

# Artificial Intelligence and Operations Research A Strategic Vision

Sanmay Das, John Dickerson, Sven Koenig,  
Ramayya Krishnan, Pascal Van Hentenryck, and Phebe Vayanos.

July 2020

## 1 The Vision

Artificial intelligence (AI) and Machine Learning (ML) have received significant attention in recent years, primarily due to breakthroughs in game playing, computer vision, and natural language processing that have captured the imagination of the scientific community and the public at large. IBM Watson, Deep Blue, AlphaGo have shown how AI can reach a level of intelligence that surpasses human performance on specific tasks. Millions of people seek assistance in everyday queries and other tasks using Apple Siri and Amazon Echo. Computer vision software is now being used in medical decision making, assisting doctors to provide improved health care to larger population segments. Although today AI has a significant focus on Machine Learning (ML), computer vision, and natural language processing, it also features research on constraint satisfaction and optimization, robotics, decision-making under uncertainty, human-machine collaboration, and multi-agent systems, to name a few.

The convergence of big data, high-performance computing, and AI has the potential to transform all sectors of the economy, including energy, health-care, manufacturing, security, supply chains, and transportation, and to address fundamental societal challenges. AI is likely to broaden its application domains, going beyond its main focus on computer vision and natural language processing and interacting with many other disciplines in engineering, the physical and social sciences, and the humanities. This will open new opportunities and challenges for AI and the corresponding disciplines. *This proposal focuses on one of these opportunities: how to exploit the synergies of the AI and Operations Research (OR) communities.* In fact, the COVID-19 pandemic has already highlighted how important AI and OR are to address fundamental societal issues, how complementary they are, and how they may address the fundamental inequalities that have been clearly exposed by this crisis. It is thus the role of this proposal to establish the joint strategic vision for AI/OR that will maximize the societal impact of AI and OR in a world that is bound to undergo significant changes.

This proposal is a collaboration between the ACM SIGAI and the INFORMS AI Strategy Advisory Committee established by past INFORMS President Ramayya Krishnan to provide a vision for INFORMS in a world where the role of AI will be prominently featured.<sup>1</sup> Its goal is to pursue a set of exciting opportunities at the intersection of AI and Operations Research (OR) that could help in delivering the economic and societal benefits just sketched. It starts with the recognition that the two communities have complementary strengths whose synergies have started to emerge. Indeed, in its quest to build intelligent agents, the AI community has always been quick to borrow methods and technologies from other fields, including from economics, OR, and statistics. In fact, in several areas, the boundaries between AI and OR technologies are blurring.

---

<sup>1</sup>One of the outcomes of the AI Strategy Advisory Committee has been an internal white paper [1].

Machine learning has adopted, refined, and expanded optimization algorithms that originated from the OR community, including stochastic gradient descent and convex optimization. Since AI aims at building rational agents maximizing a performance measure, optimization has been an integral part of AI. 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. In decision under uncertainty, there is considerable activity at the union of reinforcement learning, Markov decision processes and their generalizations, and stochastic optimization.

This proposal also recognizes that AI problems present new fundamentally new challenges for OR. Machine-learning tasks operate on massive datasets, challenging even the more scalable optimization algorithms. Robots, self-driving cars, and voice-control personal assistants operate autonomously in real time and often under uncertainty. Increasingly, AI applications are multimodal, combining visual, vocal, and textual data. These challenges present significant opportunities for O.R. researchers: They will require novel optimization algorithms, a fresh look at old algorithms, and new methodologies to address real-time requirements and uncertainty, massive datasets, streaming data, and the need for robustness and reliability. Tight integrations of machine learning and optimization also present true opportunities for the two communities, leveraging OR expertise in optimization, stochastic modeling, and problem solving. Finally, as AI moves to novel disciplines where OR has significant domain expertise, the opportunities and synergies will multiply, creating a virtuous cycle of scientific innovation and societal impact.

The COVID-19 pandemic has also highlighted the existence of critical societal problems where vast amount of data is not available initially. It reinforces our conviction that there is a strong need to blend data-driven and model-driven approaches and that an AI/OR joint initiative and strategic vision may have benefits that have not yet been fully contemplated by the two fields.

There are however barriers and difficulties in realizing this vision due to cultural differences between AI and OR. Some of the challenges are obvious. First, the two fields use different vocabulary used for the same concepts, e.g., reinforcement learning in AI and (approximate) dynamic programming in OR. Second, the missions of the two fields have also been quite distinct: whereas AI focuses on building intelligent agents, OR has traditionally focused on process improvement to enhance quality of service or efficiency. Third, the two fields have also focused on different application areas, which makes communication and knowledge transfer more difficult. Finally, there is a strong investment imbalance between the two fields, not only in funding but also in human capital. Students are massively attracted to artificial intelligence around the world, and OR, despite providing some of the underlying technologies, only captures a small share. However, the COVID19 pandemic provides a strong incentive to bridge these differences and build a strategic vision to address fundamental societal challenges by leveraging the strengths of both communities.

This proposal aims at overcoming these difficulties and providing a stepping stone for a strong and sustained collaboration between the two fields. A unified vision for AI and OR can be articulated around an AI/OR stack that generalizes the traditional AI stack along several dimensions. Figure 1 is a preliminary attempt in that direction. In perception, it addresses questions such as what data to collect, using techniques from the Design of Experiments, adaptive design/sampling, and data mining. In the learning stage, the focus may be on reliable predictions, merging techniques from machine learning, statistics, stochastic optimization, and quality control. The planning and execution level include strategic, tactical, and operational decisions, combining machine learning, optimization, Markov decision processes, and human-computer interactions (human-aware or human-centered decision-making). Finally, the evaluation layer should leverage methodologies from management science, simulation, quality control and reliability to assess the performance of intelligent systems along multiple dimensions.

It is important to emphasize again that systems implementing this AI/OR stack will combine both data- and model-driven methodologies, pushing the field in uncharted and exciting territories. The OR community has a long history in developing mathematical models for specific classes of applications (e.g., multi-commodity flows in transportation, supply chains, and disaster response) and modeling language to

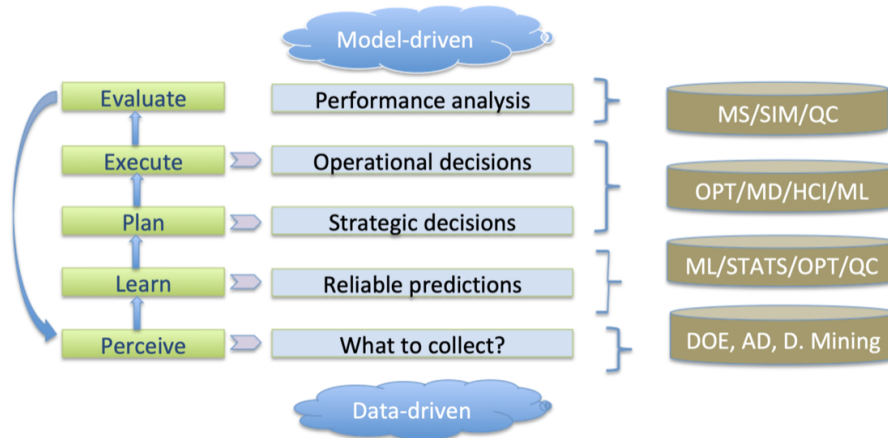


Figure 1: The AI/OR Stack for the Strategic Vision.

express them, separating the modeling of reality and potential solution techniques. The availability of massive datasets, together with data mining and unsupervised machine learning, present a key opportunity to discover the structure underlying practical applications and derive models exploiting the salient properties of the data. Tighter integrations of machine learning and optimization and, in particular, the ability to optimize over machine-learning models, may improve the reliability of decision making. Finally, and equally important, the capability to work with limited data, simulate potential futures, and optimize decision making under significant uncertainty and risk is also critical, highlighting again the synergies between the data-driven and model-driven approaches.

Last, but not least, AI and OR technologies need to go beyond their traditional objectives and promote equity and fairness in their predictive and prescriptive models and algorithms. Given that AI and OR are now being used in consequential decision-making tasks that may affect entire population segments, fairness, equity, and absence of bias should not be an afterthought: They need to become first-class citizens in the entire AI/OR stack, from data collection to machine-learning algorithm, decision-making tasks, and validation. Both communities are actively exploring these issues under various labels, including algorithmic fairness and algorithmic justice. But the role of equity and fairness needs to become more central and fully integrated in every layer of the AI/OR stack. The same holds for resilience, another critical objective, given the disproportionate impact of major disruptions on low-income and underprivileged communities.

## 2 Structure

**Workshops** The core of this proposal is a series of three workshops that can be summarized as follows.

1. *Articulating a strategic vision*: This workshop has two objectives: (1) to review the state of AI/OR along several axes; and (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; (4) to select a few topics for summer schools and summer research programs that would foster an AI/OR collaborations. The grand challenges may be in areas such as resilient supply chains, sustainable energy, health care, equitable transportation, and the modeling of

human behavior.

2. *Implementing a strategic vision*: This workshop builds on the strategic vision and turns the high-level opportunities into a small set of grand research challenges that AI and OR can jointly tackle more effectively than each field would independently. The grand challenges will be complemented by tool and data repositories that would facilitate cooperation, competition, industrial collaboration, and the grand challenges.
3. *Progress on the strategic vision*: This final workshop describes preliminary results on the grand challenges identified by the second meeting.

**Working Groups** To ensure the success of the workshops, this visioning proposal is also organized in terms of working groups. Each working group will be responsible for organizing a strategic vision around an important theme where an AI/OR collaboration seems particularly promising. Each potential participant will submit a vision statement in preparation for the first workshop. These statements will then be discussed within the working group to articulate a joint vision and a collection of potential presentations and panels for the workshops. After the first workshop, the working groups will restart their work to articulate a first set of grand challenges on their respective themes. Grand challenges may be fundamental societal problems for which an OR/AI approach has significant promise and/or smaller scientific advances that would fundamentally transform both fields. Once again, each participant will be asked to prepare a statement articulating these research challenges and potential avenues to address them. These will again be synthesized and discussed at the workshop. The third and final workshop will review preliminary progress on some of these issues and may be organized around a series of lightning talks.

A preliminary and tentative list of potential themes could include

1. The AI/OR stack including End-to-End Learning and Optimization;
2. Robustness and Reliability of Machine Learning;
3. Fair and Equitable Predictions and Decision Making;
4. Blending of Data-driven and Model-driven approaches;
5. Resilient, Risk-Aware, Planning and Operations;
6. Real-Time Decision Making (under Uncertainty);
7. Human-in-the-loop decision making.

**Timeline** Table 1 is a tentative timeline for the visioning proposal. Due to COVID-19, it is anticipated that the workshops will be online. If the virus is under control by the time of the last workshop, the last workshop will take place in Maryland to be close to the funding, scientific, and technological agencies and significantly raise awareness of the potential of an AI/OR vision.

### 3 Outcomes

The visioning proposal will deliver the following outcomes:

1. Joint White Paper on the Strategic Vision;
2. Survey Paper on Research Challenges;

September 2020	Call for Position Papers
October–November 2020	Working Group Vision
January 2021	Workshop 1: Articulating a strategic vision
February–April 2021	Working Group Research Challenges
June 2021	Workshop 2: Implementing a strategic vision
Summer 2021	Working Group Research Brainstorming
January 2022	Workshop 3: Progress on the strategic vision

Table 1: Timeline of the AI/OR Visioning Proposal.

3. Organization of Summer Schools on AI/OR;
4. Organization of Summer Research Programs on AI/OR.

## 4 Organization

### 4.1 Leadership Team

Sanmay Das (Washington University), John Dickerson (Maryland), Sven Koenig (USC), Ramayya Krishnan (CMU), Pascal Van Hentenryck (Georgia Tech), and Phebe Vayanos (USC).

### 4.2 Potential List of Invitees

Table 2 contains a list of potential invitees. This list is not intended to be exhaustive or final. It is meant to convey the wealth of backgrounds and expertise at the intersection of the two fields and the significant potential for cross-fertilization.

## 5 Budget

The meetings are planned to be online at this stage, so the costs will be largely minimal (about 5K). They will probably amount to the support of a staff person for setting up the website and the various Zoom/Bluejean sessions.

## References

- [1] INFORMS Artificial Internal White Paper. <https://pubsonline.informs.org/doi/10.1287/orms.2019.06.14/full/>, 2019.