Can Machine Learning Help in Solving Cargo Capacity Management Booking Control Problems?

EMMA FREJINGER

Associate Professor, Department of Computer Science and Operations Research, Université de Montréal

Joint work with Justin Dumouchelle and Andrea Lodi

arXiv: 2102.00092

IPAM Workshop, February 23, 2021

PARCEL DELIVERY IS BOOMING

Online retail is driving growth in the parcel delivery industry

Global parcels market approx. **US$430 bn** in 2019

Global small-package volume > **100 bn** in 2020, forecasted to **double by 2026**

In the US, transportation accounts for 27% of **GHG emissions** of those 41% emitted by trucks and 9% by aircrafts

Trucking: **empty vehicles** for a significant share of vehicle miles travelled (approx. 5-20%)

In the US, **congestion** costs trucking over US$63.4 bn annually

Sources: Pitney Bowes Index and American Transportation Research Institute — ATRI

IMPROVE CAPACITY USAGE, DECREASE NEGATIVE ENVIRONMENTAL IMPACTS, DIFFERENTIATE SERVICES AND PRICES

Revenue Management

Manage demand

Structural, price and quantity decisions

Supply Chain Management

Manage supply

Our focus lies on revenue management **quantity decisions** for **cargo capacity management**

Less studied than the passenger counterpart

Unlike passenger case, capacity is defined in several dimensions (e.g., weight and volume), capacity usage typically uncertain until the time of loading, flexible routing decisions

Talluri and Van Ryzin (2005)

OUTLINE

A stylized view of cargo capacity management and the booking control problem

Related work

Exploratory work: machine learning for combinatorial optimization

Distribution logistics application and results

Conclusion

A STYLIZED VIEW OF CARGO CAPACITY MANAGEMENT

Planning horizon

Allotment contracts

A contract fixes a customer's shipping rate and amount of reserved capacity for a given period

Infrequent bidding / negotiation process

Spot market

Booking requests with different characteristics occur continuously

Decision: accept / reject

Time of request, weight, volume, revenue, destination, etc.

Cargo capacity

Capacity, e.g., for a given, day or departure

Objective: maximize **profit** – **revenue** subtracting the **cost** (e.g., transportation and excess demand costs)

Problem definition depends e.g., on scope (allotment / spot market, network / single leg) and assumptions about uncertainty

See, for example, Levin et al. (2012)

BOOKING CONTROL PROBLEM — SPOT MARKET

Illustrative example – distribution logistics

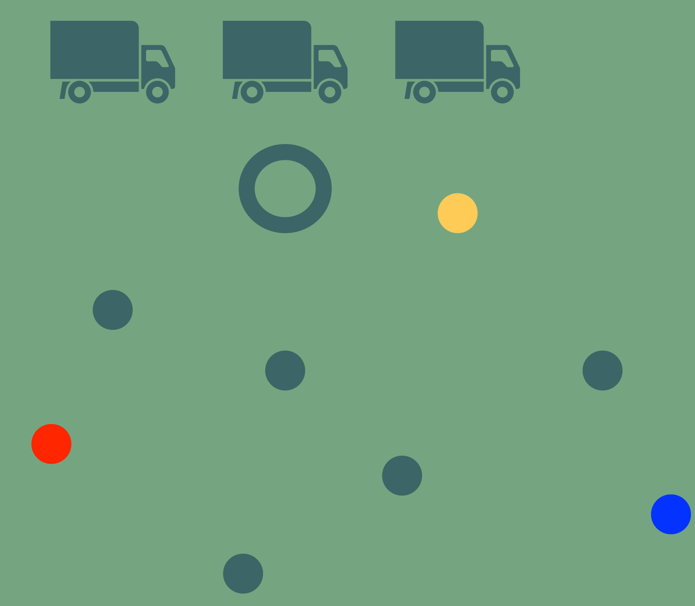

Request $j \in \mathcal{N}$ with prob. λ_j^t and revenue p_j

No request with prob. λ_0^t

Decision: $u_j^t \in \{0,1\}$

State: nb. of accepted requests before t, $\mathbf{w}_t \in \mathbb{R}^n$

Cost of solution to operational decision-making problem $\Gamma(\mathbf{w}_{T+1})$

Time $t = 1,2\ldots,T,T+1$

Profit maximization – finite-horizon Markov Decision Process

$$V_t(\mathbf{w}_t) = \lambda_0^t V_{t+1}(\mathbf{w}_t) + \sum_{j \in \mathcal{N}} \lambda_j^t \max_{u_j^t \in \{0,1\}} \{p_j u_j^t + V_{t+1}(\mathbf{w}_t + u_j^t \mathbf{e}_j)\}, \quad t = 1,\ldots,T$$

$$V_{T+1}(\mathbf{w}_{T+1}) = \Gamma(\mathbf{w}_{T+1})$$

BOOKING CONTROL PROBLEM — SPOT MARKET

Illustrative example

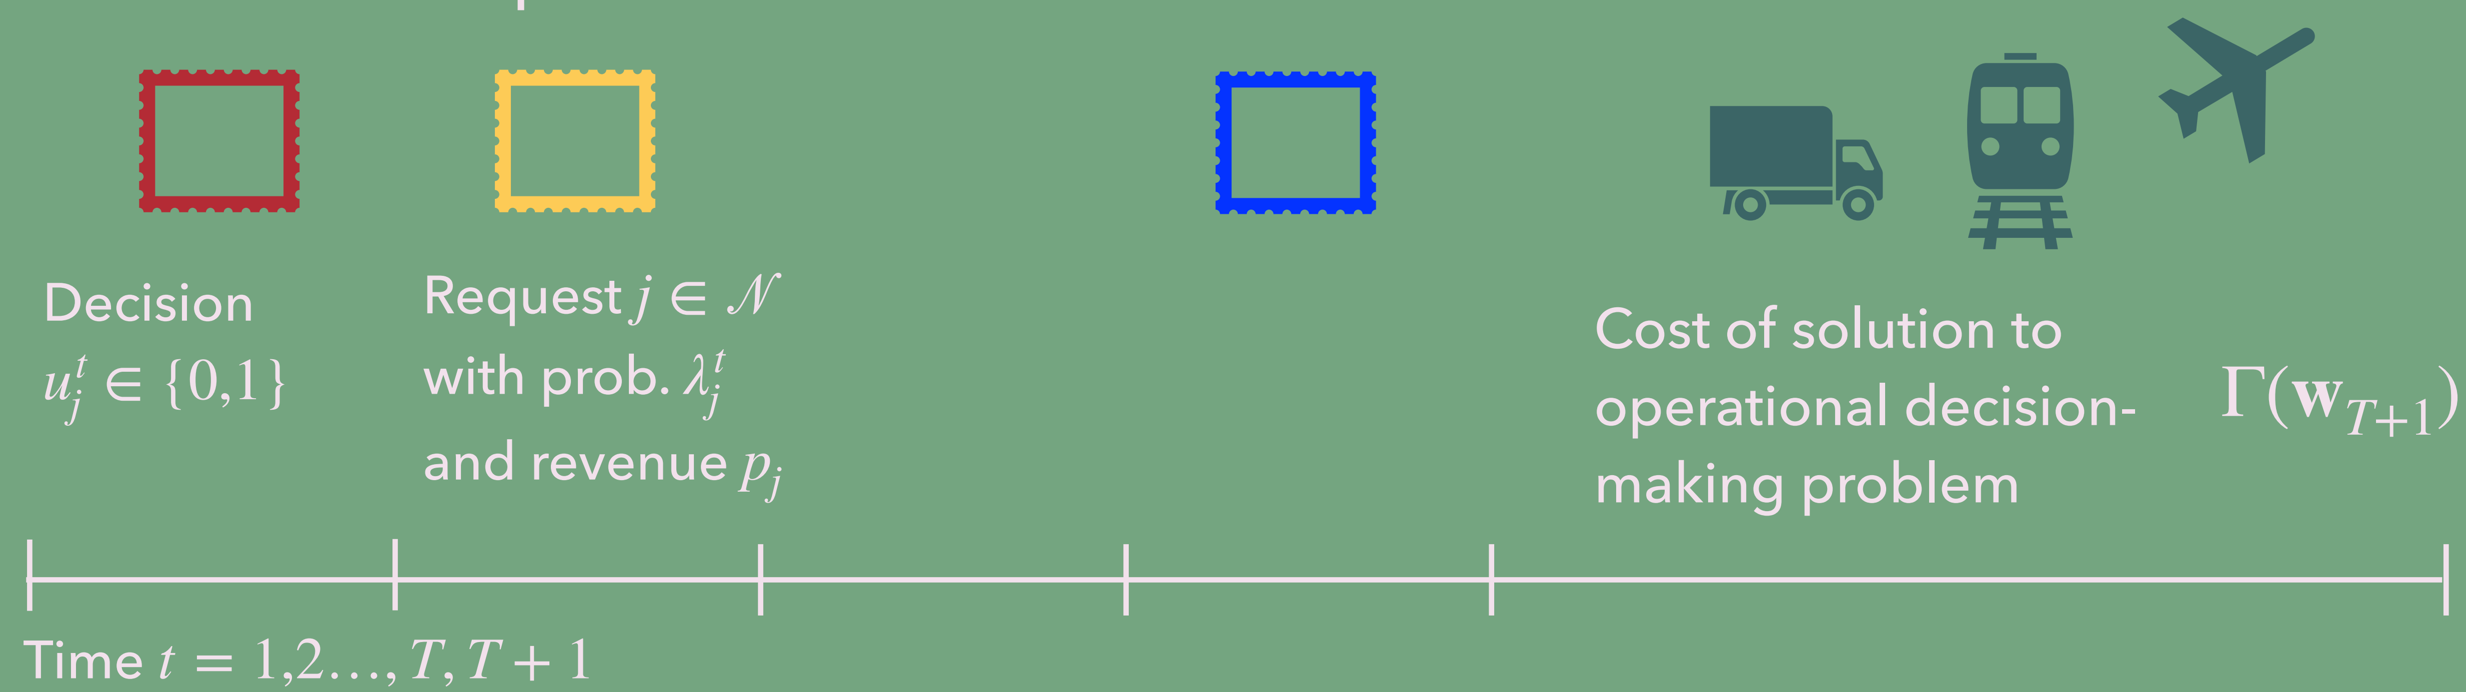

OBSERVATIONS

- End-of-horizon problem: combinatorial optimization (CO) problem, e.g., multidimensional bin-packing or routing
- High-dimensional state space, MDP intractable
- Detailed solution to the CO problem is not relevant
- Booking control under imperfect information and small computing time budget

BID-PRICE POLICIES FOR CARGO BOOKING CONTROL

Exact

MDP formulation

Threshold policy: accept request if its revenue is greater than expected opportunity cost of capacity used by the request

Linear programming (LP) formulation of optimal value function at initial state (Puterman, 1994, Adelman, 2007)

Feasible solutions provide upper bounds on optimal value functions and optimal total expected profit

Approximations

Static bid-price policy

Widely used in practice

LP ignores temporal aspects

LP resolved frequently

Williamson (1992), Talluri and van Ryzin (1998)

Dynamic bid-price policy

LP based on value function approximated with linear-in-parameters architectures

Levina et al. (2011), Barz and Gartner (2016) airline network cargo capacity management

Bid-price policy that depends on capacity usage

Lagrangian relaxation of capacity constraints in LP, Lagrangian multipliers dynamically updated

Levin et al. (2012) airline, allotments and spot market

Reinforcement learning (RL)

- Abundance of literature showing success in solving hard decision-making problems
- Reinforcement learning in revenue management: limited to **passenger** seat allocation problems in airline industry (Bondeaux et al., 2020, Gosavi et al.,2002, and Lawhead and Gosavi, 2019)
- Algorithms typically based on simulating trajectories of the system
  - Challenge for cargo booking control: the end-of-horizon problem is costly to solve which makes simulation-based algorithms prohibitively costly

Machine learning for heuristically solving CO problems

- Surge of studies, mostly focused on deterministic problems, survey Bengio et al. (2021)
- Supervised learning for predicting characteristics of the solutions (not full solution)
  - Fischetti and Fraccaro (2017): predict optimal solution value
  - Larsen et al. (2021): predict description of solutions to the second-stage problem in a two-stage stochastic program without generating second stage scenarios online. Cargo capacity management application.

EXPLORATORY WORK

Objective: high-quality booking control policy at low online computing cost

Idea: Use RL or approximate dynamic programming (ADP) to solve the booking control problem with an approximation of $\Gamma(\mathbf{w}_{T+1})$ that can be computed in very short time

In this exploratory work we use off-the-shelf ML/RL/ADP algorithms.

Approximation of $\Gamma(\mathbf{w}_{T+1})$

Mapping from *n* (state) dimensions to *m* $g : \mathbb{R}^n \to \mathbb{R}^m$

Approximation $\phi : \mathbb{R}^m \to \mathbb{R}\_$

$$\tilde{V}_t(\mathbf{w}_t) = \lambda_0^t \tilde{V}_{t+1}(\mathbf{w}_t) + \sum_{j \in \mathcal{N}} \lambda_j^t \max_{u_j^t \in \{0,1\}} \{p_j u_j^t + \tilde{V}_{t+1}(\mathbf{w}_t + u_j^t \mathbf{e}_j)\}, \quad t = 1, \ldots, T \quad (1)$$

$$\tilde{V}_{T+1}(\mathbf{w}_{T+1}) = \phi(g(\mathbf{w}_{T+1})) \quad (2)$$

$\phi(\cdot)$: solution value given by problem-specific heuristic, MIP solved to given time limit / optimality gap or predicted by a ML algorithm

SUPERVISED LEARNING

- Separate the problem of accurately predicting Γ from that of solving (1)–(2)
- Off-line data generation
    - Aim: representation of feasible final states (optimal and suboptimal)
    - Sample trajectories using stationary random policies (acceptance probability p) and compute $\Gamma(\mathbf{w}_{T+1})$
    - Labeled data $\mathscr{D} = \{(\mathbf{w}^1_{T+1}, \Gamma(\mathbf{w}^1_{T+1})), \ldots, (\mathbf{w}^N_{T+1}, , \Gamma(\mathbf{w}^N_{T+1}))\}$

DISTRIBUTION LOGISTICS

- Problem setting from Giallombardo et al. (2020)
- Booking requests: Pick-up activities
- CO problem: Vehicle Routing Problem (VRP)
  - Fixed number of vehicles K_0
  - Optimal value $z^*(\mathbf{w}, K)$ where K is the number of vehicles
  - If more than K_0 vehicles are required, then we allow for additional outsourcing vehicles at fixed cost $C \in \mathbb{R}_+$
  - Operational cost: $\Gamma(\mathbf{w}) = -\max_{K \geq K_0, K \in \mathbb{Z}} \{z^*(\mathbf{w}, K) + C(K - K_0)\}$

INSTANCES AND FEATURES

- Sets of instances
  - 4, 10, 15 and 50 locations (uniform random location, locations are partitioned into groups with different revenues)
  - Request probabilities fixed such that higher revenue requests have greater probability of occurring later in the booking period
- Features $g(\mathbf{w}_{T+1})$
  - Fixed-size: max, mean, median, standard deviation, 1st/3rd quartiles derived from capacity, depot location, total number of accepted requests per location, aggregate statistics of the locations (relative distances)

APPROXIMATION

- We seek an accurate approximation of
  $\Gamma(\mathbf{w}) = -\max_{K \geq K_0, K \in \mathbb{Z}} \{z^*(\mathbf{w}, K) + C(K - K_0)\}$
  - Predict $z^*(\mathbf{w}, K)$
  - Compute outsourcing cost $C(K - K_0)$ with a bin-packing solver, MTP (Martello, 1990)
- Supervised learning
  - Data
    - For each set of instances, generate sample trajectories using stationary random policy with different values of p
    - Labels $z^*(\mathbf{w}, K)$ computed using heuristic solver FILO (Accorsi and Vigo, 2020)
    - $|\mathscr{D}^{\text{Train}}| = 2{,}000$ and $|\mathscr{D}^{\text{Test}}| = 500$
  - Model: random forest

ALGORITHMS

- Off-the-self RL and ADP using predicted costs $\tilde{V}_{T+1}(\mathbf{w}_{T+1})$ except to evaluate final cost, then we use one call to VRP solver (FILO)
  - On-policy RL: **SARSA** with neural state approximation
  - ADP: Rollout with Monte Carlo Tree Search (MCTS) with Upper Confidence Bounds applied to Trees (UCT) predefined number of simulations at each stage X (30 or 100)
    - Base policy: Stationary random policy – **MCTS-rand-X**
    - Base policy: SARSA – **MCTS-SARSA-X**
- Baselines
  - **Booking-limit policy (BLP)** and **booking-limit policy with reoptimization (BLPR)** proposed by Giallombardo et al. (2020), solved using SCIP for instances with 4 and 10 locations
  - For the smallest instances: exact dynamic programming with VRP solver – **DP-Exact** – or with predicted costs – **DP-ML**
  - Best stationary random policy – ***rand-p***

RESULTS

- Supervised learning: low mean squared and mean absolute errors even for the larger instances
- Control policies evaluated on the same 50 realizations of requests for each sets of instances
- Intel Core i7-10700 2.90 GHz with 32 GB RAM

SOLUTION QUALITY 4 LOCATIONS

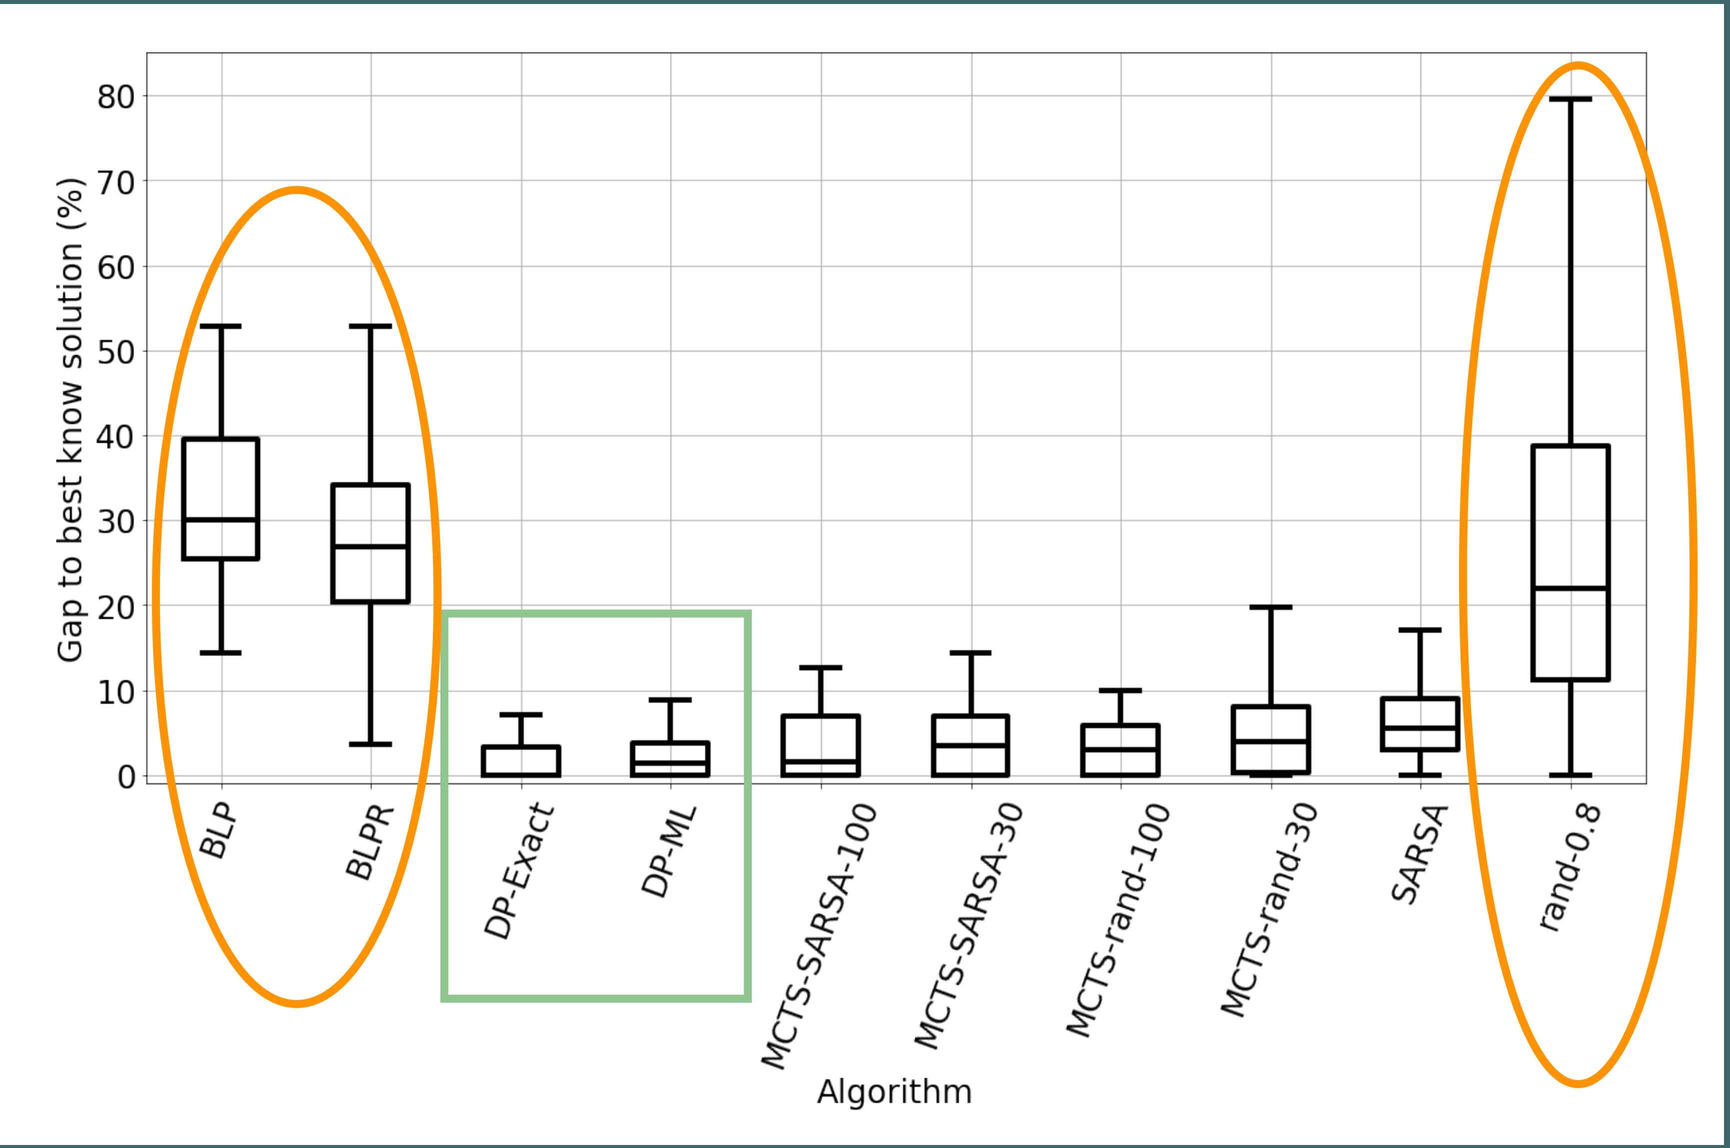

Cost approximations $\tilde{V}_{T+1}(\mathbf{w}_{T+1})$ lead to high-quality solutions as DP-ML close to DP-Exact

ADP and RL algorithms outperform baselines

SOLUTION QUALITY 10 LOCATIONS

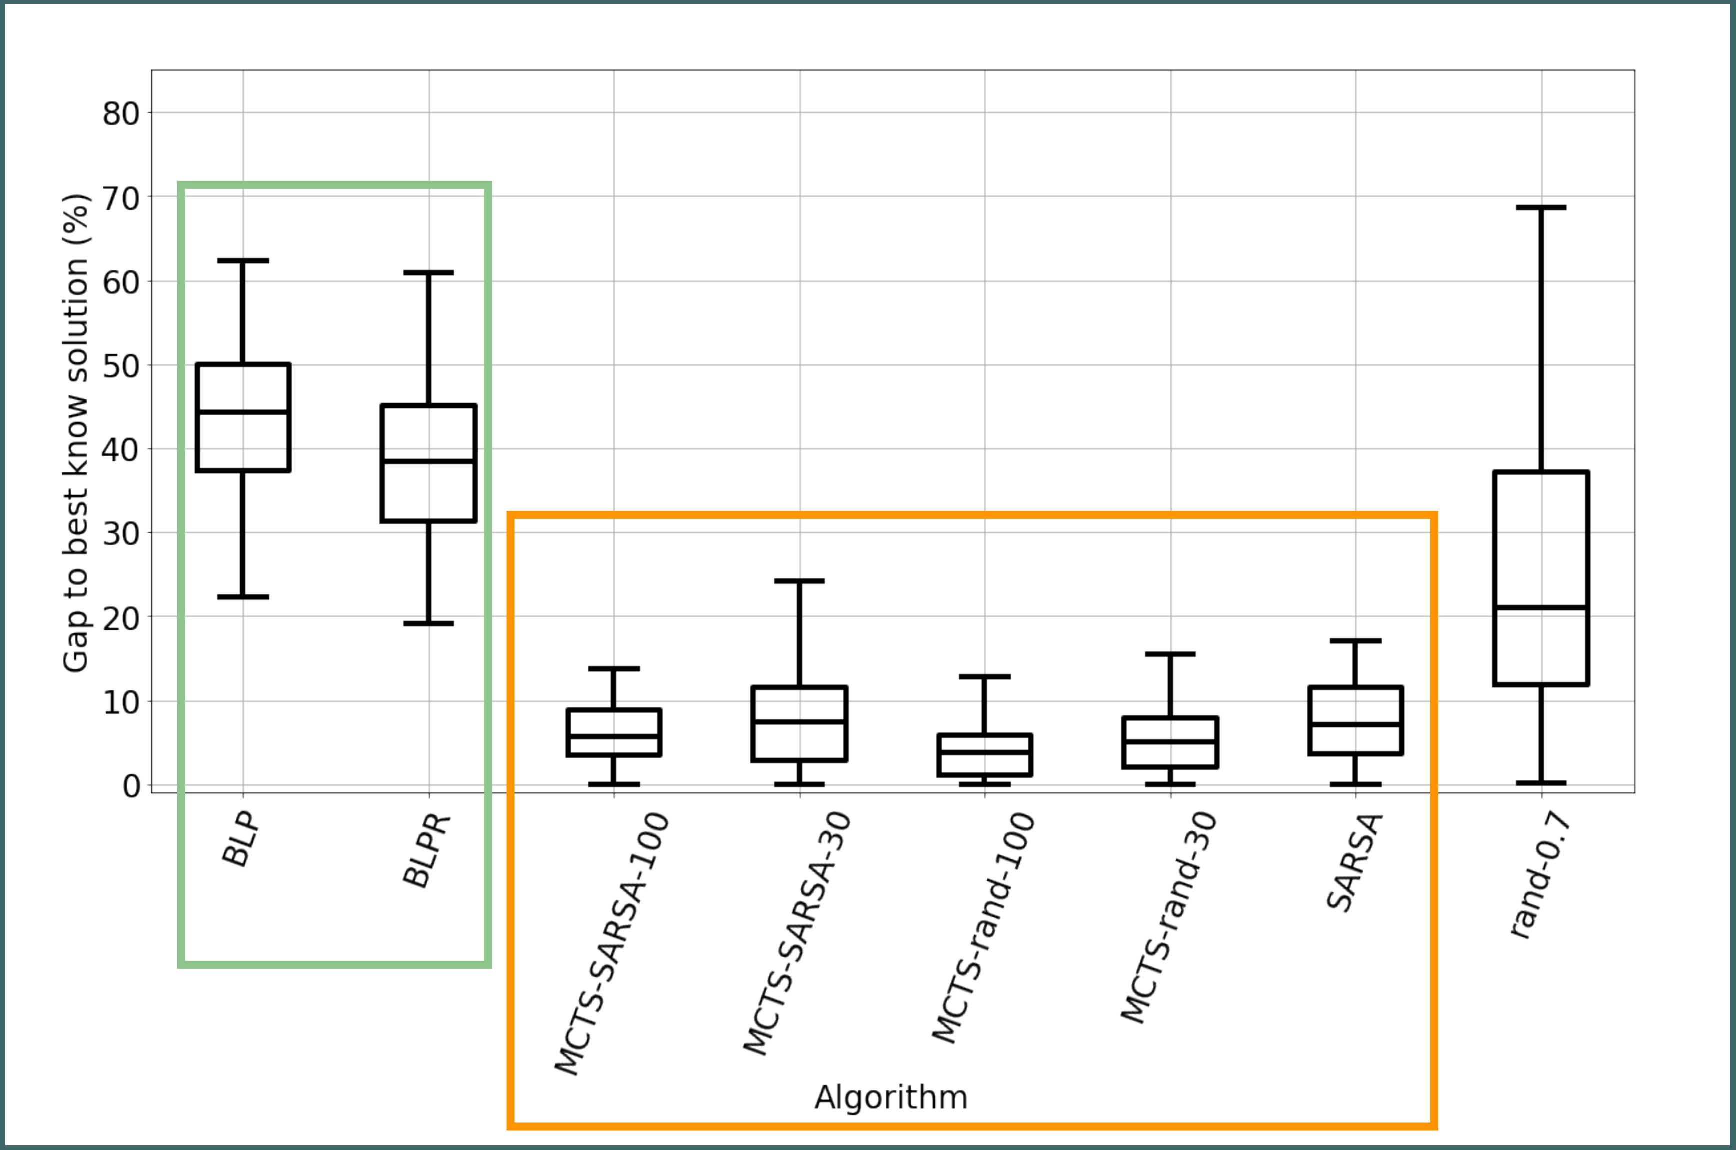

ADP and RL algorithms outperform baselines

BLP and BLRP produce poorer quality solutions than rand-0.7

SOLUTION QUALITY 15 AND 50 LOCATIONS

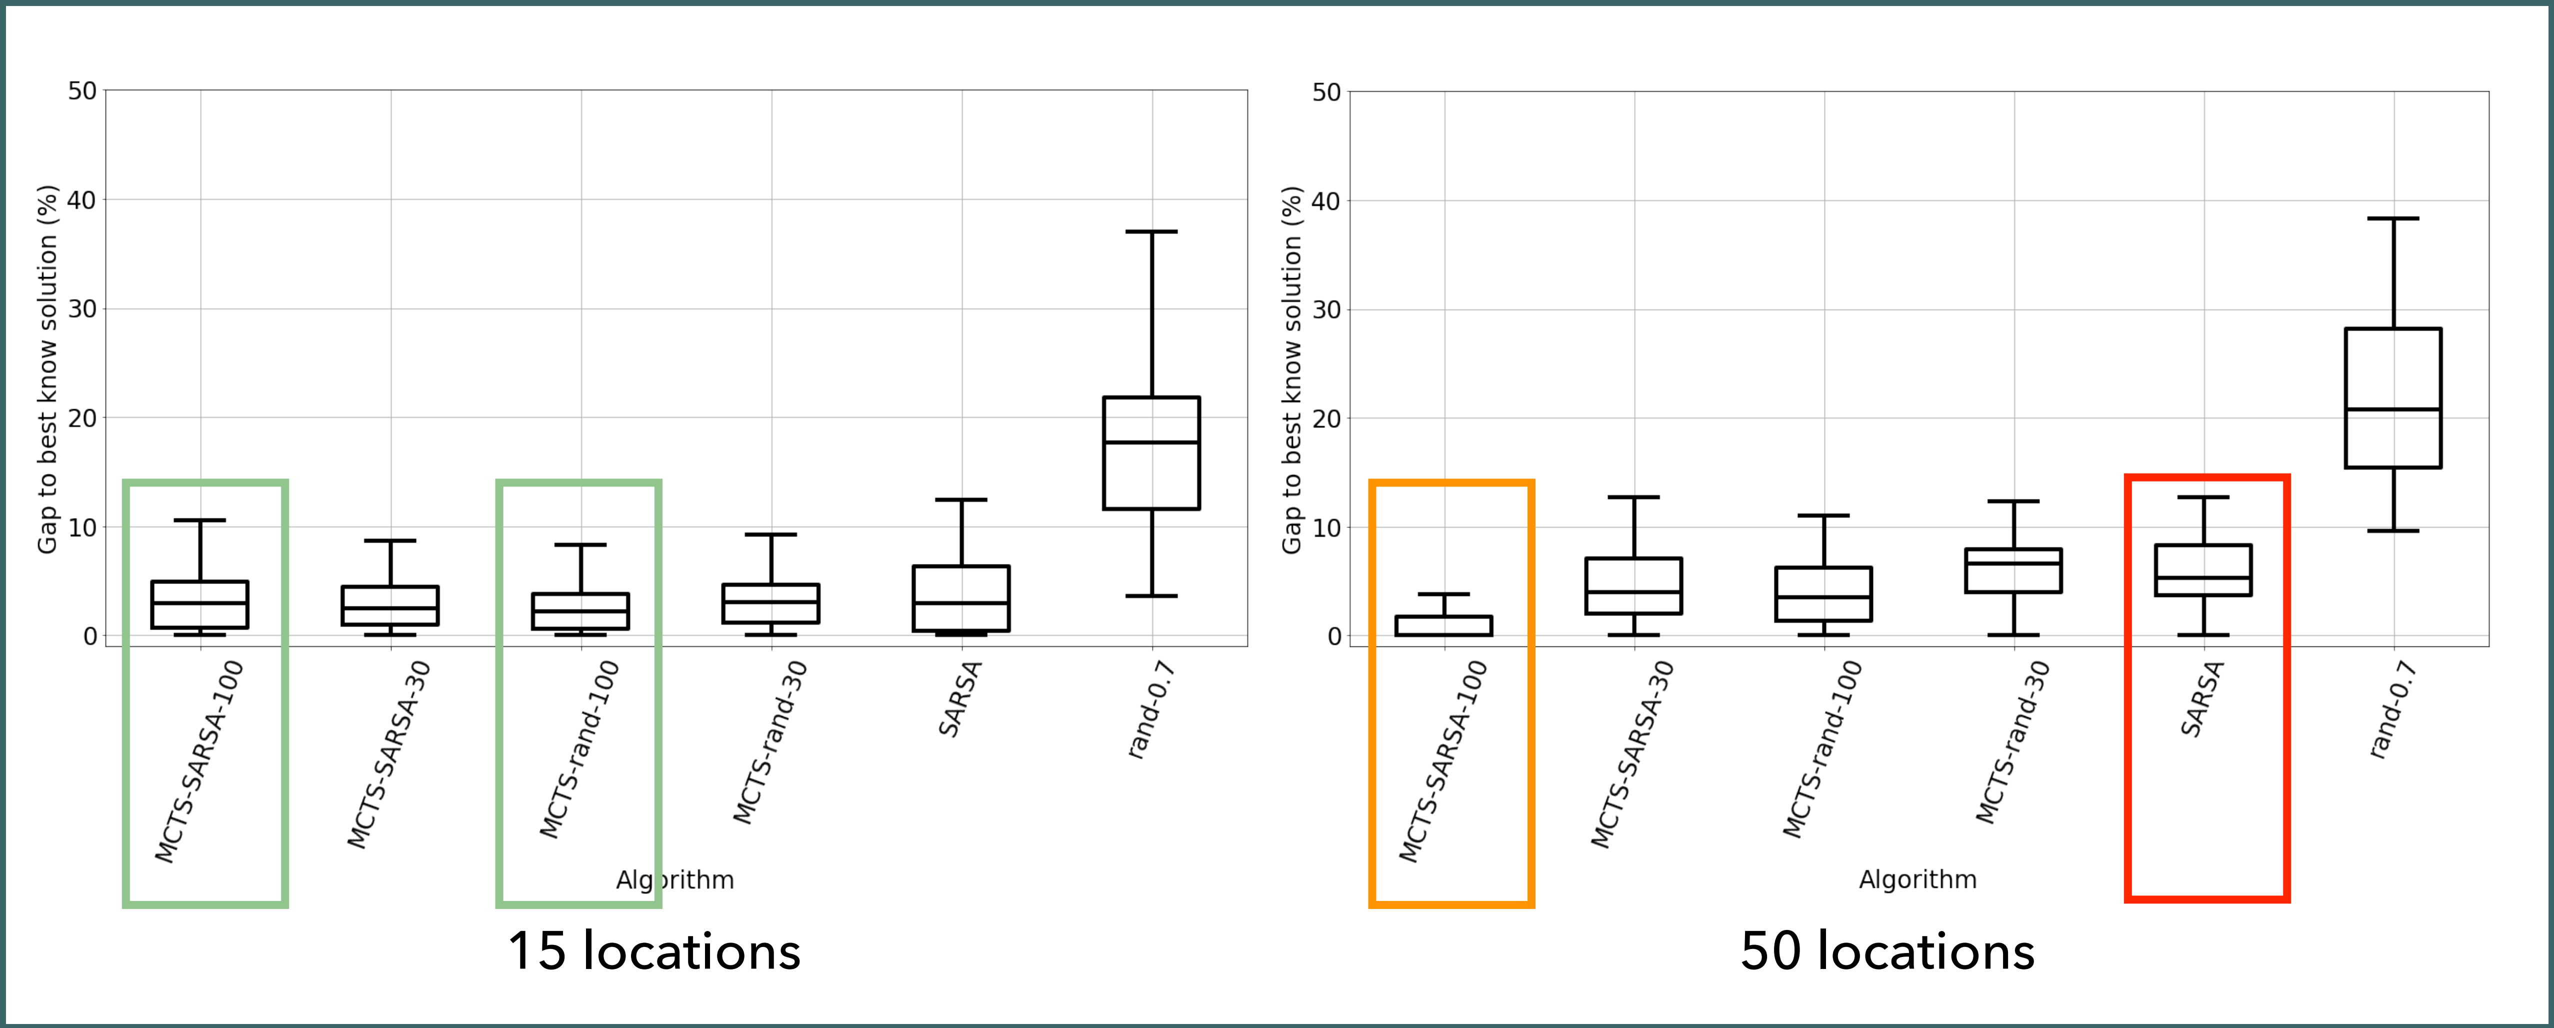

MCTS algorithms with 100 simulations (instead of 30) produce the best quality solutions

MCTS-SARSA-100 performs better on the largest instances, SARSA relatively good quality

COMPUTING TIME

Offline	Online
Cost approximation Data generation (<1h) and training ML algo (<1min) Test: ML prediction and bin-packing heuristic	VRP to evaluate final cost Cost approximation
SARSA – training using ML and heuristic < 2 h	< 2 s (50 locations)
MCTS-SARSA – same as SARSA	avg. 2.1/7.5 min (50 locations, 30/100 sim.)
	MCTS-rand-X online algo avg. 1.2/4.9 min (50 locations, 30/100 sim.)
	rand-X < 2 s (50 locations)
BLP **BLPR** Initial policy < 3 s	< 11 s < 13 s (10 locations)

Conclusion

- Idea: use machine learning to predict solution cost to the CO problem that constitutes the main bottleneck in existing solution approaches
- Exploratory work using off-the-shelf ML/RL/ADP shows promising results
- Ongoing work: second application on airline cargo capacity management
- Future work: beyond off-the-shelf RL and ADP

REFERENCES

- Accorsi & Vigo. A fast and scalable heuristic for the solution of large-scale capacitated vehicle routing problems, Tech. rep., University of Bologna, 2020.
- Adelman, Dynamic Bid Prices in Revenue Management. Operations Research 55(4):647-661, 2007.
- Barz & Gartner. Air cargo network revenue management. Transportation Science 50(4):1206-1222, 2016.
- Bengio, Lodi & Prouvost. Machine learning for combinatorial optimization: A methodological tour d'horizon, EJOR 290(2):405-421, 2021.
- Fischetti & Fraccaro. Using OR + AI to predict the optimal production of offshore wind parks: A preliminary study. In Optimization and Decision Science: Methodologies and Applications, volume 217, pages 203–211. Springer, 2017.
- Giallombardo, Guerriero & Miglionico. Profit maximization via capacity control for distribution logistics problems, arXiv:2008.03216, 2020.
- Larsen, Lachapelle, Bengio, Frejinger, Lacoste-Julien, Lodi. Predicting tactical solutions to operational planning problems under imperfect information, accepted for publication in IJOC, arXiv:1807.11876.v3, 2020.
- Levin, Nediak & Topaloglu. Cargo Capacity Mangement with Allotments and Spot Market Demand, Operations Research 60(2):351-365, 2012.
- Levina, Levin, McGill & Nediak. Network cargo capacity management, Operations Research 59(4), 1008-1023, 2011.
- Martello & Toth, Knapsack problems: algorithms and computer implementations, Wiley series in discrete mathematics and optimization, 1990.
- Puterman. Markov Decision Processes: Discrete Stochastic Dynamic Programming, Wiley series in probability and statistics, 1994.
- Talluri & Van Ryzin. The theory and practice of revenue management, Springer, 2005.
- Talluri & van Ryzin, An Analysis of Bid-Price Controls for Network Revenue Management. Management Science 44(11-part-1):1577-1593, 1998.
- Williamson. Airline network seat control. Ph.D Thesis, MIT, Cambridge, MA, 1992.

Thank you!

emma.frejinger@umontreal.ca

emmafrejinger.org