Report from Dagstuhl Seminar 24342

Leveraging AI for Management Decision-Making

Stefan Feuerriegel*[1], Foster Provost*[2], and Galit Shmueli*[3]

1 **Ludwig-Maximilians-Universität München, DE.** feuerriegel@lmu.de
2 **New York University, US.** fprovost@stern.nyu.edu
3 **National Tsing Hua University – Hsinchu, TW.** galit.shmueli@iss.nthu.edu.tw

Abstract

Artificial intelligence (AI) is transforming decision-making across industries and management functions, which leads to increased operational efficiency and creates significant economic impact. A recent surge in attention to AI in business decision-making has been driven by new AI technologies – such as deep learning, causal machine learning, generative AI and explainable AI – and their applications in areas like operations, marketing, information systems, and quality management. Yet, the potential of AI to optimize business decisions also introduces ethical, legal, and societal challenges, particularly in high-stakes business settings. This motivates our Dagstuhl Seminar, which aimed to foster interdisciplinary collaboration between scholars in management and computer science, as well as practitioners from industry. As a result, the seminar generated new suggestions for the field to evolve in the future by identifying new research opportunities with managerial relevance.

Seminar August 18–21, 2024 – https://www.dagstuhl.de/24342
2012 ACM Subject Classification Computing methodologies → Artificial intelligence; Social and professional topics → Computing and business; Computing methodologies → Machine learning; Information systems
Keywords and phrases applications, business, decision-making, management, marketing
Digital Object Identifier 10.4230/DagRep.14.8.24

1 Executive Summary

Stefan Feuerriegel
Foster Provost
Galit Shmueli

Introduction

Artificial intelligence (AI) has been named a core element of the "fourth industrial revolution" [40]. According to recent estimates by McKinsey & Company, AI has the potential to deliver the added global economic value of $13–20 trillion annually [5].

AI is increasingly being embraced for decision-making in management, both across a wide array of industries (e.g., healthcare, banking, education, manufacturing, retail) and functions (e.g., marketing, accounting, operations, IT). For example, AI can be used for modeling customer behavior [2, 6, 25, 26, 4]. These predictions can also serve as input for better decision-making. Examples include assortment optimization [24, 19], investment

* Editor / Organizer

decisions [32], scheduling [42], allocation decisions [29, 31, 46], and pricing [1]. AI can predict business failures and thus act as an early warning system for improving service quality [34]. AI can help to locate drivers of low quality and eventually improve product quality [44, 45].

Recent advancements in AI research hold great promise for decision-making in businesses and organizations. Driven by the surge in data availability, computing power, and algorithmic advancement, contemporary AI algorithms are capable of emulating human decision-making and judgment [39]. This places AI in a position to augment and automate a wide range of management decisions within companies and organizations. At the same time, the use of AI for decision making, especially in high-risk applications, poses ethical and legal challenges such as the use of algorithmic risk assessment tools in criminal justice [18]. Likewise, new implications arise from emerging AI acts (e.g., the EU AI Act) with crucial implications for how AI applications must be designed.

A key enabler for data-driven decision-making in business and organizations are new AI technologies [17]. For example, deep learning can empower better decisions in business analytics and operational decision-making [30, 28]. Causal machine learning (ML) allows for optimal targeting (e.g., of customer coupons) by estimating and subsequently leveraging individualized treatment effects [37, 20, 11, 14, 15]. Probabilistic machine learning fuses methods from a statistical foundation with flexible building blocks from neural networks to yield models that are both flexible and explainable for practitioners in risk management [38, 35, 36]. Further, explainable AI (XAI) has emerged as a principled, user-centered tool not only for explaining black-box prediction models but also for explaining the decisions that are made or recommended by AI systems [33, 45, 16, 27, 3], which can be used to identify root causes of bad quality and thereby inform better decision-making in quality management [44]. Likewise, generative AI [12] and AI fairness [9, 7, 10] offer new research opportunities. Importantly, the aforementioned examples can only be solved effectively through new AI technologies that have been developed in recent years. At the same time, the use of advanced AI on digital platforms, especially the marriage of reinforcement learning with behavior modification techniques, has spurred controversy, creating adverse effects to humans and societies, such as addiction, social discord, and political polarization [22]. Existing and envisioned combinations of prediction and causal behavior modification have implications to platforms and their business customers [23]. These technologies have also created new types of barriers for academic researchers [41, 21].

Aims of the seminar

The aim of our Dagstuhl Seminar “Leveraging AI for Management Decision-Making” (24342) was to discuss the future of research on AI/ML in businesses and organizations, and how the field should evolve. We especially sought to focus on “rethinking” the field by discussing the current state of AI/ML in businesses and organizations, discussing thought-provoking questions (e.g., is explainable AI really needed in practice? Where can generative AI actually lead to productivity gains? Does the algorithmic approach to fairness hold much for the future of AI ethics?), and maybe identifying and elevating new important research questions.

Our intended outcome was to reach a joint position (as a group) on what are important and unimportant research directions and what those directions should be going forward. What are the challenges? What are the opportunities? What research questions deserve attention? What questions are getting more attention than they really deserve? Below, we summarize our thoughts where we discuss existing research gaps and make suggestions for the field going forward.

Prior to the seminar, a survey was shared with all participants to identify key topics of interest around which we then designed our discussions. Most participants were primarily interested in topics related to method design and development, as well as the practical applications of AI. Some also expressed interest in evaluating these methods and understanding their impact on organizations. There was a strong focus on exploring the broader implications of AI, particularly in areas like ethics and governance. When asked about specific topics they would like to discuss during the seminar, participants showed the most interest in explainable AI, followed by causal ML, and generative AI. We eventually decided to prioritize the first two – explainable AI and causal ML – thus anticipating that generative AI will naturally arise in all sessions due to its prominence and thus regardless of whether it is a dedicated topic. Many indicated they were also keen to explore how AI can be applied effectively, with discussions centered around overcoming practical challenges in real-world deployments. Other popular topics include the ethics and governance of AI, its economic impact, and the implications for the future of work. However, there was less interest in topics like AI literacy and hybrid work environments. Further, participants also suggested additional topics for discussion. Some highlighted the importance of understanding the behavioral impacts of AI, such as long-term reliance on AI systems and the potential for deskilling human workers.

Organization of the seminar

We designed our Dagstuhl Seminar as an "un-conference". By following the format of an un-conference, we eliminated the traditional sequence of research presentations from our agenda. Instead, we aimed to focus on interactivity, collaboration, and co-creation, by making space for discussions of different forms regarding how to shape the field in the future. We held discussions with the full group as well as in smaller break-out groups, where subgroups changed from day to day; we obtained information from individuals via surveys; the schedule also encouraged informal one-on-one or small group discussions while socializing. These various modes of interaction were critical, because our seminar attracted participants from a diverse crowd, from academia and industry, from method research to behavioral research, from marketing to operations. Such diverse researchers typically do not meet or interact, and hence we seized the opportunity to foster novel interactions.

We aimed to learn from each other and create more impact. For example, we actively asked each participant in the get-to-know session to provide a summary statement about their current research and where they would like to go. Throughout the seminar, there were many opportunities to potentially start new collaborations. To spur discussion, we organized short "inspiration exchanges", which were designed as kick-offs to our breakout sessions. Hence, the idea was primarily to discuss the current state of research and point to gaps and needs to elicit forward-thinking. Here, we selected the topics prior to the seminar based on a survey that was sent to the participants. As a result, we identified two important breakout sessions: (1) AI and causality, and (2) AI and responsibility. We summarize the discussions and findings from both breakout sessions below.

2 Table of Contents

Executive Summary
Stefan Feuerriegel, Foster Provost, and Galit Shmueli 24

Summary of Breakout Session on "AI and Causality"
What is causal ML? 28
Key discussions 28
Conclusion 29

Summary of Breakout Session on "AI and Responsibility"
What is explainable AI? 30
Key discussions 31
Conclusion 31

Participants 35

3 Summary of Breakout Session on “AI and Causality”

This breakout session focused on exploring the interface of AI and causality by discussing open questions and emerging challenges.

3.1 What is causal ML?

The session started with an inspirational short presentation introducing the concept of causal ML [43, 14]. Causal ML is a branch of machine learning focused on identifying and understanding cause-and-effect relationships, rather than simply detecting patterns in data. Traditional machine learning models excel at predicting outcomes by recognizing correlations, but they often struggle to determine whether a specific factor directly causes a change in the target variable. Causal ML attempts to address this limitation by combining statistical techniques from causal inference with modern machine learning approaches. It allows researchers to answer questions like “What will happen if we change this feature?” or “What is the impact of a particular intervention?” This is particularly useful in fields like management where the aim is to to understand how management decisions will influence outcomes, so that the best decision for the business can be made.

In practice, Causal ML uses traditional statistical and econometric methods such as randomized controlled trials, propensity score matching, and instrumental variables to isolate causal effects. Additionally, more advanced ML-based techniques like causal forests and doubly robust estimators have been developed to handle complex, high-dimensional data where traditional methods might fall short. By leveraging these approaches, Causal ML models can go beyond predictions to inform actionable insights, such as making recommendations for business decisions in strategy, marketing, etc. Generally, causal ML follows the same frameworks as traditional causal inference, but the use of ML changes the underlying estimation strategy and thus offers several benefits in business practice, such as making personalized decisions for customers or users. Moreover, the goals of causal inference for decision making can be different from traditional causal-effect estimation [14], motivating the application of alternative machine learning approaches [15].

3.2 Key discussions

Defining Concepts and Aligning Terminology:

The discussions highlighted the need for clearer definitions and a shared vocabulary in the field. Participants acknowledged that “causality” is often interpreted differently across disciplines, which can hinder effective collaboration among different fields and between practitioners and researchers. Aligning terminology and notation was seen as a critical step for advancing research and practical applications.

One of the provocative questions debated was whether prediction (causal or non-causal) is sufficient for effective decision-making or if more granular causal inference is necessary. Participants discussed scenarios where accurate predictions might suffice (e.g., short-term business decisions), versus contexts where a deeper understanding of causal relationships is crucial (e.g., long-term strategic planning or policy-making). The consensus was that while predictive models are valuable, they may fall short in areas where understanding the “why” behind outcomes is essential. Yet, beyond that, the participants see a large need for more research and eventually evidence allowing for decision support.

As mentioned above, one of the biggest challenges in causal ML is deciding when to prioritize causal models over traditional predictive models. Predictive ML models, such as those used for classification or regression, are primarily designed to optimize accuracy based on historical patterns, without necessarily understanding the underlying relationships. However, when the goal is to make decisions that could change the environment or influence outcomes (like deciding on a new marketing strategy or a treatment plan in healthcare), causal modeling could be the "go to" approach. Yet, the challenge lies in identifying these scenarios where the extra complexity and effort of causal analysis are justified. Deep causal understanding requires carefully constructed assumptions and a deep understanding of the data-generating process, which can be resource-intensive. On the other hand, causal prediction may be sufficient (and quite effective) for decision making – and is often overlooked [13]. Without clear indicators of when causal methods are needed, organizations may either waste resources on unnecessary complexity or, conversely, miss opportunities to derive actionable insights where causality matters.

From a practical perspective, adopting causal ML techniques is often hindered by non-technical challenges such as the availability of skilled personnel, time, and budget constraints. Traditional ML has broad and accessible toolchains, tutorials, and community support, facilitating rapid training and deployment of models. In contrast, using ML to achieve causal understanding requires additional specialized training, including knowledge of statistical theory and causal inference frameworks, which are not as widely understood or accessible. Moreover, the data requirements for causal models can be more stringent, often requiring richer datasets or careful experimental designs. Together, this may create barriers for companies trying to leverage causal insights, especially those with limited resources.

A recurring question was whether algorithms for automated decision-making always require a causal framework. Opinions varied, with some arguing that in certain applications (e.g., dynamic pricing, recommendation systems), causal understanding is not always necessary, while others stressed that understanding causal relationships is crucial for ensuring fairness and avoiding unintended consequences.

Participants raised concerns about the limitations of causal inference from observational data, particularly when confounding factors and model misspecifications are present. The discussion emphasized the need for developing more robust methodologies to derive reliable causal insights, especially in real-world applications where randomized experiments are often impractical.

3.3 Conclusion

The breakout sessions successfully brought to light the complexities and nuances involved in integrating causality with AI, especially around formal definitions, guidelines when causal modeling is beneficial (and when not), and effective approaches. The discussions highlighted the importance of interdisciplinary collaboration to advance research, develop practical tools, and address the societal implications of AI-driven decision-making.

4 Summary of Breakout Session on "AI and Responsibility"

The second breakout session was centered around responsibility in AI adoption. It brought together participants to discuss the ethical, social, and governance implications of AI technologies. The breakout session was intentionally named broadly, so that it would include specific methodologies (e.g., explainable AI) as well as organizational implications. Eventually, the conversation was broad, covering explainable AI, generative AI, algorithmic fairness, and the societal impact of AI systems. A central focus was on the challenges and pitfalls of ensuring accountability in AI systems, which require a certain level of transparency in how decision-making algorithms operate.

4.1 What is explainable AI?

The starting point was an inspirational high-level short introduction and exchange discussing methods for explaining AI algorithms, models, and decisions, as well as the underlying objectives, rationales, and limitations [8]. Explainable AI (XAI) refers to a broad set of techniques and methods used to make the decisions, predictions, and/or inner workings of AI models more transparent and understandable to humans. Many AI systems, especially those that use complex algorithms like deep learning, are often seen as "black boxes" because their decision-making processes are not easily interpretable. This lack of transparency can be problematic, especially in sensitive areas like finance, credit lending, or law enforcement, where understanding how decisions are made is crucial for trust and accountability. Explainable AI aims to bridge this gap by providing insights into how or why models reach their conclusions, which can help users trust and effectively use AI systems.

The goal of XAI is to provide explanations that are not only faithful but also understandable to different stakeholders, such as data scientists, business leaders, or end-users. For example, in quality management, an XAI model might highlight which processes contributed most to a low quality level, helping manufacturing make more informed decisions about where to improve production processes. Techniques like feature importance scores, counterfactual explanations, interpretable models such as rule lists or shallow decision trees, and visualization tools are commonly used to make AI models and decisions more interpretable. By making AI decisions clearer, explainable AI could – in principle – help ensure that systems are not only accurate but also fair, reliable, and aligned with ethical standards, especially when those decisions impact people's lives. However, as we discuss below, there was large dispute among the participants whether this promise is true in practice.

Despite its benefits, XAI suffers from several limitations. One major limitation in practice is that it is often unclear for whom any particular explanation method was designed, and there is often a mismatch between the capabilities of these methods and the needs in practice. This leads to frequent cases where explanation methods fail to serve the needs of users or stakeholders who would like to use such explanations for decision-making. For example, it is often unclear how certain methods such as feature importance can help decision-makers generate insights that are actionable and that can be translated into better decisions. Another major challenge is that, for highly complex models like deep neural networks, providing simple and intuitive explanations can be difficult. Sometimes, explanations generated by explainable AI methods may oversimplify the decision process, leading to misunderstandings or even incorrect conclusions. Finally, there is little consensus on what a "good" explanation is.

4.2 Key discussions

Participants emphasized the importance of XAI for fostering transparency and trust, particularly in sectors like healthcare, finance, and criminal justice. However, the group also critically examined the current state of XAI research and its use in practice, thereby emphasizing salient limitations. A recurring concern was that many so-called "explanations" provided by AI systems are overly simplified, often failing to capture the complexities of the underlying algorithm. Another problem is that the explanations are rarely aligned with the needs of stakeholders in practice [33], so that actions derived from XAI may lead to negative outcomes. Together, the discussions highlighted that XAI can create a false sense of understanding or accountability, particularly when explanations are designed more for compliance than genuine transparency.

The participants also see a large potential for future research and thus for how the field could move forward. The discussion highlighted the need for more rigorous evaluations of the impact of AI explanations on decision-makers, particularly in contexts where human lives or rights are at stake. Currently, the evaluations are rarely aligned with how XAI systems are used in practice. For example, many evaluations are often one-sided and only focus on the role of programmers, while more holistic evaluation approaches that account for the different roles of stakeholders are still scarce. Another point for future research is to clarify what we mean under 'understandable' as this is a necessary condition to make assessments as to when one model is more interpretable than another.

The session also touched on the rise of regulatory approaches to AI (e.g., some regulatory frameworks offer a right for explanations in the context of data-driven decision-making). Here, participants emphasized the importance of regulatory and governance frameworks to provide oversight for the development and deployment of AI to prevent harm while enabling innovation.

4.3 Conclusion

The session underscored the complexity of ensuring responsible AI in practice. While explainable AI is a step towards greater transparency, participants highlighted that algorithmic explanations alone are not enough. In response to the session, the participants aim to create a commentary that offers critical reflections and guidance. We plan to discuss the following directions in our commentary: Which method should be applied by whom, in what context, and with what goal in mind? We will thus take a step back and will propose a comprehensive framework that shows the translation of real-world business problems into the model world while highlighting the critical dimensions that influence the success of XAI initiatives.

References

1 Svenja Bergmann and Stefan Feuerriegel. Machine learning for predicting used car resale prices using granular vehicle equipment information. *Expert Systems with Applications*, 263:125640, 2024.

2 Lennart Baardman, Igor Levin, Georgia Perakis, and Divya Singhvi. Leveraging comparables for new product sales forecasting. *International Journal of Operations & Production Management*, 27(12):2340–2343, 2018.

3 Kevin Bauer, Moritz von Zahn, and Oliver Hinz. Expl(ai)ned: The impact of explainable artificial intelligence on cognitive processes. *Information Systems Research*, 2023. forthcoming.

4 Ruomeng Cui, Santiago Gallino, Antonio Moreno, and Dennis J Zhang. The operational value of social media information. *Production and Operations Management*, 27(10):1749–1769, 2018.

5 Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, and Kate Smaje. *The economic potential of generative AI.* McKinsey & Company, 2023.

6 Real Carbonneau, Kevin Laframboise, and Rustam Vahidov. Application of machine learning techniques for supply chain demand forecasting. *European Journal of Operational Research*, 184(3):1140–1154, 2008.

7 Maria De-Arteaga, Stefan Feuerriegel, and Maytal Saar-Tsechansky. Algorithmic fairness in business analytics: Directions for research and practice. *Production and Operations Management*, 31(10):3749–3770, 2022.

8 Koen W De Bock, Kristof Coussement, Arno De Caigny, Roman Słowiński, Bart Baesens, Robert N Boute, Tsan-Ming Choi, Dursun Delen, Mathias Kraus, Stefan Lessmann, et al. Explainable ai for operational research: A defining framework, methods, applications, and a research agenda. *European Journal of Operational Research*, 317(2):249–272, 2024.

9 Mateusz Dolata, Stefan Feuerriegel, and Gerhard Schwabe. A sociotechnical view of algorithmic fairness. *Information Systems Journal*, 32(4):754–818, 2022.

10 Stefan Feuerriegel, Mateusz Dolata, and Gerhard Schwabe. Fair ai: Challenges and opportunities. *Business & information systems engineering*, 62:379–384, 2020.

11 Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Stefan Bauer, Niki Kilbertus, Isaac S. Kohane, and Mihaela van der Schaar. Causal machine learning for predicting treatment outcomes. *Nature Medicine*, 30:958–968, 2024.

12 Stefan Feuerriegel, Jochen Hartmann, Christian Janiesch, and Patrick Zschech. Generative AI. *Business & Information Systems Engineering*, 66(1):111–126, 2024.

13 Carlos Fernández-Loría and Foster Provost. Causal classification: Treatment effect estimation vs. outcome prediction. *Journal of Machine Learning Research*, 23(59):1–35, 2022.

14 Carlos Fernández-Loría and Foster Provost. Causal decision making and causal effect estimation are not the same… and why it matters. *INFORMS Journal on Data Science*, 1(1):4–16, 2022.

15 Carlos Fernández-Loría, Foster Provost, Jesse Anderton, Benjamin Carterette, and Praveen Chandar. A comparison of methods for treatment assignment with an application to playlist generation. *Information Systems Research*, 34(2):786–803, 2023.

16 Carlos Fernández-Loría, Foster Provost, and Xintian Han. Explaining data-driven decisions made by AI systems: The counterfactual approach. *Management Information Systems Quarterly*, 46(3):1635–1660, 2022.

17 Stefan Feuerriegel, Yash Raj Shrestha, Georg von Krogh, and Ce Zhang. Bringing artificial intelligence to business management. *Nature Machine Intelligence*, 4(7):611–613, 2022.

18 Travis Greene, Amit Dhurandhar, and Galit Shmueli. Atomist or holist? A diagnosis and vision for more productive interdisciplinary AI ethics dialogue. *Patterns*, 4(1):100652, 2023.

19 Sebastian Gabel, Daniel Guhl, and Daniel Klapper. P2v-map: Mapping market structures for large retail assortments. *Journal of Marketing Research*, 56(4):557–580, 2019.

20 Robin M Gubela, Stefan Lessmann, and Björn Stöcker. Multiple treatment modeling for target marketing campaigns: A large-scale benchmark study. *Information Systems Frontiers*, 26(3):875–898, 2024.

21 Travis Greene, David Martens, and Galit Shmueli. Barriers to academic data science research in the new realm of algorithmic behaviour modification by digital platforms. *Nature Machine Intelligence*, 4(4):323–330, 2022.

22 Travis Greene, Galit Shmueli, and Soumya Ray. Taking the person seriously: Ethically aware is research in the era of reinforcement learning-based personalization. *Journal of the Association for Information Systems*, 24(6):1527–1561, 2023.
23 Galit Shmueli, Ali Tafti. How to "improve" prediction using behavior modification, International Journal of Forecasting, Vol. 39(2), pp. 541–555, 2023.
24 Sebastian Gabel and Artem Timoshenko. Product choice with large assortments: A scalable deep-learning model. *Management Science*, 68(3):1808–1827, 2022.
25 Tobias Hatt and Stefan Feuerriegel. Early detection of user exits from clickstream data: A markov modulated marked point process model. In *Proceedings of The Web Conference 2020*, pages 1671–1681, 2020.
26 Tobias Hatt and Stefan Feuerriegel. Detecting user exits from online behavior: A duration-dependent latent state model. *arXiv preprint arXiv:2208.03937*, 2022.
27 Lukas-Valentin Herm, Kai Heinrich, Jonas Wanner, and Christian Janiesch. Stop ordering machine learning algorithms by their explainability! a user-centered investigation of performance and explainability. *International Journal of Information Management*, 69:102538, 2023.
28 Christian Janiesch, Patrick Zschech, and Kai Heinrich. Machine learning and deep learning. *Electronic Markets*, 31(3):685–695, 2021.
29 Milan Kuzmanovic, Dennis Frauen, Tobias Hatt, and Stefan Feuerriegel. Causal machine learning for cost-effective allocation of development aid. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.
30 Mathias Kraus, Stefan Feuerriegel, and Asil Oztekin. Deep learning in business analytics and operations research: Models, applications and managerial implications. *European Journal of Operational Research*, 281(3):628–641, 2020.
31 Mathias Kraus, Stefan Feuerriegel, and Maytal Saar-Tsechansky. Data-driven allocation of preventive care with application to diabetes mellitus type II. *Manufacturing & Service Operations Management*, 26(1):137–153, 2023.
32 Abdurahman Maarouf, Stefan Feuerriegel, and Nicolas Pröllochs. A fused large language model for predicting startup success. *European Journal of Operational Research*, 2024.
33 David Martens and Foster Provost. Explaining data-driven document classifications. *MIS Quarterly*, 38(1):73–100, 2014.
34 Christof Naumzik, Stefan Feuerriegel, and Markus Weinmann. I will survive: Predicting business failures from customer ratings. *Marketing Science*, 41(1):188–207, 2022.
35 Yilmazcan Ozyurt, Tobias Hatt, Ce Zhang, and Stefan Feuerriegel. A deep Markov model for clickstream analytics in online shopping. In *ACM Web Conference*, 2022.
36 Yilmazcan Ozyurt, Mathias Kraus, Tobias Hatt, and Stefan Feuerriegel. AttDMM: An attentive deep Markov model for risk scoring in intensive care units. In *ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.
37 Foster Provost and Tom Fawcett. *Data Science for Business: What you need to know about data mining and data-analytic thinking*. O'Reilly Media, Inc, 2013.
38 David Rügamer, Chris Kolb, and Nadja Klein. Semi-structured distributional regression. *The American Statistician*, 78(1):88–99, 2024.
39 Yash Raj Shrestha, Shiko M. Ben-Menahem, and Georg von Krogh. Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4):66–83, 2019.
40 Klaus Schwab. *The fourth industrial revolution*. World Economic Forum, 2017.
41 Galit Shmueli. Research dilemmas with behavioral big data. *Big Data*, 5(2):98–119, 2017.
42 Julian Senoner, Bernhard Kratzwald, Milan Kuzmanovic, Torbjørn H Netland, and Stefan Feuerriegel. Addressing distributional shifts in operations management: The case of order

fulfillment in customized production. *Production and Operations Management*, 32(10):3022–3042, 2023.
43 Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
44 Julian Senoner, Torbjørn Netland, and Stefan Feuerriegel. Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science*, 68(8):5704–5723, 2022.
45 Julian Senoner, Simon Schallmoser, Bernhard Kratzwald, Stefan Feuerriegel, and Torbjørn Netland. Explainable AI improves task performance in human-AI collaboration. *arXiv:2406.08271*, 2024.
46 Daniel Tschernutter and Stefan Feuerriegel. Data-driven dynamic police patrolling: An efficient Monte Carlo tree search. *European Journal of Operational Research*, 321(1):177–191, 2025.

Participants

- Kevin Bauer
Universität Mannheim, DE
- Margrét Bjarnadóttir
University of Maryland – College Park, US
- Jessica M. Clark
University of Maryland – College Park, US
- Theodoros Evgeniou
INSEAD – Fontainebleau, FR
- Carlos Fernández-Loria
HKUST – New Territories, HK
- Stefan Feuerriegel
Ludwig-Maximilians-Universität München, DE
- Sebastian Gabel
Erasmus University – Rotterdam, NL
- Travis Greene
Copenhagen Business School, DK
- Jungpil Hahn
National University of Singapore, SG
- Christian Janiesch
TU Dortmund, DE
- Enric Junqué de Fortuny
IESE Business School – Barcelona, ES
- Nadja Klein
KIT – Karlsruher Institut für Technologie, DE
- Mathias Kraus
Universität Erlangen-Nürnberg, DE
- Niklas Kühl
Universität Bayreuth, DE
- David Martens
University of Antwerp, BE
- Claudia Perlich
Two Sigma Investments LP – New York, US
- Joel Persson
Spotify – London, GB
- Foster Provost
New York University, US
- Galit Shmueli
National Tsing Hua University – Hsinchu, TW
- Sriram Somanchi
University of Notre Dame, US
- Wei Sun
IBM TJ Watson Research Center – Yorktown Heights, US
- Wouter Verbeke
KU Leuven, BE
- Michael Vössing
KIT – Karlsruher Institut für Technologie, DE
- Alona Zharova
HU Berlin, DE
- Patrick Zschech
Universität Leipzig, DE