Superforecasters of Tomorrow
Looking to the Future

A PhD Journey
Mistakes and Lucky Breaks

A Community of Young Researchers:
AIROYoung

ORMS TOMORROW

informs | SPRING -SUMMER | VOL.13 | NO.1





CONTENT

Letter from the Editor	3
Magazine Members	4
Superforecasters of Tomorrow — Looking to the Future	5
PhD: My mistakes, my lucky breaks, and what I tell my students now	8
Modeling a Community of Young Researchers: On the Origins of AIROYoung	11
Turning Noise into Knowledge: The Statistics Behind Diffusion Models	15
Why Supply Chains Still Need Humans in the AI Era	18
AI in Resuscitation: Smarter Help When Every Second Counts	21
OR/MS Tomorrow Mini-Poster Competition 2026	25



FROM THE EDITOR

Dear OR/MS Community Members,

"At ORMS Tomorrow, our goal is to serve students, researchers, and early-career professionals who are navigating a challenging and uncertain phase in their careers, as described by Derek Thompson in The Atlantic. Stanford economist Chad Jones notes that, despite increasingly powerful AI systems, progress ultimately depends on the weakest links—human knowledge and skills. Together, these perspectives highlight a key tension: rapid advancements demand that individuals adapt just as quickly. The question, then, is how we can actively bridge this gap.

In this issue, we bring together a range of perspectives that speak directly to this challenge. We are especially grateful to Professor Fildes for sharing his insights on what it takes to be a strong superforecaster—and the common pitfalls to avoid. Continuing in this spirit, we are delighted to feature Thiago Serra, Assistant Professor at the University of Iowa, who reflects on his journey and offers thoughtful advice for OR/MS PhD students. We also sincerely thank Veronica and Alice for sharing the story of AIRO Young, an all-student-led initiative dedicated to supporting the OR/MS community.

Alongside these contributions, our editorial team presents a diverse set of articles. Alina returns with a clear and engaging piece on the mathematics behind diffusion models such as DALL·E. This is followed by Asana's essay on the enduring importance of human judgment in increasingly AI-driven supply chain systems. We conclude the issue with thought-provoking perspectives from Marzieh and Fernando on the evolving role of AI in healthcare, along with the much-anticipated seventh installment of the Mini Poster Competition.

We hope this issue offers useful insights and helps you better navigate and make sense of this rapidly evolving landscape.

We would also like to acknowledge the following for their photographic contributions to the front matter of the issue: StockSnap for the Content page image, and Travesía_Literaria for the header photo. The cover photo was generated using ChatGPT.

Discover more about OR/MS Tomorrow by visiting our website at www.informs.org/Publications/OR-MS-Tomorrow, where you can access individual articles from this issue, explore past editions, and learn more about our editorial team. If interested in contributing an article for our Fall/Winter 2026 issue, email us at orms_tomorrow@mail.informs.org. Keep up-to-date with all our news, events, and more by following us on LinkedIn!

Happy reading!
Ronak Tiwari

OUR MEMBERS

Editor



Ronak Tiwari
WHU - Otto Beisheim School of Management

Editorial Board Members



Hasini Balasuriya
University of Louisville



Marziehsadat Rezaei
University of Louisville

Editorial Staff Writers



Alina Gorbunova
Georgia Institute of Technology

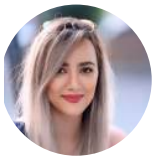


Fernando Acosta-Perez
University of Wisconsin-Madison

Editorial Staff Writers



Özgür Oskay
Eindhoven University of Technology



Asana Hosseini Dolatabad
North Carolina State University

Social Media Coordinators



Paula Penagos
University of Missouri - St. Louis



Dan Ni Lin
Texas Tech University

Production Editor



Valentina Castaneda-Torres
Georgia Institute of Technology

Advisors



Dr. Piyal Sarkar
Penn State Harrisburg



Dr. Bismark Singh
University of South Hampton (UK)



Superforecasters of Tomorrow — Looking to the Future



Robert Fildes
Distinguished Emeritus Professor, Lancaster University

How do you become a forecaster, or even better a Superforecaster? And why would you wish to take on such a role? For many years, forecasting as an academic discipline was a bit of an orphan. Were you an operations researcher, a statistician, or an econometrician? Each of these subject disciplines had their own requirements – in OR some of the premier institutions didn't even have an advanced course which included forecasting so research posts were rare.

My own training as a statistician only included stochastic processes. It was only in the 1970s that the first course books focused on forecasting started to appear, stimulated by Box and Jenkins (1970). But forecasting as a research topic was still not part of the mainstream with few university departments majoring in it. This all began to change with the emergence of the International Institute of Forecasters and the founding of two academic journals, the International Journal of Forecasting and the Journal of Forecasting which through their published papers effectively defined the discipline. A crucial component of the Institutes success was an annual international symposium welcoming all forecasters from weather to demographics but particularly focusing on business and economics. It was and is a broad methodological church with mixed parentage from the core areas but also embracing operations and psychology. Now data science with AI and Machine Learning are at its forefront. This is a formidable range for any researcher to tackle.

But there are good reasons to become a superforecaster, with shortages in academia but also in a wide range of commercial, software and public organizations.

The Forecasters' Skill Set

The early career student will have a home in one of the core areas from data science and econometrics through to behavioural science or psychology. In addition to their core curriculum, surprise, the student needs to learn the basics of forecasting: benchmark methods, comparative validation and forecasting competitions, accuracy measures, model selection. Two core forecasting texts are: Ord et al. (2017) and more technical, Hyndman and Athanasopoulos (2018). Why do I state

the obvious – because many researchers have grown up with a single methodological focus, they are often unaware of the forecasting canon. Many a machine learner still thinks that all that needs to be done is propose a new algorithm, unaware that exponential smoothing, some 70 years old, is still hard to beat when forecasting business data.

Outside academia the forecaster's job may be research focussed seeking out improvements on current practice, but it may also involve the production and presentation of regular forecasts: this latter role relies on a quite different skill set. This brings up a second area of neglect, an understanding of organizational forecasting and the psychology of forecasting. For most students the best route to understanding is a placement in an organization to discover just why it is that a fancy new algorithm is neglected in favour of current practices which depend on the expertise of the forecasters in the organization. Knowing about different decision contexts and the data available will also stimulate ideas for a research topic, with on-line data and novel logistics systems as recent examples.

In short, a deep understanding of the core disciplines employed in forecasting is not enough, the novice superforecaster has to delve into the practicalities of the subject.

What Pitfalls to Avoid?

As preparation for this wide variety of roles, the most important early pitfalls to avoid are (i) specializing too earlier, and (ii) not mastering one of the core areas (more would be too ambitious). Forecasting is not an abstract discipline but depends on its organizational context to be effective. Consequently, the novice needs a broad understanding of the context in which their specialist topic will prove valuable. The context determines both the data availability, the methods to explore and their constraints, and the loss functions, both theoretical and organizational that need to be adopted.

What makes a Superforecaster?

Superforecasting was first identified by Tetlock and colleagues (2016). The book of the same name drew lessons from a research study of around 80K longer-term political forecasts. They found that the more accurate forecasters shared characteristics in common, they were 'cautious, humble, open minded, analytical – and good with numbers!'. For readers of this piece, we can take some of these requirements for granted. But what about open-mindedness: nowadays every problem, some believe, can be solved through ML methods, forgetting a key lesson from past forecasting research that simple methods often outperform their much more complicated competitors and, in addition, are much more acceptable to users. So, the most serious pitfall is also the most attractive to the student forecaster, to immerse themselves in the technology of a particular set of techniques, neglecting the fact that the broader discipline offers lessons that need to be absorbed. The other side of this coin is that for a job in organizations such as Google or Amazon, the forecaster will need to be thoroughly knowledgeable about the technical details and opportunities (probably in ML). So early on in a research training program, even to the extent of choosing a place of study, the student faces a dilemma: do they choose the deep dive into the technical or aim for a forecasting career more broadly? It requires the most difficult subjective judgmental forecast to decide which of the two will prove most satisfying. In making my own choice I realized I preferred breadth rather than depth (despite a highly specialized PhD) and that pointed the way to work in a business school committed to engaging with organizational problems rather than, for example, a data science department. And remember, it's very hard for a university department of whatever colour to compete with large technical departments of the research-oriented technology companies. As an example, the winner of a major 'forecasting competition' worked for Uber (Smyl as discussed in Makridakis et al. (2018)).

A stark choice then: Do you choose breath or depth?

The Future of Forecasters and Forecasting

Since the founding of the Institute in 1981, forecasting has seen dramatic changes, both in research communities but also in practice. Dawes et al. (1994) looked to the future of forecasting 30 years ago. They were remarkably prescient with three themes predominant then and now:

1. The synergies (and problems) arising from the forecaster's knowledge and the forecasting algorithm in use.
2. The combining of and model selection from a range of algorithms.
3. The importance of context in forecasting which encompasses the users, the decisions which depend on the forecasts and the data.

The last thirty years have seen major changes in the forecasting discipline, particularly with the advent of ML/AI methods. Some assume that there will emerge a ML/AI method which will render these concerns irrelevant. As yet there is no sign of such an all-encompassing solution. When the M5 forecasting competition was launched with its focus on ML methods, the winner turned out to be a combination of statistical and ML. More crucially, the majority of ML methods failed to outperform critical statistical benchmarks. While advances since then have been rapid, there is no evidence that a general 'winner' will be found including Large Language Models. I personally don't believe there could be for the reasons given above. Ask yourself the question, can a general method designed for many different circumstances, however clever, prove most accurate despite the singular characteristics of the particular problem being faced?

So these three areas remain topical and relevant today. But with novel development of new ML methods, now available to any would-be superforecaster, in particular Large Language Models, amplified by the ability to process vast amounts of data, new challenges will emerge and new pitfalls, not least the trap of getting caught up in the technology whilst ignoring the problem.

The role of humans in the forecasting process is a concept that is central to trustworthy AI, forecasting. Superforecasters of the future will have their work cut out for them to find ways to streamline both the information flows, that are increasingly more complex, and the value from the decisions supported by the forecasts. This gives tremendous opportunities for the forecasting field, both in innovation and impact. As a final note, these themes are stimulating and fun, not least because they require the researcher to consider a range of perspectives, and that it is perhaps the best training for a 'superforecaster.'

Robert Fildes is Distinguished Emeritus Professor of Management Science in the Management School, Lancaster University UK, Founding Director of the Lancaster Centre for Marketing Analytics and Forecasting and recipient of the Beale Medal, the UK Operational Research Society's highest accolade. He was co-founder in 1981 of the Journal of Forecasting and in I985 of the International Journal of Forecasting. His current research interests are concerned with the comparative evaluation of different forecasting methods, the role of the forecaster, and the implementation of improved forecasting procedures and systems in organizations. He is well-versed in what superforecasters can be!

References

Box, G.E.P., Jenkins, G.M., 1970. Time Series Analysis: Forecasting and Control. Holden-Day, San Francisco.

Dawes, R., Fildes, R., Lawrence, M., Ord, K., 1994. The past and the future of forecasting research. International Journal of Forecasting 10, 151–159.

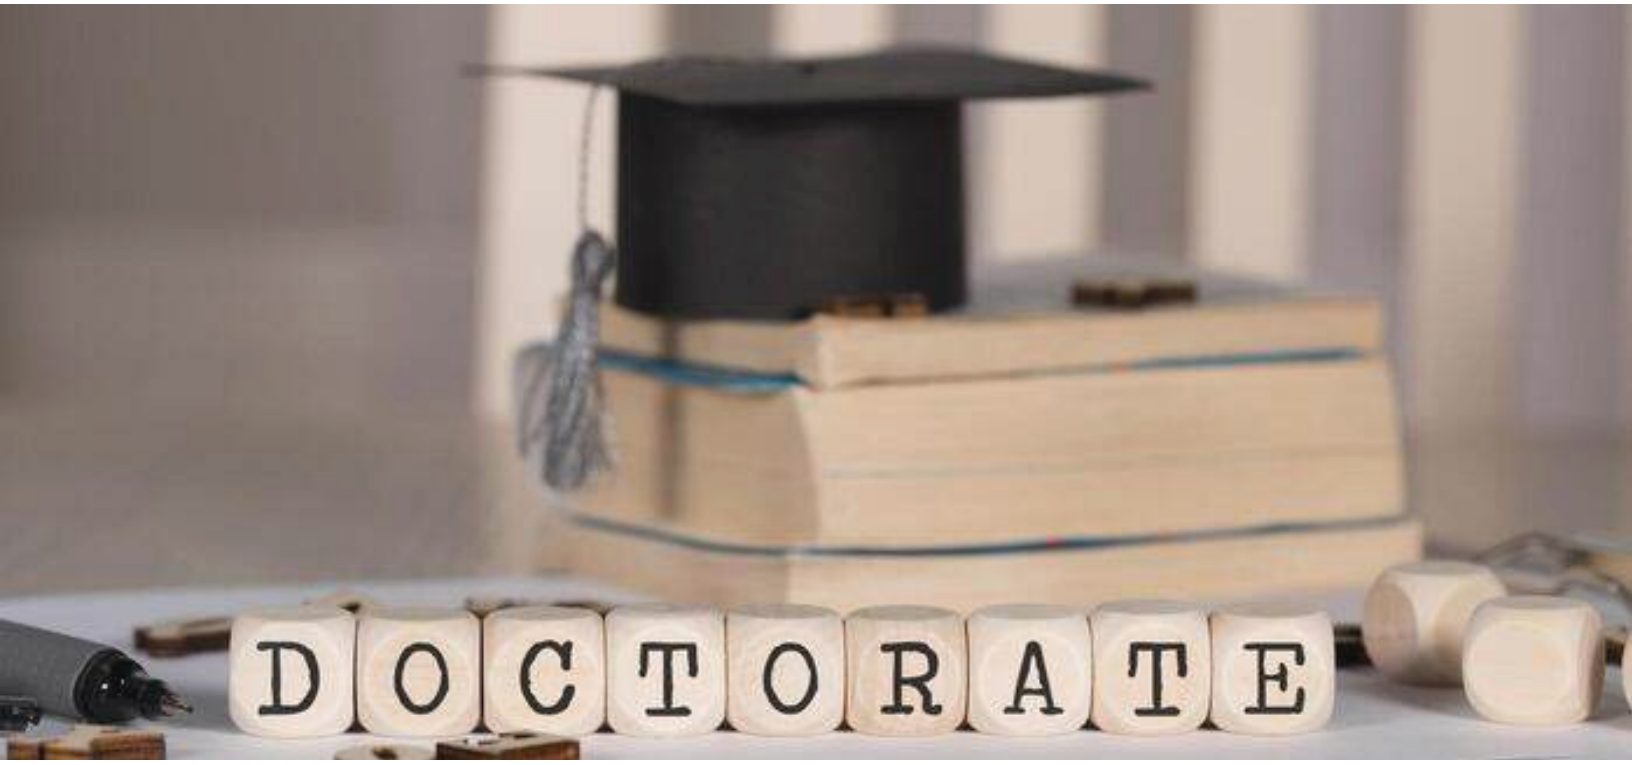
Hyndman, R.J., Athanasopoulos, G., 2018. Forecasting: Principles and Practice. OTexts. URL: https://otexts.com/fpp2/.

Makridakis, S., Spiliotis, E., Assimakopoulos, V., 2018. The m4 competition: Results, findings, conclusion and way forward. International Journal of Forecasting 34, 802–808. doi:10.1016/j.ijforecast.2018.06.001.

Ord, J.K., Fildes, R., Kourentzes, N., 2017. Principles of Business Forecasting. 2 ed., Wessex Press.

Acknowledgements: *We would like to thank Kara Combs for taking time to review this article. Photo credit goes to Jakub Zerdicki for the header photo and Stephen Dawson for the footer photo.*





PhD: My mistakes, my lucky breaks, and what I tell my students now



Thiago Serra
Assistant Professor, University of Iowa

Almost 8 years after defending my PhD, I am still figuring out what is the right way to pursue one. In retrospect, I was lucky to stay in academia given how I completed mine. I have also seen other students doing the right things, certainly more often than I did only to run out of luck at critical points. But now I must tell my own students something and, more importantly, direct them in a way that would help them succeed and hedge against that luck component. So let's start with what I think I should have done, but which I didn't.

The Thesis

When I start working with a PhD student, I tell them to think about where they want to be in 5 years, and then we work backwards from there. If they want to stay in academia, then the best outcome would be landing a tenure-track position that starts right after their PhD (or with a gap year, or two, to go for a postdoc in-between; but let us not chase unicorns here). In operations research, landing that position means being ready for the academic job market by early fall of their 5th year. But what does "being ready" even mean?

With a pool of hundreds of candidates for a tenure-track job, the first screening is often just a quick glance at their CVs. Are the papers there? Where were they submitted? Did they receive a first response at least? Thinking in terms of journal papers, we can work with the assumption of three good pieces of work, which look great on CVs. For instance, consider a scenario such as the following:

1. One paper published in a reputable journal, though not necessarily among the most prestigious venues, perhaps with substantial contribution from collaborators.
2. A second paper in advanced stage of review at a journal that is the best technical fit for the work, with the student having contributed significantly to all aspects.
3. A third paper under review and circulating as a working paper, which is largely the student's own idea and demonstrates growing independence.

Every paper has a purpose. With the first paper, the student has experienced the full cycle of a research project, including resubmissions to the same or different journals, as well as dealing with reviews that were constructive and others that were unfair to their work. Learning that rejection is part of the parcel. With the second paper, the student signals the ability to produce strong work that is mostly their own. This paper is the most important because there is still enough time to reasonably

expect to have a journal response prior to entering the job market. With the third paper, the student shows that they have established a research pipeline and that more work will follow in due time. By that stage, the student is ready to take off. So let's talk about that backwards plan again, now in more detail:

1. To land a tenure-track job in five years, the student should ideally be ready for the market in four years.
2. But to have that second paper at least in major review, assuming an expected review time of six months and the occasional setback of rejection, whether for justified constructive feedback or unfair reviews, the paper likely needs to be ready within three years.

With that in mind, I came up with the following timeline:

1. Semesters 1 to 3: the student is primarily taking courses, while gradually ramping up research activities by contributing to existing projects and carrying out specific tasks. This is when the first paper is produced.
2. Semesters 4 to 6: having completed the coursework, the student becomes the main contributor of a project while relying on guidance from the advisor to identify a direction that is both productive and publishable. This leads to the second paper.
3. Semesters 7 to 9: if everything goes well, the student begins proposing the direction for the next step. This becomes the third paper, which will either appear on the CV as another submission (if completed by the end of semester 8) or as a working paper.
4. Semester 10: thesis defense, campus interviews, all the nail-biting that happens between interviews and hearing back, and pursuing a back-up plan.

In practice, however, projects defy our initial expectations of how long they would last. Sometimes we also get stuck, bored, or both. In addition, it is difficult to confine ourselves to a single line of work, especially when we are not entirely sure we want to commit to it in the long run. As a result, most students eventually want to have at least a second project going on. When my students reach this point, I have a second conversation with them. This is when I tell what I actually did, which was admittedly much messier and riskier.

The Antithesis

In the second conversation with my students, I show them what happened in my case. First, I go to my Google Scholar profile, sort the papers by year, and scroll down to the period in which I was a PhD student. There is some relevant work in 2012, the year before I started: my master's thesis and a paper at the CP conference, which grew out of that thesis work. Then nothing else relevant until 2018, the year I defended: my PhD thesis and conference papers at ICAPS and ICML — and the ICML paper is not even part of my thesis work.

So what happened? I chose to work with two advisors on different topics, working on a main project with each of them as well as a spin-off project driven by my own questions about those projects. Then I pursued summer internships after my third and fourth years. Heading to my final year, I also started something entirely new on the side. That probably accounts for 80% of my PhD effort. I will skip the other 20%, which was even less focused.

For my PhD defense, all that I could show as evidence of external validation was the ICAPS paper from my fourth-year internship and a major revision at *Mathematical Programming* for one of my PhD projects, which had only been submitted shortly before I started my fifth year. I went to the job market a year later, already out of the PhD, and the only additional evidence was the ICML paper from that fifth-year exploration. Nevertheless, that was enough to get me a few interviews and eventually land a tenure-track position in 2019.

Continuing that conversation with my students, I go back to my Google Scholar profile and roll up to 2020, which I like to call the "miraculous year": both projects with my advisors became journal papers, one spin-off also became a journal paper and the other became a conference paper. That fifth-year exploration yielded another two conference papers. Although none of that arrived in time to help me on the academic job market, the publication pattern becomes somewhat stable afterward — maybe I have learned my lesson!

In other words, I was very lucky to land a job before most of that materialized in 2020: every department that considered hiring me took a considerable leap of faith. They had probably seen similar cases before, and I appreciate their willingness to take the risk with me and others who similarly spread themselves too thin during their PhD years. However, the journey could have been much safer if I had completed some of that work earlier. But this entire discussion misses the main point: why is it that we do what we do?

The Synthesis

> Every time I send a student away with objective steps for moving a project forward and the student comes back with questions and side explorations (and, of course, none of those objective steps implemented), I am reminded of why we do what we do: we are in this for the learning and the discovery. Indeed, if we do not pay attention, we could miss that a secondary issue to what we are exploring is actually a far more important problem to solve. In fact, it could be the best thing that we are ever do — but only if we notice it!

From that perspective, my choices were not so erratic as they may seem:

1. I chose to work with two advisors, John Hooker and Egon Balas, because I admired their work and wanted to learn from both. And I did learn a lot!

2. By pursuing those spin-offs, I learned to explore new ideas, had opportunities to fail, and developed a better sense of the kinds of problems I can — and should — tackle.

3. I chose to do internships because I wanted to continue thinking about how to connect my work with practice. That definitely helped ground my focus.

4. I returned to the same place for a second internship because I had a great manager, Arvind Raghunathan. He also left the door open for me to return as a one-year visitor after my PhD, which became a lifeline for my academic aspirations.

5. Through Arvind, I also met his friend Srikumar Ramalingam, who turned out to have a taste for problems that was very similar to mine. He eventually sent me a paper about deep learning, a topic which I kept ignoring because it seemed disconnected from my background. That paper sparked the fifth-year exploration with my PhD colleague and friend Christian Tjandraatmadja, which ultimately became the major theme of my research going forward.

6. Through Srikumar, I met his then student Xin Yu, who would lead the computational side of many of my projects involving deep learning in the following years. Together, we were able to do a lot more than we would have done on our own.

We do what we do because we love it. And the best way to do it well is by finding like-minded peers who become part of our success network. But without something tangible to show regularly as our contribution, we may put our long-term academic prospects in jeopardy. Hence, we need to have a plan — even if only as a benchmark against what we end up doing along the way.

Research is like cooking with too many pans on a two-burner stove top: leave a pan unattended for too long and its contents are lost; switch pans too frequently and nothing gets done; but stick with the same pans forever and you risk cooking something that will never become edible. So go ahead and flirt with another project if you must, but remember that there is only so much that you can cook before the guests arrive. And if the food is not ready, ill-mannered guests may take a bite out of you!

Note: I would like to thank Sabrina Harris, my partner in life, for many comments on how to improve this essay — and for bearing with me through the ups and downs of this journey!

Thiago Serra is an assistant professor of business analytics at the University of Iowa. His scholarship focuses on the theory, practice, and integration of machine learning and mathematical optimization. Previously, he was an assistant professor at Bucknell University, a visiting research scientist at Mitsubishi Electric Research Labs, and an operations research analyst at Petrobras. He has a Ph.D. in operations research from Carnegie Mellon University, a master's degree in computer science from the University of Sao Paulo (USP), and a computer engineering degree from the University of Campinas (Unicamp). He serves as an associate editor for the journals INFORMS Journal on Computing and INFORMS Journal on Data Science.

Acknowledgements: *We would like to thank Pavithra Sripathanallur Murali for taking the time to review this article. Photo credit goes to Ecaterina Tolicova for the header photo and to Christian Lendl for the footer photo.*





Modeling a Community of Young Researchers: On the Origins of AIROYoung



Veronica Dal Sasso
Siemens Mobility, Italy

Alice Raffaele
University of Padova, Italy



AIRO Young (AY) is the youth section of the Italian Association of Operations Research (AIRO), dedicated to fostering collaboration among students and early-career researchers in Operations Research (OR). Through annual workshops, networking events, and career development initiatives, AY creates opportunities for young researchers to advance their careers, expand their professional networks, and bridge the gap between academic and industry opportunities in the OR field.

AY has celebrated its first ten years. In February 2026, the tenth edition of the AY annual workshop (10AYW) took place in Padua, Italy. **More than 120 participants** from Italy and abroad attended a **five-day event** whose main theme was "*Optimization between determinism and uncertainty*", and included several OR topics, such as routing and transportation, scheduling and timetabling, railway optimization, but also advances in mathematical programming, interplay with machine learning, and many more.

The event was composed of a one-day Ph.D. school, with four masterclasses by Marco Di Summa, Matteo Fischetti, Francesco Rinaldi, and Domenico Salvagnin (all from the University of Padua), and a four-day workshop, with **58 presentations, five sponsor talks, and four plenary lectures** by Martine Labbé (Université Libre de Bruxelles), Stefan Røpke (DTU), Wolfram Wiesemann (Imperial College Business School), and Hande Yaman Paternotte (KU Leuven). Nonetheless, there have been several opportunities to discuss, connect, and expand one's network. Moreover, there has been a **special roundtable session** with AY's founding members and former board members to commemorate the milestone anniversary, reflect on AY's evolution, discuss future steps, and share emotions while watching a celebratory video. But what is the story behind the origins and the formation of this youth chapter?

The First AY Decision Makers

Just as mathematical modeling requires optimizers, the soon-to-be AY was shaped by four individuals, depicted in Figure 1, who identified a problem and decided to tackle it with their best tools.

The first AY promoter was **Alberto Santini**. In 2015, Alberto was pursuing his Ph.D. at the University of Bologna. Among his activities, he was also organizing a series of webinars for



Figure 1: AY founders at the 4th AY Workshop in Bozen, Italy, in February 2020. From the left, Veronica Dal Sasso, Martina Fischetti, Lavinia Amorosi, and Alberto Santini.

young researchers that he streamed online. Nowadays, this is quite common. However, at that time, it was actually a challenge, as tools like Microsoft Teams or Zoom were either not even developed or were still in their early stages; for sure, they were not everyone's bread and butter. Nevertheless, with the help of the European Association of Operations Research (EURO), Alberto was able to buy some necessary equipment, including a microphone. These webinars were highly successful, underscoring a concrete need in our community for events for young researchers.

Among the attendees of these EURO young OR researchers webinars was **Veronica Dal Sasso**, who began her Ph.D. in Padova in 2014. At that time, she was the only PhD student in Operations Research in Padova, not only in the Mathematics Department but also in the Engineering Department. The group in Padova was rather small, and peers were in short supply. This was the main motivation that pushed Veronica not only to present her work during the webinars, but also to attend the other webinars in person in Bologna. On this occasion, she met Alberto.

In October 2015, **Lavinia Amorosi** went to the Optimization @Work Ph.D. School at ZIB, in Berlin. There she met Alberto and Veronica. It was the beginning of a long-term friendship (including hosting Veronica during the first weeks of her future move to Rome).

Last but not least, **Martina Fischetti** met Alberto Santini at the Technical University of Denmark, where she was pursuing her Ph.D. and he was a visiting Ph.D. student. Some time later, in January 2016, she also met Veronica and Lavinia at the Winter School on Network Optimization in Estoril.

Formulating a Community Problem

Therefore, the idea of a chapter dedicated to young researchers within AIRO stemmed from these four people coming together at different academic events.

Indeed, by sharing how much they enjoyed meeting the same and new people at different conferences and Ph.D. schools, Alberto, Lavinia, Martina, and Veronica also realized that Ph.D. students and early-career researchers often lack their own funding or may not even be aware of such events. **They identified a community problem**: the need to promote initiative to **foster networking and collaboration** among young

researchers who often face difficulties during their doctoral paths, and thus to **share experiences** with others in the same situation.

These aspects soon became their objective function to be maximized, by taking into account their own constraints, such as being just a small group of scholars. But as for the EURO young OR researchers webinars, they were supported in taking their first steps, that is, making their first decisions about what to do in practice to move from a nice idea raised while chit-chatting at some dinners to effectively creating an opportunity and engaging people to participate. A mention is needed here: Alberto's initial proactivity, and later the development of the AY chapter, were also strongly supported by his supervisor, Daniele Vigo (University of Bologna), who helped bring the idea of a young OR community to life through valuable suggestions and assistance in securing funding.



Figure 2: 1AYW and 10AYW group pictures.

The Root Node: The 0^{th} AY Workshop

Not long after the Winter School in Estoril, Alberto sent an email to Lavinia and Veronica inviting them to organize a pilot event for AY. As this idea of community met a well-sought-after need, he received only positive responses. The first-ever gathering of *AIROYoungers* – as we like to call AY members – was at the AIRO 46th Annual Conference in Trieste in September 2016, where Alberto suggested that Lavinia host a poster session dedicated to young researchers and organized by young researchers. There, Martina also joined what would become the founders' team: AY was formally born.

The first official AY workshop (1AYW) was held later, in February 2017, at Sapienza University of Rome, on the topic "Emerging optimization problems in complex networks". On the official 1AYW website, the preamble said, about AY: "*Its aim is to foster collaboration between students and early-career researchers interested in the field of OR, and to provide them with new opportunities to advance their career and expand their network. It also strives to connect the demand and the offer in the OR job market, both in academia and in the industry.*" By relaxing everything unnecessary, this was and still is the **core problem** the whole AY chapter has been addressing ever since.

Branching Since 2017

At the 1AYW, which lasted two days, there were 35 participants, 13 presentations, and two plenary lectures, delivered by Justo Puerto (University of Sevilla) and Maria Grazia Scutellà (University of Pisa). Figure 2 shows the 1AYW and 10AYW group pictures. Comparing these as well as the numbers of the 1AYW with those of the 10AYW at the beginning, **we can appreciate how much AY has grown** over the past ten years.

Besides regularly organizing its annual workshop, AY has **expanded in several directions**: from creating a mailing list (currently with 325 members) to establishing a presence on major social media platforms; from organizing special sessions at Italian and international conferences to developing the *AY Dissertation Award*; from running webinars in collaboration with the LOG group of the Universidade Federal da Paraíba to overseeing the publication of special issues and volumes; and many more.

The growing involvement of young researchers soon led Alberto, Lavinia, Martina, and Veronica to establish an **official board**, initially composed of four members: two joint coordinators (Lavinia and Martina), one treasurer (Valentina Morandi), and one social media manager (Serena Fugaro), who was later joined by another social media manager/web curator (Alice Raffaele). Starting in 2021, regular elections have been held every three years thereafter. As participation and activities continued to grow, the second board (2021–2024) was **expanded to six members**: two joint coordinators (Lorenzo Peirano and Alice Raffaele), one treasurer (Veronica Dal Sasso), two social media managers (Serena Fugaro and Giusy Macrina), and one web curator (Michele Barbato). The current board, elected in 2024 and active until 2027, maintains the same structure and includes Matteo Cosmi and Andrea Mancuso as joint coordinators, Martina Doneda as treasurer, Sara Stoia and Marta L. Tessitore as social media managers, and Alessandro Druetto as web curator.

Who will form the next AY board? How will AY change, adapt, and transform itself to embrace future generations? We do not yet know. Nevertheless, even though some of us are now part of the AIRO *Forever* Young chapter, we look forward to finding out the new – hopefully infinite, for once – branches of enthusiasm that this youth chapter will explore.

References

AIRO Young, 2017. Young Researchers' Chapter of AIRO. http://www.airoyoung.airo.org. Accessed: 2026-04-28.

Amorosi, L., Fischetti, M., 2020. The New Generation of OR Enthusiasts: the AIRO Young Experiment, in: SN Operations Research Forum, Springer. p. 15.

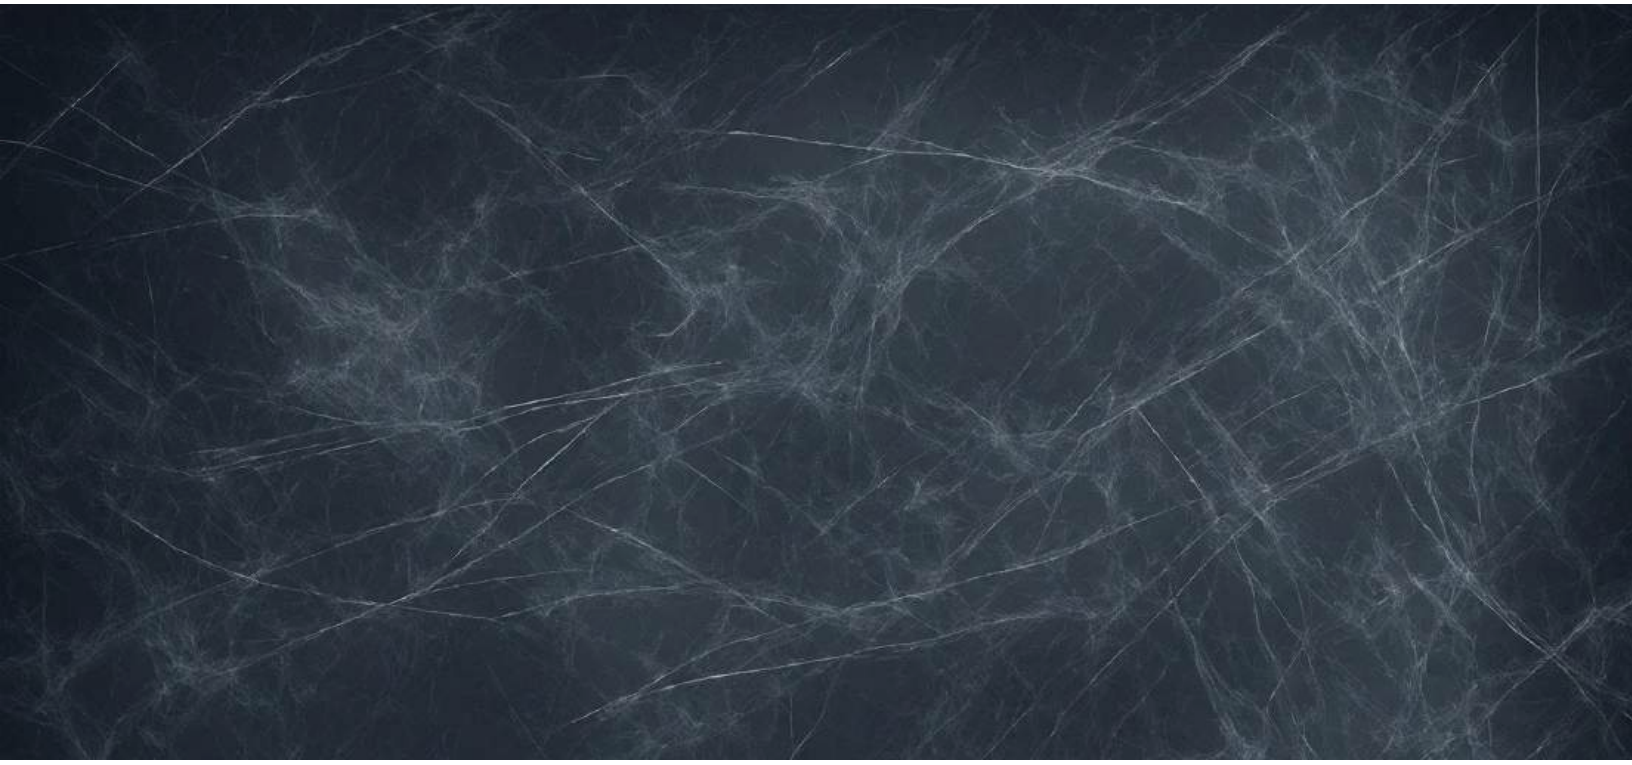
Cavagnini, R., Dal Sasso, V., Morandi, V., Raffaele, A., 2020. How to Organize an OR Workshop: the AIROYoung Experience, in: SN Operations Research Forum, Springer. p. 34.

Italian Association of Operations Research, 1961. http://www.airo.org/. Accessed: 2026-04-28.

Raffaele, A., 2026. 10 Years of AIRO Young. https://drive.google.com/file/d/1YdJoAPOO-yHCKxOzHLN5xqU2n7k1kIZN/view. Video file. Accessed: 2026-04-28.

Acknowledgements: We are grateful to Lavinia Amorosi, Martina Fischetti, and Alberto Santini for reviewing this column, helping us recall and rediscover how AY was born, and offering invaluable advice. A special thank you goes to all the people, participants, and sponsors who have contributed to the AY community over all these years. Photo credit goes to Broesis for the header photo and Florendia for the footer.





Turning Noise into Knowledge: The Statistics Behind Diffusion Models



Alina Gorbunova
Georgia Institute of Technology

Diffusion models *such as OpenAI's DALL·E and Stability AI's Stable Diffusion have dazzled users with their ability to output everything from silly cat images to sophisticated diagrams of mechanical parts. Pulling back the curtain on these models reveals an elegant and sophisticated statistical framework that at its core is a combination of stochastic processes, variational inference, and score-based estimation: the same language that we as researchers in operations research speak.*

A diffusion model learns in two phases: a forward process that slowly adds bits of Gaussian noise over T time steps to a data sample, such as an image, and a reverse process that undoes this added noise to recover the original data sample. The goal of this forward and backward process is to learn how to convert a Gaussian distribution to a different target (data) distribution.

Going Forwards and Backwards

The mathematical foundations for diffusion models were first presented by Sohl-Dickstein et al. (2015). During the forward process a small Gaussian perturbation is added to the sample from the previous time step so that x_t depends only on x_{t-1}. Does this sound like a familiar property? If you thought "Markov Property", you are right on the money. A diffusion model's forward process $q(\mathbf{x}_{1:T}|\mathbf{x}_0)$ can be defined as a Markov chain that gradually adds Gaussian noise:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) \coloneqq \prod_{t=1}^{T} q(\mathbf{x}_t|\mathbf{x}_{t-1}) \tag{1}$$

where

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) \coloneqq \mathcal{N}\left(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}\right) \tag{2}$$

The term $\{\beta_t\}_{t=1}^{T}$ is the diffusion rate, also called the variance scheduler. It determines how much noise is added at each time step to the previous sample to get the next sample, and it increases as t increases. There are several ways to determine how β_t increases from a linear to an exponential schedule, with each schedule having a different impact on the performance and precision of the model. Examples of different variance schedulers and their impacts can be found in Table 1. Guo et al. (2025) provides a more comprehensive list of schedulers.

Schedule	$\bar{\alpha}_t$ Behavior	Practical Effect
Linear	Quick drop	Wastes capacity at high t
Cosine	Slow, smooth decline	More informative time steps, better quality than linear
Learnable	Data-adaptive, optimized jointly	Best flexibility, higher training complexity

Table 1: Different Variance Schedulers and Their Impacts

A key analytical convenience is that at any time step t, the joint distribution can be computed in closed form. This is called reparameterization and it defines $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$ to write

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) \coloneqq \mathcal{N}\left(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_{t-1}, (1-\bar{\alpha}_t)\mathbf{I}\right) \tag{3}$$

When $t = T$, the sample is pure Gaussian noise such that $\mathbf{x}_T = \mathcal{N}(\mathbf{0}, \mathbf{I})$. Then the reverse process occurs. During this process, a small amount of noise is removed from the sample with the goal of returning to the original data sample $\mathbf{x}_0$ using a Markov chain defined as

$$p_\theta(\mathbf{x}_{0:T}) \coloneqq p(\mathbf{x}_T) \prod_{t=1}^{T} p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \tag{4}$$

where

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \coloneqq \mathcal{N}\left(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \mathbf{\Sigma}_\theta(\mathbf{x}_t, t)\right) \tag{5}$$

The parameters θ are typically a neural network.

Training the Model

To successfully train a diffusion model, the reverse Markov chain needs to match the actual reversal of the forward Markov chain. In other words, the parameters θ need to be learned such that the joint distribution of the reverse process closely approximates the forward process. And what metric can be used to measure the similarity between two distributions? If Kullback-Leibler (KL) divergence popped into your mind, you are on the right track!

The training loss combines minimizing the KL divergence with maximizing the log-likelihood of predicting the correct data $p_\theta(\mathbf{x}_0)$, which can also be written as minimizing the negative log-likelihood, as

$$\mathbb{E}_q \left[-\log(p_\theta(\mathbf{x}_T)) - \sum_{t \geq 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{q(\mathbf{x}_t|\mathbf{x}_{t-1})} \right] \tag{6}$$

The $\log(p_\theta(\mathbf{x}_T))$ is difficult to calculate, so the variational lower bound (as known as the Evidence Lower Bound, or ELBO) is used to approximate it. Ho et al. (2020) showed that the ELBO can be written as a sum of KL divergences between each Gaussian distribution. This simplification means the loss function can be rewritten as a denoising loss function such that given noisy sample $\mathbf{x}_t$, predict the noise ϵ that needs to be added:

$$\mathbb{E}_{t,\mathbf{x}_0,\epsilon}\left[||\epsilon - \epsilon_\theta(\mathbf{x}_t, t)||^2\right] \tag{7}$$

This is a mean squared error loss, a simple equation at the heart of diffusion model's training. Using this simple loss function, helps to prioritize more challenging tasks throughout the training, such as removing large amounts of noise at higher values of t.

A diffusion model's entire training objective reduces to a single mean squared error loss: given a noisy sample $\mathbf{x}_t$, predict the noise $\boldsymbol{\epsilon}$ that was added. This deceptively simple equation is the statistical engine at the heart of modern generative AI and it is one that any researcher fluent in probability and optimization is already well-equipped to understand, extend, and apply.

Sampling and Inference

Once a diffusion model is trained, it can be used to generate new samples by running the reverse process from scratch. In other words, the model starts from pure Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and iteratively denoises across T time steps using the learned reverse transition $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$. This is called ancestral sampling and mirrors the reverse Markov chain defined in Equations 4 and 5.

At each time step, the model predicts the noise $\epsilon_\theta(\mathbf{x}_t, t)$ present in the current sample and uses it to estimate the mean $\mu_\theta(\mathbf{x}_t, t)$ in order to draw the next sample $\mathbf{x}_{t-1}$. The mean can be written as

$$\mu_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(\mathbf{x}_t, t)\right) \tag{8}$$

Using this equation, the next sample $\mathbf{x}_{t-1}$ can be written as the mean plus some noise level σ_t injected at time step t

$$\mathbf{x}_{t-1} = \mu_\theta(\mathbf{x}_t, t) + \sigma_t \mathbf{z} \tag{9}$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. It's this $\mathbf{z}$ noise injection that makes the diffusion model stochastic. Two new samples from the same $\mathbf{x}_T$ will produce different results. Repeating this noise injection over T time steps will produce a new sample $\mathbf{x}_0$. This process is illustrated in Figure 1.

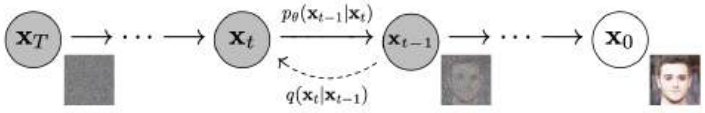


Figure 1: An example of a new sample $\mathbf{x}_0$ generated by a diffusion model Ho et al. (2020)

A critical drawback is the cost. If $T = 1000$, generating one new sample requires 1000 sequential neural network computations, creating a significant bottleneck. This is addressed by Song et al. (2020) with the introduction of denoising diffusion implicit models (DDIM). This model defines a subsequence of time steps $\tau_1 < \tau_2 < \ldots < \tau_S$ with $S << T$ to construct a deterministic rule to skip the intermediate steps between each τ_s. DDIM

also does not require retraining the model, it uses a different solver applied to the same trained network. This reduces the number of neural network computations from $T = 1000$ to as few as 10-50 with a minimal loss in quality, but with a 20-100x speedup.

As these models find their way into applications from molecular design to demand forecasting to combinatorial optimization, researchers equipped with this statistical vocabulary will be well-positioned — not just to use these tools, but to improve them.

Conclusion

Diffusion models are one of the most statistically principled frameworks in modern generative AI, and their foundations read like a syllabus from an OR/IE graduate program: Markov chains, variational inference, KL divergence, MSE estimation, and stochastic processes. From the elegant reparameterization trick that makes training tractable, to the ELBO decomposition that reduces a complex likelihood problem to a simple noise prediction task, to the schedule design problem that is at its heart a constrained optimization, every component of a diffusion model has a mathematical story worth understanding.

Diffusion models are an open research frontier with problems that speak directly to the operations research toolkit. Schedule design is a constrained optimization problem. Sampling efficiency is a resource allocation problem. Constrained generation during the reverse process remains largely unsolved. For graduate students looking for a foothold in machine learning research, diffusion models offer something rare: a state-of-the-art method where the most important open questions are not about network architecture, but about probability, optimization, and mathematical rigor.

References

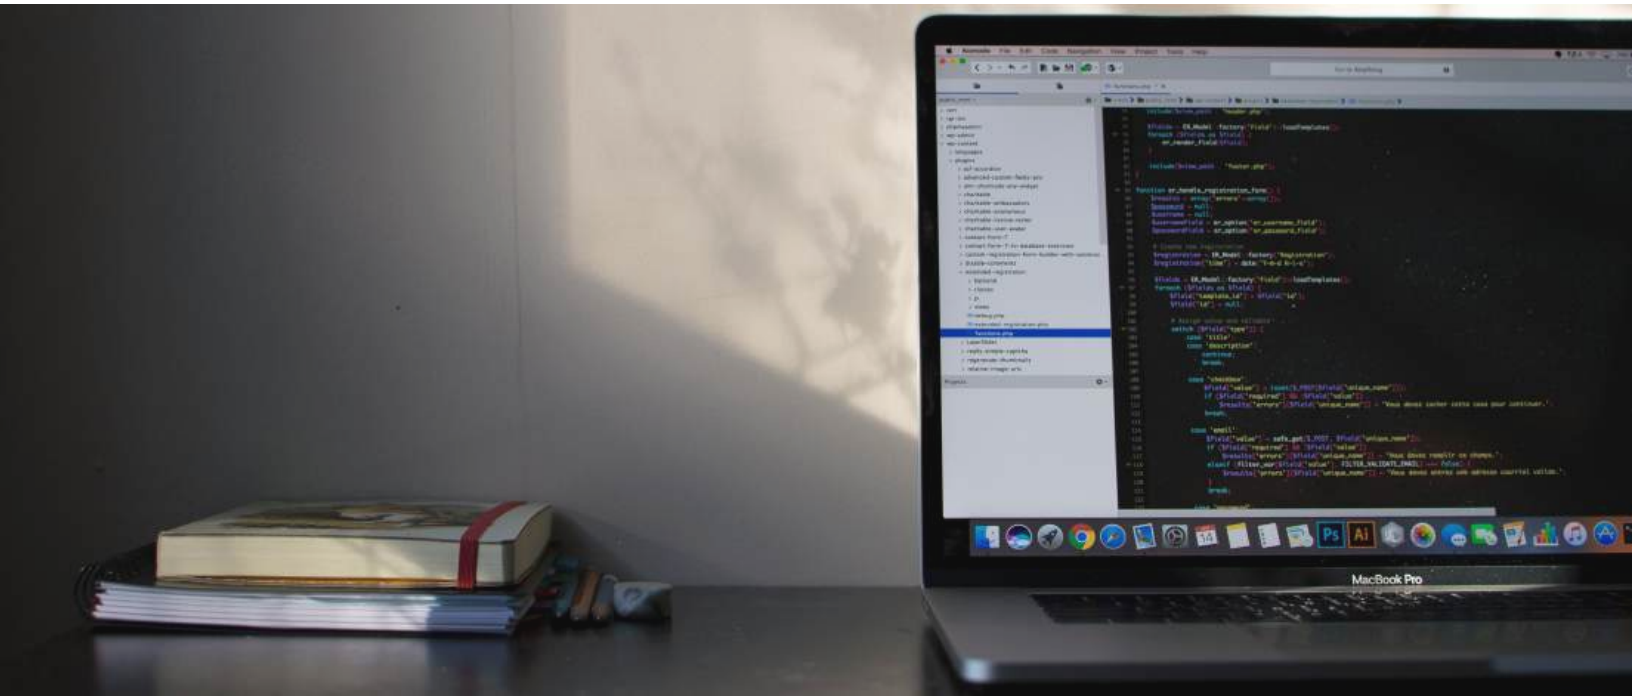
Guo, Z., Lang, J., Huang, S., Gao, Y., Ding, X., 2025. A comprehensive review on noise control of diffusion model. arXiv preprint arXiv:2502.04669 .

Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. Advances in neural information processing systems 33, 6840–6851.

Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S., 2015. Deep unsupervised learning using nonequilibrium thermodynamics, in: International conference on machine learning, pmlr. pp. 2256–2265.

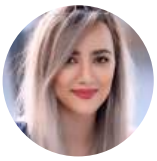
Song, J., Meng, C., Ermon, S., 2020. Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 .

Acknowledgements: *We would like to thank Ronak Tiwari for taking time to review this article. Photo credit goes to TheDigitalArtist for the header photo and Emile Perron for the footer photo.*





Why Supply Chains Still Need Humans in the AI Era



Asana Hosseini Dolatabad
North Carolina State University

Artificial intelligence is rapidly transforming supply chain management, automating decisions that once required significant human expertise. Yet as algorithms become more capable, a critical question emerges: can optimization alone build supply chains that are sustainable, resilient, and socially responsible? This article argues that human judgment remains indispensable in the AI era. Drawing on recent operations management research, it examines three key areas where humans continue to add irreplaceable value — defining sustainability objectives that algorithms cannot set on their own, providing contextual understanding that historical data fails to capture, and maintaining accountability in an environment where algorithmic authority is quietly growing. Rather than viewing AI as a replacement for human decision-making, this article suggests that the future of supply chain management lies in designing effective collaboration between humans and intelligent systems.

Artificial intelligence is changing supply chains at an incredible pace. Tasks that once relied heavily on managers' experience and intuition—such as forecasting demand, managing inventory, or coordinating suppliers—are increasingly being supported by algorithms. Companies are investing heavily in machine learning, large language models (LLMs), and automated decision systems in the hope of becoming faster, leaner, and more responsive.

> As this transformation accelerates, however, a bigger question begins to emerge: if algorithms become good enough, what role will humans still play?

A recent vision statement published in *Manufacturing & Service Operations Management* by Cohen et al. (2026) argues that this question is becoming central to the future of operations management. The paper discusses how AI is reshaping supply chains across five interconnected layers: intelligence, execution, strategy, human, and infrastructure. Yet what makes the paper especially interesting is that it does not treat AI simply as a technical tool. Instead, it raises a broader concern: modern supply chains may become increasingly efficient and automated, while at the same time becoming more dependent on human judgment, trust, and oversight.

This article focuses on the "human layer" of that discussion. The key issue is not whether AI can replace managers entirely. In many cases, it probably can automate large parts of operational decision making. The more important question is whether optimization alone is enough to build supply chains that are sustainable, resilient, and socially responsible.

The Sustainability Paradox

One of the most common assumptions about AI-driven supply chains is that better forecasting should naturally reduce waste. If firms can predict demand more accurately, they should be able to produce only what customers actually want. From an operations management perspective, however, the story may not be so simple.

AI systems dramatically improve responsiveness. In industries such as fast fashion, firms can react to changing trends almost instantly. This strengthens what Cachon and Swinney (2011) describe as "quick response" systems, where companies rapidly replenish inventory and adapt products to uncertain demand.

But faster response does not necessarily eliminate waste—it may simply move waste somewhere else.

Research by Long and Nasiry (2022) suggests that as firms become faster and more responsive, they often rely more heavily on upstream flexibility. To stay ready for sudden demand changes, firms may hold larger amounts of raw materials, unfinished products, or excess production capacity. As a result, stores may end up with fewer unsold garments, while significantly more fabric, semi-finished inventory, and production resources are wasted earlier in the supply chain.In other words, AI may reduce visible waste downstream while quietly increasing invisible waste upstream.

This is exactly where human judgment becomes critical. AI systems cannot decide on their own what *should* be optimized. Algorithms simply optimize the objective function they are given. If a company prioritizes speed, profit, and rapid trend adaptation above everything else, AI will pursue those goals extremely efficiently—even if the result is greater environmental pressure and more unsustainable consumption patterns.

That is why humans still matter. Managers define the objectives, constraints, and trade-offs that shape algorithmic behavior. Sustainability does not automatically emerge from optimization. It has to be intentionally built into the system.

The Context Gap

Another reason supply chains still require humans is that real-world operations contain far more context than algorithms can easily capture.

AI systems are trained on historical data and statistical patterns. But supply chains operate in environments filled with uncertainty, incomplete information, and rapidly changing conditions. A manager may recognize rising political tension in a supplier region, understand local labor issues, or anticipate operational disruptions long before those risks appear in formal datasets.

Humans therefore continue to provide something algorithms struggle with: contextual understanding.

Cohen et al. (2026) emphasize that AI contributes pattern recognition and predictive capabilities, while operations management still depends heavily on structural thinking, trade-offs, and contextual reasoning. Even highly accurate systems can fail when environments suddenly change in ways the data did not anticipate.

Ironically, more advanced AI systems do not always make managerial work easier. As algorithms become integrated into operational decisions, employees must spend more time interpreting outputs, evaluating recommendations, and deciding when systems should or should not be trusted. In many cases, the challenge is no longer whether AI can make decisions, but whether humans know when to rely on those decisions.

Trust, Responsibility, and Algorithmic Authority

One of the most interesting parts of the paper involves the problem of trust calibration between humans and algorithms. Research discussed by Cohen et al. (2026) suggests that people often react inconsistently to AI systems. Sometimes employees lose confidence in algorithms after seeing even a small mistake. In other situations, workers become overly dependent on algorithmic recommendations simply because disagreeing with the system feels professionally risky.

Over time, this can quietly shift authority inside organizations from human judgment toward algorithmic recommendations.

Employees may follow AI-generated decisions not because they genuinely trust them, but because relying on the algorithm feels safer than challenging it. In highly sensitive supply chain decisions—such as supplier selection, sourcing, or disruption management—this can become problematic. Decisions that are environmentally harmful, strategically weak, or ethically questionable may receive less scrutiny simply because they came from a system that appears objective or data-driven.

The Changing Role of Managers

The paper also discusses a future in which organizations may rely on multiple AI systems working together. Instead of a single algorithm making decisions, firms may increasingly use one system to generate recommendations and another system to evaluate them.

For example, one AI system might recommend a low-cost supplier, while another system reviews whether that supplier creates sustainability risks, geopolitical exposure, or policy violations. In this structure, one algorithm acts almost like a decision maker while another behaves more like a reviewer or supervisor.

Even in this type of system, however, humans remain essential. Managers still define goals, decide which trade-offs matter most, and intervene when systems conflict or fail under unexpected conditions. The role of managers may therefore shift from directly making every operational decision to supervising and guiding networks of intelligent systems.

Conclusion: Optimization Alone Is Not Enough

A simple example from retail workforce scheduling illustrates why human involvement still matters. Algorithms can create mathematically optimal schedules based on labor costs, staffing constraints, and demand forecasts. Yet store managers often adjust these schedules using local knowledge that never appears in the data. They may know which employees work well together, which workers are unreliable, or when customer traffic tends to become unusually heavy. Interestingly, these human adjustments often improve actual operational performance.

This example captures a broader lesson about AI-enabled supply chains. Supply chains are not purely technical systems. They are sociotechnical systems shaped by human judgment, organizational incentives, environmental constraints, and contextual understanding.

AI will almost certainly continue improving forecasting accuracy, responsiveness, and operational efficiency. But the future of sustainable supply chains may ultimately depend less on eliminating humans from decision making and more on understanding how humans and intelligent systems can complement one another effectively..

> AI can optimize supply chains, but it cannot decide what should truly matter. Human managers remain essential because they define priorities, balance trade-offs, and recognize risks and context that data alone cannot capture.

References

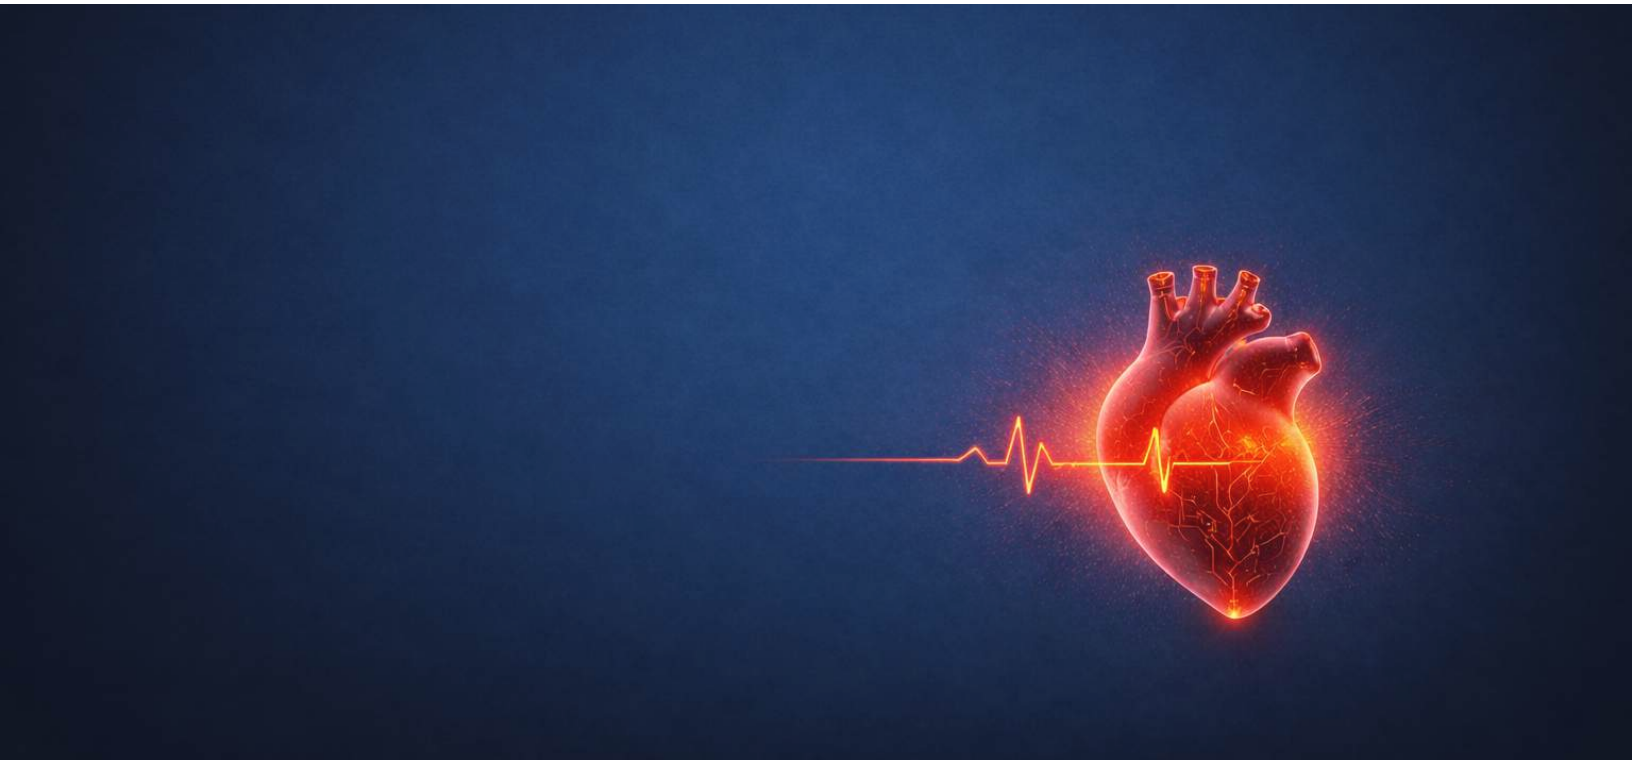
Cachon, G.P., Swinney, R., 2011. The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. Management science 57, 778–795.

Cohen, M.C., Dai, T., Perakis, G., Agrawal, N., Allon, G., Boute, R.N., Cachon, G.P., Chen, Z., Cohen, M., Cristian, R., et al., 2026. Om forum—supply chain management in the ai era: A vision statement from the operations management community. Manufacturing & Service Operations Management .

Long, X., Nasiry, J., 2022. Sustainability in the fast fashion industry. Manufacturing & Service Operations Management 24, 1276–1293.

Acknowledgements: *We would like to thank Ronak Tiwari for taking time to review this article. Photo credit goes to assistance of ChatGPT for the header and footer photos.*





AI in Resuscitation: Smarter Help When Every Second Counts



Marziehsadat Rezaei
University of Louisville

Fernando Acosta-Perez
University of Wisconsin-Madison



A Race Against Time: *Imagine a situation where a person suddenly collapses and their heart stops beating. In those first few minutes, everything depends on speed: recognizing what is happening, starting cardiopulmonary resuscitation (CPR), using a defibrillator if needed, and making the right decisions under intense pressure. This is the reality of cardiac arrest, one of the most serious medical emergencies in the world. Despite major advances in emergency medicine, survival rates after cardiac arrest are still low (Perkins et al., 2021). The reason is simple: resuscitation is not just difficult, it is a race against time.*

AI Steps Into Resuscitation

When a person suffers a cardiac arrest, survival depends on decisions that need to be made in seconds. Medical personnel or emergency responders should recognize the event quickly to dispatch resources and begin care coordination. Although cardiac arrest constitutes a medical emergency, in practice, it is a system challenge at which detection, communication, and execution must go hand in hand to maximize patient outcomes.

Today, artificial intelligence (AI) is opening a new and exciting chapter in this field. AI is already being used in many areas of medicine to find patterns in large amounts of data and support better decisions (Haug and Drazen, 2023). In resuscitation, the primary form of AI is machine learning (ML), which comprises a set of models that can learn from data to make predictions under uncertainty. Its promise is especially powerful because the challenge is not only to treat cardiac arrest, but also to recognize danger earlier, act faster during the event, and guide care afterward. In other words, ML may help at every stage of the resuscitation process: before, during, and after.

One reason this area is so interesting is that it is not limited to one method or one setting. For instance, a cardiac arrests can occur either in hospital or out of hospital. In hospital cardiac arrests benefit from continuous monitoring and immediate access to care. In contrast, out of hospital cardiac arrest remains a logistical challenge, with slower detection and constrained resources significantly affecting patient outcomes. These challenges highlight that improving outcomes requires not only better prediction, but also better coordination and decision-making across complex systems.

To go from predictions to actions, is where operations research (OR) can make an important contribution. If ML helps answer questions such as “Is this patient deteriorating?” or “Is this emergency call likely to involve cardiac arrest?”, OR helps answer the next question: “What is the best action to take?”

In this article, we examine how can ML and OR support

resuscitation across the three different stages of cardiac arrest (i.e., before, during, and after), and discuss the critical role of designing systems capable of leveraging the support from these tools and translating it into better patient outcomes.

From Prediction to Decision Making

Before a cardiac arrest even happens, ML can be used to identify patients who are at high risk. In hospitals, ML models can analyze clinical data such as vital signs, lab results, and monitoring information to detect signs of deterioration (Ueno et al., 2020). This kind of early warning system could alert care teams before a patient reaches the point of collapse. Preventing a cardiac arrest, of course, is even better than responding to one.

Out of hospital, preparing for cardiac arrest events include **location and coverage planning**. In this setting, survival often depends on whether an ambulance, AED, or nearby responder can reach the patient fast enough. OR models can help determine where to place these resources in order to maximize coverage, reduce response time, and improve access. This is typically modeled as a location-allocation problem, with common formulations including *maximal covering*, *set covering*, *and p-median* models. For example, optimization studies of public AED placement have reported better OHCA coverage under optimized deployment than under existing placement strategies (Chan et al., 2018). These are not simple placement problems. Cardiac arrests are rare, highly time-sensitive, and geographically uneven, so planners must often balance efficiency with fairness, especially across urban and rural areas.

During the arrest itself, ML can support some of the hardest real-time decisions. Researchers have developed tools that can classify heart rhythms, detect whether a patient has a pulse, evaluate CPR quality, and even help determine the best timing for defibrillation (Zace et al., 2025). Some systems use sensors or video-based deep learning models to measure chest compression depth and rate, offering feedback without interrupting CPR (da Costa et al., 2025). In emergency dispatch centers, natural language processing has also shown promise by helping identify cardiac arrest from emergency calls faster than traditional methods, allowing CPR instructions to begin sooner (Byrsell et al., 2021).

During an emergency in the out of hospital setting, emergency systems must quickly decide which resources to send, in what order, and along which routes. Should the closest ambulance be dispatched, or should it be preserved for another nearby high-risk zone? Should a drone carrying an AED be sent alongside EMS? Should fire, police, or trained volunteers also be activated? These are classic OR decision problems involving routing, real-time assignment, and dynamic resource coordination under uncertainty. Related work by Gao et al. (Gao et al., 2024) showed that intelligent real-time UAV–ambulance dispatch policies can outperform simpler benchmark strategies in time-critical EMS settings, especially when capacity is limited.

After the return of spontaneous circulation (ROSC), when the patient's heart starts beating again, ML models can answer prognosis questions: will the patient survive, and what kind of neurological recovery is likely? Researchers have used clinical records, EEG signals, imaging, and registry data to build models that predict survival and brain outcomes (Pham et al., 2022). These predictions could help clinicians make more informed decisions and improve communication with families during a very critical period.

Similarly, in the out of hospital setting, after an event emergency responders face the problem of **redeployment after a call**. When one ambulance or response team is sent to a cardiac arrest, the system immediately becomes less prepared for the next emergency. OR can help determine how to reposition the remaining resources to maintain system readiness. For example, Essus et al. (Essus et al., 2024) proposed a real-time optimization framework for EMS relocation and dispatching and found that explicitly accounting for workload balance leads to more effective redeployment decisions. This is especially important in cardiac arrest because even small delays can substantially affect survival. A good emergency system, therefore, must not only respond well to the current patient, but also stay prepared for the next one.

So, what have researchers found so far? Overall, the results are encouraging. Many ML models have shown strong performance in predicting outcomes, detecting deterioration, and improving CPR monitoring. Deep learning models that combine multiple data types sometimes report especially high accuracy, while ML methods such as XGBoost and random forests perform well across many different applications. OR studies have also reported promising results, particularly for improving coverage, dispatch efficiency, and redeployment decisions. At the same time, the field is still developing. Many studies are retrospective, use different outcome definitions, and are difficult to compare directly. In the OR literature as well, results are often highly dependent on local geography, demand patterns, and system design, which can limit generalizability across hospitals or EMS systems. A model or optimization policy that looks impressive in one dataset or system may not work equally well in another hospital, EMS network, or country.

These research findings lead to the most important practical question: ***how can ML and OR help in real systems?*** The answer is not that they will replace doctors, nurses, paramedics, or dispatchers. Instead, they can act as smart support tools. ML can warn clinicians earlier, identify high-risk patients, or recognize cardiac arrest during emergency calls. OR can then help determine how to allocate and dispatch resources, improve CPR-related logistics, and support better coordination across the chain of survival.

In a setting where every second matters, even a small improvement in speed or decision quality could translate into saved lives. For example, a more accurate prediction model is valuable, but it becomes far more powerful when linked with an OR model that determines where ambulances should wait, which responder should be activated, or whether an AED drone should be launched. In that sense, the real promise lies not in prediction alone, but in connecting prediction to decision

making.

Challenges on the Road Ahead

Although there is substantial evidence highlighting the power of ML in predicting risk across different resuscitation tasks, only a few studies have evaluated the potential of these systems beyond their predictive accuracy (Blomberg et al. (2019, 2021)). For example, in the context of out of hospital cardiac arrest, results from randomized control trials suggest that although ML can help in detecting the cardiac arrest earlier than emergency medical dispatcher by analyzing voice recordings, this did not necessarily translate into improved dispatching decisions (Blomberg et al. (2021)). This highlight the nuanced challenges behind implementing these models in practice.

This limited impact may be explained by several factors: lack of trust from medical dispatchers, limited integration with existing workflows, or the fact that current protocols were not designed to incorporate ML-based decision support. More broadly, this is where many open OR problems remain. It is not enough to develop accurate predictive models if emergency systems do not know how to use them operationally.

To this end, several open challenges are especially important. First, many existing studies stop at prediction and do not evaluate whether the predictions improve real decisions or patient outcomes. Second, uncertainty remains a major issue: demand varies over time and space, travel times can change suddenly, and not all responders have the same availability or skill level. Third, fairness matters. If models are trained mostly on data from certain hospitals or regions, they may not perform equally well everywhere, and optimized systems may unintentionally favor some populations over others.

Another open challenge is **human–AI interaction**. Even the best decision-support model will have limited impact if clinicians, dispatchers, or EMS leaders do not trust it or cannot easily incorporate it into practice. Future systems must therefore be not only accurate, but also transparent, usable, and aligned with real protocols.

To overcome these challenges, ML and OR systems must be tested in real clinical environments, not only in retrospective studies or simulations. Future progress will depend on stronger validation, more diverse datasets, better reporting standards, and close collaboration between clinicians, engineers, and policymakers (Zace et al., 2025). It will also depend on developing integrated frameworks that combine prediction, optimization, and implementation, rather than treating them as separate tasks.

Conclusions

AI in resuscitation is still evolving, but its potential is hard to ignore. From predicting cardiac arrest before it happens, to guiding CPR during the event, to supporting recovery decisions afterward, AI offers a vision of resuscitation that is faster, smarter, and more personalized. At the same time, the field should not stop at prediction alone. Some of the most important unanswered questions are operational: where to place resources, how to dispatch them, how to manage uncertainty, and how to ensure fairness and readiness across the entire emergency response system.

This is why OR has an important role to play. It helps move the field from knowing more to doing better. In the end, the goal is not to make medicine less human. It is to give healthcare professionals better tools, and better decisions, when human lives are on the line.

AI and operations research can make resuscitation faster, smarter, and more coordinated. Their greatest value is not replacing doctors, nurses, or emergency responders, but helping them make better decisions under extreme time pressure. When every second matters, the right prediction combined with the right action can help save lives.

References

Blomberg, S.N., Christensen, H.C., Lippert, F., Ersbøll, A.K., Torp-Petersen, C., Sayre, M.R., Kudenchuk, P.J., Folke, F., 2021. Effect of Machine Learning on Dispatcher Recognition of Out-of-Hospital Cardiac Arrest During Calls to Emergency Medical Services: A Randomized Clinical Trial. JAMA Network Open 4, e2032320. URL: https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2774644, doi:10.1001/jamanetworkopen.2020.32320.

Blomberg, S.N., Folke, F., Ersbøll, A.K., Christensen, H.C., Torp-Pedersen, C., Sayre, M.R., Counts, C.R., Lippert, F.K., 2019. Machine learning as a supportive tool to recognize cardiac arrest in emergency calls. Resuscitation 138, 322–329. URL: https://linkinghub.elsevier.com/retrieve/pii/S0300957218309754, doi:10.1016/j.resuscitation.2019.01.015.

Byrsell, F., Claesson, A., Ringh, M., Svensson, L., Jonsson, M., Nordberg, P., Forsberg, S., Hollenberg, J., Nord, A., 2021. Machine learning can support dispatchers to better and faster recognize out-of-hospital cardiac arrest during emergency calls: a retrospective study. Resuscitation 162, 218–226. doi:10.1016/j.resuscitation.2021.02.041.

Chan, T.C., Shen, Z.J.M., Siddiq, A., 2018. Robust defibrillator deployment under cardiac arrest location uncertainty via row-and-column generation. Operations Research 66, 358–379. doi:doi.org/10.1287/opre.2017.1660.

da Costa, L.M.A., Scheidewind, R.O., Peron, R.F., Timmerman, T., de Cerqueira Lario, F., Moreira, F.R., Nunes, R.A.B., Scudeler, T.L., 2025. Effectiveness of a smartwatch-based feedback system in improving cardiopulmonary resuscitation quality: a simulation study. Resuscitation Plus , 101114. doi:10.1016/j.resplu.2025.101114.

Essus, Y., De La Fuente, R., Venkitasubramanian, A., 2024. Real-time optimization for relocation and dispatching of emergency medical services with balanced workload and outsourced ride-hailing services. Computers & Industrial Engineering 187, 109823. doi:10.1016/j.cie.2023.109780.

Gao, X., Kong, N., Griffin, P., 2024. Shortening emergency medical response time with joint operations of uncrewed aerial vehicles with ambulances. Manufacturing & Service Operations Management 26, 447–464. doi:10.1287/msom.2022.0166.

Haug, C.J., Drazen, J.M., 2023. Artificial intelligence and machine learning in clinical medicine, 2023. New England Journal of Medicine 388, 1201–1208. doi:10.1056/NEJMra2302038.

Perkins, G.D., Gräsner, J.T., Semeraro, F., Olasveengen, T., Soar, J., Lott, C., Van de Voorde, P., Madar, J., Zideman, D., Mentzelopoulos, S., et al., 2021. European resuscitation council guidelines 2021: executive summary. Resuscitation 161, 1–60. doi:10.1016/j.resuscitation.2021.02.003.

Pham, S.D., Keijzer, H.M., Ruijter, B.J., Seeber, A.A., Scholten, E., Drost, G., van den Bergh, W.M., Kornips, F.H., Foudraine, N.A., Beishuizen, A., et al., 2022. Outcome prediction of postanoxic coma: a comparison of automated electroencephalography analysis methods. Neurocritical care 37, 248–258. doi:10.1007/s12028-022-01449-8.

Ueno, R., Xu, L., Uegami, W., Matsui, H., Okui, J., Hayashi, H., Miyajima, T., Hayashi, Y., Pilcher, D., Jones, D., 2020. Value of laboratory results in addition to vital signs in a machine learning algorithm to predict in-hospital cardiac arrest: A single-center retrospective cohort study. PloS one 15, e0235835. doi:10.1371/journal.pone.0235835.

Zace, D., Semeraro, F., Schnaubelt, S., Montomoli, J., Ristagno, G., Fijačko, N., Gamberini, L., Bignami, E.G., Greif, R., Monsieurs, K.G., et al., 2025. Artificial intelligence in resuscitation: a scoping review. Resuscitation Plus 24, 100973. doi:10.1016/j.resplu.2025.100973.

Acknowledgements: *We would like to thank Ronak Tiwari for taking time to review this article. Photo credit goes to assistance of ChatGPT for the header and footer photos*





OR/MS Tomorrow Mini-Poster Competition 2026

Share your research idea with the world... on one page!

About: The competition mainly seeks to nurture student researchers in developing the skills to summarize and articulate their research ideas effectively and efficiently. OR/MS Tomorrow is running the 2026 mini poster competition for a seventh consecutive year!

Eligibility: The competition is open to all students, including current INFORMS members and non-members. Candidates are required to design one letter-sized page to communicate their research topic to the broader audience in the OR/MS community. We accept submissions in three groups: undergraduate, master's, and Ph.D. students.

Criteria: Winners will be chosen from each group based on the novelty of the topic, maturity of the content, quality of the information, engagement, and aesthetics.

Prizes:

- The first-place winner in each group will receive $250 USD, and their poster will be featured in the Fall/Winter 2026 issue of OR/MS Tomorrow.
- The second-place winner in each group will receive $150 USD.
- Honorable mention in each group will receive $100 USD.
- Non-INFORMS member winners will also receive a free one-year INFORMS student membership.
- For undergraduate students entering the competition as a team, the cash prize will be divided among all team members.
- All winners receive an electronic certificate celebrating their achievements.

Submission Deadline: **October 15, 2026 at 11:59 PM AoE**

Submission requirements:

- The mini poster must be original.
- The poster must be submitted as a PDF file with a **MINIMUM** font size of 12 pt. Posters with smaller font sizes will have a score penalization.
- Portrait orientation is preferred.
- The author must be a student at the time of submission and provide a letter from their advisor or a certificate of current enrollment to indicate their enrollment as an undergraduate, master's, or Ph.D. student.
- Undergraduate students can submit individually or as a team of at most five (5) students. Still, the submission must go under the name of only one author. The remaining students' names can be added at the end of the submission form.
- Master's and Ph.D. students must submit only single-author posters. Collaborators or advisors can be acknowledged at the end of the submission form.
- Submissions are limited to one poster per student or group of students.
- Judging is blind, so the author must **NOT** include their name, address, email, or any other identifying information or acknowledgments on the poster itself.

Submit your poster here!

Check out the 2020, 2021, 2022, 2023, 2024, 2025 Fall/Winter issues to see the previous mini poster competition winners in our OR/MS Tomorrow Recent Issues page.

If you have questions about the competition, feel free to contact orms_tomorrow@mail.informs.org or check page here. Photo credit goes to TheDigitalArtist for header.