Predictions for Better Decisions: Towards Integrated Prediction and Optimization

EMMA FREJINGER

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

CIRRELT, Mila

Applied Machine Learning Days, Lausanne, Switzerland

AI and Mobility Track

March 29, 2022

Based on joint work with Greta Laage, Robin Legault, Andrea Lodi, Mike Hewitt, Guillaume Rabusseau and Gilles Savard

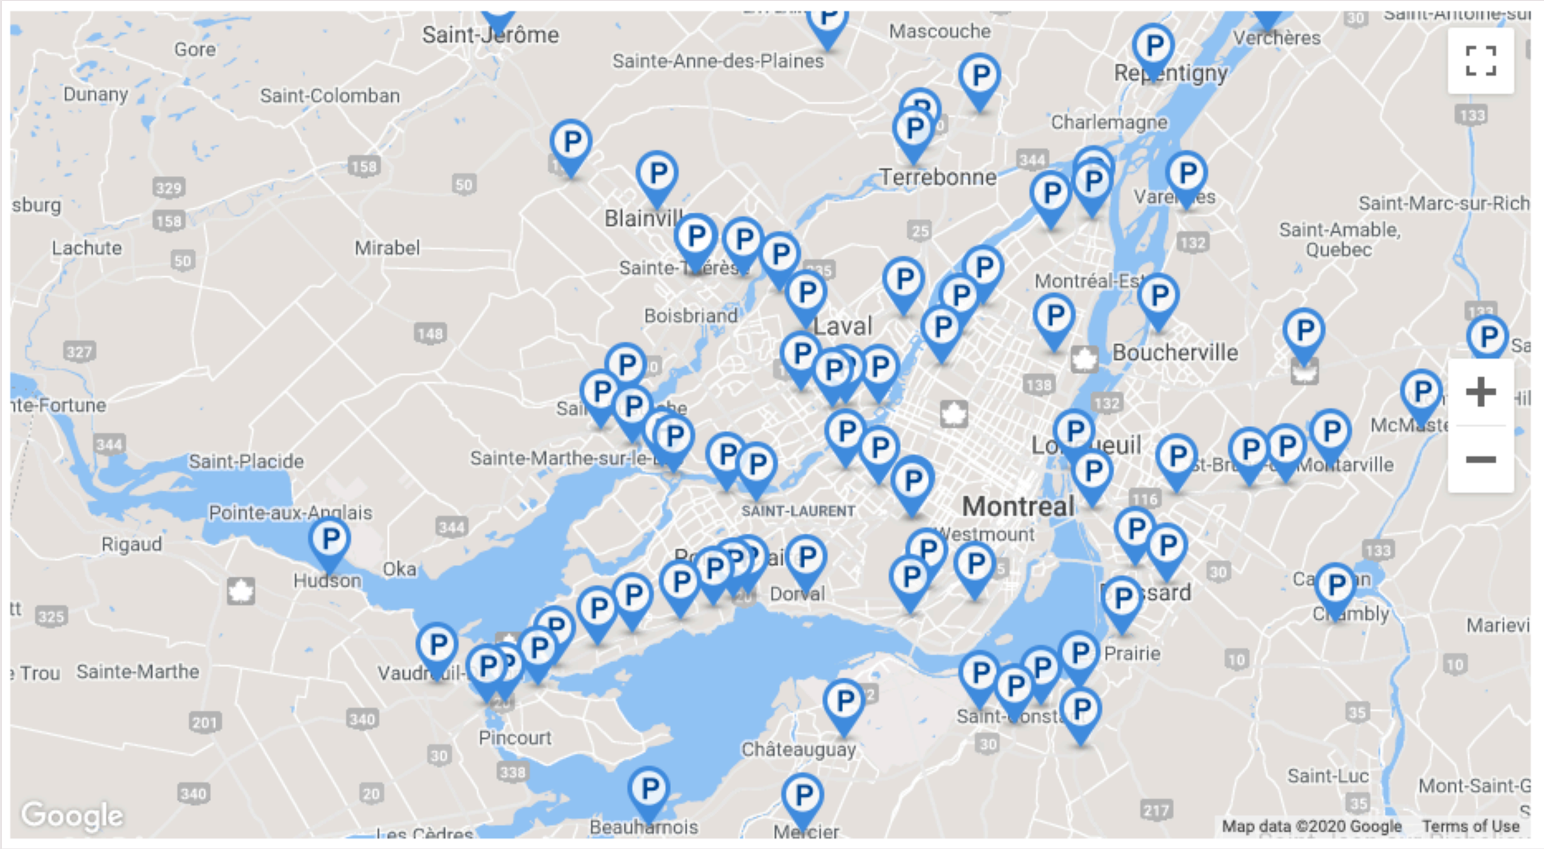

Demand for park-and-ride facilities

Domestic, import and export container traffic

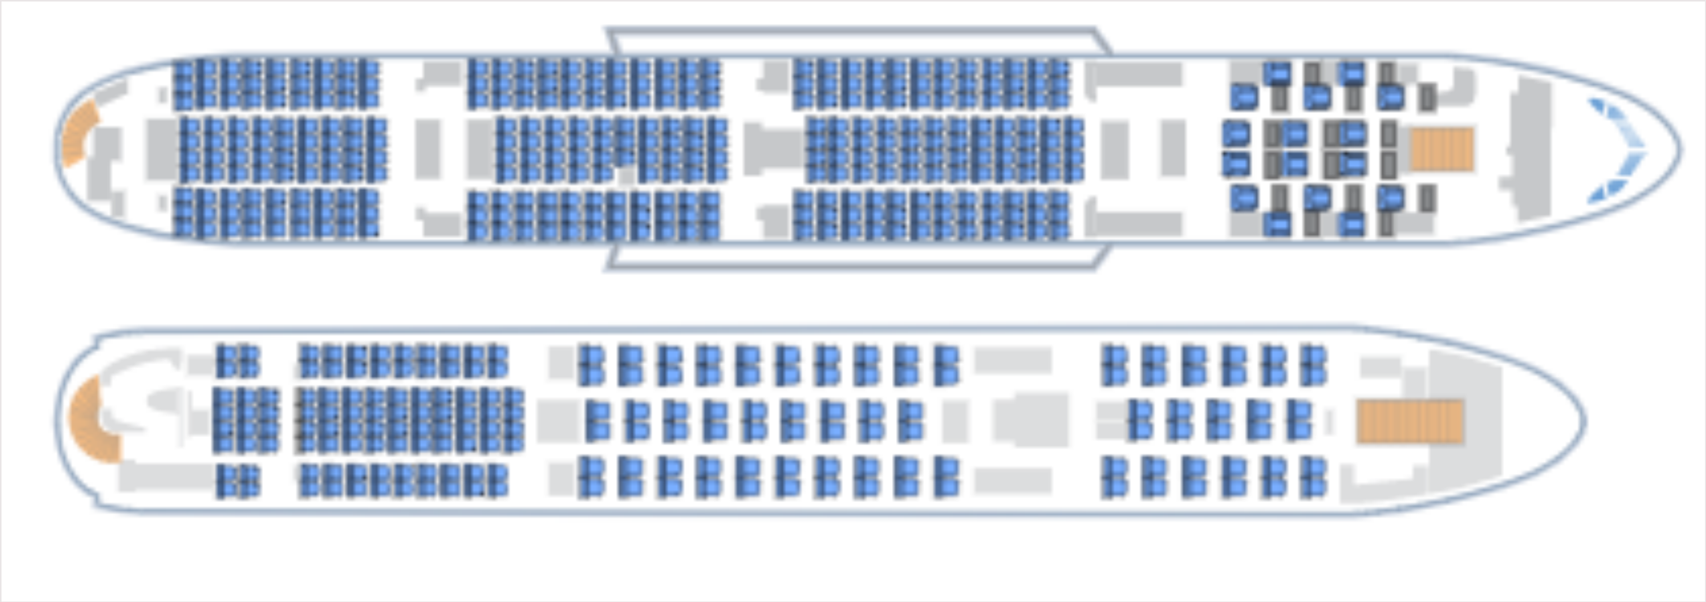

Passenger demand for air travel

Demand forecasting for transportation and mobility:

When, where, how much?

Human behaviour drives demand

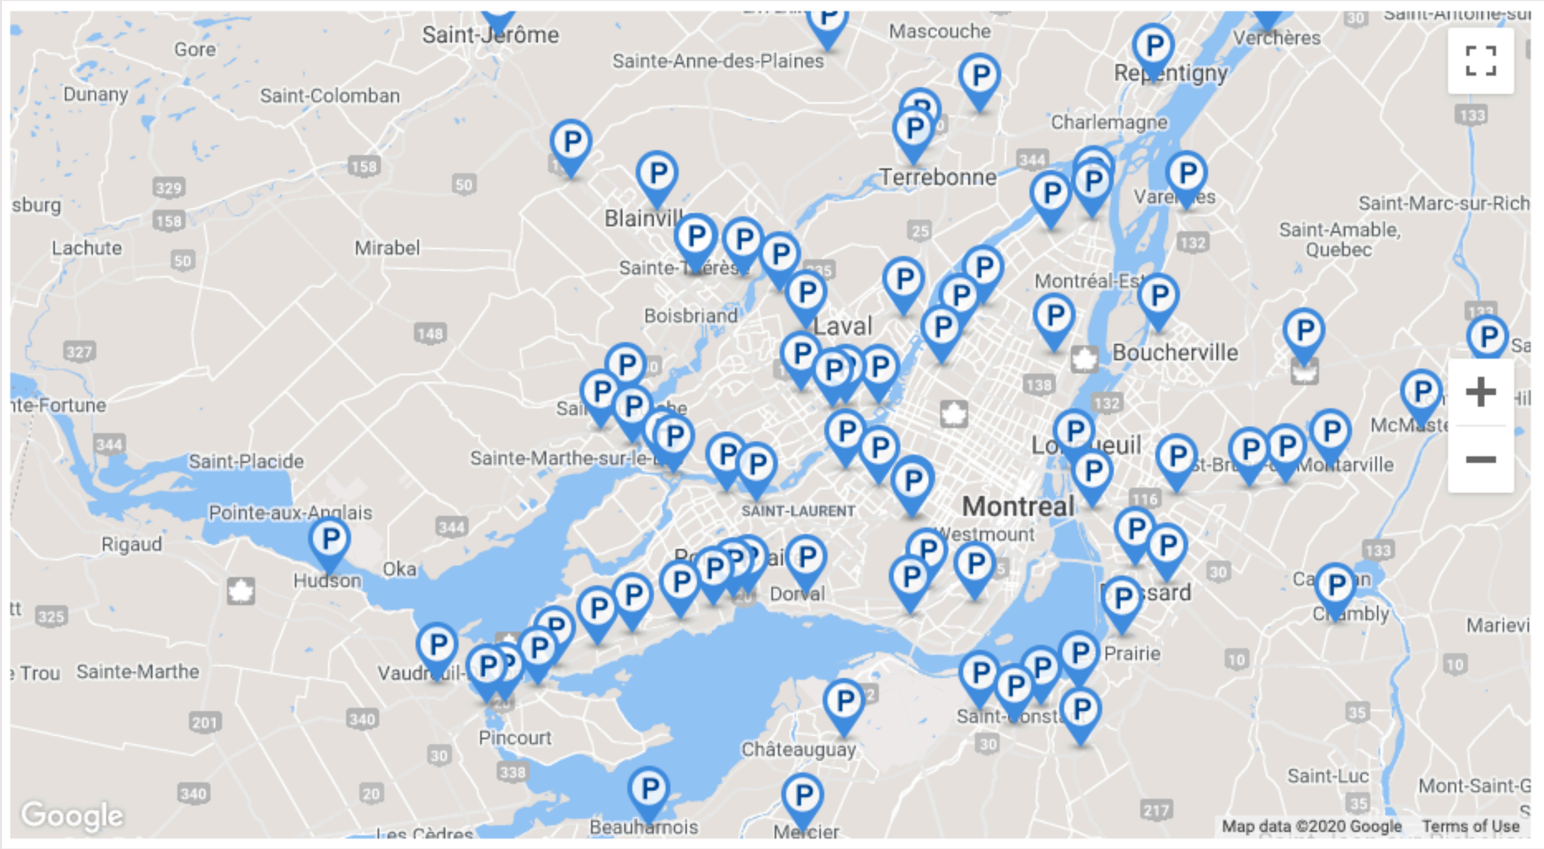

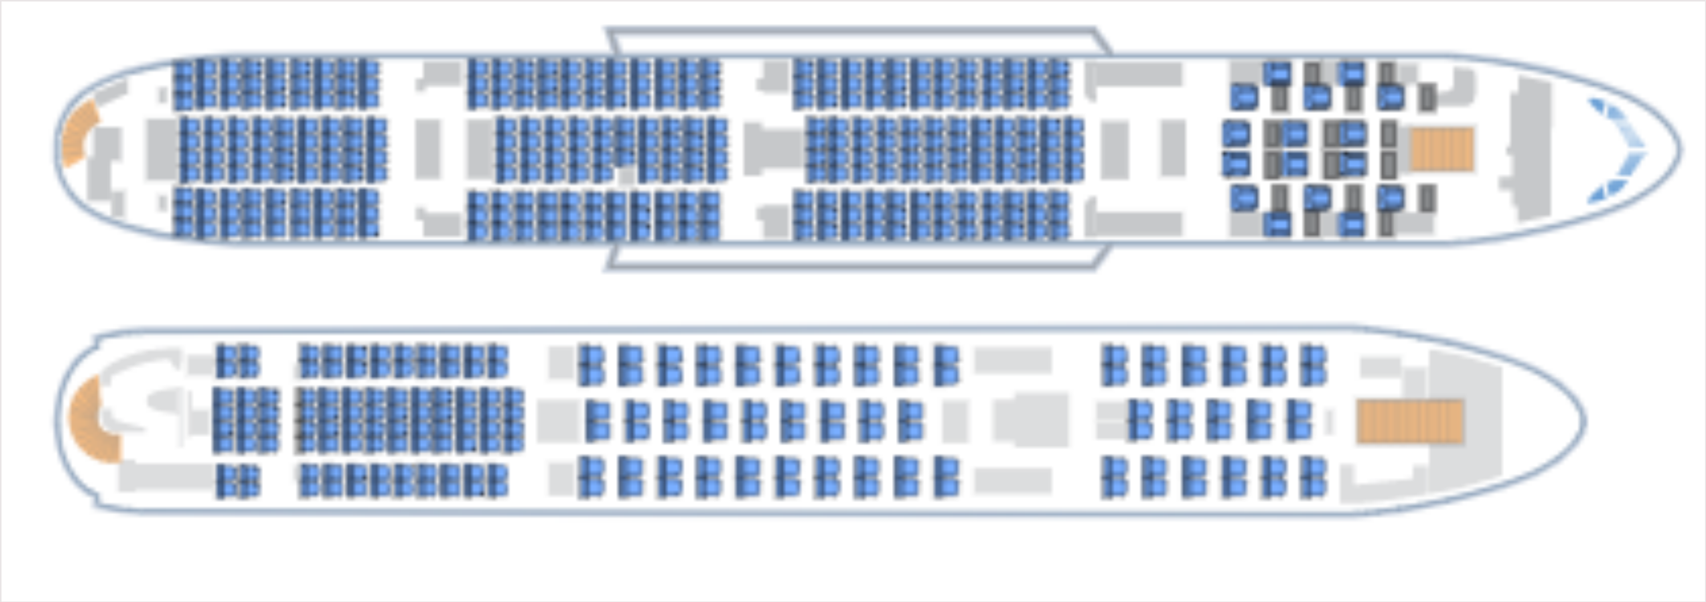

Demand for park-and-ride facilities

Domestic, import and export container traffic

Passenger demand for air travel

Decide location and capacity of facilities to maximize captured demand

Plan transport services and their capacity to satisfy delivery requirements at minimum cost

Price tickets to maximize revenue

Demand predictions are rarely useful on their own –

Used to make **decisions**

Supply optimization and demand management

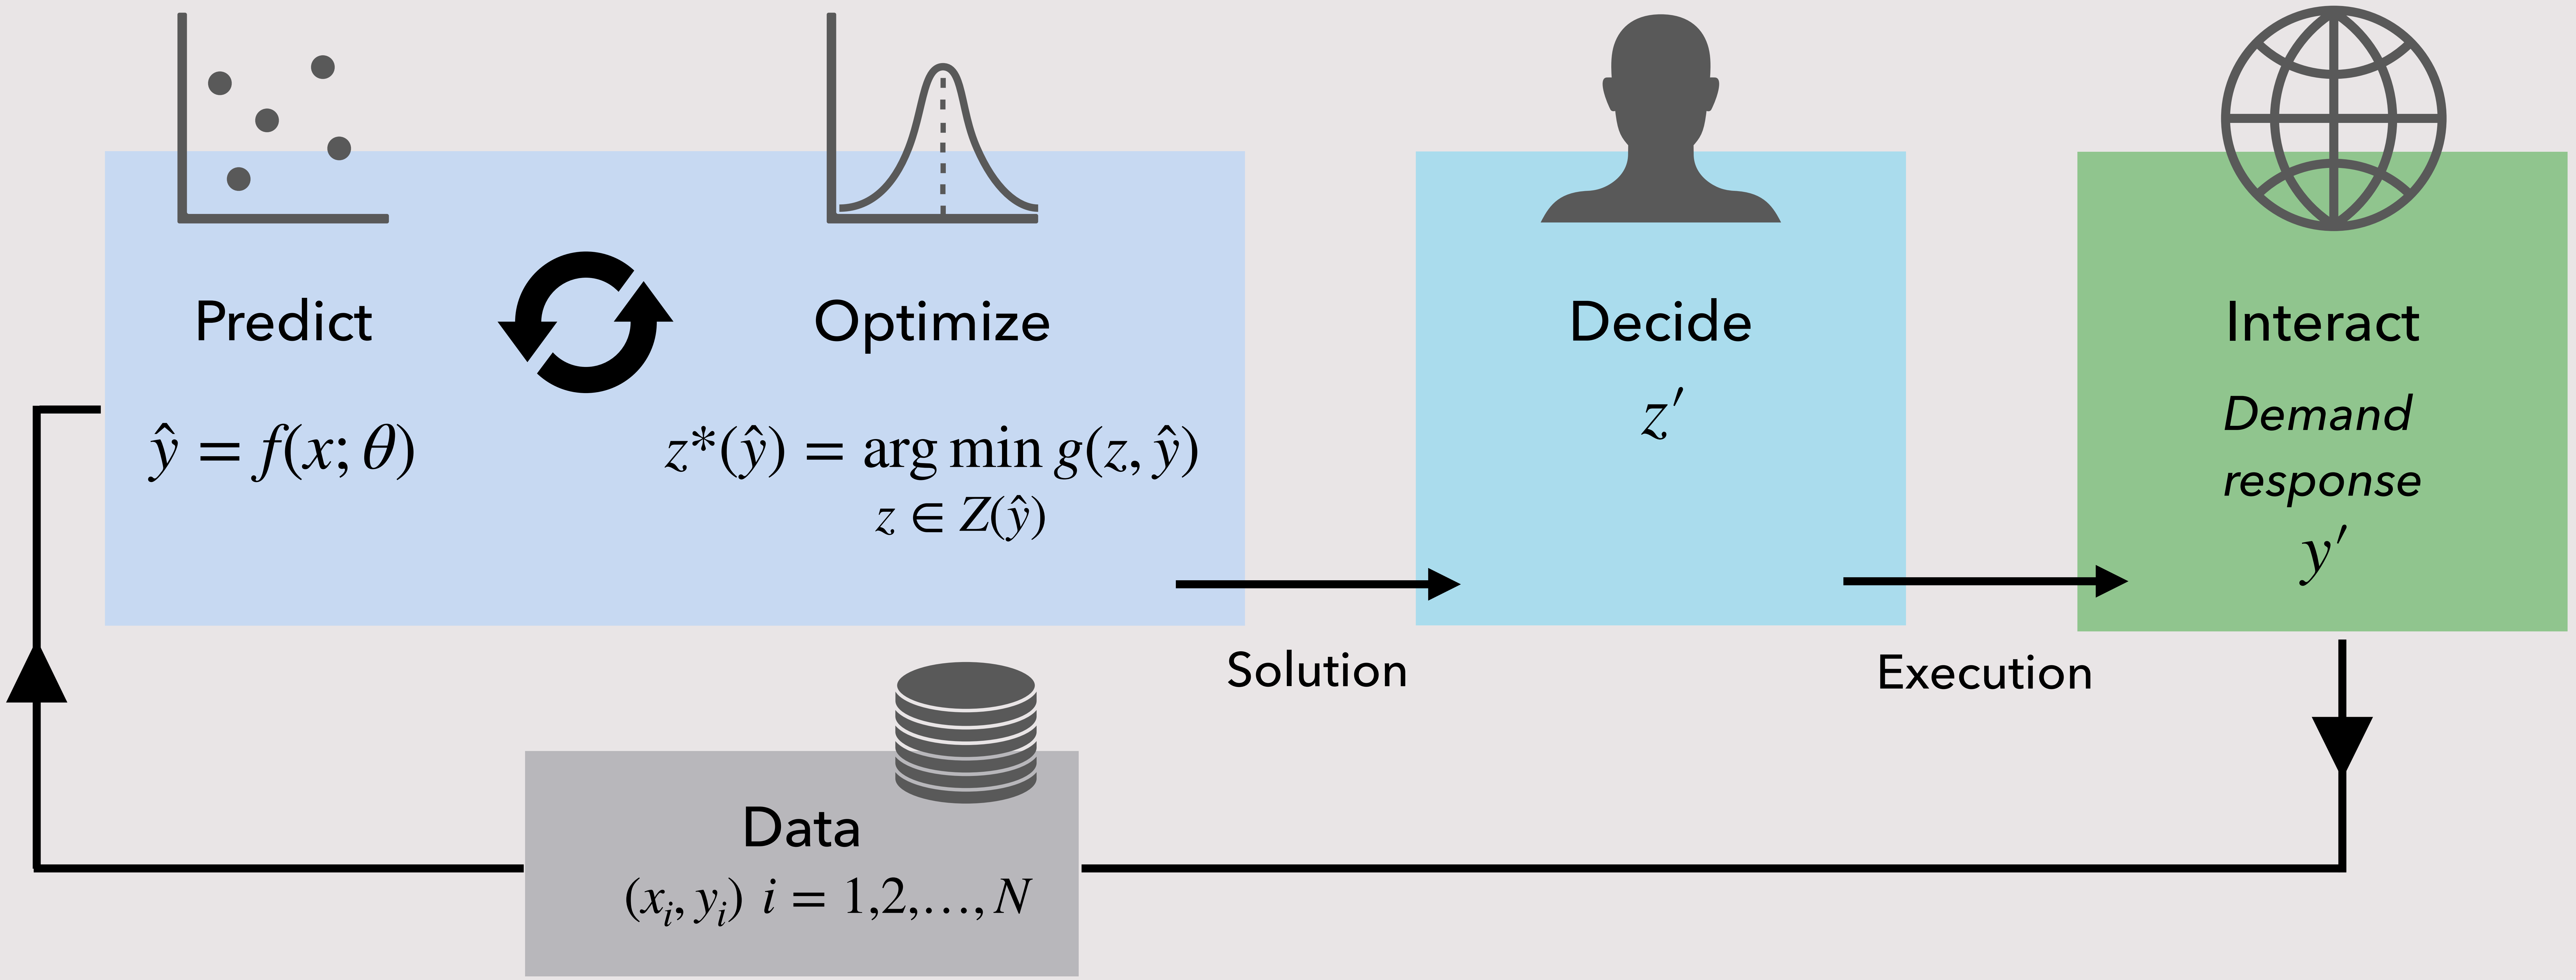

In general, decision-making problems in mobility and transportation are recurrent, involve a human in the loop, occur in complex uncertain environments

MACHINE LEARNING FOR DECISION-MAKING IN TRANSPORTATION AND MOBILITY

– Supply optimization and demand management

Models

Integrate prediction and optimization models

Solving methods

Speed-up solving methods of deterministic or stochastic discrete (or combinatorial) optimization problems through learning

- ML augmented CO
- Predicting CO solutions

Surveys: Bengio et al., 2021, Kotary et al., 2021

OUTLINE

1. Integrate prediction and optimization to improve decisions and anticipate demand response
2. Decision awareness in learning can be of high value
3. Measuring actual impact may not be as easy as it sounds

1

Predict Optimize

Most optimization models for transport and mobility assume that demand is fixed and known.

Reality:

- Decisions impact demand. Need for anticipation.
- User preferences are heterogeneous and we have imperfect knowledge thereof.

OPTIMIZATION WITH ENDOGENOUS DEMAND (UNCERTAINTY)

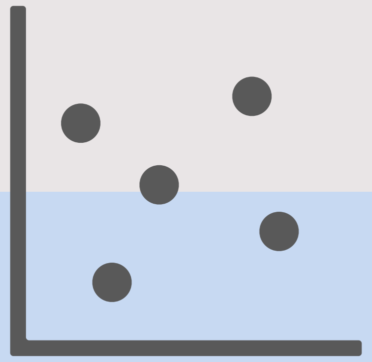

Predict

$$\hat{y} = f(x; \theta)$$

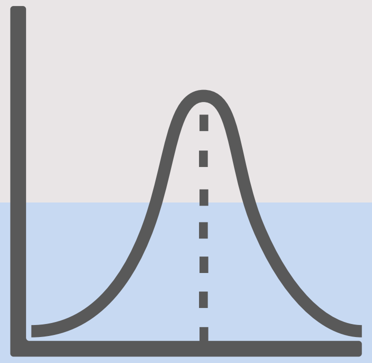

Optimize

$$z^*(\hat{y}) = \arg\min_{z \in Z(\hat{y})} g(z, \hat{y})$$

$$x = (x', z)$$

Features used for prediction include exogenous variables x' and decision variables z

- Stochastic programming: decisions impact the probability distributions of uncertain model parameters (e.g., Bhuiyan et al. 2020)
- Robust optimization: decision-dependent uncertainty sets
- Optimization with random utility maximizing users/customers

OPTIMIZATION WITH RANDOM UTILITY MAXIMIZING (RUM) USERS

Optimize Decisions

$$z^*(\hat{y}) = \arg\min_{z \in Z(\hat{y})} g(z, \hat{y})$$

Predict

Users optimize

Maximize random utility

$$\hat{y} \in \arg\max \mathbb{E}_{\varepsilon}[u(y, x', z, \varepsilon)]$$

Bilevel programming formulation

Leader makes a decision z anticipating followers' reactions

Followers react to z choosing an option that maximizes their utility u (modelled as a random variable)

The two objectives are conflicting

- Bilevel programming: important in different domains, e.g., pricing problems in transportation
- NP-hard even when leader and follower problems are linear programs
- Most work assume deterministic follower model. Few exceptions, e.g.,
  - Network pricing (Gilbert et al., 2014, 2015), competitive facility location (Dan and Marcotte, 2019)

Source: Wikipedia

COMPETITIVE FACILITY LOCATION

A simulation approach to deal with any type of random utility maximization (RUM) discrete choice model

Robin Legault and Emma Frejinger, A Simulation Approach for Competitive Facility Location with Random Utility Maximizing Customers, arXiv:2203.11329, 2022.

COMPETITIVE FACILITY LOCATION

- Locate facilities in a competitive market to maximize captured customer demand
- Generative perspective: simulate customers' utilities instead of using probabilities (Paneque et al., 2021)
- Sample average approximation: flexible, but requires a lot of scenarios
- Clustering heuristic: aggregate customers according to preference profile – reduces the number of scenarios without affecting the optimal solution

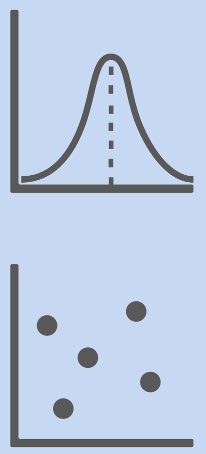

$$\max_{z \in Z} \mathbb{E}_\theta \left[\mathbb{P}_\varepsilon \left[\arg\max_{c \in C(z)} \{u_c(\theta, \varepsilon)\} \in D(z) \right] \right]$$

$\{\psi_n\}_{n \in N}$ $\qquad$ $\{\xi_{ns}\}_{n \in N, s \in S}$

Any model requires simulation to evaluate $\mathbb{E}_\theta$ if the support of θ is infinite

We compare the performance of our simulation approach to MOA (Mai and Lodi, 2020)

$$\max_{z \in Z} \frac{1}{|N|} \sum_{n \in N} \mathbb{P}_\varepsilon \left[\arg\max_{c \in C(z)} \{u_c(\psi_n, \varepsilon)\} \in D(z) \right] \tag{1}$$

$$\max_{z \in Z} \frac{1}{|N||S|} \sum_{n \in N} \sum_{s \in S} \mathbf{1} \left[\arg\max_{c \in C(z)} \{u_c(\psi_n, \xi_{ns})\} \in D(z) \right] \tag{2}$$

Trade-off: number of simulated customers $|N|$ and number of scenarios for each customer $|S|$ approximating their behaviour

Higher Entropy

More scenarios required to obtain high-quality solutions

Leads to harder instances

Simulation approach offers computational advantage over state of the art for MNL for most instances

Large-scale problems (New York City)

Can effectively solve for mixed MNL

Mixed MNL when θ has infinite support:

- For the *same number* of simulated customers: solving (1) provides better solutions than solving (2), but at a *large computational cost*

- If entropy is not too high, (2) can be solved with $|N| > 100{,}000$ and $|S| = 1$ in seconds

RESULTS

- Interpretability: Information-theoretic characterization of instances - entropy
- Outperforms state of the art when observable attributes are strong predictors of customers' behaviour (relatively low entropy)
- Large number of simulated customers is required to close the relative generalization gap – favours the simulation approach
- Large $|N|$ seems more important than large $|S|$

2

Predict Optimize

Standard practice: Predict, then optimize

Decision awareness can be of high value

PREDICT, THEN OPTIMIZE

- Training according to a prediction criterion
- Minimize a loss function $L(y, \hat{y})$
  - E.g., distance between predicted and observed ground truth values

Predict

$$\hat{y} = f(x; \theta)$$

Predictions

Optimize

$$z^*(\hat{y}) = \arg\min_{z \in Z(\hat{y})} g(z, \hat{y})$$

$$L(y, \hat{y}) = \left(y - f(x, \theta)\right)^2$$

Train error = $\frac{1}{N} \sum_{i=1}^{N} L(y_i, f(x_i, \hat{\theta}))$

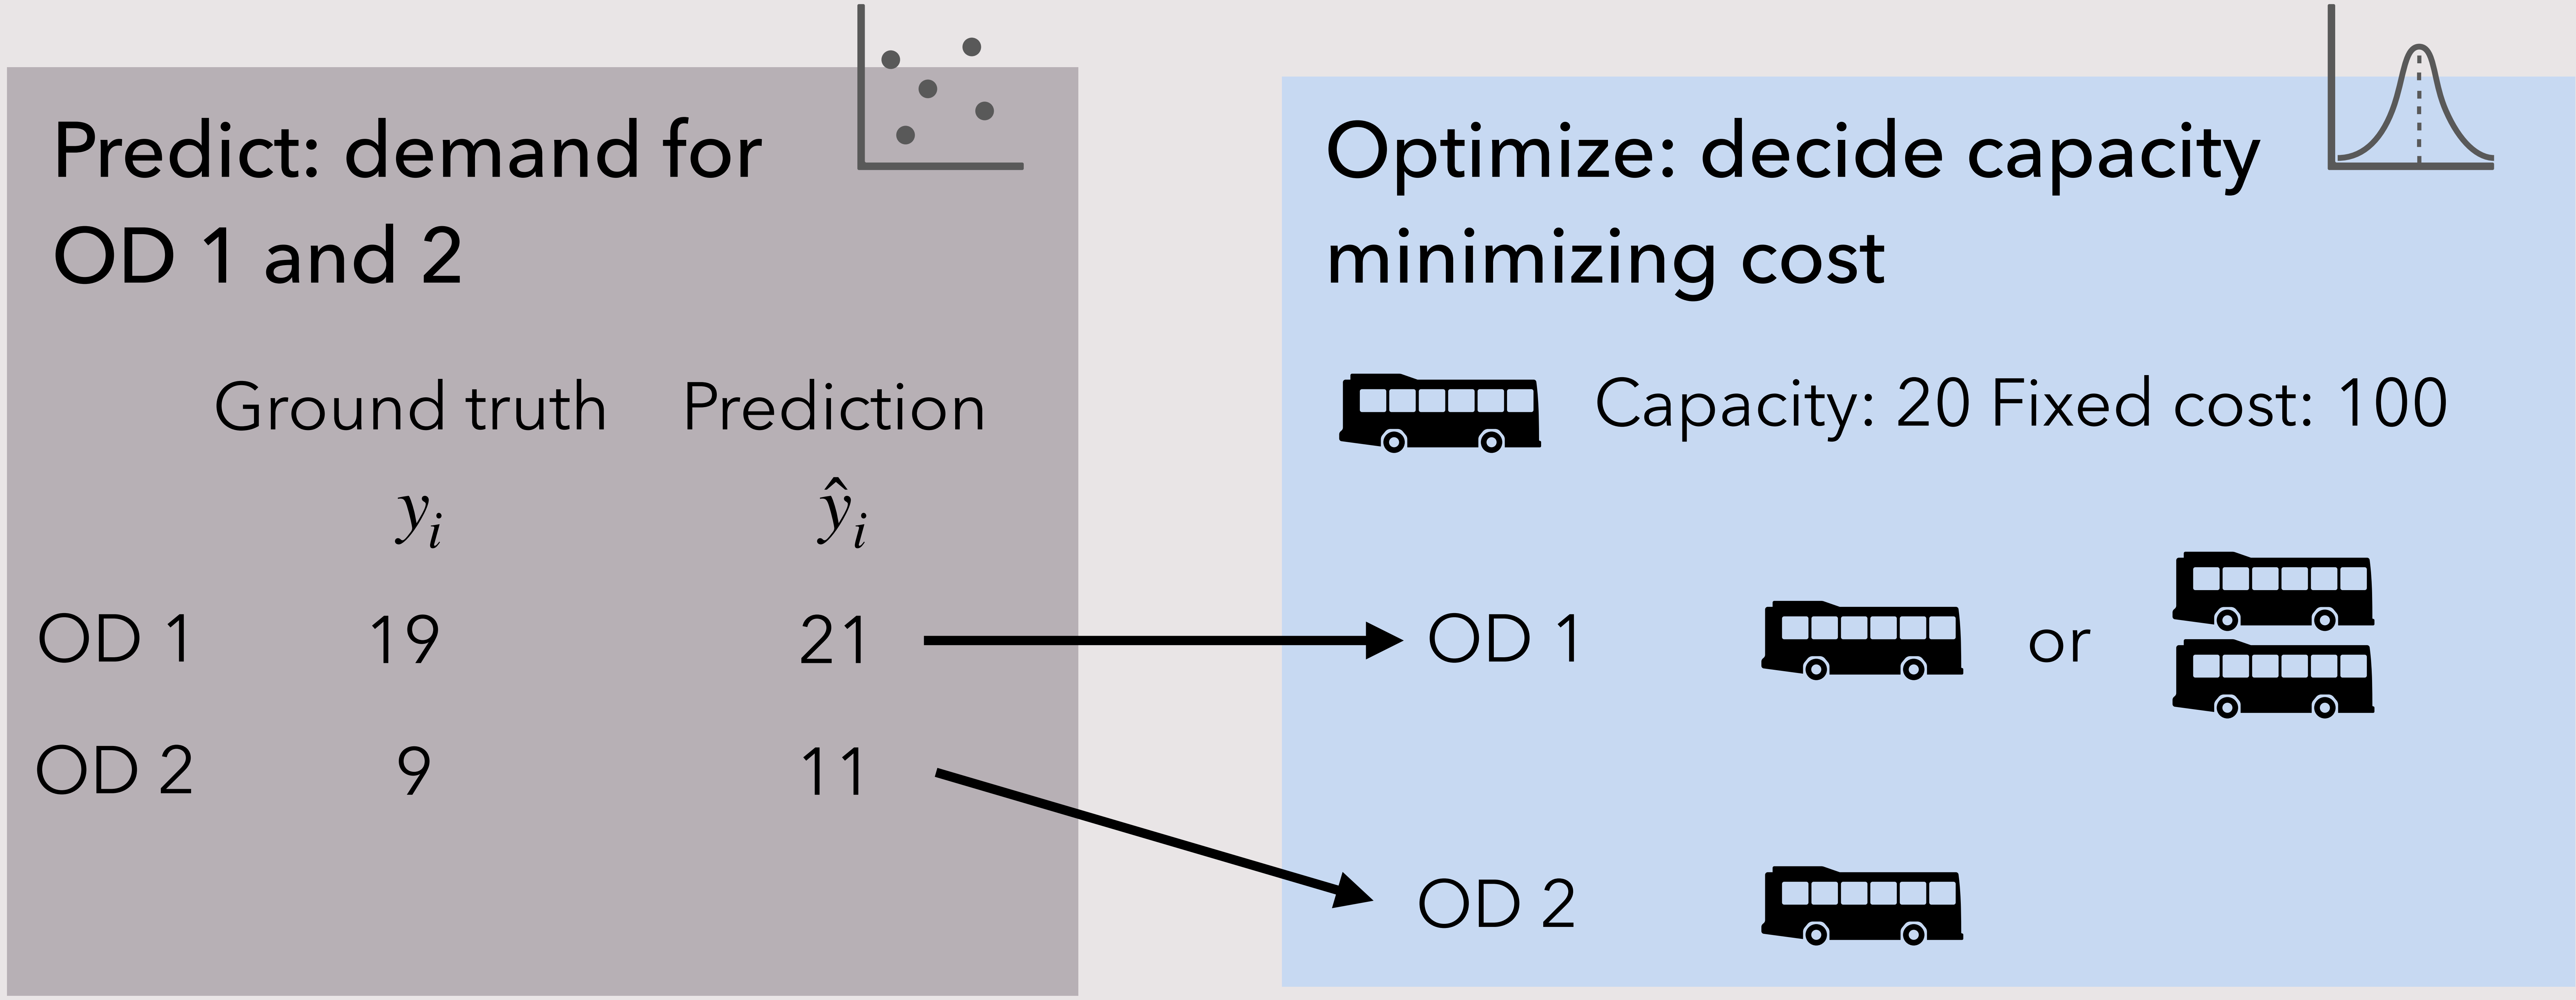

Illustrative example:

Equal prediction errors but different decision costs

END-TO-END LEARNING / DECISION AWARENESS

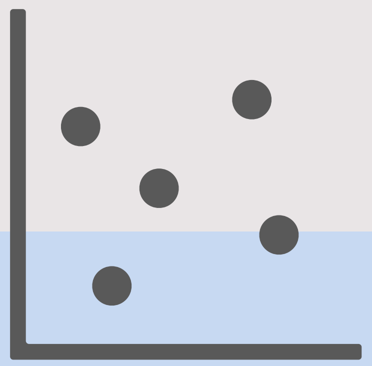

Predict

$$\hat{y} = f(x;\theta)$$

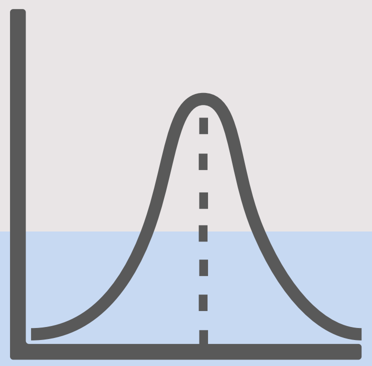

Optimize

$$z^*(\hat{y}) = \underset{z \in Z(\hat{y})}{\arg\min}\, g(z,\hat{y})$$

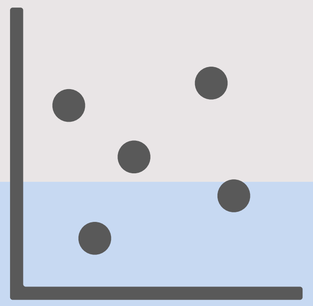

$$L_{\mathrm{r}}(y,\hat{y}) = g(z^*(\hat{y}),\hat{y}) - g(z^*(y),y)$$

Difference in objective function value from using $\hat{y}$ as opposed to y

- Training using regret minimization
- Challenge: differentiate through argmin operator
- Transport and mobility: $z^*(\hat{y})$ is typically a solution to a discrete optimization problem with predictions occurring in contraints and objective
  - E.g., Mixed Integer Linear Programs

END-TO-END LEARNING

Gap in the literature:

How to deal with large MILPs predictions in objective and constraints?

- Discrete (deterministic) optimization with unknown parameters in objective function only (Elmachtoub and Grigas, 2021, Ferber et al., 2020, Mandi et al., 2020, Pogančić et al., 2020)
  - Rely on linear programming results. E.g., Ferber et al. (2020) use cutting planes, Mandi et al. (2020) focus on MILPs having strong continuous relaxations
- Survey: linear programs and beyond (Kotary et al., 2021)

END-TO-END LEARNING RELATED TOPICS

- Model-based reinforcement learning: decision-aware model learning (e.g., Grimm et al., 2020) for sequential decision making problems formulated as Markov Decision Processes
- In case of observations both y and solutions z (optimal or suboptimal)
  - Data-driven inverse optimization with noisy data: very few results on discrete optimization with noisy data (Moghaddass and Terekhov, 2021)
  - Inverse reinforcement learning (Ng and Russell, 2000), dynamic discrete choice modeling (Aguirregabiria and Mira, 2010), system identification for control (Gevers, 2005)

WHAT TO DO IN PRACTICE?

Large-scale mixed integer linear program with predictions in constraints and objective function

Tactical planning and the periodic demand estimation problem for freight transportation

Collaboration with the Canadian National Railway Company (CN)

Laage, Frejinger and Savard, Periodic Freight Demand Estimation for Large-scale Tactical Planning, arXiv:2105.09136v2, 2021.

Laage, Frejinger and Savard, A Two-step Heuristic for the Periodic Demand Estimation Problem, arXiv:2108.08331, 2021.

PERIODIC DEMAND ESTIMATION

Tactical plan → **Operational plan**

Cyclic

Repeats in each period (week)

Adjusted tactical plan

Tactical plan	Operational plan
Origin: Montreal	Origin: Montreal
Mon, Wed, Fri, 9AM	Mon, Wed, Fri, 9AM
Toronto, Cap:300	Toronto, Cap:350
Vancouver, Cap:500	Vancouver, Cap:450
Tue, Thu, Sat, 5PM	Tue, Thu, Sat, 5PM
Quebec, Cap:100	Quebec, Cap:200
Halifax, Cap:400	Halifax, Cap:300
	Extra service, Cap:200

- A cyclic tactical plan (service network design) is in place over a given time horizon (e.g., a season)
  - Satisfy demand at minimum cost
- Input: *periodic* demand
  - Demand expected to repeat in each period (e.g., week)

ILLUSTRATIVE EXAMPLE

x-axis: time periods (week) over time horizon T

y-axis: demand y_{k_0} for commodity k_0

Each line: a mapping from per period forecasts to periodic demand

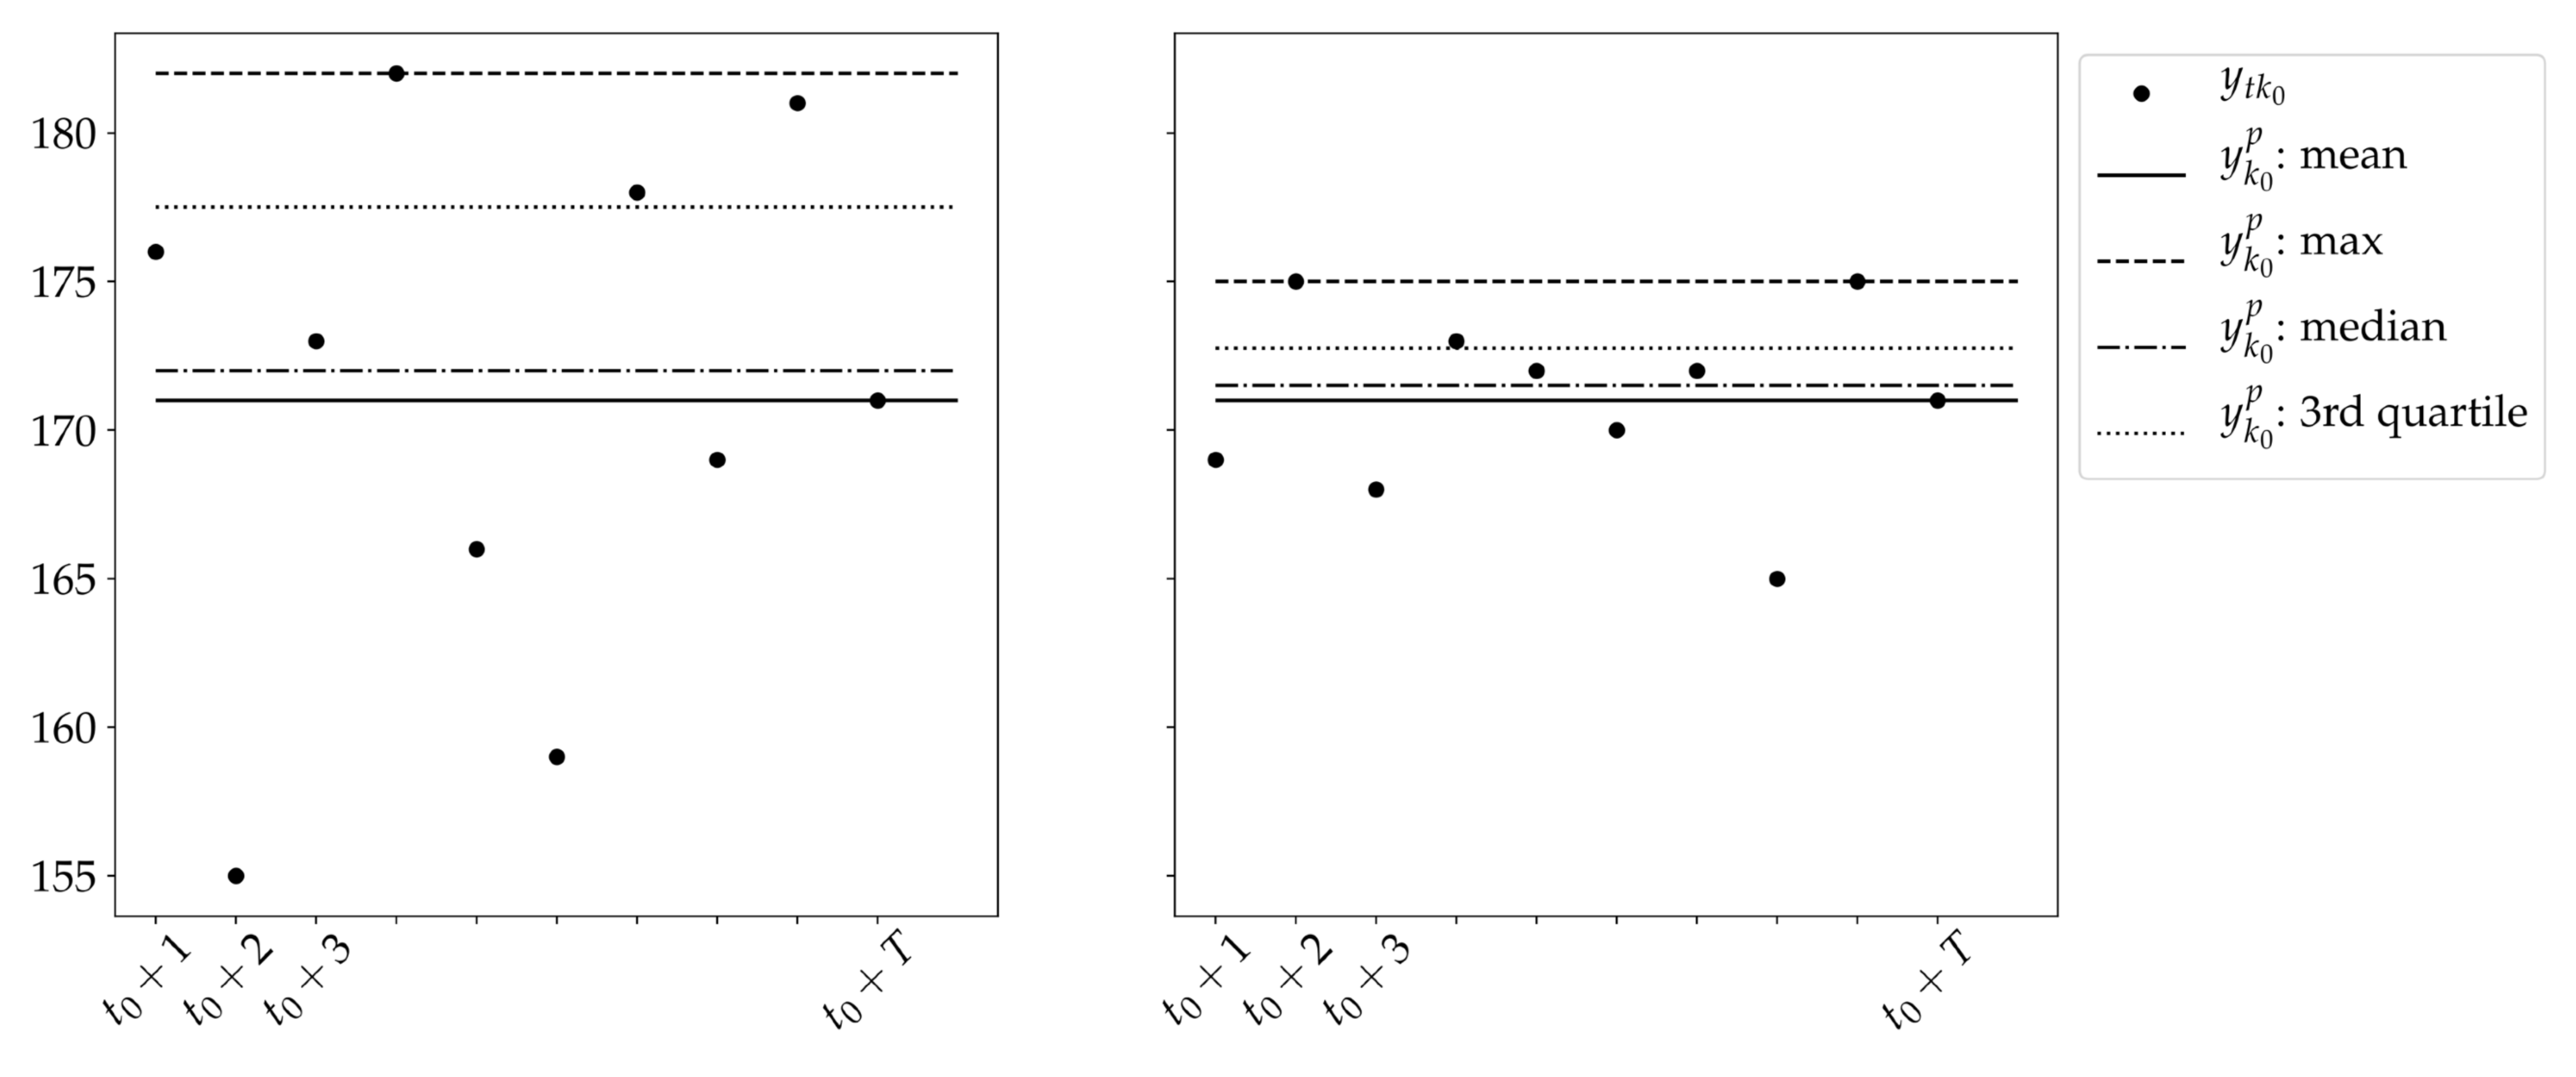

- How to map demand forecasts *per period* to *periodic* demand? I.e., what is a good periodic demand scenario?
  - The mean is typically used in practice
  - Use a distribution instead of a single value per commodity and period: discrete optimization problem under uncertainty (e.g., Crainic et al., 2020) - computationally costly to apply to real large-scale problems

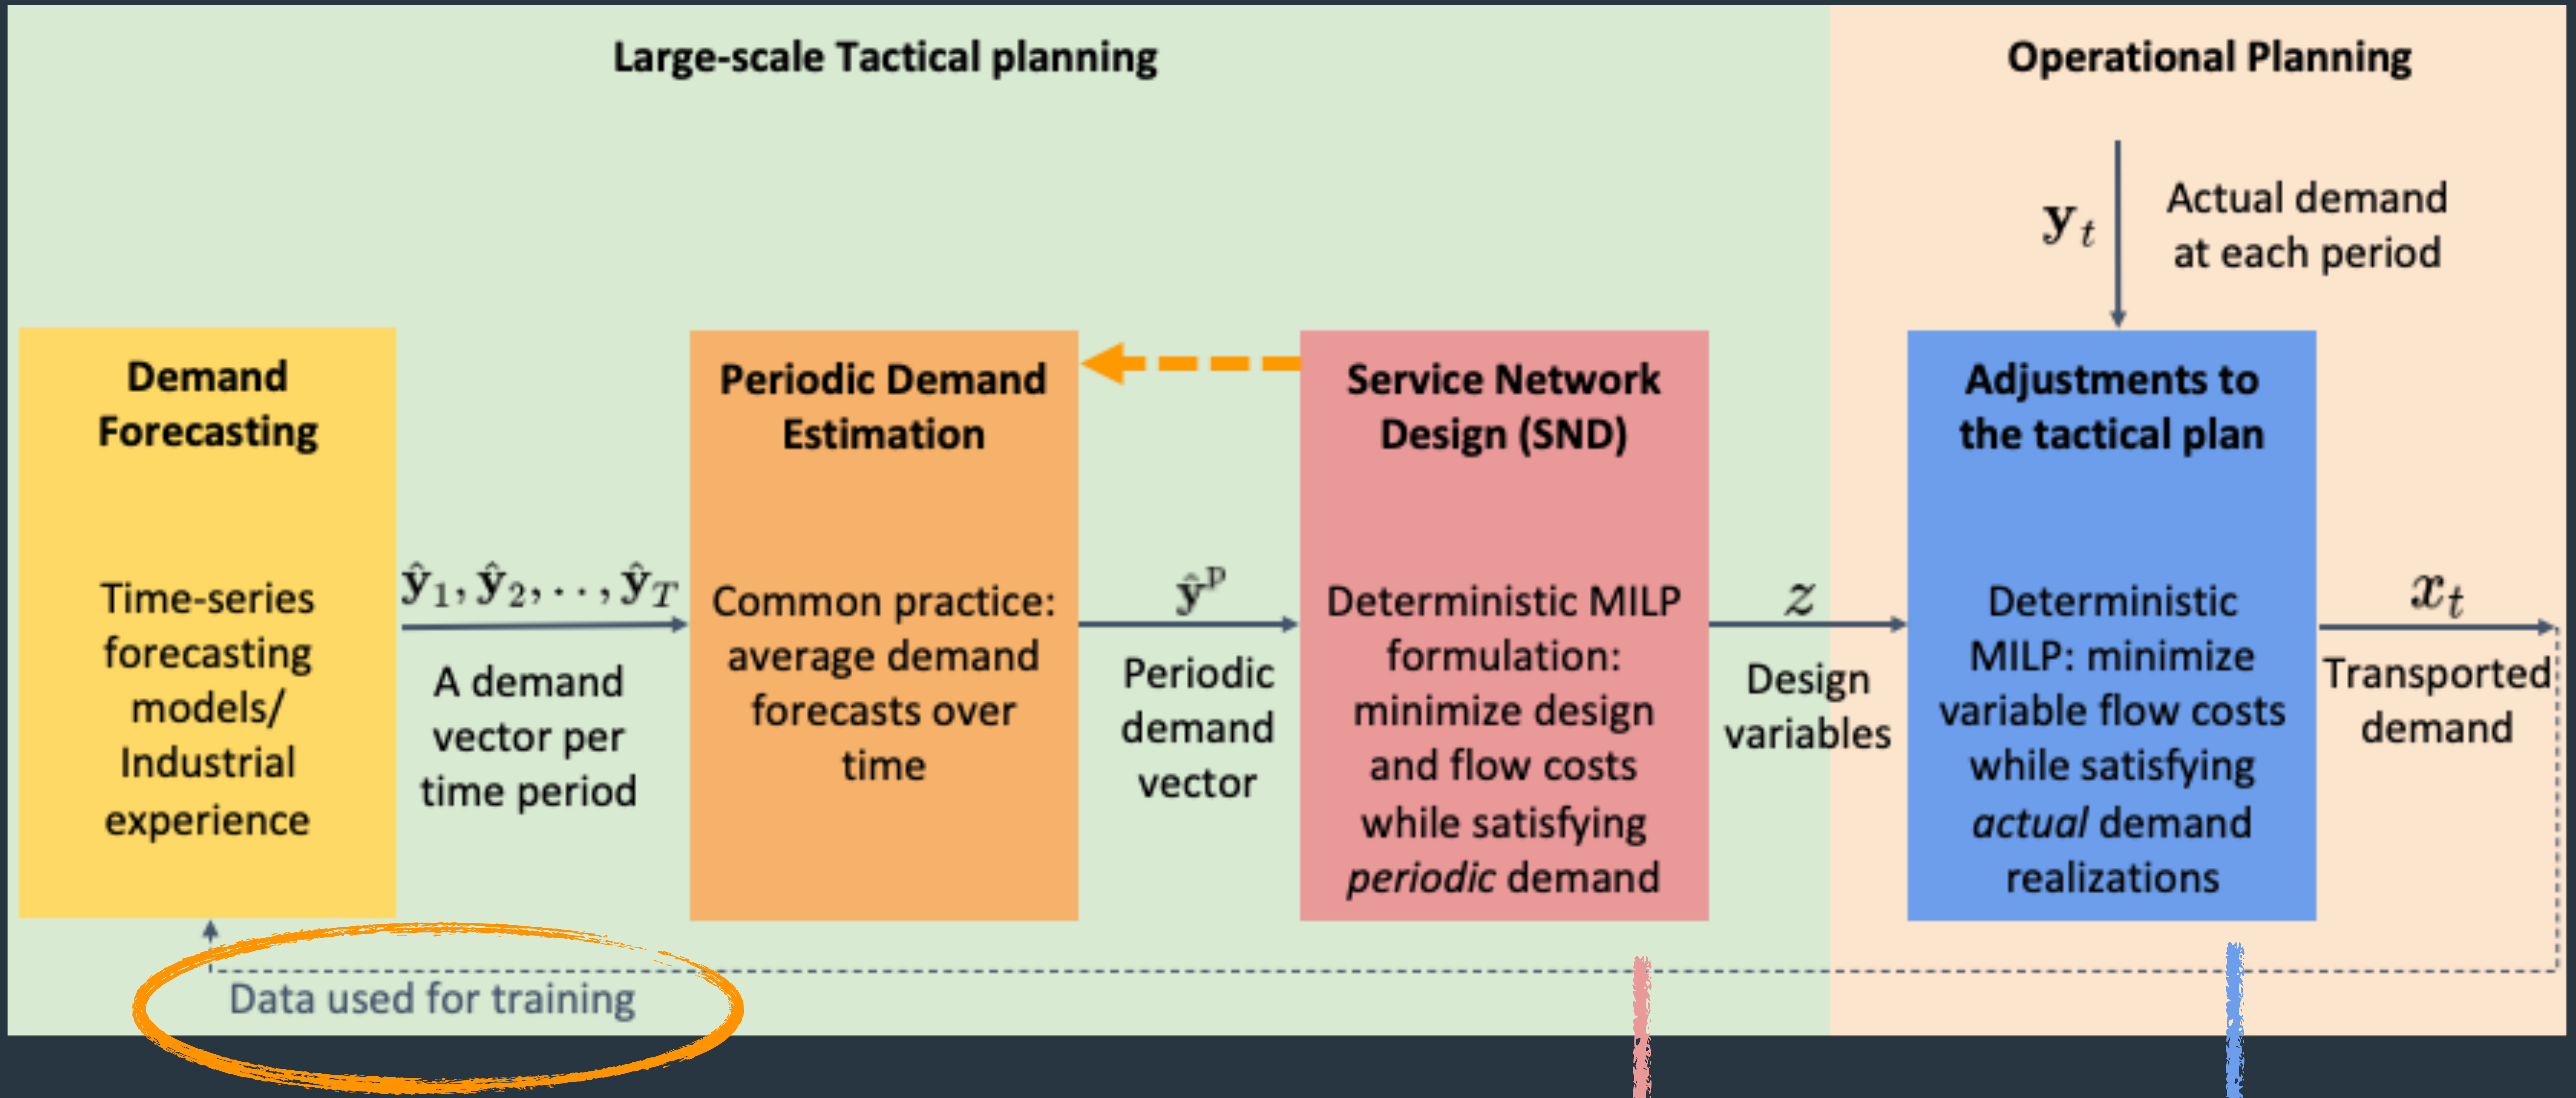

Data of transported demand is at the operational level

May be constrained by the supply – censored (left or right) or truncated data

MCND: Multicommodity Capacitated Fixed-charge Network Design

wMCND: MCND with fixed design variables

THE PERIODIC DEMAND ESTIMATION PROBLEM

Average of demand forecasts

$$\mathbf{y}^{\mathrm{p}} = \boldsymbol{\alpha} \odot \mathbf{y}^{\mathrm{p}}_{\mathrm{mean}}$$

Vector of deviation coefficients

$\boldsymbol{\alpha} \in \mathbb{R}^K_+$

Element-wise multiplication: $y^{\mathrm{p}}_k = \alpha_k \, y^{\mathrm{p}}_{\mathrm{mean},\,k}$

$$\textbf{ePDE} \quad \left| \begin{array}{ll} \min_{\boldsymbol{\alpha}} & C^{\mathrm{PDE}}(\boldsymbol{\alpha}, \mathbf{z}, \mathbf{x}, \mathbf{x}_1, .., \mathbf{x}_T) \\ \text{s.t.} & \mathbf{y}^{\mathrm{p}} = \boldsymbol{\alpha} \odot \mathbf{y}^{\mathrm{p}}_{\mathrm{mean}} \\ & \boldsymbol{\alpha} \leq \boldsymbol{\alpha}_{\max} \\ & \boldsymbol{\alpha} \geq \boldsymbol{\alpha}_{\min} \\ & (\mathbf{z}, \mathbf{x}, \mathbf{x}_1, .., \mathbf{x}_T) \in \operatorname{argmin}_{\mathbf{z}', \mathbf{x}', \mathbf{x}'_1, .., \mathbf{x}'_T} \mathbf{MCND} - \mathbf{wMCND}(\mathbf{y}^{\mathrm{p}}, \mathbf{z}', \mathbf{x}', \mathbf{x}'_1, .., \mathbf{x}'_T) \end{array} \right.$$

- With $\mathbf{y}^{\mathrm{p}}_{\min} = \min_{t=1,..,T} \{\mathbf{y}_t\}$ $\quad \mathbf{y}^{\mathrm{p}}_{\min} = \boldsymbol{\alpha}_{\min} \odot \mathbf{y}^{\mathrm{p}}_{\mathrm{mean}}$
- With $\mathbf{y}^{\mathrm{p}}_{\max} = \max_{t=1,..,T} \{\mathbf{y}_t\}$ and $\mathbf{y}^{\mathrm{p}}_{\max} = \boldsymbol{\alpha}_{\max} \odot \mathbf{y}^{\mathrm{p}}_{\mathrm{mean}}$

- Based on per period demand forecasts, estimate periodic demand as a deviation from average forecasts
- Intuitive interpretation
- Solve problem using clustering techniques and a heuristic combined with a general purpose MIP solver

A HIGH-VALUE PROBLEM

Case study from the Canadian National Railways

170 commodities

10 weeks planning horizon

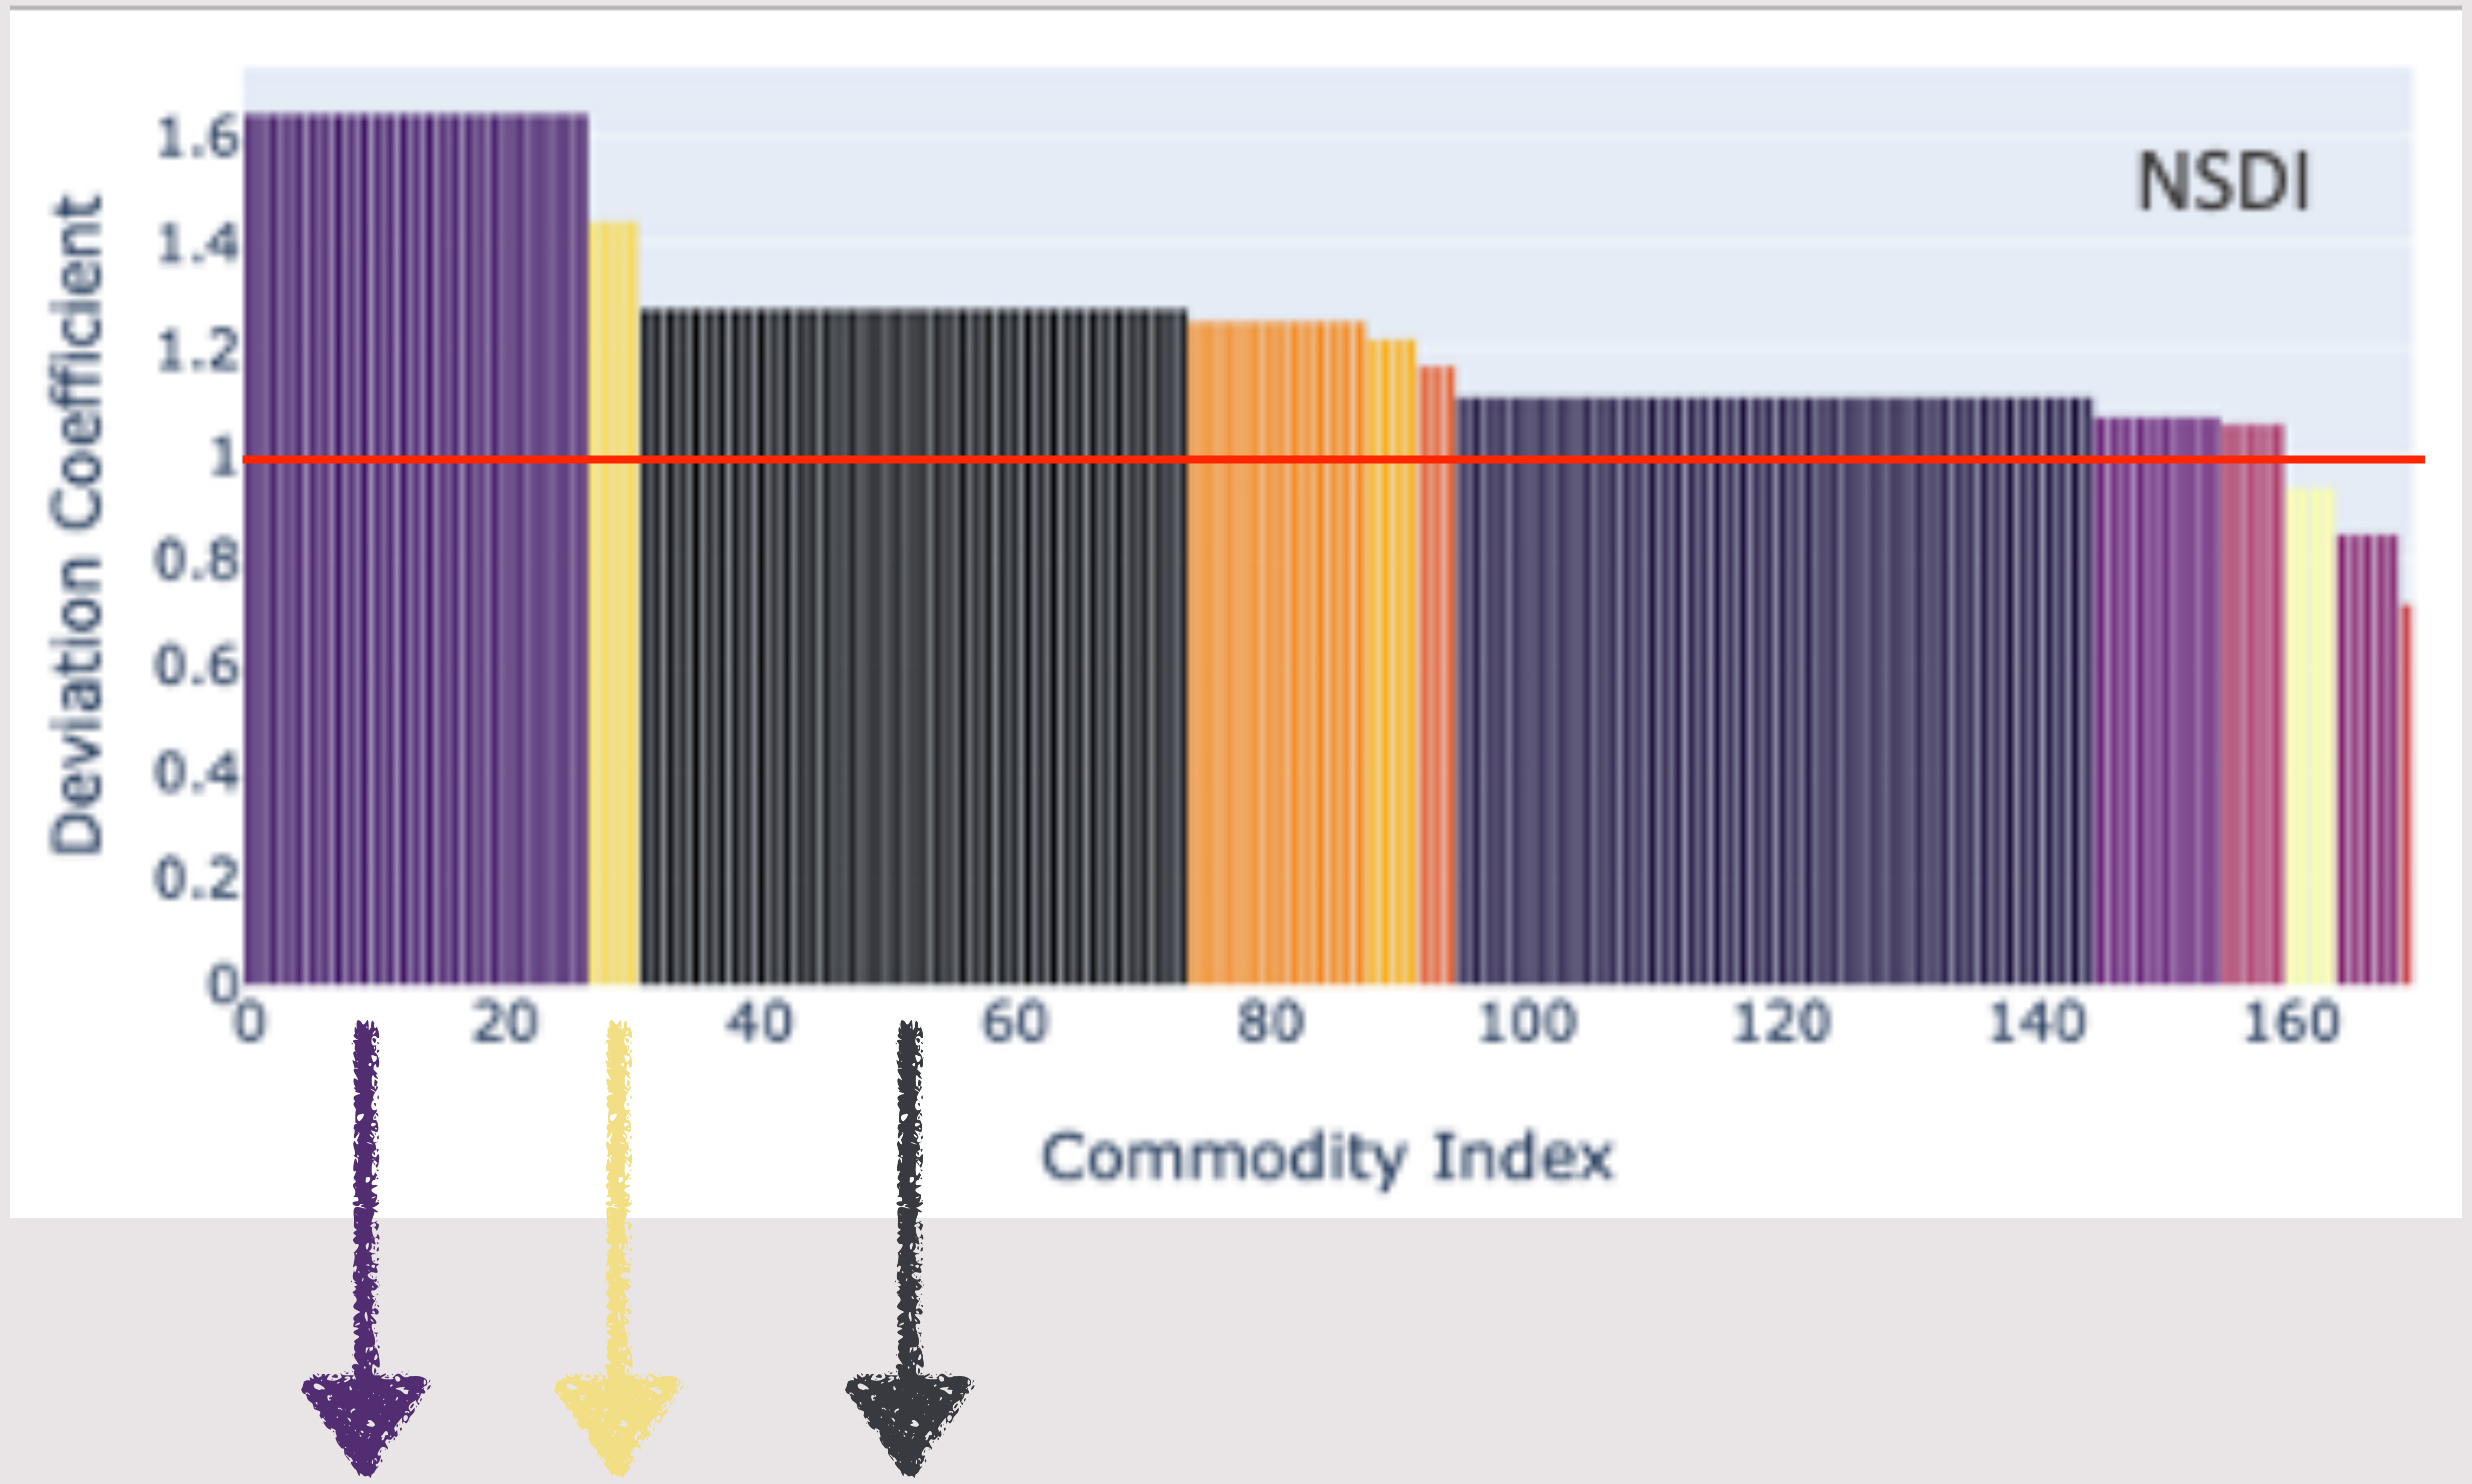

Average

Each colour: cluster of commodities having the same value of α_k

One bar per commodity

- High value using information from downstream decision-making problem when identifying the demand scenario
- Large cost reductions – more than 15% – compared to using average forecasts
- For commodities where the problem is sensitive to large demand values: $\alpha_k > 1$

3

Predict Optimize Decide Interact

Measuring the impact

A case of counterfactual prediction

Companies involved: IVADO Labs and Air Canada

Laage, Frejinger, Lodi, Rabusseau, Assessing the impact: Does an Improvement to a Revenue Management System Lead to an Improved Revenu?, arXiv:2101.10249

Greta Laage, 2nd place for the Anna Valicek Award from the Airline Group of the International Federation of Operational Research Societies

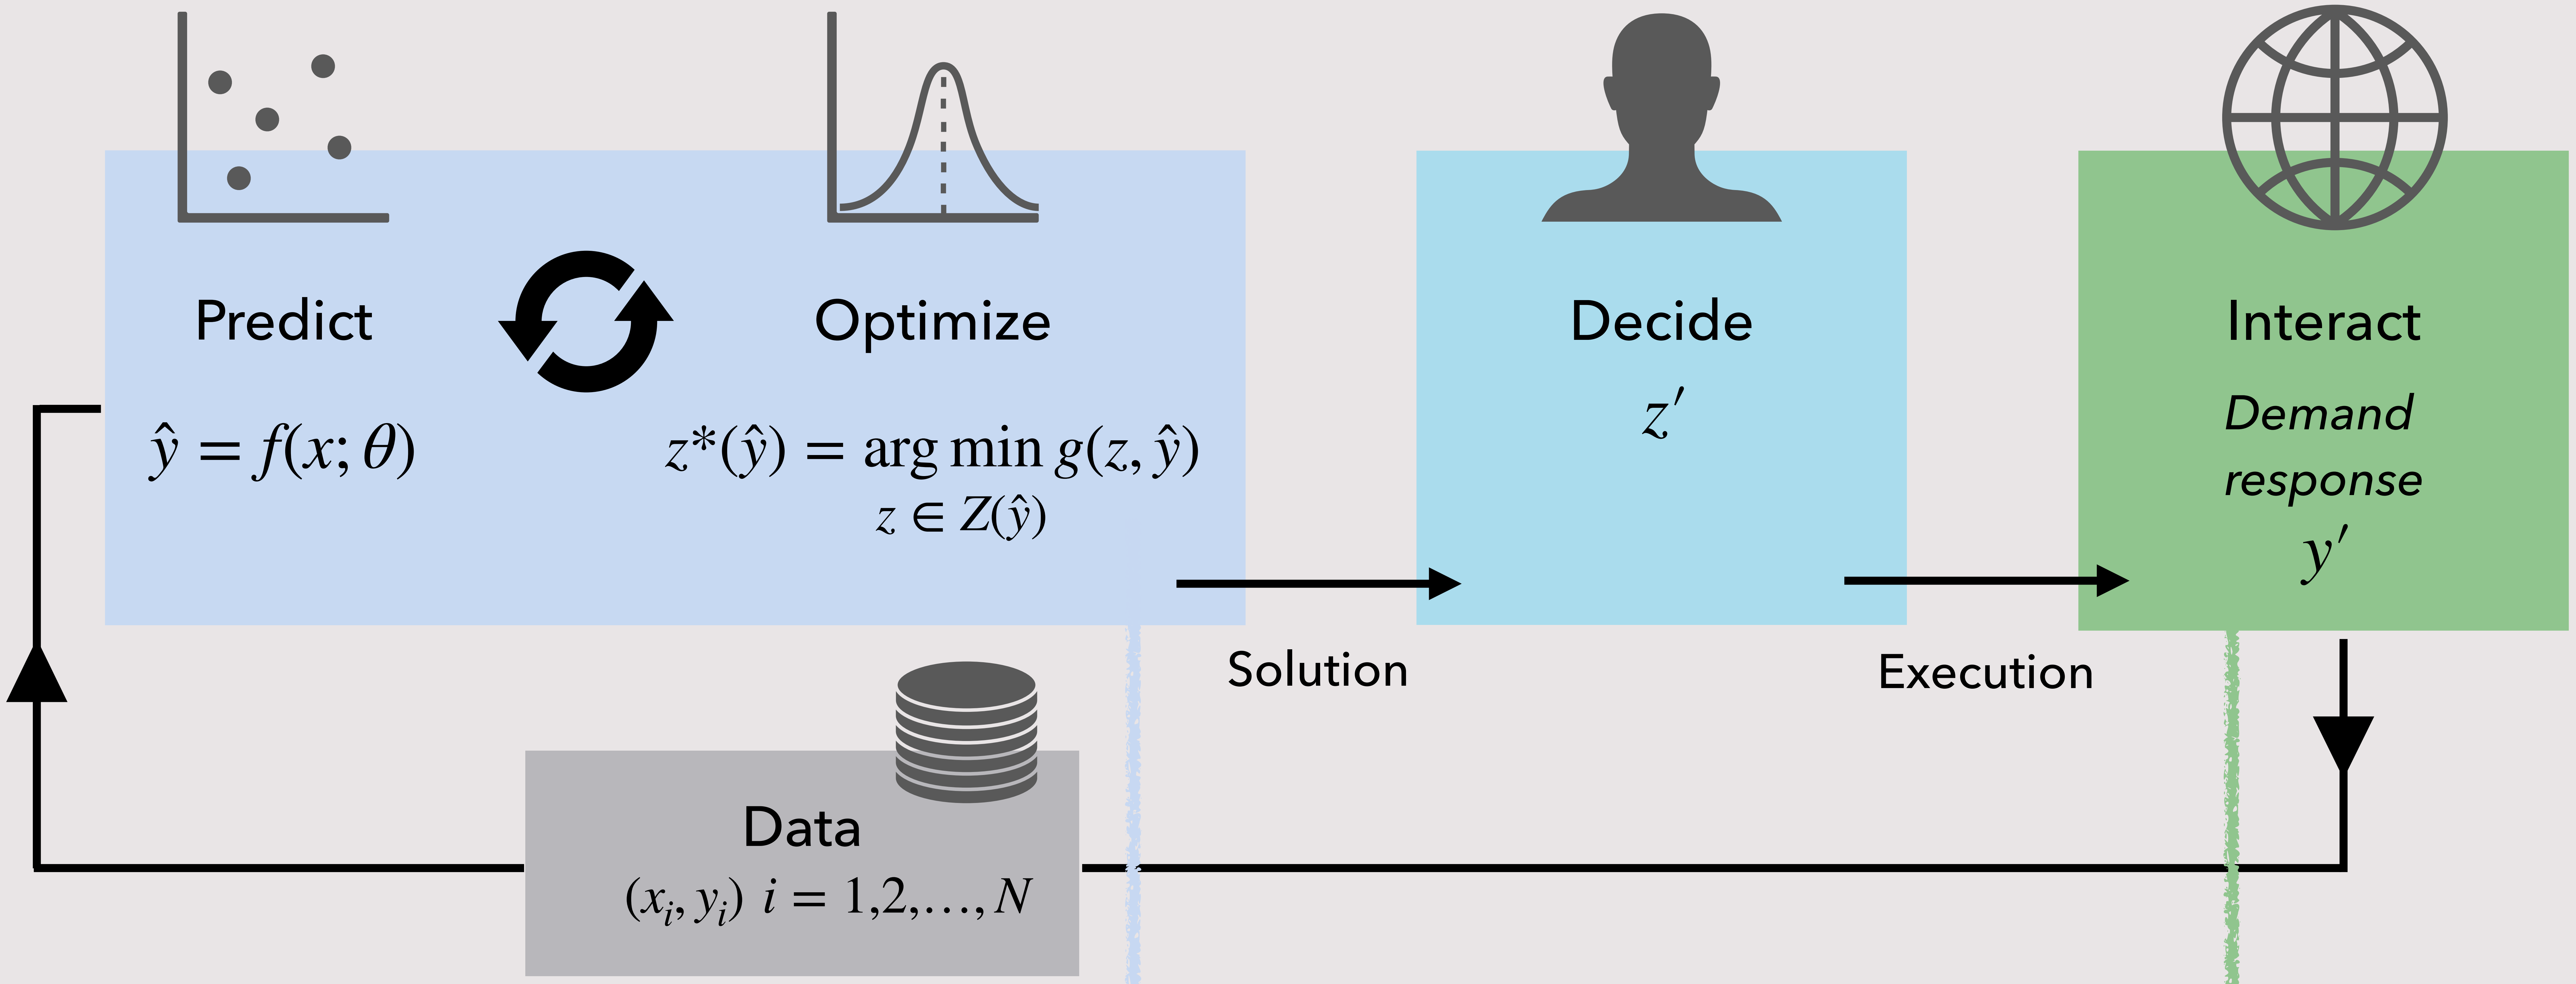

Value/cost of the solution assuming that the models are correct.

Actual value/cost in the real environment.

How can we assess the actual impact of a new (or modified) decision-support system?

Proof of concept: test system on a limited scale (e.g., subset of origin-destination pairs) and compare performance to what would have been the performance business as usual

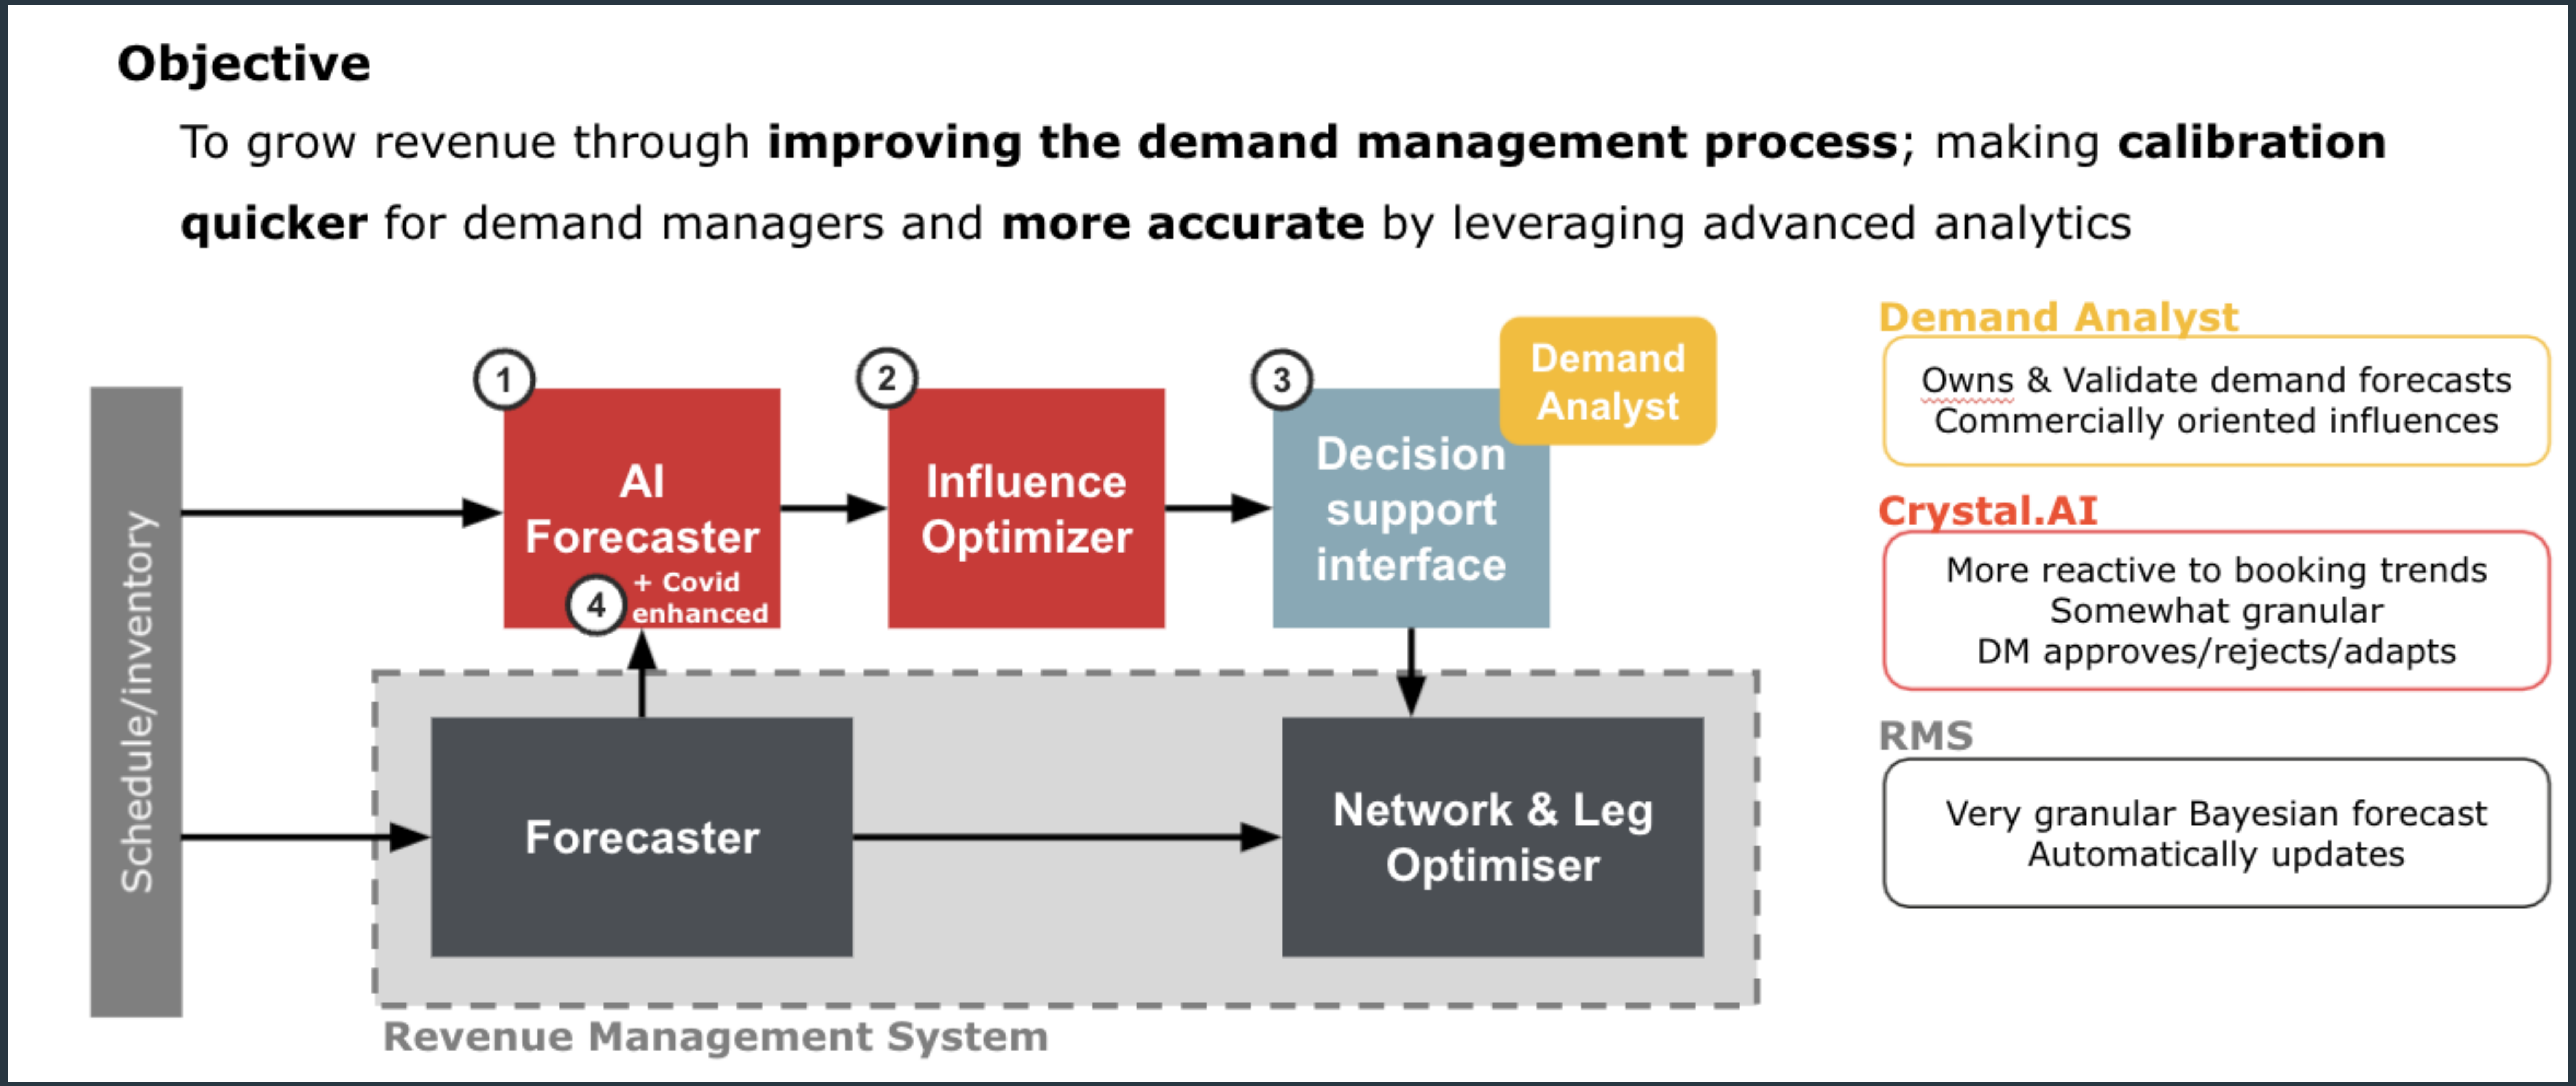

Source: IVADO Labs, Air Canada

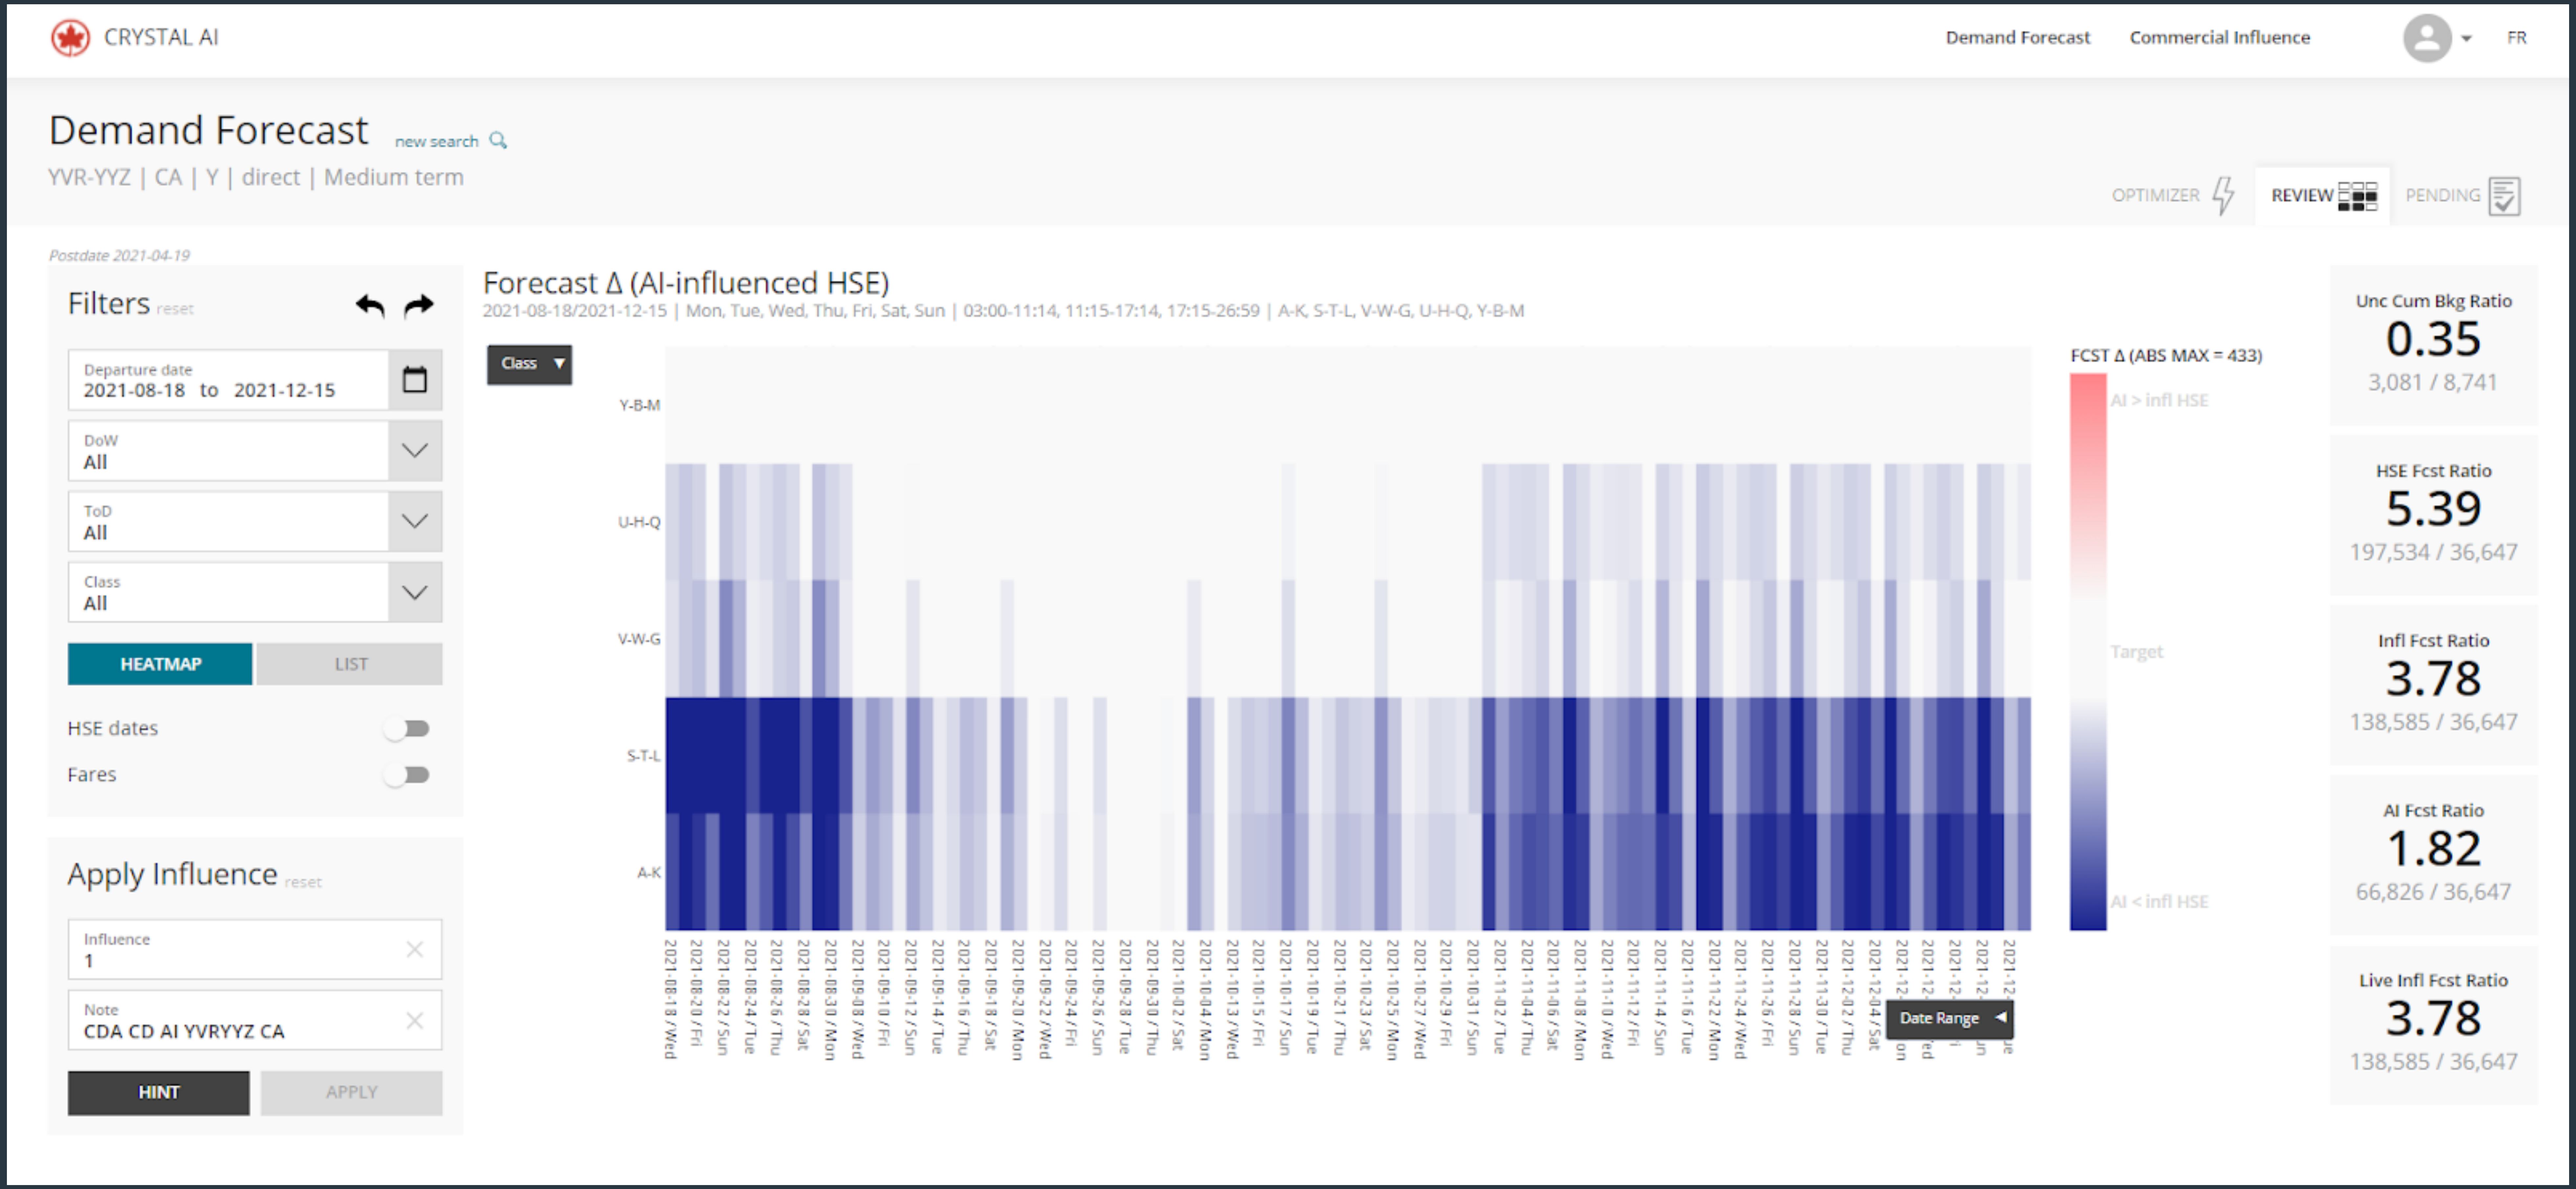

Source: IVADO Labs, Air Canada

Demand analyst

User interface: display discrepancies between base forecasts and AI forecaster

Decide influences for the system (impact on optimization)

ASSESSING THE IMPACT

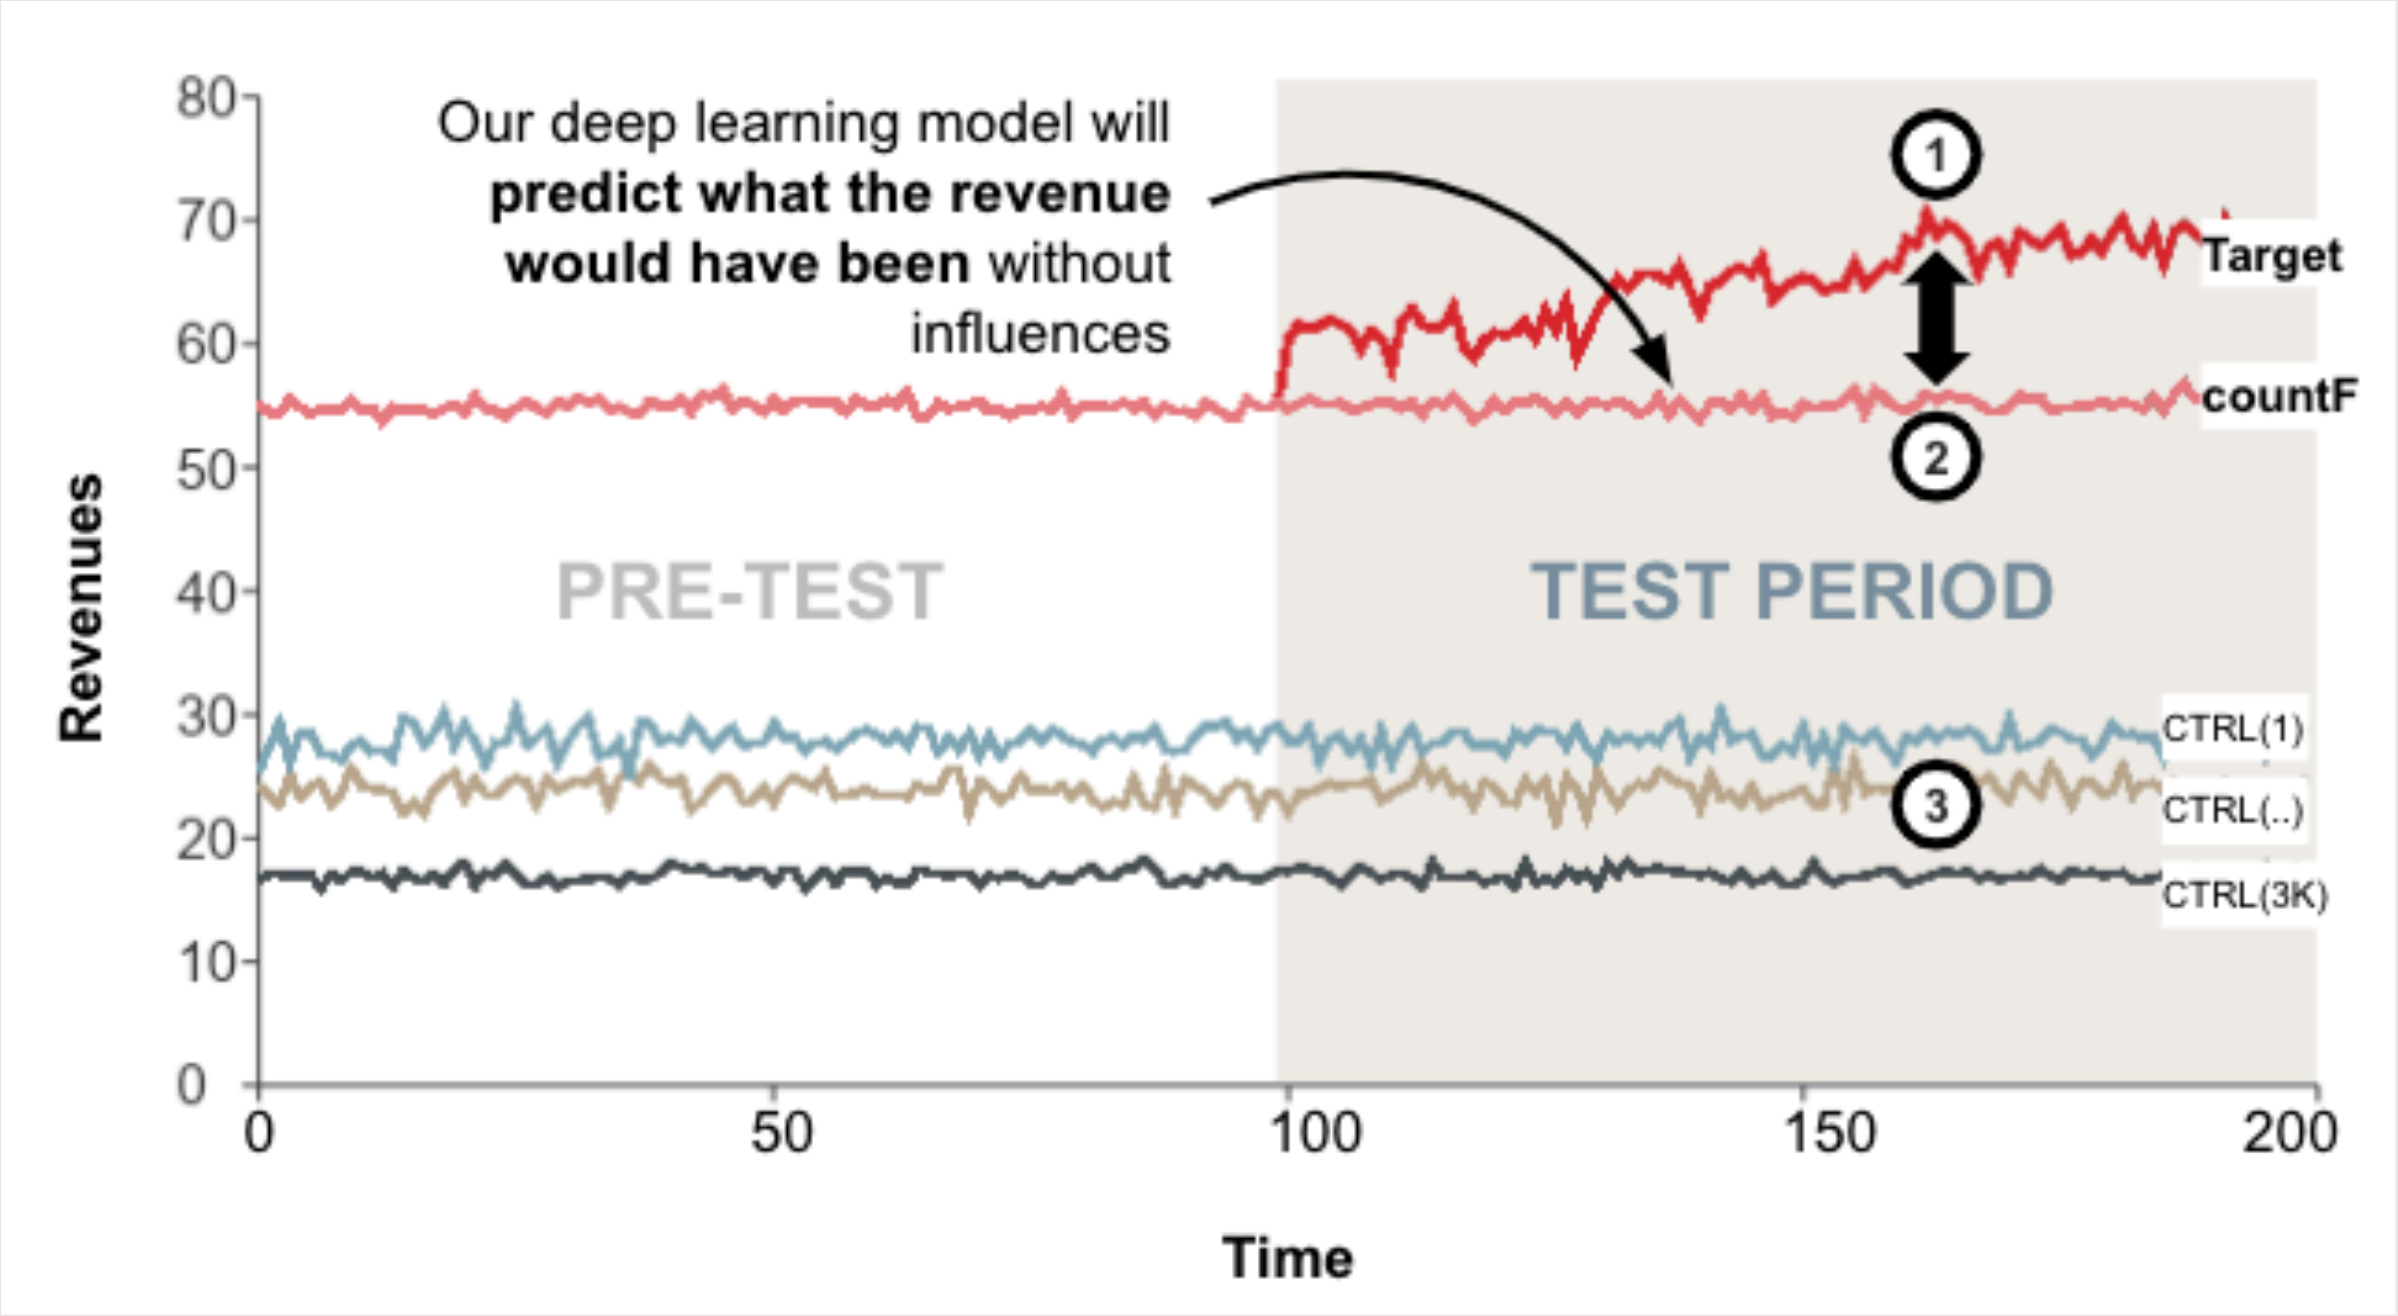

Source: IVADO Labs, Air Canada

1. Observed: actual revenue of treated OD
2. Unobserved: untreated revenue for treated OD, Y_t
3. Observed: 3,000 control ODs, X_t

Objective: Estimate total impact (all treated ODs over the whole treatment / test period

Counterfactual prediction: $Y_t = F(X_t)$

- Examples of existing approaches
  - Simulation (Weatherford and Belobaba, 2002, Fiig et al., 2019): does not assess *actual* impact
  - Year over year change: easy to compute but unreliable
  - A-B testing: can be noisy and adequate control ODs may not exist
- Counterfactual prediction: adaptable to various treatment lengths, could measure relatively small impacts

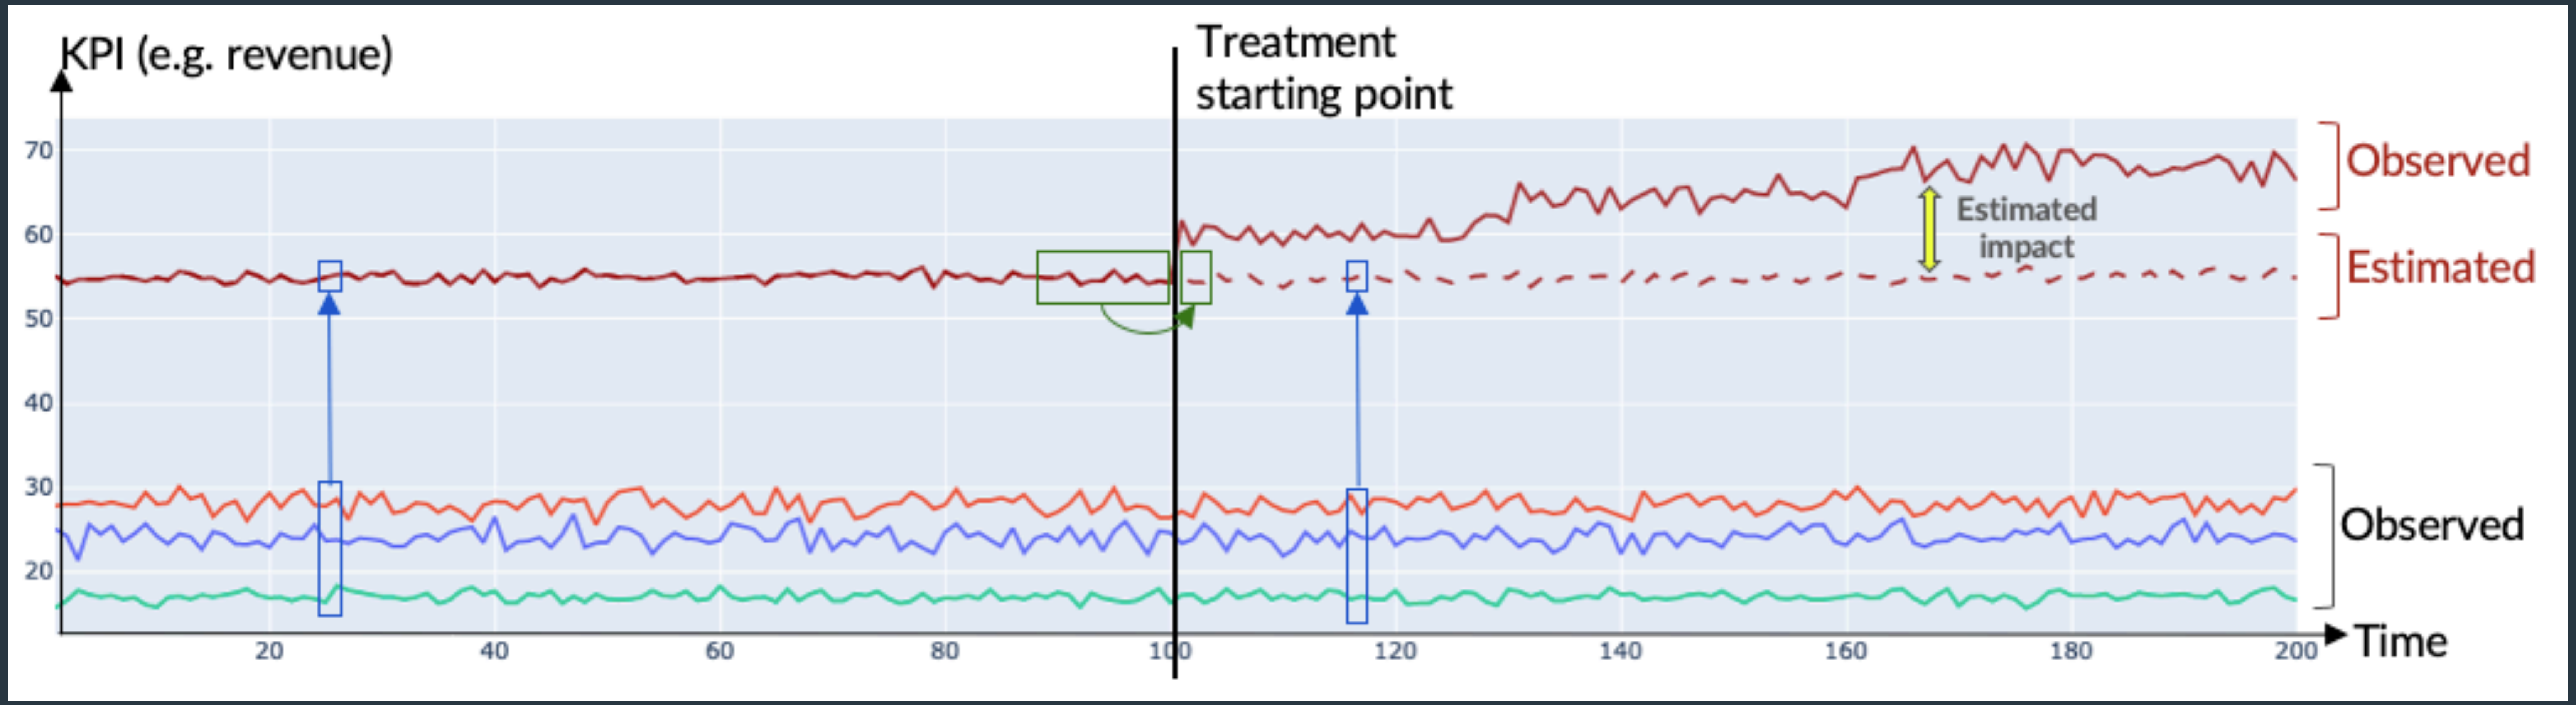

Our setting: Multiple treated units, a large set of controls, relatively small impact

Literature: mostly focused on macro enomoic settings. E.g., impact of the German reunification

COUNTERFACTUAL PREDICTION MODELS

- Synthetic control
  - Differences-in-differences (Ashenfelder and Card, 1985)
  - Abadie-Diamond-Hainmueller Synthetic Control Method (Abadie and Gardezabal, 2003, Abadie et al., 2010)
  - Constrained regression (Doudchenko and Imbens, 2016)
- Robust synthetic control (Amjad et al., 2018)
- Matrix completion with nuclear norm (Athey et al., 2018)
- Feed-forward neural network (can deal with multiple treated units)

Percentage error per period

Post-treatment periods

Simulated impact (random variable with known mean μ_ϵ and variance)

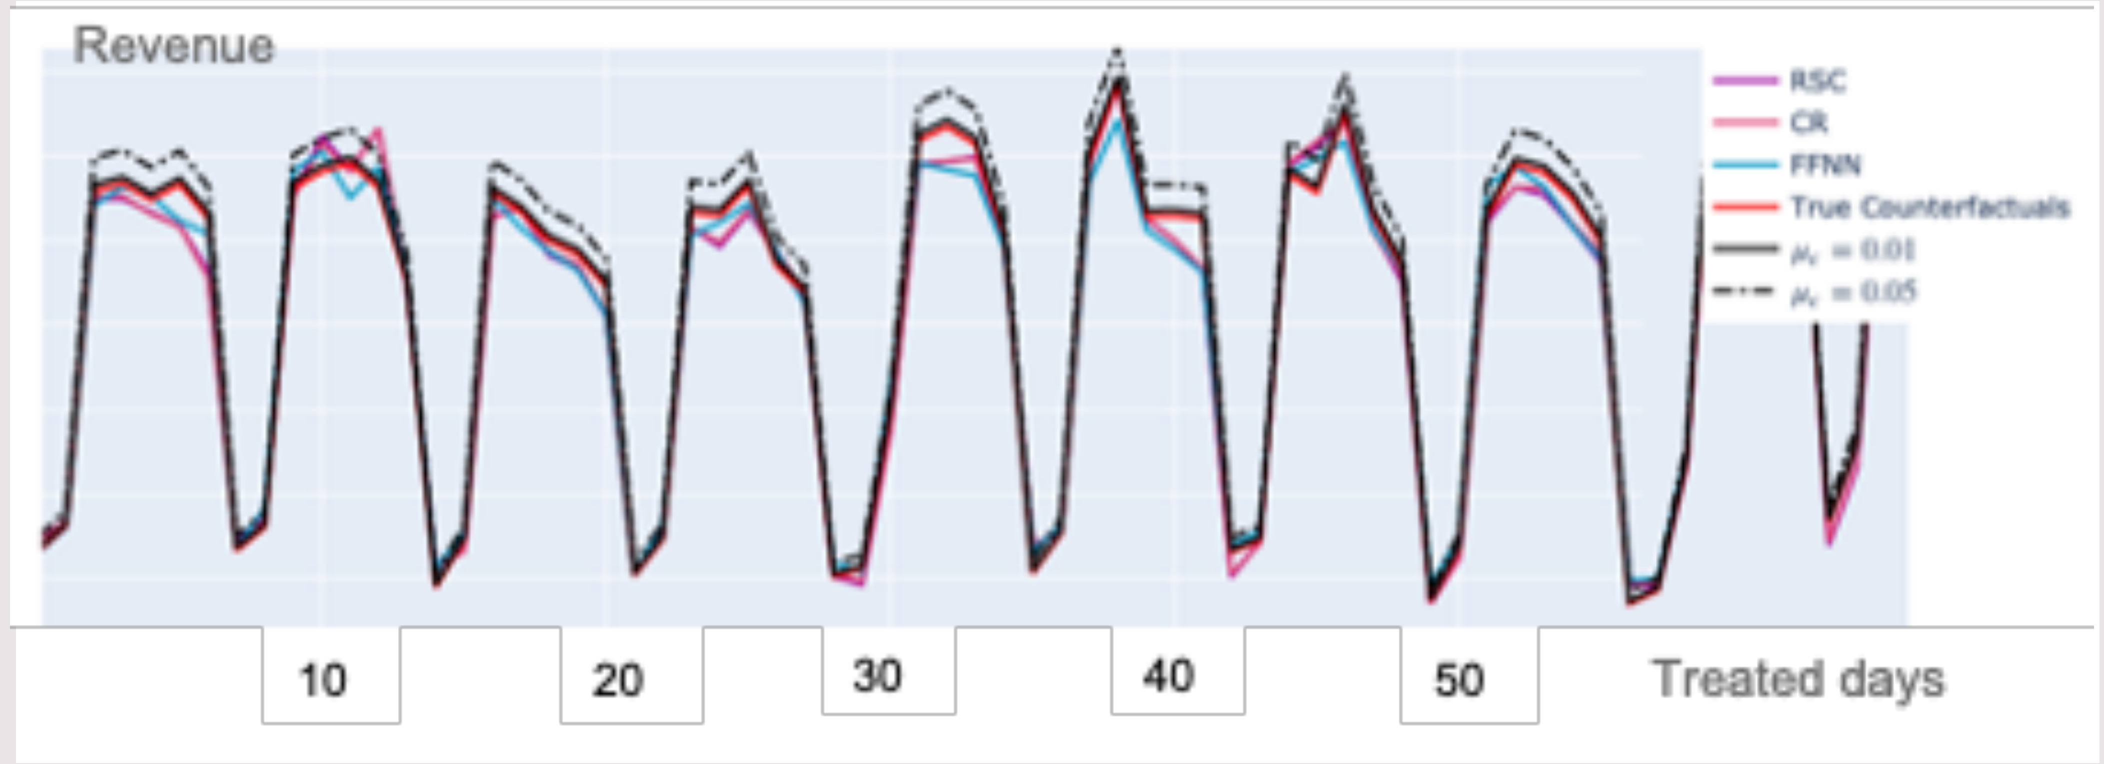

μ_ϵ	0.01	0.02	0.03	0.05
True	1.0%	2.0%	3.0%	5.1%
RSC	1.7%	2.6%	3.7%	5.7%
CR	1.5%	2.5%	3.5%	5.6%
FFNN	0.6%	1.6%	2.6%	4.7%

ACCURATE RESULTS

- 30 treated ODs (15 non directional)
- 317 control ODs carefully selected (unaffected by treatment)
- Observations January 2013 - February 2020
- 15 pseudo-treatment periods of 6 months
- Several counterfactual prediction models have similar performance
- Best performing models predict total revenue with total percentage error of less than 1%
- Accurate estimation of (simulated) impact

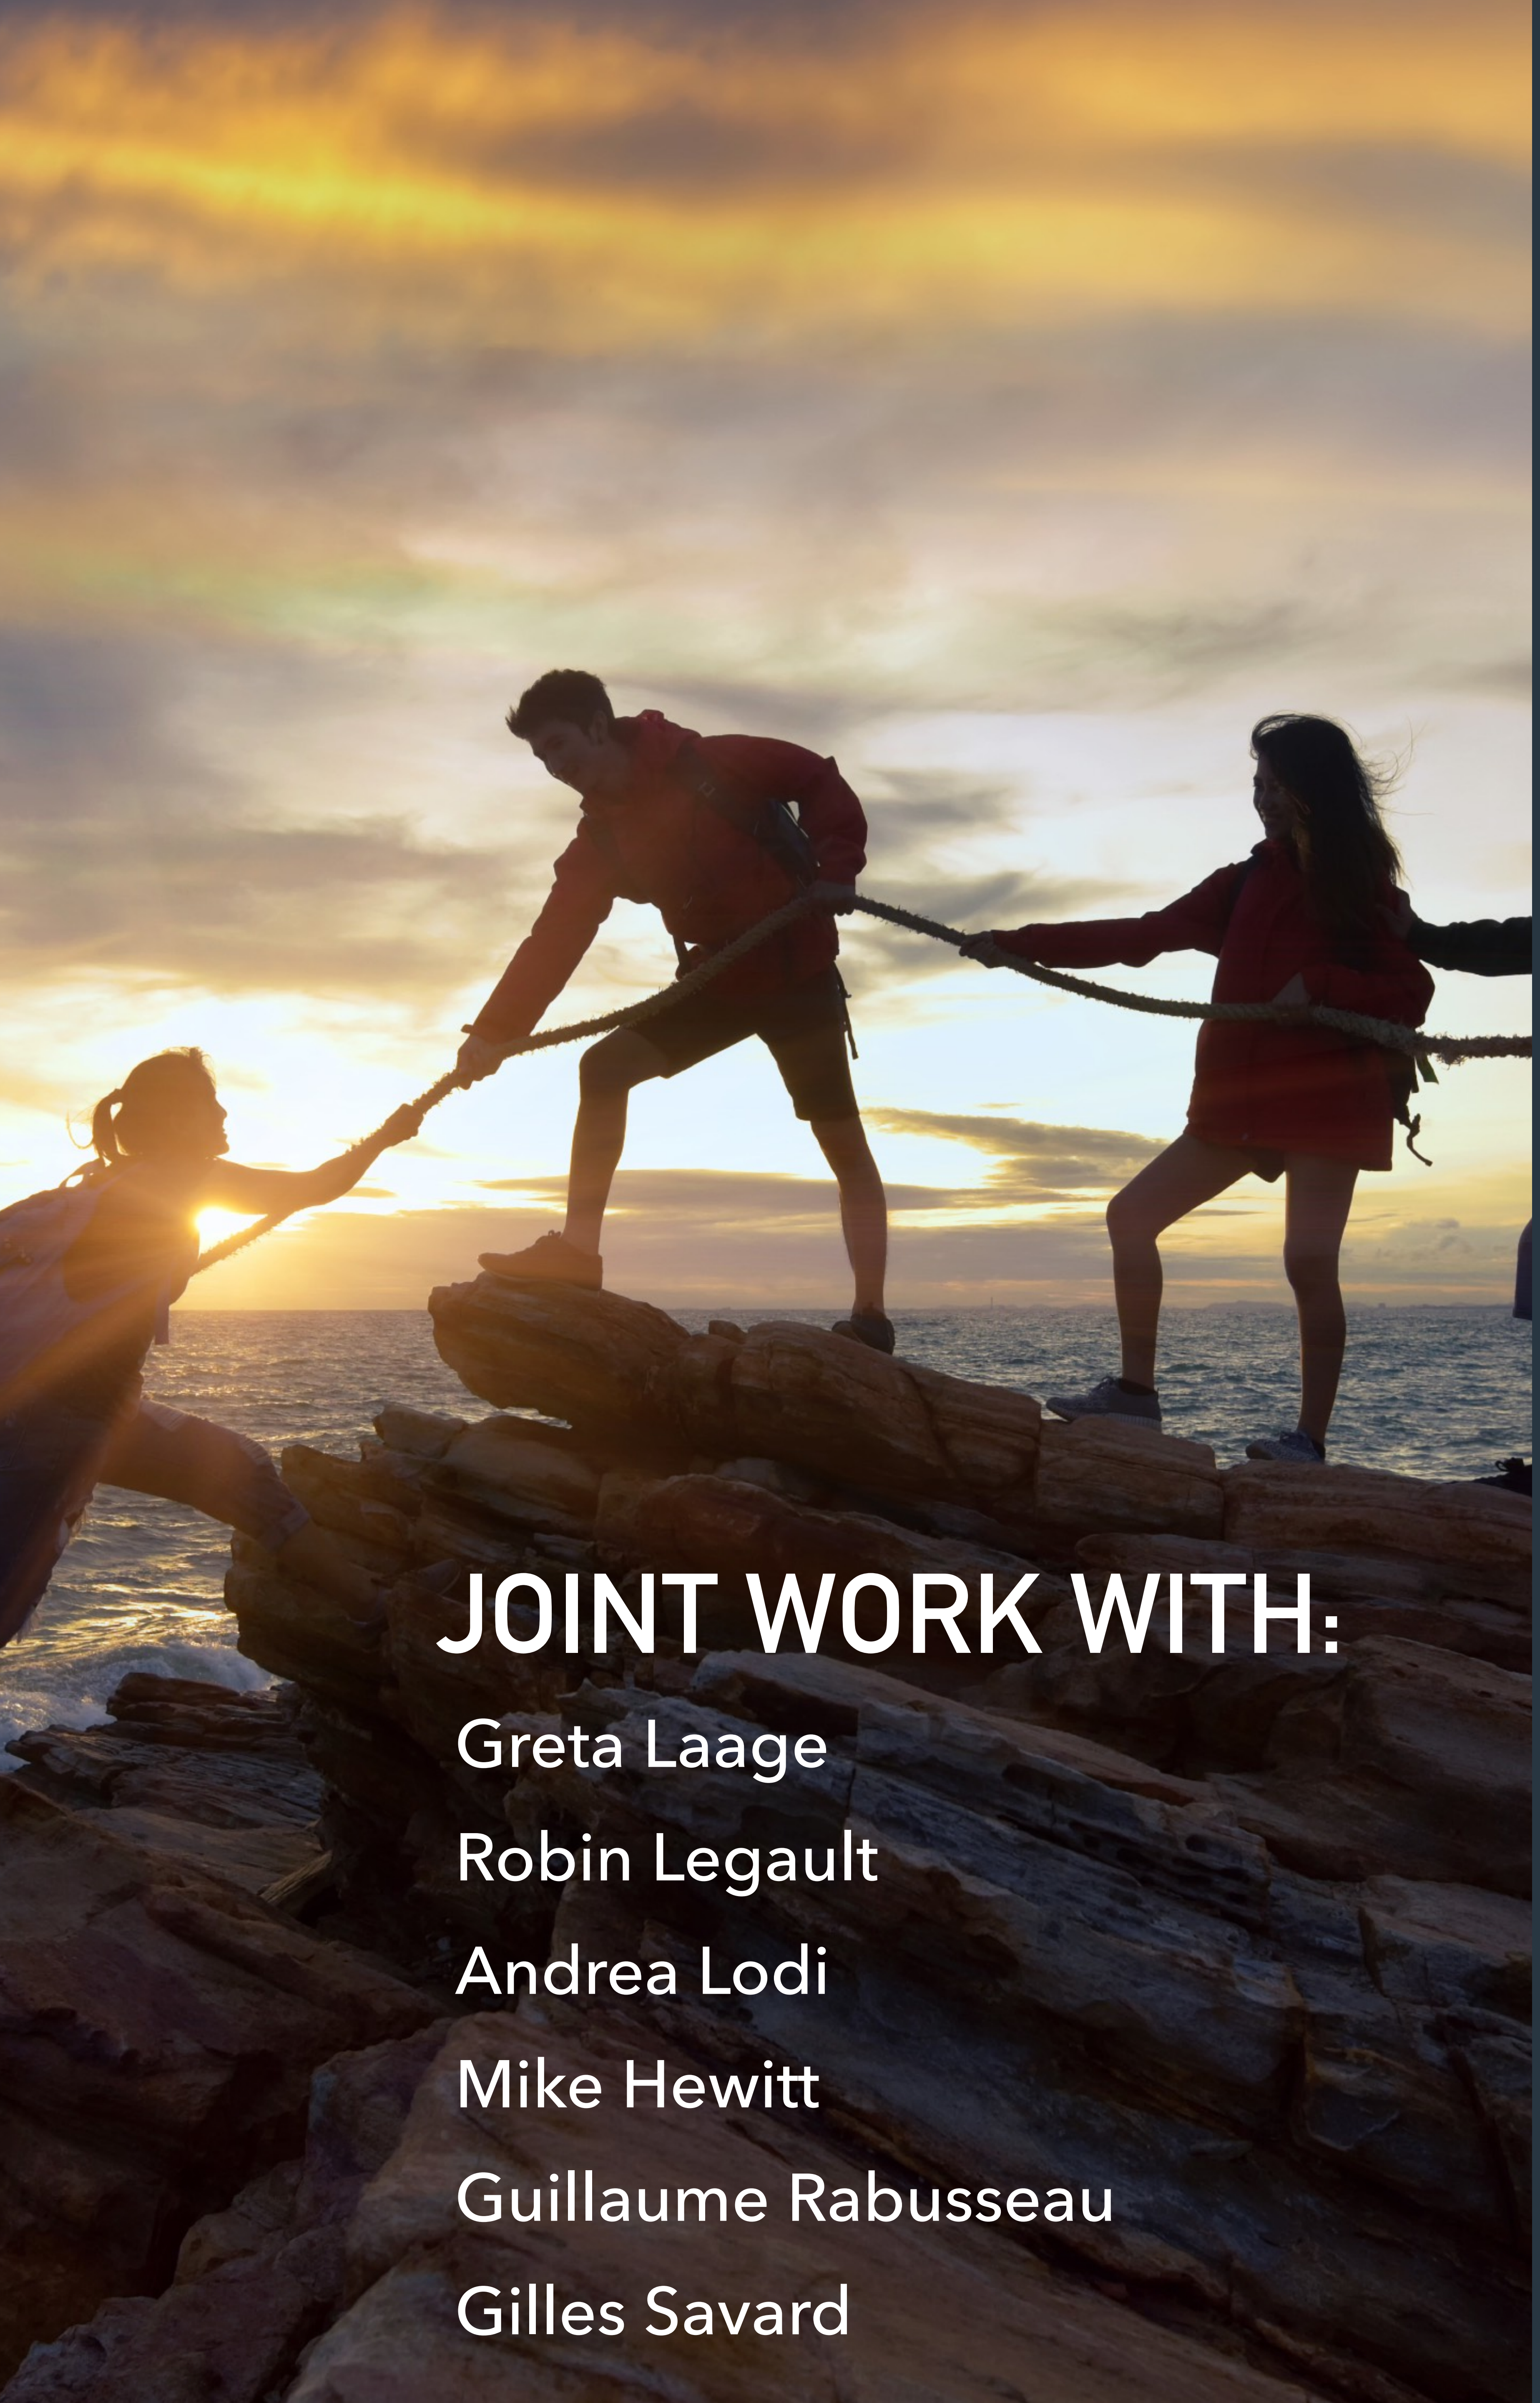

JOINT WORK WITH:

Greta Laage
Robin Legault
Andrea Lodi
Mike Hewitt
Guillaume Rabusseau
Gilles Savard

CONCLUSION

- Predictions often used to make decisions
- Integrating prediction and (discrete) optimization can be of high value to transport and mobility application
- Important problems arise in this context
  - Challenging discrete optimization problems with endogenous demand uncertainty
  - Decision awareness in learning
  - Measuring actual impact and reducing post-decision disappointment
- Research: several open research questions
- Practice: innovative pragmatic solutions related to decision awareness

Thank you!

emma.frejinger@umontreal.ca

emmafrejinger.org

- Abadie, A., Diamond, A., and Hainmueller, J. Synthetic controlmethods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association, 105(490):493–505, 2010.
- Aguirregabiria, V. and Mira, P. Dynamic discrete choice structural models: A survey. Journal of Econometrics, 156(1):38–67, 2010.
- Ashenfelter, O. and Card, D. Using the longitudinal structure of earnings to estimate the effect of training programs. The Review of Economics and Statistics, 67(4):648–660, 1985.
- Amjad, M., Shah, D., and Shen, D. Robust synthetic control. The Journal of Machine Learning Research, 19(1):802–852, 2018.
- Athey, S., Bayati, M., Doudchenko, N., Imbens, G., and Khosravi, K. Matrix Completion Methods for Causal Panel Data Models, 2018.
- Bengio, Lodi, Prouvost, Machine learning for combinatorial optimization: A methodological tour d'horizon, EJOR (290):405-421, 2021.
- Bhuiyan, Medal, Harun, A stocahstic programming model with endogenous and exogenous uncertainty for reliable network design under random disruption, European Journal of Operational Research 285(2):670-694, 2020.
- Crainic, T. G., Hewitt, M., Maggioni, F., and Rei, W. Partial benders decomposition: General methodology and application to stochastic network design. Transportation Science, 2020.
- Dan and Marcotte, Competitive facility location with selfish users and queues, Operations Research 67(2):479-497, 2019.
- Doudchenko and Imbens, Balancing, regression, difference-in-differences and synthetic control methods: A synthesis, Technical report, National Bureau of Economic Research, 2016.
- Elmachtoub, Grigas, P. Smart "predict, then optimize". *Management Science*, 2021.
- Ferber, A., Wilder, B., Dilkina, B., and Tambe, M. MIPaaL: Mixed integer program as a layer. Proceedings of the AAAI Conference on Artificial Intelligence, 34(02):1504–1511, 2020.
- Gevers, Identification for Control: From the Early Achievements to the Revival of Experiment Design, European Journal of Control 11(4-5):335-352, 2005.
- Gilbert, Marcotte, Savard, Mixed-logit network pricing. Comput Optim Appl, 57:105-127, 2014.
- Gilbert, Marcotte, Savard, A numerical stody of the logit network pricing problem, Transportation Science 49(3):706-719, 2015.
- Grimm, C., Barreto, A., Singh, S., and Silver, D. The value equivalence principle for model-based reinforcement learning. NeurIPS, 2020.
- Kotary, J., Fioretto, F., Hentenryck, P. V., and Wilder, B. End-to-end constrained optimization learning: A survey, arXiv: 2103.16378, 2021.
- Mandi, J., Demirovic, E., Stuckey, P. J., and Guns, T. Smart predict-and-optimize for hard combinatorial optimization problems. Proceedings of the AAAI Conference on Artificial Intelligence, 34(02):1603–1610, 2020.
- Moghaddass, M. and Terekhov, D. Inverse integer optimization with multiple observations. Optimization Letters, 15:1061–1079, 2021.
- Ng, A. Y. and Russell, S. J. Algorithms for inverse reinforcement learning. In ICML, 663–670, 2000.
- Paneque, Bierlaire, Gendron, Sharif Azadeh, Integrating advanced discrete choice models in mixed integer linear optimization, Transportation Research Part B (146):26-49, 2021.
- Pogančić, M. V., Paulus, A., Musil, V., Martius, G., and Rolinek, M. Differentiation of blackbox combinatorial solvers. In International Conference on Learning Representations, 2020.