ARTIFICIAL INTELLIGENCE INTERNAL WHITE PAPER

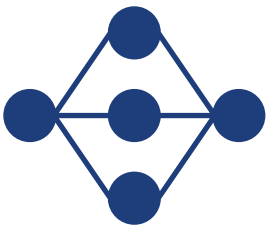

DECEMBER 2, 2019

Artificial Intelligence Internal White Paper

Purpose

Artificial intelligence (AI) has received significant attention in recent years, primarily due to breakthroughs that have captured the imagination of the scientific community and the public at large. AI is at the center of a fundamental transformation of the economy, driven by the availability of data, high-performance computing, and seminal scientific and technological contributions. This internal white paper addresses the INFORMS members, as well as other operations research (O.R.) professionals. Its goal is to provide a vision for INFORMS in a world where the role of AI will be prominently featured. This includes: increasing awareness of the synergies and opportunities for O.R. with and within AI, better positioning INFORMS and the profession to contribute to and benefit from AI, and build on and enhance INFORMS capabilities in AI and O.R.

AI originated in computer science and is now a significant part of education and research missions of colleges of computing and computer science departments. Today AI has a significant focus on machine learning (ML), computer vision, and natural language processing, but it also features research on constraint satisfaction and optimization, robotics, decision-making under uncertainty, human-machine collaboration, computer games, and multi-agent systems, to name a few. AI adopts techniques from a wide variety of fields, including economics, O.R., and statistics in its quest to build intelligent agents. When looking at the future of INFORMS, these recent developments raise existential questions: Where does operations research fit in this landscape? Does INFORMS have a role to play? And, if so, how does INFORMS claim its place in this area? Are there cultural differences that need to be surmounted to achieve this? Where do we begin?

In early 2019, INFORMS President Ramayya Krishnan appointed an expert AI Strategy Advisory Committee to consider these questions. The committee members have reached out to several prominent researchers and industry professionals in INFORMS and related societies and worked with the INFORMS Board to get input on some of the above questions. Individuals were selected to cover a wide range of expertise in optimization, AI/ML, and statistics. The following theme was often repeated: the field of O.R. and analytics contributes techniques essential to the success of AI and its applications. In addition to advancing artificial intelligence, the field of O.R. and analytics has been a leader in augmenting intelligence of human actors involved in decision-making in every sector of the economy. Yet, O.R. and analytics lack the mind share that is commensurate with its contributions to date as well as its potential future contributions.

This white paper explores the aforementioned issues from various angles, with the goal of educating the O.R. community. It starts by defining AI for this purpose. It then examines synergies between O.R. and AI and explores ways in which the two areas can benefit from one another. The goal is to determine ways of increasing awareness of the following.

- The role of O.R. within the AI community
  - What are the ways in which O.R. is already contributing to research in AI?
  - How can INFORMS and the profession be in a better position to contribute to and benefit from AI?
  - How does INFORMS get the word out about its interest in collaboration?
- The role of AI within the INFORMS community
  - What opportunities do AI/ML provide for the O.R. community?
  - Does AI pose a threat to O.R.? Can INFORMS find its rightful place in this space? How does INFORMS ignite urgency around this concern?
  - Does INFORMS need to increase its efforts in education, publications, workshops, and other activities that illustrate the synergies between O.R. and AI?
- The role of O.R. for AI to the public and government
  - How can we increase awareness of how important O.R. is to AI with funding agencies, employers, and the public at large?

In addition to answering the above questions, this internal white paper highlights the importance of interdisciplinary collaboration across the various research areas needed to succeed with AI. Preliminary outreach to professional organizations (e.g., AAAI, SIAM, EURO, NISS, etc.) has shown that collaboration across the societies is uniformly seen as a critical activity that will be supported. Now is the time for INFORMS to capitalize on this interest.

This white paper takes the view that AI is much more than the latest buzzword: It recognizes that AI is making fundamental advances in science, engineering, and technology that will eventually transform every sector of the economy. AI, together with high-performance computing, provides the engine behind the data revolution. INFORMS' efforts to enter the analytics arena shared a similar motivation but it is fair to say that the AI movement has been gaining greater external recognition and funding. Indeed, AI has captured the public's imagination, received both congressional and presidential backing, and is raising considerable interest from industry for the potential in improving their operations. There is also an international threat as China has pledged to become the world leader in AI, a situation that both Europe and the United States are now addressing.

Background

The convergence of big data, high-performance computing, and AI has the potential to transform all sectors of the economy. IBM Watson, Deep Blue, and AlphaGo [1] have shown how AI can reach a level of intelligence that surpasses human performance on specific tasks. Watson defeated human Jeopardy! champions by using a question-answering system combined with cognitive computing over a large knowledge base. AlphaGo used a combination of pattern recognition (driven by deep learning), advanced search algorithms, and reinforcement learning to defeat the best human players at what is widely regarded as the most involved strategy board game

ever devised. Millions of people seek assistance in everyday queries and other tasks using Apple Siri and Amazon Echo. While Jeopardy! and AlphaGo demonstrated AI advances on games, the same underlying technology is being applied to or contemplated for critical applications in many sectors with potentially significant impact on society. More generally, it can be argued that every industry is seeing an explosion of applications that demand human-like intelligence with super-human speed, accuracy, and performance. For instance, AI is now being applied to medical decision-making, including in breast cancer detection or distinguishing cancerous lesions in dermatology, ride-hailing companies to match riders and drivers, driverless cars paired with autonomous drones for faster food delivery, and transforming the factory floor. This proliferation of impactful applications in almost every field has given rise to several initiatives in government, across universities, and in research organizations everywhere. They include an executive order at the highest level of the government with President Trump's initiative on Artificial Intelligence (see these documents for more detail), a billion-dollar investment at MIT in the area of artificial intelligence, and UVA's largest investment to fund a School of Data Science.

It is important to emphasize that the vast majority of machine-learning methods use O.R. and statistics at their core: As previously mentioned, the field of AI is quick in adopting and refining technologies to pursue its goal of building intelligent agents. At the same time, many AI problems present new challenges for optimization and related O.R. methods and hence could expand the potential impact of O.R. on society. Given this symbiotic relationship between AI and O.R., it is valuable to explore opportunities for collaboration between the two fields and determine ways to promote strong interaction between the two communities.

What is AI?

The Association for the Advancement of Artificial Intelligence (AAAI) (i.e., the counterpart of INFORMS for AI) defines its mission as "advancing the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines." To fully grasp the AI community and its evolution over the years, it is important to recognize its focus on understanding the fundamentals of human and machine intelligence, which goes beyond its potential business and engineering applications.

In recent years however, the term "Artificial Intelligence" has become a household word outside the AI community. In the business press, it is widely used to describe the use of computers to make real-time decisions on the basis of vast amounts of data. Bernard Mars provides six definitions of artificial intelligence and describes how some of the industry's leaders are focusing their AI research efforts. In its report "State of Artificial Intelligence for Enterprises," Teradata uses the following definition: "AI is the ability to automate enterprise decisioning using human-to-machine cognitive interactions where machines are able to augment and assist human capabilities by sensing and continuously learning, reasoning and inferring, deciding, and acting to drive a business outcome."

A few ingredients are common across all definitions: Clear definition of the problem to be solved, use of vast amounts of data, high-performance implementations of mathematical algorithms to arrive at decisions that can be deployed in the required environments. In essence, AI is presented as a highly scalable, actionable, automated data-driven decision-making technology using a wide

and deep arsenal of mathematical algorithms from many disciplines. These definitions in fact are well-aligned with the one adopted in the most popular AI textbook by Russell and Norvig [2] who take the view that the goal of AI is "to create rational agents who can perceive and act such that some objective function is optimized." They also capture the so-called "AI stack" that integrates perception, learning, decision-making, and action.

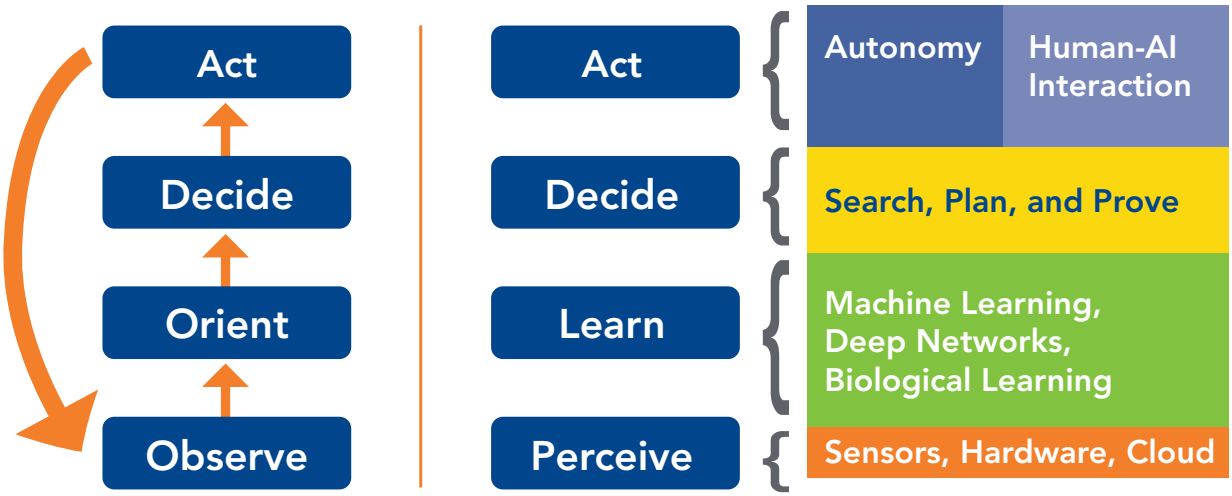

Note that the focus of AI on autonomous systems also contrasts with the focus of O.R. on decision support, including for strategic and tactical decision-making.

These definitions also try to articulate the potential of current AI technologies in applications that benefit society, but they do not capture the long-term, scientific goal of the AI community. This distinction is important because it highlights some of the fundamental differences with the INFORMS community, which presents itself as the study of "Operations Research & Analytics" that "are proven scientific mathematical processes that enable organizations to turn complex challenges into substantial opportunities by transforming data into information, and information into insights that save lives, save money and solve problems." Viewed in this light, it becomes easier to understand the existential question of the INFORMS community and how it fits with recent AI developments: AI has a broad vision of understanding and building intelligent agents but these agents can then be used to address the challenges that are the focus of the INFORMS community.

To emphasize the similarities and differences between AI and INFORMS communities, it is useful to consider the following graph (reproduced from the opening session of AAAI 2019, a premier conference in artificial intelligence) that reports the number of papers in the various areas commonly associated with AI. It clearly highlights the current focus of AI on machine learning (ML), computer vision, and natural language processing (NLP), areas that have substantially benefited from advances in deep learning.[1] Note, however, that other areas of AI have also seen

[1] *It is important to emphasize that machine learning submissions may include a significant portion of optimization papers applied to machine learning. The above classification does not capture this.*

significant increases in submitted and accepted papers in recent years. Primary areas such as constraint satisfaction and optimization, reasoning under uncertainty, planning, routing, and scheduling strongly overlap with research in the INFORMS community. Others, such as knowledge representation, robotics, cognitive modeling, and interactive entertainment are more distant. [Note: an Appendix 1 will be added to provide resources for additional information on AI.]

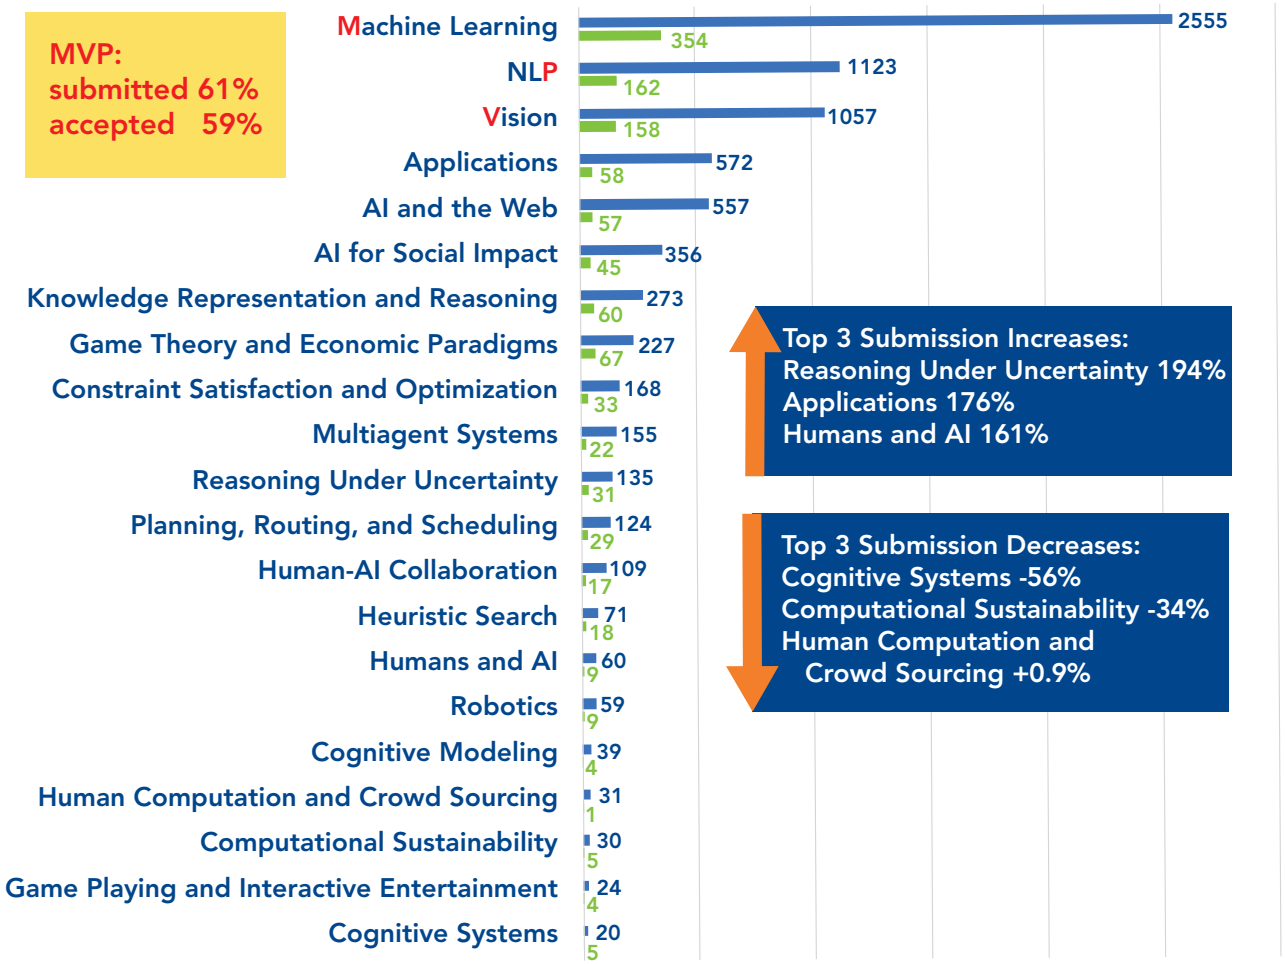

The Impact of AI on O.R. as a Discipline

The previous section emphasized that AI and INFORMS communities have different foci and overall objectives. Yet it also highlighted that AI provides methodologies and technologies that are directly applicable to the mission of INFORMS. Moreover, in many cases, the boundaries between AI and O.R. technologies are blurring. For instance, machine learning algorithms leverage optimization algorithms that originated from the O.R. community, constraint programming and mathematical programming have been merging for about two decades, and both communities investigate how to make decisions under uncertainty. As the graph presented above indicates, AI has a significant focus on computer vision and natural language processing, areas that have not been widely studied by the INFORMS community. But the AI and business communities have been quick to realize that AI can now be applied to a wide range of industrial and societal problems.

A significant challenge for the INFORMS community comes from AI's important place in the public perception. Since AI technologies now display human or superhuman performance on tasks that were strongly associated with human intelligence, AI has attracted significant attention in scientific communities, business and industry, and government. Although the expectations in some circles are probably too high at this point (and the AI community is careful in calibrating them), it is clear that AI will have a sustained and significant impact on society. It is thus important for the INFORMS community to understand its role in an AI world, especially since many prominent O.R. professionals and researchers agree that O.R. does not have a similar brand recognition. Why is that and how can we change it?

In recent years, machine learning and, in particular, deep learning have received considerable attention. Machine learning techniques such as clustering, random forests, and neural networks were already highly popular in the 1980s and have been successfully used in forecasting, predictive modeling, classification, and identification problems for decades. What has changed is the ability of machine learning, sometimes in conjunction with other techniques (including search and optimization), to achieve or surpass human performance on tasks that are strongly associated with human intelligence. The performance of AI on computer games, computer vision, and natural language understanding appears "magical" for the public at large. Moreover, the performance of these AI systems creates opportunities in many applications that capture the imagination and are simple to understand: virtual assistants conversing in English, video surveillance, automated sensing of faults in manufacturing, rapid evaluation of images to detect cancerous lesions, and self-driving cars are some well-known examples. Interestingly, not a single factor is responsible for the achievements of machine learning in the last two decades. Rather, it is the result of the convergence of better algorithms/models, large-scale high-performance computing (including GPUs), and the availability of massive data sources enabled by communication technologies.

It is thus not surprising that computer science is often the first discipline associated with artificial intelligence. However, it is also important to recognize that AI successes build on a variety of old and new O.R. techniques, some of which are reviewed in the next section. The AI community is well aware of the value of O.R. and is largely technology agnostic: It has adopted and expanded O.R. and other technologies freely in order to pursue its goal of building intelligent agents. Many in AI now have a deep understanding of O.R. technologies and some actively seek collaborations.

But these contributions typically do not enter the public awareness and have not even been fully appreciated within the INFORMS community until recently. More importantly perhaps, these developments are rarely highlighted as successes of applying O.R. techniques to AI problems by the INFORMS community, as they should be. Indeed, emphasizing the impact of O.R. within AI has the following significant benefits:

- It helps O.R. to be recognized as an important contributor for AI, creating potential investments in the field.
- It elevates the brand recognition for O.R.
- It provides opportunities for interaction with other disciplines.
- It opens a wide range of new problems for research in O.R., including traditional topics in AI and entirely novel business applications that feature real-time decision-making under uncertainty with access to wide varieties and large volumes of data.

INFORMS as an organization should highlight these opportunities and ensure that O.R. fully realizes its potential to expand and enhance AI technologies and methodologies. The INFORMS community is ideally positioned to play a leading role in some areas of great significance. Indeed, the community has traditionally focused on areas such as manufacturing, supply chains, healthcare, energy, and transportation; it has acquired significant domain expertise and often has deep engagements and collaborations in these sectors. The AI community has been focusing more on areas such as computer vision, natural language processing, robotics, social networks, and computational game theory, although AI is increasingly looking for novel applications. Moreover, optimization, statistics, and stochasticity are areas of significant strength at INFORMS and are three of the foundational disciplines of AI. There is also a recognized need for interdisciplinary collaborations at all levels, i.e., not just in the different mathematical sciences, but also across the ethical, legal, and business dimensions. It is thus an ideal time for the O.R. community to jump in and take advantage of this recognition.

It is important, however, to recognize the cultural differences between O.R. and AI that may complicate these opportunities. Some of the obvious challenges are:

- The two fields use different vocabulary used for the same concepts, e.g., reinforcement learning in AI and (approximate) dynamic programming in O.R. The proliferation of closely related names, such as data science and analytics, has further exacerbated this issue.
- The missions of the two fields have been quite distinct: While AI focuses on building intelligent agents, O.R. has traditionally focused on process improvement to improve quality of service or efficiency. The criteria for the Edelman Award illustrate this point particularly clearly.
- As mentioned, the two fields have focused on different application areas, which makes communication and knowledge transfer more difficult.
- 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 O.R., despite providing some of the underlying technologies, only captures a small share (see the Teradata report on the "State of Artificial Intelligence for Enterprises").

The Role of O.R. in an AI World

Given the visibility of AI, it is important to consider the role and relevance of O.R. in an AI world. Historically, O.R. emerged from the search for a problem-solving framework grounded in sound mathematical principles: Data-driven decision-making based on the optimization of an objective function subject to constraints over a set of controllable variables. Nowadays, most machine learning problems, and many other AI tasks, are expressed as optimization problems. This is the case of supervised, semi-supervised, unsupervised, and reinforcement learning, path planning in robotics, and algorithmic configuration. Moreover, Stochastic Gradient Descent (SGD), a fundamental O.R. contribution, has become the workhorse for training deep neural networks and probably accounts, at this point in time, for more computing cycles than any other technology.

It is perhaps surprising for the INFORMS community to realize that natural language processing has been cast as an optimization problem for several decades now: What has changed is the nature of these optimization problems and the technology to solve them. Research areas in AI such as planning and scheduling, computational sustainability, AI for social good, and computational game theory often leverage discrete and stochastic optimization techniques. Moreover, 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. O.R. has thus provided some of the engines that have powered the AI revolution, with SGD being by far the most widely used at this point.

AI applications are pushing the boundaries of O.R. technologies. Many machine-learning tasks operate on massive datasets, challenging even the more scalable convex 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 INFORMS community, leveraging its expertise in optimization, stochastic modeling, and problem solving.

O.R. has a unique opportunity to reinvent itself in the age of AI. Possible avenues to perform this transformation include, but are not limited to:

- contributing new optimization algorithms that will provide the engine of the next-generation of AI tools;
- applying AI technologies, in conjunction with traditional O.R. tools, to provide better solutions and tackle new problems in areas such as manufacturing, logistics and supply chains, healthcare, energy, retail, and transportation, which have been the backbones of the O.R. community;
- developing a new generation of machine learning and optimization algorithms that tightly integrate AI and O.R. technologies, addressing the scalability, real-time, and multimodal requirements mentioned earlier;
- jointly working with the AI community to address the daunting challenges faced by autonomous agents, including robustness and reliability;
- investigating how AI and O.R. can cooperate on decision-support systems and human-in-the-loop decision-making, where O.R. has much to contribute;
- collaborating with the AI community to define the ethical and policy frameworks needed to deploy autonomous agents in a variety of settings; and
- contributing to the theoretical understanding of the performance and limitations of deep learning systems, and algorithms for training them.

There is no doubt that O.R. is relevant in an AI world and has significant scope for contributing foundational and algorithmic breakthroughs and novel applications that may have sustained impact on society. The question is how to rally the community to explore this opportunity and be a part of the AI revolution. One approach is to follow the eminently practical three-step plan outlined by Stephen Wright:

1. Demonstrate to the INFORMS community how O.R. is already used in AI. In machine learning, there is already considerable recognition of the importance of O.R. in AI. In planning and scheduling, sustainability, and constraint satisfaction, there has already been significant integration of the underlying technologies.
2. Explain to the INFORMS community that O.R. has more unique expertise to contribute to AI, in particular in discrete, robust, and stochastic optimization.
3. Recognize that AI is fundamentally multidisciplinary and that O.R. should fully participate and collaborate in this research space going forward, without being too quick to stake its own territory within this space. AI researchers have already contributed important advances to optimization algorithms, and similar contributions to other areas of O.R. can be expected in the future.

Simultaneously, the INFORMS community can work on shaping its own vision for O.R. in an AI world. Building on the AI stack presented earlier, the O.R. community has a unique opportunity to expand the research agenda for O.R. and leverage all the unique skills of the community. Figure 1 is a first attempt in that direction: It describes a possible O.R. stack, expanding the AI stack in multiple directions. 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 O.R. systems along multiple dimensions. It is also important to emphasize that systems implementing the O.R. stack will combine both data- and model-driven methodologies, pushing the field in unchartered and exciting territories.

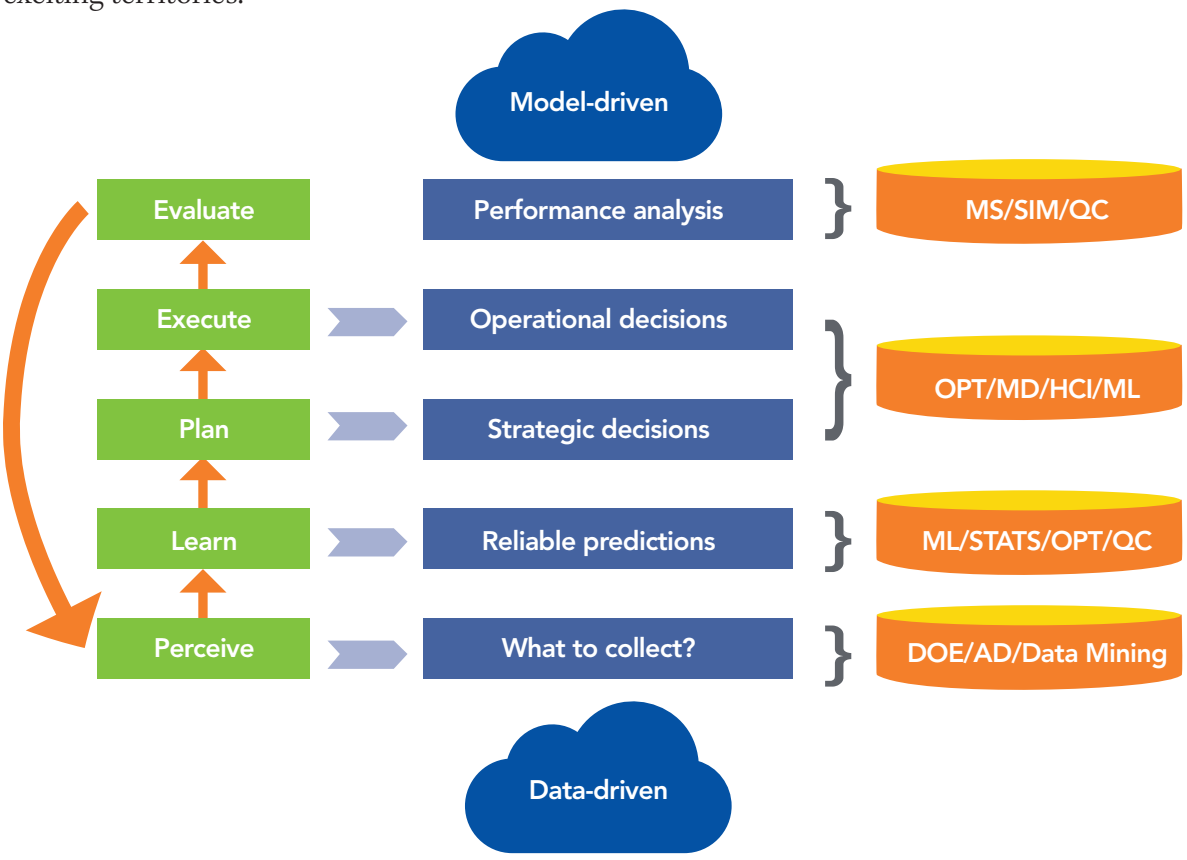

Figure 1: Toward an O.R. Stack

Engagement

This section summarizes engagement activities that can address the challenges and opportunities identified in the previous sections. It is important to emphasize that this section should not be regarded as a roadmap or an endorsement of specific action items, rather, it provides a collection of potential engagements from which INFORMS could derive such a roadmap.

O.R. Education: Threats and Opportunities resulting from AI Focus

The emergence of AI and data science is potentially an existential threat to OR/IE departments. It could be averted by pursuing the following three goals:

1. Raise awareness of AI and its impact on O.R. in the community.
2. Increase the presence of O.R. in data science initiatives.
3. Increase inclusiveness within the INFORMS community.

There are many ways to achieve these goals. Without trying to be exhaustive, they range from inviting prominent researchers from AI to INFORMS meetings, exercising leadership in universities to break the barriers to multidisciplinary research, creating new conferences, journals, awards, and competitions that encourage cross-fertilization between AI and O.R., to suggesting new curriculum guidelines and providing a forum for more applied and computational research.

There was significant agreement that, within OR/IE departments, there is a need to include AI courses in the O.R. curriculum and adapt existing classes to an AI-driven world. The need for advanced data science and computing skills (including in senior design projects) was a recurrent theme. The inclusion of O.R. content in computer science and data science programs (including the ACM curriculum) was also seen as an avenue to foster collaborations with O.R.

The issue of conference versus journal publications is important to debate. AI research is disseminated through conference proceedings and public repositories such as ArXiv. It results in a highly dynamic and exciting field with short reviewing times. In general, the pace of publication in O.R. is significantly slower but many in the community primarily value the rigorous review process of existing journals. It is important to evaluate whether other communities (e.g., IEEE) provide INFORMS with a middle ground approach, where journals review papers (with page limits) quickly. Most importantly, it is critical to recognize leading AI and computer science conferences in hiring, promoting, and tenure decisions.

The number of opportunities to improve inclusiveness is unlimited: they range from new journals dedicated to data science and/or bridging the gap between the more theoretical journals and and application journals, to special issues, tutorials, and *Editor's Cuts* focusing on AI. The lack of conferences and journals publishing applied work was repeatedly mentioned as an impediment to bridge the gap between the communities.

However, perhaps the most significant opportunity lies in the education realm: To pursue large-scale educational activities that will train the *next generation* of practitioners and researchers in AI, O.R., and statistics. For example, a general-audience basic O.R. course would be valuable to increase awareness of O.R. methods to a larger body of students (similar to basic statistics courses that are available across many campuses). These students will then be in a position to apply AI techniques to traditional O.R. problems, O.R. techniques to traditional AI problems, and – most importantly – to develop the science and technology that will blend AI, O.R., and statistical techniques in novel ways to address the numerous challenges that modern society faces.

O.R. Research: The Need to Tackle High-Visibility Challenges?

On the research side, it is important to stress the opportunities emerging from AI: They open new avenues for O.R. research and raise new challenges for the community, often in novel applications. Research in AI is often driven by a number of high-visibility challenges (e.g., computer vision, games, robot cup) and their associated datasets can be used for measuring progress. It demonstrates success and captures the public imagination when AI exhibits super-human performance in areas that were considered outside the realm of their capabilities. The INFORMS community may draw some inspiration from this style of goal-oriented research. Student & early-career competitions and hackathons are also valuable steps in this direction.

Fostering Research Collaboration

Several sister societies have been proactive in integrating AI research into their conferences. AI and deep learning are significant topics in SIAM conferences, ASIS&T has prominently featured AI in keynotes of their annual conferences, and EURO had a stream on optimization and data science in 2019. Statistical organizations, although they have better brand recognition than O.R., are also aware of the need to foster collaborations and are considering workshops and training programs to ensure that AI uses the best statistical methods available.

INFORMS can play a leadership role in fostering a multidisciplinary culture by providing funding incentives (including through university-wide initiatives), organizing multidisciplinary workshops with funding agencies and multidisciplinary summer schools to educate the next generation of O.R. researchers. Interactions with the ACM SIGAI organization seems particularly desirable.

Branding

There is a wide consensus about the lack of mindshare and brand recognition for O.R. As previously mentioned, AI has captured the public imagination by solving applications that are associated with human intelligence. There is a certain halo around AI and it will not disappear in the near future, given the substantial science and technology underlying AI applications. O.R. on the other hand does not reach the public awareness despite significant and sustained contributions to society. The community, however, has not reached a consensus about how best to tackle this branding dilemma. Some suggest viewing O.R. as a building block for AI, while others would like to present O.R. as the next "big" thing. There is, however, a large agreement about the need to demonstrate:

1. the unique capabilities of O.R. to advance AI research, including discrete and robust optimization and a variety of statistical and simulation methods; and
2. the numerous application domains of O.R., including manufacturing, supply chains, energy, revenue management, inventory management and control.

There is also a need to make O.R. more accessible by promoting, and removing any barriers to, entry to the fields through TED Talks, Fireside chats, videos, and tutorials. These initiatives will provide the basic literacy of the field and INFORMS can take a leading role in curating these.

INFORMS has another significant opportunity: *to take a bold stand in helping to define AI and what it really means for businesses*. Most companies have recognized the need to invest in AI and face challenges in how to adopt and deploy the technology. O.R. practitioners are ideally positioned to facilitate the adoption of AI for business and industrial problems, highlighting the fact that succeeding with AI requires a strong suite of analytical methods and capabilities, all of which come under the umbrella of advanced analytical methods.

Efforts with Funding Agencies

The funding gap between AI and O.R. is substantial and will further grow in the coming years. It is critically important to raise awareness of O.R. in funding agencies and at OSTP, and to organize agenda-setting workshops at the intersection of AI and O.R., including at NSF and the CCC. There is a significant opportunity in articulating far-reaching *multidisciplinary* research programs that can benefit society as a whole.

References

1. David Silver et al., 2017, "Mastering the game of Go without human knowledge," *Nature*, Vol. 550, pp. 354-359.
2. Russell, Stuart and Peter Norvig, 2010, "Artificial Intelligence: A Modern Approach," Pearson: New York.