Recording
Opening Slide Presentation

Krishnan's opening remarks covered three topics. He began with the brief description of the AI initiative that began during his Presidency of INFORMS and how the work done by colleagues on the AI strategy committee at INFORMS led to the outreach and discussions with ACM SIGAI and CCC which culminated in these AI/OR workshops. His substantive remarks began with definitions of AI (Russell and Norvig) and Operations Research which clearly points to the close relationships between the two fields. Both AI and OR are “umbrella fields” and have adapted and used tools, methods and models from adjacent fields of decision theory, economics and probability/statistics to name a few.

Krishnan noted that historically the two fields focused on different applications (OR on business critical and societal domains like supply chains, clearing energy markets vs. highly visible consumer facing AI apps like image recognition, google translate, games), recognized achievements and advances differently and historically developed different ways of disseminating advances (e.g., journals vs, conferences, open-source code vs. closed, benchmark data sets and competitions). Recently, there is greater recognition of the synergy between the fields and more opportunities such as joint conferences have emerged (e.g., CPAIOR), though much remains to be done.

Krishnan then highlighted three opportunities for the communities to collaborate. The first opportunity is well underway and that is collaboration at the methods layer. There has been tremendous progress in the AI community (specifically, ML) by way of open-source libraries and projects (e.g., Apache Spark) and creating python callable libraries of OR and AI algorithms is one concrete way for the communities to collaborate and build on each other's progress on methods. The second collaboration opportunity is at the data layer. Just like Imagenet pioneered advances in new methods for image recognition, data sets (synthetic, simulated or curated) are required to power progress in societal domains such as supply chain management and sustainable energy to name a few. ARPA-E, the synthetic energy grid data sets and the GO competition offer a good model for what is possible. Advances at the data layer will pioneer advances in the methods layer as well. Krishnan concluded with a call to collaborate on the “policy” layer with the idea that AI and OR can learn by collaborating with the social sciences to extend the art and science of policy interventions based on a causal understanding of the decision variables and their inter-relationships. Using a child welfare decision support example, Krishnan pointed out the opportunities for AI and OR to create new

algorithmic decision tools whose use could be evaluated as decision support intervention using tools like randomized controlled trials and A/B testing. These themes of collaboration between the AI and OR communities were reinforced in his remarks with Sven Koenig during the closing session.

Session 1: Methods

This session set the stage with thought-provoking presentations from leading researchers from OR and AI with a review of key technological innovations at the intersection of OR and AI. Two of the presentations were on the intersection between AI and optimization, with a focus both on how AI can help optimization algorithms perform better without sacrificing rigor and correctness, as well as on contributions that optimization techniques from OR can make in improving the state of the art in AI. There was discussion on interface areas like multi-agent systems, which require practical advances in areas like game theory and planning, and also on developments in areas like sequential decision-making and reinforcement learning which have co-evolved in the AI and OR communities, with some differences in approach and techniques, but as importantly, in language used to communicate about the area.

Session 1 Video Recording

Slide Presentations
Stephen Wright - Research in OR & AI: Continuous Optimization
Andrea Lodi - Machine Learning for Discrete Optimization
Katia Sycara - Multi-agent Systems

Breakout Sessions Summary
Breakout 1 – Incentives for Collaboration + Overcoming Cultural Differences
Moderator: Katya Sheinberg
Notetaker: Zihao Yang

- Discussed key obstacles preventing collaboration between AI and OR. Some key problems identified were publications happening in different venues (OR in journals and AI in conferences), lack of incentive structure to collaboratively publish, and not recognizing each other's contributions and therefore not knowing about each other's work.
- Some solutions are to: encourage both areas to publish in both venues, encourage attendance to each other's conferences, ensure that faculty in departments should recognize publications in both areas, and encourage PhD students to take courses across departments to help them to understand languages from different communities.

- More specific solutions: IEEE Journals might work as a venue to publish work from both communities due to the fact that they have a similar rigor as OR publications and also allow cross track review process; holding competitions with AI and OR techniques might illuminate problems such as community differences and language barriers.

Breakout 2 – Multi-Disciplinary Collaborations
Moderator: Maria Gini
Notetaker: Jiaoyang Li

- We discussed how to attract people from both AI and OR to attend each other's conferences. We first looked at some current conferences: INFORMS is not very popular for AI researchers; NeurIPS has successfully attracted some top OR researchers; and CPAIOR is a joint venue but has a very small community. We then talked about several possible ways to improve this. Can we get names of the speakers from one area to be invited by the conferences of the other area? Can, for example, INFORMS fund OR speakers to attend AI conferences and similarly AI conferences fund the opposite? Can we get funding for joint workshops or summer schools?
- We also discussed how to overcome the language barrier. We need the community to agree on standard language mapping. We can involve senior members to help in activities with students to map the language. Examples include doing an overview of sets of topics for the other community, creating tutorial slides and writing papers to describe problems that can be shared.
- We also discussed how to get students to work in the joint area of AI and OR. We want to lower the barrier for the students to join the joint area by, for example, shared course materials, summer schools, and the language mapping mentioned above. NSF might have educational funding that would be appropriate.

Breakout 3 – Interface Areas
Moderator: Kate Larson
Notetaker: Enpeng Yuan

We discussed several areas where it was believed that AI-OR collaboration would benefit and support progress in each area.

- Insights from AI to enrich OR: We discussed a number of examples where insights and progress in AI might be useful in better understanding and addressing a number of optimization problems. This included whether it might be possible to learn when it is worthwhile to solve an optimization problem, learning from similarly solved problems, and recognizing when learning has promise with respect to optimization problems and when it does not. We also discussed whether the discrete optimization

community might be able to gain insights from the types of problems where deep learning works well in order to better recognize when structure is and is not important.
- Insights from OR to enrich AI: We also discussed how OR can enrich AI. Examples we considered included whether insights from OR could help guide architecture design decisions in ML, as well as a recognition that interactions and collaborations with the behavioral OR community would be useful for work on Human-AI teams and AI for supporting human-decision makers more broadly.
- Insights required from both fields: We identified concerns around fairness and fair machine learning as a topic of great interest inside the AI community recently. It was observed that there can be a gap between the abstract models of fairness that have been studied and what is needed to operationalize them in applications. Furthermore, how one should incorporate fairness directly into decision models is an interesting open question which needs insights from both disciplines.

Breakout 4 – Non-Standard Research Areas
Moderator: Ilya Hicks
Notetaker: Amanda Kube

- We discussed several projects of interest: relations to combinatorial optimization and causal inference and building better debugging tools for detecting problems before and after deployment whether they are automatic, semi-automatic, or continuously monitored.
- We also discussed challenges of deployment including how to find problems with an optimization function that is based on theory when your theory may have issues (what constitutes a deviation and how much deviation is acceptable) and how to maintain consistency when the world is always changing (when have we changed to a new equilibrium).
- We also talked about learning from near misses and the use of counterfactual reasoning to learn what would have happened if we had done it differently. In addition, finding if the problem was the objective function or the uncertainty set. Relating to the constant changing of the real world, there could be changes in the environment or environmental unknowns that are unaccounted for.

Session 2: Applications

The first applications session covered key emerging application areas enabled by modeling, algorithmic, and technological innovations at the interface of AI and OR. We heard from leading researchers experienced in working with organizations to help them tackle some of the most important problems they face. Application domains and partners included:

- Commercial applications: problems related to supply chain resiliency (e.g., with Ford Motor company), inventory, transportation, and procurement optimization (e.g., at Mango), supply chain digitization (e.g., at InBev), price optimization (e.g., Rue-La-La), personalized offering (e.g., Ryanair), online resource allocation (e.g., IBM), and routing, pickup and delivery operations (e.g., at UPS).
- Societal and policy-making applications: problems related to public housing allocation (e.g., at LAHSA), substance use prevention (e.g., at Chestnut health systems), biodiversity conservation (e.g, at Panthera).

The permeating theme was the need to blend data-driven approaches with OR modeling, the importance of accounting for end user needs and/or values to ensure adoption of the solutions, and the challenges and opportunities in working with real data (e.g., offline and/or online data, observational data). Speakers showcased concrete improvements in the organizations they worked with that were brought by this research. The power of blending AI and OR techniques to design interpretable and fair solutions that align with human value judgements and that do not exacerbate and even correct existing biases was also showcased.

Session 2 Video Recording

Slide Presentations
David Simchi-Levi - The MIT Data Science Lab
Ranga Nuggehalli - UPS ORION Project: A Case Study in Behavioral Optimization
Robert Hampshire (Slide show presentation and video recording omitted for privacy reasons)
Phebe Vayanos - Data-driven Decision-making for Social Impact

Breakout Sessions Summary
Breakout Room 1 – Typical Application Domains
Moderator: Swati Gupta
Notetaker: Zihao Yang

- We discussed domains where AI and OR are applied. In our discussion we noticed that boundaries are merging significantly. In OR, to name a few, there is transportation, energy, power systems, supply chain management, finance, manufacturing, workforce management, forestry, food production, oil and gas, disaster management, telecom, high performance computing, advertising, ship designing and emerging applications like quantum computing. In AI there are applications in healthcare, robotics, manufacturing, customer service (chat bots). There is a lot of increasing clinical decision making, NLP in marketing, machine

translation, retail, recommendation systems, etc. There is increasing overlap between the industries covered by AI and those covered by OR.
    - There are a lot of techniques AI researchers can bring to OR domains. Specifically, transportation is creating a lot of real time multivariate data, data in multiple dimensions and infrastructures and the need to work around how these different components interact with each other. AI can help optimize OR systems by understanding the data and understanding the spatial temporal forecasting models.
    - We discussed that there is a growing trend looking at equity finance and accessibility in OR and AI systems**.**

Breakout Room 2 – Benchmark Data
Moderator: Thiago Serra
Notetaker: Jiaoyang Li

- We discussed the challenges of obtaining data from industry practitioners. The bar of proof of privacy is high. Sometimes even if one can provide a theoretical proof, it is still hard to convince the companies that it is safe and will not reveal their privacy. One possibility is to bring solvers to (private) data instead of bringing data to solvers, although it has a limitation that this method is more suitable for model-driven approaches.
- We also discussed the risks about the benchmark-driven culture as people keep resolving the same (old) instances using the same metrics. What if these instances and metrics are misleading? Is there consensus on the benchmarking and evaluation methodology in AI/OR?
- We also discussed similar issues for solvers instead of benchmark data. Popular OR solvers, like Gorubi/CPLEX and even SCIP are not permissively open-source. Only COIN-OR has an EPL (Eclipse public license). This raises challenges for research.

Breakout Room 3 – Automation vs. Human in the Loop
Moderator: Ramayya Krishnan
Notetaker: Enpeng Yuan

- Algorithms can recommend decisions to humans. To improve algorithm adoption, it helps to communicate the cost of not following the algorithm’s decisions to its users, and train/educate the algorithm users. Visualization is a powerful tool to help non-technical people to understand the solutions, since business leaders only follow interpretable decisions. However, quantifying a human being is hard, so “human in the loop” may be inevitable.
- Need clear explainability measures.
- Can use model distillation to replace a complex model with a simpler one.

- AI-made decisions could seem more fair than human-made decisions. A study of a factory in China shows that workers are happier to receive scheduling instructions from a machine than a human.
- Survey shows that people are willing to rely on AI for scheduling/planning, but not for decision-making.

Breakout 4 – Collaboration at the Policy Level
Moderator: Amy Greenwald
Notetaker: Amanda Kube

- At the policy level/outside academia, it seems that AI is more well-known than OR. We discussed why this might be. Perhaps it is due to the global competition surrounding AI, leading to the allocation of more government resources. It might also be due to the flashy applications associated with AI. In contrast, people may take OR for granted: e.g., assuming that their packages will be delivered on time.
- We discussed policy decisions that we feel an AI or OR expert should be included in making, including: supply chain risk mitigation (in regards to COVID-19), transportation including multimodal transport and placement of EV charging stations, public health risks/epidemiology, etc.
- Lastly, we mentioned that a major stumbling block is to predict how well we expect new policies will perform, as we do not always know the counterfactual. We mentioned a need to "do no harm," and to be sure to consult the relevant stakeholders.

Session 3: Methods

The second methods section was focused on issues of fairness, accountability, transparency, and ethics. These topics are critical to ensuring that technological developments actually benefit society, whatever their intended application. At the same time, several groups are working on direct application of AI and OR techniques to high-stakes societal decision-making domains like social service provision and allocation of resources in healthcare settings. We heard from leaders in this area on how their work drew from techniques across AI, machine learning, OR, and theoretical computer science to make advances in both the theory and practice of socially informed, fair, and ethical algorithmic decision-making.

Session 3 Video Recording

Slide Presentations
John Dickerson - AI & OR for Matching Markets in Healthcare
Subbarao Kambhampati - Explainable Plans & Decisions
Aaron Roth - Multivalid Learning: Meaningfully Quantifying Uncertainty

Breakout Sessions Summary

Breakout 1 – Resilience and Validation
Moderator: Katya Sheinberg
Notetaker: Zihao Yang

- Discussed what resilience and validation means for AI and OR. Discussed a lot of examples where it is necessary. In general almost all applications require resilience because the world is changing and data is changing so it is never applied in a static setting. Techniques developed in academia for ML tend to be for static settings. The need for resilience mainly comes from industry and that is not easy to quantify and not easy to collaborate with academia.
- Robust optimization in OR is in a way more mature than the current approaches to resilience in AI. AI can borrow from OR in this respect, but these techniques go somewhat unused because it is hard to measure the uncertainty that defines uncertainty sets and related trade-offs.
- We need to see if we can feed off each other in these communities. AI can borrow more advanced OR approaches and provide rich settings for techniques to improve and AI/ML will allow OR with the opportunity to automatically estimate uncertainties and provide tradeoffs in addition to providing OR with the setting to formulate the true goals. Do we need to find the optimal solution or is there a larger question of what we're trying to solve and model?

Breakout 2 – Opportunities for Better Understanding of Each Others' Language
Moderator: Maria Gini
Notetaker: Jiaoyang Li

- We discussed the similarities and differences between AI and OR in academia. There are a lot of overlaps between AI and OR. For example, machine learning, planning, NLP, and even distributed systems in AI have counterparts or overlaps in OR. AI people use many OR techniques. AI and OR people use different terminologies. For example, the word "model" has different meanings in AI and OR. OR people use "model-based" and "data-based", while AI people use "model-based" and "model-free".
- We also discussed AI and OR in industries because industries do not distinguish between AI and OR too much and naturally work on them together. Industries advertise more AI techniques than OR techniques, although their AI techniques sometimes contain OR components.
- We also discussed different methods for helping people understand each other's language. For example, we can create examples to show the differences between AI and OR and the differences between how AI and OR researchers frame the

problems. We need to encourage people to get interested in each other's area. We can create some workshops to bridge the gap. We can start with building a dictionary that maps the terminology. We eventually might want to merge the terminology.

Breakout 3 – Best Practices in Publication
Moderator: Sven Koenig
Notetaker: Enpeng Yuan

- AI/ML and OR have different publication models (quick turnaround on small ideas in conferences versus slow turnaround on big ideas in journals). Also, AI/ML conferences might value applications more than OR journals. To get OR researchers to publish more in AI/ML, evaluations of OR researchers (for example, for promotions and tenure) need to recognize publications in top AI/ML conferences.
- Some CS departments adopted an evaluation model where conference publications count as much as journal publications (which is becoming accepted practice in many areas of CS). Of course, there are good and not-so-good conferences, just as there are good and not-so-good journals. The letter writers are thus asked to provide information on the quality of the conferences and journals that the evaluee has published in. OR departments could use this model.
- One issue with conference publications in ML and ML-heavy CS areas (such as vision and natural language processing) is that the average review quality and thus the average paper quality in conferences has decreased in recent years as the number of paper submissions has skyrocketed and good reviewers have become scarce. This might explain why OR departments are hesitant to provide credit for publishing in conferences. ML conferences are working hard on fixing the issue. Also, accepted journal papers in some AI areas can be presented in conferences and top conference papers are (automatically) referred to journals.

Breakout 4 – Challenges and Opportunities in Fairness
Moderator: Maria de Arteaga
Notetaker: Amanda Kube

- We started by discussing how to operationalize a fairness problem. How do we adapt metrics for resource allocation or overcome reporting bias in the data? Additionally, should we think of fairness in the decision problem or fairness in the prediction problem? How do we elicit preferences and emphasize explainability, as different stakeholders have different definitions of fairness for their domain?
- Next we talked about going beyond constrained allocation and including governance and stakeholders playing a role. But are stakeholders (normal citizens) always better? Additionally, misspecified objective functions are the cause of many AI

failures and we need to balance out incentives as researchers and ease of measurement with understanding what is the true best objective for the problem.
- Lastly, we discussed connections with HCI and participatory design for eliciting preferences and deciding what the system is trying to achieve

Session 4: Applications

The final session of the workshop brought together speakers that are experts in collaborating with communities and organizations to address important societal problems. The session was entirely focused on issues related to fairness, interpretability, and social impact. The first speaker discussed how AI and multi-agent research can be used to address problems in public health (e.g., HIV prevention), conservation (e.g., to prevent poaching), and security (e.g., to protect borders) and presented concrete improvements through real world deployments. The second speaker focused on the importance of building interpretable models (as opposed to explaining uninterpretable/black box models). The third speaker discussed the use of AI and OR techniques to uncover inequities in pay.

Session 4 Video Recording

Slide Presentations
Milind Tambe - AI for Social Impact
Cynthia Rudin - Interpretable Machine Learning: Fundamental Principles and 10 Grand Challenges
Margrét Vilborg Bjarnadóttir - HR Tech & Pay Equity

Breakout Sessions Summary
Breakout 1 – Good Challenges for teams of AI and OR researchers
Moderator: Swati Gupta
Notetaker: Zihao Yang

- We discussed which good challenge problems are ideal for teams of AI/OR.
- Sparse data and sensitivity lead to brittle inferences -- AI, OR for social good problems.
- Three prominent problem areas emerged from our discussion, for important applications such as healthcare, and applications where development with only AI or only OR wouldn't be possible (energy, chip design, networks). (a) **Utility-based AI/ML**: What is the value of accuracy in different regimes or parts of data? How does it impact pipelines? Can it be applied to conservation problems? What about counterfactual reasoning? Can risk level in decisions be quantified? (b) **Iterative Predict and Optimize AI/OPT.** Iteratively exploring the space of

discrete decisions for iterative predict-and-optimize, and to use AI to quantify intermediate states. (c) **Data aggregation techniques** using AI for scaling OPT so that performance of OPT models can be scaled up and bottlenecks can be removed.

Breakout 2 – Data-sharing Practices
Moderator: Thiago Serra
Notetaker: Jiaoyang Li

- We first discussed the recent efforts in the AI community of encouraging people to share data in the paper reviewing process (like AAAI).
- We also discussed data issues regarding the student internships. Students go to internships and work on a lot of data but they cannot use the data after the internship or publish papers based on the data.

Breakout 3 – Ethical Considerations for Teams of AI/OR systems and Consequences for AI/OR Education
Moderator: Phebe Vayanos
Notetaker: Enpeng Yuan

- We first discussed the question of whether AI/OR systems raise specific challenges (compared to, say, pure statistics / data science). We concluded that they are a lot more powerful (in their scale, magnitude, and capabilities) and are increasingly being deployed in settings that directly impact people and society and thus they do raise more difficult/different challenges. In particular, they may have unintended consequences due to e.g., biases (think image recognition), unmodelled phenomena, uncertainty, lack of transparency, etc.
- Thus, we all agreed that there is an urgent need to train our students (both undergrads and grads/PhD) on ethics around AI so that they can both understand, anticipate, and communicate the potential impacts and unintended consequences of the AI systems they develop or use. Yet, the vast majority of AI/OR/IE degrees focus more on algorithms and math and do not discuss ethics.
- The group thus made several recommendations to address this issue:
  a) We need to integrate ethics in core classes and even the “Intro to AI/OR” classes should have a session on ethics of AI/OR -- these lectures could be taught by professors in schools of philosophy, if possible. Indeed, it was noted that AI/OR researchers may not be suited to teach these courses. The group agreed that we should actively engage with ethicists in philosophy that complement the technical material: there is an urgent need for teaching courses jointly with other disciplines.
  b) In addition, the group recommended that basic fairness metrics should be taught even in undergrad classes on say ML (just like we teach students to look at AUC or Accuracy).

c) Finally, the group recommended that we should teach case studies (similar to what is done in business schools) that illustrate potential failures of AI/OR systems to spur discussion.

Breakout Room 4 – Suggestions for Next Workshop
Moderator: John Dickerson

- All members of this breakout session appreciated the slate of speakers in the first workshop. With this as an initial basis, though, everybody agreed that more diversity in the speakers' backgrounds *outside* of AI/ML and OR/MS would be useful. Many areas were discussed, including HCI & participatory design, cognitive psychology & behavioral modeling, macroeconomists, and econometricians. To get value from "outsiders" speaking, having a concrete ask (e.g., for cognitive psychologists, asking them to describe state-of-the-art human behavioral models and then asking how AI/ML/OR/MS folks can make that better).
- Including (more) folks from industry and the public sector would be good. Hosting a workshop on an individual topic, such as robustness of supply chain, and then targeting industry and public sector practitioners in that single topic area was one promising idea.
- Understanding how to get data, especially large public sector datasets, was discussed---specifically, hosting an in-person/hybrid workshop in the DC area to attract members of the intelligence community, statistical agencies such as Census and SAMHSA, and then better understand how data flows, who can share what, and how we can build out repeatable methods for accessing and analyzing that data.

Closing Session – Next Steps

Prof. Sven Koenig and Prof. Ramayya Krishnan
Closing Session Video Recording
Slide Presentation - Closing Remarks (an AI Perspective)

In his closing remarks, Koenig listed several common misconceptions about the relationship of AI and OR and pointed out that they disappeared during the workshop in favor of the opinion, clearly stated by one of the invited speakers, that both disciplines have very similar objectives, except that they make different assumptions and have thus developed different tools due to the different use cases they are interested in. In fact, large parts of AI are concerned with decision making and optimization, as was also mentioned in the opening remarks. For example, as noted by Krishnan in the opening session, the most popular AI textbook (AI: A Modern Approach by Russell and Norvig) views AI as the study of rational agents, which select actions that are expected to

maximize their given performance measures -- which is similar to the OR way of solving problems with optimization as the goal.

AI has developed its own tools for decision making and optimization, but also adopted tools from related disciplines, for instance:

- utility theory and multi-attribute utility theory from decision theory,
- game theory and auctions from economics, and
- stochastic dynamic programming techniques (such as value iteration and policy iteration for Markov decision processes) from OR.

Given this adaptation and adoption of ideas from one discipline into another, the question is: “How best do we keep both areas abreast of such progress?” The discussions during this workshop explored some possible ideas to do so.

In terms of conferences at the interface of AI and OR, CPAIOR (the International Conference on the Integration of Constraint Programming, AI, and OR) is one such example. It would be worthwhile to consider conferences that include AI and other areas within OR. A good model for how to merge different research communities into a homogeneous community at their intersection can be found in the context of computational economics, with conferences such as EC (the ACM Conference on Economics and Computation) and WINE (the Conference on Web and Internet Economics) at the intersection of AI and economics.

Overall, it seems to make sense to create a science of making good decisions (by optimizing given objectives), to combine ideas from different decision-making and optimization disciplines and provide a larger set of tools to practitioners than any one discipline can provide in isolation. Such a science would not only be at the interface of AI and OR but also include economics, control theory, decision and utility theory, and a variety of other disciplines. Some recent conferences try to create such communities, such as ADT (the International Conference on Algorithmic Decision Theory) and EAAMO (the ACM Conference on Equity and Access to Algorithms, Mechanisms, and Optimization).

This workshop was very exciting because it showed that several exciting and topical research directions are of interest to both AI and OR researchers, including machine learning; human-machine teams; and fairness, accountability, and transparency. We will continue to identify the next steps to bring AI and OR researchers together. These should ideally be “baby steps,” that can be realized relatively quickly and with relatively few resources so that we can point to first successes quickly.

Some ideas to consider are:

- We could make data sets for a variety of application domains available to both research communities and/or hold competitions of interest to both research communities. The AI research community can often easily be influenced by competitions and is already looking at competition domains that are typically studied in OR. For example, the top conference on AI planning and scheduling (ICAPS) this year promoted competitions such as train scheduling, learning to run a power network, automatic reinforcement learning for dynamic job shop scheduling, and planning for the dynamic pickup and delivery problem. Many of these applications are typically of interest to OR practitioners.
- We could also start an AI/OR summer school for AI and OR Ph.D. students, where for each topic, such as local search or probabilistic planning, both an AI expert and an OR expert lecture.
- In the longer run, we should start to think about AI/OR MS and Ph.D. programs (like the ACO programs in the past) and perhaps about an ACM/INFORMS AI/OR conference or journal.

Of course, some funding is needed for such activities, but the National Science Foundation plus the societies organizing this workshop would probably be interested in finding a way to work towards these interdisciplinary activities. These specific suggestions offer concrete examples about how to realize the collaboration potential at the data, methods and policy layers that were outlined during the opening session.

* *The material in this report is based upon work supported by the National Science Foundation under Grant No. 1734706. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.*
* *This report was published on October 22, 2021.*

Appendices

A. Workshop Participants

** Workshop organizers*

Name	Last	Affiliation
Jacob	Abernathy	Georgia Institute of Technology
Alper	Atamturk	University of California, Berkeley
Dimitris	Bertsimas	Massachusetts Institute of Technology
Sudip	Bhattacharjee	University of Connecticut
Margret	Bjarnadottir	University of Maryland
Nadya	Bliss	Arizona State University and CCC
Elizabeth	Bradley	University of Colorado Boulder and CCC
* Sanmay	Das	George Mason University
Maria	de Arteaga	University of Texas at Austin
*John	Dickerson	University of Maryland
Thomas	Dietterich	Oregon State University
Bistra	Dilkina	University of Southern California
Adam	Elmachtoub	Columbia University
Xiao	Fang	University of Delaware
Elena	Gerstmann	INFORMS
Maria	Gini	University of Minnesota and CCC

Amy	Greenwald	Brown University
Swati	Gupta	Georgia Institute of Technology
Robert	Hampshire	US Department of Transportation
Peter	Harsha	CRA
Ilya	Hicks	Rice University
Kartik	Hosanagar	University of Pennsylvania
Stefanie	Jegelka	Massachusetts Institute of Technology
Andrew	Kahng	University of California, San Diego
Subbarao	Kambhampati	Arizona State University
Amin	Karbasi	Yale University
*Sven	Koenig	University of Southern California
*Ramayya	Krishnan	Carnegie Mellon University
Amanda	Kube	Washington University - PhD Student
*Radhika	Kulkarni	SAS Institute, Inc (retired)
George	Lan	Georgia Institute of Technology
Kate	Larson	University of Waterloo
Jiaoyang	Li	University of Southern California - PhD Student
Andrea	Lodi	Cornell Tech
Daniel	Lopresti	Lehigh University and CCC
Ross	Maciejewski	Arizona State University

Arya	Mazumdar	University of California, San Diego
Ranga	Nuggehalli	UPS
Marek	Petrik	University of New Hampshire
Ahmad	Ridley	NSA
Justin	Romberg	Georgia Institute of Technology
Aaron	Roth	University of Pennsylvania
Edward	Rothberg	Gurobi
Cynthia	Rudin	Duke University
Ann	Schwartz	CCC
Thiago	Serra	Bucknell University
Jay	Shan	Miami University
Katya	Sheinberg	Cornell University
David	Simchi-Levi	Massachusetts Institute of Technology
Satinder	Singh	University of Michigan Ann Arbor
Alice	Smith	Auburn University
Peter	Stone	University of Texas at Austin
Nick	Street	University of Iowa
Katia	Sycara	Carnegie Mellon University
Milind	Tambe	Harvard University and Google

Madeline	Udell	Cornell University
Berk	Ustun	University of California, San Diego
*Pascal	Van Hentenryck	Georgia Institute of Technology
*Phebe	Vayanos	University of Southern California
Stephen	Wright	University of Wisconsin-Madison
Zihao	Yang	Carnegie Mellon University - PhD Student
William	Yeoh	Washington University
Enpeng	Yuan	Georgia Institute of Technology - PhD Student
Jerry	Zhu	University of Wisconsin-Madison

B. Workshop Agenda: September 23-24, 2021
Day 1 (Thursday, September 23, 2021)

11:00 AM	Welcome / Introductions / Opening
11:15 AM	Opening address *Ramayya Krishnan is the W. W. Cooper and Ruth F. Cooper Professor of Management Science and Information Systems at Heinz College and the Department of Engineering and Public Policy at Carnegie Mellon University.*

11:30 AM	Sessions 1 - Methods Research at the intersection of OR & AI; review of key technological innovations ● *Stephen Wright, University of Wisconsin (continuous optimization)* ● *Andrea Lodi, Cornell Tech (discrete optimization)* ● *Katia Sycara, Carnegie Mellon University (multi-agent systems)* ● *Satinder Singh, University of Michigan (reinforcement learning and decision-making under uncertainty)*
12:30 PM	BREAK (10 minutes)
12:40 PM	Moderated Q&A
01:00 PM	Breakouts
01:40 PM	Report back from breakouts
02:00 PM	LUNCH BREAK (1 hour)
03:00 PM	Session 2 - Applications Coose application domains at the intersection of OR and AI; review key emerging application areas enabled by the technological innovations;mix of areas that are OR strengths and AI strengths ● *David Simchi-Levi, MIT (supply chain)* ● *Ranga Nuggehalli, UPS (Edelman winner - UPS)* ● *Robert Hampshire, (US Department of Transportation)* ● *Phebe Vayanos, University of Southern California (Data-driven decision-making for social impact)*

04:00 PM	BREAK (10 minutes)
04:10 PM	Moderated Q&A
04:30 PM	Breakouts
5:10 PM	Report back from breakouts
5:30 PM	Wrap up Day 1

Day 2 (Friday, September 24, 2021)

11:00 AM	Day 2 Welcome/Recap
11:05 AM	Session 3 - Methods FATE Session ● *John Dickerson, University of Maryland (fairness in machine learning and mechanism design)* ● *Subbarao (Rao) Kambhampati, Arizona State University (explainable plans / decisions)* ● *Aaron Roth, University of Pennsylvania (private data analysis, fairness in machine learning)*
11:50 AM	Break (10 minutes)
12:00 PM	Moderated Q&A
12:20 PM	Breakouts

01:00 PM	Report back
01:20 PM	Lunch (60 minutes)
02:20 PM	Session 4 - Applications Fairness ● *Milind Tambe, Harvard University (AI for social good)* ● *Cynthia Rudin, Duke University (interpretable machine learning)* ● *Margret Bjarnadottir, University of Maryland*
03:05 PM	Break (10 minutes)
03:15 PM	Moderated Q&A
03:25 PM	Breakouts
04:05 PM	Report Back
04:25 PM	Closing Session
05:30 PM	End of Day 2