

Data-Driven Combinatorial Optimisation

Emma Frejinger^{*1}, Andrea Lodi^{*2}, Michele Lombardi^{*3}, and
Neil Yorke-Smith^{*4}

1 University of Montreal, CA. emma.frejinger@umontreal.ca

2 Cornell Tech – New York, US. andrea.lodi@cornell.edu

3 University of Bologna, IT. michele.lombardi2@unibo.it

4 TU Delft, NL. n.yorke-smith@tudelft.nl

Abstract

Machine learning’s impressive achievements in the last decade have urged many scientific communities to ask if and how the techniques developed in that field to leverage data could be used to advance research in others. The combinatorial optimisation community is one of those, and the area of data-driven combinatorial optimisation has emerged. The motivation of the seminar and its design and development have followed the idea of making researchers both in academia and industry belonging to different communities – from operations research to constraint programming, from artificial intelligence to machine learning – communicate, establish a shared language, and ultimately (try to) set the roadmap for the development of the field.

Seminar 23–28 October 2022 – <http://www.dagstuhl.de/22431>

2012 ACM Subject Classification Theory of computation → Constraint and logic programming; Computing methodologies → Machine learning; Theory of computation → Mathematical optimization; Theory of computation → Reinforcement learning

Keywords and phrases combinatorial optimisation, constraint programming, machine learning, Mixed integer programming, operations research, Reinforcement learning

Digital Object Identifier 10.4230/DagRep.12.10.166

1 Executive Summary

Emma Frejinger (University of Montreal, CA)

Andrea Lodi (Cornell Tech – New York, US)

Neil Yorke-Smith (TU Delft, NL)

License  Creative Commons BY 4.0 International license
© Emma Frejinger, Andrea Lodi, and Neil Yorke-Smith

In the last five years, an area now being influenced in a new way by machine learning (ML) is combinatorial optimisation (CO). Combinatorial optimisation is studied for both its importance in theory, since CO problems are NP-hard problems, and for its importance in real-world decisions, for example, planning drivers and routes for a fleet of delivery vehicles. CO problems are studied in operations research (OR) and also traditionally in symbolic artificial intelligence (AI) such as constraint programming (CP) and satisfiability modulo theories.

This Dagstuhl Seminar built on the fast-growing interest in combining ML with ‘traditional’ AI methodologies like CP, and with OR more generally [1, 2]. Surveying the scattered initiatives, the seminar had the ambition to set the agenda for constraint-based ‘Combinatorial Optimisation 2.0’. Historically, several communities have focussed on different approaches to CO, mostly in a disjoint manner. This division between, on the one hand, the OR and

* Editor / Organizer



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Data-Driven Combinatorial Optimisation, *Dagstuhl Reports*, Vol. 12, Issue 10, pp. 166–174

Editors: Emma Frejinger, Andrea Lodi, Michele Lombardi, and Neil Yorke-Smith



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

symbolic AI communities, and on the other, the ML and functional AI communities, is historically strong. While in recent years a dialogue between symbolic and functional AI communities has emerged, there remains too little connection between the discrete OR and ML communities.

The seminar was organised by Emma Frejinger (Canada), Andrea Lodi (USA), Michele Lombardi (Italy) and Neil Yorke-Smith (Netherlands). Michele was unable to attend in person, due to last minute circumstances, and joined plenary parts of the seminar online. Similarly, it was necessary for Pierre-Luc Bacon to give his tutorial remotely.

Seminar Overview

The seminar opened with four tutorials, whose abstracts are given in this report, on the topics of CP (by Tias Guns), mixed (non)-linear integer programming (MIP) (by Ruth Misener), end-to-end ML for CO (by Ferdinando Fioretto), and reinforcement learning (RL) (by Pierre-Luc Bacon).

The seminar included a set of informal short introductory and topical talks, and sessions of collaborative planning. The overarching questions that structured this planning are, on the one hand, (1) how ML can help in modelling or solving CO problems – or both modelling and solving – and in particular constraint-based models and solving; and on the other hand, (2) how CO can help in tasks approached using ML, including ML training and algorithms. Then, (3) what problems and tasks can be addressed (only) by the synergistic combinations of these methodologies?

Through discussions, the participants identified jointly six topics to be approached in smaller working groups: (i) self-supervised representation learning for combinatorial optimisation, (ii) uncertainty, prediction, optimisation and decision-focussed learning, (iii) OR for ML, (iv) vehicle routing and the role of ML, (v) ML-augmented MIP solvers, and (vi) fairness. The groups discussed challenges, existing work and identified open research questions with promising future avenues at the intersection between OR and ML. The working groups are summarised below.

The outcomes of the seminar in furthering the development of a community at the intersection of OR and ML are expected to be felt in the coming couple of years. Already, however, there are tangible outcomes in terms of roadmap ideas, an open online discussion forum (Slack)¹, multiple new collaborations, and a research grant submitted. A special issue of the journal *Frontiers in Applied Mathematics and Statistics* is organised by one of the participants.

The scientific programme was beautifully facilitated by the surroundings and academic services of Schloss Dagstuhl. Further, on the opening evening of the seminar, volunteers among the participants took part in “slide bingo”. During this humorous session they improvised presenting the slides of others. On Wednesday afternoon, the participants took a walk to a nearby village in unexpectedly fine sunny weather for October.

¹ On Slack, ML<>CO

Reflections on the Week

All communities present at the seminar found benefit from the interactions and discussions. Meeting in person at the scale of a Dagstuhl Seminar was much appreciated! Participants were aware of the differing emphases, mindsets, and publication practices of different communities. In general, it was felt that strengthening the connection between ML and OR helps in bridging the gap between predictive and prescriptive analytics, which can benefit industrial or government actors and citizens alike.

Working Groups

In this section, we briefly summarise the discussions in the six working groups.

Self-supervised representation learning for combinatorial optimisation

This working group enjoyed lively discussions around the concept of a “foundational model” for CO. The motivation is to avoid retraining from scratch when there is a relatively small change. The group discussed transfer learning in terms of problem formulation, downstream task and instance distributions.

The group identified open questions:

- What is the equivalent of saving models/checkpoints in ML or natural language processing (NLP) in data-driven CO (DDCO)?
- Can we share pre-trained models to generate SAT/CP/MIP embeddings without training again?
- NLP has the concept of a tokeniser that preprocesses the text before it gets fed into the network; in DDCO we would need similar pre processors that transfer the problem instances into the model’s expected (graph) structure.
- What is the equivalent of “large” aspect from large language NLP models for DDCO?
- Is there a (super) GLUE benchmark equivalent for DDCO?

Uncertainty, prediction, optimisation and decision-focussed learning

Decision-focussed learning aims at training prediction models against a loss reflecting the quality of decisions instead of a classic prediction loss. This working group weighed up the questions: when is decision-focussed learning (DFL) better than (traditional) alternatives? The group gave energy into thinking about stochastic formulations, data perturbation and interpretability of DFL. The group also identified a connection with RL, in particular contextual bandits (single-stage decision). Since there has been some confusion around the terminology, the group recommended to use “decision-focussed learning” instead of “predict and optimise” or “predict + optimise”.

The group wrote down example problems for three settings of decision focussed learning for CO. First, unknown parameters in the objective. This is the most studied case and there are several applications in the literature. Second, unknown parameters in the right-hand side of the constraints. For example, transport network planning where demand predictions occur in capacity constraints. Third, unknown parameters in the left-hand side of the constraints. For example, healthcare scheduling problems where treatment durations are predicted and should not exceed a given schedule length.

Operations research for machine learning

This working group was provoked by the feeling in the OR community that the ML community is seldom happy with discrete optimisation. In other words, how can CO have an influence on problems that generally come from the ML community? Those from OR background in the group expressed that they want to make real contributions to ML.

The group proceeded to outline three major obstacles:

- Scalability. OR methods tend to be limited when dealing with extremely large datasets that are often associated with ML applications.
- Optimality. OR methods have been designed in general to provide guarantees. Computing confidence bounds for ML applications would be excellent but one major obstacle is that the loss function is computed over samples from an unknown distribution. In other words, there is uncertainty with respect to the real objective function, which would require a re-interpretation of what confidence bounds are.
- Software. Whereas the ML community is used to working with self-installing open-source software, the OR community uses more cumbersome, often commercial software.

The recommendation was for research into a new generation of optimisation-based heuristics. The group identified four areas in which discrete optimisation methods are likely relevant: (i) optimal transport problems, (ii) neural-network verification, (iii) the broad area of fairness, explainability and interpretability, and (iv) training for Gaussian processes with tree kernels.

Vehicle routing and the role of machine learning

This working group ascertained exciting research on using ML to help solve routing problems. Two main aspects are, first, that the current state of (deployed) routing software has no idea whether it has seen a problem before, the types of problems being solved, and so forth; and second, anticipating the future – for instance dynamic settings, demand estimation, service time estimation, and so forth – would make the solution of routing problems even more relevant in practice. The working group felt that leveraging ML in both aspects could lead to significant improvements.

The group identified open questions:

- Does the ML model output individual actions or does it output an instance-specific heuristic?
- Can we learn insights about the problem from the predictions?
- Where does or could deep RL work best?
- Can we make a unified routing model (a ‘foundation model’)?

Machine learning augmented MIP solvers

This working group started from the general questions: which is the big challenge in MIP solving? And, will ML-augmented MIP solvers be ever significantly better than improved versions of the current solvers?

The group found that one significant motivation to go for ML-augmented MIP is ‘democratisation’: ML could allow the more general use of MIP technology by automatising some steps of MIP development and solution that depend on the specific a class of problems in hand without requiring the intervention of experts in the loop. Such a democratisation would require the definition of a robust pipeline on how to learn – from data – tasks like branching, cutting, preprocessing, etc. on the specific class of instances in hand, i.e., characteristic

of the application one wants to solve. Such a robust pipeline does not exist yet though there is strong evidence of successful stories where ML-augmented tasks are performed more efficiently than in classical MIP solvers. Some of those successful examples have been already integrated into the solvers, even commercial ones.

The group identified a number of interesting directions, as yet unexplored in the fast-moving subfield of ML-for-MIP. Among these are hypothesis generation, defining appropriate performance metrics, learning for cutting plane generation and selection. The group also discussed benchmark libraries.

Fairness

This working group sought to learn more about data-driven CO models for fairness. The group recognised that an issue with fairness is already in its definition. While in the ML community, the concept of fairness is related to a tradeoff between overall accuracy and group accuracy, in CO for decision making, there is not a clear definition of fairness.

The group discussed an application in the online scheduling of radiologists and neuro-radiologists of CT scans. In this context, because of the scarcity of the resources, fairness is associated with their correct and ‘fair’ use.

Fairness at large is also related to explainability and interpretability and the working group discussed the use of classical CO methods that tend to be more interpretable of the ML ones that are often perceived as black-boxes. Further, a potential important area in this context is that of integrating ML and OR to achieve a higher level of explainability, for example by improving methods like decision trees and ML classification algorithms.

References

- 1 Yoshua Bengio, Andrea Lodi, Antoine Prouvost: *Machine learning for combinatorial optimization: A methodological tour d’horizon*. Eur. J. Oper. Res. 290(2): 405-421 (2021). <https://doi.org/10.1016/j.ejor.2020.07.063>
- 2 James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, Bryan Wilder: *End-to-End Constrained Optimization Learning: A Survey*. IJCAI 2021: 4475-4482. <https://doi.org/10.24963/ijcai.2021/610>

2 Table of Contents

Executive Summary

Emma Frejinger, Andrea Lodi, and Neil Yorke-Smith 166

Overview of Talks

End-to-end constrained optimization learning

Ferdinando Fioretto 172

Data-driven combinatorial optimisation, with a CP flavour

Tias Guns 172

Machine learning for mathematical optimization and mathematical optimization
for machine learning

Ruth Misener 173

An overview of reinforcement learning and learning for control

Pierre-Luc Bacon 173

Participants 174

3 Overview of Talks

3.1 End-to-end constrained optimization learning

Ferdinando Fioretto (Syracuse University, US)

License © Creative Commons BY 4.0 International license
© Ferdinando Fioretto

Main reference James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, Bryan Wilder: “End-to-End Constrained Optimization Learning: A Survey”, in Proc. of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021, pp. 4475–4482, ijcai.org, 2021.

URL <https://doi.org/10.24963/ijcai.2021/610>

This tutorial reviews the recent advancements made in using constrained optimization to incorporate structural information and domain knowledge into machine learning models. We start by reviewing how to convert optimization into differentiable layers to use in machine learning models. Such integration enables to enforce structural information and domain knowledge into machine learning models. Next, we focus on extending this setting by integrating constrained optimization to enforce structure in the outputs of learned embeddings, leading to end-to-end decision-focused learning, that trains models to directly optimize the performance in targeted applications. Finally, we review techniques to learn constrained optimization surrogates by leveraging a distribution of optimization problems and their solutions leading to enhanced optimization modeling technology for operations research decision tasks. The tutorial concludes with a discussion of challenges and open questions.

3.2 Data-driven combinatorial optimisation, with a CP flavour

Tias Guns (KU Leuven, BE)

License © Creative Commons BY 4.0 International license
© Tias Guns

Joint work of Tias Guns, Rocs Canoy, Jayanta Mandi, Maxime Mulamba, Victor Bucarey Lopez, Emilio Gamba, Ignace Bleukx, Michelangelo Diligenti, Michele Lombardi, Bart Bogaerts

Main reference Mohit Kumar, Samuel Kolb, Tias Guns: “Learning Constraint Programming Models from Data Using Generate-And-Aggregate”, in Proc. of the 28th International Conference on Principles and Practice of Constraint Programming, CP 2022, July 31 to August 8, 2022, Haifa, Israel, LIPICs, Vol. 235, pp. 29:1–29:16, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

URL <https://doi.org/10.4230/LIPICs.CP.2022.29>

We first provided a general overview of different constraint solving technologies, including constraint programming. We then discussed the general trend of using machine learning to either learn (part of) the model, the problem specification, or to learn how to obtain solutions faster. Both topics were well covered in the seminar overall.

The rest of the talk then focussed on recent work in learning part of the model, more specifically learning the constraints using passive or (inter)active constraint learning, as well as learning the objective using pre-training networks and integrating constraint solving in the inference, with examples in visual constraint solving, learning preferences in vehicle routing and more.

The tutorial ended with how using learning part of the model also increases the need for explainable models, both at the machine learning and the constraint solving side.

3.3 Machine learning for mathematical optimization and mathematical optimization for machine learning

Ruth Misener (Imperial College London, GB)

License © Creative Commons BY 4.0 International license
© Ruth Misener

Joint work of Ruth Misener, Calvin Tsay, Francesco Ceccon, Jordan Jalving, Joshua Haddad, Alexander Thebelt, Carl Laird, Miten Mistry, Jan Kronqvist, Radu Baltean, Pierre Bonami, Andrea Tramontani

Main reference Francesco Ceccon, Jordan Jalving, Joshua Haddad, Alexander Thebelt, Calvin Tsay, Carl D. Laird, Ruth Misener: “OMLT: Optimization & Machine Learning Toolkit”, CoRR, Vol. abs/2202.02414, 2022.

URL <https://arxiv.org/abs/2202.02414>

We consider how machine learning can be used for expediting mathematical optimization solvers. We also consider how mathematical optimization can contribute to machine learning. This presentation is biased towards the kind of optimization I understand, so it mostly concerns mixed-integer nonlinear optimization.

3.4 An overview of reinforcement learning and learning for control

Pierre-Luc Bacon

License © Creative Commons BY 4.0 International license
© Pierre-Luc Bacon

Main reference Evgenii Nikishin, Romina Abachi, Rishabh Agarwal, Pierre-Luc Bacon: “Control-Oriented Model-Based Reinforcement Learning with Implicit Differentiation”, in Proc. of the Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 – March 1, 2022, pp. 7886–7894, AAAI Press, 2022.

URL <https://ojs.aaai.org/index.php/AAAI/article/view/20758>

The success of the field of reinforcement learning hinges upon its multidisciplinary nature. This tutorial highlights significant contributions from other disciplines, such as optimal control, operations research and simulation, to build a more robust theoretical understanding of modern algorithms. We begin our overview by studying the class of temporal difference learning algorithms as an application of the stochastic approximation method. We then see how sample average approximation and the ‘stochastic counterpart’ method in stochastic optimization can offer insights into the class of fitted value methods behind the latest advances in deep reinforcement learning. Using the tools for sensitivity analysis in simulation, we then explore how ideas in derivative estimation give rise to the class of policy gradient methods in reinforcement learning. Finally, we conclude our tour d’horizon by broadening our perspective on reinforcement learning by merging optimization techniques in optimal control with learning through end-to-end or decision-aware methods.

Participants

- Karen Aardal
TU Delft, NL
- Claudia D’Ambrosio
Ecole Polytechnique –
Palaiseau, FR
- Bistra Dilkina
USC – Los Angeles, US
- Ferdinando Fioretto
Syracuse University, US
- Emma Frejinger
University of Montreal, CA
- Maxime Gasse
Polytechnique Montréal, CA
- Stefano Gualandi
University of Pavia, IT
- Oktay Gunluk
Cornell University – Ithaca, US
- Tias Guns
KU Leuven, BE
- Serdar Kadioglu
Brown University –
Providence, US
- Lars Kotthoff
University of Wyoming –
Laramie, US
- Hoong Chuin Lau
SMU – Singapore, SG
- Pierre Le Bodic
Monash University –
Clayton, AU
- Andrea Lodi
Cornell Tech – New York, US
- Marco Lübbecke
RWTH Aachen, DE
- Sofia Michel
NAVER Labs Europe –
Meylan, FR
- Andrea Micheli
Bruno Kessler Foundation –
Trento, IT
- Ruth Misener
Imperial College London, GB
- Laurent Perron
Google – Paris, FR
- Sebastian Pokutta
Zuse Institut Berlin, DE
- Louis-Martin Rousseau
Polytechnique Montréal, CA
- Helge Spieker
Simula Research Laboratory –
Oslo, NO
- Kevin Tierney
Universität Bielefeld, DE
- Pashootan Vaezipoor
University of Toronto, CA
- Pascal Van Hentenryck
Georgia Institute of Technology –
Atlanta, US
- Stefan Voß
Universität Hamburg, DE
- Neil Yorke-Smith
TU Delft, NL
- Yingqian Zhang
TU Eindhoven, NL

