QUANTUM OPTIMIZATION

Quantum Computing Meets Optimization

CORONA-LIU COLLABORATION

Real Problems, Real Students, Real Results

REMEMBERING RONALD A. HOWARD

The Man Who Engineered Good Decisions

CELEBRATING 50 YEARS

ORMS TODAY

2026 | Volume 53, Issue 1

informs. Smarter Decisions for a Better World

HOW ANALYTICS IS REMAKING BASKETBALL

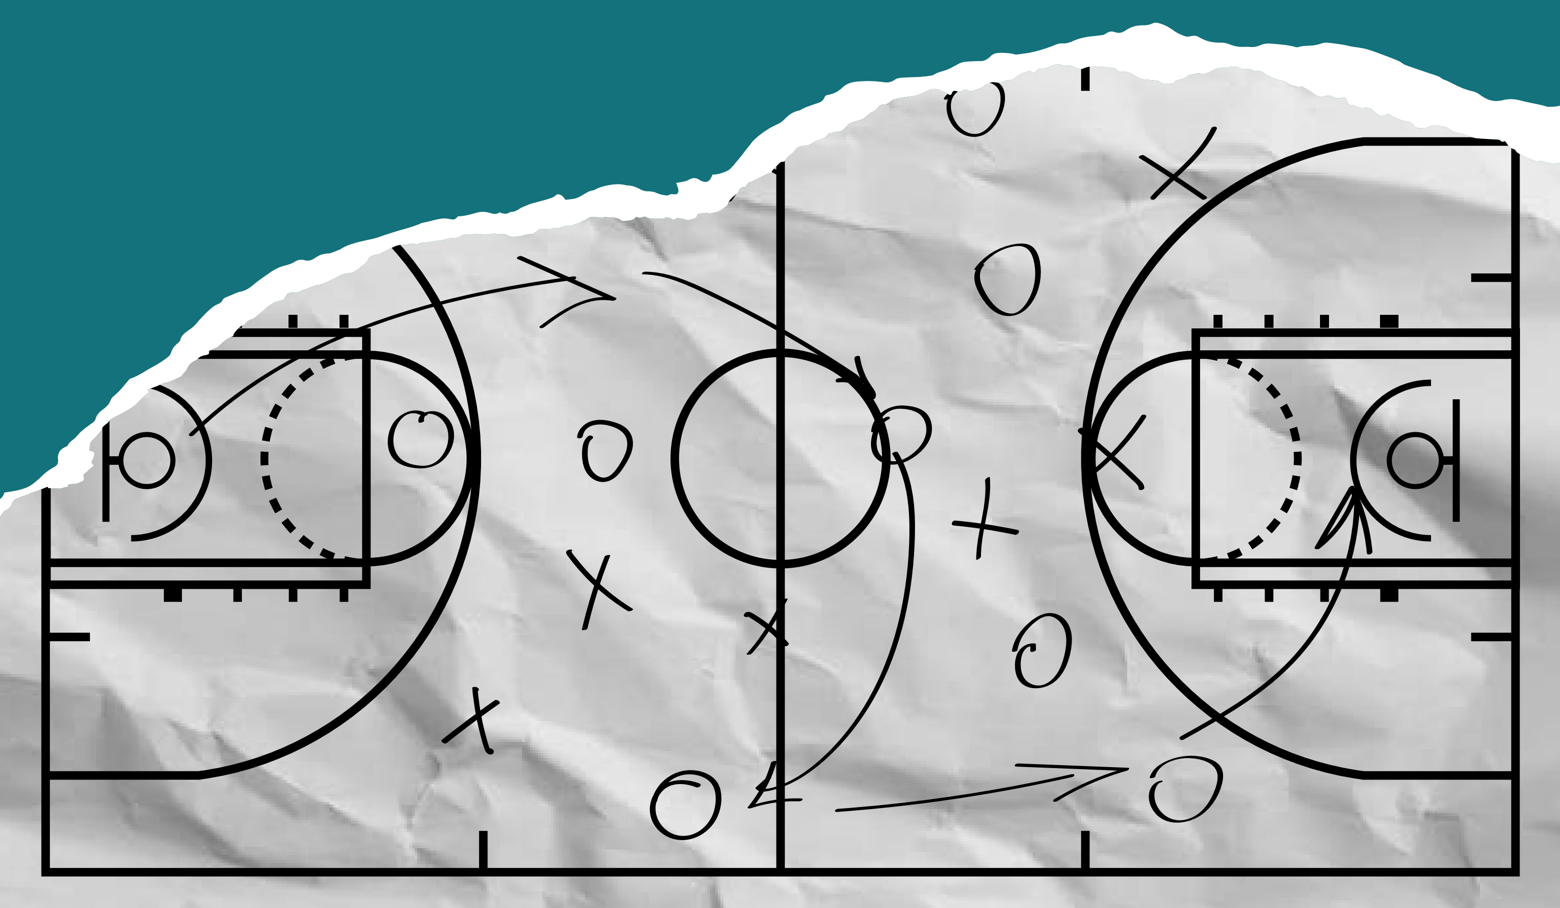

OR/MS TODAY ADVERTISING AND EDITORIAL OFFICE

5521 Research Park Drive, Suite 200, Catonsville, MD 21228
Telephone: 443.757.3500
Email: informs@informs.org

Editor
Kara Tucker, ktucker@informs.org

Director of Marketing
Mary Leszczynski, mleszczynski@informs.org

Director of Publications
Matthew Walls, mwalls@informs.org

Advertising Sales
advertising@informs.org

Production
Emily Arnold, earnold@informs.org
Eilyn Cubillo, ecubillo@informs.org

Magazine Editorial Advisory Board
Anna Nagurney, Committee Chair

INFORMS Online
www.informs.org
ormstoday.informs.org

ORMS TODAY

OR/MS Today (ISSN 1085-1038) is published quarterly by the Institute for Operations Research and the Management Sciences (INFORMS). Canada Post International Publications Mail (Canadian Distribution) Sales Agreement No. 1220047. Deadlines for contributions: Manuscripts and news items should arrive no later than six weeks prior to the first day of the month of publication. Address correspondence regarding editorial content to the director, Mary Leszczynski, at mleszczynski@informs.org. The opinions expressed in *OR/MS Today* are those of the authors and do not necessarily reflect the opinions of INFORMS, its officers or the editorial staff of *OR/MS Today*. Membership subscriptions for *OR/MS Today* are included in annual dues. INFORMS offers nonmember subscriptions to institutions; the rate is $62 for USA, $85 all other countries. Single copies can be purchased for $25 plus postage. Periodicals postage paid at Catonsville, MD, and additional mailing offices. Printed in the United States of America. POSTMASTER: Send address changes to *OR/MS Today*, INFORMS, 5521 Research Park Dr., Suite 200, Catonsville, MD 21228.

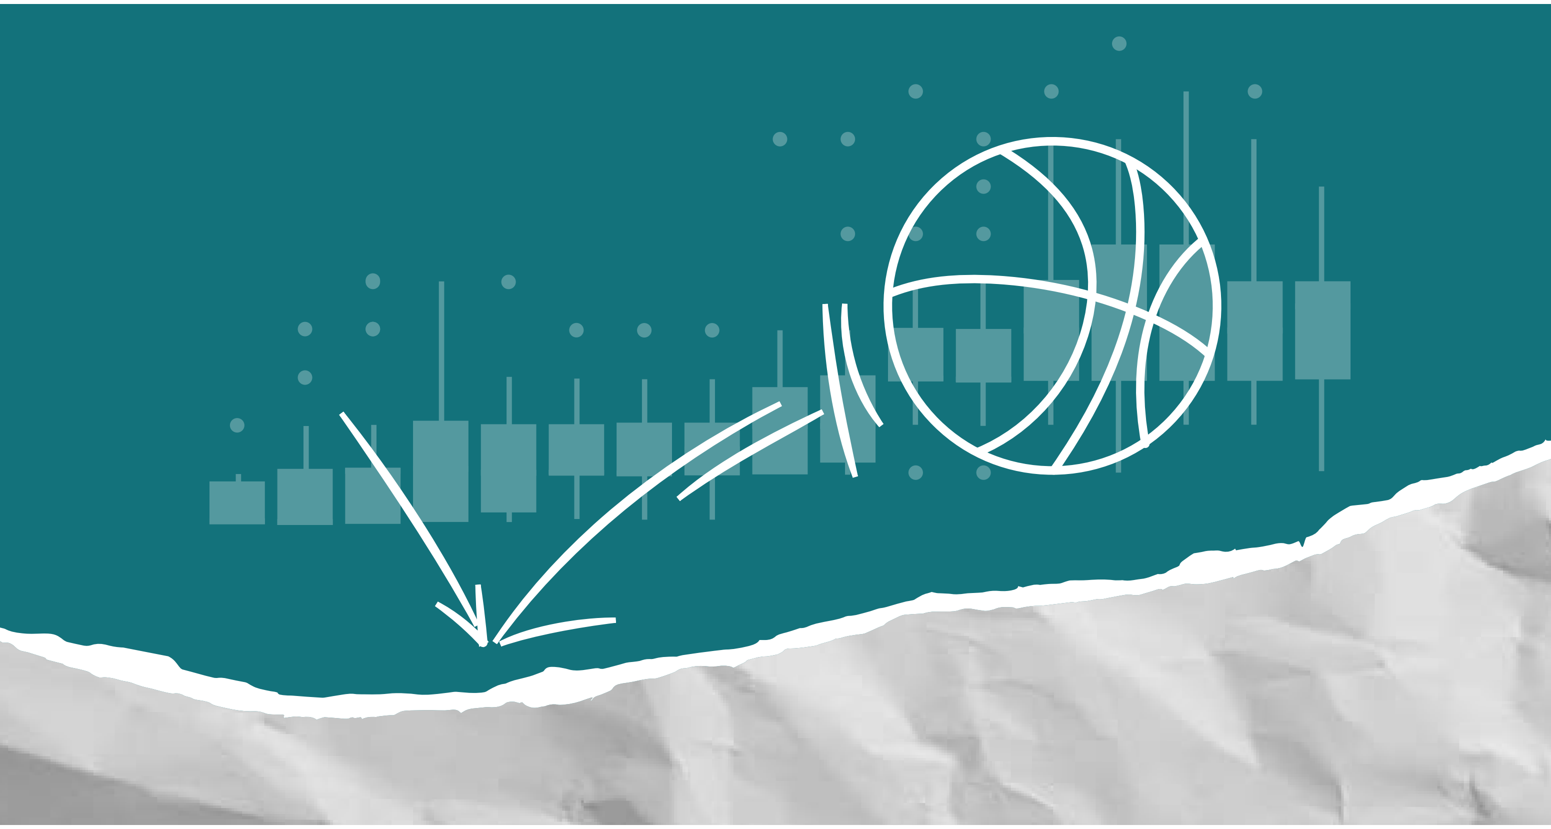

DEPARTMENTS

MORE THAN A MAGAZINE

BY KARA TUCKER

AFTER MORE THAN 13 YEARS WITH INFORMS, this issue of *OR/MS Today* will be my last as editor. Writing that sentence feels both surreal and deeply meaningful. For nearly a decade, this magazine has been more than a publication to me – it has been a front-row seat to the ingenuity, curiosity and impact of a remarkable community.

When I first started working at INFORMS, I understood the importance of the organization's mission (sort of). At the time, I was a production editor for a handful of INFORMS' peer-reviewed journals, and although I was not well versed with what operations research and management science were, I knew that OR/MS was indeed important in making smarter decisions for a better world. What I did not yet fully appreciate – and what I couldn't fully comprehend until I switched gears to work on the magazines – was the depth of generosity within this community. Members entrusted me with their stories – stories of breakthrough models; complex implementations; hard-won lessons; unexpected career paths; and bold ideas about the future of analytics, AI and quantum optimization. That trust has been one of the greatest privileges of my career.

From my very first issue, I quickly learned that *OR/MS Today* is not just about methodologies, technologies or case studies. It is about people. It is about the members who devote their time and expertise to solving complex problems, advancing industries and improving lives, often without fanfare. My role as editor gave me the privilege of helping share those stories, and I am profoundly grateful to everyone who trusted me to help tell them.

One of the great joys of this position is the opportunity to work directly with members from across disciplines, career stages and geographic regions. Whether collaborating on feature articles, highlighting innovative applications, summarizing award-winning projects or showcasing emerging voices in the field, I was continually inspired by the willingness of members to share not only their successes but also their failures and lessons learned, as well as their passion for advancing analytics and operations research.

Throughout the years, I also had the additional joy of connecting with many of you through What's Your StORy? interviews (which began as a short written Q&A and has evolved into a full-production studio sit-down interview, complete with a green screen and INFORMS coffee mug!) and the newer Member in a Minute video shorts on the INFORMS YouTube channel. These conversations allowed us to step beyond professional accomplishments and explore personal journeys and fun facts, motivations and sometimes even unexpected paths that lead individuals into OR/MS. Sitting down with members to hear their stories

The INFORMS Marketing team poses with board members during board meeting on January 14, 2026.

reinforced something I have long believed: the strength of INFORMS lies not only in the excellence of its research and practice but in the authenticity and generosity of its people – and creating a welcoming environment for these people to come together and connect.

Through these interactions, one theme became clear: INFORMS is more than a meeting or just your *professional* association. It is a community that aspires to be a lifelong *home* – a place where members can grow, connect and contribute across every stage of their careers. Over time, INFORMS became that kind of home for me as well. I grew here – as an editor, a working mom, a teammate and collaborator, and as someone continually learning from a field that never stands still. I formed deep friendships here and found purpose in helping elevate the work of others.

Editing this magazine also gave me a unique vantage point on the evolution of our discipline. Over 13+ years, I have watched analytics expand into new industries, witnessed emerging technologies transform practice and seen our community respond thoughtfully to societal challenges that demand rigorous, ethical decision-making. The tools may change, the platforms may evolve and the buzzwords may rotate, but the underlying commitment of INFORMS members to clarity, impact and intellectual rigor has remained constant.

Of course, none of this work happens alone. I am deeply grateful to the authors who met deadlines (and occasionally asked for extensions to which I always obliged, ahem, Jeff Cohen); to the columnists who reliably delivered insight and perspective; to Pete Horner, Mary Leszczynski, Jim Cochran, Anna Nagurney, Tinglong Dai, Anne Robinson and many others for being a guiding post in my role and lending their expertise when I needed it; and finally, to my INFORMS colleagues – old teams and new – who care deeply about quality and service to our members. Producing each issue of *OR/MS Today* has always been a collaborative effort, and I have been fortunate to work alongside individuals who prioritize quality, style and accuracy. And who have made my days/weeks/months/years/decades *fun* and allowed me to fix their typos and grammar without taking it personally :).

I am also thankful for the many members who reached out over the years with story ideas, thoughtful feedback or a simple note of appreciation. Those interactions consistently reminded me why this publication matters.

For me, this transition is bittersweet. It is never easy to leave a place that has shaped you so profoundly. But I leave with gratitude – for the trust you placed in me, for the stories you shared and for the opportunity to help reflect this community back to itself in the pages of this magazine.

Although I am stepping away from my role at INFORMS, I remain a strong believer in the organization's mission and the value it provides to members around the world. I have no doubt that *OR/MS Today* will continue to evolve and thrive, reflecting the innovation and energy of the community it serves.

To everyone who contributed an article, agreed to an interview, appeared on camera, shared a story idea or simply read the magazine, thank you. It has been a privilege to help document and celebrate your work.

As I close this chapter, I do so with immense gratitude, pride in what this publication represents and excitement for what lies ahead for INFORMS and its members – and me!

Good night, and good luck.

KARA TUCKER is editor of *OR/MS Today* and *Analytics* magazines.

SETTING THE STAGE

BY MARK LEWIS

FROM ITS INCEPTION, INFORMS HAS BEEN a place for people interested in the theory and practice of operations research (O.R.) to connect over shared ideas and to actualize their potential through community. As this year's INFORMS president, I am honored to continue this tradition while also leading our community to consider the most pressing issues and innovations affecting the field. I would like to use my first President's Desk column in *OR/MS Today* to communicate my priorities over the coming year. First, the board and I aim to elevate several areas growing in prominence across the OR/MS landscape where we believe our community can (and should) expand its impact. My second intention is to reaffirm our commitment to inclusivity within and beyond the INFORMS community. There is significance in being first in any role, and I am keenly aware of this as the first African American INFORMS president. Although some might assert that personal identities have little impact on how you show up and engage in work, I believe, in fact, that our personal identities significantly inform how we see the world; it is this unique perspective, in addition to our technical education and training, that elevates and enhances the work that we engage in and the outcomes that we pursue. We are a vibrant community looking to the science of OR/MS to make smarter decisions for a better world.

New Committees and More

INFORMS has a longstanding committee on artificial intelligence (AI). Today, AI has an impact in every arena. There remain exciting technical questions in AI in terms of improving machine learning (ML) and reinforcement learning (RL) algorithms. In addition, there are implementations that use AI to improve or enhance a user's experience. For example, recommender systems [1] are necessarily equipped with ML and RL algorithms. There are also still questions about how AI will influence education, both in how lessons are designed and in how they are delivered. Last year, thanks to my predecessor David Hunt, INFORMS started an ad hoc committee on quantum computing [2]. In the world of large-scale computational needs, building computers that can meet these needs is crucial and, I think, within reach.

This year, I will seek board approval to institute two further ad hoc committees. The first is partially in response to the needs for computing power and partially in response to what appears to be a reduced focus of the federal government. This committee will consider an INFORMS-wide focus on **energy and sustainability**. Although clearly not a new topic, I feel that the OR/MS community is poised to meet new challenges. For example, what should be the response for an increased demand for energy for computing power? Should there be further incentive programs for homes to move toward solar (decreasing demands on the grid)? What lessons can be learned from our international colleagues on battery exchange programs for electric vehicles (and what new challenges arise)? What can logistics companies do to ensure they have a reduced carbon footprint, given the expectation of near-next-day delivery of goods?

The second committee I would like to start is on **humanitarian operations and logistics**. This, again, is not a new area for the OR/MS community, but new policies require a renewed vision. With the U.S. Agency for International Development cut worldwide and Supplemental Nutrition Assistance Program (SNAP) benefits reduced domestically [3], what should be the response in terms of supply procurement and dissemination? Given the increased demand for benefits, how might the community respond in places without a long-term infrastructure in place? These topics have natural intersections with the scope of work already advanced by the Energy, Natural Resources, and the Environment (ENRE) and Public and Societal Operations Research (PSOR) communities within INFORMS. On the other hand, much like AI, I imagine with the scale of each of the problems, the ad hoc committees will find cross-cutting applications. Indeed, the battery exchange issue previously mentioned could be closely connected to Location Analysis (another INFORMS community). Humanitarian operations will often find synergies with the Health Applications Society as well.

What I hope to accomplish by instituting each of these committees is to raise the importance of these areas community-wide. In doing so, I hope to encourage even further cross-pollination, whether it be from academics in adjacent fields or between academia and industry. On that last point, I believe that all of these areas are closer than they have been since the ORSA/TIMS days. That is to say, the research done in academia can be virtually directly applied to the real world, which would then influence more problems considered by academia in a cycle.

Circling back to my second intention for this article, I hope that every member of INFORMS feels welcomed as a part of this community. If you have yet to find your place, I invite you to explore not only our societies and sections but also our affinity groups – Minority Issues Forum (MIF), Women in OR/MS (WORMS), PRIDE Forum, the Military Veterans Interest Forum, etc. – for both mentorship and friendship.

MARK LEWIS is the Maxwell M. Upson Professor of Engineering at Cornell University. He is the 2026 INFORMS President.

REFERENCES

1. S. Raza, M. Rahman, S. Kamawal, A. Toroghi, A. Raval, F. Navah and A. Kazemeini, 2025, "A Comprehensive Review of Recommender Systems: Transitioning from Theory to Practice," *Computer Science Review*, Vol. 59, Art. no. 100849.
2. D. Hunt, 2025, "Introducing the INFORMS Quantum Computing and Operations Research Ad Hoc Committee," *OR/MS Today*, Vol. 52, No. 1, January 31, https://doi.org/10.1287/orms.2025.01.02.
3. J. Joffe-Block, 2025, "The Agriculture Secretary Says SNAP Changes Are Coming. Here's What We Know," *NPR*, December 1, https://www.whro.org/2025-12-01/the-agriculture-secretary-says-snap-changes-are-coming-heres-what-we-know.

THANK YOU TO OUR 2026 INFORMS ANALYTICS+ CONFERENCE SPONSORS!

informs ANALYTICS+ CONFERENCE

DIAMOND LEVEL

hexaly — A TRADEMARK OF LOCALSOLVER SAS

PLATINUM LEVEL

FICO

GUROBI OPTIMIZATION

informs ROUND TABLE

SILVER LEVEL

GAMS

BRONZE LEVEL

oat

ADDITIONAL SPONSORS

Day 2 Keynote Speaker — informs ANALYTICS SOCIETY

UPS George D Smith Prize — ups

TOP SOCIAL POSTS

Check out the top 8 most clicked items on INFORMS' social media channels last month. Connect with us on your favorite channel!

1. INFORMS staff were thrilled to host our INFORMS Board members at the home office today!
INFORMS X

2. Calling All Analytics Innovators! [Submit a Presentation] 2026 INFORMS Analytics+ Conference
INFORMS LinkedIn

3. This episode of What's Your StORy? features David E. Bernal Neira from Purdue University
INFORMS YouTube

4. Lego uses bricks to demystify AI
Axios

5. Discover how effective communication and strong interpersonal skills can drive your professional success
INFORMS LinkedIn

6. Free Webinar: Master Uncertainty in Decision-Making
INFORMS YouTube

7. INFORMS Member Spotlight at CES! INFORMS Fellow Sridhar Tayur, Professor at Carnegie Mellon University
INFORMS LinkedIn

8. INFORMS Names Six Finalists for the 2026 Franz Edelman Award
INFORMS Newsroom

RECENT PRESS RELEASES

Rethinking Restroom Equity: How Mixed-Gender Designs Improve Access in Public Venues

INFORMS Names Six Finalists for the 2026 Franz Edelman Award, the World's Leading Honor in Analytics, Operations Research and Management Science

New Study Reveals People Judge Lines by What's Ahead – Not How Long They Wait

When Businesses Can't Silence Reviews, Consumers Tell the Truth, New Study Finds

MEMBERS IN THE NEWS

Are the Shutdown-Era Flight Disruptions Over?, *Time*, featuring Sheldon Jacobson

A Seismic Shift in Computing Is on the Horizon (and It's Not AI), *CNN Business*, featuring Sridhar Tayur

Are We Ready or the Tidal Wave of AI-Generated Video?, *Forbes*, by Anjana Susarla

Staffing Shortages May Ground Air Travel Even Post-Shutdown, *WORT*, with Laura Albert

RESOUNDINGLY HUMAN: The INFORMS Podcast

December
We wrapped up an exciting year for the INFORMS podcast with Pascal Van Hentenryck, Georgia Institute of Technology, a 2025 INFORMS Annual Meeting Plenary Speaker. Pascal discusses the impact of machine learning and optimization working hand-in-hand … as well as the best cookie recipe according to AI! We also share one final update from our 2025 INFORMS President Dave Hunt with a look back over the past year.

January
We kicked off 2026 with long-time INFORMS member Jim Cochran, University of Alabama, who shares his personal INFORMS journey from his first meeting to becoming an INFORMS Fellow. Then, 2026 INFORMS President Mark Lewis shares his goals and objectives for INFORMS in the coming year.

February
This month, we shine a spotlight on a growing buzzword – quantum computing – with Tamás Terlaky, Lehigh University, and explore its incredible potential. Then, we kick off our coverage of the upcoming 2026 INFORMS Analytics+ Conference with Keynote Speaker Scott Taylor, The Data Whisperer. Then, we touch base with the 2026 INFORMS President Mark Lewis, who gives us a look ahead at his goals and objectives for INFORMS in the coming year.

@INFORMS

@INFORMS_orms

EXECUTIVE DIRECTOR'S LOG: NOTES FROM DAY 24 OF 2026

BY ELENA GERSTMANN

AT BREAKFAST THIS MORNING, I WAS thinking about what to write for my first *OR/MS Today* column of 2026, and I felt an unexpected mix of joy and pride bubble up. That doesn't always happen when you're staring at a blank screen eating a bowl of cereal, so I paused and asked myself where it was coming from.

Three things came to mind immediately. None of them are flashy on their own, but together they tell a story that feels worth sharing. A story that should make you as members proud.

First, our journals are reaching the world in a way they never have before. Late last year, INFORMS completed the transition of its institutional distribution to EBSCO Business Source, significantly expanding access to our journal content through libraries and institutions worldwide. INFORMS continues to serve members directly; EBSCO now handles all online institutional access. At the time, we believed this shift would increase the reach and impact of our community's work. What we didn't yet have was evidence.

Now we do.

When I reviewed download data from November 2025, I saw that approximately 80% of article downloads came from readers at institutions that were not previously subscribers to INFORMS content. I double-checked the numbers with colleagues to be sure I was interpreting them correctly.

When INFORMS hosted and sold its own content, our journals were available to about 2,500 subscribing institutions. Through EBSCO, they are now accessible to more than 30,000 institutions worldwide, in over 150 countries. In 2025, total downloads of INFORMS journal content neared 8 million, about a 30% increase over 2024, which was also a record year.

Downloads are not a perfect measure of impact, but they are a meaningful signal of discoverability and access. More people are finding our work – and using it. This is exactly what we mean by our strategic goal to "advance the science and technology of decision-making and elevate its impact." This community does extraordinary work and getting that work into the hands of people who can learn from and build on it is part of our responsibility as a professional association.

Second, the Franz Edelman Award finalists were announced. As always, the list is impressive. Organizations are applying analytics, operations research, artificial intelligence and management science to challenges ranging from supply chains and sustainability to public food distribution systems. The work is rigorous, creative and consequential.

This year's six finalists include Chewy; the Department of Food and Public Distribution, Government of India; ECCO; Google; Microsoft; and NVIDIA. As with every year, the range of organizations and problem domains is striking. Each was selected through the same rigorous review process, and each represents outstanding applied work.

What caught my attention this year is that three of the finalists – Google, Microsoft and NVIDIA – are among the largest and most influential technology companies in the world. These organizations are shaping global markets, computing infrastructure and the future of artificial intelligence. And they chose to submit their work to the Edelman Competition.

Essentially, they are asking volunteer experts from the INFORMS community to evaluate their applications alongside those from other finalists. They are choosing to be assessed through a process grounded in rigor, evidence and documented impact. **That matters.** It speaks to the standards this community has set, and the role INFORMS continues to play in defining what excellence looks like in applied analytics.

Finally, there is a quieter development that is just as significant. Through the U.S. FY2026 appropriations process, funding for the National Science Foundation (NSF) was preserved despite a budget proposal that would have cut it by more than 50%. Those proposed cuts would have been devastating for basic research and the broader ecosystem that supports scientific discovery.

Congress ultimately rejected those cuts. NSF funding was stabilized to near current levels, well above what had been proposed.

This didn't happen easily or automatically. It reflected sustained advocacy across the research community, including professional associations, coalitions, institutions and individual researchers who made the case for why public investment in science matters. INFORMS played a role in that effort through advocacy tools for members, letters to agency leaders, partnerships with peer societies and participation in Capitol Hill visits. Many of you added your voices directly. Thank you.

For many of our academic members, NSF funding makes it possible to do the work they were trained to do – work that leads to new methods, new insights and real-world impact. Preserving that support helps ensure the next generation of research and researchers can move forward.

So here I am, on day 24 of 2026, feeling genuinely encouraged heading into the rest of this year.

Our knowledge is reaching more people than ever before. Some of the most influential organizations in the world are seeking evaluation through our community's stamp of approval. And collective advocacy helped protect the scientific infrastructure that so many of our members – and frankly, the world – rely on.

None of this happened by accident. It reflects strategic choices, volunteer and staff leadership, member engagement and a shared commitment to smarter decision-making for a better world.

That feels like a great way to start the year.

ELENA GERSTMANN, Ph.D., FASAE, CAE (egerstmann@informs.org), is the executive director of INFORMS (5521 Research Park Drive, Ste. 200, Catonsville, MD 21228). She can be reached via email.

2025 ANNUAL SUMMARY OF INFORMS CODE OF CONDUCT COMPLAINTS

In the spirit of transparency and as part of our commitment to provide a welcoming and safe environment for all, the executive director of INFORMS provides an annual summary of the previous year's code of conduct complaints.

	QUANTITY
TOTAL COMPLAINTS FILED VIA LIGHTHOUSE SERVICES	3
A. INVESTIGATIONS	
Complaints Not Requiring Internal Investigation	2
Complaints Requiring Internal Investigation	1
Complaints Requiring Internal Investigation **and** a 3rd Party External Investigation	0
B. RESOLUTIONS	
Complaints Still Under Investigation	0
Complaints resolved and closed	1
C. APPEALS	
Appeals Filed	0

For more information about INFORMS Code of Conduct, scan the QR Code or visit: www.informs.org/About-INFORMS/ Governance/INFORMS-Code-of-Conduct

CREATING A BUSINESS ANALYTICS DEPARTMENT RESEARCH RANKING

BY NICK STREET, SAMUEL BURER AND KANG ZHAO

Which academic business analytics (BA) departments produce the most research?

THIS IS A DECEPTIVELY SIMPLE QUESTION that many in our field are interested in – from prospective students to academic administrators and industry partners. But finding a rigorous, reliable answer is not easy.

Existing rankings are of little help. Reputation-based program rankings such as the *U.S. News & World Report* are notoriously subjective. For example, before the 2010s, many institutions in the *U.S. News* business analytics rankings did not even offer a BA degree! Meanwhile, research-specific rankings based on existing lists of publication outlets, such as the University of Texas at Dallas (UT-Dallas) and *Financial Times* (*FT*), focus more on traditional business disciplines including management information systems and supply chain management.

At the University of Iowa's Tippie College of Business, we decided to design a ranking that is based on data and would answer our question by acknowledging that BA programs encompass a wide swath of academic fields. We unveiled our ranking and the methodology behind it at the 2025 INFORMS Annual Meeting in Atlanta last October.

As an emerging, multidisciplinary academic field, BA even being defined as a discipline – or what counts as a BA department or BA publication – is nontrivial. We fortunately had lots of experience grappling with these issues over the past decade.

In our design of BA academic programs [1], for which we won the 2021 UPS George D. Smith Prize, we adopted a methods-centered approach that combines computing, data science and applied mathematics (optimization and statistics) with domain-specific work in operations management.

We also created our own evidence-based list of A-level publication outlets [2], primarily for use by our students and early-career faculty to focus their work in the most impactful directions. This list incorporates the top outlets from data science (e.g., KDD and NeurIPS), optimization (e.g., *SIAM Journal on Optimization*) and statistics (e.g., *Journal of the American Statistical Association*), along with leading traditional business outlets. The result was a comprehensive set of publication targets for the BA field.

This brings us back to the question of which BA departments produce the most and best research. We realized that our list of A-level outlets would provide solid ground for a comprehensive ranking of BA departments. However, multiple practical problems remained.

Given the multidisciplinary nature of the A-list, and the fact that publication metadata often does not include departmental affiliations, we realized that we can't simply search each outlet for authors from a particular university. This would result in publication counts that measured, or even emphasized, the quality of the respective computer science, statistics or mathematics departments at those universities, which is very much not the goal. Instead, we want to focus on BA departments housed in business schools.

So, our data collection would have to proceed from the department level rather than the outlet or university level. Which reveals the second problem: What even counts as a BA department? Every business school divides the academic landscape in its own way, based on everything from collegiate priorities to historical legacies. Particularly in a relatively young and multidisciplinary field such as BA, there is an enormous range of department names and faculty compositions.

We overcame these barriers with four primary steps:

1. **Identify the right departments.** We started with a collection of about 110 universities, gathered from various sources: *U.S. News* rankings, UT-Dallas and FT50 rankings, Association of American Universities membership, etc. (Note that this set of universities is intended to serve as a seed set. If we missed your school and you would like to be included, please just let us know!) Within the college or school of business at each seed university, we identified the department having the best alignment with our definition of the field.
2. **Collect faculty names.** Within each of the identified departments, we manually scanned departmental/collegiate websites to collect the faculty roster. We chose to include only tenured and tenure-track faculty in this step, which has its own set of design choices and limitations – e.g., research active instructional track or emeritus faculty were excluded, as were Ph.D. students, but faculty with primary administrative appointments were included if their name appeared on the faculty roster. This collection was done in the middle of the fall semester so that all departments had enough time to update their web rosters for the latest academic year.
3. **Gather publication data.** For each faculty member, we queried the Scopus publication database. Scopus was chosen because of its relatively complete indexing of both journal and conference publications (cf Web of Science) and because of the ease of direct queries (cf Google Scholar). We chose a five-year window of publication data, corresponding to the period used by the Association to Advance Collegiate Schools of Business (AACSB). We note here another limitation: similar to AACSB academic productivity reporting, we are gathering data on each department's *current* faculty. Counts can therefore be dramatically affected by recent hirings, departures and retirements.
4. **Collate results**. Finally, departmental totals for both total publications and A-list publications were aggregated. In addition to total publications of both types, we computed per-faculty-member averages.

Our initial set of rankings was circulated internally and sent to a few academic colleagues around the country for feedback. Then, the first official rankings reflecting the time period 2020-2024 went live just in time for the 2025 INFORMS Annual Meeting in late October. Faculty rosters were gathered in fall 2025. The ranking can be found at https://tippie.uiowa.edu/faculty-research/research-expertise-interests/business-analytics-research-rankings and can be sorted in multiple ways, although we suggest that A-list papers per faculty member is the most informative.

Much of the top end of our rankings confirm conventional wisdom. It's no surprise to see traditional powerhouse departments such as the Massachusetts Institute of Technology, the University of Arizona and University of Florida near the top, but one might not have expected to see the University of California, Los Angeles or Temple University quite so high on the list. As a shameless plug, Iowa comes in tied with University of North Carolina at Chapel Hill at No. 15, one step below New York University.

We are already working on the next iteration of rankings and are expecting significant improvements. Basic data collection errors – Did we miss a university? Include the wrong department? Fail to identify the right faculty? – will be corrected, and we ask the community to send those corrections directly to tippie-research@uiowa.edu.

We will improve other aspects, too. The current ranking is limited to U.S. schools; the decision of whether to cast a geographically wider net is under consideration. The web-based publication does not currently offer filtering to change the list of journals to fit individual needs. Finally, the next release will incorporate citation information to quantify academic impact as well as productivity.

We hope this tool proves a useful step forward, clarifying the definition of this multidisciplinary field and providing comparative evaluation for academics, administrators and prospective students. We look forward to your feedback.

NICK STREET is a professor of business analytics and associate dean for research and Ph.D. programs at the University of Iowa.

SAMUEL BURER is the Tippie-Rollins Professor in Business Analytics and department executive officer at the University of Iowa.

KANG ZHAO is the Leonard Hadley Professor in Business Analytics and senior associate dean at the University of Iowa.

REFERENCES

1. A. M. Campbell and N. Street, 2021, "The evolution of analytics education at Tippie," *OR/MS Today*, August 5, https://doi.org/10.1287/orms.2021.04.12.
2. B. W. Thomas and A. M. Campbell, 2022, "Creating a business analytics publication list," *OR/MS Today*, February 4, https://doi.org/10.1287/orms.2022.01.11.

WHAT IS LEARNING? THE TWO SIDES OF LINEAR REGRESSION

BY KEN FORDYCE

AS FIRMS WORK TO MOVE ARTIFICIAL intelligence (AI) into effective decision support [1] for managing demand supply networks, the hot topic is, "What is learning?" The purpose of this article is to generate discussion on this topic and provide my two cents.

The linear statistical model, aka linear regression (LR) or "best fit" straight line, is a well-known statistical modeling method and a foundational method in machine learning (ML). The core of LR is to find the straight line that best fits the data points collected. Our example is lot size versus man-hours, where we want to predict man-hours based on lot size. This data is provided in Table 1 (columns 1-3) and Figure 1.

A primer on linear regression can be found in the 2024 Q1 article in the *International Journal of Applied Forecasting* [2]. We want to find values for β_1 and β_0 (coefficients) such that the estimated man-hours best match the actuals, where estimated hours = $\beta_0 + \beta_1 \times$ lot size. The "best" values are called b_0 and b_1. "Best" is defined as minimizing the difference between the actual value and predicted value.

Within the ML community, LR is often the starting point for supervised or error-based learning. Chapter 7 of "Fundamentals of Machine Learning for Predictive Data Analytics" [3] provides a detailed description of LR from an ML perspective. ML relies on a method called gradient descent [4] that iteratively [5] "learns" its way to the best values for β_0 and β_1. This sounds exciting and clearly is learning.

But maybe there is an inconvenient fact – LR was invented in the early 1800s [6] and has been a staple of statistical modeling for more than 100 years. So how did they find the best values for β_0 and β_1 in the "before times?" To answer this question, we go to Chapter 7 in the classic text "Applied Linear Statistical Models" [7]. It turns out that, when using partial derivatives and simultaneous equations, there is a simple formula in which b_1 is a weighted average (see Table 1, column 6) of the Y (man-hours).

The variables and equations are:

- $X(i)$ are the lot size values, where I = 1-10.
- $XAVE$ is the average value for X, which is 50.
- $Y(i)$ are the actual man-hours, where I = 1-10.
- $YAVE$ is the average for Y, which is 110.90.
- $K(i)$ are the weights.
- b_1 is the value of β_1 that is the best match.
- b_0 is the value of β_0 that is the best match.

$$K(i) = \frac{X(i)=XAVE}{\prod_{i=1}^{n}(X(i)-XAVE)^2}.$$

$$b_1 = \sum_{i=1}^{n} k(i) \times Y(i).$$

$$b_0 = YAVE - b_1 \times XAVE.$$

Columns 4-6 in Table 1 show the calculations. With this approach, finding b_0 and b_1 (best values for β_0 and β_1) looks more like an algorithm than learning. The "best" β_0 and β_1 values are 12.2235 and 1.9735, respectively.

What if there are multiple X variables? For example, we add weight and material type to lot size. This is multiple linear regression, and there are standard "matrix or linear algebra" formulas [7] to find the best coefficients. Additionally, linear least squares primitives have been available in programming environments [8] since the early 1970s.

To add to the "mystery":

1. If β_0 and β_1 have the respective values of 0 and 2.185, the fit is almost as good. A slight change in the slope (β_1) enables a large change in the y intercept (β_0).
2. The values b_0 and b_1 can be found using the optimization method linear programming that supports a diversity of constraints [9].

Ongoing Challenge

Is linear regression "learning," statistical modeling or optimization with linear programming? Does it matter for a company which one it is? Let's discuss learning and how AI can help drive smarter decisions to improve organizational performance – the ongoing challenge.

id	x, Lot size	yact, Man-hours	X(i) - XAVE	(X(i)-XAVE)^2	K(i) - weights
001	30	73	-20.00	400.00	-0.0059
002	20	59	-30.00	900.00	-0.0088
003	60	128	10.00	100.00	0.0029
004	80	179	30.00	900.00	0.0088
005	40	87	-10.00	100.00	-0.0029
006	50	108	0.00	0.00	0.0000
007	60	135	10.00	100.00	0.0029
008	30	69	-20.00	400.00	-0.0059
009	70	148	20.00	400.00	0.0059
010	60	123	10.00	100.00	0.0029
ave	50.00	110.90			
sum				3400.00	0.0000
b1	1.9735				
b0	12.2235				

TABLE 1: Lot size vs. man-hours for linear regression demonstrating weighted average solution from the "before times."

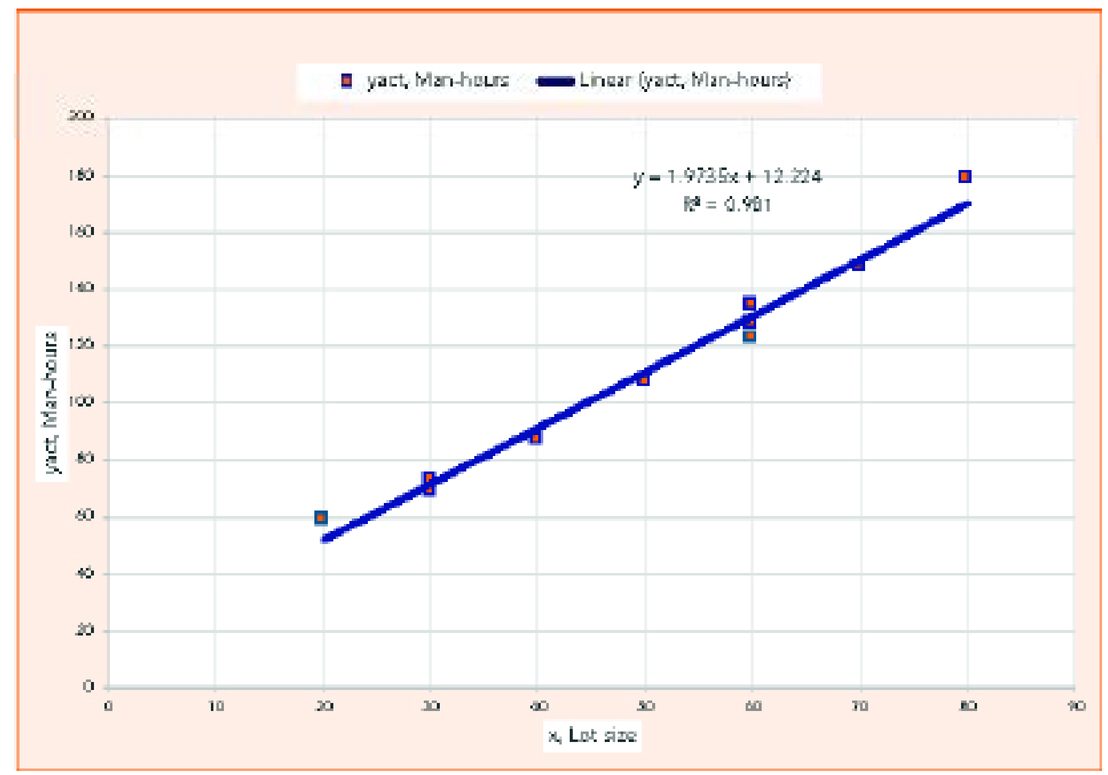

FIGURE 1: Lot size vs. man-hours scatterplot with best fit straight line.

KEN FORDYCE (kfordyce@arkieva.com) is a consultant on analytics without borders at Arkieva. He joined Arkieva in 2013 after completing a 36-year career with IBM covering a wide range of analytics applications from supply chain to colon cancer. Ken was lucky enough to be part of teams that altered the landscape of best practices in central planning, dispatch scheduling, artificial intelligence, applied optimization and data science. These teams received major awards from IBM, IBM customers, INFORMS, AAAI and POMS.

Note: *References can be found online: https://doi.org/10.1287/orms.2026.01.09*

MAKING SENSE OF AMBIGUITY IN APPLIED ANALYTICS

BY ANEESH SAJAN

WE OFTEN ASSUME THAT APPLIED analytics begins with clear questions, reliable metrics and well-defined data. In practice, however, many valuable analytical projects start in ambiguity. I see teams operate in situations in which the problem is loosely defined, the data reflects past system behavior rather than objective truth and success metrics are unclear or openly debated. Not all hard problems are ambiguous. Many are complex but clear: We know the goal, even if reaching it is difficult. Ambiguity exists when that clarity is missing. We may not know what should be measured, which signals to trust or what a good outcome looks like. Applied analytics often begins here: turning vague questions into clearer ones before any serious analysis occurs.

To illustrate, consider a hypothetical social media platform trying to decide whether it has "enough" influencers for a particular domain, such as home fitness. At first glance, the problem looks simple – count the number of creators in the domain, compare supply against audience demand and decide whether to recruit more influencers. In practice, ambiguity appears immediately. What qualifies as an influencer? How do we define the domain? How do we infer demand? What does "enough" actually mean? The challenge here is not computational but conceptual. Before running any analysis, we should first decide what question truly needs answering.

Ambiguity in the Problem-Definition Space

In ambiguous settings, the first step is to state the problem as it is currently understood, with the acknowledgment that this framing may be incomplete or incorrect. Stakeholders typically describe symptoms rather than root causes. In some cases, no underlying problem exists. What appears to be an issue may instead require clearer definition, validation or ongoing monitoring rather than immediate intervention.

In the influencer example, the observed symptom might be that users repeatedly see the same influencer or that engagement growth slows in certain interest domains. From this single observation, several plausible problem statements can emerge. One interpretation is that the platform lacks enough influencers in the domain. Another possibility is a discovery or ranking issue, where the recommendation system exhibits popularity bias, overamplifying a small group of established creators and reinforcing a rich-get-richer dynamic due to exploit-heavy optimization and limited exploration. A third possibility is that the influencer pool itself is skewed. If most creators share similar demographics, formats or production styles, the platform may appear saturated while still underserving meaningful segments of user interest. These explanations are not mutually exclusive. The apparent imbalance may reflect a mix of influencer availability, algorithmic bias, creator incentives, user-interest inference and measurement artifacts, rather than a single root cause.

When the problem itself is unclear, applied analytics should resist locking into a single explanation too early. Instead of treating the first problem statement as correct, it helps to pause and examine it. This means separating what is being observed from what might be causing it, being clear about assumptions and considering more than one possible explanation. At this stage, the goal is not to optimize or fix anything but to build understanding. Applied analytics under ambiguity becomes an iterative process of refining the problem before deciding how to solve it.

Ambiguity in the Solution Space

As analysts, we rarely arrive at a solution in one step. Each stage involves choices about methods, assumptions and how different pieces fit together. Even after a solution is built, it often needs testing, comparison and refinement, along with discussion with business partners. In ambiguous settings, solving a problem is not a single decision but a series of connected choices that gradually shape the final outcome.

In the influencer example, the solution may look simple at first: Define the topic, estimate audience interest and count how many relevant influencers exist. In practice, each step involves judgment. Changing how the domain is defined changes which influencers are included. Interest estimates often reflect what users were shown in the past, not what they truly want. Measures of supply depend on whether influence is defined by activity, reach or engagement. Additional factors such as recommendation bias, creator incentives and uneven visibility also matter. Over time, seasonality, language, geography and behavior changes further affect the results. For these reasons, the solution develops through iteration rather than appearing fully formed at once.

In practice, this works best when the solution is broken into smaller hypotheses, and each part is tested step by step. Iterative data exploration helps build intuition along the way. Regular check-ins and proactive updates with stakeholders help keep the work aligned with the real objective. Even then, the final solution may have limitations. Good applied analytics acknowledges these limits, explains the trade-offs and shows why the chosen approach makes sense given the constraints. Conclusions should be grounded in evidence and communicated clearly.

Recommendations based on my experience working through ambiguity in applied analytics:

1. **Listing hypotheses and building intuition:** When a problem is unclear, start by writing down the different hypotheses, which helps separate what we observe from what we assume. Early analysis stays simple; basic summaries, rough comparisons with other business reports and quick visual checks are often enough to build intuition. At this stage, focus less on accuracy and more on understanding how the system behaves and where signals may be misleading.
2. **Designing for iteration and guarding against blind spots:** Ambiguous problems rarely have answers in a single pass. As understanding improves, definitions, metrics and methods often need to be refined. At the same time, repeated iteration can create blind spots if we keep looking at the data in the same way. To avoid this, we should deliberately explore alternative views, test different definitions and ask what might be missing.
3. **Engaging in peer review, collective reasoning and trust building:** Ambiguity makes results easier to question. Sharing work early with other analysts, scientists or domain experts helps surface hidden assumptions and strengthen reasoning. This process also helps build trust. When stakeholders understand how conclusions were reached and what limitations remain, they are more likely to place trust in the results – even when the answers are not definitive.
4. **Avoiding analysis paralysis through deliberate stopping rules:** Ambiguity can lead to endless exploration. Although some uncertainty is unavoidable, applied analytics must still support decisions. Ask whether additional analysis would meaningfully change the decision. If not, it may be time to stop.
5. **Writing executive summaries with actionable recommendations:** Ambiguous analysis only creates value if decision-makers know what to do with it. A short executive summary should be provided that clearly states what was learned, what remains uncertain and what actions are recommended. The recommendations should be concrete and feasible, not just analytical observations. Even with limited confidence, proposing thoughtful next steps is better than leaving stakeholders unsure how to proceed.
6. **Treating execution as part of analytics:** After decisions are made, track outcomes and feed them back into future analysis. Over time, this feedback loop reduces ambiguity and improves judgment. In practice, this works best when analysts partner closely with product teams and other stakeholders to execute recommendations, observe real-world effects and refine understanding together.

Ambiguity is a normal part of real-world problems. The most useful applied analytics work often starts before any modeling, by being clear about what problem we are trying to solve and what we do not yet know. Many challenges are not about complex computation but about understanding the situation and helping people make better decisions. In this sense, applied analytics is the art of solving business problems using science.

ANEESH SAJAN is a lead data scientist with 15+ years of experience in applied analytics, experimentation and machine learning, including 10+ years at Amazon. He builds scalable data science solutions and translates complex, large-scale data into practical, decision-ready insights.

AI THAT ACTUALLY SHIPS

A product operating model for moving from prototype to production

BY RAM KUMAR NIMMAKAYALA

MOST ARTIFICIAL INTELLIGENCE (AI) FAILS in translation – not from model to metrics but from demo to daily use. Here's a practical playbook that leaders can use to move AI from prototype to pilot to production with clear gates, simple metrics and accountable ownership.

The Uncomfortable Truth

The most common AI artifact in large organizations is a demo. Teams prove something is possible, executives applaud and then … nothing changes on Monday.

The pressure has never been higher. Generative AI hype has put AI on every board agenda; shadow tools are proliferating faster than governance can track; and executives are demanding return on investment from experiments. Without a deliberate operating model, most organizations will continue producing impressive demos that never touch a real workflow.

Why do so many efforts stall?

- There is no explicit business pain point or owner.
- "Success" is not defined in measurable terms.
- Data is messy or restricted.
- Compliance and security arrive at the end.
- No human is accountable after launch.

The cure isn't a bigger model. It's a better operating model.

A Three-Stage Path with Hard Exits

Every AI initiative should live in one of three states: **Prototype → Pilot → Production** – with nothing in between. Each stage has a purpose and explicit exit criteria (see Figure 1).

1. **Prototype (discovery):** Build the smallest working version to prove a signal. Define what "good" looks like before building. If the signal is weak, kill it fast.
2. **Pilot (controlled exposure):** Put the prototype in a narrow, real workflow with real users. Measure real usage and surface real risk. Integrate just enough to learn.
3. **Production (scaled ops):** Harden data access, privacy and lineage. Name an owner. Rehearse rollback. Monitor cost, drift and behavior.

Stage language creates shared expectations across product, engineering, compliance and operations. Everyone knows what "good enough to graduate" means.

The AI Product Manager's Real Job

An excellent AI product manager (PM) is not "the prompt person." The AI PM is the system architect for value, feasibility, trust, adoption and safety:

- **Value – Why build this?**
  Name the painful workflow and the metric to move (time saved, risk reduced, revenue protected).
- **Feasibility – Can we ship it here?**
  Do we have the data and access? Choose the correct pattern for the problem: prompting, retrieval-augmented generation (RAG), fine-tuning or agent with tools – not "just use GPT."
- **Trust – Will people believe it?**
  Fit the user experience (UX) to the job (inline assist, sidebar copilot, background automation). Constrain answers to an approved, source-of-truth context.
- **Adoption – Will it stick?**
  Make the first value less than 30 seconds. Start with one hungry team; expand by evidence, not enthusiasm.
- **Safety/quality – Can we defend it?**
  Run live evaluations for hallucinations, bias and policy violations. Provide a human escalation path when confidence is low. Adhere to this rule: no evaluations, no launches.

Mini-example: Instead of "AI tutor," ship "advisor copilot that answers the 40 most repetitive questions from approved policy text, auto-links sources and escalates everything else to a human."

The Promotion Gate: Progress Must Be Earned

Moving from **prototype to pilot** or **pilot to production** requires passing a simple, auditable gate:

- Quantified value with real users (not opinions)
- Usage evidence (engagement, task completion, deflection)
- An approved data path (privacy, security, access control)
- A named owner post-launch
- Monitoring and rollback defined and rehearsed
- Policy and risk review cleared; unmitigated high-risk use cases stop here

Graduation isn't a vibes-based decision – it's a checklist you can audit.

What "In Production" Actually Means

Deployment is not successful; *stability* is. Production AI is continuously observed across three layers:

1. **Technical health:** uptime, latency, error rate
2. **Behavioral quality:** accuracy/relevance, hallucination, human-escalation rates, safety/bias boundaries
3. **Business value:** cost to run versus outcome delivered

If health, quality and value aren't on dashboards with named owners, then it isn't production – it's an unmanaged experiment wearing production clothes (see Figure 2).

Five Levers to Scale (It's a System, Not a Hero)

Stop hunting for the unicorn engineer and build the machine:

1. **People:** AI PM, data engineering, machine learning (ML) engineering, platform/infra, risk/compliance and operations to clear lanes and escalation paths
2. **Process:** Scrappy in discovery, disciplined in delivery (change control, approvals, lineage)
3. **Data:** Source of truth, retention and access governed; always answers "where did this answer come from?"
4. **Technology:** Infrastructure that supports iteration and scale; avoids fragile one-offs
5. **Production ops:** Monitoring, alerting, retraining cadence and incident playbooks – the launch is the beginning

You don't productionize a model; you productionize the system around it (see Figure 3).

Three Pragmatic Use Cases (with Simple Numbers)

These patterns ship reliably across education, customer service and back-office operations. They're specific enough to copy and general enough to adapt.

1. **Support the copilot with human escalation:**
   - **Pain:** Frontline teams drowning in repeat questions
   - **Pattern:** A policy-grounded copilot answers first-line questions; uncertain cases escalate with a transcript and suggested following actions
   - **Metrics:** First-contact resolution, median handle time, user satisfaction and percentage escalated
   - **Target:** A 20%-40% deflection with <5% unsafe responses (auto-flagged), human-in-the-loop for the rest

   One-liner: Narrow domain + approved content + clear escalation path = durable wins.

2. **Look for early-warning retention signals:**
   - **Pain:** Outreach too late
   - **Pattern:** A supervised model flags at-risk users, the copilot drafts personalized outreach from customer relationship management notes and activity, and staff approve/send

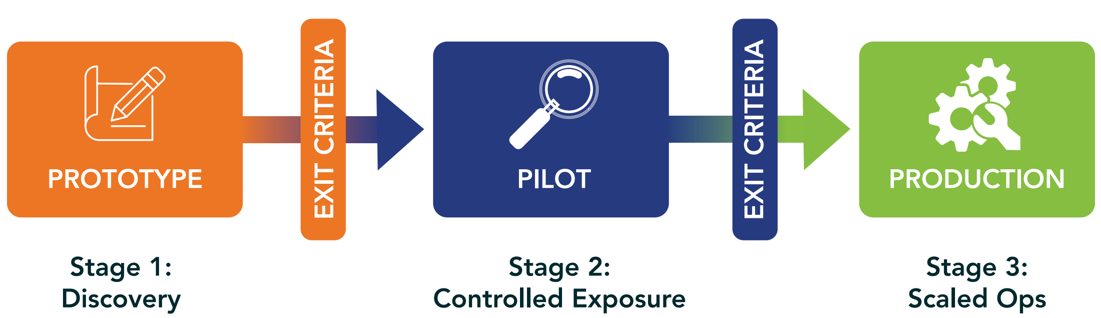

FIGURE 1: Every AI initiative lives in one of three states with explicit exit criteria between each. No initiative graduates without evidence.

FIGURE 2: Production AI is monitored across three layers. If all three aren't on dashboards with named owners, it's not production; it's an unmanaged experiment.

FIGURE 3: Scaling AI requires five interlocking levers. You don't productionize a model; you productionize the system around it.

- **Metrics:** Precision/recall on risk, time-to-outreach and conversion to "stayed engaged"
- **Target:** A 10%-15% lift in proactive outreach within 48 hours of the risk signal; staff keep authorship and accountability

One-liner: Predict, then assist the human – don't auto-decide.

3. **Document intake and routing with an audit trail:**
   - **Pain:** Manual extraction, validation and routing slowing the back office
   - **Pattern:** AI extracts fields, checks policy, classifies and routes with confidence scores; low-confidence and exceptions go to humans
   - **Metrics:** Cycle time, percentage auto-approved, exception rate by reason and audit completeness
   - **Target:** A 30%-50% cycle-time reduction while lifting audit completeness to ~100% via structured logs

One-liner: Automate the boring; preserve the judgment.

These patterns work but only if you avoid the organizational pitfalls that quietly erode AI investments over time.

Anti-patterns That Quietly Create AI Debt

- **Shadow AI:** This includes unsanctioned tools with sensitive data. Fix with an internal AI register and clear "what not to paste" guidance.
- **One-person stacks:** Only one engineer can deploy or fix it. Fix with shared infra, handoffs and documentation.
- **"Metrics later":** You'll never cleanly retrofit measurement. Establish success metrics at the outset.
- **Compliance at the end:** This is a rework tax. Fix by embedding legal/risk at the stage definition.

AI debt is primarily organizational, rather than technical, incurred when ownership, policy and measurement are afterthoughts.

Responsible AI, Translated for Operators

Responsible AI shouldn't be a philosophy; it should be controls that anyone can run:

- **Owner:** Every AI system has a named human responsible for performance, compliance and escalation from day 1.
- **Risk control:** Impact is reviewed before rollout; high-risk uses trigger deeper review (ethics, legal, compliance). Human override and a tested shutdown path exist.
- **Transparency:** An operator can document what the AI does, the data it uses and whom it affects; disclose AI involvement in consequential decisions; and maintain an internal register of systems.

"No policy, no production" isn't red tape – it's how you earn the right to scale.

Here's an executive playbook you can steal for Monday:

1. **Start with a painful workflow, not a shiny model.** If no one is begging for help, don't build it.
2. **Define success before you sprint.** Agree on the single number that proves impact.
3. **Stand up the cross-functional team early.** Embed the product, data/ML, platform, compliance and ops – don't "throw it over the wall."
4. **Enforce the promotion gate.** Do not advance without value evidence, policy sign-off, monitoring and a named owner.
5. **Treat post-launch as stewardship.** Implement dashboards, evaluations, retraining, rollback and cost control. Make "still valuable six months later" the brag.

Closing Thought

Enterprises don't need more prototypes; they need more production-grade AI that teams trust and leaders will fund again. When you make the operating model explicit through three stages (promotion gates, three-layer monitoring and a policy with a named owner), the conversation changes. AI stops being theater. It becomes infrastructure.

And when it does, something actually changes on Monday.

RAM KUMAR NIMMAKAYALA is Principal Product Manager – AI & Data Strategy at Western Governors University, one of the largest online universities in the U.S. He leads AI products that support more than 190,000 learners and is frequently invited speaker and panelist on pragmatic, responsible AI for large enterprises.

INSIDE THE NUMBERS GAME: ANALYTICS AND THE NEXT ERA OF BASKETBALL

BY TRISTAN ROBBINS, LANDON MCINTOSH, WILL ROBINSON AND SCOTT NESTLER, CAP-X

The Lineage of Analytics in the NBA

For much of its history, professional basketball has been defined by a blend of strategy and athleticism [1]. As player skills have continued to improve, so has the decision-making that shapes the sport. For decades, coaches and scouts leaned almost entirely on the "eye test" and basic statistics such as points, rebounds and assists to evaluate player performance. In the mid-1970s, the box score was expanded, introducing statistics such as offensive rebounds, turnovers, blocks and steals. Combined with gut instinct and film study, decisions hinged more on traits like athleticism and leadership than on data alone. These traditional methods shaped roster moves, starting rotations and playing styles in the NBA for decades.

In the late 1990s, however, the NBA began a dramatic transformation with the introduction of analytics. Play-by-play data was introduced in 1997, ushering in the era of real statistical analysis. Through the early 2000s, statistical analysis steadily gained traction, shaping both on- and off-court decisions – a shift accelerated by Moneyball – and by the late 2000s, the use of play-by-play data helped bolster a leaguewide analytical surge.

In 2013, the NBA became the first U.S. professional sports league to implement player tracking for every game through its partnership with STATS SportVU [2]. Using high-resolution cameras in every arena and advanced computer vision technology, SportVU provided real-time tracking data based on speed, distance, player separation and ball movement, delivering the most detailed player and team analysis the league had ever seen.

The rise of advanced analytics through player-tracking technology led to developments in spatial analysis and new advanced metrics, used to assist coaches and organizations to make informed decisions that were previously based on instinct. This shift from counting statistics to tracking analytics has allowed the game of basketball to grow, improving player evaluation techniques as well as the quality of the game for fans and viewers. Statistical learning models heavily influence today's NBA decision-making and are responsible for the modern revolution of emphasis on three-point shooting, offensive rebounding and load management.

Recent Technological Innovations

Hawk-Eye Optical Tracking

After the early SportVU era of tracking basketball, Second Spectrum was the NBA's optical tracking provider beginning in 2016. At the same time, Hawk-Eye technology was being widely used in tennis and soccer, making line calls in grand slams and providing referees with quick offsides decisions during the 2022 World Cup [3]. In 2023, the NBA adopted Hawk-Eye, replacing Second Spectrum as the lead provider of optical tracking technology. Genius' Second Spectrum still works with the NBA but in a more entertainment-focused way through League Pass [4]. These alternative telecasts feature advanced team and player statistical insights integrated directly into the stream, giving analytics-focused fans an ideal viewing experience.

At the time, Second Spectrum had a "dot" tracking system using six cameras to track player torsos in 2D, which created a single point for each player at any given moment [5]. Hawk-Eye's implementation of the "pose" tracking system uses 14 cameras that track 29 points on players' bodies (arms, legs, hands, etc.) plus the ball in full 3D in near-real time and supports more accurate output data than Second Spectrum. As a result, the NBA announced that the leaguewide implementation of Hawk-Eye would begin in the 2023-2024 season.

Hawk-Eye's main impact at the NBA level is twofold. First, the impact on the game is that Hawk-Eye's technology supports better and faster officiating and reviews because it increases both the accuracy and the speed of calls. In Game 4 of the Nuggets vs. Clippers series in 2025, Hawk-Eye's review showed that Gordon's dunk went in just about a millisecond before the buzzer [6].

But perhaps Hawk-Eye's biggest impact is at the team level [5]. Synergy Sports, which has long been one of the most widely used analytics platforms in basketball, is now owned by Sportradar, one of Hawk-Eye's partners, and has begun combining its play-tracking tools with Hawk-Eye's pose data. Nearly every basketball metric imaginable will be improved, and many new metrics will be created using this tracking technology. Over the past several years, NBA teams have consistently expanded their analytics departments, with advantages going to early adopters of new tracking tools and those able to integrate and utilize them most effectively (see Figure 1). In turn, the demand for analytically skilled basketball professionals has reached new heights.

Training and Player Development

Player development in professional basketball has been revolutionized by blending innovation and technology with forward-thinking applications of analytics. Modern coaching staffs rely on tracking data, advanced visualization analysis and film study to develop players. Rather than replacing coaching instincts with analytics alone, new technologies are aimed at augmenting traditional coaching methods, providing teams with real-time performance improvement feedback.

Modern player development is dependent on the interpretation of data tracking by team analysts. By applying principal component analysis (PCA) to years of NBA player tracking and box score data, Abraham Montalvo demonstrates how players can be grouped into style clusters [8]. Some players are designated as shooters, defenders or ball handlers, which allows coaches to look beyond broader designations such as "guard" and "forward." In turn, coaches can emphasize the specific traits necessary to be successful relative to others in their cluster. Newer tracking data focuses on player performance during real game scenarios; PCA can provide teams with complete feedback on areas of strength and weakness, allowing coaches to develop tailored training regimes for individual players. Drills based on a player's statistical weaknesses represent an advanced analytics evolution from the antiquated, non-data-driven methods of the past.

Recently released research that simulated more than 11,000 NBA games highlights how player versatility can be quantified according to lineup-based performance levels, illustrating how different players alter their playing style based on teammate combinations and situational context [9]. Through this analysis, coaches can anticipate how differently players will perform based on specific lineup combinations.

Injury Prevention

Preventing injuries is critical for NBA teams hoping to compete in today's data-driven league. Injuries have been on the rise in recent years, raising the stakes for keeping key players available [10]. Beyond the clear human and financial costs, each game that a starter misses could harm a team's record, ultimately affecting seeding and playoff outcomes. Front offices and coaches now view injury prevention as a critical part of gaining a competitive edge on opponents, focusing on specific and measured rest, with unique regimens for each player [11].

Coaches regularly engage in load management to cycle players in and out of games, aiming to maximize productive output while minimizing the potential for injury. Although load management was originally centered around off days, it is now incredibly targeted. Teams monitor how each player's load accumulates in games and throughout the season, carefully adjusting practice intensity and other training aspects to avoid injury with the goal of preserving player ability without sacrificing performance or health. Deciding which players to rest for each game is a decision that artificial intelligence (AI) has had increased control over. Some machine learning algorithms have erred on the side of caution, suggesting to rest multiple key players in individual games more frequently [12]. This practice, although beneficial for teams,

FIGURE 1: NBA analyst headcounts by season, 2009-2025 [7].

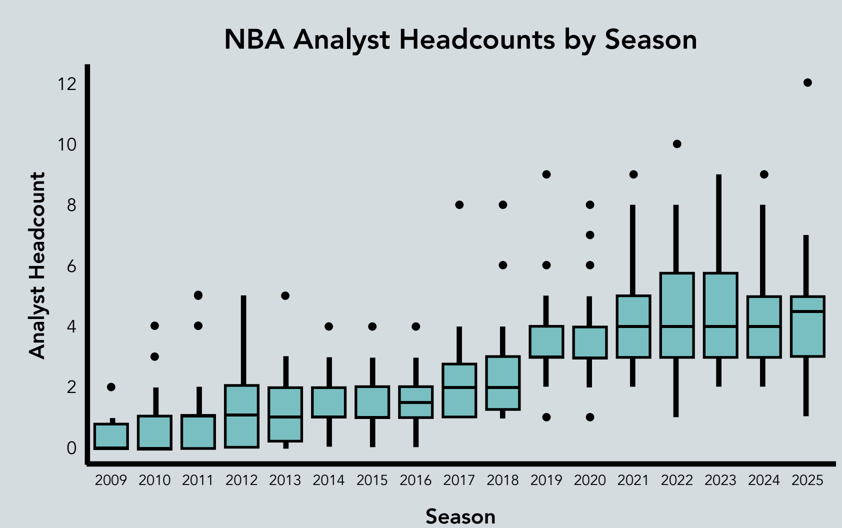

As the sport continues to evolve, teams will find new ways to gain a competitive edge, and improvements in analytics represent a prime opportunity.

frustrated many fans and the league, which issued a new rule to limit how many stars could be rested in marquee games. The goal of the increased rest is to ensure players are available and refreshed for key matchups that can affect playoff probabilities and ultimately maximize each team's chances of winning the championship.

The league and individual teams have explored a variety of technologies to improve today's load management, including wearables, optical tracking and sleep policies. Teams now use wearables to monitor player vitals and health signals during practices, flagging signs of fatigue for the coaching staff. Optical tracking of camera feeds in games translates individual player movements into categorized metrics such as sprints and jump landings. These specific profiles of a player's in-game exposure to fatigue also produce flags for dangerous mechanics and abnormally high workloads on specific joints and the overall body [13]. Player rest is also critical for avoiding injuries; studies have shown that napping on game days improves play performance [14]. Many teams have established sleep minimums to ensure that their athletes are well rested each day of the season, with an emphasis on game days [15]. The abundance of data for each player enables teams to make informed decisions about load management, aiming to minimize injuries while keeping a full-strength team on the court.

Effect of Analytics

In-game Strategy

The most direct application of analytics is on the court, where coaches constantly make decisions. A head coach faces about 33 tactical decisions, on average, in a single game through substitutions, timeouts and challenges (see Figure 2, which does not include decisions that were not recorded in play-by-play data such as post-timeout plays, defensive strategy, starting lineups and many others). Over the course of a season, these choices alone add up to nearly 3,000 decisions that can be shaped by analytics. Models that quantify win probability or expected points added because of an action provide a framework for these decisions that need to be made quickly. Tracking data enhances lineup analysis, showing which combinations of players defend or score most efficiently. Automated decisions found through prescriptive analysis can act as a guide for consistent strategy for coaches.

One recent example of analytics driving in-game decision-making is the lineup rotation involving Isaiah Hartenstein of the Oklahoma City Thunder during the 2025 NBA Finals. Head coach Mark Daigneault alternated Hartenstein between the starting group and the bench, based on lineup tracking data that measures which units possess the greatest efficiency. With Chet Holmgren in a double-big frontcourt, the Thunder had one of the best point differentials of their season, with a +7.0% offensive rebound advantage over nearly 500 minutes of regular season and playoff action [16]. They also saw a defensive rating of 109.9 points per 100 possessions on 679 possessions [17]. Analytics revealed this tall duo held back opponent field-goal percentage at the rim and extended the Thunder's second-chance scoring on offense. Conversely, removing Hartenstein in earlier games to insert guard Cason Wallace allowed the Thunder to play with smaller, faster lineups that tracking data showed were better at defending the perimeter and forcing pace on the Indiana Pacers' guard-oriented rotations [18]. By rotating through these different data-driven looks each game, the Thunder optimized their contests, demonstrating how lineup efficiency metrics and monitoring systems have become central to rotation and substitution planning.

FIGURE 2: Average substitutions, timeouts and challenges made by a coach each game.

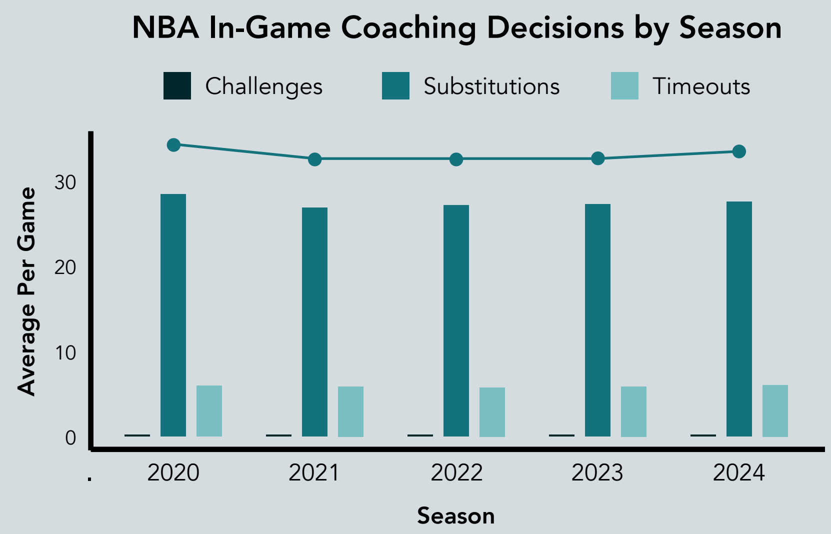

Front Office and Long-term Decision-making

Off the court, analytics shape the choices that define team trajectories for years. General managers (GMs) and front offices face fewer decisions, but each often comes with more long-term impact. Across an entire season, a GM may only make about 50 key personnel decisions drafting players, signing free agents, negotiating extensions and executing trades (see Figure 3). Unlike the coaching staff, who manage countless choices every game, including substitutions, timeouts, challenges and other decisions, the GM's fewer choices have a longer time horizon. Draft evaluations now incorporate computer visualization tracking data from college broadcast feeds, and player acquisition through trades focuses on which players will provide the best fit to a team's structure. Similarly, analytics helps front offices evaluate contract negotiations when teams are trying to leverage additional value beyond a player's market value.

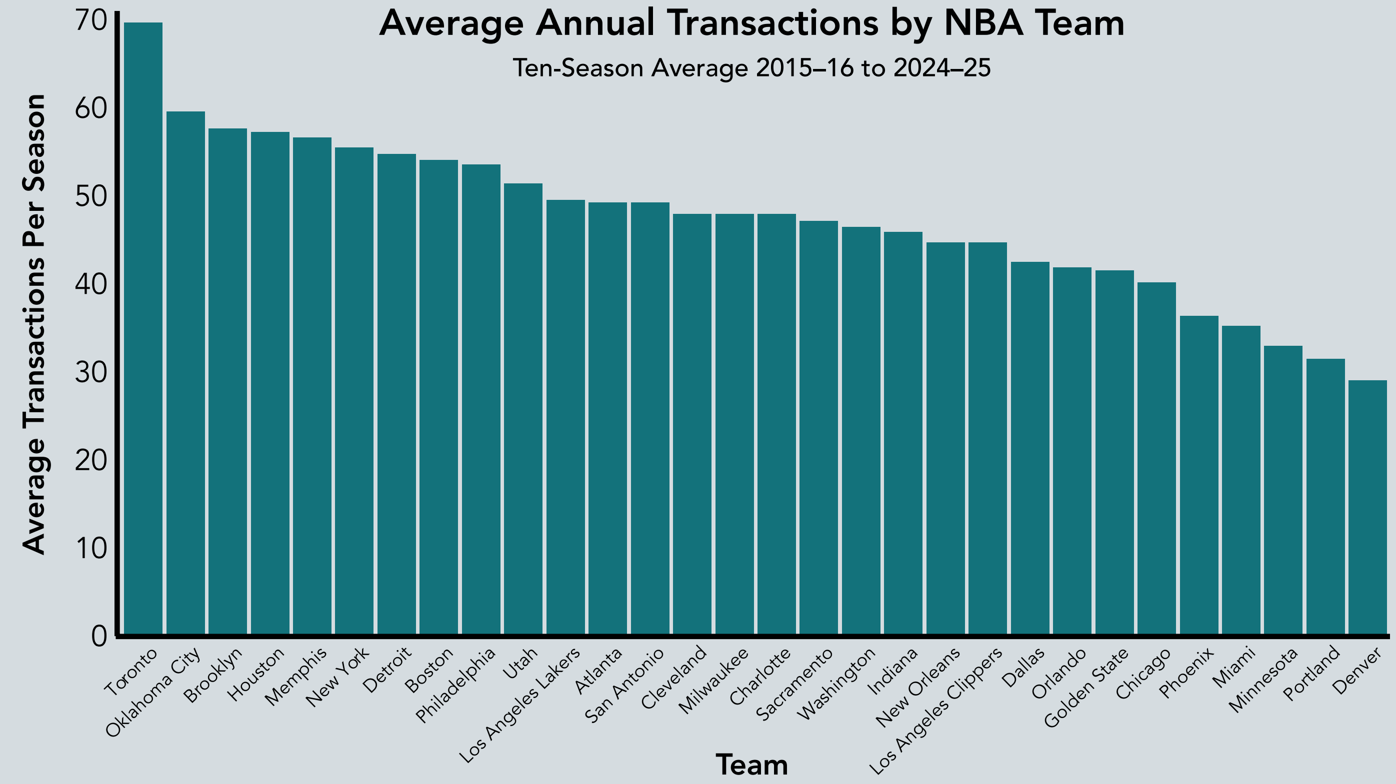

FIGURE 3: Average transactions made by a front office over the last 10 years. It is important to note that these also include G League transactions, Exhibit 10 contracts, two-way contracts and waiving players. [19]

A recent example of analytics in front-office decision-making is when the Denver Nuggets signed Jamal Murray to a four-year, $208 million maximum extension during the 2024-2025 offseason [20]. Although Murray's strong statline (21.2 points, 6.5 assists and a 42.5% 3-point average) provides a reasonable baseline for his value, Denver's front office supplemented traditional statistics with more advanced figures derived from computer visualization tracking data [21]. Using Second Spectrum data on two-man actions and assist pairings, the Nuggets analyzed Murray's production during possessions with superstar Nikola Joki, demonstrating their increased lineup efficiency when they share the court [22]. League reporting this season emphasized the Nuggets' reliance on these metrics to project fit, highlighting how consistent scorers stabilize team offenses [23]. To increase production, high-impact scorers like Murray provide long-term lineup stability. Murray is a top 40 player when looking at volatility analysis, a measure of how consistent a player's offensive production is across games. When an organization commits to a future-defining transaction, it leans on advanced analytics, forecasting a desired player's fit and production well into the future.

The Future of Analytics

The next era of basketball analytics is emerging at the college level. Although the NBA has had more than a decade of access to advanced tracking data, most collegiate programs have historically lacked this data and in-depth analytics. However, that landscape is beginning to shift. More affordable tracking providers and an increase of students and professionals interested in this field have made advanced analytics increasingly more popular with NCAA men's basketball programs. Among the elite of Division I basketball, some schools have already integrated tracking data through providers like Sportradar, SkillCorner and ShotTracker [24].

Whether through broadcast feed or in-arena cameras, the recent rise of tracking data in college basketball allows coaches to mirror projects and approaches that the NBA has been using for years.

As the sport continues to evolve, teams will find new ways to gain a competitive edge, and improvements in analytics represent a prime opportunity to find value in players and ideas that other teams pass on. By incorporating analytics into all levels of decision-making, NBA teams are making smarter, quicker and more efficient decisions, enabling them to compete at a higher level year after year. In a seemingly endless sea of data, the teams that can separate signal from noise will find themselves winning more games and becoming champions as the sport matures.

TRISTAN ROBBINS is an undergraduate student studying statistics, economics and sport management at the University of Florida (UF).

LANDON MCINTOSH is an undergraduate student studying computer science and economics at UF.

WILL ROBINSON is an undergraduate student studying statistics at UF.

SCOTT NESTLER, Ph.D., CAP-X, Stat, is professor of practice in sport management at UF, and director of the UF Sports Analytics Lab.

Note: *References can be found online: https://doi.org/10.1287/orms.2026.01.04*

OPTIMIZATION, THE QUANTUM APPROACH

BY SARA LODHA

If you think you understand quantum mechanics, you don't understand quantum mechanics.
– Richard Feynman, Nobel Prize-winning physicist

AS CONFUSING AS THIS SOUNDS, THE strange rules of quantum physics are exactly what make quantum computing, and quantum optimization, so powerful. They turn the weirdness of nature into a new way to solve problems faster than ever before. Let's break it down from the beginning.

Defining Quantum (Science and Computing)

The word quantum literally means "how much." Quantum physics studies how energy and matter behave at the smallest possible scale, where the normal rules of nature no longer apply. At that level, everything is made up of incredibly small particles, photons (light), electrons and quarks that behave in ways that seem random and unpredictable to us. These particles can act like both waves and solid objects, exist in several states at once and even instantly affect one another from a distance. These strange principles – superposition, entanglement and uncertainty – have opened the door to new technologies, including quantum computing, which uses these same natural behaviors to process information in ways classical computers never could [1].

Quantum computing takes the principles of quantum physics and turns them into a new way of processing information. A classical computer uses bits (0s and 1s) to represent data. Every operation, whether adding numbers, comparing routes or finding the best option, happens one step at a time, following a clear path through all possible answers. A quantum computer, on the other hand, doesn't follow a single line of reasoning. It uses qubits, which can exist in a superposition of 0 and 1, meaning each qubit can represent multiple possibilities at once. When many qubits are connected through entanglement, they start to influence each other, creating a web of possibilities that the system can explore simultaneously.

Imagine you're at Disneyland, trying to find the fastest route through all the rides. A classical computer would test one path, time it, then try the next and so on, checking every possible route until it finds the best one. A quantum computer, by contrast, could consider all possible routes at the same time, using superposition and entanglement to compare outcomes in parallel.

That said, quantum computing isn't about magically finding all the answers at once; it's about guiding probabilities toward the most likely solution. While the qubits are in their superposed, in-between states, they represent many potential answers, each with a probability of being correct. Quantum algorithms are designed to nudge those probabilities so that when the system finally "collapses" into a single answer (when we measure it), the most probable outcome is also the best, most optimal one [2].

Quantum in Everyday Lives

The idea of *quantum* physics has been around for decades. What began as abstract equations is now beginning to shape the world around us. For example, GPS technology allows us to coordinate navigation, logistics and delivery routes across the globe down to the second.

Across industries, researchers are already testing what happens when quantum thinking meets real-world challenges. For example, in finance, quantum computing is used to analyze risk or optimize large portfolios; in machine learning, it's used to help algorithms uncover patterns faster in massive datasets. But one of the most promising, and relatable, areas of all is optimization. Optimization problems are everywhere: finding the best route for a delivery truck, scheduling flights, balancing inventory or designing energy-efficient supply chains. Meeting the kinds of challenges that grow exponentially more complex as the number of variables increases is where quantum computing shows its greatest promise: not by replacing classical systems but by working alongside them to explore the vast numbers of possibilities in parallel and guide us more quickly toward the best answers [3, 4].

Defining Optimization

Optimization is about finding the best option among many, while balancing constraints such as time, cost or resources. It's what we do every day without realizing it, from planning the fastest route to work to deciding the most efficient way to run errands on a busy afternoon. At its heart, quantum optimization doesn't just compute faster – it thinks differently. Rather than processing one option after another, it reframes the entire challenge as a physical system that naturally seeks balance.

Optimization Using Quantum Computers

In quantum optimization, the aim is to find the best solution from among many possible options. This could be the most efficient delivery routes, the optimal production schedule or the most balanced resource allocation. To do this on a quantum computer, the problem is translated into a form the hardware can use: Each possible configuration of variables becomes a quantum state, and each state is assigned an energy value that corresponds to how good (or bad) that configuration is.

The goal is then to find the configuration with the lowest energy – that is, the best solution under the constraints of cost, time, resources, etc. The quantum system evolves so that lower-energy states become more probable. When we measure the system, we get a candidate solution. Because quantum results are inherently probabilistic, we repeat this process, and frequently, a classical computer then refines the most promising results. This hybrid quantum-classical loop is what allows quantum optimization to be practical right now even though the hardware is still maturing.

The Algorithms Driving Quantum Optimization

Now that we understand quantum optimization as a way of reframing problems as energy landscapes, the next step is to look at how those landscapes are actually explored. Quantum annealing is a method designed to drive a system toward its lowest-energy state by mapping the problem onto a physical system that naturally seeks low energy. According to the vendor D-Wave Systems, quantum annealing "uses quantum physics to find low-energy states of a problem and therefore the optimal or near-optimal combination of elements" [5]. In practice, this method is especially suited for combinatorial optimization, characterized by problems where many discrete decisions interact (e.g., route choices, scheduling orders or binary allocations). The quantum annealer makes use of quantum fluctuations or tunneling to escape suboptimal local minima and aim for the true global minimum.

The Quantum Approximate Optimization Algorithm (QAOA) is a hybrid quantum-classical algorithm that approaches optimization by alternating between quantum operations and classical parameter tuning. The quantum computer encodes the cost function (the measure of how good each solution is) into a Hamiltonian and then applies quantum operations; next, a classical optimizer reviews the result and adjusts parameters for the next cycle. This iterative loop continues until a desirable solution is found.

Because of this structure, QAOA is well suited to near-term quantum hardware (often called noisy intermediate-scale quantum devices) and is applied for combinatorial optimization problems in which an approximate but high-quality solution is acceptable.

Although both methods aim to find low-energy, high-quality solutions, they work very differently in practice. For example, if you were optimizing delivery routes for a fleet of trucks, a quantum annealer would encode the entire routing problem into a physical energy landscape and let the system relax into a good solution QAOA would approach that same problem iteratively, running a quantum circuit, checking the result, adjusting parameters and repeating until the solution improves. This difference makes annealing better suited for large, highly structured combinatorial problems, whereas QAOA offers more control and flexibility on today's gate-based quantum devices [5-7].

> [Optimization] is what we do every day without realizing it, from planning the fastest route to work to deciding the most efficient way to run errands on a busy afternoon. Quantum optimization is still in its early stages, but its direction is unmistakable.

What Does the Industry Think?

Quantum computing is emerging as the next frontier in operational efficiency, promising speed and scale far beyond even today's most powerful supercomputers. According to a 2025 survey conducted by Wakefield Research on behalf of D-Wave [8], the vast majority of business leaders recognize both the urgency and opportunity in optimization: 81% believe they've already reached the limits of what classical optimization can deliver, and 88% say their companies would go above and beyond for even a 5% improvement in efficiency. Yet many remain held back by outdated technology (39%), budget constraints (38%) and reliance on classical methods (36%).

Despite these hurdles, momentum is building. Nearly three in four organizations (72%) that have invested in or plan to adopt quantum optimization within the next two years anticipate returns of at least $1 million, and 85% believe they risk falling behind if they fail to act soon. Already, 22% of executives surveyed report quantum optimization making a huge impact for early adopters, whereas half (50%) expect it to be disruptive within their industries. These figures underscore that the value of quantum optimization is no longer theoretical; it's being recognized as a differentiator. That shift from potential to practice is becoming visible across industries today. Quantum optimization is one of the fastest-growing domains of quantum computation, with both academia and business investing heavily in early solutions.

The clearest evidence of this shift comes from the companies already experimenting with quantum optimization in their day-to-day operations:

- BMW Group used recursive QAOA to tackle complex supply chain "partition problems," achieving results comparable to leading heuristic algorithms.
- Volkswagen applied quantum-based routing in Lisbon to dynamically adjust public transport routes based on real-time traffic data.
- Toyota Central R&D Labs developed quantum-inspired traffic control methods to ease congestion by optimizing signal systems.
- Coca-Cola Bottlers Japan Inc. leveraged quantum optimization to streamline deliveries to more than 700,000 vending machines, improving efficiency and reducing fuel use.
- ExxonMobil explored quantum algorithms for maritime routing, reducing travel time and energy consumption.

Together, these insights and examples illustrate a clear reality: Although large-scale quantum advantage is still developing, quantum optimization is already producing tangible returns and performance gains for early adopters willing to experiment [8, 9].

But Is It Perfect?

Quantum optimization holds incredible promise, from solving complex logistics to reshaping manufacturing and design, but the road to practical, scalable use is still full of hurdles. Even as early pilots show encouraging results, several challenges continue to slow real-world adoption.

Current quantum hardware is still constrained by noise, limited qubit counts and instability, making it difficult to run large-scale commercial workloads. Competing approaches, from annealing to gate-based to hybrid architectures, also create uncertainty, with no clear winner yet. As Gurobi notes, quantum optimization still lacks a definitive, exact algorithm, and reproducible results remain difficult because of the technology's probabilistic nature. In practice, this means that although quantum optimization is beginning to deliver value in controlled pilots, most applications today remain hybrid, experimental and domain specific. As McKinsey (via Gurobi) observes, quantum computing is advancing quickly but remains "largely experimental and hypothetical at this early stage" [10].

Quantum optimization is still in its early stages, but its direction is unmistakable. As classical systems reach their computational limits, quantum approaches are emerging as the next step forward – not to replace traditional computing but to extend it. Over the next few years, most real progress will come from hybrid systems, where quantum processors explore vast spaces of possibilities and classical computers refine the results into precise, actionable solutions. The evolution will be gradual. Algorithms will improve, hardware will stabilize and industries will learn how to express their challenges in quantum terms. Companies that start building familiarity, experimenting with pilot projects or developing internal expertise now will be best positioned when quantum capability scales. At its heart, optimization is about making better decisions, balancing what's possible with what's practical. Quantum optimization simply expands that horizon. It allows us to look beyond the limits of sequential thinking and glimpse the full landscape of potential outcomes. In the end, quantum optimization isn't just a technological shift: It's a new way of thinking about problem-solving itself. Step by step, it's redefining how we understand efficiency, possibility and the very nature of choice.

SARA LODHA is a senior data analyst focusing on supply chain and operations. She holds a master's degree in business analytics from UC San Diego and has experience in supply chain strategy and manufacturing operations. She is passionate about solving business problems using data-driven solutions.

REFERENCES

1. https://scienceexchange.caltech.edu/topics/quantum-science-explained/quantum-physics
2. https://www.ibm.com/think/topics/quantum-computing
3. https://www.geeksforgeeks.org/physics/real-life-applications-of-quantum-mechanics/
4. https://pubsonline.informs.org/do/10.1287/orms.2025.02.16/full/
5. https://docs.dwavequantum.com/en/latest/quantum_research/quantum_annealing_intro.html
6. https://milvus.io/ai-quick-reference/how-does-quantum-annealing-work-in-solving-optimization-problems
7. https://quantum.cloud.ibm.com/docs/tutorials/quantum-approximate-optimization-algorithm
8. https://www.dwavequantum.com/media/sdylxhd0/quantum-computing-the-key-to-addressing-today-s-complex-business-problems.pdf
9. https://lingarogroup.com/blog/when-quantum-computing-meets-supply-chain-management
10. https://www.gurobi.com/faqs/quantum-optimization/

LESSONS FROM A LONG-TERM STRATEGIC PARTNERSHIP

Colombia's Organización Corona and Long Island University tackle real-world problems together

BASED ON YEARS OF COLLABORATION, this article highlights the strategic partnership between Organización Corona, a Colombian multinational specializing in home improvement, and the Data Analytics and Strategic Business Intelligence program at Long Island University (LIU).

This mutually beneficial relationship allows LIU students to tackle real-world projects from Corona, ranging from machine learning forecasting models to artificial intelligence (AI) chatbots. These initiatives provide students with high-level technical experience while delivering actionable business solutions.

The following interview captures the perspectives of Rodrigo Estrada (director of innovation, Corona), Juan F. Monsalvo (lead data scientist, Corona) and Associate Professor Juan R. Jaramillo (LIU) on the success of this academic-industry synergy.

Interview with the Director of Innovation

In a fast-paced industry, why is it worth the company's time to invest a full semester into a student collaboration?

Rodrigo Estrada: For us, this investment of time goes beyond simple collaboration; it is about building a three-way shared value ecosystem. We view this as a strategic symbiosis between the university, the student and the company.

First, we provide students with real, current industry problems, allowing them to apply their knowledge in live scenarios – something irreplaceable for their development as the talent of the future. Second, for the university, it is a critical opportunity to validate the relevance of its curricula. And third, for Corona, it is a strategic commitment: it allows us to see what's happening "out there" and build internal capabilities regarding new technologies, ensuring our future relevance.

How does this partnership function as an "external R&D lab" for exploring ideas that your internal roadmap can't prioritize?

Estrada: This alliance functions as a "safe harbor" for uncertainty. In day-to-day corporate operations, resources are tied to short-term goals and an immediate, tangible return on investment (ROI). However, disruptive innovation requires oxygen and freedom to fail.

By collaborating with academia, we create an environment in which the primary objective shifts from profitability to learning and capability building. It gives us access to fresh minds with the time and intellectual curiosity to explore emerging technologies and unexplored paths that our internal roadmap could not financially justify.

Furthermore, we gain what we call cognitive diversity: We break the company's internal bias by incorporating the perspective of young students, often from different cultures and countries, who tackle problems without the usual corporate preconceptions.

As we often say at Corona, "What the company can do depends on what the company knows"; this initiative expands that frontier of knowledge.

When you sit in those final presentations, what specific qualities are you looking for in a student's delivery and solution?

Estrada: Beyond the technical solution, we seek a mindset. In these final presentations, we evaluate four fundamental pillars that distinguish a good executor from a future leader:

1. **Problem Redefinition:** The critical capacity to understand, question and, if necessary, redefine the problem to address the root cause.
2. **Capacity for Synthesis:** The ability to connect scattered dots between the problem and the solution and communicate them in a concise and practical manner.
3. **Human Skills (Soft Skills):** Empathy, leadership and assertive communication are just as vital as technical knowledge.
4. **Critical Thinking and Fundamentals:** Solid foundation and the ability to demonstrate "learning how to learn" and constant reinvention.

How does this collaboration help the company stay at the forefront of innovation compared to competitors who don't collaborate with academia?

Estrada: The competitive advantage lies in the culture of continuous learning and reskilling, which these programs inject into our organization. By interacting with academia, our internal teams are challenged to stay current and look beyond their day-to-day operations. It grants us capacity for foresight: helping us see and understand technological futures before they become industry standards, allowing us to prepare our internal capabilities while the technology matures and costs are optimized. Fundamentally, it is a tool to ensure Corona's relevance and survival in the long term.

Beyond the technical results, how does Corona view its role in giving back to the academic community and helping shape the next generation of data professionals?

Estrada: We view this as an inescapable ethical and social responsibility. As an organization grows and receives value from society, it has a duty to give back. Our philanthropic role here is to act as a facilitating bridge.

We aim to smooth the learning curve and ease the students' transition into the workforce by exposing them to corporate realities before they graduate. At the same time, we help universities remain relevant by providing feedback on the real challenges we face, bridging the gap between academic theory and industrial practice. However, there is also valuable introspection in this act: By helping shape the next generation, we help Corona recognize its imperative need to transform. It is a way of generating social value that, ultimately, strengthens the industrial fabric upon which we all depend.

Interview with the Data Scientist

How do you select which internal projects are "student-ready" versus those that stay strictly in-house?

Juan F. Monsalvo: The selection is based on a key distinction: exploration versus exploitation.

For students, we select projects of *high uncertainty*, where the main objective is not immediate efficiency but validating technical feasibility and developing new capabilities. These are challenges designed to navigate different technological paths and understand which one holds the greatest potential. We look for the student to act as an explorer, helping us to "see and understand futures," preparing us for the moment when those technologies mature and their costs are reduced.

Conversely, the projects we keep strictly *in-house* are those in which the technology, value capture and solutions are clear and predictable and therefore must be executed with speed and operational efficiency.

What has been the most surprising technical solution or approach a student team has proposed that your internal team hadn't considered?

Monsalvo: We had a revealing case in a sentiment analysis project for product reviews. The student team proposed a preprocessing layer using a large language model (LLM) before applying sentiment analysis.

The LLM was tasked with breaking down and segmenting the text to identify which parts specifically referred to the product, point-of-sale service and delivery logistics. This was brilliant because we discovered that more than 40% of our reviews contained mixed information; a customer could love the product (positive) but hate the delivery (negative). Previously, this might have been erroneously classified as a general negative product review. Thanks to this granular approach, we managed to better understand the opportunities for improvement.

Data science is 80% data cleaning. How do you prepare students for the "messiness" of your company's actual data?

Monsalvo: Our philosophy is: zero isolation from reality. We provide students with raw, real company data, complete with all its imperfections, noise and gaps. We believe that providing "clean" or synthetic data would do a disservice to the students professional preparation.

However, we don't leave them alone in the chaos. Because we have already scrubbed this data internally, we know the traps and necessary actions. We act as guides: Instead of telling them what to clean, we pose challenges and strategic questions that highlight inconsistencies. We aim to trigger their critical thinking, leading them to question: Is this data useful? Does this correlation make sense? We want the students to learn that data cleaning is not an operational task but a fundamental analytical process.

> The collaboration between Organización Corona and LIU demonstrates that a long-term academic-corporate partnership creates a unique innovation ecosystem.

Beyond the final code, how does interacting with these students benefit your team's technical culture or your approach to problem-solving?

Monsalvo: The most valuable technical ROI is *mental openness*. Interacting with these teams exposes us to a cognitive diversity that breaks our corporate tunnel vision. Students – lacking company biases – explore alternative paths and ask fundamental questions that we sometimes take for granted.

Additionally, there is a benefit of *technical reciprocity*: To guide students who have fresh knowledge, our internal team is forced to update and renew its own knowledge. It challenges us to be better technical interlocutors and keep up with the innovative solutions they propose. It is an exercise that raises the technical bar for the entire organization.

Interview with the Professor

What sparked the initial conversation with Corona, and why did you feel their projects were a good fit for your curriculum?

Juan R. Jaramillo: I am a strong advocate for experiential learning; real-world projects provide students with a motivation level that a textbook simply cannot replicate. I've observed that when external partners enter the classroom, student engagement and attitude improve significantly.

Previously, I worked with entities that provided single-semester problems. Although valuable, these short-term projects had two main drawbacks: Organizations expected immediate solutions that weren't always feasible, and the constant rotation of partners made it difficult to build meaningful momentum.

The spark for this specific partnership happened in 2018. While chairing the INFORMS Analytics Society's Innovative Applications in Analytics Award (IAAA), I invited Rodrigo Estrada to join our judging panel. During our discussions, he mentioned Corona's interest in creating a dedicated space to evolve their AI and analytics capabilities. I shared my vision for a long-term academic partnership with LIU. We quickly realized our goals were aligned to foster sustainable growth for both the company and our students.

How does the relationship work?

Jaramillo: From the outset, we eliminated financial ties and focused on prioritizing the learning process over immediate commercial benefit. This framework provides the flexibility necessary for academic exploration and relieves student pressure when results don't meet corporate expectations.

During this time, members of both organizations have developed a deep professional rapport. Our dialogue has evolved beyond simply sourcing semester projects; we now engage in a genuine exchange of ideas. We introduce Corona to emerging AI and machine learning methodologies, and they provide us with invaluable insights into how these theories are implemented in a complex corporate environment. This synergy has made project selection seamless; Corona understands our students' profiles and identifies challenges that perfectly align with our curriculum.

How does working on a live corporate project change the students' engagement compared to using clean or synthetic datasets?

Jaramillo: Working with live corporate projects introduces a level of complexity that synthetic datasets simply cannot match. Real-world data is inherently unstructured, messy and incomplete; navigating these "unclean" datasets is a vital, challenging skill that most students are experiencing for the first time.

Beyond the technical challenge, the level of engagement significantly increases when the work has the potential for real-world impact. Our students meet with Corona's data scientists biweekly – approximately eight times a semester. This creates a shift in accountability: students often feel a greater sense of responsibility to deliver quality work to a professional peer.

Furthermore, we foster a culture of "collaborative competition." By allowing students to view each other's progress, they are inspired to learn from their peers while striving to develop unique, high-value solutions.

How has the collaboration evolved over the semesters to keep up with the fast-moving pace of data science?

Jaramillo: Our collaboration coincided with the rise of transformer models and the integration of generative AI into the industry. This timing allowed us to create a rapid innovation cycle: On the academic side, we identify emerging techniques and algorithms, which Corona then reviews and tests within their production environments.

The relationship has since evolved beyond the scope of individual projects. We now meet periodically to explore broader ideas and strategic opportunities. These sessions help Corona implement innovative solutions while allowing LIU to continuously update our curriculum and ensure our teaching remains relevant.

Conclusions

The collaboration between Organización Corona and LIU demonstrates that a long-term academic-corporate partnership creates a unique innovation ecosystem. By prioritizing exploration over immediate profit, Corona gains fresh technical insights and cognitive diversity, and students develop essential professional capabilities by solving messy, real-world data challenges in a fast-evolving industrial environment.

RODRIGO ESTRADA is the director of innovation at Organización Corona, leading R&D, product lifecycle management, design thinking, new product development, entrepreneurship business models, innovation digital strategy, growth platforms and culture.

JUAN F. MONSALVO is the lead data scientist at Organización Corona, delivering AI, IIoT, computer vision and GenAI products scaling recognition and conversational agents across Organización Corona through robust MLOps.

JUAN R. JARAMILLO is an associate professor of data analytics and Applied AI in the Department of Innovation and Entrepreneurship at the Long Island University School of Business. He is the 2026 INFORMS Analytics Society President and Innovative Applications in Analytics Award Co-chair.

ENGINEERING GOOD DECISIONS: REFLECTIONS ON THE LIFE AND WORK OF RONALD A. HOWARD

BY ALI ABBAS, ERIC HORVITZ, CARL SPETZLER AND JIM MATHESON

THE 2025 INFORMS ANNUAL MEETING marked a special session dedicated to the first anniversary of the passing of Ronald A. Howard, a founder of the field of decision analysis and a pioneer of operations research (O.R.). We'd like to reflect on and share the life and work of our dedicated colleague.

An extensive biography of Howard's career was published in *Profiles in Operations Research* [1] by co-author Jim Matheson. We will provide further reflections about his legacy and multiple contributions, including a background on his formative years, connections between his interests, and how interactions with his academic advisors, Philip Morse and George Kimball, helped to catalyze the development of decision analysis that we know today.

Background and Influences

Howard is recognized for his breakthrough work in dynamic probabilistic systems, including the introduction of policy iteration [2]. He is widely celebrated for his contributions to and leadership in transforming statistical decision theory into the modern discipline and practice of *decision analysis* [3] – a phrase he coined to describe methods and practices for engineering good decisions.

Howard pursued his work on probability, decision theory and decision analysis at Stanford University and the Stanford Research Institute [4]. For decades, he delivered influential classes in decision analysis at Stanford and served as advisor to numerous doctoral students. He co-authored *The Foundations of Decision Analysis* with Ali Abbas [5], a comprehensive text in the area. Beyond the pursuit of maximum expected utility, Howard thought deeply about values, organizing classes and seminars on ethics and writing on an ethical code to guide decisions [6, 7].

Howard was passionate about communicating key ideas to colleagues and students. He shared his thoughts widely in many fora, including publishing his reflections over the years in *OR/MS Today* [8, 9]. Driven by a belief that decision analysis should be made available to young people, he co-founded the Decision Education Foundation, an organization aimed at adapting professional methodologies in decision analysis to people without expert training. He provided guidance to the organization for many years.

Childhood Through Undergraduate Education

Howard grew up during World War II in Rosedale, Queens, near his parents' candy store. Across the side street, he discovered the Rosedale public library – "where you didn't even have to pay," as he would recall. He shared that he read as many as six books a week during that time. Howard entered the Massachusetts Institute of Technology (MIT) as an electrical engineering student but wanted to ensure that he included the human interaction element in his education. He sought out a program at MIT in engineering and economics that required an additional year of course work, leading to a second bachelor's degree. Through this program, Howard took courses in probability under Robert Solow (future Nobel laureate) and labor relations under George Schultz (future U.S. Secretary of State).

While taking a probability class at MIT, Howard did not find the homework problems challenging enough. He engaged with the professor and offered to compose homework problems and solutions for future classes, in exchange for being excused from the standard assignments. The professor agreed. Howard graduated in 1955 with B.S. degrees in both electrical engineering and engineering economics.

Doctor of Science at MIT and Policy Iteration

Howard remained at MIT for his D.Sc. work. Phillip Morse was Howard's official thesis advisor, with George Kimball serving as his de facto advisor. Morse led the first Operations Research group in the U.S. during World War II and was head of the new MIT interdepartmental Operations Research Center, the ORC. His colleague, George Kimball, was a chemical physicist from Columbia University. Morse and Kimball wrote the first seminal book on O.R., *Methods of Operations Research* [10].

Howard often mentioned that Kimball and Morse had different perspectives on O.R. Morse pursued mathematical elegance, writing equations and working on innovative solutions. Kimball sought applied solutions to practical problems. Howard enjoyed both perspectives, which provided him with a balanced lens on both the theory and practice of O.R.

Howard worked closely with Kimball on projects with the Operations Research Group of Arthur D. Little, Inc. At MIT, Howard also met with Bob Sittler (a student of Bill Linvill) and Richard Bellman, who had invented dynamic programming. During that time, Howard was interested in long-run optimal policies within dynamic programming. He invented the policy iteration algorithm during his doctoral work, leading to the dissertation, "Studies in Discrete Dynamic Programming." Later, Bellman commented, "Of these methods of successive approximation, the most important one is due to R. Howard" [11]. Ten years later, Howard published the two-volume set *Dynamic Probabilistic Systems* [3].

Influences of Jaynes and Laplace

Howard's views on probability – specifically, that probability is subjective and reflects an individual's state of knowledge rather than an objective frequency – were influenced by the writings of Edwin Jaynes [12] and Pierre-Simon Laplace. It was Myron Tribus, then dean of engineering at Dartmouth, who gave Howard a manuscript written by Jaynes. Jaynes used inferential notation, which conditioned every probability assignment on the state of an individual's information. This notation was later adopted by Howard in many writings, where the state of information is represented by an ampersand (&).

Jaynes was a scholar of the work of Laplace. Howard noted that whenever someone brought up a criticism of Laplace, Jaynes would go back to the original source and find that Laplace had been misquoted. In this spirit, Howard obtained a copy of Laplace's original work in French, "Théorie analytique des probabilités," and learned French to ensure that he captured the precise nuances of Laplace's thought without translation.

Stanford and Decision Analysis

After his postdoctoral work at MIT, including numerous interactions with Howard Raiffa at Harvard University, Howard joined the new department of Engineering-Economic Systems (EES) at Stanford University in 1965. Bill Linvill, then department chair, advocated that all Ph.D. students gain practical experience through internships at companies. Linvill and Howard focused their attention on having students exposed to the work at Stanford Research Institute (SRI) (now SRI International). SRI was closely associated with Stanford University, though there were few faculty connections at the time. The EES program at Stanford was established with seed funding from two SRI departments.

In early 1965, Howard had invited Jim Matheson to take a consulting position with SRI. A year later, the team at SRI created the Decision Analysis Group (DAG), headed by Jim, with Howard as the key academic member. Their vision for DAG was to be the "teaching hospital" for the advancement of decision analysis. In this role, Howard actively collaborated with Jim, and they compiled and

> Howard's legacy will continue to inform the way we approach the most difficult decisions as we move into the future.

edited the "Blue Books" [4], a collection of articles referred to by generations of Howard's students. The DAG team became a hub of international work on the practice of decision-making, attracting visitors such as Tribus.

Stepping back, decision analysis is built on the foundations of statistical decision theory. While giving an executive seminar to General Electric, a participant asked Howard whether what he was teaching could help them with a decision about nuclear waste. Howard replied, "Of course," and he spent that night wondering how his teachings could help with the practicalities of such a massive, complex problem. He later coined the term "decision analysis" to bridge the gap between theory and practice [2]. Since that time, numerous tools and practices have been developed for decision analysis.

Along with Howard and Matheson, Tribus became a champion for the practice of decision analysis. The DAG consulted on projects when Tribus served as Undersecretary of Commerce for Standards and Technology, including the "Project Stormfury" study on whether to seed hurricanes [13] to diminish their destructiveness. When Tribus joined Xerox, he continued to be a major client of the DAG team, where another member, Carl Spetzler, had the opportunity to work closely with Tribus for many years. They, alongside Carl-Axel Staël von Holstein, Amos Tversky and Daniel Kahneman, conducted workshops on probability assessment aimed at minimizing bias.

Strategic Decisions Group

In 1981, Howard, Spetzler, Matheson and Jeff Foran formed the Strategic Decisions Group (SDG), a decision-analysis consulting business. SDG attracted outstanding staff, including many of Howard's students and alumni from the SRI DAG teams.

Tribus continued to play a strong role as client and supporter. He became a convert to Deming's "total quality management" and challenged leaders at SDG on their approach to practice. Although decision analysis centered on reaching clarity on the best action under uncertainty, Tribus helped establish the "decision quality" (DQ) paradigm. With DQ, practitioners recognized that clarity of intent is distinct from commitment to action. In large, complex organizations, getting true commitment to action requires significant participation and conflict resolution to gain alignment among actors.

Ethics

Howard believed that ethical difficulties tended to remain unaddressed because it is much easier to avoid complex discussions about sensitive values and trade-offs than to reflect and converge with clarity on points of view. He taught that ethical difficulties can be avoided by following three practices: (1) decline being part of organizations that have ethical codes and behavior inconsistent with your own, (2) avoid participating in ethically objectionable activities and (3) treat all people as you would those you care about.

He developed the first ethics course in the Stanford School of Engineering and titled it "The Ethical Analyst." During the class, Howard would provide examples and distinctions to help students form their personal ethical code. Many derivatives of this work were later published, including *Ethics for the Real World* (co-authored with Clint Korver) [6] and what is believed to be Howard's last work, "A Hippocratic Oath for Technologists," co-authored with Abbas [7].

Decision Analysis for All: Decision Education Foundation

In 2001, the early pioneers of decision analysis, including leaders at SDG and several of Howard's students (David Heckerman, Eric Horvitz, Tom Keelin and many others) reflected that the power of decision analysis and decision quality should extend beyond the corporate world. How could the insights, practices and tools be leveraged more universally? This thinking led to the formation of the Decision Education Foundation (DEF), with Howard as its first president. Howard continued as a DEF director until his passing. It was this initiative that led to numerous *OR/MS Today* articles, including the prospect and promise of teaching decision skills to youth and troubled teens [8, 9].

A Legacy Extending Into the Era of AI

Howard has mentored many students all over the world, largely within statistics and decision-making. In the 1980s, several of his students became leaders in the blossoming area of leveraging probability and statistical decision theory to solve artificial intelligence (AI) challenges. At the time, the work of David Heckerman, Eric Horvitz and colleagues [14] was considered countercultural; the dominant paradigm for building systems had become focused on logic-based methods, often skipping over the critical importance of probabilistic representations and reasoning under uncertainty. Howard's mentorship and support had an important influence on the growth of uses of decision-theoretic methods and concepts in AI, adapting decision theory and tools designed for human decision-making to the challenge of automated reasoning under great uncertainty and incompleteness. A branch of AI became dominated by probabilistic graphical models and decision-theoretic optimization, and AI evolved to more centrally rely on probability and decision theory, including the wide use of "reinforcement learning," which builds directly on the work of Richard Bellman and Howard's contributions to policy iteration.

Talks given by several of Ron's former students at his retirement celebration provide a sense for the influence that he has had on the lives of colleagues and students around the world [15, 16].

Concluding Reflection

We hope that these reflections about the life and legacy of Ronald A. Howard will shed light on the influences on his early work, the founding of the vibrant field of decision analysis and the ripple effect that his mentorship had on students, colleagues and the world. Howard's legacy will continue to inform the way we approach the most difficult decisions as we move into the future.

ALI ABBAS is a professor of industrial and systems engineering at the University of Southern California. He is the recipient of research awards from INFORMS and IEEE and is a distinguished lecturer for the IEEE Systems Society. He has analyzed numerous homeland security decisions including the TSA Pre decision with the TSA's chief risk officer. He completed his doctoral work at Stanford University under Ronald Howard, and together, they authored *Foundations of Decision Analysis* and the forthcoming book, *Advanced Topics in Decision Analysis.*

ERIC HORVITZ is Microsoft's chief scientific officer, leading company-wide initiatives in science and technology. He is a recipient of the Feigenbaum Prize and Allen Newell Award for contributions to AI, a member of the National Academy of Engineering and American Academy of Arts and Sciences, as well as a fellow of AAAI and ACM. He completed his doctoral work at Stanford University under Ron Howard, developing normative models of bounded rationality for high-stakes, time-critical decision-making – work that laid the foundation for a lifelong pursuit of decision-theoretic approaches to AI.

CARL SPETZLER has been one of Ron Howard's partners and friends for over 50 years. Together with Jim Matheson, they founded SDG (Strategic Decisions Group) and DEF (The Decision Education Foundation). Carl is the lead author of "Decision Quality – Value Creation from Better Business Decisions" (Wiley, 2016). He is a recipient of the Ramsey Medal from the Decision Analysis Society of INFORMS and the Pioneer Award from the Society of Decision Professionals.

JIM MATHESON founded and led the world's first Decision Analysis Group at SRI International, co-invented influence diagrams with Ronald Howard, and co-founded the Strategic Decisions Group with Howard, Carl Spetzler and Jeff Foran. Matheson later co-founded SmartOrg, Inc. and served for over four decades as a Consulting Professor at Stanford University. He is a recipient of the Ramsey Medal of INFORMS and a fellow of both INFORMS and the Society of Decision Professionals.

Note: *References can be found online: https://doi.org/10.1287/orms.2026.01.02.*

Ron Howard and Eric Horvitz.

DIGITAL FIRST RESPONDER: A SCALABLE AI TO UNCOVER HIDDEN PATIENT RISK

BY AARON LAI AND LI-LIN LIANG

A horse! A horse! My kingdom for a horse!
– William Shakespeare, *Richard III*

PRIORITIZING PATIENTS FOR TREATMENT is often a life-and-death decision. The frontline clinicians must make decisions with incomplete information, and an incorrect decision could result in delay or even death. An artificial intelligence (AI) assistant could "triage at scale" to help clinicians in challenging environments, such as disaster response after an earthquake.

> This LLM-enabled scalable triage tool could save time and lives, especially under intense pressure; our approach supports a safer priority rule by surfacing hidden high-risk signals from longitudinal history in seconds.

Gaps in Current Triage

Consider this scenario: A 69-year-old male with a minor cut on his head from a collapsed ceiling is conscious, reports mild dizziness and walks around to check on others. He appears to be a low-risk patient to the field clinician. Unfortunately, he is taking warfarin (a blood thinner that prevents his blood from clotting). His head could have an internal hemorrhage and thus need a head scan immediately. This is when an AI triage assistant could synthesize his full clinical history to inform the frontline staff and reduce their cognitive loads.

What Is a Triage, and Why Is It Important?

Emergency departments (EDs) are often overcrowded. The triage clinicians often prioritize patients based on observations. Therefore, we propose a large language model (LLM) tool to inform users of the relevant clinical history of a patient to achieve "triage at scale," which is particularly critical during large-scale crises.

Beyond crisis management, this technology is vital for navigating the daily complexities of an aging population. It offers physicians the most relevant information at the point of care. Ultimately, deploying such a tool enhances the health system's allocative efficiency, providing a scalable strategy to tackle rapidly growing health care expenditures.

The Components and Rationale

The proposed solution contains a data layer, a triage decision engine and a validation layer.

The first element is a data layer that uses Fast Healthcare Interoperability Resources (FHIR), a modern standard for electronic health information exchange. The triage LLM needs to ingest and understand data in its native format for standardization and auditability.

The second element is a triage decision engine that uses an open-source or open-weight LLM if possible. Clinical data is sensitive, and it may not be possible to be accessed by a commercial LLM. Another benefit is that the system will not be affected by unplanned changes to ensure consistency.

The third element is a validation layer (LLMs-as-judges) that acts as a quality assurance of the triage LLM suggestions. During the development phase, the reasoning of the recommendations from the triage LLM will be compared with the explanations (ground truth) from the test data to establish validity. During the implementation phase, the reasonings are checked for consistency. It is not practical to run validations in every decision, but it is advisable to run validation on regular intervals to check against unexpected model degradation. Figure 1 shows an illustrated diagram on using an ensemble (LLMs from different families for diversity to minimize inherent bias) to validate the output of a specialized LLM as "AI Peer Review."

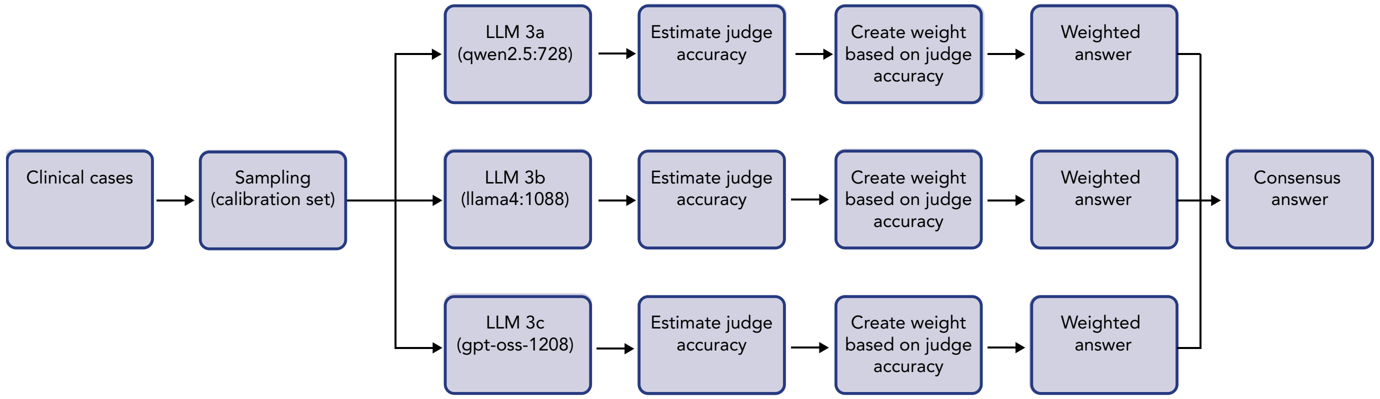

FIGURE 1: An illustrative diagram on using LLMs-as-judges (AI Peer Review).

Prototype Performance Comparable to Baseline Human Accuracy

The lead author created a proof-of-concept prototype using 100 clinical cases from an emergency medicine handbook [1] transformed into FHIR using Gemma 3 27B (an open-weight LLM). A MedGemma 27B (4-bit quant) model (a clinically trained open-weight LLM from Google) is used because it understands clinical terms and can directly ingest FHIR data.

The overall Emergency Severity Index (ESI) prediction accuracy (between actual and predicted ESI) is 56% (exact ESI), which is on par with the 59% from humans [2]. The LLM tends to be more cautious (guardrail bias), with only 2 out of 100 cases being undertriaged. The clinical reasoning of the predicted severity is consistent with the handbook explanations in 90% of cases. Consistency was assessed by checking whether the model cited the same risk factors as the handbook rationale. The clinicians are decision-makers and need to know and trust the clinical reasoning. This is also a profound point about explainable AI.

Patient safety is the priority because emergency triage errors have asymmetric costs: a false negative (underestimating risk) can be fatal, whereas a false positive (overestimating risk) may result in additional tests. Therefore, the model is tuned to minimize adverse outcomes and prioritize patient safety over resource efficiency. We accept increased utilization of precautionary checks to prevent the catastrophic failure of missing life-threatening conditions. This is why "human-in-the-loop" as a decision-maker is paramount, even though the process is more efficient.

To move from prototype to practice, it is essential to follow the responsible AI principles to ensure that the prototype will not inadvertently discriminate against any group. The reasoning output should be detailed enough to understand but simple enough to avoid cognitive overload. Clinical decisions are high stakes: Extensive testing on various situations and edge cases are needed to ensure consistency. In production, the judge ensemble should regularly run on a stratified sample and route high-disagreement cases to clinical review.

To sum up, this LLM-enabled scalable triage tool could save time and lives, especially under intense pressure; our approach supports a safer priority rule by surfacing hidden high-risk signals from longitudinal history in seconds.

AARON LAI, CFA, serves as a senior fellow of the Krenicki Center for Business Analytics and Machine Learning at Purdue University. He has extensive experience within the healthcare industry and is currently in the Doctor of Technology program at Purdue, with research concentrating on LLM triage in critical care scenarios. The perspectives presented herein represent his personal views and may not necessarily align with those of his employer or other affiliations.

LI-LIN LIANG, Ph.D., is an associate professor at the Institute of Public Health and Deputy Director of the Health Innovation Center at National Yang Ming Chiao Tung University. Her research focuses on evaluating health care reforms, with a current emphasis on applying machine learning methods to simulate the potential effects of interventions and policy changes.

REFERENCES

1. ESI Triage Research Team, "Emergency Severity Index, Version 4: Implementation Handbook," AHRQ Publications, https://www.govinfo.gov/app/details/GOVPUB-HE20_6500-PURL-gpo23161.
2. www.ena.org/education/triage.

IMPROVING RESPONSE THROUGH ADVANCE OPERATIONAL PLANNING AND ACADEMIC COLLABORATION

BY LOUIS LUANGKESORN

THE MISSION OF THE AMERICAN RED Cross is to prevent and alleviate human suffering during emergencies by mobilizing the power of volunteers and the generosity of donors. The Red Cross responds to major disasters such as hurricanes, floods and earthquakes. However, large-scale responses are complex, especially in the early days when information is incomplete or prone to error. Rapid action depends on estimating need, but early estimates of needs and resource requirements are challenging to produce with limited data. Red Cross planners work with academic researchers to research and develop models that are used as part of the planning process in disaster response. Currently, Red Cross planners have been working with Carnegie Mellon University through student projects and with researchers at Virginia Tech in long-term research work.

To improve disaster response speed and efficiency, the Red Cross is developing doctrine, procedures and tools for advance operational planning (AOP), which considers the expected area and intensity of disaster impact, as well as the demographics of the impacted area. This approach allows the Red Cross to anticipate demands and mobilize resources before complete information is available. This work has been ongoing since 2015.

One step of AOP is to make an estimate of the disaster-caused need, which is based on the geographic area of the disaster, intensity of the disaster and demographics of the impacted area. In the days prior to a notice disaster (such as a hurricane) or the day of a no-notice disaster (such as an earthquake), the Red Cross makes an estimate of the need; then the AOP process translates that into estimates of the human and material resources required to meet that need. Whereas many of these estimates rely on expert intuition informed by experience, the AOP process involves developing standard processes incorporating predictive models that capture the knowledge of subject matter experts to make these estimates more robust. The Red Cross can then modify the plan as needed as more information becomes available.

The Red Cross has developed and put into use predictive models for hurricanes, and with an AOP process that incorporates predictive models, it could incorporating models for other disasters, such as earthquakes, wildfires, tornados and tsunamis. Developing these models would require significant resources beyond what Red Cross staff and volunteers alone can provide, creating an opportunity for collaboration with academic partners.

Academic Partnerships: Collaboration Benefits and Structure

When the Red Cross works with academic partners, the goal is to develop an understanding and analytics capability that can be incorporated into the Red Cross Concept of Operations. Therefore, when starting a partnership with an academic researcher, the Red Cross introduces the researcher to its operational framework and processes. This helps the researcher understand the time and resource constraints and focuses on any research and development toward results that can be incorporated into Red Cross practices.

Red Cross staff and volunteers have some ability to research methods for early disaster impact estimation, but operational demands limit their bandwidth. Academic researchers and students help fill this gap by

- identifying additional resources and datasets for assessing disaster impacts,
- staying up to date with the latest research and modeling technologies and
- implementing models using accessible platforms.

In one student project, students used Google Colab for an earthquake impact prediction project. Google Colab is a web-based analytics platform that allows users to run analytical code on the cloud, allowing nontechnical staff and volunteers to utilize analytical applications. This introduced new technical capabilities that the Red Cross can use in other contexts.

Another line of research that Red Cross academic partners follow is looking at integrating other types of data into the predictive model. Red Cross academic partners have experience in working with data such as social vulnerability indicators in other contexts and can evaluate a range of different types of models that integrate these data sources with data already in use.

One important consideration when working with academic partners is to understand how work on the project impacts their career development in the context of their institution. For long-term results, the work must advance the career of the academic partner. This is complicated because the products developed for the Red Cross do not typically contribute to an institution's faculty performance review, so there need to be additional outcomes that can contribute to their career/tenure. In addition, essential activities such as framing the problem in a way that contributes to both Red Cross and academic career goals can take time to fully resolve, which may cause problems with an academic's career timeline.

How a project can help an academic partner's institutional advancement varies by institution. For example, Mariana Escallon-Barrios is a teaching faculty member at Carnegie Mellon University's Heinz College of Information Systems and Public Policy. Although published research papers are not specifically required in her position, research that involves engaging in real-world problems is encouraged and can be documented through project outcomes. Academics can work with industry partners through industrial sponsored capstone or course projects. For example, projects can involve Red Cross volunteers and staff engaging with the students and school staff, providing practical engagement with real-world situations in which students use their analytic skills gained through their master's programs on projects that impact disaster response.

Christopher Zobel and Andy Arnette, faculty members at Virginia Tech's Pamplin College of Business focusing on operations management, find that their engagement with the Red Cross, while not directly contributing to their institutional evaluations, has valuable side effects. The practical problems they address with the Red Cross ground their work in real-world applications. This is particularly true when working with data, because the Red Cross data reflects the operational issues the organization faces, anchoring the research in ways that benefit the Red Cross and attract interest from certain publication outlets. Additionally, the social impact of faculty collaborating with a nonprofit such as the Red Cross is personally meaningful and enriches their teaching, because the operational environment of the Red Cross in disaster response differs from the examples typically found in textbooks. Incorporating these real-world scenarios into class discussions allows for the exploration of different goals and details that lead to varied solutions, even though the fundamentals of operations remain the same.

Another major factor in making collaborations work is the communication of expectations for deliverables and timeline. As with the Red Cross planners developing and refining its AOP tools on an ongoing basis, the work done with its academic partners was not on a strict timeline. Similarly, because the academic faculty involved were not on strict timelines with their institution, they could take time to develop understandings that lead to results. Although work that involves students needs to be more well defined to fit student academic schedules, the overall relationship has time to develop, so it can be beneficial for the academic partner careers as well as meet Red Cross goals.

Future Directions: Ongoing Model Improvement and Collaboration

Fruitful collaboration with academics and practitioners depends on several key characteristics. First, open communication is essential – aligning goals and incentives. Additionally, flexibility and adaptability are crucial to success, allowing both parties to achieve outcomes in practice and benefiting the practitioner, academic faculty and involved students. In this case, the collaboration benefits from being part of the Red Cross' ongoing program of research supporting the development of the AOP process. It is this ongoing process that can incorporate the results of academic partnerships into its operations. This research and development improves the Red Cross' ability to fulfill its mission to alleviate suffering in the face of disaster.

Acknowledgments

I would like to acknowledge Dr. Escallon-Barrios, Dr. Arnette, Dr. Zobel and Mr. Whitehead for taking the time to write and talk about their experiences in the Red Cross collaboration with academia as part of the development of the Red Cross Advance Operational Planning process.

LOUIS LUANGKESORN is the lead data scientist at Highmark Health and a volunteer national planner for the American Red Cross.

LEARN, GROW AND CONNECT AT THE ANALYTICS+ CONFERENCE

Up your game, now and for the future

BY ZOHAR STRINKA, CAP-X

> We consistently hear that the best part of Analytics+ is the opportunity to connect with others.

IN JUST A FEW SHORT WEEKS, YOU WILL have an exceptional opportunity to sharpen your skills, power up your future career and network with the leading experts in the field of analytics.

Don't miss it.

The 2026 INFORMS Analytics+ Conference takes place April 12-14 in National Harbor, Maryland. Wherever you are in your analytics journey, joining us at this event will help you learn and grow.

Attending this conference year after year helped me launch my career, offered me new opportunities and continues to keep me at the cutting edge of analytics. As chair of this year's conference, I'm a little bit biased. But experience has shown me just how important these few days a year are in sharpening my analytics thinking.

Ten years ago, I was wrapping up my education in operations research (O.R.) and had an opportunity to go to a new-to-me INFORMS conference in the spring. I had two important insights at that conference. First, the profession of analytics is inherently industry-agnostic. Second, this conference is where that cross-industry community of professionals comes together to learn how to do our jobs better.

During that first conference, I made many new connections with whom I'm still in touch today. Later, at the 2022 INFORMS Business Analytics Conference, I met a future colleague who led to an inflection point in my career.

Each year I attend, I learn new things – even as an experienced professional. I come out of talks ready to use the state-of-the-art where I should and knowing more best practices every time. I return to my work both inspired and more equipped to do analytics correctly, the first time.

Inspiring Keynotes

This year, we have two fantastic keynote speakers lined up for National Harbor. We know they will both inspire you and help you level up in analytics.

On Monday, April 13, Scott Taylor, a passionate data evangelist colloquially known as "The Data Whisperer," will help us understand why and how to tell a great data story that will get buy-in from leaders and managers. We're excited to bring this kind of message to you because effective communication is so vital – not only to your own career but also to help everyone get more value from analytics.

On Tuesday, April 14, Cynthia Rudin will inspire you to think beyond maximizing one fixed goal and instead how optimizing for interpretability also helps adoption. She has led the field in thinking about how to minimize the trade-off between transparency and accuracy, while showing just how easy it is to improve on the default approach. This message is crucial not only for analytics professionals but for anyone trying to judge the *value* of using analytics.

Training & Networking for Every Level

Early in our planning for the 2026 Analytics+ Conference, we settled on a philosophy of "training for every level" [1] because no matter where you are in your career, you can benefit from continued learning.

With that in mind, we have reshaped Tuesday's agenda to end with a training workshop block. Here, you'll have the opportunity for hands-on learning, deeper discussion and even more networking with peers. *Each session will have a take-home tool or resource you can bring back to your organization and start using immediately!*

You will find panels for early career, mid-career and senior leaders at our Sunday evening parallel workshops, which will each include valuable networking opportunities. If you are a first-time attendee trying to learn how to make the most of the conference, these sessions are especially useful.

We consistently hear that the best part of Analytics+ is the opportunity to connect with others. If you or your organization are thinking of trying something for the first time, chances are you'll find an attendee who has done it before and can tell you what's worked, what hasn't and where the "gotchas" are.

Talking with your peers can be the best way to learn what's happening at other organizations. Many of our attendees have decades of experience in analytics and view attending the conference and sharing guidance as a crucial way to give back. But no matter how much experience you have, analytics is so varied and broad that everyone will go home with new insights and tools.

New for 2026!

Hands-On Analytics Training Workshops
Tuesday, April 14, 2:30–4:30pm EST
You asked. We listened.

In 2026, we're going deeper with interactive training workshops designed to turn analytics concepts into real-world impact.

These sessions go beyond theory, giving you practical tools, expert insights and immediately usable takeaways.

With sessions designed for a range of experience levels, this workshop series offers something for everyone.

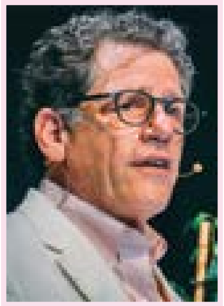

LEFT: Scott Taylor
RIGHT: Cynthia Rudin

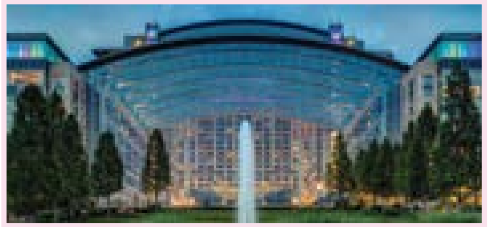

Gaylord National Resort

The Decision-First Mindset

INFORMS has always brought a decision- and impact-first perspective to what it means to work in analytics. Whereas other conferences focus on methods or technologies, INFORMS looks to real-world actions and decisions. Some of the ways you will see this in action at the Analytics+ Conference include the following:

- The Edelman Competition, in which Chewy, the Department of Food and Public Distribution (DFPD) of India, ECCO Shoes, Google, Microsoft and Nvidia will compete to convince the judges that they have delivered real value for their organizations using analytics. These presentations are open to all and will give you new ideas for how to measure the value of the work we do. Then head to the Edelman Gala Monday evening to find out who won our "Nobel Prize of O.R."!
- The INFORMS Analytics Framework™ as a structured way of navigating the conference and finding exactly the right skills to accelerate your career.
- Networking with Certified Analytics Professionals® (CAPs) who have earned one of the three levels to see how the Essentials, Pro or Expert certification can help you do more in your work. These certifications validate that you know what it takes to get stakeholder buy-in or sustain a model for the long haul.

Register Today!

Registration is now open! Visit https://meetings.informs.org/analytics/register/ to register. If you have at least six people from your organization attending, optimize your prices by using the group rate.

Take advantage of the hotel block at the Gaylord National Resort so you can make the most of your time at the conference to take the shortest path there.

For more information and registration details, visit the Analytics+ website at https://meetings.informs.org/analytics/.

On behalf of INFORMS and the 2026 INFORMS Analytics+ Conference organizing committee, we look forward to you joining us in National Harbor as we explore how analytics continues to drive smarter decisions for a better world.

ZOHAR STRINKA, CAP-X, Ph.D., is a consultant focused on solving data and process problems for her clients in a variety of industries. She is the general chair of the 2026 INFORMS Analytics+ Conference. Check out meta-problem.com.

REFERENCE

1. https://meetings.informs.org/wordpress/analytics/training-opportunities/

FROM CONTENT FACTORY TO CAPITAL ALLOCATOR: REFRAMING MARKETING AS A SEQUENTIAL DECISION SYSTEM

BY PAVAN KUNCHALA

Marketing's Hidden O.R. Problem

Artificial intelligence (AI) now sits at the center of many marketing decisions. Systems generate content, adjust bids and offers, and modulate customer contact in near real time, all tracked through increasingly granular performance metrics. These capabilities, however, mask a deeper issue. The problem is not execution speed or model complexity but decision framing. In many cases, marketing AI is solving a short-term optimization problem in a setting that is inherently dynamic.

In practice, marketing decisions are treated as independent, short-term optimization tasks. Algorithms are built to maximize immediate response, attribution-adjusted return, quarterly return on ad spend or short-horizon revenue. This amounts to repeatedly solving a single-period optimization problem in an environment that is fundamentally dynamic. This framing treats short-term performance improvements as separable from longer-term enterprise value. Customer behavior, however, is not memoryless. Actions taken today shape future states, including engagement, fatigue, trust, price sensitivity and long-term profitability.

This article reframes marketing AI as a sequential decision system in which budgets, promotions, offers and exposures are allocated over time under uncertainty. Using an operations research (O.R.) lens, it focuses on state dynamics, reward formulation and policy evaluation rather than campaign-level optimization. The goal is not to introduce new algorithms but to clarify the decision structure that marketing AI systems are implicitly solving.

The Efficiency Trap: When Optimization Inflates KPIs but Erodes Profit

Most marketing AI systems are highly efficient at what they are asked to do. They efficiently maximize response rates, click-through or attributed revenue, when instructed to do so. In many organizations, this produces a familiar pattern:

- Frequent discounting to responsive segments
- Escalating contact intensity for high-engagement customers
- Reallocating budgets toward channels with short-term lift

Standard performance metrics improve, key performance indicators (KPIs) rise and attribution models report incremental impact. Short-term profitability may improve as well, but longer-term customer contribution often deteriorates.

This efficiency trap is particularly visible in mature retail and consumer packaged goods (CPG) environments, in which loyal customers are repeatedly targeted because they respond reliably. Over time, discounts cannibalize margin, frequent outreach accelerates fatigue and high-quality customers become conditioned to wait for incentives. The system behaves rationally within its objective function, even as enterprise value deteriorates.

The core issue is not algorithmic sophistication. It is problem formulation.

Why Attribution Models Fail as Decision Systems

Attribution models remain the backbone of many marketing analytics stacks. They estimate incremental contribution by channel, campaign or touchpoint and are often used to guide budget allocation. Although useful for retrospective analysis, attribution models are poorly suited to governing decisions over time.

These models exhibit several structural limitations:

1. Static framing that ignores state persistence
2. Independence assumptions in environments with strong interactions
3. Absence of transition dynamics
4. Reward proxies that substitute revenue for contribution margin

As a result, attribution-driven optimization behaves like a sequence of myopic decisions. Each period's allocation may be locally optimal, but the policy is not. This is analogous to managing an inventory system by optimizing daily sales without accounting for stockouts, replenishment cycles or holding costs. The system appears responsive, but performance degrades over time.

Marketing as a Sequential Decision Problem

A more appropriate representation of marketing decisions is a sequential decision system, where actions today influence future states and rewards.

Marketing can be modeled as a Markov decision process or a related sequential framework:

- **State:** Customer engagement, fatigue, tenure, price sensitivity, margin profile
- **Action:** Offer level, contact decision, channel selection or holdout
- **Reward:** Incremental contribution margin
- **Transition:** How actions move customers between states over time

This framing immediately resolves many pathologies of conventional marketing AI. High-frequency discounting is no longer attractive if it degrades future state value. Excessive outreach is penalized when it increases the probability of disengagement or churn.

Importantly, this does not require perfect state observability. Even coarse state definitions outperform flow-based optimization when decisions are repeated over long horizons. Once marketing is reframed as a sequential decision problem, the objective naturally shifts from maximizing activity to allocating capital across time under uncertainty.

Marketing AI as a Capital Allocation System

Recent advances in agentic AI make this reframing operationally feasible. In this context, *agentic* refers to policy-driven systems that autonomously select and repeat actions over time based on observed state and predefined objectives.

Rather than treating AI as a content generator or tactical optimizer, these systems can be designed as policy learners that allocate marketing budgets across customers, segments and time periods. They behave like portfolio optimizers, allocating finite capital such as attention, incentives and marketing spend, while balancing exploration and exploitation to maximize expected long-term return.

System performance is driven far more by how states and rewards are defined than by the choice of algorithm. A simple policy optimizing the right objective will outperform a more sophisticated model optimizing the wrong one. This is where O.R. and AI intersect most productively. O.R. defines the decision structure, and AI supplies the learning capacity.

> A simple policy optimizing the right objective will outperform a more sophisticated model optimizing the wrong one.

Policy Behavior Under Different Objectives

Customers differ in engagement decay, price sensitivity and contribution margin, and their behavior evolves as a function of prior exposure and purchase history. Two decision policies operate under identical information and budget constraints. The only distinction between them is the objective used to evaluate actions.

One policy prioritizes short-term response, allocating greater exposure to customers who have historically engaged. The other prioritizes long-term contribution, explicitly accounting for how current actions affect future customer value. The learning mechanism and available data are held constant.

Under repeated decision cycles, capital allocation behavior diverges. The response-focused objective concentrates capital on historically responsive customers, whereas the value-focused objective spreads capital to preserve future contribution.

Figure 1 illustrates how different decision objectives shape policy behavior over time in a marketing environment.

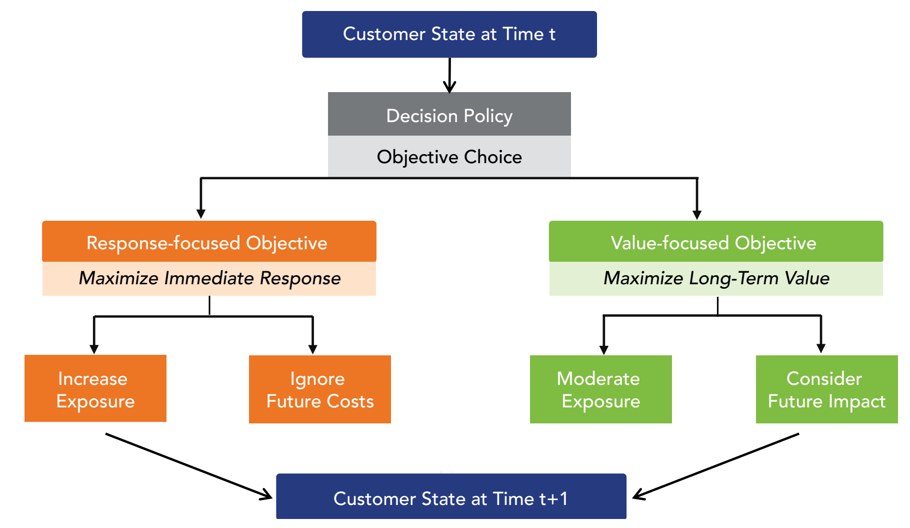

FIGURE 1: How objective choice shapes policy behavior over time.

Reframing marketing as a sequential decision system has direct implications for how analytics are designed, deployed and governed at scale.

Implications for Decision System Design

Reframing marketing as a sequential decision system has direct implications for how analytics are designed, deployed and governed at scale.

First, **problem formulation matters more than algorithm choice.** In sequential settings, how customer states are defined and updated has a greater impact on outcomes than whether the underlying policy is implemented with bandits, reinforcement learning or heuristic optimization.

Second, **reward design determines behavior.** Optimization systems reliably pursue what they are asked to maximize. If reward functions emphasize short-term response, systems will rationally favor actions that extract immediate value, even when those actions degrade future outcomes. Aligning rewards with long-term contribution is therefore a core design responsibility, not a modeling detail.

Third, **uncertainty is structural rather than incidental.** Customer behavior evolves because of prior decisions. Treating uncertainty as noise to be averaged away, rather than as a state-dependent feature of the system, leads to policies that overexploit transient signals and underinvest in long-term value.

Finally, **governance must be embedded in the decision process itself.** In sequential systems, oversight is expressed through constraints such as frequency caps, budget limits and protections against irreversible outcomes. These are not post hoc monitoring tools; they are first-order design decisions that shape how automated policies behave over time.

Conclusion

Automation in marketing is no longer optional. What remains a choice is how automated decisions are framed and evaluated over time. Marketing has always involved allocating scarce resources under uncertainty. What has changed is the scale, speed and persistence with which those allocations are now executed by AI-enabled systems.

Policy-driven decision systems grounded in operations research make it possible to treat marketing as what has long been in practice: an intertemporal capital allocation problem. When decisions are modeled explicitly as sequential and state-dependent, many familiar inefficiencies arise less from execution failures than from misaligned objectives. Over time, the quality of decision models is likely to matter more than content generation or algorithmic novelty in determining how marketing systems perform at scale.

PAVAN KUNCHALA leads applied AI and decision science work at Tiger Analytics, focusing on large-scale decision systems for marketing, pricing and revenue growth management in CPG and retail. His work applies operations research and decision modeling to customer lifetime value and enterprise decision systems. He holds an MBA from the University of California, Berkeley.

WHAT GenAI TAUGHT US ABOUT THE FUTURE OF GLOBAL OPTIMIZATION

BY TIMO BERTHOLD AND IMRE PÓLIK

ARE YOUR OPTIMIZATION MODELS nonlinear yet? Recent advances in nonlinear optimization technology are paving the way for the use of full-scale nonlinear models in production, thereby mitigating the need for linearized models.

Nonlinear Optimization Coming of Age

For decades, operations research (O.R.) practitioners have been taught to bend reality by modeling applications as linear optimization problems. Nonlinear relationships were routinely approximated, relaxed or ignored, not because they were unimportant but because they were computationally intractable. That situation is currently fundamentally changing. In this article, we share our aha moment that occurred when looking deeper into a recent generative artificial intelligence (AI) success story.

A Wake-up Call from AI

In May 2025, Google DeepMind released a short but striking report [1] on a notable experiment: the researchers used AlphaEvolve, an ensemble of large language models (LLMs) based on Google's Gemini, to generate Python code to find solutions to various mathematical problems. Many of the benchmark problems it addressed were, from an O.R. perspective, classical nonlinear optimization problems. The codes generated by the LLMs found improving solutions for many problems without relying on classical optimization solvers or algorithms, which is an interesting and innovative use of AI.

Classical Algorithms Strike Back

Because many of these problems can be modeled directly as nonlinear optimization problems, we decided to test how far modern global optimization solvers could go. We formulated several of the AlphaEvolve problems as nonlinear optimization models and solved them using FICO® Xpress Optimization. Somewhat to our surprise, the solver, out of the box, was not only able to match the solutions found by the AlphaEvolve team but also, in some cases, found even better solutions.

One of the test problems was a variant of the classical circle packing problem: packing unit circles into a unit square (or, more generally, into a rectangle of perimeter 4) to maximize the sum of their radii. Figure 1 shows a newly discovered configuration for this problem. A key advantage of modeling-based optimization became immediately apparent. Changing the container from a square to a rectangle of fixed perimeter requires adding just one decision variable and one constraint; the remainder of the model is unchanged. By contrast, custom heuristics must often undergo a major refactoring (or retraining) when the underlying problem changes.

In joint work with Gioni Mexi and Dominik Kamp from Zuse Institute Berlin, we also obtained new solutions for the minimum distance ratio and the hexagon packing problem from the latest AlphaEvolve paper [2]. We summarized our findings in a blog post [3] and a couple of conference talks, focusing on the strengths and weaknesses of the two approaches: direct modeling and optimization with a general-purpose solver versus code generation targeting the specific instance. (The preprint of an upcoming publication is on arXiv [4].)

Based on this experience, we started to look at other packing problems in Erich Friedman's packing problem collection [5]. We settled on the problem of packing unit circles into an ellipse of minimal area. The most challenging part of the model is characterizing the inclusion of a circle within an ellipse, which we addressed using the S-lemma [6]. Figure 2 shows solutions for ellipses packed with nine and 15 circles.

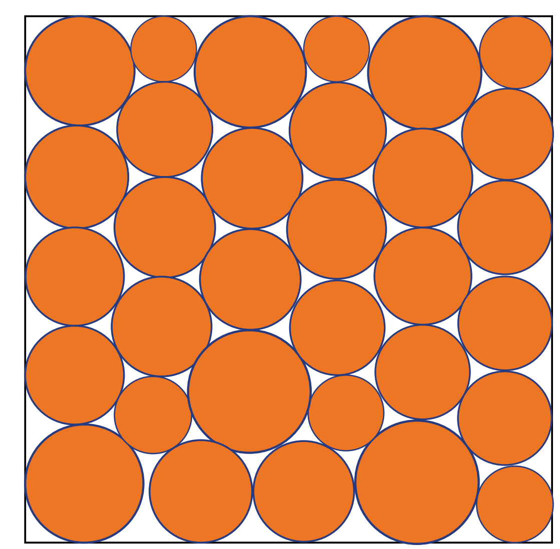

FIGURE 1: An improved circle packing configuration for 32 circles in a square.

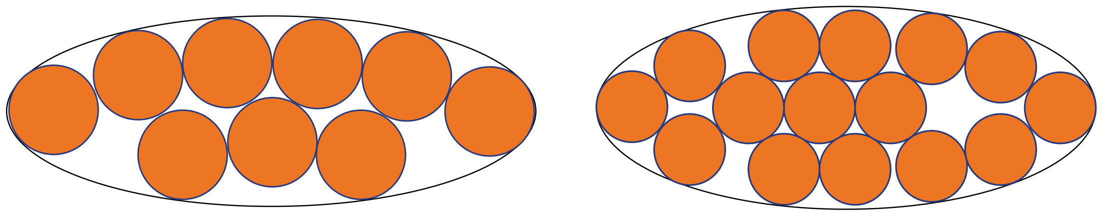

FIGURE 2: New solutions for nine circles and 15 circles, each a single unit in size, packed in an ellipse of minimum area, improving the existing solutions from 12.403 to 12.27 and from 19.824 to 19.763, respectively.

MINLP: The New MIP?

The usual way to introduce mixed integer nonlinear programming (MINLP) or global optimization is that, unlike local optimization, these algorithms guarantee convergence to a globally optimal solution. Although true, this does not acknowledge the power of modern global optimization solvers, and potential users may be deterred by the perceived difficulty of this class of problems. In our view, a better way to think about these algorithms is that they provide the same kind of guarantees and solutions that we have become accustomed to over the decades from mixed-integer programming (MIP) solvers: Along with a primal solution, they also generate a valid dual bound and report a duality gap to quantify how much more room there is for optimization. You can use a good global optimization solver the same way as you would use a mixed-integer linear optimization solver, but the constraints and objective can be nonconvex and nonlinear and can even include integralities. For practitioners, this means you can now directly model nonlinear aspects of your business problems without sacrificing reliability or decision confidence.

This difference is manifested in the comparison against AlphaEvolve: The heuristics it generated provide only primal solutions without any quality guarantees. Even though the solutions tend to be good, the user is left wanting, not knowing how much is left on the table. In business decision-making, that missing information matters. There are a few remedies for that. One is presented in an independent, open-source replication of the AlphaEvolve study, where in one case, the LLM ended up writing an optimization model and feeding it into HiGHS through SciPy, which it had access to. This can be augmented by the approach explored in a follow-up to the AlphaEvolve study [2], where the relevant literature was provided to the LLM, and it was prompted to provide mathematical insights. This could be used to add valid inequalities to the generated optimization model, strengthening the dual bound. These hybrid directions are promising yet only further reinforce the central role of mature optimization technology.

However, be careful when it comes to validating approaches that affect the dual side of a problem. On the one hand, when the LLM output is a primal heuristic, you can independently verify the solutions it provides. On the other hand, if the output is an optimization model or a custom bound-providing algorithm, then claims about the dual bound are difficult to verify externally without examining the details of the generated model or code.

The Global Optimization Software Landscape

Although BARON had long been the reference in global optimization, the last few years saw a new wave of commercial global optimization solver releases, led by FICO Xpress (2022) and followed by Gurobi (2023) and Hexaly (2024). On the open-source front, SCIP is a strong offering. These tools are multithreaded, support callbacks and solver controls familiar from MIP, and are built around outer-approximation frameworks that leverage the full power of modern MIP engines. (See [7] for more details on how such an implementation works in practice and how the separate components come together.) In addition to the proliferation of new implementations, the performance of existing ones has improved substantially: FICO Xpress Global has achieved a 4.5x speedup since its initial release and a 17.8x speedup on hard instances.

The Future Is Nonlinear

So, are we there yet? Can nonlinear global optimization solvers be used just as if they were mixed-integer linear optimization solvers? We believe the answer is a resounding yes. The technology is mature. The tools are available. The number of possible applications is enormous. It is time for practitioners to revisit their optimization models and ask a simple question: Which of my linear approximations can now be replaced by reality?

Note. *The reference list can be found online: https://doi.org/10.1287/orms.2026.01.01.*

TIMO BERTHOLD is a director at the global analytics software firm FICO, leading the Mixed-Integer Optimization R&D team, which is responsible for the development of the mixed-integer solver component in the FICO Xpress Solver. Timo is also a lecturer at Technische Universität Berlin, where he earned his Ph.D. in 2014 and his habilitation in 2022, and he is a member of the research campus MODAL. Actively involved in the academic and business community, he continues to bridge industry, research and teaching in computational optimization.

IMRE PÓLIK is a director at FICO, leading the Continuous and Nonlinear Optimization R&D team, which is responsible for the development of the global optimization solver component in the FICO Xpress Solver. An optimization solver developer since 2010, he is passionate about extending the scope and power of optimization software

ATTEND WITH YOUR TEAM TO GET THE MOST OUT OF #2026ANALYTICS

informs ANALYTICS+ CONFERENCE

More Than a Conference — A New Way to Think About Analytics

Sharpen your skills, broaden your horizons and accelerate your career at Analytics+, where data meets decision-making. Join analytics leaders, practitioners, and decision-makers at the 2026 INFORMS Analytics+ Conference, April 12–14, 2026, at the Gaylord National Resort & Convention Center in National Harbor, Maryland. The Analytics+ Conference is where attendees learn to move beyond tools and models and become a decision-centric force that drives real organizational impact.

Why Attend Analytics+?

If you are looking for a conference that helps you think differently about analytics, not just how to build models, but how to design better decisions, Analytics+ delivers. You will learn from peers who are applying analytics at scale, connect with leaders shaping the future of the field, and leave with ideas you can apply immediately in your organization.

Featured Keynote Speakers

Scott Taylor, The Data Whisperer

Kick off the conference with a dynamic keynote focused on the strategic ‘why’ behind data and analytics. Scott Taylor is known for his engaging storytelling and practical guidance on how to communicate data in ways that influence executives, build trust, and secure buy-in for analytics initiatives.

Cynthia Rudin, Duke University

Close the conference with insights from a leading voice in interpretable machine learning. Cynthia Rudin’s keynote explores how transparent, understandable models can improve trust, accountability, and real-world decision-making—especially in high-stakes environments.

Explore Focused Analytics Tracks

Choose from a rich set of tracks designed to reflect the full analytics lifecycle, from strategy to execution:

- The Analytics Ecosystem – Tools, infrastructure, governance, and technology that support successful analytics programs.
- Analytics Leadership – Strategies for leading teams, aligning business goals, and building an ethical analytics culture.
- Analytics Decision-Making – How problem formulation, methodology, and process choices shape outcomes.
- Smarter Decisions for a Better World – Inspiring applications of analytics and how it can be used to improve our world.

Training and Skill-Building Opportunities

Analytics+ offers training opportunities for every career stage. From hands-on workshops and exhibitor-led sessions to framework-based learning aligned with the INFORMS Analytics Framework, you can tailor your experience to deepen both technical and business skills. Dedicated workshops, networking opportunities, and career-focused sessions help you grow faster by learning directly from peers and industry experts.

Be Part of the Analytics Community

Beyond sessions and workshops, Analytics+ is where the analytics community connects. Network with practitioners across industries, attend special events and receptions, and experience the renowned Edelman Competition showcasing the world’s most impactful analytics applications.

Register and join us at 2026 Analytics+, where analytics meets decisions, and ideas turn into action.