Journal of the Operational Research Society

ISSN: 0160-5682 (Print) 1476-9360 (Online) Journal homepage: www.tandfonline.com/journals/tjor20

Commentaries on A. Mor, C. Orsenigo and M.G. Speranza "optimization and supervised learning for decision making: competitors or partners?"

Okan Arslan, Carmine Cerrone, Richard Eglese, Bissan Ghaddar, Bruce L. Golden, Martin Kunc, Gilbert Laporte, Francisco Saldanha-da-Gama, Said Salhi, Kate Smith-Miles & Daniele Vigo

To cite this article: Okan Arslan, Carmine Cerrone, Richard Eglese, Bissan Ghaddar, Bruce L. Golden, Martin Kunc, Gilbert Laporte, Francisco Saldanha-da-Gama, Said Salhi, Kate Smith-Miles & Daniele Vigo (2026) Commentaries on A. Mor, C. Orsenigo and M.G. Speranza "optimization and supervised learning for decision making: competitors or partners?", Journal of the Operational Research Society, 77:6, 1653-1664, DOI: 10.1080/01605682.2026.2653919

To link to this article: https://doi.org/10.1080/01605682.2026.2653919

Published online: 28 Apr 2026.

JOURNAL OF THE OPERATIONAL RESEARCH SOCIETY
2026, VOL. 77, NO. 6, 1653–1664
https://doi.org/10.1080/01605682.2026.2653919

COMMENTARY

OPEN ACCESS Check for updates

Commentaries on A. Mor, C. Orsenigo and M.G. Speranza "optimization and supervised learning for decision making: competitors or partners?"

Okan Arslan[a], Carmine Cerrone[b], Richard Eglese[c], Bissan Ghaddar[d], Bruce L. Golden[e], Martin Kunc[f], Gilbert Laporte[a,g], Francisco Saldanha-da-Gama[h], Said Salhi[i], Kate Smith-Miles[j] and Daniele Vigo[k]

[a]HEC, Montreal, Canada; [b]University of Genoa, Genoa, Italy; [c]Lancaster University, Lancaster, UK; [d]Western University, London, Canada; [e]University of Maryland, College Park, MD, USA; [f]University of Southampton, Southampton, UK; [g]University of Bath, Bath, UK; [h]Sheffield University Management School, Sheffield, UK; [i]University of Kent, Kent, UK; [j]University of Melbourne, Melbourne, Australia; [k]University of Bologna, Bologna, Italy

ABSTRACT

This article comprises a set of commentaries provided by scholars in the field for the discussion paper by Mor, Orsenigo and Speranza. Though all appreciate the integration of optimisation and Machine Learning, and considered the topic to be timely and challenging their views also highlight several aspects that ought not to be ignored. The authors then produced their responses to these challenging commentaries in a concise and forward looking spirit. For simplicity, the commentaries are put in alphabetical order of the authors surnames.

KEYWORDS

Optimisation; learning; computation; practice

Okan Arslan—HEC Montréal, Canada

Gilbert Laporte—HEC Montréal, Canada and School of Management, University of Bath, UK

Mor, Orsenigo and Speranza take on a question that sits behind many hallway conversations in our community: is supervised learning moving towards *replacing* optimisation in decision support, or are we entering an era where the two can be designed as true complements? The authors answer carefully: they distinguish between what optimisation and supervised learning fundamentally do in decision-making pipelines, and then map the increasingly rich set of interaction patterns between them, from "predict-then-optimize" to decision-focused learning and, at the far edge, end-to-end learning.

What is most compelling is the paper's insistence on separating "goals" from "tools". Optimisation, as framed here, is a problem-model-algorithm workflow, where modelling is a deliberate act of abstraction and value-setting: choosing decision variables, constraints, and objectives that reflect what we believe should happen. Supervised learning, by contrast, is primarily a pattern-extraction workflow, turning historical labelled instances into a trained predictor that can be updated as new data arrive. This distinction is not semantic; it reveals why "replacement" is a misleading benchmark. Even if learning systems could output feasible solutions quickly, the question "optimal with respect to what?" does not disappear—it merely becomes implicit, buried inside labels, loss functions, and the biases of past decisions. It also clarifies a second, practical fault line the authors emphasise: predictive accuracy need not translate into decision quality, motivating decision-focused learning that trains against decision loss rather than prediction loss.

The paper's emphasis on distribution signals a broader point: what instances are we training on, benchmarking on, and ultimately deploying on? The authors propose "ML-enhanced instance generation" as an underexplored direction: learn generators that better reproduce the structure of real instances, so that algorithm selection and evaluation become less stylised and more transferable. This direction is promising, yet it exposes a tension in today's OR-ML interaction. Much of learning-augmented optimisation implicitly assumes that training distributions are stable and representative. Yet many operational environments are not: demand regimes and networks change, and what counts as "typical" can shift abruptly. In such settings, the most valuable contribution of OR may be exactly what supervised learning struggles with: explicit feasibility, safety margins, robustness to worst cases, and a capacity to reason about counterfactuals rather than correlations.

OR's distinctive contribution is not only solving hard deterministic models, but designing decisions that remain acceptable when travel times, demands, and costs are random or change regime—through robustness, stochastic optimisation, and explicit risk measures.

CONTACT Said Salhi s.salhi@kent.ac.uk Journal of the Operational Research Society, Kent Business School, University of Kent, Kent, UK

Supervised learning can strengthen when *it moves beyond point predictions* towards calibrated uncertainty that optimisation can use: predictive distributions, scenario generators that preserve correlations, and signals of "out-of-distribution" instances where the safest action may be to hedge or defer. This also reframes the predict-then-optimise critique in practical terms: the goal is not the lowest mean-squared-error, but uncertainty information that is *decision-relevant*.

The section on end-to-end learning is appropriately cautious: amortising expensive solve time into training is tempting, but feasibility, variable-size outputs, the need for labelled solutions, and generalisation outside the training dataset are not side issues, they are the main points. Many OR problems are hard because the decision space is combinatorial and heavily loaded with constraints, so that small prediction errors can create infeasibility or large regret. This is why the most credible near-term direction is not "learning instead of optimization," but "learning-guided optimization with optimization safeguards", so that learning proposes and optimisation verifies, repairs, and certifies.

The paper's conclusion with "partners, not competitors" lands at a moment when OR-ML hybrids are becoming the default expectation in many applications, forcing us to be explicit about what is normative (objectives/constraints) and what is learned (data-driven structure). The future, then, should not be measured by whether learning can "beat" a solver on a benchmark, but by whether hybrid pipelines can deliver decisions that are faster, safer, more transparent, and more aligned with what decision makers (and societies) actually value. Congratulations to the authors for a timely and clarifying synthesis that helps set a constructive agenda for the OR-ML interface.

ORCID

Kate Smith-Miles http://orcid.org/0000-0003-2718-7680

Carmine Cerrone—Department of Economics and Business Studies, University of Genoa, Italy

Bruce L. Golden—Robert H. Smith School of Business, University of Maryland, U.S.A

At a high level, Mor et al. (2026), in their provocative article, explore the connection between operations research and machine learning/artificial intelligence. To make their analysis more specific, they focus on the relationship between optimisation and supervised learning in decision making. They point out that optimisation can be used to speed up or improve supervised learning algorithms and vice versa. A fundamental question is: Are optimisation and supervised learning competitive or cooperative methodologies?

In machine learning (ML), three common and important examples of supervised learning are linear regression, nonlinear regression, and neural networks. In regression, we find the best model (i.e., coefficients) by minimising the sum of squared errors. In linear regression, this is a convex function, and a global minimum is easily obtained. In nonlinear regression, we are often minimising a non-convex function in which case there is no guarantee that we will find the global minimum. In training a neural network to obtain the best edge weights, we need to minimise a non-convex loss function. Again, there is no guarantee of obtaining the global minimum. Clearly, optimisation is essential in the implementation of supervised learning.

Ideas from machine learning have recently been incorporated into heuristics to obtain near-optimal solutions to a wide variety of NP-hard problems in combinatorial optimisation. But starting in the 1980s, Tabu Search used memory to guide the search process (see Glover, 1989, 1990). In particular, the search was influenced by the history of the intermediate solutions it encountered. This was an early example of learning and, for several decades, Tabu Search, developed by Fred Glover, was one of the most powerful metaheuristics available.

The Carousel Greedy (CG) algorithm was proposed in 2017 (Cerrone et al., 2017). Unlike learning- based approaches, CG exploits the information generated during the greedy construction phase to guide subsequent decisions. In contrast, a recent paper by Dragone et al. (Dragone et al., 2025) introduces a new metaheuristic, called Data-Dependent Carousel Greedy (DDCG). This procedure enhances the original CG framework by incorporating adaptive, data-driven mechanisms that explicitly depend on the observed data. The performance of this metaheuristic is impressive, as it outperforms existing metaheuristics and runs in a reasonable amount of computational time.

With respect to exact methods in integer and mixed integer linear programming (MILP), more formal supervised learning methods have been shown to improve performance and numerical stability. For example, Zhang et al. (2026) design a reinforcement learning/neural network model that learns how to select cutting planes in MILP.

Of course, there are other articles in which a supervised learning model, typically a neural network, is embedded within an MILP. For example, see Li et al. (2025), Stinchfield et al. (2025), and Moreno-Palancas et al. (2025).

In Li et al. (2025), The authors present an MILP model in which a deep neural network (DNN) is embedded for fault prediction and production optimisation in tablet pressing machines. The DNN estimates failure probabilities from operational parameters and these predictions are used within the MILP to optimise process conditions, production scheduling, and maintenance planning in order to maximise profit. ReLU activations are linearised using Big-M constraints, ensuring global optimality. A case study reveals improved profitability, resource utilisation, and robustness to changes in energy and maintenance costs.

In Stinchfield et al. (2025), the authors introduce a chemical process design approach that tightly integrates machine learning (mainly neural networks and decision trees) and MILP optimisation. The machine learning models are used to build piecewise-linear surrogate functions that approximate system costs and performance metrics, replacing the complex nonlinear relationships of the original problem. These ML-based surrogates are embedded within an MILP formulation derived from an initially intractable nonlinear model. In this paper, ML captures this underlying process behaviour, while the MILP efficiently optimises the selection and design of chemical process configurations. The approach is validated in several case studies, achieving a high level of computational efficiency and solution quality.

In Moreno-Palancas et al. (2025), the authors study stochastic bilevel optimisation problems. They propose an approach in which the lower-level problem is approximated using a neural network. The network is trained over multiple realisations of uncertain parameters, enabling a computationally tractable single-level reformulation of stochastic bilevel problems. The methodology is validated on a batch chemical process scheduling problem, demonstrating the efficient solution of complex bilevel models. Looking forward, there seem to be ample opportunities to embed supervised learning within exact optimisation models.

Over the last ten years, Dimitris Bertsimas and his colleagues have argued for greater integration of machine learning models and optimisation models. A good starting point is their foundational 2019 book entitled Machine Learning Under a Modern Optimisation Lens (Bertsimas & Dunn, 2019). The authors demonstrate that many machine learning models can be formulated as convex optimisation problems or as mixed-integer optimisation problems. Given the current power of hardware and optimisation software, this can lead to solutions to decision problems in reasonable computing times along with global optimality guarantees. This point of view is supported by Sun et al. (2019) and Gambella et al. (2021). Two related and significant research projects are presented by Bertsimas and Dunn (2017) and Bertsimas and Kim (2023).

At present, there are many articles which show that machine learning models can be directly embedded into MILP formulations. The key point here is that optimisation and machine learning (especially supervised learning) are already merging naturally and bearing fruit.

In some optimisation applications, such as strategic planning exercises, it is not necessary to obtain the optimal solution or even a near-optimal solution. For example, in logistics planning, making decisions about the placement of depots, the size of service regions, and the number of trucks required by a service region does not require knowing the exact routes that will be traversed on a given day (Kou et al., 2024). Instead, the goal might be to determine an accurate estimate of optimal route lengths. This observation provides another important connection between supervised learning and optimisation. In particular, it involves the estimation of optimal solution values in combinatorial optimisation problems such as the Travelling Salesman Problem (TSP), the Vehicle Routeing Problem (VRP), and the Split Delivery Vehicle Routeing Problem (SDVRP). In the last five years, researchers have used randomisation, linear regression, neural networks, Kolmogorov-Arnold networks, and random forests to estimate optimal solution values to within one to three percent with respect to mean absolute percent error; for example, see (Kou et al., 2026). Simultaneously, Varol et al. (2024) have developed sophisticated neural network models that reliably estimate optimal TSP lengths with impressive accuracy.

In Mor et al. (2026), it is suggested that learning characteristics of optimal solutions based on input instances might effectively reduce the search space and, consequently, the computational effort. One simple example of this is to explore the contribution that supervised learning can provide with respect to estimation, using estimates to improve stopping conditions in heuristic search. In preliminary work, Kou et al. (2024) apply this idea and demonstrate that a set-covering-based heuristic for the VRP can reduce the average optimality gap while minimally increasing computational time. A more challenging example is one in which quantitative metrics might not be fully known in advance. Nearly 10 years ago, in Lum et al. (2017), we worked on partitioning a street network into compact, balanced, and visually appealing truck routes. How do we quantify the notion of visually appealing? Can machine learning enable us to learn from practical examples of visually appealing routes in order to quantify what a "good solution" looks like? This question might be connected to the emerging area of inverse optimisation.

In conclusion, we view optimisation and supervised learning to be cooperative and mutually reinforcing methodologies. Furthermore, we expect the machine learning and optimisation communities to continue to be creative in suggesting new synergies that will benefit both communities as well as decision making, in general. Finally, we thank Mor et al. for stimulating our thinking.

References

Bertsimas, D., & Dunn, J. (2017). Optimal classification trees. *Machine Learning*, *106*(7), 1039–1082. doi:10.1007/s10994-017-5633-9

Bertsimas, D., & Dunn, J. (2019). *Machine learning under a modern optimization lens*. Dynamic Ideas LLC, Waltham.

Bertsimas, D., & Kim, C. W. (2023). A prescriptive machine learning approach to mixed-integer convex optimization. *INFORMS Journal on Computing*, *35*(6), 1225–1241. doi:10.1287/ijoc.2022.0188

Cerrone, C., Cerulli, R., & Golden, B. (2017). Carousel greedy: A generalized greedy algorithm with applications in optimization. *Computers & Operations Research*, *85*, 97–112. doi:10.1016/j.cor.2017.03.016

Dragone, R., Cerrone, C., & Golden, B. (2025). The data-dependent carousel greedy algorithm. Submitted for Publication.

Gambella, C., Ghaddar, B., & Naoum-Sawaya, J. (2021). Optimization problems for machine learning: A survey. *European Journal of Operational Research*, *290*(3), 807–828. doi:10.1016/j.ejor.2020.08.045

Glover, F. (1989). Tabu search – part 1. *ORSA Journal on Computing*, *1*(3), 190–206.

Glover, F. (1990). Tabu search – part 2. *ORSA Journal on Computing*, *2*(1), 4–32.

Kou, S., Golden, B., & Bertazzi, L. (2024). Estimating optimal split delivery vehicle routing problem solution values. *Computers & Operations Research*, *168*, 106714. doi:10.1016/j.cor.2024.106714

Kou, S., Golden, B., & Bertazzi, L. (2026). Accurately estimating optimal SDVRP solution values using regression models. *Computers & Operations Research*, *188*, 107366. doi:10.1016/j.cor.2025.107366

Li, J., Wu, L., Qin, Y., & Zhi, H. (2025). Optimal operation of a tablet pressing machine using deep- neural-network-embedded mixed-integer programming. *Inventions*, *10*(2), 29. doi:10.3390/inventions10020029

Lum, O., Cerrone, C., Golden, B., & Wasil, E. (2017). Partitioning a street network into compact, balanced, and visually appealing routes. *Networks*, *69*(3), 290–303. doi:10.1002/net.21730

Mor, A., Orsenigo, C., & Speranza, M. G. (2026). Optimization and supervised learning for decision making: Competitors or partners? *Journal of the Operational Research Society*, 1–10. forthcoming doi:10.1080/01605682.2026.2616419

Moreno-Palancas, I. F., Femenia, R. R., Diaz, R. S., & Caballero, J. A. (2025). Solving bilevel problems under uncertainty with embedded neural networks: Incorporating scenario sets as inputs. *Digital Chemical Engineering*, *16*, Article 100253.

Stinchfield, G., Khalife, N., Ammari, B. L., Morgan, J. C., Zamarripa, M., & Laird, C. D. (2025). Mixed- integer linear programming formulation with embedded machine learning surrogates for the design of chemical process families. *Industrial & Engineering Chemistry Research*, *64*(16), 8299–8311.

Sun, S., Cao, Z., Zhu, H., & Zhao, J. (2019). A survey of optimization methods from a machine learning perspective. [ArXiv:1906.06821 [Cs, Math, Stat].]. http://arxiv.org/abs/1906.06821

Varol, T., Özener, O. Ö., & Albey, E. (2024). Neural network estimators for optimal tour lengths of traveling salesman problem instances with arbitrary node distributions. *Transportation Science*, *58*(1), 45–66. doi:10.1287/trsc.2022.0015

Zhang, X., Chen, L., Yang, Z., & Zeng, Z. (2026). Learning to select cutting planes in mixed integer linear programming solving. *Expert Systems with Applications*, *298*(Part D), 129924. doi:10.1016/j.eswa.2025.129924

Richard Eglese—Lancaster University, Lancaster, U.K

I found this an interesting paper exploring the differences and relationship between two approaches to finding solutions to well defined problems. There is considerable research taking place to improve the effectiveness of both optimisation and machine learning (ML), and the authors have chosen to restrict themselves to focus on supervised learning within the ML approach.

I shall confine myself to a few comments, mainly illustrated by reference to vehicle routeing problems.

The most effective method to use to solve a problem depends on how the problem is structured and the context for its use. For example, for a strategic problem involving depot location or districting, an approximation of the costs of travelling salesman problems (TSPs) may be needed, without the requirement to specify the detailed routes. In such a situation, a supervised learning approach to find the costs of the TSPs may be effective. Akkerman and Mes (2025) illustrate this approach in their paper which includes an example based on dynamic waste collection in Amsterdam. Their methodology still requires solutions from a training set and identification of features that may help the supervised learning model provide close approximations to the costs. However, if there are important constraints that the routes must respect, the ML approach does not guarantee that its cost estimates can be achieved by routes which satisfy the constraints. As it says in the final section of the discussion paper, "For example, in a vehicle routing problem the violation of operational constraints, such as vehicle capacity or road restrictions, may generate dangerous routes." The effects and consequences of failing to satisfy constraints in a problem is an important consideration when deciding on the approach to use.

If full solutions to routeing problems are needed, and not just approximations of their costs, then seeking to use a solely machine learning approach is much more challenging. "End-to-end learning" (E2EL) is discussed in Section 4 of the paper, but it is admitted that "in general, ensuring the feasibility of a solution in the presence of constraints remains a major challenge". Taking the example of solving TSPs again, there has been progress in using machine learning approaches, and a recent survey can be found in Alanzi and Menai (2025). This includes Graph Neural Net and other ML methods. Alanzi and Menai describe advancements that have improved the scalability of ML-based approaches including some that combine ML with traditional optimisation methods and conclude "Despite these advancements, challenges remain in generalization across problem distributions, maintaining solution quality in large-scale instances, and efficiently handling diverse real-world constraints."

The introduction to the discussion paper poses the question "…could ML one day replace traditional optimization?" In my opinion, the answer to that is negative, because there are many problems for which traditional optimisation is already able to provide an exact or close-enough solution in an acceptable computing time and guarantee to satisfy relevant constraints. As described in the paper, there are many ways in which optimisation and ML approaches can cooperate with each other to provide an effective solution method. Many types of cooperation have been proposed, but an example of one approach that provides a formal framework for such cooperation can be found in Bergman et al. (2022).

No doubt further research in optimisation and ML, as well as the availability of more powerful computing hardware, will increase the scope of problems that can be tackled by these approaches in the future.

References

Akkerman, F., & Mes, M. (2025). Distance approximation to support customer selection in vehicle routing problems. *Annals of Operations Research*, *350*(1), 269–297. doi:10.1007/s10479-022-04674-8

Alanzi, E., & Menai, M. E. B. (2025). Solving the traveling salesman problem with machine learning: A review of recent advances and challenges. *Artificial Intelligence Review*, *58*(9), 267. doi:10.1007/s10462-025-11267-x

Bergman, D., Huang, T., Brooks, P., Lodi, A., & Raghunathan, A. U. (2022). JANOS: An Integrated Predictive and Prescriptive Modeling Framework. *INFORMS Journal on Computing*, *34*(2), 807–816. doi:10.1287/ijoc.2020.1023

Bissan Ghaddar—Ivey Business School, Western University, Canada

I was very excited to read this discussion paper as it offers a timely and much-needed clarification at the interface of optimisation and supervised learning, two paradigms that increasingly coexist in modern decision-making systems. As machine learning continues to expand its influence across applied domains, the question of whether learning-based approaches might replace optimisation has become unavoidable. The authors engage with this debate by showing that optimisation and supervised learning are largely complementary, while also acknowledging that in certain settings such as end-to-end learning machine learning may act as a potential competitor to classical optimisation approaches.

A central highlight of the paper is its systematic and clear discussion of how supervised learning can enhance optimisation across the full decision pipeline. The authors outline several promising directions, including ML-assisted instance generation, learning-based algorithm design, and adaptive heuristics. Conversely, the paper also clearly illustrates how optimisation helps supervised learning itself. The paper provides relevant references and survey articles that situate these contributions within the broader literature.

Overall, this paper is of particular interest to the OR community because it reframes the ML-optimisation interaction as an opportunity to strengthen optimisation rather than sideline it. It encourages OR researchers to engage with supervised learning not as a replacement technology, but as a source of new methodological tools. By doing so, it contributes to shaping a future in which optimisation remains the backbone of decision-making, enriched by data-driven learning rather than dominated by it. Follow-up discussions could expand the scope beyond supervised learning to include, for instance, reinforcement learning and large language models for modelling and algorithmic reasoning, among other methods. I hope many future works build on this foundation, strengthening OR's role in modern decision making.

Martin Kunc—Southampton Business School, University of Southampton, UK

The paper provides a critical examination of the interplay between optimisation and supervised learning within contemporary decision-making frameworks.

Historically, optimisation has served as the cornerstone of operations research, offering mathematically rigorous models and algorithmic solutions for complex, constraint-driven problems. In contrast, supervised learning, propelled by advances in data availability and computational power, has emerged as a dominant paradigm for predictive analytics. Rather than positioning these approaches as mutually exclusive, the authors advocate for a synergistic perspective, emphasising how machine learning can augment optimisation processes through improved model calibration, algorithmic design, and instance generation.

A salient contribution of the paper lies in its discussion of integrated methodologies, such as decision-focused learning, which prioritise decision quality alongside predictive accuracy. Nevertheless, significant challenges persist, particularly regarding computational tractability, constraint satisfaction, and generalisation across heterogeneous problem instances. End-to-end learning approaches, which bypass traditional optimisation, offer efficiency gains but raise concerns about feasibility and robustness in high-stakes applications.

Future Directions: Advancing this research domain requires the development of hybrid frameworks that reconcile data-driven adaptability with optimisation's structural rigour. Promising avenues include ML-enhanced instance generation, representation-aware modelling for cross-domain transferability, and algorithmic design informed by learned patterns. Such innovations could redefine decision-making as an integrated paradigm, bridging predictive and prescriptive analytics.

Francisco Saldanha-da-Gama—Sheffield University Management School, UK

I thank the Editor of JORS for inviting me to comment on this paper, and I congratulate the journal on this initiative of pairing a discussion paper with several commentaries.

A central question guiding the paper is whether machine learning (ML) could one day replace "traditional" optimisation. As the authors emphasise, in the context of problem solving to support decision making, researchers working in optimisation increasingly seek to leverage ML to enhance their methods, while the ML community is becoming more engaged with optimisation to strengthen their own techniques. Given the breadth of ML as a field, the authors focus the discussion on supervised machine learning (SML), that is, on approaches based on an explicit training scheme. They clearly describe the goals, concepts and components underlying both optimisation and SML. This greatly facilitates the discussion. In particular, it becomes evident—if any doubt remained—that these are fundamentally different technologies. In Section 3, the authors examine how optimisation supports SML and vice versa. This analysis culminates in a discussion of their integration, which is effectively summarised in Figure 7. End-to-end learning as a possible competitor to optimisation is analysed in Section 4, where the authors convincingly highlight the several challenges that make such a possibility unlikely. The final section synthesises the arguments supporting the conclusion that SML is unlikely to replace optimisation, at least in the foreseeable future, given the current knowledge. Instead, the two appear destined to remain strong partners in the development of effective quantitative tools for decision-making.

I read this paper with great interest. I find the topic compelling, timely, and highly relevant. Although this is not a new theme for debate (see, e.g., de Lange, 2020), it remains of substantial importance. I count myself among those who view optimisation and SML as complementary technologies, each addressing distinct aspects of problem solving. The fact that (S)ML can support optimisation may sometimes give the impression that it constitutes an optimisation tool in itself. As the authors clearly explain, this is not the case. They exemplify two key aspects that distinguish these fields: algorithm selection and model updating. Furthermore, even when an SML tool aims to identify viable solutions to a decision-making problem, it often relies on optimisation methods, as emphasised in Section 3.1, where several illustrative examples are provided.

One important aspect highlighted in the paper is the role of human knowledge and experience in the critical phase of optimisation model definition—an activity that arguably "cannot rely on the support of any automated technique". While I largely agree with this assessment, it is worth acknowledging ongoing attempts to replace or reduce human expertise in this phase, as illustrated by Huang et al. (2025). Where this trend may ultimately lead, particularly in the context of artificial general intelligence, remains an open question. Regardless, SML is already playing an increasingly prominent role in optimisation model specification. Beyond forecasting model parameters and helping identify the most effective formulations (see Section 3.2), other notable examples can be found in areas such as distributionally robust optimisation, where SML can guide the identification of suitable ambiguity sets of distributions (see, e.g., Kannan et al., 2024). In this regard, I expect significant advances in the near future, particularly at the intersection of optimisation modelling, data science and decision-making under uncertainty.

The growing relevance of SML in the design of algorithms for solving optimisation problems is another important topic discussed by the authors towards the end of Section 3.2 (ML-augmented optimisation). Here, again, the long-term implications remain unclear. In addition to the references cited in the paper, I would add Huang et al. (2024), who discuss the role that large language models may play in algorithm selection and design. Although such models are not based exclusively on SML—this is just one component in their training—this emerging trend is nevertheless worth noting.

In my view, Section 3.3 is a pivotal part of the paper. The authors discuss several modern trends that intertwine optimisation and SML. The summary presented in Figure 7 is particularly insightful. The loop depicted in this figure is especially noteworthy: optimisation can support SML, which in turn can support optimisation. This reciprocal relationship reinforces the idea that the two technologies can strengthen themselves through mutual interaction.

When addressing end-to-end learning in Section 4, the authors focus on a scenario in which SML could potentially compete with optimisation. This is particularly true when a decision maker is satisfied with an approximate solution to the optimisation problem. Nevertheless, I agree with their conclusion that, given the current knowledge, it is unlikely to think that SML will be a "successful competitor to the classical optimization models and algorithms". Indeed, even today, many significant advances in SML continue to rely on optimisation. Beyond the examples discussed in Section 3.2, additional cases include: (i) least-quantile-of-squares regression (e.g., Bertsimas and Mazumde, 2014; Puerto and Torrejon, 2025), (ii) support vector regression (e.g., Martínez-Merino et al., 2025), (iii) multi-class classification (Bertsimas and Dunn, 2017; Blanco et al., 2020), and (iv) reclassification (Blanco et al., 2026). These developments in data science owe much to optimisation. Moreover, they illustrate that many SML tools keep relying on a loss function to be minimised, which quantifies the discrepancy between the model predictions and the actual values. As long as this remains the case, optimisation will retain a central role and the conceptual framework illustrated in Figure 1 will remain valid. Overall, the above references and those presented by the authors demonstrate that substantial research from the optimisation perspective continues to enhance the effectiveness of SML tools.

Although the paper focuses on optimisation and SML, many of its conclusions, in my view, extend to other branches of ML, such as unsupervised learning and reinforcement learning. Clustering provides a clear example. Optimisation has played a role in clustering for decades: as early as the 1980s, Kusiak (1984) was already reviewing optimisation models for clustering. Nowadays, some of the most popular clustering methods, such as k-means or k-medoids, are rooted in well-established optimisation models. Even more complex clustering problems, like districting—which incorporate a geographical content—continue to rely heavily on optimisation models and methods (see, e.g., Pomes et al., 2025, and the references there in).

Summing up, I fully agree with the authors when they conclude that optimisation and (S)ML will remain strong partners rather than competitors in the development of efficient and effective tools for decision-making. This enduring partnership stands to benefit both communities and the broader field as a whole.

References

Bertsimas, D., & Dunn, J. (2017). Optimal classification trees. *Machine Learning*, *106*(7), 1039–1082. doi:10.1007/s10994-017-5633-9

Bertsimas, D., & Mazumde, R. (2014). Least quantile regression via modern optimization. *The Annals of Statistics*, *42*(6), 2494–2525. doi:10.1214/14-AOS1223

Blanco, V., Japón, A., & Puerto, J. (2020). Optimal arrangements of hyperplanes for SVM-based multiclass classification. *Advances in Data Analysis and Classification*, *14*(1), 175–199. doi:10.1007/s11634-019-00367-6

Blanco, V., Japón, A., Puerto, J., & Zhang, P. (2026). Optimal probabilistic feature shifts for reclassification in tree ensembles. *Pattern Recognition*, *174*, 112980. doi:10.1016/j.patcog.2025.112980

De Lange, G. (2020). Combining machine learning, optimization for better decision-making. *Analytics, INFORMS Publication Record*, doi:10.1287/LYTX.2021.01.06

Hewitt, M., Ortmann, J., & Rei, W. (2022). Decision-based scenario clustering for decision-making under uncertainty. *Annals of Operations Research*, *315*(2), 747–771. doi:10.1007/s10479-020-03843-x

Huang, C., Tang, Z., Hu, S., Jiang, R., Zheng, X., Ge, D., Wang, B., & Wang, Z. (2025). ORLM: A customizable framework in training large models for automated optimization modeling. *Operations Research*, *73*(6), 2986–3009. doi:10.1287/opre.2024.1233

Huang, S., Yang, K., Qi, S., & Wang, R. (2024). When large language model meets optimization. *Swarm and Evolutionary Computation*, *90*, 101663. doi:10.1016/j.swevo.2024.101663

Kannan, R., Bayraksan, G., & Luedtke, J. R. (2024). Residuals-based distributionally robust optimization with covariate information. *Mathematical Programming*, *207*(1-2), 369–425. doi:10.1007/s10107-023-02014-7

Kusiak, A. (1984). Analysis of integer programming formulations of clustering problems. *Image and Vision Computing*, *2*(1), 35–40. doi:10.1016/0262-8856(84)90042-8

Martínez-Merino, L. I., Puerto, J., & Rodríguez-Chía, A. M. (2025). Ordered Weighted Average Support Vector Regression. *Expert Systems with Applications*, *274*, 126882. doi:10.1016/j.eswa.2025.126882

Pomes, A., Diglio, A., Nickel, S., & Saldanha-da-Gama, F. (2025). Multi-stage stochastic districting: Optimization models and solution algorithms. *Annals of Operations Research*, *346*(3), 2225–2251. doi:10.1007/s10479-024-06459-7

Puerto, J., & Torrejon, A. (2025). A fresh view on Least Quantile of Squares Regression based on new optimization approaches. *Expert Systems with Applications*, *282*, 127705. doi:10.1016/j.eswa.2025.127705

Said Salhi—Kent Business School, University of Kent, UK

This is an interesting and provocative discussion paper that opened up several challenges that could exist either at the intersection of optimisation and supervised ML where one can take advantage of the other though with different level of magnitude.

There are many applications especially those combinatorial optimisation problems that are considered NP hard problems with large problem size where ML can provide useful information by identifying promising as well as useless attributes both of which are critical as these could then be removed from the huge number of configurations and hence reduce the problem size for which optimisation can then flourish. This is well known in OR as reduction techniques or neighbourhood reduction in heuristic search arena. This contribution of ML for optimisation is well established in several combinatorial problems such as knapsack, location and routeing where some decision variables are set to zero or to one due to patterns discovered by ML leaving optimisation to tackle the rest of the decision variables. This, however, may not always be helpful if feasibility is not maintained especially in tight constrained problems. In addition, learning *via* heuristics/ML is now incorporated within B&B to tighten the whole process by improving upper bounds and turning infeasible solutions (i.e., Lower Bounds) into feasible ones. So, it is clear that ML somehow can enhance optimisation and many research is ongoing to foster this rich avenue. The other way around may appear at the first glance to be not obvious at all but that is not the case if one looks at the core of ML which is made up of some sort of optimisation, mostly continuous optimisation. For example, the k-mean is one of the ML but the algorithm itself is similar to Cooper's algorithm (Cooper, 1964) used to solve the multi-source Weber problem (or the planar p median problem) except it uses larger dimension instead of the two-dimensional space for locating the facilities. This locate-allocate procedure is now even much more improved than its original version produced by Cooper which is also used for the k-mean as part of ML (Drezner and Salhi, 2017). Optimisation is also used to identify good values for the parameters of some ML algorithms as well as being the hidden part of their engine. For instance, over two decades ago, embedding heuristic optimisation-based adaptive learning within neural network, one of the mostly currently adopted ML algorithm, was shown to be rather powerful for solving a class of classification using UCI repository of ML databases (Smithies et al. (2004).

It is also worth noting that whenever ML are included, optimality may not be guaranteed. This is the elephant in the room which is not made clearer to the public that the solution provided not only is unlikely to be optimal but also could turn out to be relatively poor. This side of the coin can be made clearer. There is however light at the end of the tunnel as there are hot research topics within ML, that are currently being explored to alleviate this issue and attempt to turn ML to a more robust tool by attempting to understand/incorporate the effect of big data, some of which could be misleading, while enhancing its quality through learning *via* some form of optimisation. Optimisation and particularly heuristic optimisation can play an important role in enhancing this part of ML. One of the issues one needs to be aware of is that ML rely heavily on past data to forge patterns etc. The question could arise if a new problem with different characteristics etc is to be solved, ML ought not to be the first option to tackle it as it will not be able to predict whereas in this particular case it is more relevant and also exciting for optimisation as new model need to be developed and then solved either exactly or *via* heuristics. In these circumstances one may argue that the former may utilise some benchmarks if these can be found either as existing solutions in companies or derived from heuristics. These benchmarks could obviously be used as target at the learning stage but these data may be limited, poor and hence mislead the ML so some thinking along these lines could be one way forward. Having said that, ML applications are immense and well received by users in most domains given its massive successes. This reinforces the old saying that it is better to have good solutions to the right problem than no solution at all or optimal ones for the wrong problem! In other words, optimisation has to embrace ML and be part of it so to enhance it and make it more robust and reliable by exploring keys aspects of optimisation like how to deal with tight constraints and soft/hard objectives for which ML may be at the moment relatively behind.

I would like to thank the authors for carefully and critically highlighting the several aspects and challenges that exist when trying to complement these two vast areas for which we are all grateful to have to our disposal. This discussion paper will, without any doubt, entice many researchers to explore research avenues that would not have been left in the dark and this will enhance not only the integration of optimisation and ML but also the visibility of OR in general.

References

Cooper, L. (1964). Heuristic Methods for location-allocation problems. *SIAM Review*, *6*(1), 37–53. doi:10.1137/1006005

Drezner, Z., & Salhi, S. (2017). Incorporating neighbourhood reduction for the solution of the planar p-median problem. *Annals of Operations Research*, *258*(2), 639–654. doi:10.1007/s10479-015-1961-y

Smithies, R., Salhi, S., & Queen, N. (2004). Adaptive hybrid learning for neural networks. *Neural Computation*, *16*(1), 139–157.

Kate Smith-Miles—University of Melbourne, Australia

What a timely discussion on the many facets of interaction between supervised learning and optimisation. As the world grapples with the implications of artificial intelligence (AI) solving problems previously tackled by humans, with and without technological assistance, it is natural to ask whether humans with expertise in solving optimisation problems will eventually be replaced with AI that can:

- accurately formulate a mathematical **model** from a written **problem** description (Li et al, 2023; Liang et al, 2026; Xiao et al. 2025);
- automatically select the best existing **algorithm** (and tune its **parameters or design components**) for solving a given model, given the instance characteristics (Chen et al., 2022; Hutter et al., 2009; Kerschke et al., 2019) and explain its choice (Smith-Miles et al., 2014);
- design test **instances** that can adequately "stress-test" an algorithm for trustworthiness (Lopes & Smith-Miles, 2013; Smith-Miles & Bowly, 2015; Smith-Miles & Hemert, 2011);
- or even find **solutions** to the problem without the need for any optimisation algorithm at all (Abassi et al., 2020; Sultana et al., 2022).

It is important to have an accurate assessment of the current state-of-the-art to understand whether such AI is imminent, whether it can be trusted, and furthermore, if it should be viewed as a *competitor* to human expertise, or a *partner* sitting alongside humans using a suite of algorithms (optimisation and AI) to find the best trusted solutions.

The authors have presented a useful framework for describing the opportunities for supervised learning (as one style of AI) to support or replace pillars of the generic optimisation process. They have arranged the literature within this common framework of the various stages: problem→instance→model→algorithm→solution. The references I have added above fit comfortably within this framework. Augmenting the reviewed literature with these additional references suggest to me that humans with expertise in optimisation will be required to stay in the loop a while longer! Nevertheless, we would be well advised to keep abreast of developments in AI and machine learning, since it is a rapidly developing space.

Personally, I was not expecting much success when my collaborators first proposed to use supervised learning to directly predict near optimal solutions to an optimisation problem (Abbasi et al., 2020) without the need for a model or any optimisation algorithm (cf. Figure 8, "end-to-end learning"). However, I was surprised by the impressive performance that can be achieved by training a machine learning model to recognise characteristics of good quality solutions.

While there are more references, beyond those listed above, that could be added to the paper within the presented framework, I'd like to focus my commentary on some recent advances of relevance to the section "Supervised learning for optimisation" that appear to have been overlooked, and provide some clarity about existing state-of-the-art and remaining opportunities.

Specifically, the authors write:

> *"Tests are usually performed on a set of synthetic instances which are commonly taken from the literature or randomly generated … The issue of what instances to use for the testing of algorithms is constantly discussed in the OR community, but no better ways than the ones mentioned above have been identified so far."*

In fact, there have been methods proposed to tackle the challenge of rigorously testing optimisation algorithms, following urging in the mid-1990s for better methods (Hooker, 1994, 1995). My contributions to this effort commenced with recognising that the machine learning community had been using supervised learning to learn about machine learning algorithm performance (so called "meta-learning), and the core ideas were generalisable to learn about other kinds of algorithms, including optimisation algorithms (Smith-Miles, 2009).

Based on the key idea that optimisation problem instances (like machine learning datasets) could be characterised by a feature vector that summarises their difficulty (Smith-Miles & Lopes, 2012), we have developed *Instance Space Analysis* (Smith-Miles et al., 2014; Smith-Miles & Muñoz, 2023): a methodology that uses supervised learning—of the relationship between instance characteristics and algorithm performance—to support:

1. creation of novel test instances required for rigorous algorithm testing;
2. insights into strengths and weaknesses of algorithms for different instance characteristics;
3. automated algorithm or parameter or model selection.

By projecting instances from a high dimensional feature space onto a 2D map we define and visualise an "instance space" of all possible test instances for an optimisation problem within a mathematically defined boundary, and show how comprehensively an optimisation algorithm has been tested. Within the instance space we can see gaps where new test instances should be designed, with target characteristics, to ensure trustworthy and unbiased testing with instances that are diverse, challenging, discriminating, and even real-world-like (Lopes & Smith-Miles, 2013). Evolutionary algorithms are used to fill these gaps to produce comprehensive suites of test instances (Smith-Miles et al., 2021; Smith-Miles & Bowly, 2015). The performance of algorithms (or models or parameter configurations) can also be viewed across the instance space to understand how performance depends on instance characteristics, cf. location in the instance space ((Liu et al., 2024). Automated algorithm selection is achieved using supervised learning in the instance space to predict the best performing algorithm for a given instance.

Instance Space Analysis (ISA) offers a solution to John Hooker's call for a more empirical science of algorithm testing. It is well accommodated within the framework of supervised learning for optimisation proposed by the authors. It remains an open opportunity to use ISA to rigorously "stress-test" how well AI methods can compete with classical OR methods for solving optimisation problems, or if the best results come from when they work in partnership!

References

Abbasi, B., Babaei, T., Hosseinifard, Z., Smith-Miles, K., & Dehghani, M. (2020). Predicting solutions of large-scale optimization problems via machine learning: A case study in blood supply chain management. *Computers & Operations Research*, *119*, 104941. doi:10.1016/j.cor.2020.104941

Chen, T., Chen, X., Chen, W., Heaton, H., Liu, J., Wang, Z., & Yin, W. (2022). Learning to optimize: A primer and a benchmark. *Journal of Machine Learning Research*, *23*(189), 1–59.

Hooker, J. N. (1994). Needed: An empirical science of algorithms. *Operations Research*, *42*(2), 201–212. doi:10.1287/opre.42.2.201

Hooker, J. N. (1995). Testing heuristics: We have it all wrong. *Journal of Heuristics*, *1*(1), 33–42. doi:10.1007/BF02430364

Hutter, F., Hoos, H. H., Leyton-Brown, K., & Stuetzle, T. (2009). ParamILS: An automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, *36*, 267–306. doi:10.1613/jair.2861

Kerschke, P., Hoos, H. H., Neumann, F., & Trautmann, H. (2019). Automated algorithm selection: Survey and perspectives. *Evolutionary Computation*, *27*(1), 3–45.

Li, Q., Zhang, L., & Mak-Hau, V. (2023). Synthesizing mixed-integer linear programming models from natural language descriptions. *arXiv Preprint arXiv:2311.15271*.

Liang, K., Lu, Y., Mao, J., Sun, S., Yang, C., Zeng, C., … Teo, C. P. (2026). LLM for large-scale optimization model auto-formulation: A lightweight few-shot learning approach. *arXiv Preprint arXiv:2601.09635*.

Liu, C., Smith-Miles, K., Wauters, T., & Costa, A. M. (2024). Instance space analysis for 2D bin packing mathematical models. *European Journal of Operational Research*, *315*(2), 484–498. doi:10.1016/j.ejor.2023.12.008

Lopes, L., & Smith-Miles, K. (2013). Generating applicable synthetic instances for branch problems. *Operations Research*, *61*(3), 563–577. doi:10.1287/opre.2013.1169

Smith-Miles, K. A. (2009). Cross-disciplinary perspectives on meta-learning for algorithm selection. *ACM Computing Surveys*, *41*(1), 1–25. doi:10.1145/1456650.1456656

Smith-Miles, K., & van Hemert, J. (2011). Discovering the suitability of optimisation algorithms by learning from evolved instances. *Annals of Mathematics and Artificial Intelligence*, *61*(2), 87–104. doi:10.1007/s10472-011-9230-5

Smith-Miles, K., & Lopes, L. (2012). Measuring instance difficulty for combinatorial optimization problems. *Computers & Operations Research*, *39*(5), 875–889. doi:10.1016/j.cor.2011.07.006

Smith-Miles, K., Baatar, D., Wreford, B., & Lewis, R. (2014). Towards objective measures of algorithm performance across instance space. *Computers & Operations Research*, *45*, 12–24. doi:10.1016/j.cor.2013.11.015

Smith-Miles, K., & Bowly, S. (2015). Generating new test instances by evolving in instance space. *Computers & Operations Research*, *63*, 102–113. doi:10.1016/j.cor.2015.04.022

Smith-Miles, K., Christiansen, J., & Muñoz, M. A. (2021). Revisiting where are the hard knapsack problems? via instance space analysis. *Computers & Operations Research*, *128*, 105184. doi:10.1016/j.cor.2020.105184

Smith-Miles, K., & Muñoz, M. A. (2023). Instance space analysis for algorithm testing: Methodology and software tools. *ACM Computing Surveys*, *55*(12), 1–31. doi:10.1145/3572895

Sultana, N., Chan, J., Sarwar, T., & Qin, A. K. (2022). Learning to optimise general TSP instances. *International*

Journal of Machine Learning and Cybernetics, *13*(8), 2213–2228. doi:10.1007/s13042-022-01516-8

Xiao, Z., Xie, J., Xu, L., Guan, S., Zhu, J., Han, X., … Zhang, D. (2025). A survey of optimization modeling meets LLMs: Progress and future directions. *arXiv Preprint arXiv:2508.10047*.

Daniele Vigo—Department of Electrical, Electronic and Information Engineering "Guglielmo Marconi", University of Bologna, Italy

I read this paper with great interest and I see it as a valuable instrument to stimulate the future research on the integration between optimisation and learning, and, even more importantly, to favour the discussion and collaboration between different scientific communities, which may results in relevant future advances in the support of decision making.

Both Operations Research and Machine Learning are indeed mature disciplines, which have fully demonstrated their capability of providing valuable and effective toolsets to support decision making. However, several shortcomings, such as the rigidity of the underlying assumptions required by the construction of solvable optimisation models, or the difficulty in the generalisation of prediction methods, often constitute a major obstacle in their actual deployment in the real world context. Most optimisation problems which arise in practice have characteristics which are hardly compatible with the settings in which optimisation algorithms can quickly provide good, if not optimal, solutions: for example input data are generally non-deterministic, and objectives and constraints are either not fully defined or are non-linear. On the one hand these issues are generally not an obstacle for the development of machine learning approaches but, on the other hand, end-to-end methods based uniquely on machine learning algorithms often produce lower quality solutions when applied to well-cast problems. It is, therefore, clear the potential benefit of the integration between optimisation and learning, that may lead to approaches better suited to capture the difficulties of real-world problems which, at the same time, may provide better solutions to decision makers. However, to make step forward in this direction the two research communities need to collaborate and not to compete as it is often seen in the current literature.

A fundamental prerequisite to favour the collaboration between disciplines, is to create a common language, to highlight the potentials and the shortcomings of either discipline, and to define correct ways in which the contribution of different methods may be assessed.

In my view (which is that of a member of the optimisation community), the authors have indeed performed an excellent job in this respect, by providing a clear, easy to read, and comprehensive illustration of the main areas in which supervised learning and optimisation may be integrated within frameworks which better suits the need of modern decision support systems.

Clearly, the space limitation of an article has prevented to deepen the analysis and the illustration of some important details, as for example, the computational overhead associated with feature extraction, which often reduces the effective use of prediction algorithms in real-time contexts. Nevertheless, the paper highlights several inspiring ideas, such as the use of machine learning to create and improve benchmarks, and points to the relevant literature where the reader can broaden the understanding.

Probably, another unavoidable consequence of the space limitation is that that most of the emphasis of the paper is devoted to illustrating what machine learning can do to improve optimisation rather than to highlight which are the areas in which better optimisation algorithms may have a role in improving machine learning methods.

In conclusion, I think this paper constitutes an important contribution in creating the ground in which the future innovation of decision making will root, and I am sure it will stimulate new and rigorous research in this direction.

Response to Commentaries

A. Mor, C. Orsenigo—Department of Management, Economics and Industrial Engineering Politecnico di Milano, Italy

M.G. Speranza—Department of Economics and Management University of Brescia, Italy

We would like to express our sincere gratitude to the Journal of the Operational Research Society for accepting our discussion paper, and to Professor Said Salhi for his meticulous coordination of the entire process, from handling the submission to inviting and collecting commentaries from leading researchers. We are also deeply grateful to the colleagues who agreed to read and comment on our work. Their insights and stimulating reflections, along with the relevant references they provided, significantly complement our original study.

Several key themes emerged from these commentaries. First, they confirm that the topic is highly timely and that the central question of our paper resonates with many in the field. Second, there is a shared conviction that, while machine learning is evolving rapidly, real improvements in decision support tools will stem from a partnership between optimisation and supervised learning. Furthermore, as noted by some contributors, although we focused on supervised learning for clarity, much of our discussion is equally applicable to other machine learning paradigms.

An interesting topic that emerged in the comments over which we debated for its inclusion in the paper is inverse optimisation. By integrating observed behaviour into the model objective or constraints, inverse optimisation ensures that the resulting mathematical formulations are not just theoretically sound, but deeply aligned with the underlying real-world processes they aim to represent.

As highlighted by several comments, another domain worth of discussion is Large Language Models (LLMs), which are already showing potential in automating the formulation of mathematical models from natural language descriptions. LLMs may play a transformative role in algorithm selection, design, and reasoning, with the potential to redefine the OR-ML interface. We believe that LLMs should be viewed as powerful enablers rather than (potential) replacements for human expertise. While they can significantly lower the entry barrier for modelling and accelerate the prototyping phase, the human component remains fundamental in decision-making. Human intuition, ethical judgement, and ability to interpret context-specific nuances are indispensable for validating model assumptions and ensuring that the outputs of hybrid OR-ML systems are both responsible and actionable.

This discussion reinforces our belief that optimisation and machine learning are destined to become strong partners, being inspired by different principles that can promote reciprocal support. Optimisation emphasises the clear definition of the problem structure and seeks for provable properties, such as the convergence of the algorithm to an optimal, or near-optimal, solution. Supervised learning embraces uncertainty and creates the model by learning from data. Yet both are, at their core, powerful approaches for making decisions under complexity. As each field advances, a shift is needed to foster their deepening synergy: formulation and learning should no longer be viewed as separate processes, but rather as components of a unified adaptive framework. To this end, the adoption of a shared mindset is essential, where rigour meets empiricism and constraints are no longer imposed on prior assumptions but are instead informed and refined by observed data, making them more realistic and relevant to the problem at hand. The convergence of optimisation and supervised learning to a common ground will ultimately enable the development of decision-making systems that are both data-aware and structurally sound, combining adaptability with reliability.

We hope that our discussion paper, enriched by the above comments of prominent researchers, has initiated a dialogue that will continue to evolve.