



Applying Mathematical Optimization in Practice

A Note on Insights from MO Projects

Jens Schulz¹

Received: 6 August 2020 / Accepted: 11 November 2020 / Published online: 7 January 2021
© The Author(s) 2021

Abstract

Applying mathematical optimization (MO) in industry is still very often done by MO experts. In this contribution, we use an outside-in approach to describe the key requirements for optimization projects to be successfully instantiated within the businesses. We focus on the development and usage of decision support tools for operational, tactical, and strategic planning. They are to be used by business experts with little or no background in mathematical optimization. These people are still experts, e.g., in planning and managing complex supply chains or global production assets. While they are open to new tools and methods, in reality, many optimization projects fail to be handed over to them. We discuss the requirements and best practices for MO projects to overcome the underlying challenges. Several of the findings can be transferred to future applications of MO. These arise around the Industrial Internet of Things (IIoT), particularly focused on real-time applications, large, connected, and interdependent systems.

Keywords Mathematical optimization · Business applications · Future of OR

1 Introduction

In the last centuries, mathematical optimization (MO) has been applied to a rich and diverse amount of practical applications throughout the industry. MO methods have proven records in (i) one-time case studies, (ii) decision support tools, and (iii) real-time decision engines. Yet, taking into account articles in Operations Research journals, the offerings of consultancy companies, and the tools delivered by large software companies, most MO projects are still around one-time case studies and tools for these. These projects can be of high business impact, w.r.t. gross margins,

✉ Jens Schulz
schulz-gor@gmx.net

¹ Fair Isaac Germany GmbH (FICO), Stubenwald-Allee 19, 64625 Bensheim, Germany

footprints, operational efficiency, and so on. They usually require quite a huge investment of time from model developers and on the business side. Yet, often they remain one-time efforts, whereas businesses seek to develop re-usable components. They may have started with the goal to develop a decision support tool that is to be integrated into the business processes and to be used by the business. In practice, the majority of such projects starves from data quality, model accuracy, or organizational changes such as people switching roles, outsourcing, or divestments. Surviving such situations is the best seal of quality for any MO tool.

There are many optimization tools out there that are still used in practice even after disruptive events. What had been done right? What went wrong in other projects? What does it take to successfully implement a mathematical model and deploy it as a tool that is used by decision-makers? In this article, we will shed some light on these questions, share opinions, and present key learnings from a diverse set of industry projects on how to apply mathematical optimization in industry—sustainably.

Section 2 sets the scene of this contribution, the perspective taken, and highlights the challenges and benefits of MO projects in practice. Section 3 presents key user requirements for decision support tools and relates these to the available MO methods. Section 4 provides some insights on the process of conducting MO projects, and in particular highlights that interdisciplinary collaboration matters. Section 5 provides an outlook on main industry drivers for future directions of MO.

2 High-Level Perspective on MO Projects

In Section 2.1, we will shape the key user persona and assumptions made in this contribution because a general assessment about *Applying MO in Practice* may fill multiple books. Section 2.2 introduces the approach of this contribution to think from the perspective of the decision-makers. Section 2.3 provides an important aspect to this, namely that the process for making a decision is more important to decision-makers than having a mathematical model that is able to calculate an optimal solution. Sections 2.4 and 2.5 present the main benefits and challenges for MO projects when developing decision support tools.

2.1 Focus on Decision Support Tool and Two Key User Persona

When developing optimization tools and bringing them into the business processes, there are many people involved on the business side as well as on the software development side. We focus on two key user persona involved, two domain experts. One is the mathematical optimization expert and the other one is the business expert, the manager, planer, or simply called decision-maker who knows the business very well.

For simplicity, we assume that the MO expert is capable to create user interfaces (UIs) and to set up the IT infrastructure. This is partly realistic when it comes to analytic model development such as optimization models or forecasting models. Here, changes to the model impose needs and restrictions on the UI. It is almost impossible to sketch and define the whole UI and workflow upfront to UI designers before the underlying models and data are finalized.

We further focus on learnings for decision support tools that are to be used by the practitioners, who are usually no MO experts. Such requirements are only partly applicable when sufficient background in MO is present through education or self-studies. Many companies nowadays start hiring MO experts and their importance increases with the maturity of digitization of the company itself. Yet, especially in manufacturing and safety-related systems, experts in chemistry, pharmacy, engineering, mechatronics, and so on are the decision-makers. Many MO experts may not be interested to work in such roles and do not have sufficient technical expertise to take on the responsibility.

2.2 The Narrative Perspective of This Contribution

Most MO papers are written from the perspective of an MO expert—what does he/she need to do to successfully develop and deploy a mathematical model. MO experts discuss how to develop better algorithms, how to talk about optimization in order to make it easier to understand for their industrial partners, how to get the right data or how to get the data right, how to convince the partners that the algorithmic solution is optimal, how to figure out what is meant by optimal,... All of these are important questions, and still they are focused on the needs of the MO expert.

In this contribution, we take an outside-in approach, by discussing what are the user requirements, what capabilities need to be provided by a decision support tool and the models, instead of what is needed to solve the optimization problem itself. The interdisciplinary character of optimization projects requires several soft skills from optimization experts to successfully apply the beauty and power of mathematical optimization in industry. The most important soft skill is the ability to put oneself in the position of the decision-maker—not only when using the new tool but also with respect to the organizational boundaries and limitations through processes, data, and knowledge.

We look at optimization projects from the perspective of the industry experts: what are their challenges when confronted with MO projects, which are often a totally new discipline to them. The focus in industry is not on decision variables, not on a set of constraints or about a specific objective function. It is about decision-making and many use cases in which decisions need to be made. Identifying the underlying use cases often results in the insight that there need to be multiple models, hopefully a generic one just triggered with different parameter settings to satisfy each use case accordingly.

2.3 Optimality—Does It Matter?

Optimization calculates an optimal solution among a vast set of alternatives. Considering a huge amount of interdependencies at once using MO is powerful. It automates manual calculations and enables users to incorporate interdependencies that one could otherwise impossibly handle manually. Thereby, saving a lot of time. Yet, it is one piece of the puzzle. A decision support tool is supposed to support decision-making, not necessarily making the final decision. Optimization highly depends on the input data, and parameters a user needs to set before any algorithm provides a

solution. The objective function usually combines quantitative and qualitative aspects that cannot be easily aggregated. Decision-making is a process to come to a decision based on insights and an understanding of the interdependencies. It is about making tradeoffs, between quantitative and qualitative objectives, about the known and the unknown, the underlying risks.

Business thinking is clearly, there is not one optimal solution. Providing a set of alternatives encourages interaction of the user with the tool. Alternative solutions need to differ in relevant aspects, such as cost, risk, or production stability.

Resiliency, i.e., the ability to recover quickly in case of unforeseen events, leads to business constraints such as multi-sourcing, fallback plans, and the need to purchase, store, and produce goods in a way that is far from optimal from a pure cost perspective. Resiliency of networks is an important consideration when making planning decisions. A solution where risks are well-distributed rather than centralized in one warehouse, at one supplier or on a single production line is often preferred over a cost optimal one. Similar notions of fairness exist under various circumstances. For example, it is better to run two machines at 80% load than one at 100% and the other at 60%. Models need to take that into account to generate realistic solutions.

This imposes some user requirements on the tools that can be very well supported by MO. MO has to think about a suite of models and algorithms to actually support the decision-making process, instead of the one model and algorithm that generates an optimal solution.

2.4 Benefits from Using Decision Support Tools

From a business perspective, the benefits of MO models are far beyond optimality of actionable plans. It is about transparency, managing complexity, single source of truth, integrated views, and time to decision-making.

A clear articulation why and how decision support tools help the planner to make decisions is important. Otherwise, the feeling of getting replaced by a tool resembles too strongly. Transparency on data and processes is crucial before any data can be plugged into a model. A tool needs to cover the interdependencies and hence reveals them to the various users. Often, the targets and complexity of the production planner are not fully understood by the supply chain planner. Therefore, transparency within the tool reduces the silos and enables an integrated planning across units and responsibilities.

Business decisions are usually complex as they have short-, mid-, and long-term effects. The production schedule impacts storage requirements, inventory, and finally decides whether market demand can be fulfilled. Cash flow, capital binding, and cost are waged at once. It impacts other planning departments, from raw material procurement, supply chain, and sales.

It is well accepted in business that this complexity exists, and that not all interdependencies are sufficiently transparent. Most companies have been structured into procurement, supply chain, operations, and sales. The structure makes complexity manageable, but due to sharpened responsibilities, they have become unable to make integrated decisions.

The journey along developing a decision support tool usually starts by structuring data and processes into similar units, to make the complexity manageable. Yet, MO models go this one step further to link the silos by incorporating the interdependencies into the model. Hence, manageability of the overall complexity via clusters is achieved, while the integrated view on the decisions is preserved.

Planners are too often challenged for their decisions. Why had the production rates been reduced? It turned out to be a bad decision. The historical data, the assumptions, and knowns from those days are still available and may reveal that the root cause is in some unforeseen demand. Only the perfect knowledge about the future would have made a difference. One tool can become the single source of truth where all the relevant data is stored. That data must have been validated by all stakeholders. Long lasting meetings in which the source of data or the overall knowledge about the interdependencies are in question are not required anymore. Having all of this in place reduces time to decision-making and allows for reliable and traceable decisions.

Decision-making is a continuous process along all the small or big changes of the real world, from sudden supplier shortage, unplanned maintenance, demand peaks, or any logistic bottlenecks. Planning and execution need to have a closed loop to adapt quickly in case of changes. At the same time, learnings need to be incorporated into upcoming planning phases.

2.5 Challenges

While all the benefits are promising, there are challenges around getting decision support tools into practice. The strongest one is fear, the fear of a planner to be replaced by some piece of software. Possibly, the algorithm can generate a solution that is 5 to 10% better and proves him a fool.

For business managers, the cost and return on investment play an important role. The tool needs to provide sufficient business value. Given that additional skills are required, and many people need to contribute and invest time to set up the tool initially. Often, fear persists that the business logic and its complexity cannot be covered sufficiently, or that the business cannot manage it afterwards themselves.

For the users of the tool, the planners in operations, supply chain, or any other department, the challenges can be partly overcome by the way such projects are conducted, the way benefits are communicated and the capabilities of the tool. Doing it wrong will lead to successful one-off projects but not to a persistent usage of the tool. User acceptance needs to be achieved by fulfilling key user requirements and through engagement of the users. This is the focus of the next two sections.

3 User Requirements on Decision Support Tools

In this section, some key requirements with respect to a decision support tool and its usage are discussed and related to the overall philosophy to provide a tool beyond proposing an optimal solution.

More Than Just KPI-Reporting Nowadays, user-friendly reports with interactive charts and maps are quickly set up. They are key to understand the data fed into the system and to explore the calculated results. Yet, a user interface with fancy dashboards is not the main need. The user interface needs to support the standard workflows, support what-if analysis for specific use cases and needs to be intuitive. To stand this challenge, user experience (UX) design principles are key to not pollute the UI.

A planner usually observes three to five key performance indicators (KPIs). If these are in an expected range, it is worth to drill-down into the details of a solution. Here, it is common practice to watch out for three to five main product streams. To give some examples: Are the products allocated in the correct amount to the machines expected? Are the best machines fully utilized? If the UI and reports support such usage patterns, planners will find the tool helpful and as a time saver.

Tradeoffs and Scenario Comparison In practice, there is usually no single objective function. It is a mixture of qualitative and quantitative KPIs that are based on different units of measures. While there is a tendency to use a weighted objective function to combine all these measures, it is very fruitful for the decisioning process to expose a fair amount of the KPIs explicitly. A range of two to five is usually manageable—the more, the harder to understand.

Opening additional warehouses may decrease transportation cost and increase service levels, while overall inventory and the work to manage more locations increase. Users want to explore such trade-offs, may impose restrictions per objective, and compare solutions that are sufficiently distinct. Multi-objective optimization or goal programming is easy to implement and should be part of any toolbox. While the calculation of an (pareto-optimal) efficient frontier may take up a considerable amount of computation time, the benefits for the user are key to success. Another option is to provide a user driven workflow. Here, the user defines what-if scenarios based on a set of rules per scenario. In either case, the ability to compare such scenarios and to understand the drivers in the input data that trigger specific outcomes in the solutions is key to generate insights.

Uncertainty and Sensitivity MO models are an abstraction of the real world. The data provided may just be an assumption. It may be just a best guess, possibly derived from sophisticated forecasting algorithms. Yet, no technique can predict when exactly or whether at all a machine failure, sickness, or economic crisis may arise. Planners are well aware of this, generating buffers into their plans, or leaving some spare production or storage capacity. It depends on the risk attitude of the decision-maker. In mathematical models, we can cater for uncertainty. It is possible to present solutions that are easier to recover for a broad range of events. Depending on the types of uncertainty, stochastic programming, (recoverable/light) robust optimization, or a Monte-Carlo simulation may be the technique of choice.

A planner is often interested to see the distribution over the possible outcomes. Especially in engineering and production, people are used to simulation and want to

look at the distribution function that tells the likeliness to achieve the KPIs and its standard deviation. They analyze how often production capacities may run over, and when supply falls short. With this in mind, they prepare fallback plans, adjust safety stocks, or plan for alternative storage facilities. This can be supported by running the model sequentially with slight adjustments to the optimization directions.

Accuracy and Infeasibility Detection In practice, decision-making is driven by operational rules, and capacities are numbers aligned to internal standards. It has become common that “production runs above 100% capacity.” Here, different definitions for capacity exist, such as technical, practical, and available capacity. While connected, they may rely on estimated or averaged values. Hence, a solution must not meet all capacity constraints accurately to be realistic. Being close enough is often sufficient. Furthermore, a fair distribution of excess workload across machines, and even more when it comes to over time of workers, is key to deliver realistic solutions.

Another extreme case arises when the user wants to explore why a certain production plan is not feasible. Mathematically, the user wants to detect the source of infeasibility. Which constraints to define as soft constraints and how to set the penalty factors under such circumstances is getting an art on its own. Soft constraints and penalties need to be applied with care, e.g., as a pre-processing step to identify inconsistencies in the input data, and by actively asking whether the planner wants to start calculating a solution, or wants to investigate the root cause in the data first before continuing with any calculations.

Bottleneck Analysis Related to the former, a tool needs to support bottleneck analysis. As an example, in multi-stage production environments, it is key to easily figure out why certain products cannot be produced. Is there a shortage in raw materials, limited storage capacities, or are the machines of a pre- or post-production step running at maximum capacity? The root cause may be based on missing data, wrong assumption (e.g., consumption rates), a flaw in the model logic, or simply due to lower profit per unit. There are two ingredients to resolve this. On the one hand, in a post-processing step, the dependencies are traversed and the set of tight or violated constraints per product is presented to the user. On the other hand, a tailored report for analysis can be generated. Such analysis reports cannot be generic. It depends on the use case and the user. A sales planner may start from unsatisfied demands, while the production planner looks at machine utilization and products first. There are common concepts to overcome the complexity of such situations. KPIs that summarize major findings, and the possibility to quickly jump into a table or chart where data can be sorted and filtered accordingly.

Interactivity Often, a planner wants to make small changes to a solution. This can be either to fix a decision for operational reasons or to explore slightly different solutions. The goal is to gather information why the choice may be bad overall or that it may be a well-justified adjustment if it lowers the number of overall changes. More generally, this boils down to the capability of providing a manual solution in order to understand how it compares to the optimized solution.

Such interactivity with a solution is highly requested, as it increases engagement of the planner with the tool and finally increases his motivation to use the tool. The ability to drag and drop jobs within a Gantt chart, to adjust capacities or demands via sliders, or to open wizards for quick configuration ensure a smooth user experience. For the developers, this may require some UI programming knowledge. In return, it supports acceptance of the tool.

Business Use Cases Reporting on major KPIs as described before is not self-sufficient. The purpose must not be just related to reporting. It needs to be tied into the business use cases. To improve business outcomes, specific operational actions need to be taken. Here are some use cases from the planners point of view where the model needs to be flexible enough to generate solutions. In case raw materials arrive late, the planner wants to know about the impact on production order and how to schedule best to meet demands. In case unplanned maintenance of a bottleneck production line must be carried out, it is important to understand the impact on intermediate stocks, other production lines, the amount on finished goods, and finally on service level agreements. The planer may want to explore multiple options how to recover.

The algorithms or models may require to be run in a different configuration per use case, especially when soft constraints are present. Otherwise, they generate confusing results and thereby limit user acceptance.

Implementation Plan—Do Not Mess Around with the Plan Everyday Last but not least, the decisions need to be put in place. Changes to the production plan or to warehouse locations need to be implemented. To this end, a tool needs to highlight which actions need to be taken to swap from an existing plan over to a new one. For example, worker 1 will move over to machine B, while machine A will be run by worker 2 who now contributes an unplanned working shift.

A planner will usually sort the required changes by their business impact. While a manager looks top-down from the gains per initiative, the planner usually takes a bottom-up point of view. In both cases, the actions should be grouped into clusters where the most promising initiatives are weighted against their risks. Some decisions may not be executed due to time or other resource constraints. In general, the optimization model should be aware of the fatigue of change. From a planner perspective, a good solution contains a minimum set of disruptions to the current plan. Hence, minimizing the deviations, be it measured in the total number of changes or a percentage, is important to provide realistic and manageable proposals.

In summary, mathematical models and tools need to support the various use cases of the planner, not just a single one. Compared to simply proposing an optimal solution, decision-making is done iteratively. It is a journey of gathering knowledge and awareness along the way, ending up in the ability to articulate the various trade-offs and to understand the underlying risks. Hence, optimization models need to be embedded into these usage patterns to generate insights. This way, decisions will be made based on the data and the models not being perceived as a decision by the optimization model but rather by the planners and managers themselves.

4 User Engagement

In the previous section, the user needs with respect to the tool capabilities to achieve user acceptance were highlighted. In this section, the focus is turned towards best practices to make the user feel engaged. Developing a tool is a journey where unknown complexities or needs are revealed, and challenges need to be overcome. If intrinsically motivated, users start thinking how to apply the model and reveal their (otherwise unknown) needs more easily.

4.1 The Collaborative Development Approach

Many projects still apply a waterfall model. All requirements for the mathematical model and the user interface are gathered upfront in workshops and in a written and signed requirement specification document. Then, data is collected and validated. After a few months, the developer comes back with a UI and a model to gather feedback. The user needs some time to get acquainted with the tool, validate the data and that all requirements are covered, even more he needs to understand how these are captured. It is quite time consuming and generates a relationship that requires to sign off tool capabilities. While it is good practice to have clear goals and requirements, in this approach, the user feels less engaged.

Finally, in almost each project, it turns out that some tiny details are missed or not well understood in this first phase. Failing to incorporate learning loops in the development process generates a huge amount of effort for adjustments. It is key to success to bring these two user persona together early, in the beginning of a project and throughout the development to quickly identify gaps and needs.

A more collaborative and iterative development approach often proves to be far more successful than a waterfall method. An early engagement between the user and the developer ensures that both learn from each other. They can explore and understand the needs of each other. This way, they exchange feedback early and both feel engaged.

Collaborative and rapid development process works as follows: the user explains the practical restrictions and provides the corresponding data assets. If data is not available, a sketch may suffice. Then, the developer enhances the model and the UI step by step. The user can provide feedback how this data should be displayed to suit his needs. Indirectly, he gets educated how to maintain the data and how to use the tool. Gaps in the data, the model, and needs for the UI can be explored, clarified, and overcome early in the process. This approach works far better because the user feels engaged on the whole journey, it becomes his tool. Intrinsically motivated, he starts thinking how to apply the model and reveal his needs. Needs he may not have been aware of before.

A clever balance between the waterfall model and the rapid application development approach is required. Projects need to run in clearly defined boundaries and time frames to ensure they are not everlasting point-wise extensions. Wireframes for model dependencies and for the UI can help to avoid meaningless development. Yet, learning from making mistakes needs to be embraced, not only deferred through additional processes.

4.2 Agility

Nowadays, many companies are working in an agile manner. It is important to note here that agility is not to be confused with flexible or ad hoc development. In agile development, strict timelines, priorities, roles, and responsibilities are present. It is partly inflexible. Developers are working in sprints, e.g., a 2- or 3-week cadence. The target is to deliver at the end of each sprint an enhancement to the software that can be tested. The feedback can be analyzed and prioritized for the upcoming sprints.

While collaborative development has been proposed for initial development of a tool, this recommendation changes when the software has reached a certain maturity. Then, it is recommended to follow an agile approach. Changes perceived as small additions can have a tremendous impact on the model or the UI and should be thought through before implementing. Otherwise, one ends up in endless rounds of fixing unforeseen side effects. A user must be enabled to provide feedback and get a response when the feedback can be incorporated.

4.3 Maintenance and Testing

Once an optimization tool has been finalized and gets used in an operational setting, it needs to be maintained and adjusted. Users usually have lots of feature requests once they start using the tool. There may be new business use cases popping up where new charts, fine-tuning of the algorithm or additional algorithmic steps are required. Optimization models live, some tuning, some parameters to turn on or off certain features will always come. Having test cases and an automated testing framework that verify implementation is a key and helpful to the developer already in the initial iterations.

4.4 Organizational Considerations

When applying MO in practice, it is important to openly discuss the intention of each party. Within companies and everyday work, there are lots of tasks that require alignment and prioritization. Active involvement needs to be part of the annual targets, and if the person is new to MO overall, it should be recognized as a development goal. The planner needs to articulate his goals for the tool and the priorities of the tool development versus other operational and organizational needs.

Of similar importance is awareness that people may change positions. New people will take over. This is often a showstopper for using the tool continuously. To overcome this, the tool itself needs to have proven to be useful. It must be fit for purpose. Then consequently, the user will make recommendation to the successor, provide training and handover sessions. Even better to be a sparing partner in the first weeks or months. Furthermore, the personal relation of the developer and the user are key. If that proved well during development and go-live phase, the user will make recommendation to his successor—technically about the tool and personally about the collaborative work. *Feel free to ask* is easily said, but only done if earned through trusted relationships. Collaboration and interdisciplinary thinking is one of those hidden soft skills that make the difference between failure and success.

5 Outlook

We have focused on key requirements for decision support tools to be accepted and used by non-MO professionals in practice. These provide the basics for MO to also play a bigger role in future application areas as they arise in initiatives around the Industrial Internet of Things (IIoT). In IIoT, production units and whole value chains of multiple companies are connected closely or loosely. They exchange information. Decisions need to be made taking into considerations the various stakeholder needs and constraints.

Using a decision support tool for planning purposes has been discussed in detail. The concepts presented need to be transferred to real-time applications. Most money is still lost during ad hoc decisions that miss the bigger picture. With IIoT, we will see more and more real-time applications that can benefit from being evaluated by known and trusted mathematical models. Autonomous driving of cars, trains, or other vehicles are among the hot trends in mobility. Real-time control and autonomy of robots such as robot vacuum cleaners or lawn-mowing robots, automatic guided vehicles at container terminals or in large retail warehouses are already an inherent part of our lives. Routing algorithms are part of such system. Systems in manufacturing are starting a similar journey. As soon as the safety considerations can be managed, they will become real.

Production plants get more and more automated. The factory in the box is a hot topic. Being fully automated, it comes with software for automated replenishment. The machines themselves are equipped with sensors and act like autonomous agents. They detect their need for maintenance. Again, multiple machines may compete for scarce maintenance resources, and a decision with respect to material and worker availability needs to be made. Such decisions can be very well supported by MO models and capabilities discussed earlier. If planners have gained trust in the solutions and MO methods, then more likely people will trust that these work in automated and autonomous systems.

MO has the potential to play a big role there. Yet, it needs to be able to derive realistic solutions in reasonable time, interact, and adapt if decisions turned out to be sub-optimal or wrong. Analyzing the root causes and identifying the external drivers which would have made another algorithmic parametrization better will be important. The notion of learning is not that prominently present in MO. Failures and sub-optimality can be well accepted if there is trust that the same error is not made again. Here, the MO community needs to learn from the various concepts known in machine learning.

Real-world systems are distributed, complex systems. Hence, decisions are often distributed and yet, they need to be merged in order to derive system optimal business outcomes. Hence, in future, we need to think about a system of algorithms, models, and methods each tailored to tackle localized decisions aware of the boundaries of the overarching system. An overarching system needs to set the right boundaries within which the local systems can make decisions autonomously. Information exchange is key. Several key user requirements discussed earlier translate to such integrated needs. MO needs to find answers for different use cases. A single model or a single parameter set per model will not suffice. Models need to adjust data and control

parameters according to the use cases. Some decisions can be fully made automated, and others may require some human supervision.

In general, the types of production and distribution systems envisioned for the future will require a diverse set of skills to come together. From rule-based simulation approaches, statistics, machine learning, and mathematical optimization, the right methods need to be applied to tackle these challenges. Hybrid methods will strongly outperform single methodologies. This is already happening as the communities speak to and learn from each other.

Mathematical optimization experts will work in larger teams with experts from other areas. Not only experts on the business side but also from other disciplines such as machine learning, rules, simulation, UI development, and software architecture. Cloud computing is getting a standard and algorithms need to leverage the potential. Quantum computing is another research and investment area with the potential to change how we tackle business problems algorithmically. At least, it provides a potential for new ways to think how to develop hybrid solution methods. Going forward, working on a joint purpose and across disciplines, making challenges transparent, successes visible, and learning in such interdisciplinary environments will become more and more part of the MO expert's life in industry.

Compliance with Ethical Standards

Conflict of Interest The author is employed at an analytics and optimization software company, FICO (Fair Isaac Germany GmbH), and president of the GOR working group, Real-World Optimization.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.