Machine Learning and Optimization for Enhanced Decision-Making Under Uncertainty

EMMA FREJINGER

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

CIRRELT, Mila

Scientific Advisor, IVADO Labs

EPFL, May 1st, 2025

Climbing to the moon, Chris Rowney www.devianart.com

My moonshot: From data to real-world impact

Solve large-scale discrete optimization problems subject to uncertainty (exogenous or endogenous) while leveraging recent relevant data

Outline

1. Contextual stochastic optimization lens: Foundational concepts and motivation
2. From data to decision
   - Two decision pipelines: from context to decision
   - Three training pipelines: how to use data to improve prescriptive performance
3. From data to real-world impact
   - What do we know and what topics require more research

ILLUSTRATIVE EXAMPLE: CONTEXTUAL STOCHASTIC OPTIMIZATION AND EXOGENOUS UNCERTAINTY

- Contextual stochastic optimization (CSO): For any context $x \in X$ identify path $\boldsymbol{z}^*(\boldsymbol{x})$ from East to West Los Angeles minimizing expected travel time

$$(\text{CSO}) \quad z^*(\boldsymbol{x}) \in \underset{\boldsymbol{z} \in \mathcal{Z}}{\operatorname{argmin}} \, \mathbb{E}_{\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})} \left[c\left(\boldsymbol{z}, \boldsymbol{y}\right) \right]$$

- Contextual information / features $\boldsymbol{x}$ revealed before solving the problem
  - Time of day, weather, visibility, …
- Random parameters: Travel times $\boldsymbol{y}$ follow conditional distribution $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})$

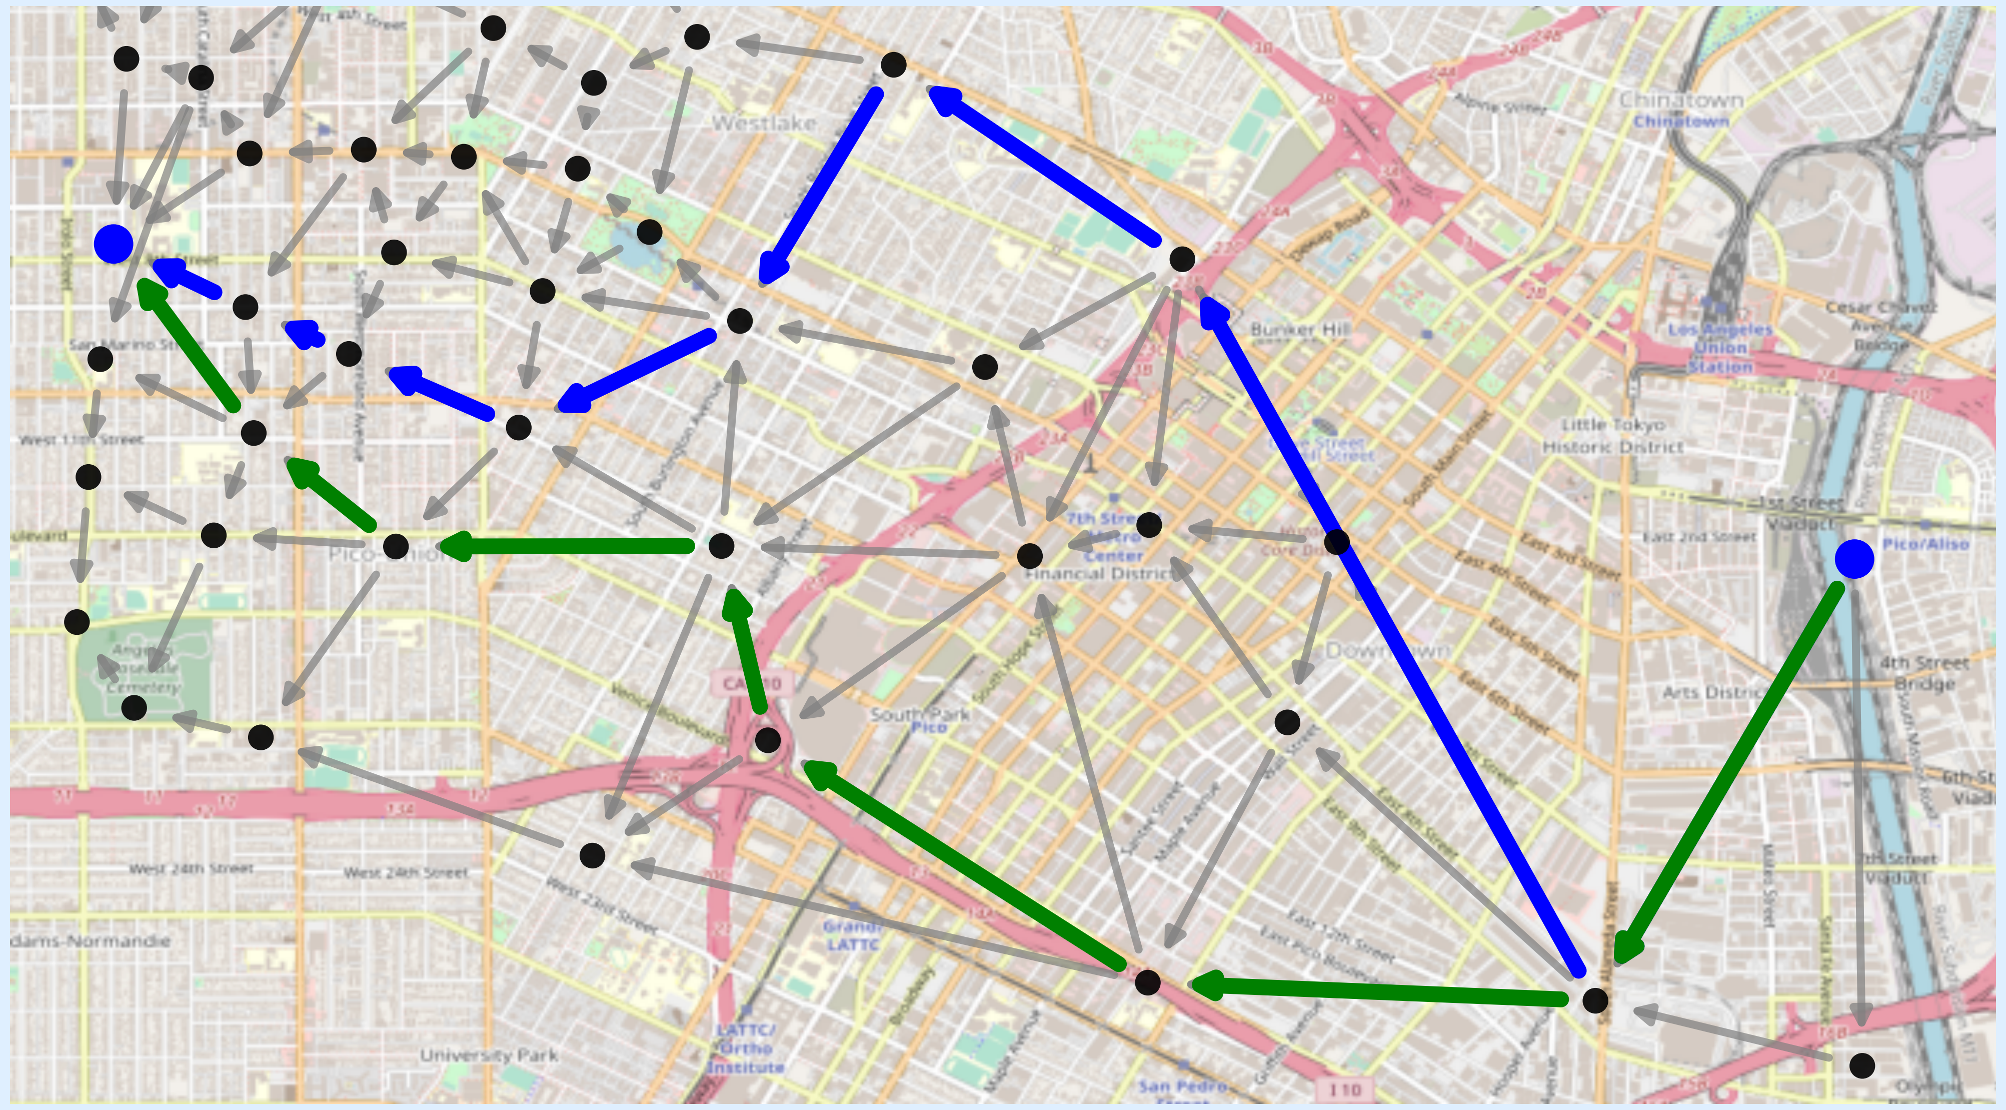

Depending on the context, the green or the blue path is optimal

Forel, Parmentier, Vidal, Explainable Data-driven Optimization: From Context to Decision and Back Again, In International Conference on Machine Learning, 2023.

ILLUSTRATIVE EXAMPLE: CONTEXTUAL STOCHASTIC DISCRETE OPTIMIZATION AND ENDOGENOUS UNCERTAINTY, COMPETITIVE FACILITY LOCATION

1

- CSO problem: For any context $x \in X$ identify a set of locations to open $\boldsymbol{z}^*(\boldsymbol{x})$ to maximize expected demand capture

$$z^*(\boldsymbol{x}) \in \underset{\boldsymbol{z} \in \mathcal{Z}}{\operatorname{argmax}} \, \mathbb{E}_{\mathbb{P}(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{z})} [r(\boldsymbol{z}, \boldsymbol{y})]$$

- Contextual information / features $\boldsymbol{x}$ revealed before solving the problem
  - Locations amenities, …
- Random parameters: Demand attracted to different locations $\boldsymbol{y}$ follow conditional distribution $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{z})$ that is endogenous, i.e. decision-dependent

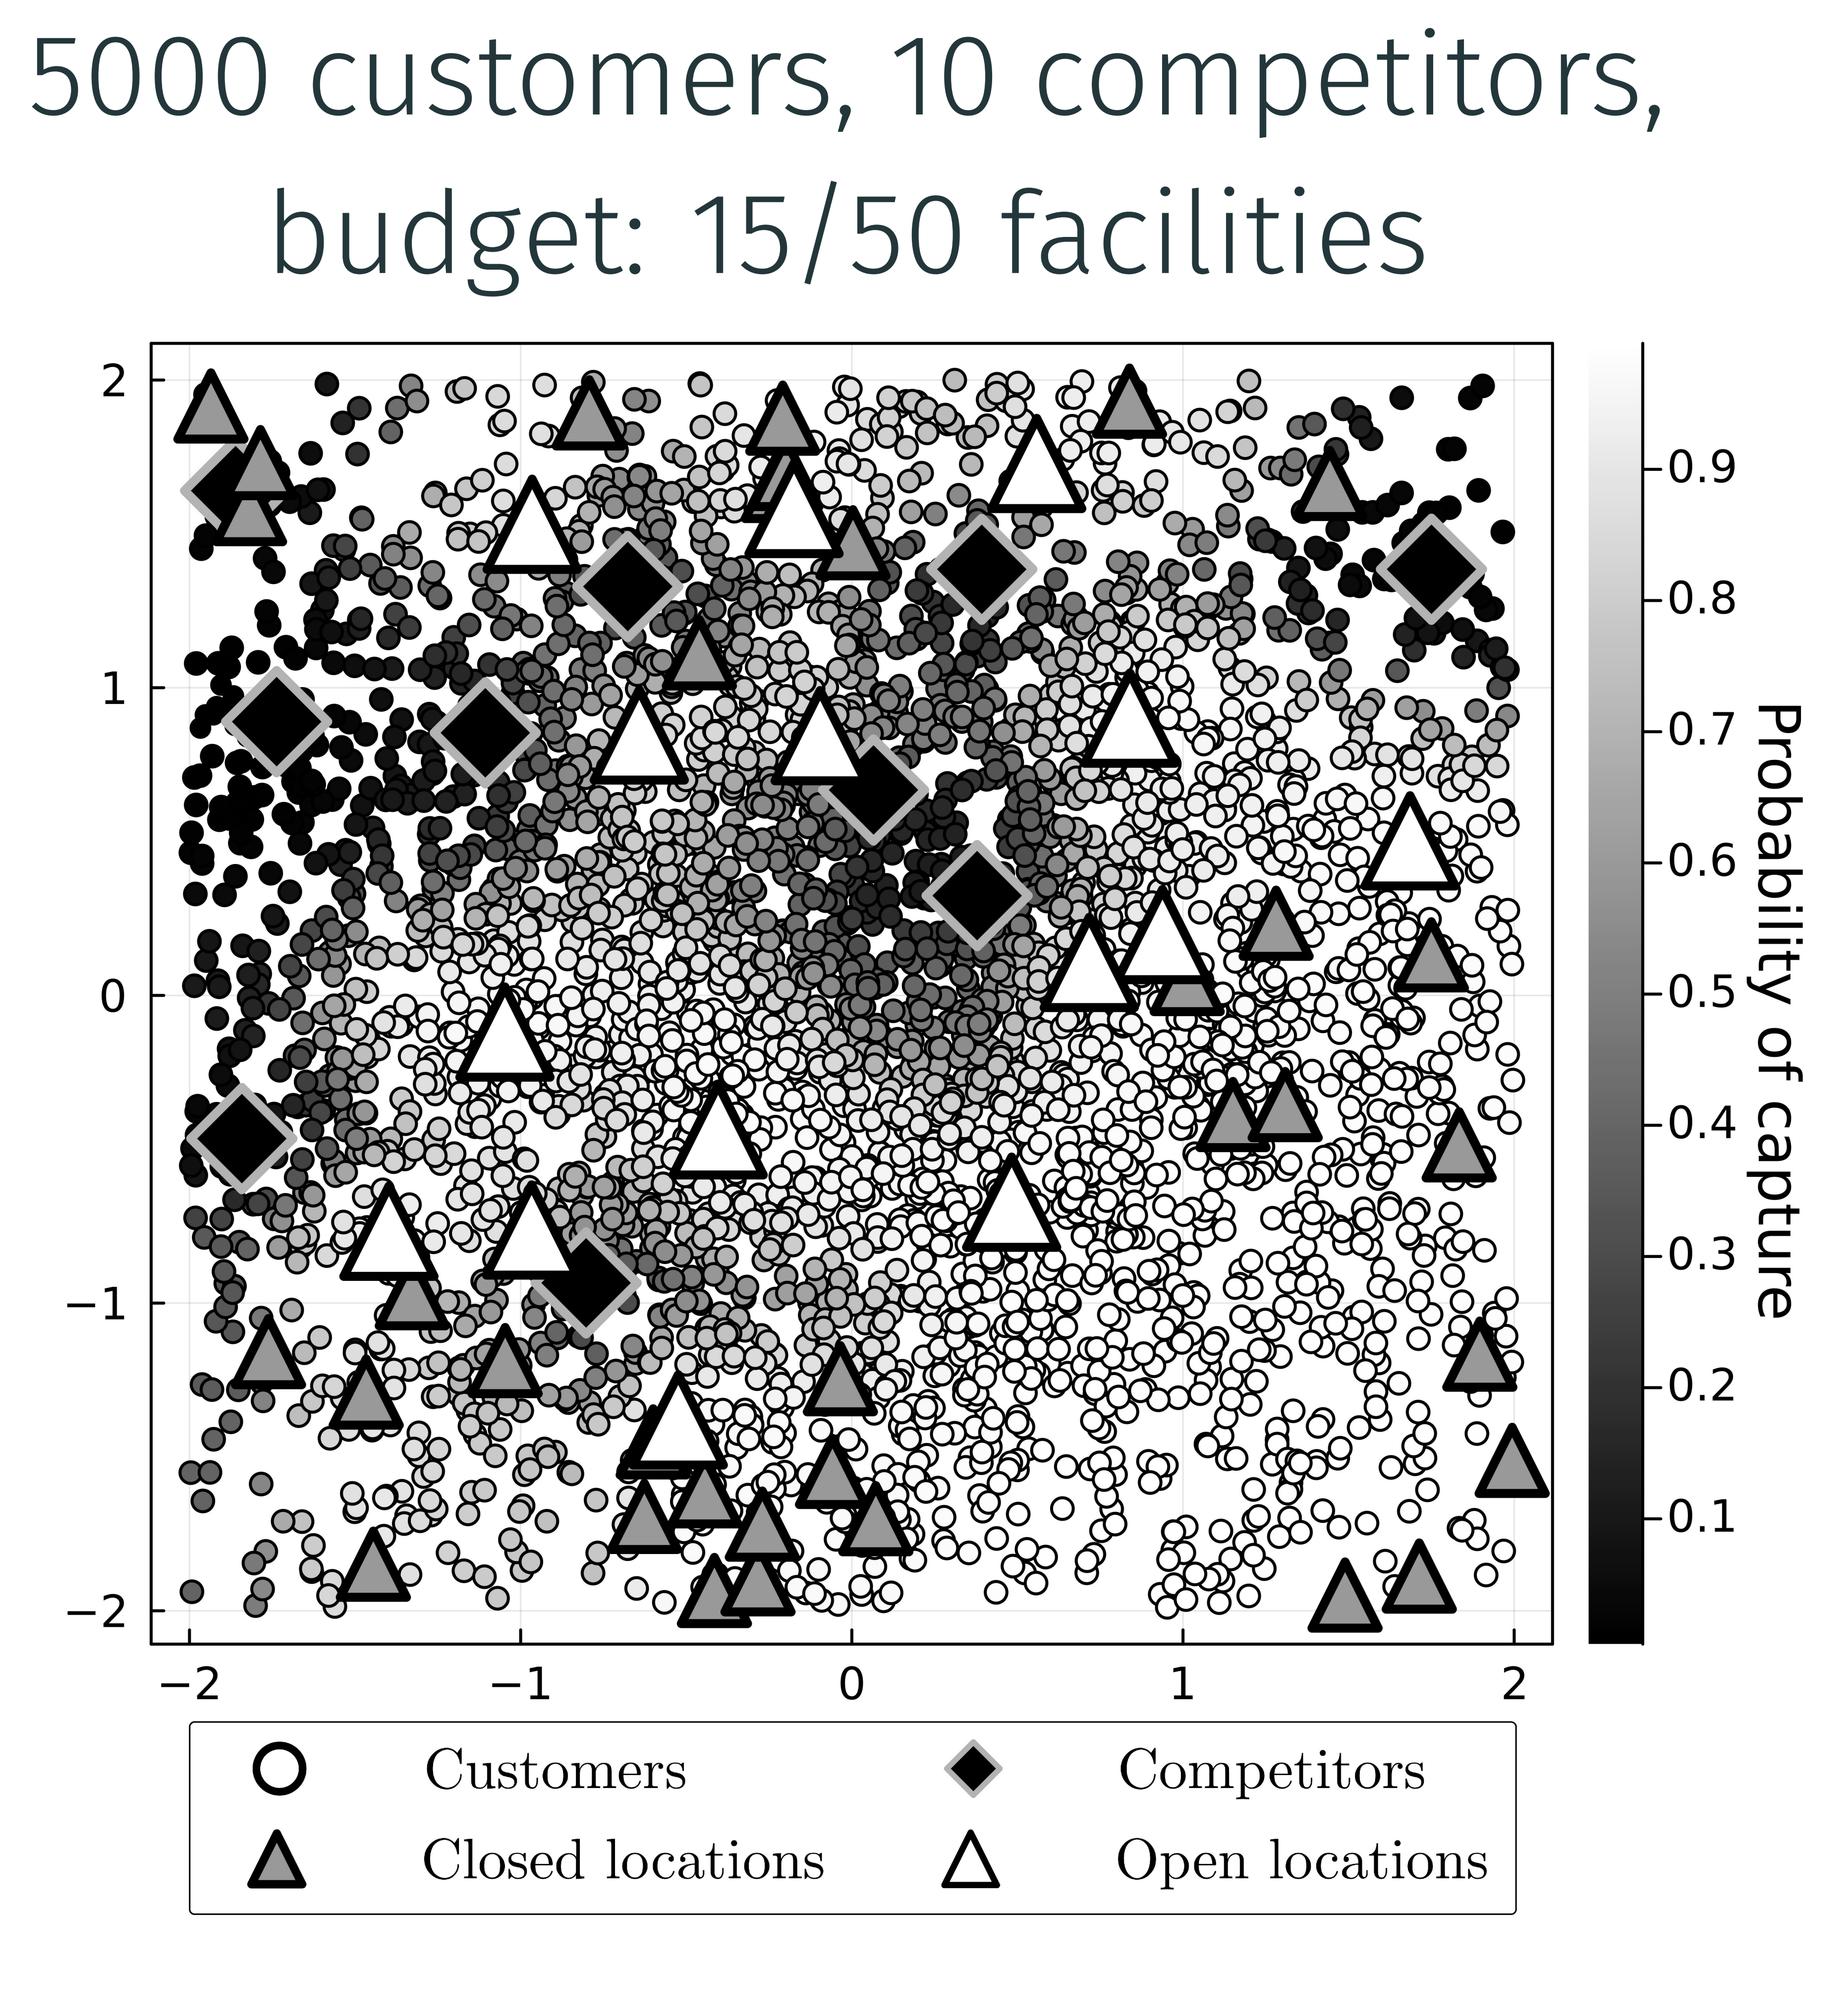

Legault and Frejinger, A Model-free Approach for Solving Choice-based Competitive Facility Location Problems Using Simulation and Submodularity, INFORMS Journal on Computing, 2024.

ILLUSTRATIVE EXAMPLE: ENDOGENOUS UNCERTAINTY CHANGES DUE TO CONTEXT

1

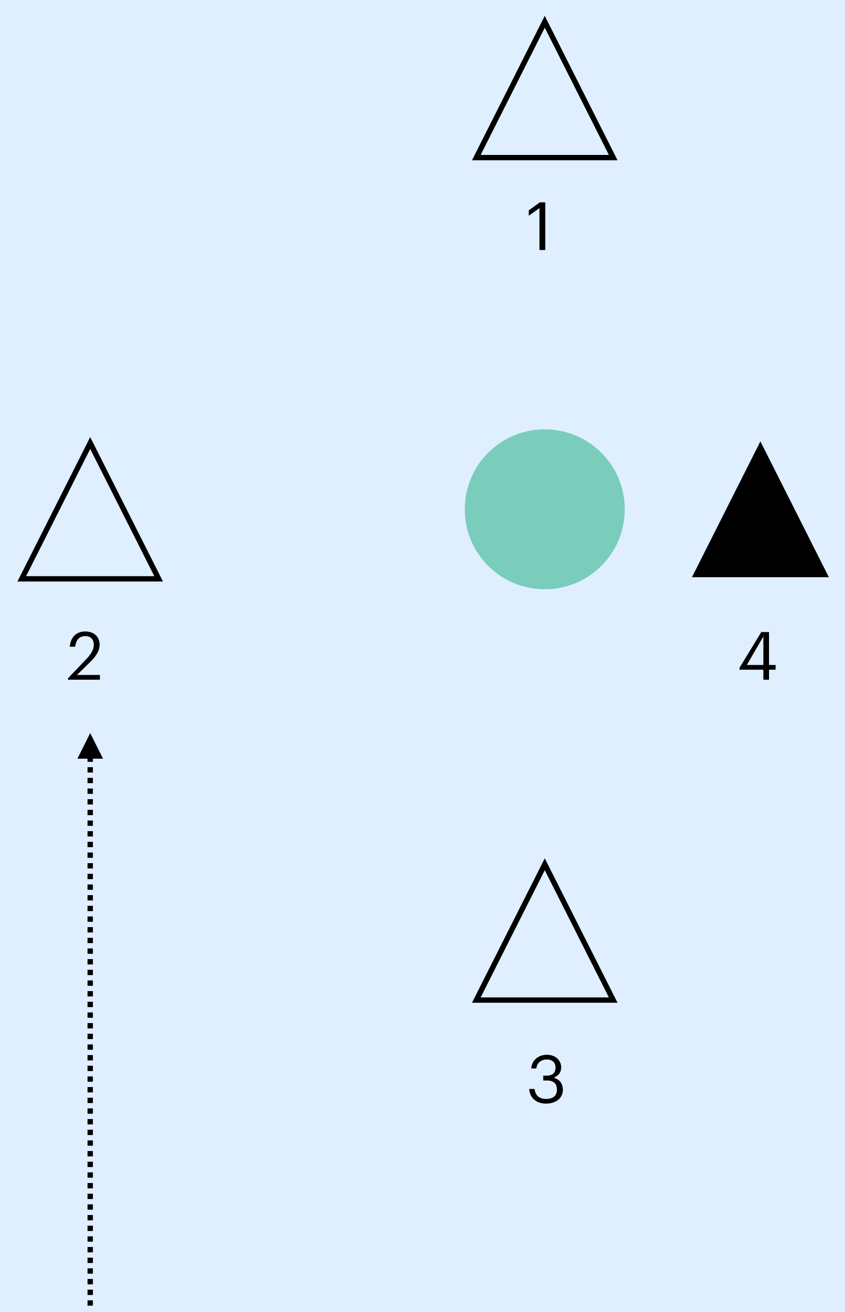

$\boldsymbol{x}$: *presence or not of an attractive feature*

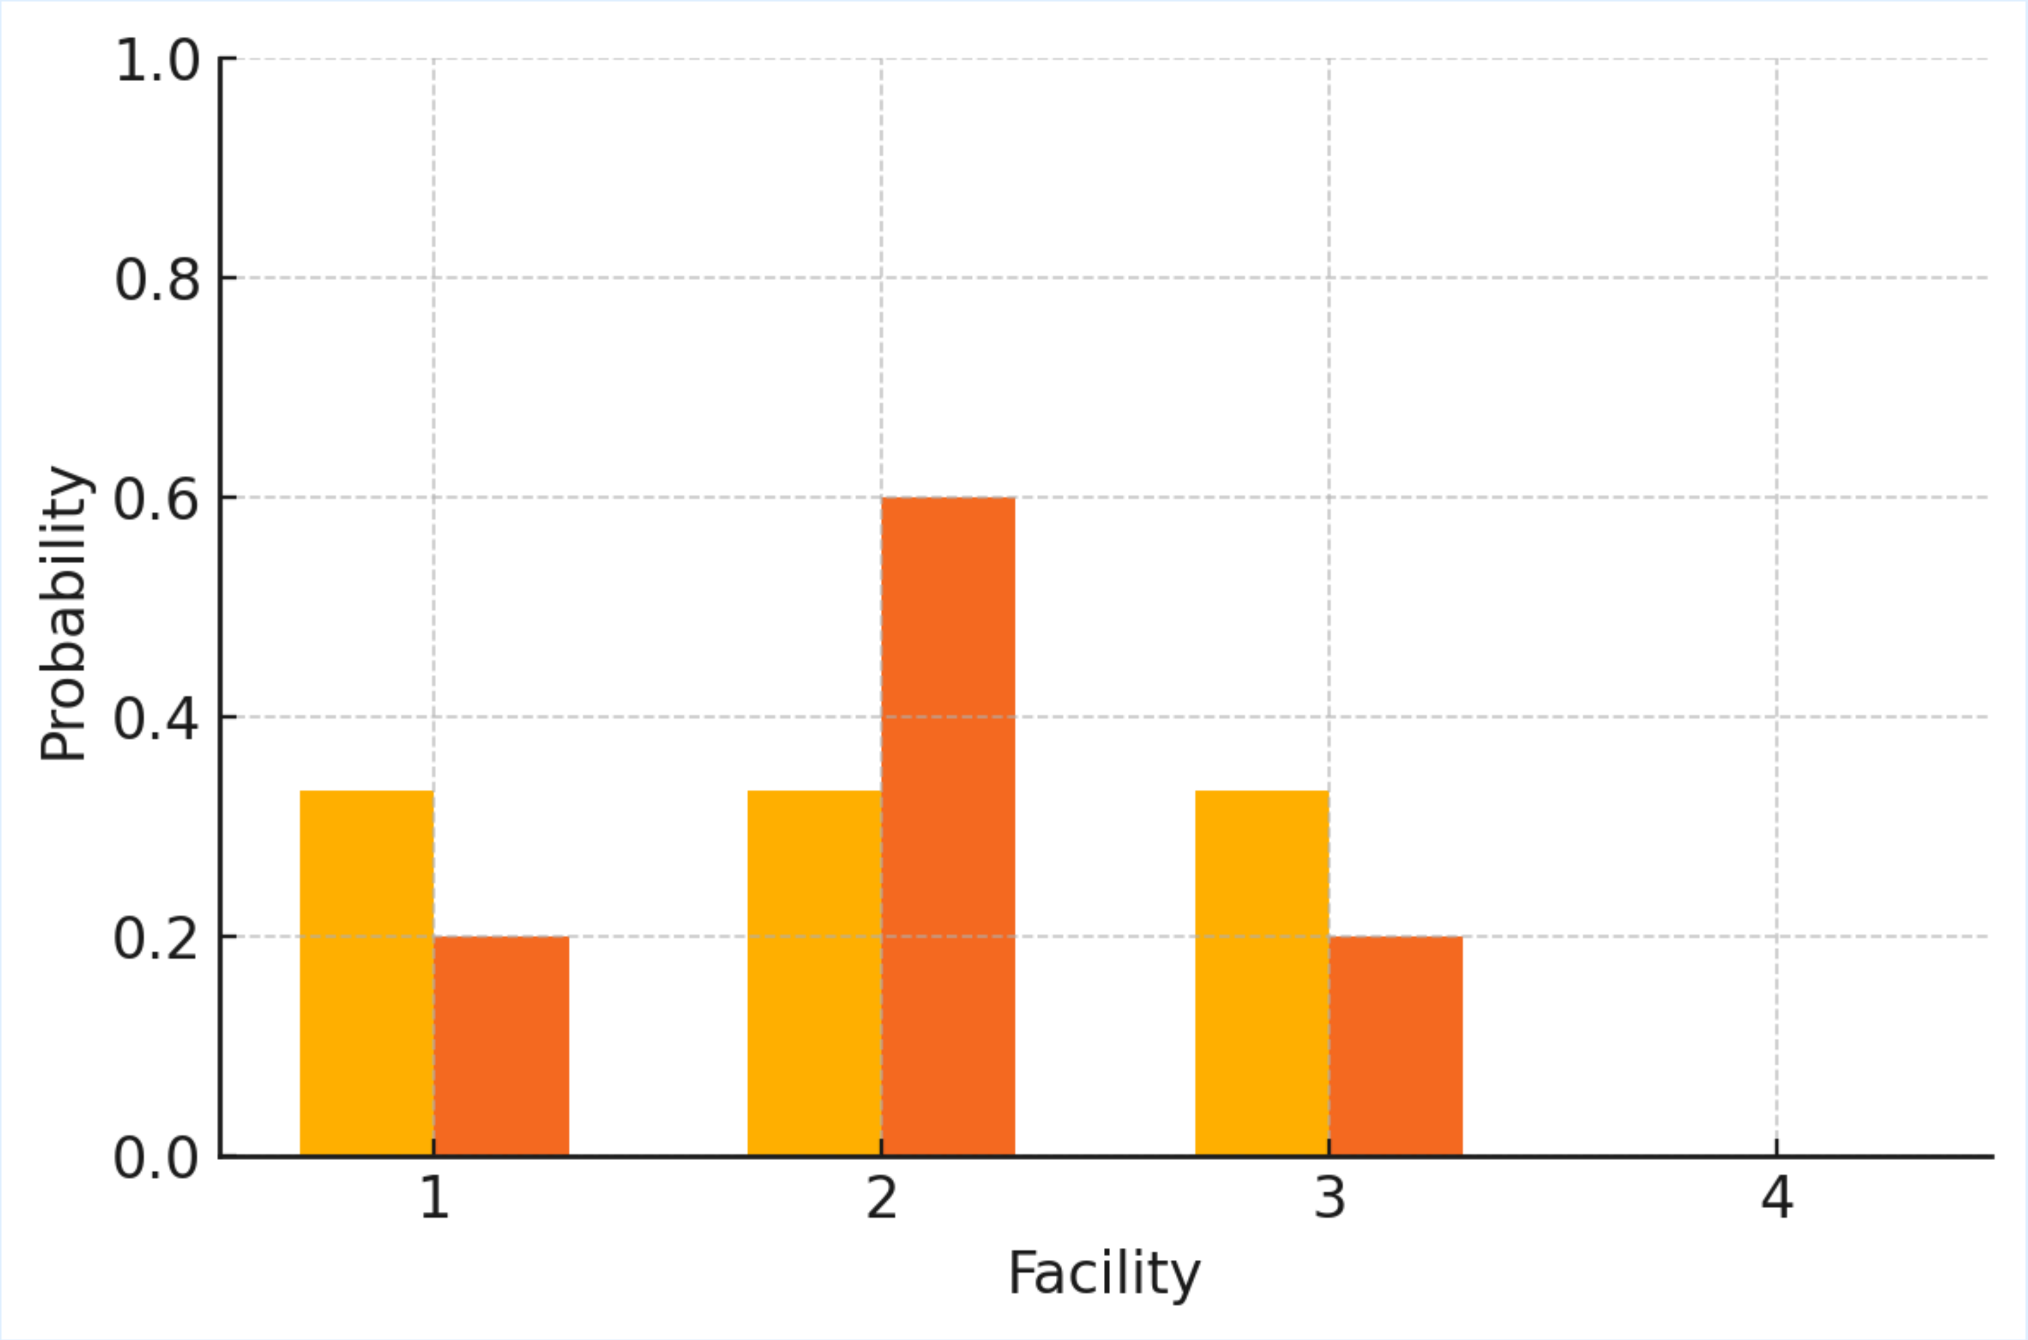

Conditional distribution of attracted demand, i.e., customer represented in green

$\mathbb{P}(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{z})$

Decision $\boldsymbol{z}^1$: open 1, 2 and 3. Close 4.

Uniform distribution (in yellow) over 1, 2 and 3 when 2 does not have an attractive feature

Higher probability for 2 (in orange) when it has an attractive feature

ILLUSTRATIVE EXAMPLE: ENDOGENOUS UNCERTAINTY CHANGES DUE TO DECISIONS

1

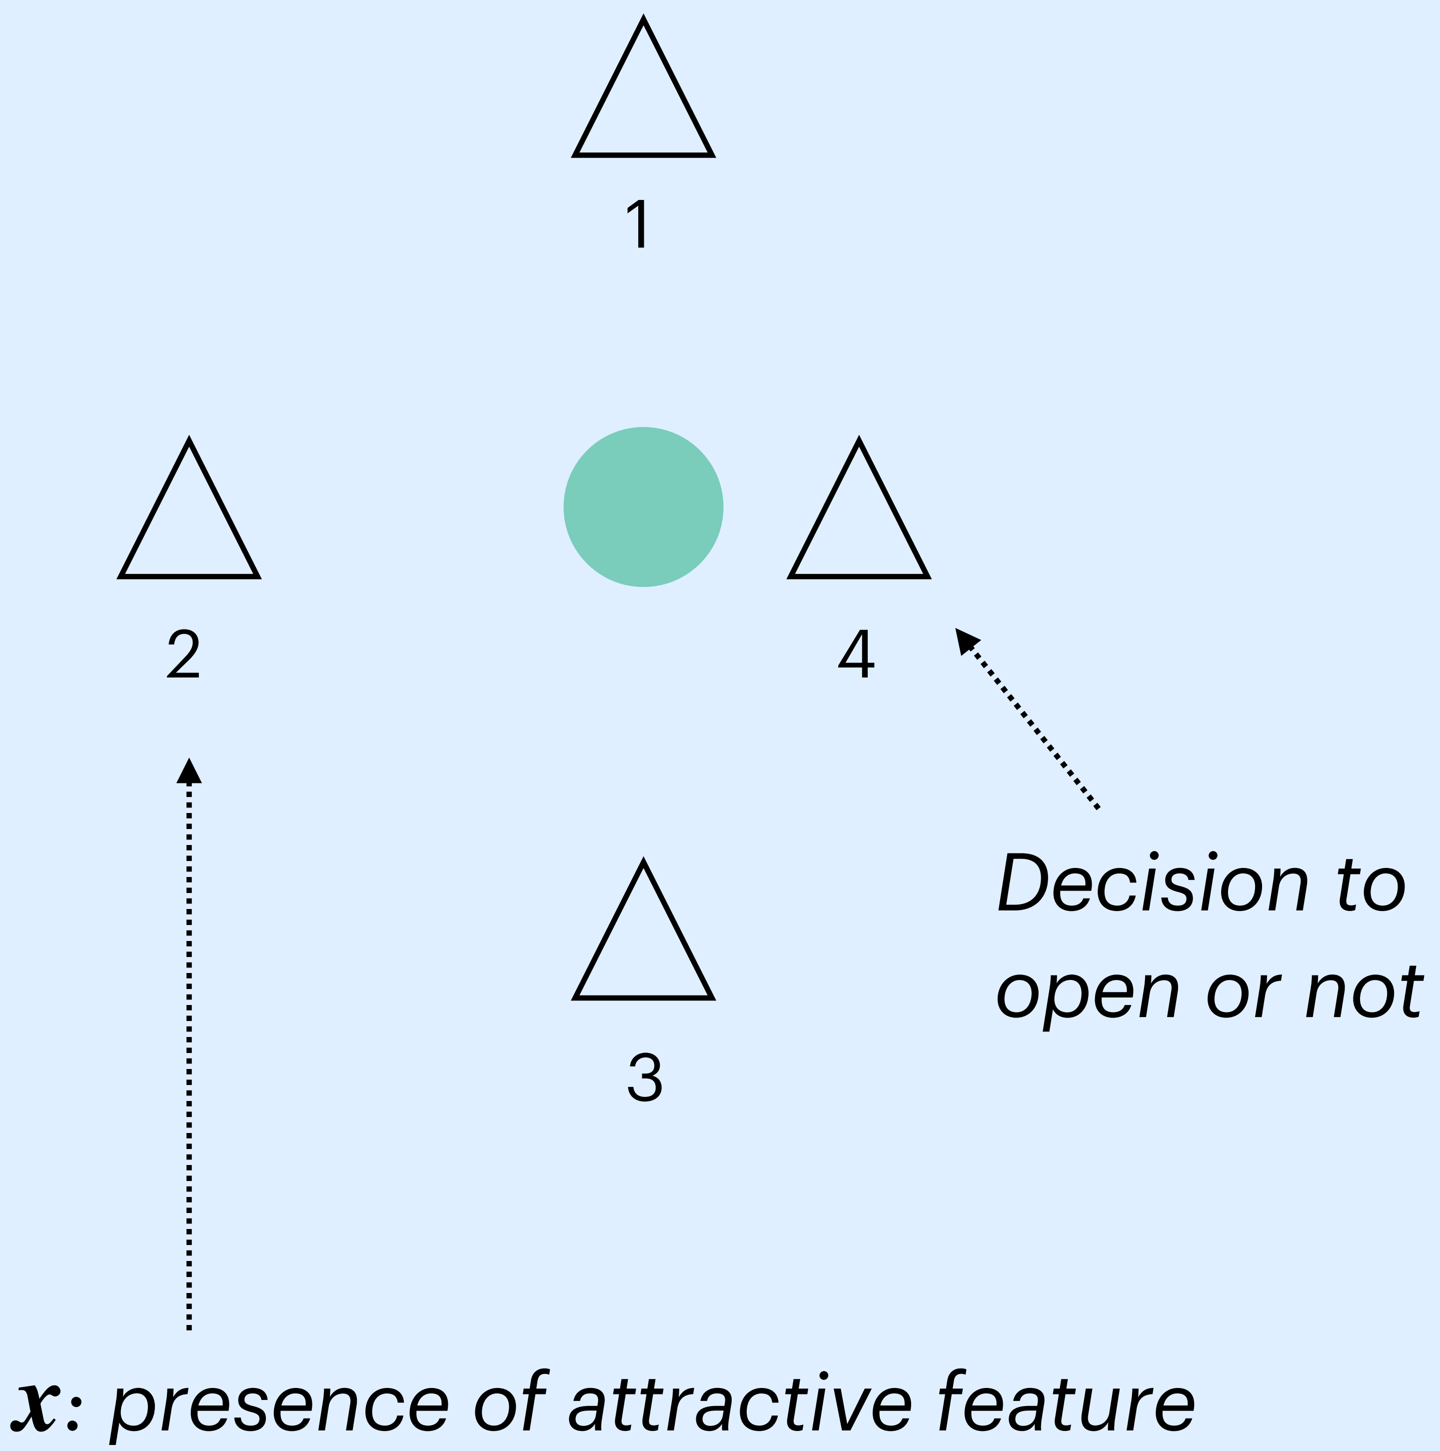

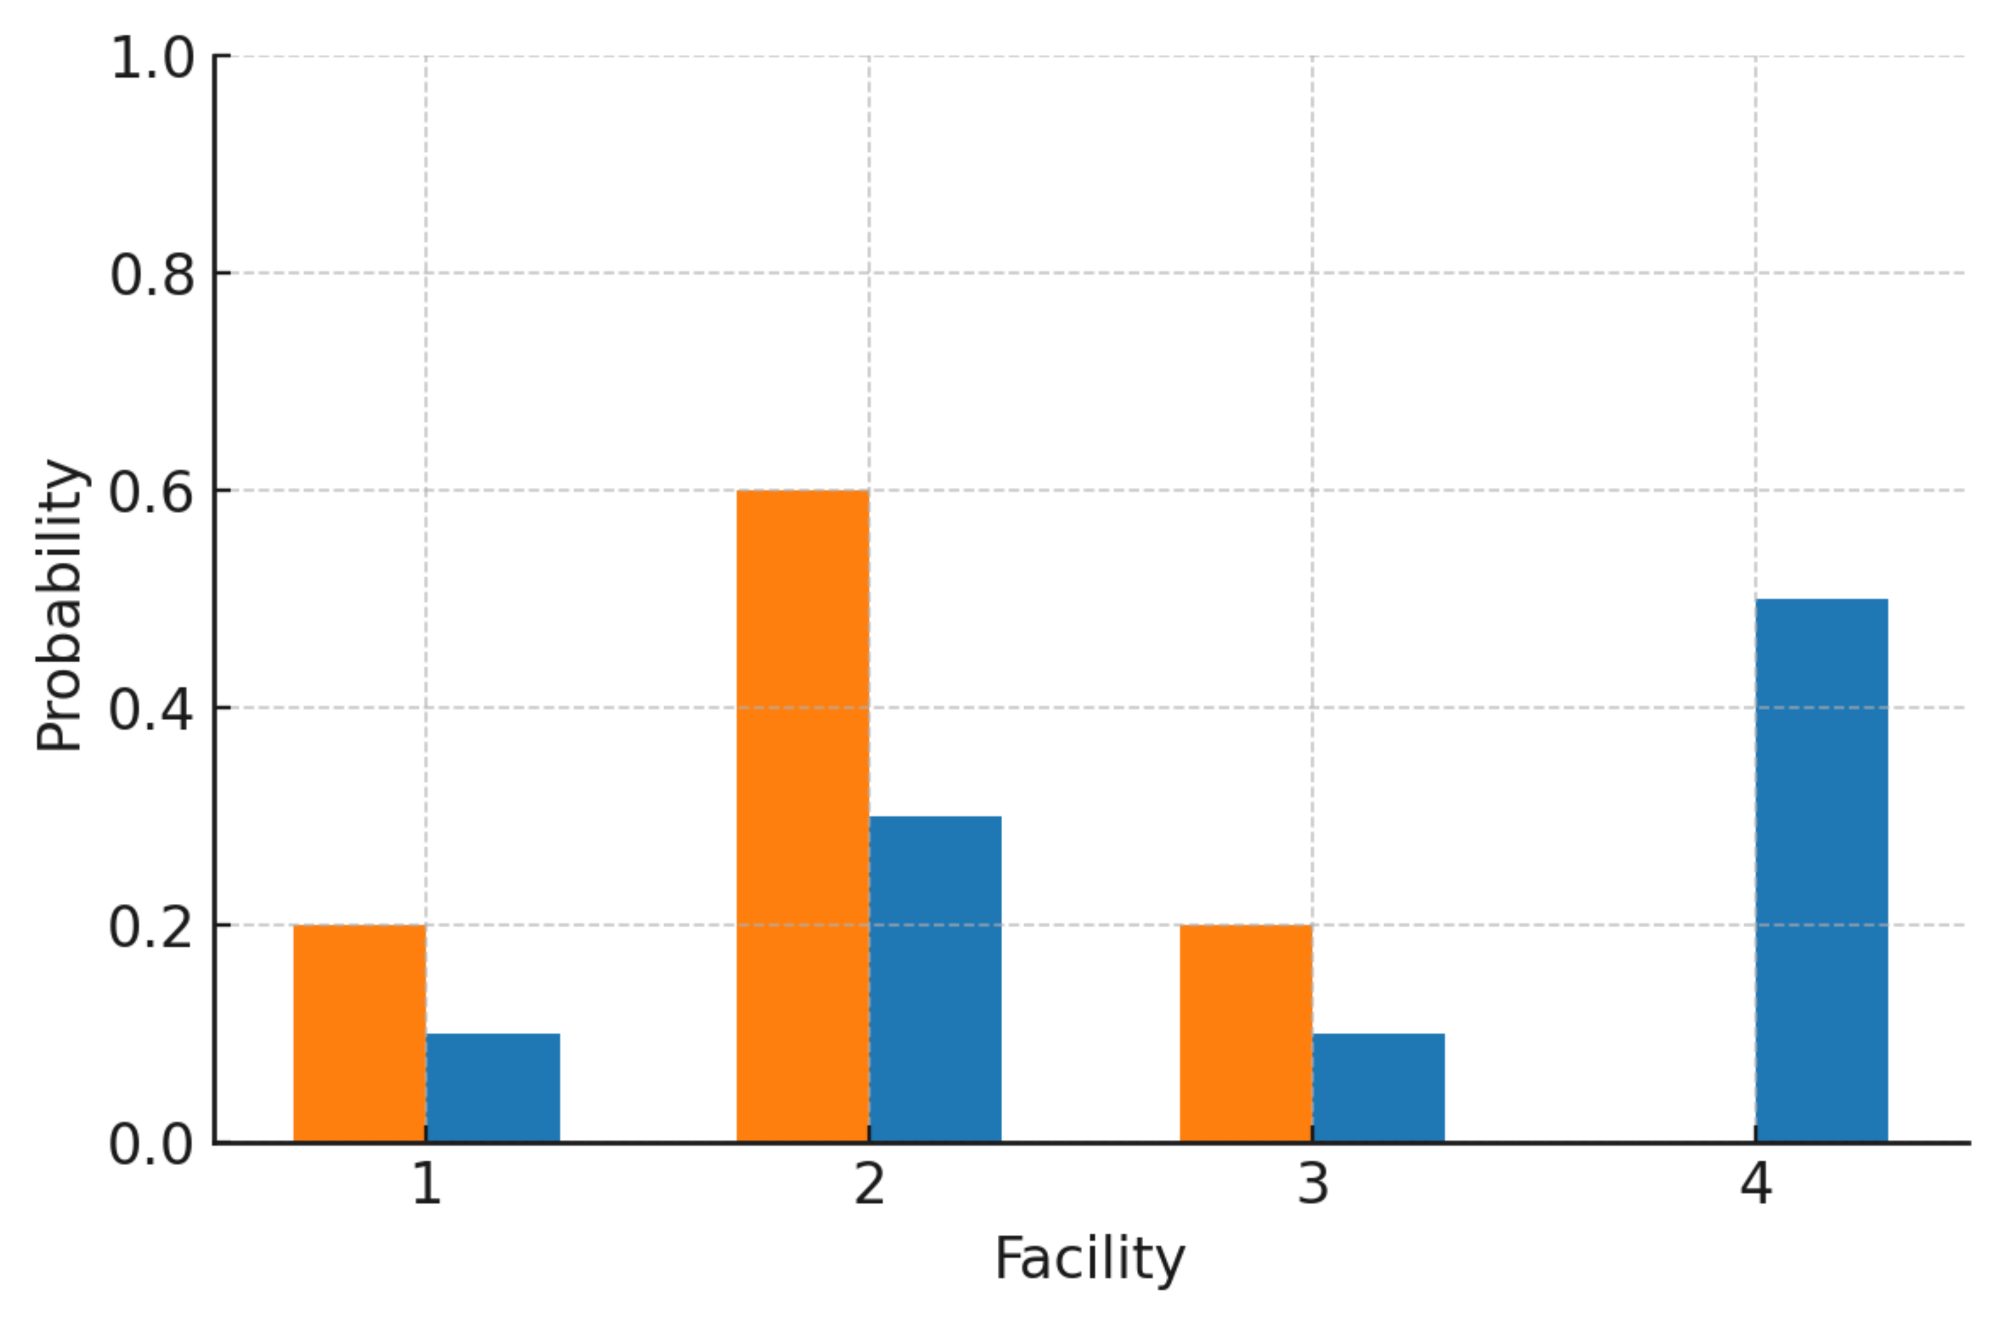

Conditional distribution of attracting demand, i.e., customer represented in green

$\mathbb{P}(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{z})$

Orange distribution same as before: 2 has attractive feature and decision to open 1, 2 and 3 only (decision $\boldsymbol{z}^1$)

In blue a distribution for decision $\boldsymbol{z}^2$ to open all four facilities

Level of uncertainty — entropy — matters for solution quality of sample average approximation and computational performance.

Legault and Frejinger, A Model-free Approach for Solving Choice-based Competitive Facility Location Problems Using Simulation and Submodularity, INFORMS Journal on Computing, 2024.

MOTIVATION: CSO PROBLEMS ARISE IN DIFFERENT APPLICATION DOMAINS

Energy, supply chain management, transportation, mobility, finance, infrastructure planning, maintenance...

Common features: contextual, large-scale, solved repeatedly over time, discrete variables, non-trivial constraints

Our focus: single or two-stage formulations

HUMAN BEHAVIOUR AND ENDOGENOUS UNCERTAINTY

- Often a high degree of uncertainty related to demand
- Exploit endogenous distributions (bilevel formulations)
  - Demand management, e.g., through pricing
  - Supply management, e.g. offer services that are attractive to customers
  - Encourage sustainable behaviour

Lamontagne, Carvalho, Frejinger, Gendron, Anjos, Atallah, Optimizing electric vehicle charging station placement using advanced discrete choice models, INFORMS Journal on Computing 35(5):1195-1213, 2023.

Pinzon Ulloa, Frejinger, Gendron, A logistics provider's profit maximization facility location problem with random utility maximizing followers, Computers & OR 167, 2024.


Climbing to the moon, Chris Rowney www.devianart.com

My moonshot: From data to real-world impact

Solve large-scale discrete optimization problems subject to uncertainty (exogenous or endogenous) while leveraging recent relevant data

Outline

1. Contextual optimization lens: Foundational concepts and motivation
2. From data to decision
   - Two decision pipelines: from context to decision
   - Three training pipelines: how to use data to improve prescriptive performance
3. From data to real-world impact
   - What do we know and what topics require more research

AIM: FIND AN OPTIMAL POLICY FOR ANY $\boldsymbol{x} \in X$

$$\pi^* \in \operatorname*{argmin}_{\pi \in \Pi} \mathbb{E}_{\mathbb{P}}\left[c(\pi(\boldsymbol{x}), \boldsymbol{y})\right]$$

- Class of feasible policies
- Joint distribution of features and uncertain parameters

Optimal policy, lowest long-term expected cost

$$z^*(\boldsymbol{x}) \in \operatorname*{argmin}_{z \in \mathcal{Z}} \mathbb{E}_{\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})}\left[c\left(\boldsymbol{z}, \boldsymbol{y}\right)\right]$$

- Deterministic feasible set
- **Exogenous uncertainty**

Solving the CSO problem for any $\boldsymbol{x}$ naturally identifies an optimal policy

(Theorem 14.60, Rockafellar and Wets, Variational Analysis, Vol. 317, 2009)

The talk follows our survey: **Utsav Sadana**, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

FROM CONTEXT TO DECISIONS: TWO DECISION PIPELINES

Decision rule

Context $\boldsymbol{x}$ → Decision rule → Decision $\pi(\boldsymbol{x})$

Predict/estimate, then optimize

Context $\boldsymbol{x}$ → Prediction model → Optimization model & solver → Decision $z^*(\boldsymbol{x})$

We need information about the uncertainty to optimize, but the distributions are unknown to us.

How to use data $\mathcal{D}_N = \{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^N$ to find a policy that approximates well the optimal policy?

Sadana, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

MOTIVATING EXAMPLE (PREDICT, THEN OPTIMIZE DECISION PIPELINE)
AIM OF LEARNING PIPELINE: MAXIMIZE PRESCRIPTIVE PERFORMANCE

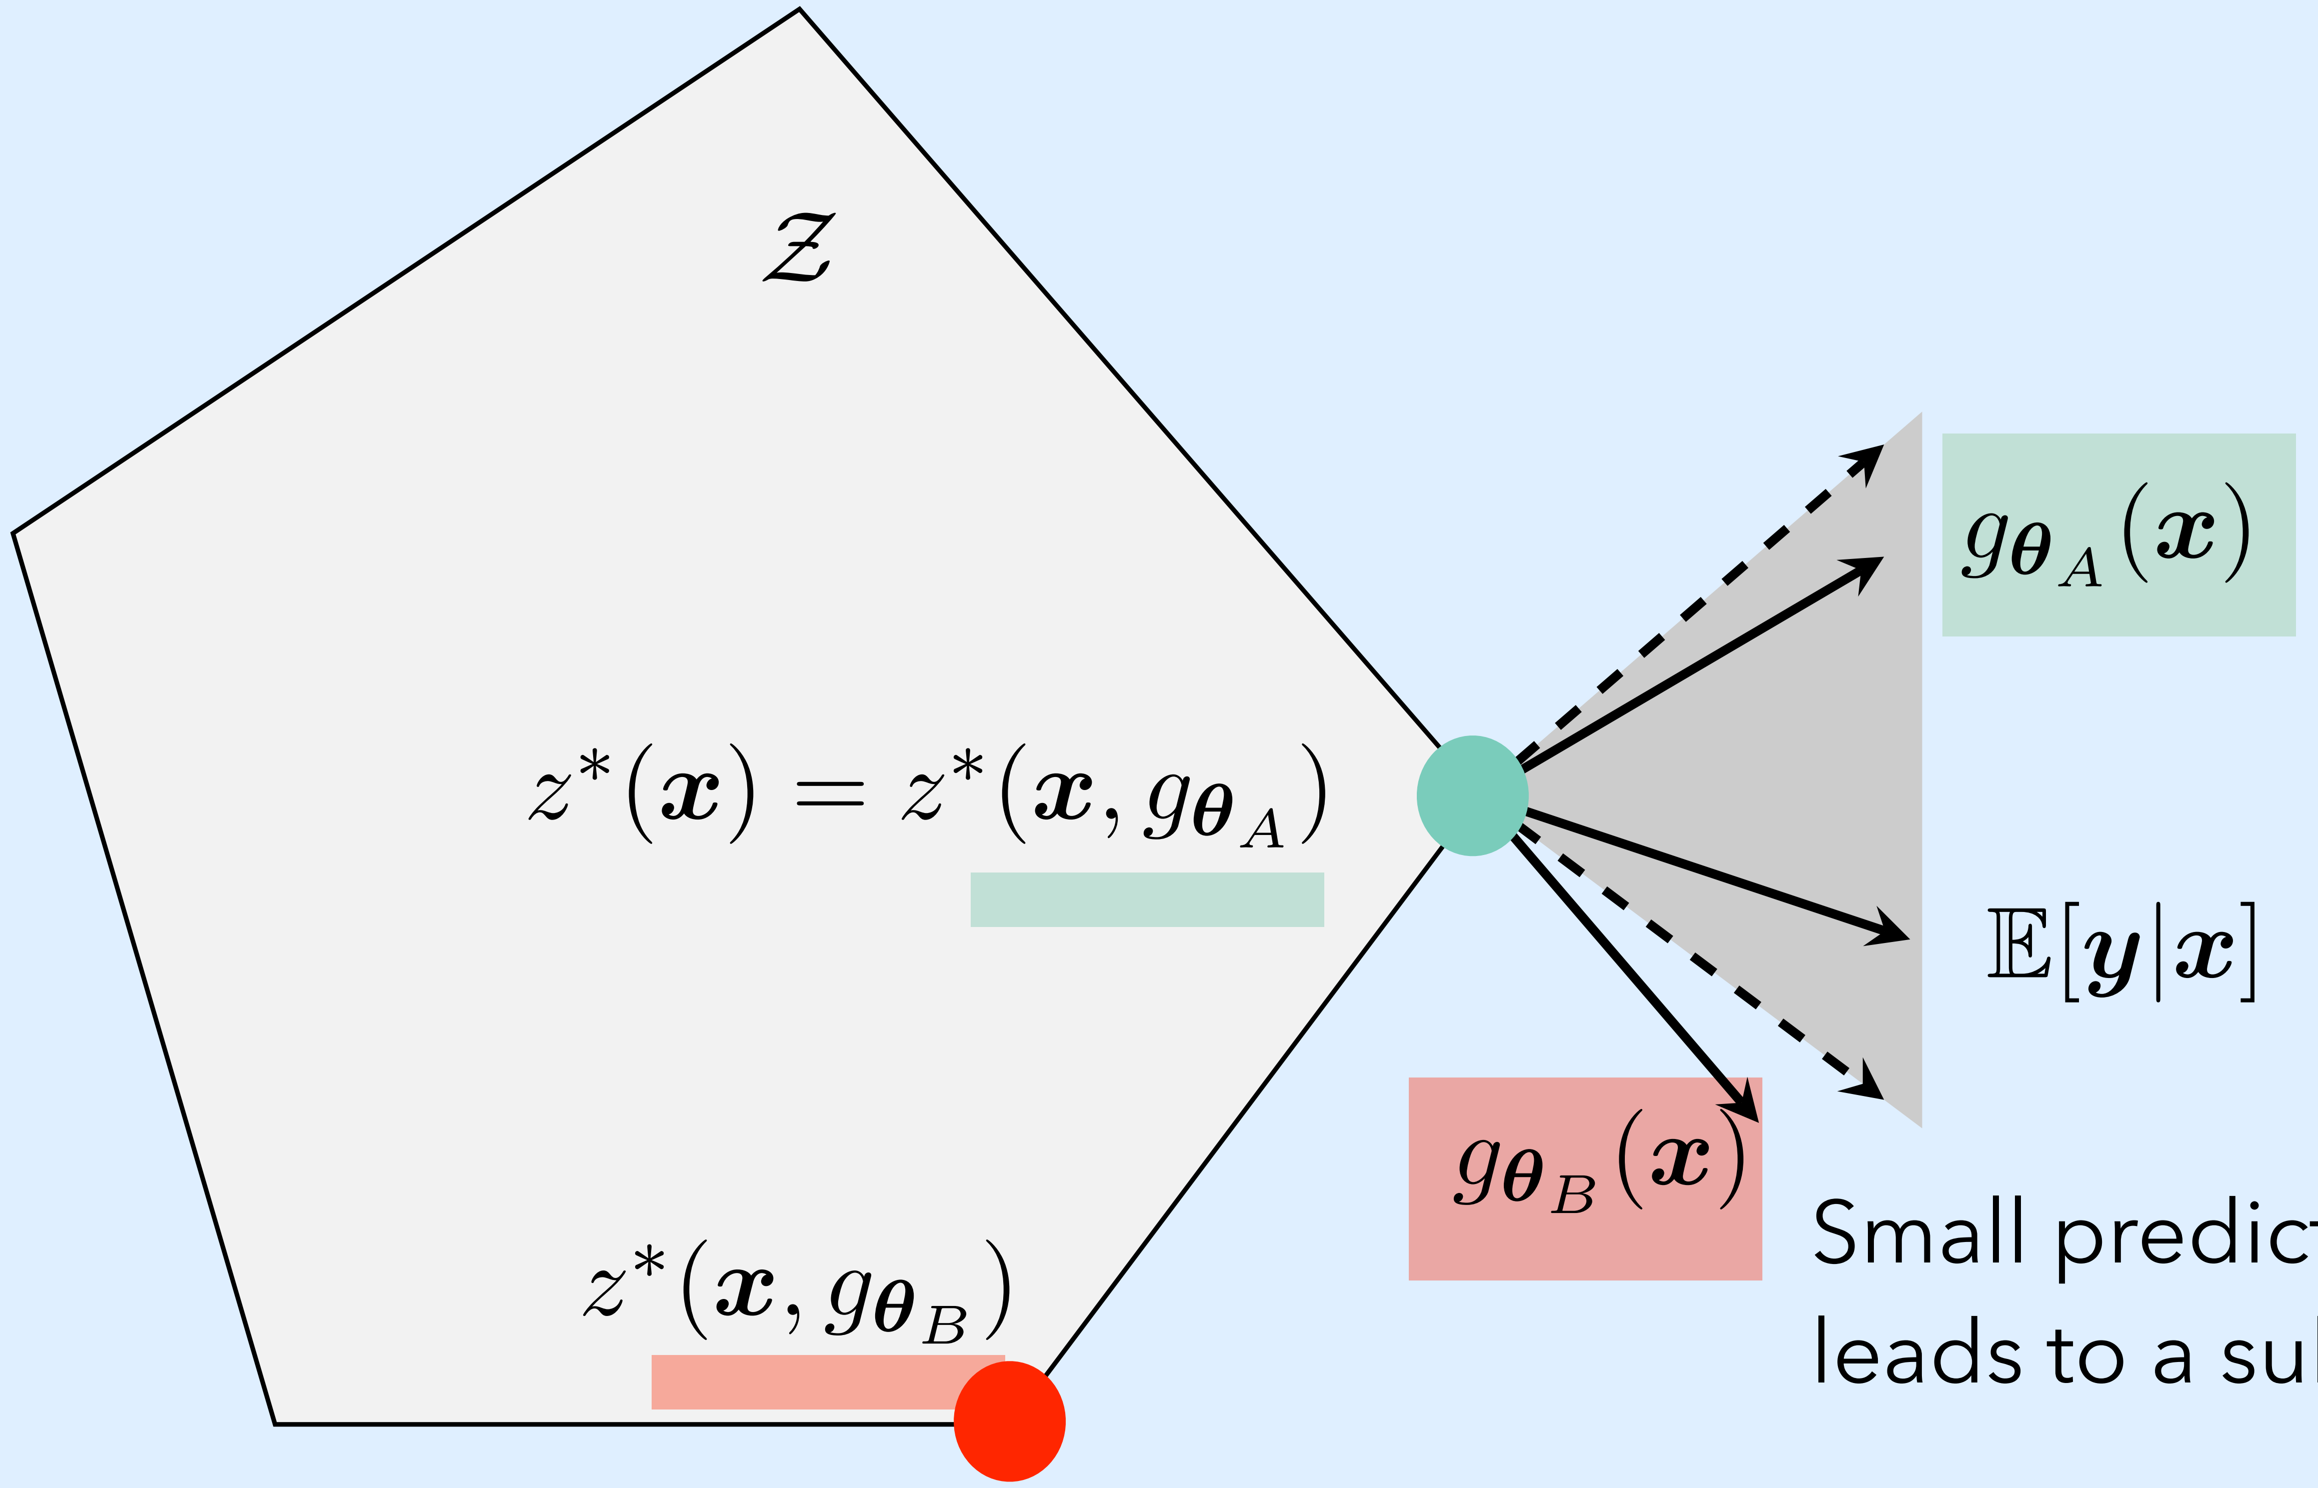

Prediction with error, but leading to the optimal solution

True cost $c(\boldsymbol{x}, \boldsymbol{y}) = -\mathbb{E}[\boldsymbol{y}|\boldsymbol{x}]^\top \boldsymbol{z}$

Small prediction error that leads to a suboptimal solution

Optimal decision

Suboptimal decision

Potential value: Improved prescriptive performance at the expense of higher prediction error.

Example adapted from: Elmachtoub and Grigas, Smart “Predict, then optimize”, Management Science 68(1):9-26, 2022.

INTEGRATED LEARNING AND OPTIMIZATION (ILO) TRAINING PIPELINE

Task loss: Train predictive component to minimize downstream cost incurred by the decision (bilevel problem)

Machine learning model $f_\theta(\boldsymbol{x})$ approximates $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})$

Training pipeline for a given training example $(\boldsymbol{x}_i, \boldsymbol{y}_i)$:

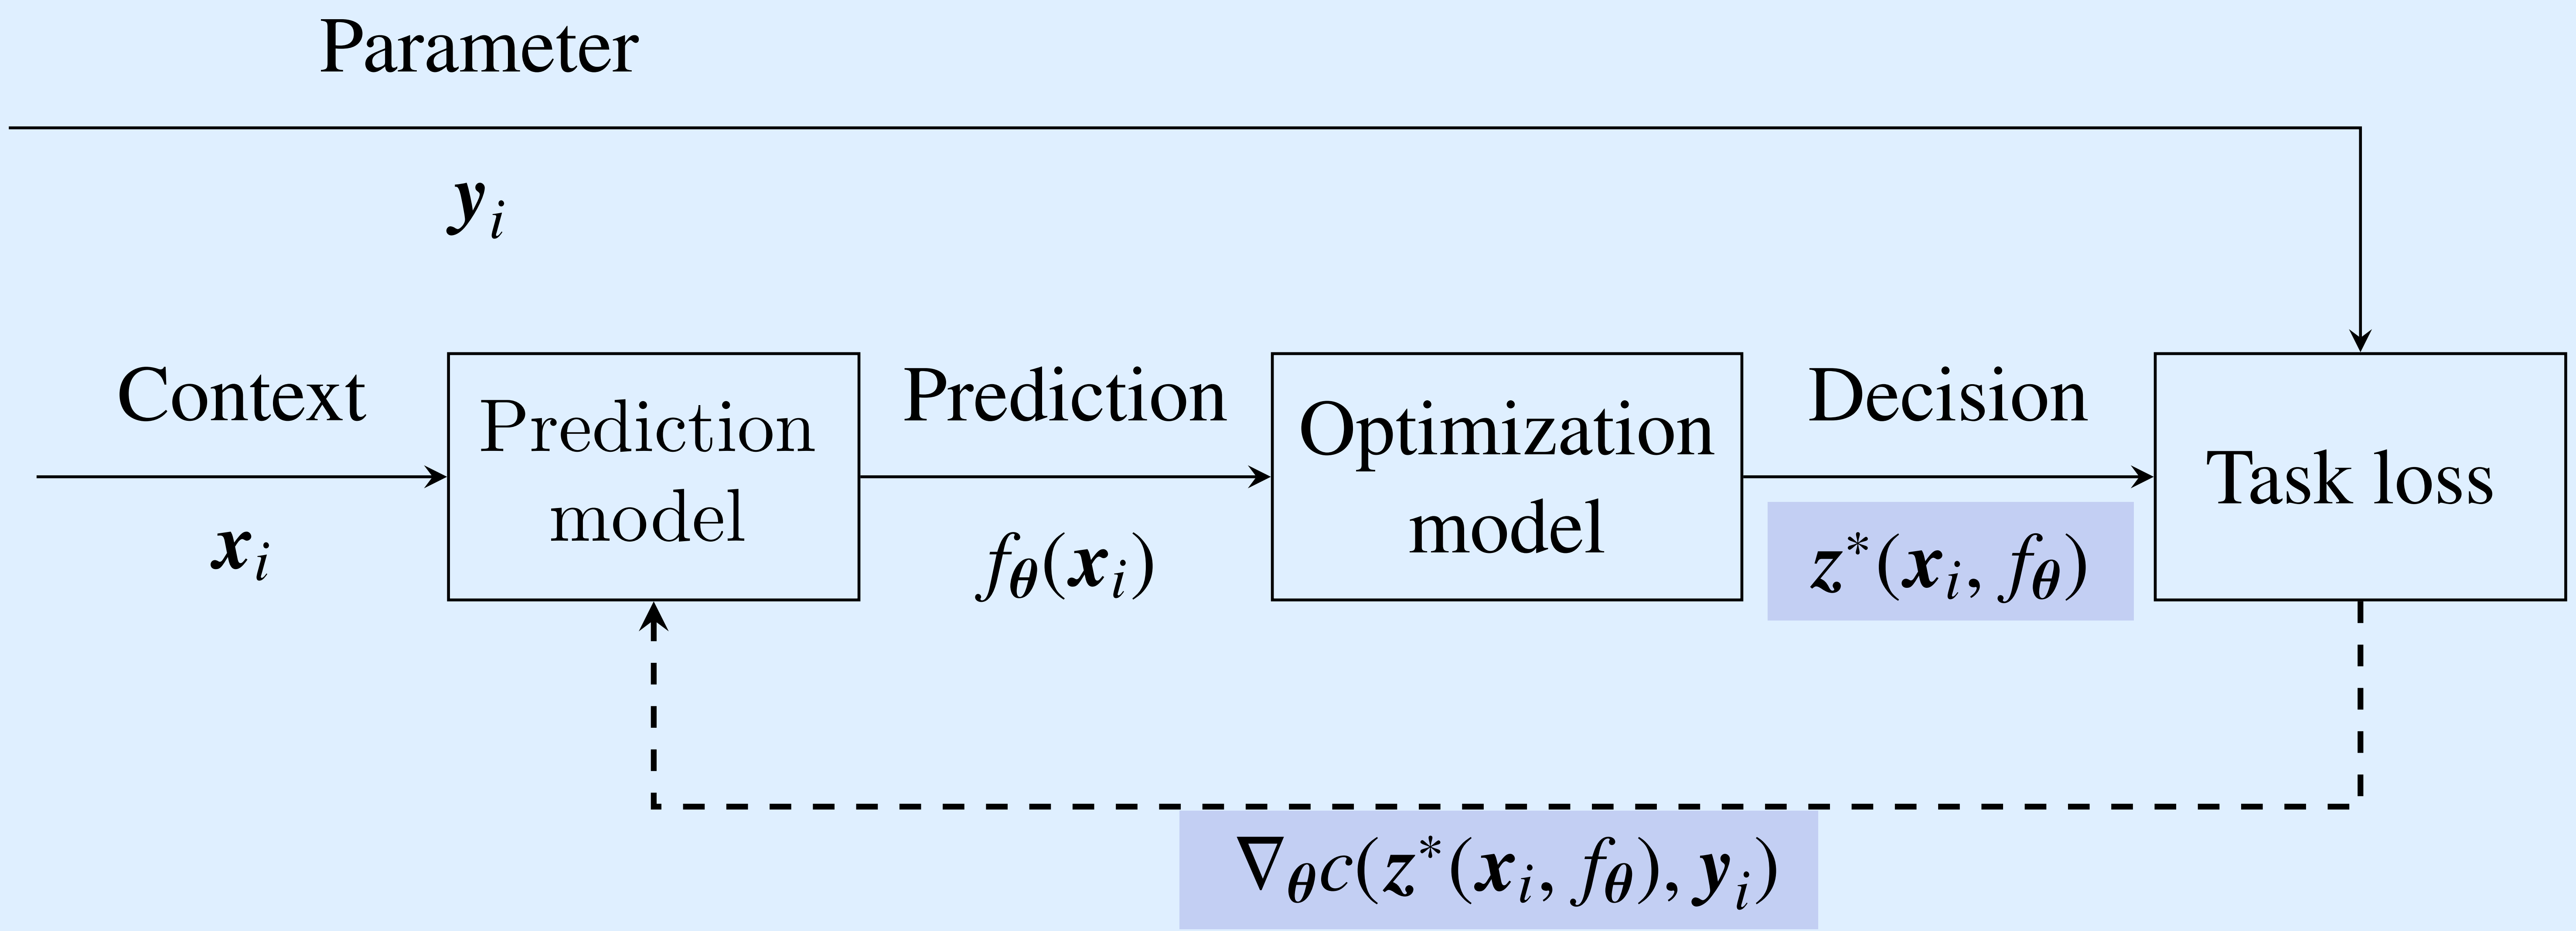

Policy induced by ILO training pipeline minimizes downstream expected cost over empirical distribution.

Computationally costly! Optimization problem solved for each data point and differentiation through argmin.

Sadana, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

THREE ALTERNATIVE TRAINING PIPELINES FOR CSO

Decision rule

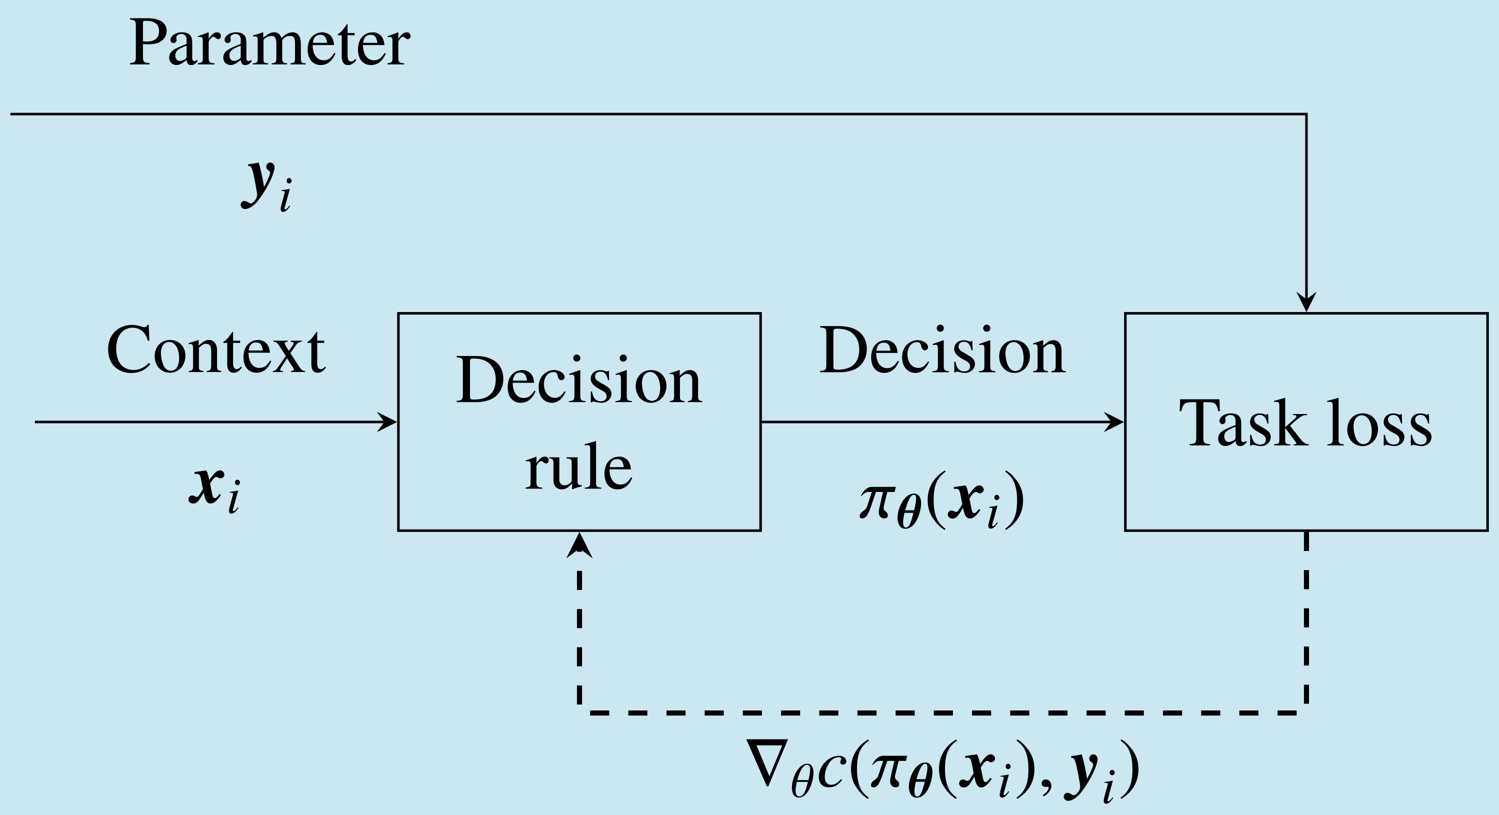

ILO

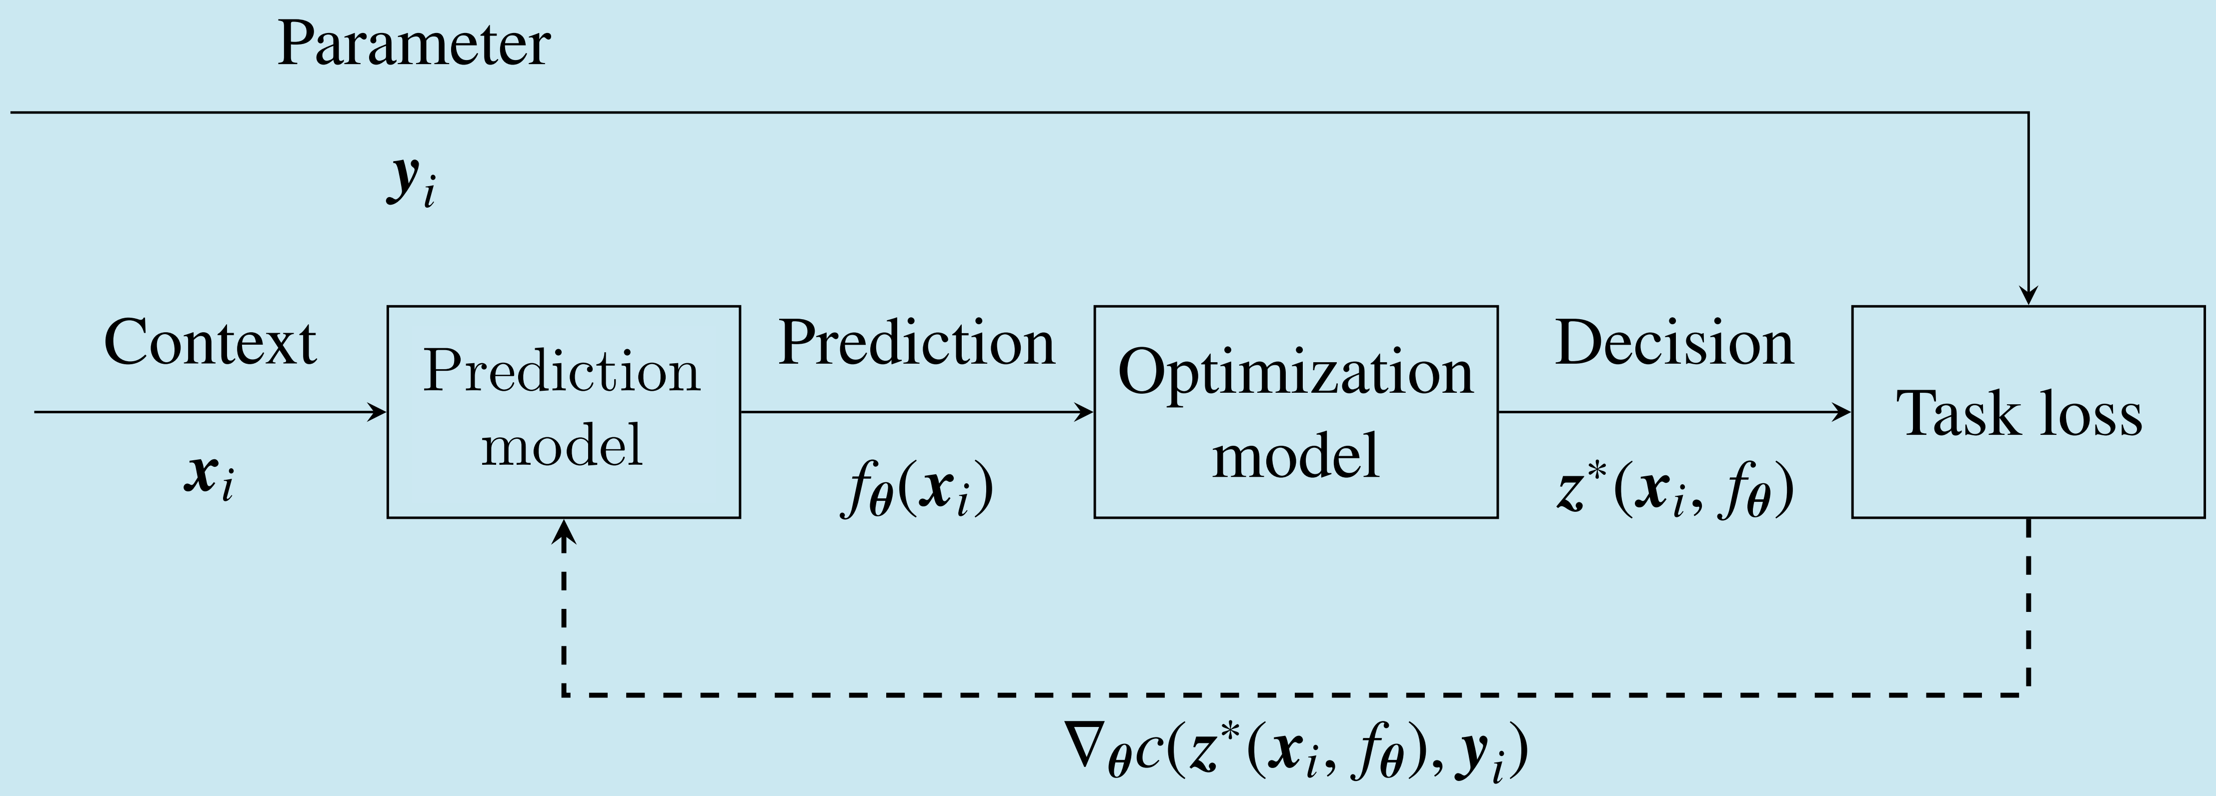

Sequential learning and optimization (SLO)

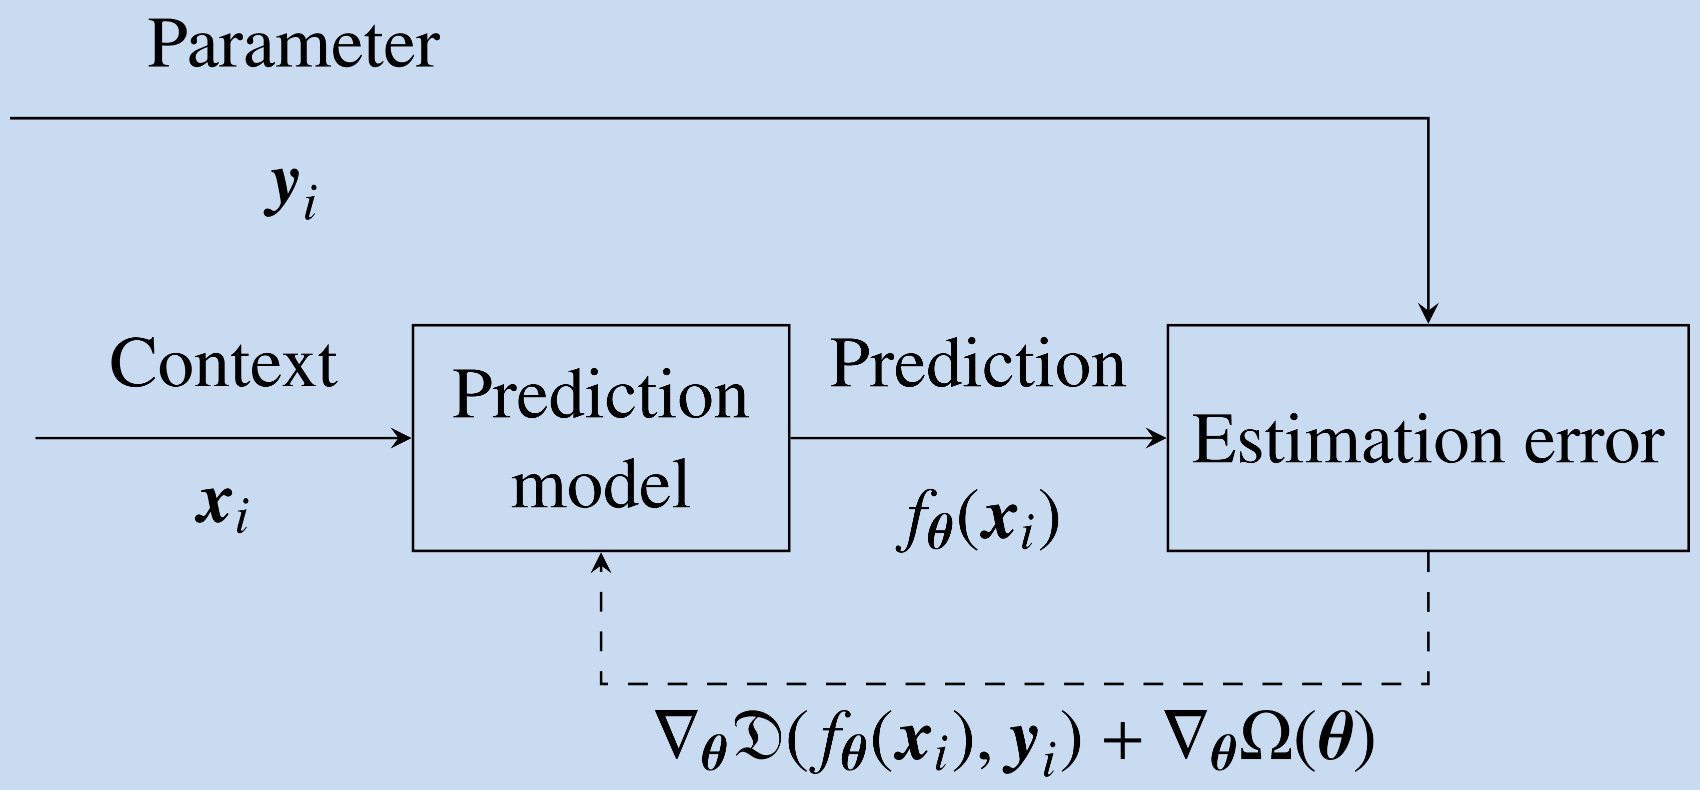

"Classical" training for *predictive* performance

E.g. negative log-likelihood as divergence function

Sadana, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

CHOICE DEPENDS ON REQUIREMENTS AND DATA AVAILABILITY

Aim: train algorithm for prescriptive performance

Requirements at deployment/ inference

Decision pipeline

Data characteristics

Training pipeline

CHOICE DEPENDS ON REQUIREMENTS AND DATA AVAILABILITY

Aim: train algorithm for prescriptive performance

Requirements at deployment/ inference

Feasibility guarantee

May be difficult to guarantee feasibility
Very short computing time budget

Decision pipeline

Predict/estimate, then optimize

Use to speed up algorithms

Decision rule

Data characteristics

Training pipeline

CHOICE DEPENDS ON REQUIREMENTS AND DATA AVAILABILITY

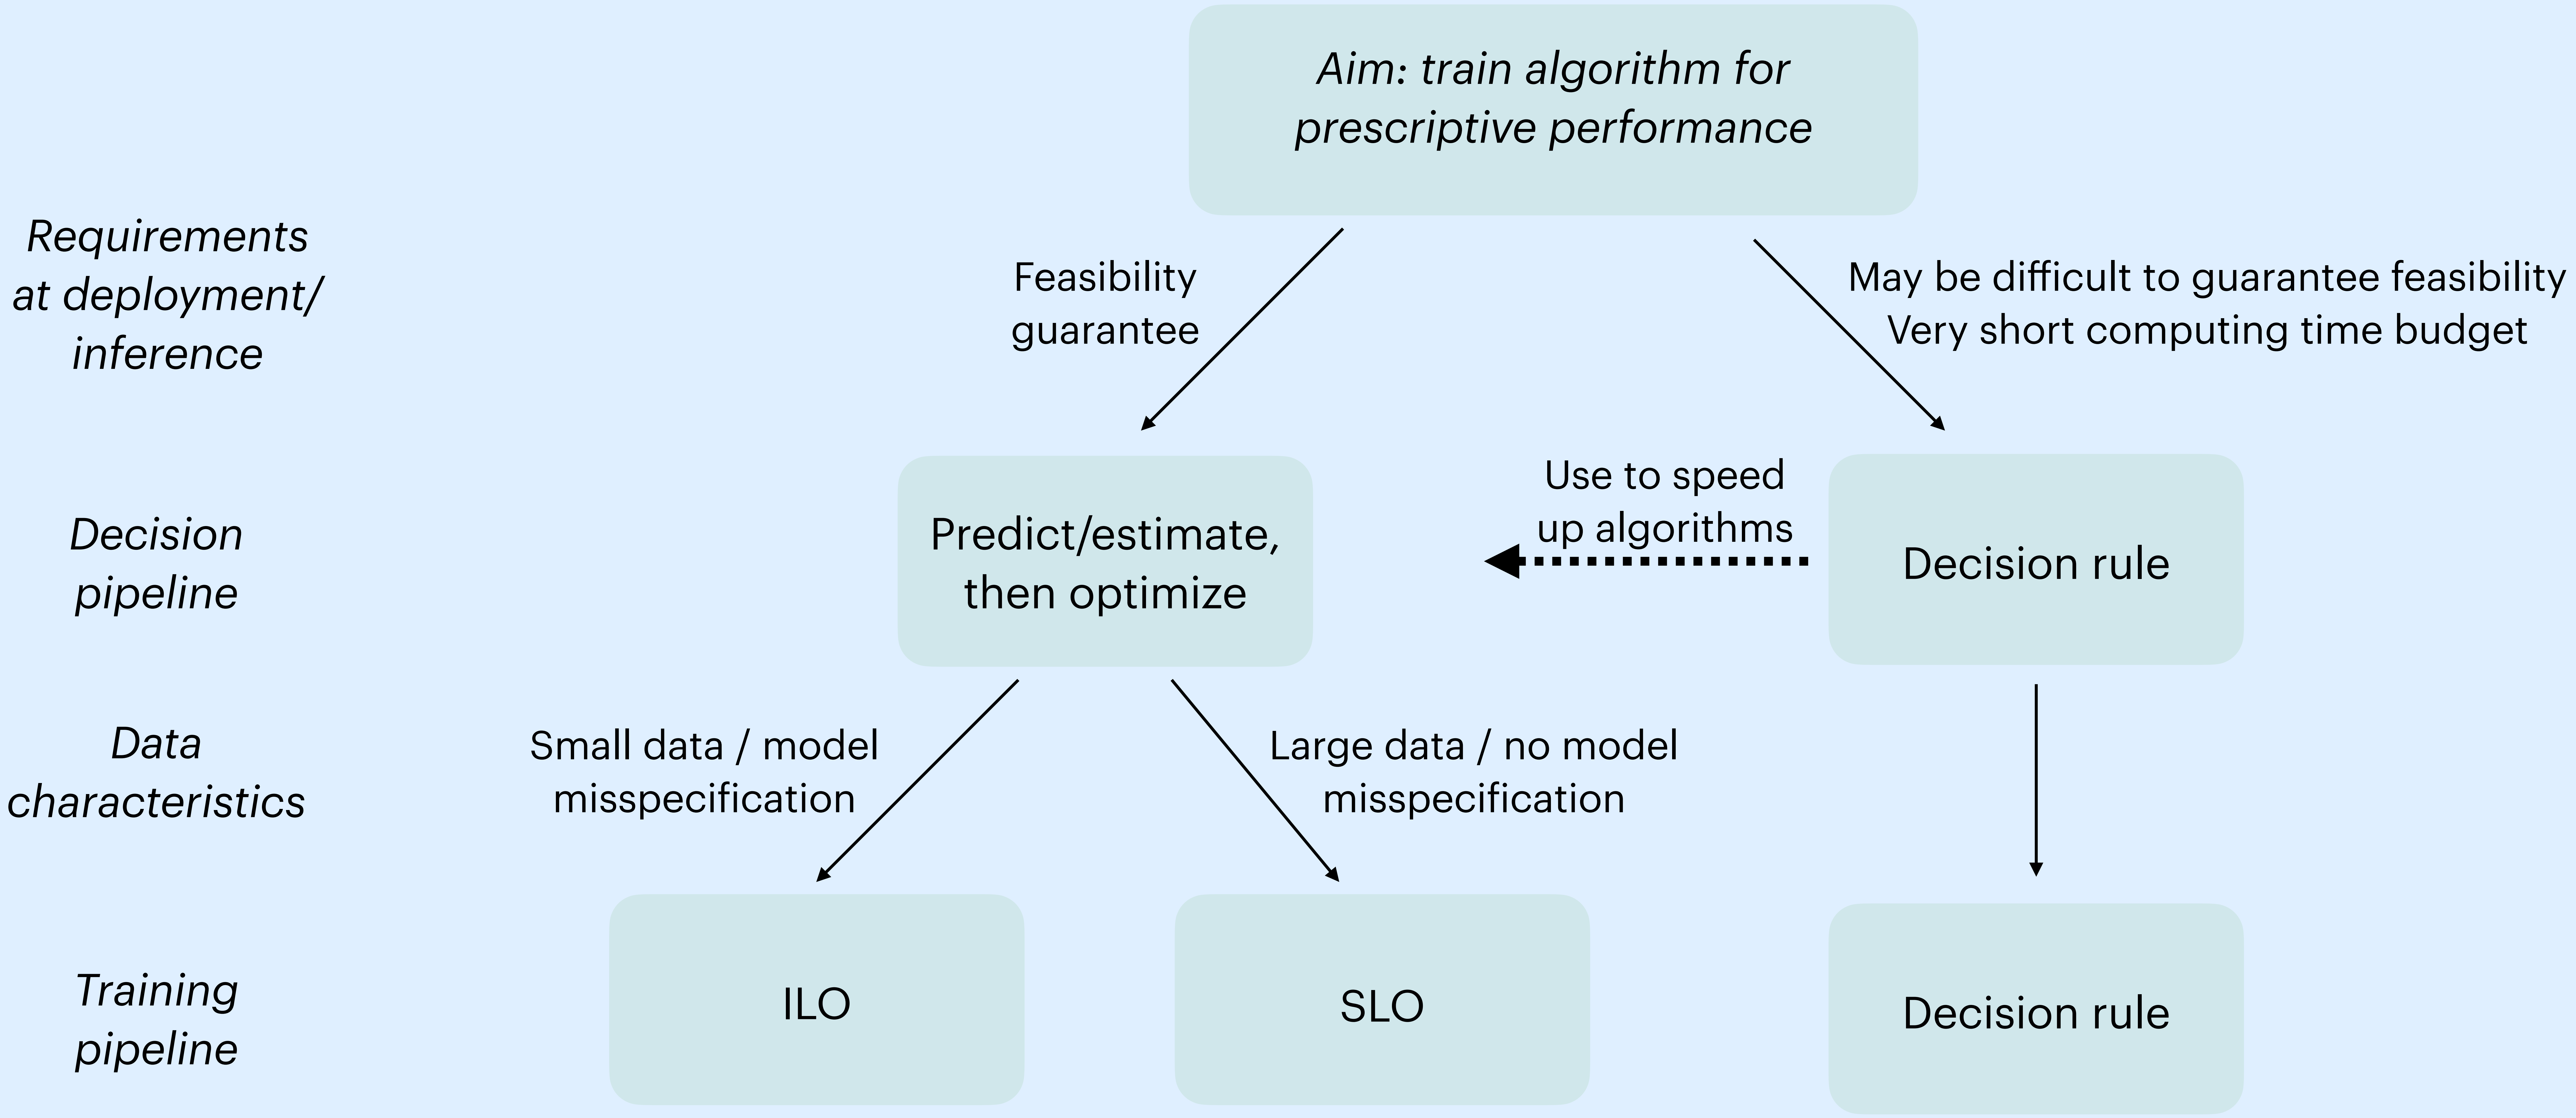

TOPICS WE COVER NEXT (FOCUS ON CONSTRAINED PROBLEMS)

2

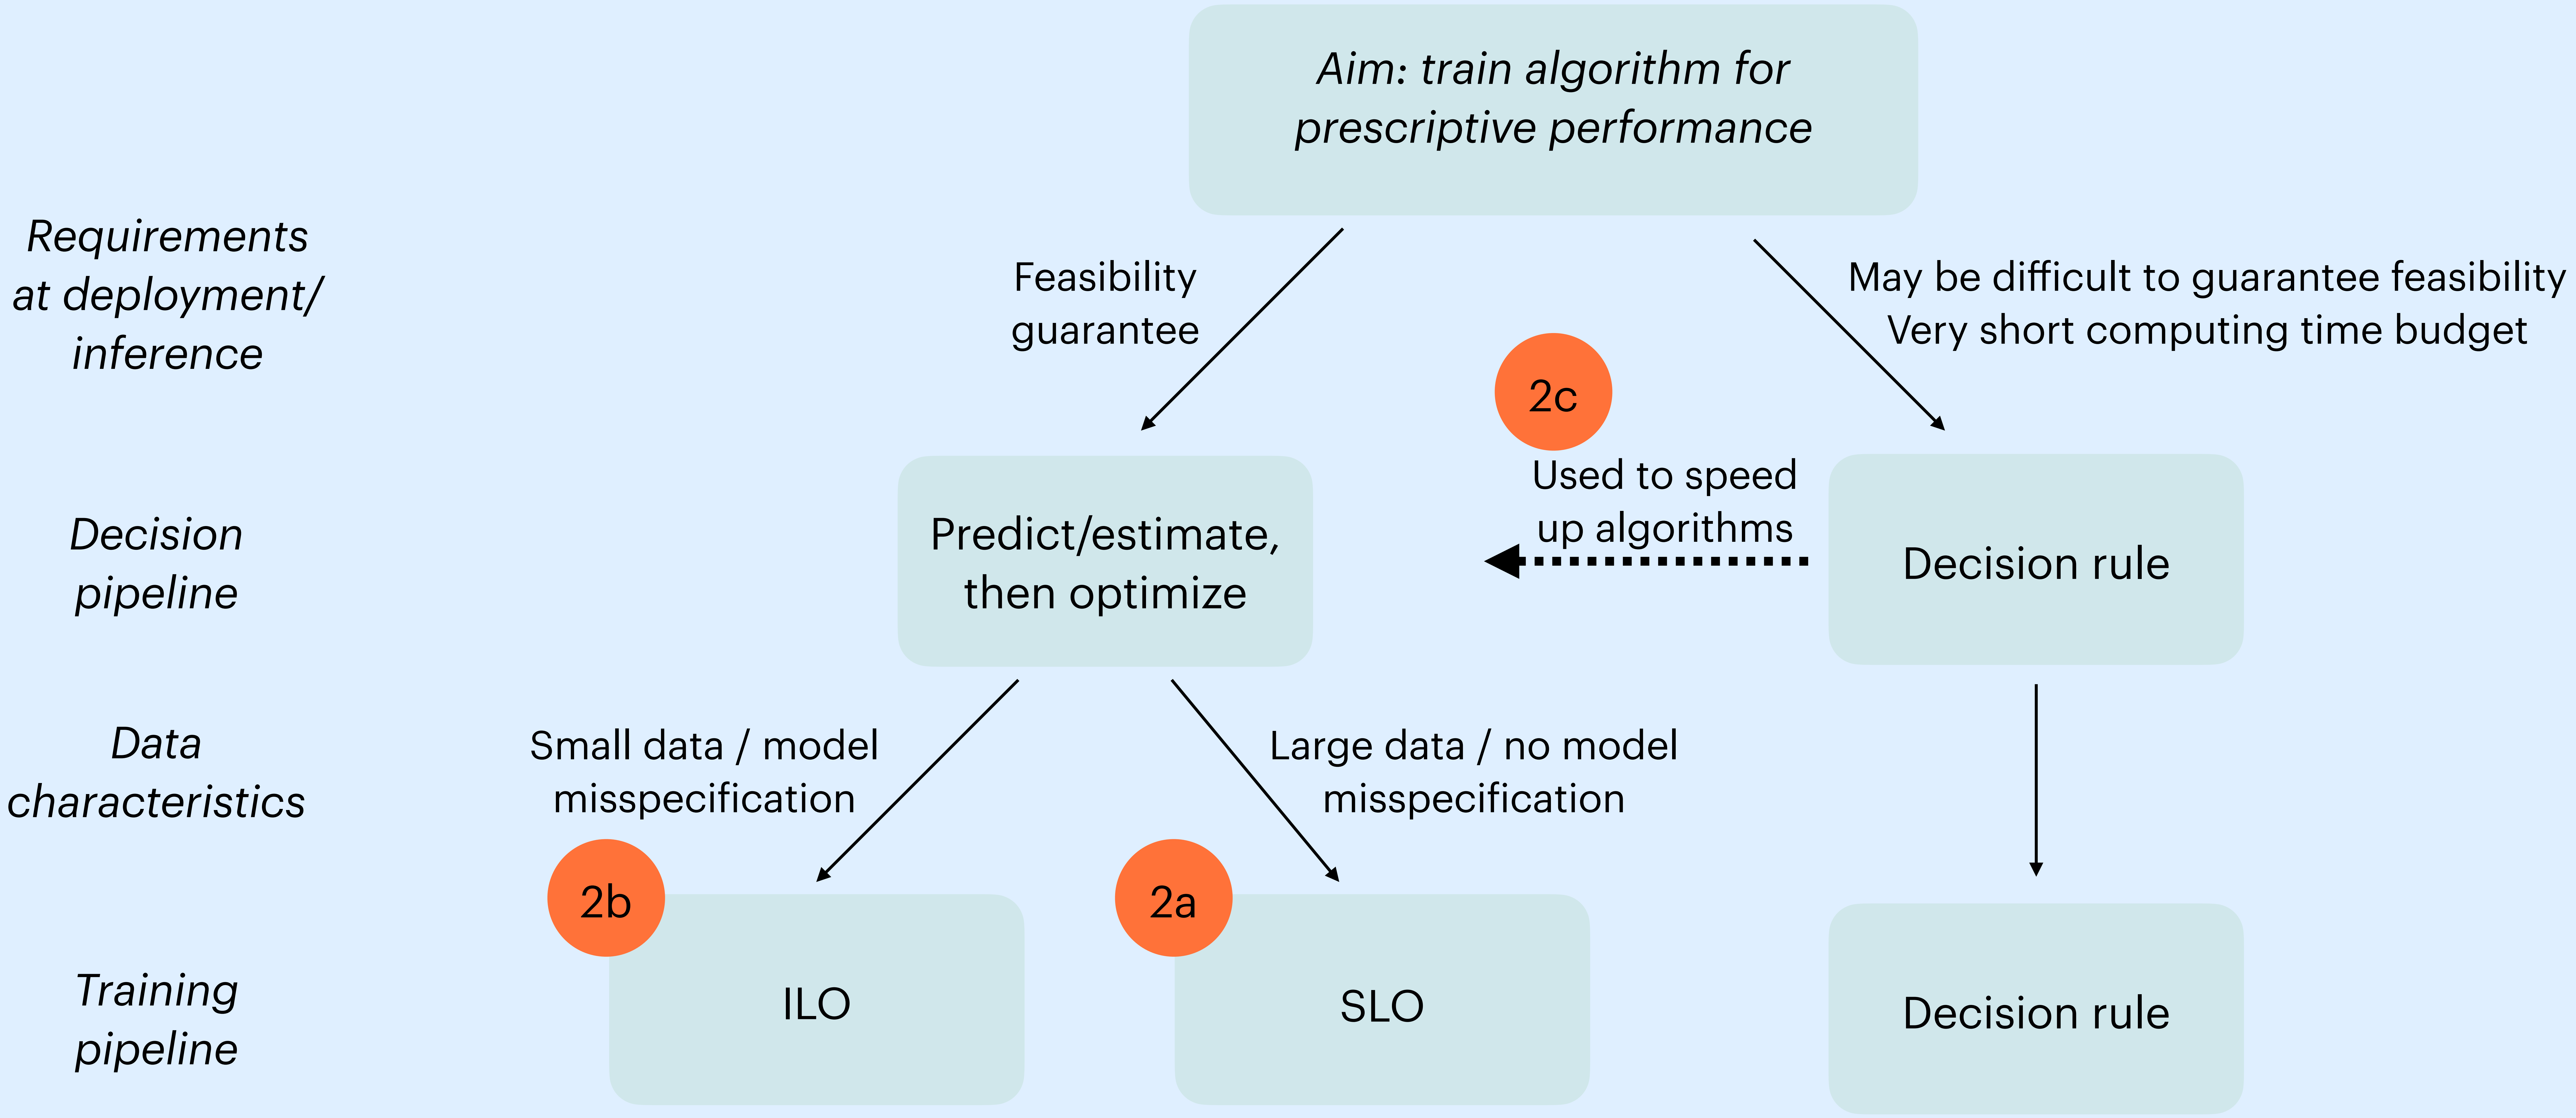

2a

SLO: ESTIMATE DISCRETE CONDITIONAL DISTRIBUTION AND SOLVE A SAA PROBLEM

Residual-based distribution

Measure residuals (errors) of a trained regression model on historical data to construct a conditional distribution

Pointers: Ban et al. (2019), Kannan et al. (2020), Deng and Sen (2022)

Weight-based distribution

Assign weights to historical data

Weights based on **proximity** (between $\boldsymbol{x}$ and historical observations using e.g. kNN), or using **supervised learning** (e.g. random forests)

Pointers: Bertsimas and Kallus (2020), Ban and Rudin (2019)

Many works in the literature, see survey for references.

E.g., studies on how to prevent overly optimistic policies (Bertsimas and Van Parys, 2022) through regularization and distributionally robust optimization.

Sadana, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

EXPECTED VALUE-BASED MODELS: POINT PREDICTION SUFFICE

Decision pipeline

Conditional distribution-based models

Predicted distribution $f_{\boldsymbol{\theta}}(\boldsymbol{x})$ that approximates $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})$

Solve CSO under $f_{\boldsymbol{\theta}}(\boldsymbol{x})$

$$\xrightarrow[\boldsymbol{x}]{\text{Context}} \boxed{\text{Prediction model}} \xrightarrow[f_{\boldsymbol{\theta}}(\boldsymbol{x})]{\text{Prediction}} \boxed{\text{Optimization model}} \xrightarrow[\boldsymbol{z}]{\text{Decision}}$$

Cost function of optimization model is **linear** in $\boldsymbol{y}$

$$\mathbb{E}_{\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})}[c(\boldsymbol{z},\boldsymbol{y})] = c(\boldsymbol{z}, \mathbb{E}_{\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})}[\boldsymbol{y}])$$

Decision pipeline

Expected value-based models

Predicted distribution $g_{\boldsymbol{\theta}}(\boldsymbol{x})$ that approximates $\mathbb{E}_{\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})}[\boldsymbol{y}]$

$$\xrightarrow[\boldsymbol{x}]{\text{Context}} \boxed{\text{Prediction model}} \xrightarrow[g_{\boldsymbol{\theta}}(\boldsymbol{x})]{\text{Prediction}} \boxed{\text{Optimization model}} \xrightarrow[\boldsymbol{z}]{\text{Decision}}$$

Sadana, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

Survey on expected value-based models: Mandi, Kotary, Verden, Mulamba, Bucarey, Guns, Fioretto, Decision-Focused Learning: Foundations, State of the Art, Benchmark and Future Opportunities, Journal of Artificial Intelligence 81:1623-1701, 2024.

THREE GRADIENT-BASED METHODS FOR ILO

Implicit differentiation

Quadratic programs with uncertain feasible domain Donti et al. (2017), Amos and Kolter (2017)

Linear objective, integer variables Wilder et al. (2019), Ferber et al. (2020)

Surrogate differentiable loss function

Regret minimization

$$c\left(\boldsymbol{z}^*(\boldsymbol{x}_i, g_\theta(\boldsymbol{x}_i)), \boldsymbol{y}_i\right) - c\left(\boldsymbol{z}^*(\boldsymbol{x}_i, \boldsymbol{y}_i), \boldsymbol{y}_i\right)$$

Linear objective in predicted parameters, continuous or integer variables, seminal work by Elmachtoub and Grigas (2022) “smart predict, then optimize”

Surrogate differentiable optimizer

Stochastic perturbations of parameters

Lead to differentiable optimization problems

Berther et al. (2020), Dalle et al. (2022)

Majority of studies focus on **expected value-based models.**
Many different methods, each with its specific tricks to deal with **gradient** being zero almost everywhere, and **computational challenges** under specific **restrictions on the models** considered.
Rem: In case of differentiable model it allows for **“optimization as a layer”** in neural networks.

Sadana, Chenreddy, Delage, Forel, Frejinger, Vidal, A survey on contextual optimization methods for decision-making under uncertainty, European Journal of Operational Research 320:271-289, 2025.

BACK TO THE OVERVIEW: TRADEOFF BETWEEN ILO AND SLO

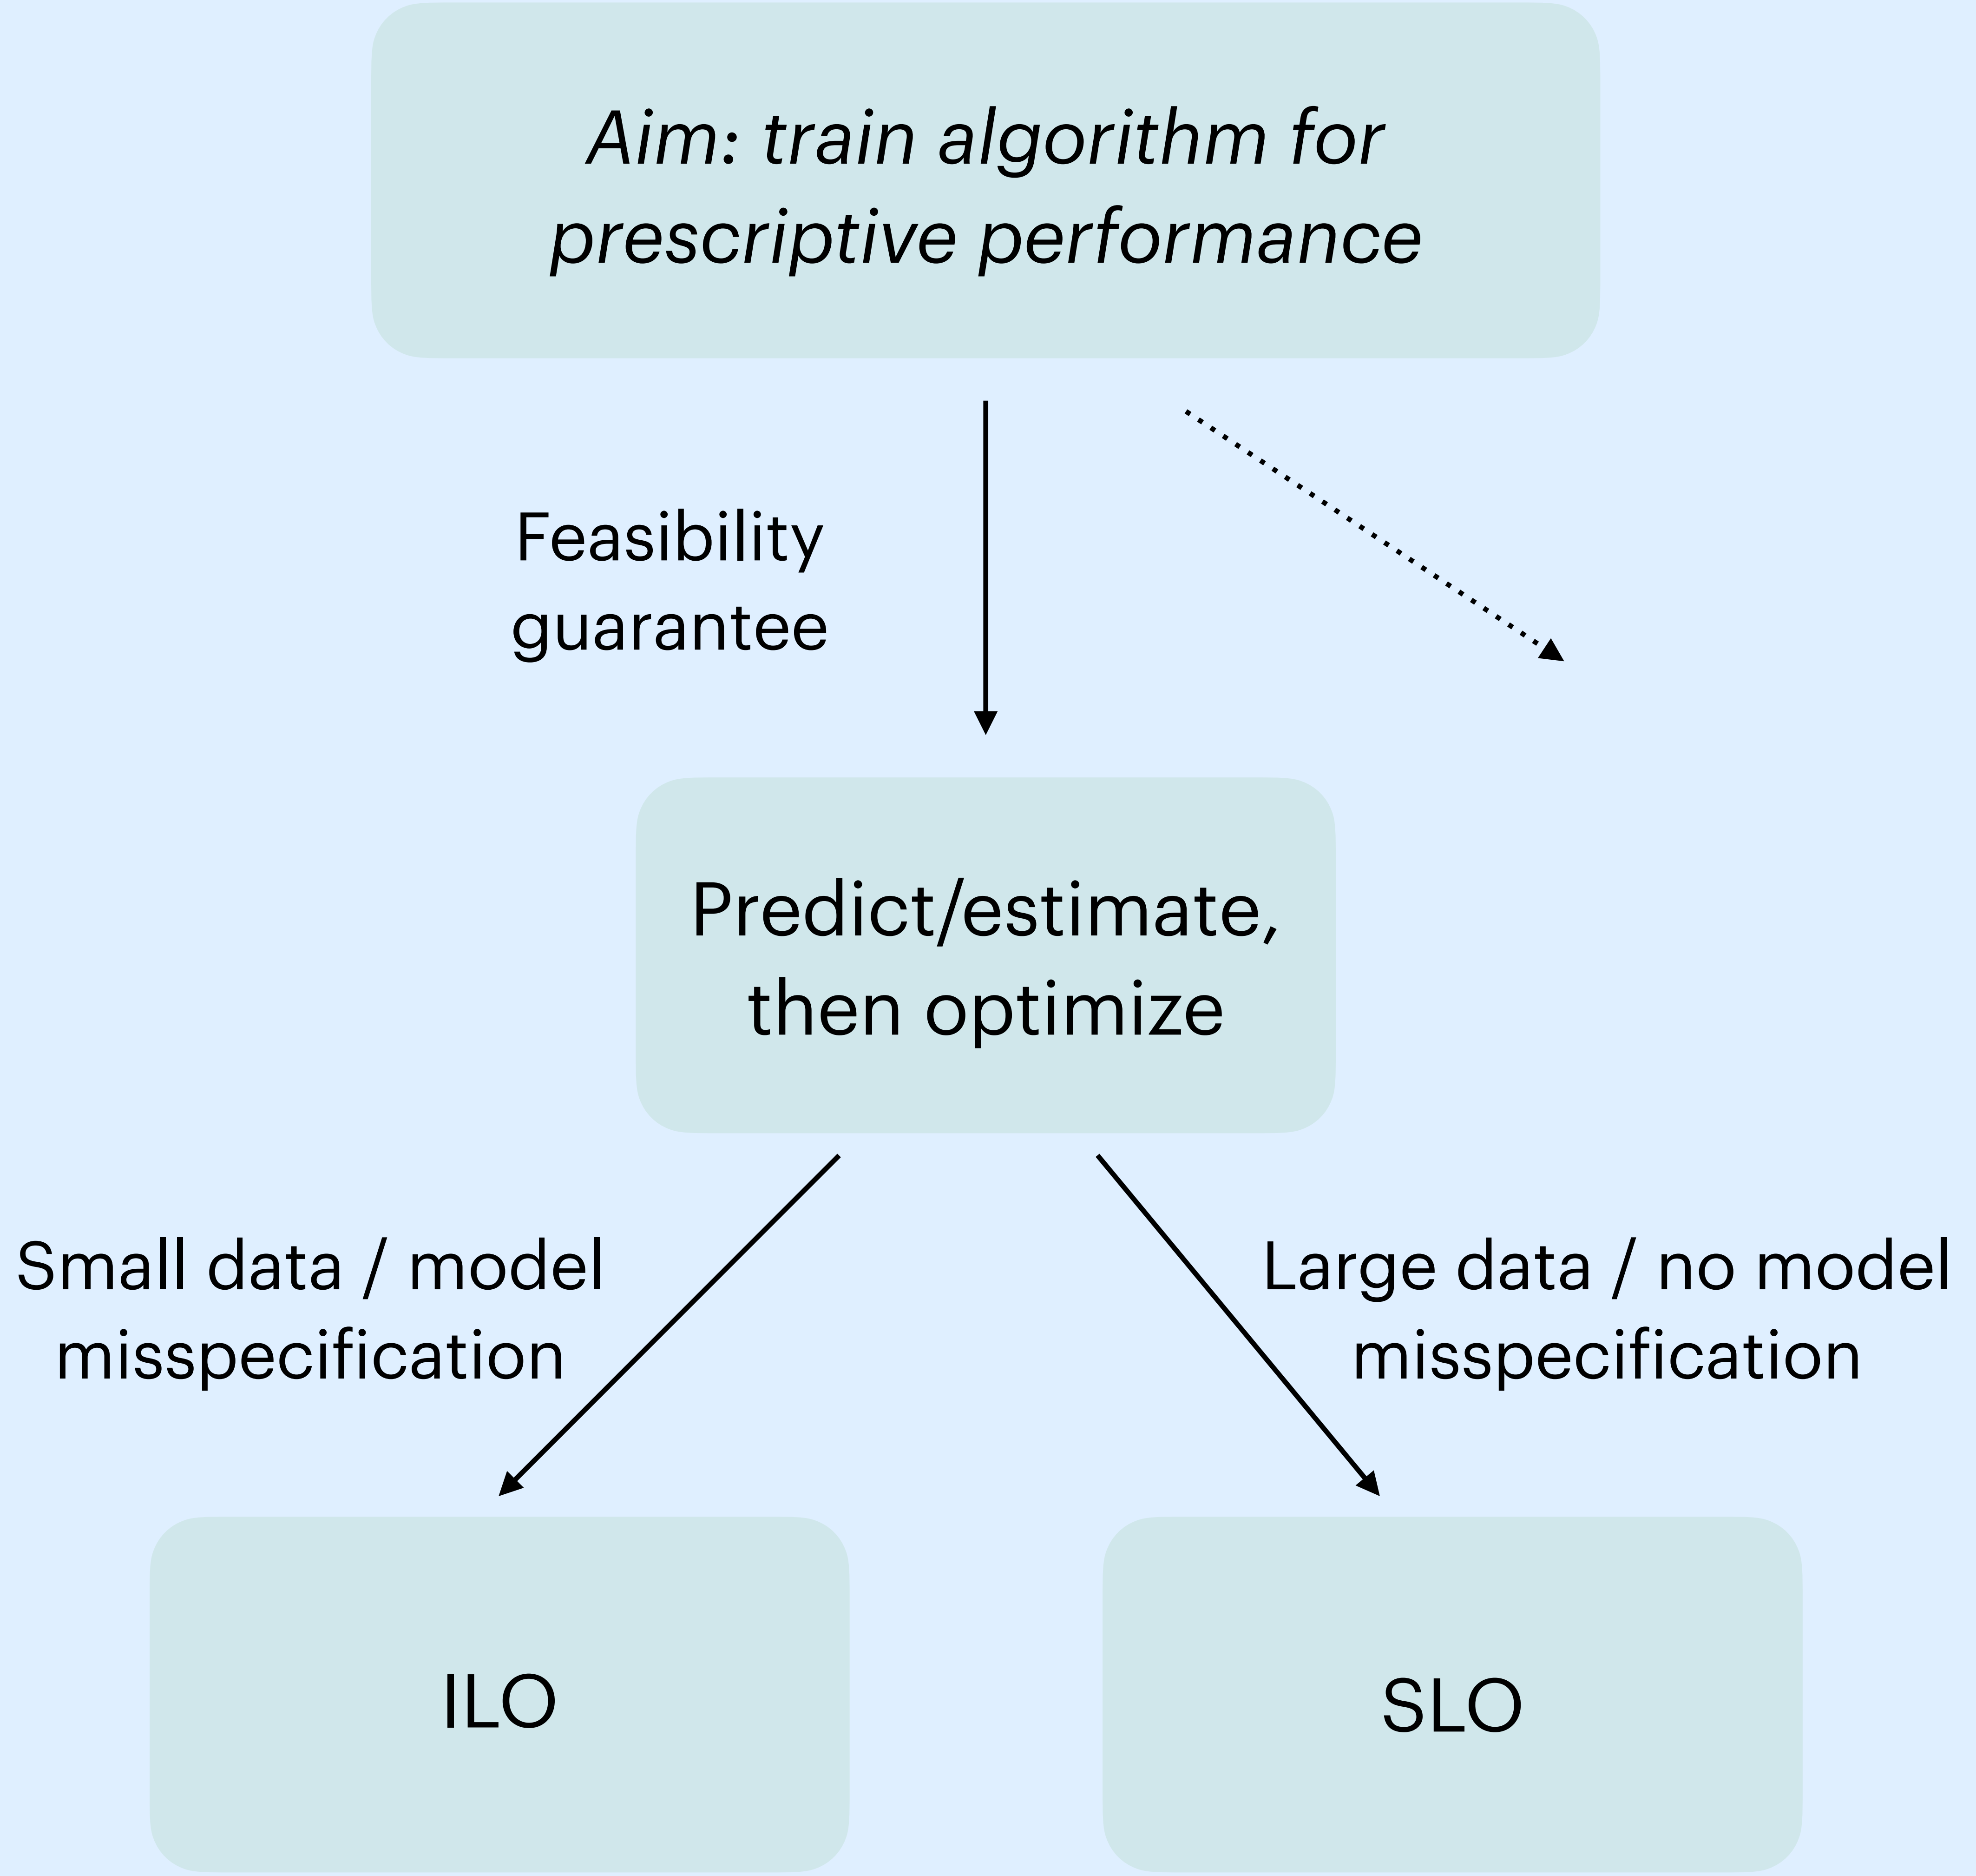

Feasibility guarantee because optimization solver is used at inference/deployment (decision pipeline)

Empirical studies: ILO generally performs better than SLO when there is model misspecification

Theoretical studies: Non-linear CSO, *asymptotic* results, stochastic dominance of *regret distribution*: SLO dominates ILO for well-specified models, ILO dominates SLO for misspecified models (Elmachtoub et al., 2023, arXiv:2304.06833v3)
Finite sample performance (based on Taylor series expansion of *regrets*): ILO has "universal double benefit" over SLO under model misspecification, SLO may be advantageous for nearly well-specified models (Elmachtoub et al., 2025, arXiv:2503.00626v1)

NEXT: SPEED UP ALGORITHMS USING DECISION RULES

2

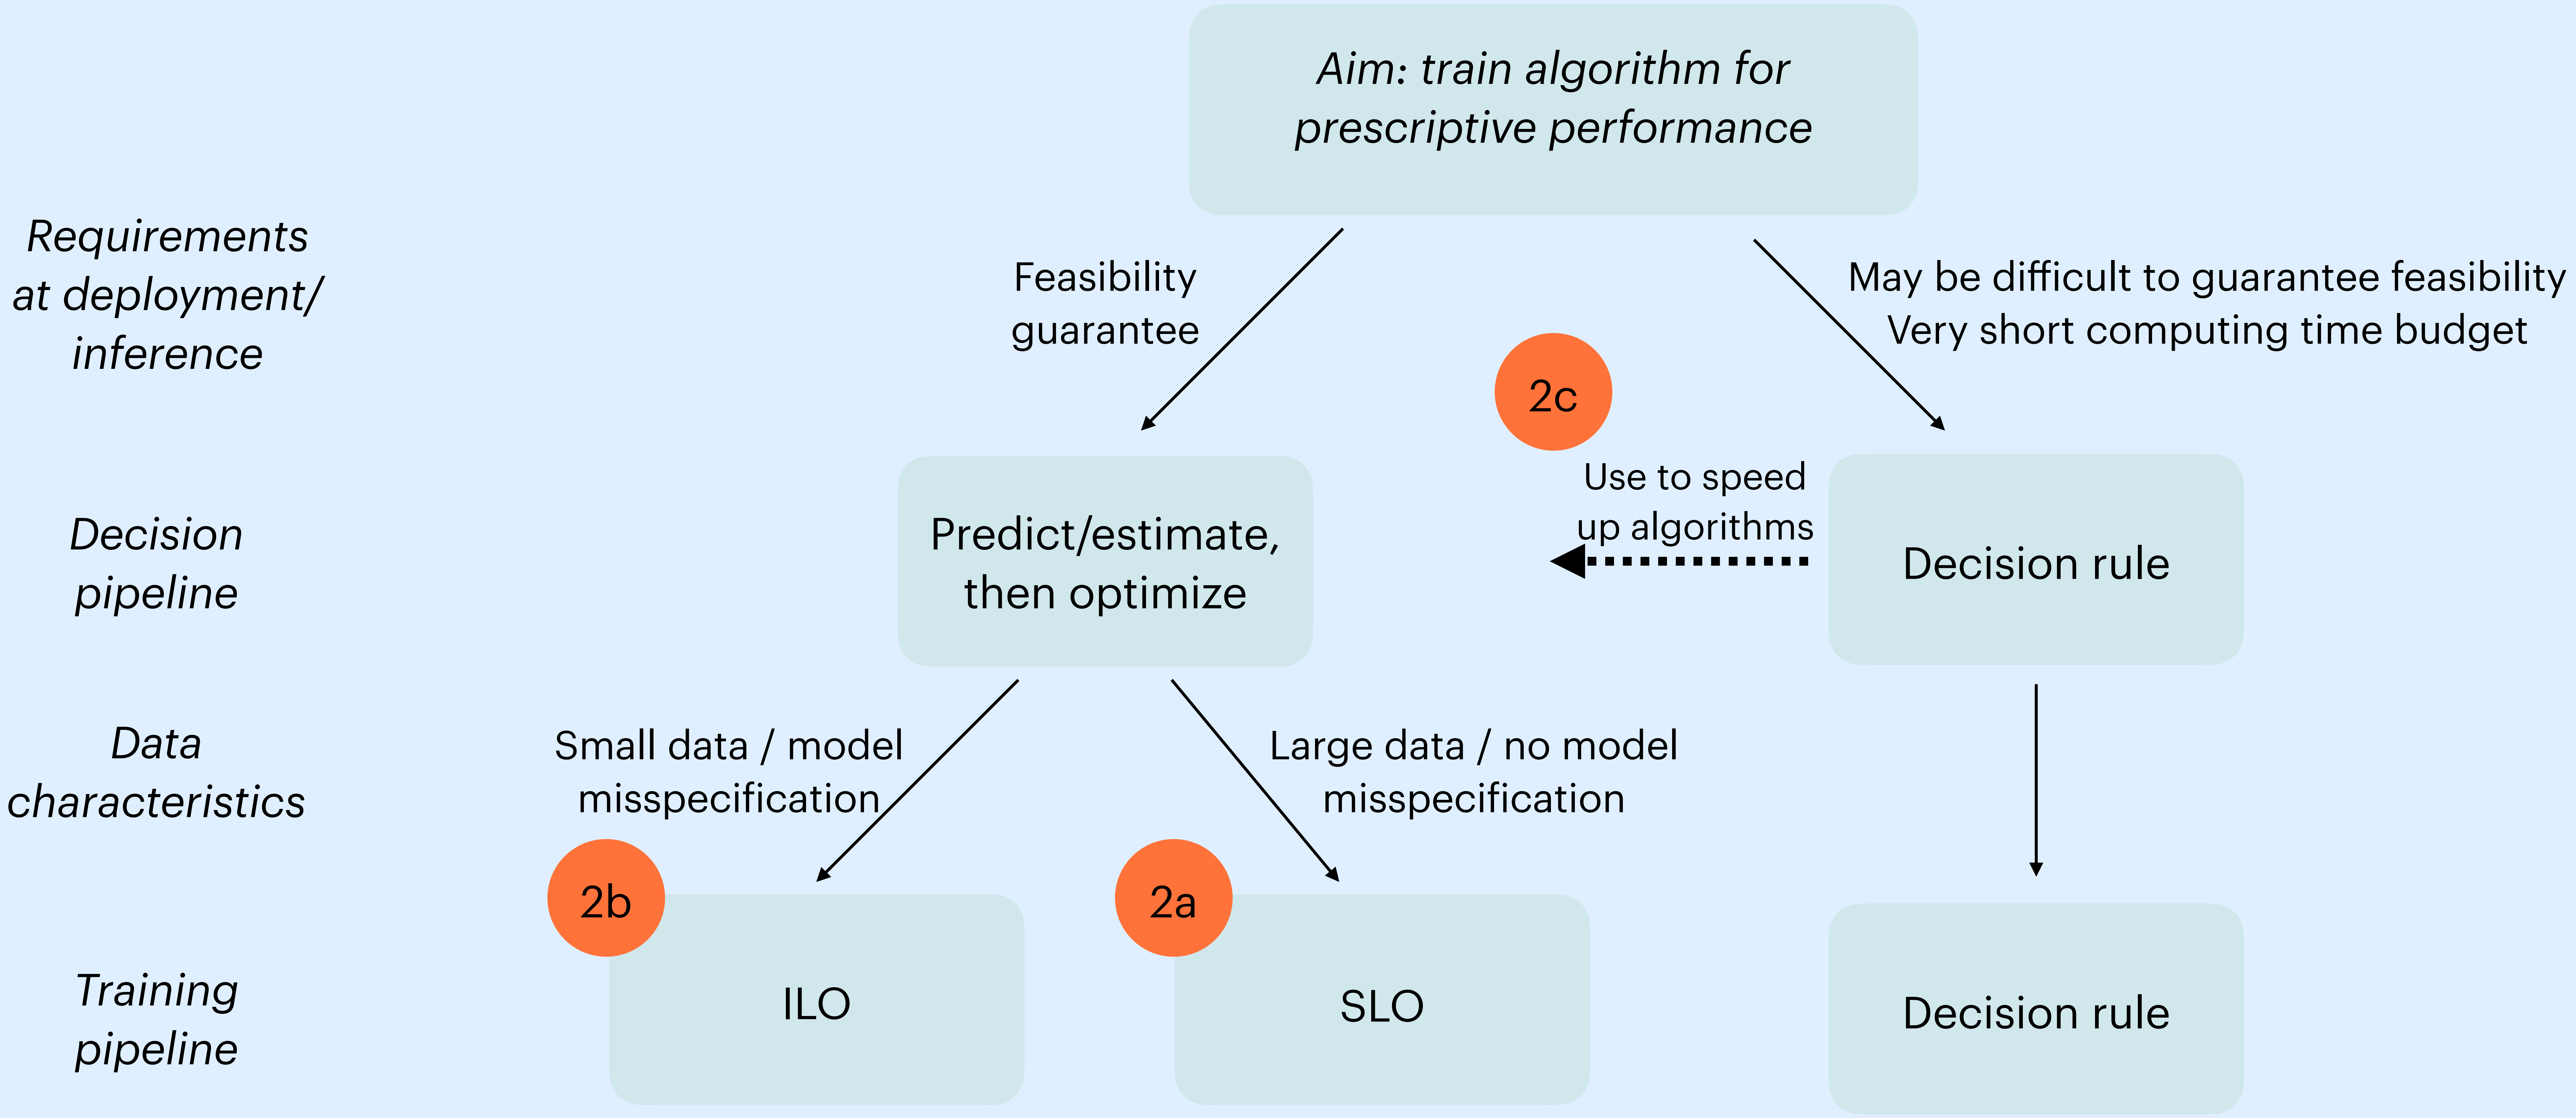

DECISION RULES: USEFUL BEYOND CSO, ON THEIR OWN AND IN OTHER ALGORITHMS

2c

Parameters defining problem instance → Decision rule → Solution

Machine learning oracle replace solver at inference

- Setting: Similar instances solved repeatedly over time
- Aim: *Exploit shared structure and use machine learning to predict high-quality solutions in very short compute time (online / inference)*
- Used extensively for continuous unconstrained problems (Amos, **Tutorial** on amortized optimization, 2022, arXiv:2202.00665v3)
- Challenging in presence of constraints and discrete variables (Kotary et al., End-to-End Constrained Optimization Learning: **A Survey**, 2021, arXiv:2103.16378v1)

How to guarantee feasibility in the presence of non-trivial constraints?

“OPTIMIZATION PROXIES” SUCCESSFULLY APPLIED IN POWER SYSTEMS

Decision pipeline

Parameters defining problem instance → [Machine learning model predicts a solution → Solution repaired by solving a system of equations] → Feasible Solution

Optimization Proxy of a linear program

Van Hentenryck, Optimization Learning, arXiv:2501.03443v1, 2025.

"OPTIMIZATION PROXIES" SUCCESSFULLY APPLIED IN POWER SYSTEMS

Decision pipeline

Parameters defining problem instance → [Machine learning model predicts a solution → Solution repaired by solving a system of equations] → Feasible Solution

Optimization Proxy of a linear program

Training pipeline

Example instance → Optimization proxy (neural net) $\xrightarrow{\pi_\theta}$ Task loss Obj. fn. value

Gradient $\nabla_\theta c(\pi_\theta)$

No precomputed labels (solutions), only access to objective function

Van Hentenryck, Optimization Learning, arXiv:2501.03443v1, 2025.

"OPTIMIZATION PROXIES" SUCCESSFULLY APPLIED IN POWER SYSTEMS

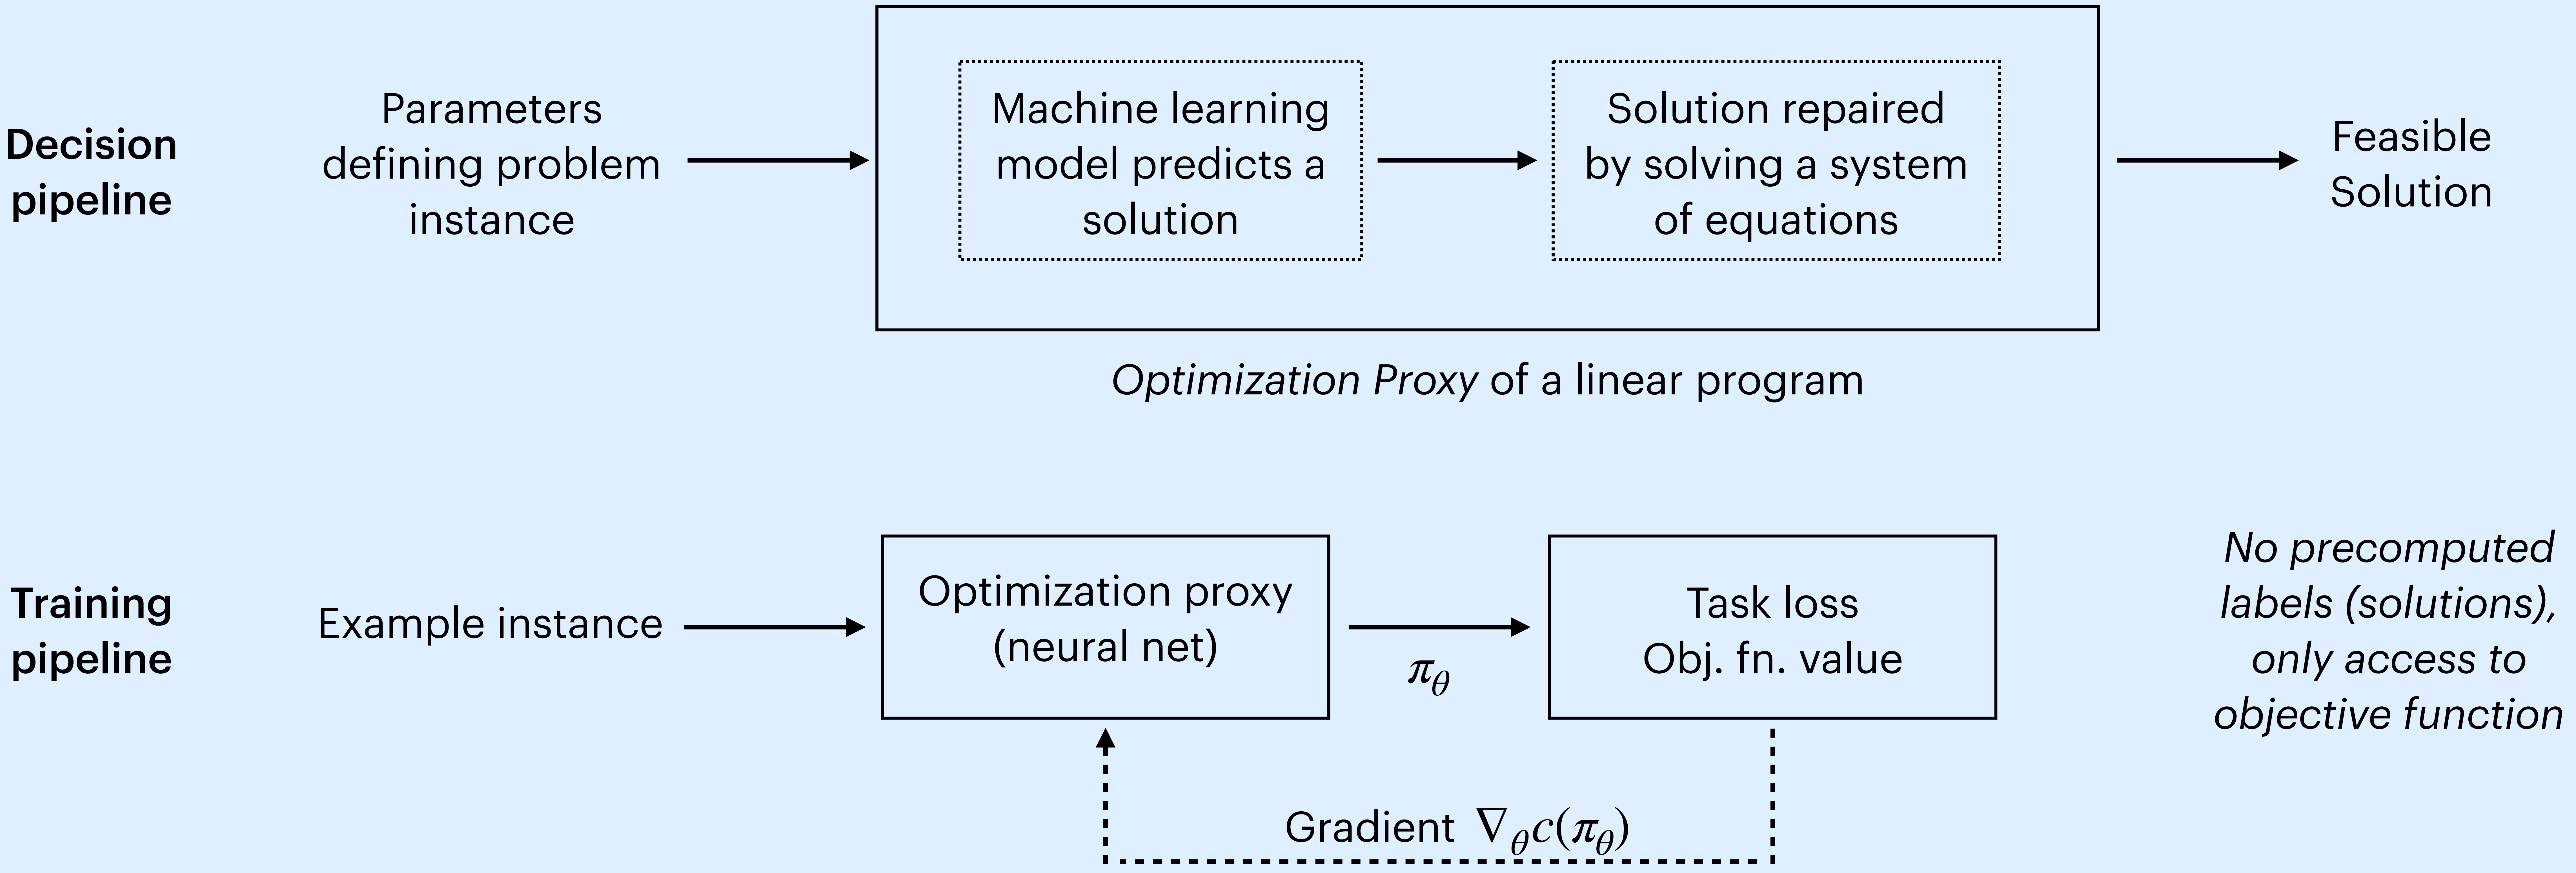

Feasibility guaranteed thanks to problem-specific "repair" layers.

Training "end-to-end" leads to better performance compared to only repairing solution at inference time.

Order of magnitudes improvements in efficiency! Makes e.g. real-time risk assessment possible.

Van Hentenryck, Optimization Learning, arXiv:2501.03443v1, 2025.

SLO AND "DECISION RULES" SUCCESSFULLY USED IN A MATHEURISTIC TO SOLVE TWO-STAGE STOCHASTIC PROGRAMS (MIXED INTEGER)

2c

Decision pipeline

Context $\boldsymbol{x}$ → Prediction model → Optimization model & solver → Decision $z^*(\boldsymbol{x})$

Two-stage stochastic program

Integer second stage: context includes first-stage decision variables

Benders decomposition method (L-Shaped)

- Computational bottleneck: iteratively computing information about second-stage solution when solving first-stage problem

Larsen, Lachapelle, Bengio, Frejinger, Lacoste Julien, Lodi, Predicting Tactical Solutions to Operational Planning Problems Under Imperfect Information, INFORMS Journal on Computing 34(1):227-242, 2022.

Larsen, Frejinger, Gendron, Lodi, Fast Continuous and Integer L-Shaped Heuristics Through Supervised Learning, INFORMS Journal on Computing 36(1):203-223, 2024.

SLO AND "DECISION RULES" SUCCESSFULLY USED IN A MATHEURISTIC TO SOLVE TWO-STAGE STOCHASTIC PROGRAMS (MIXED INTEGER)

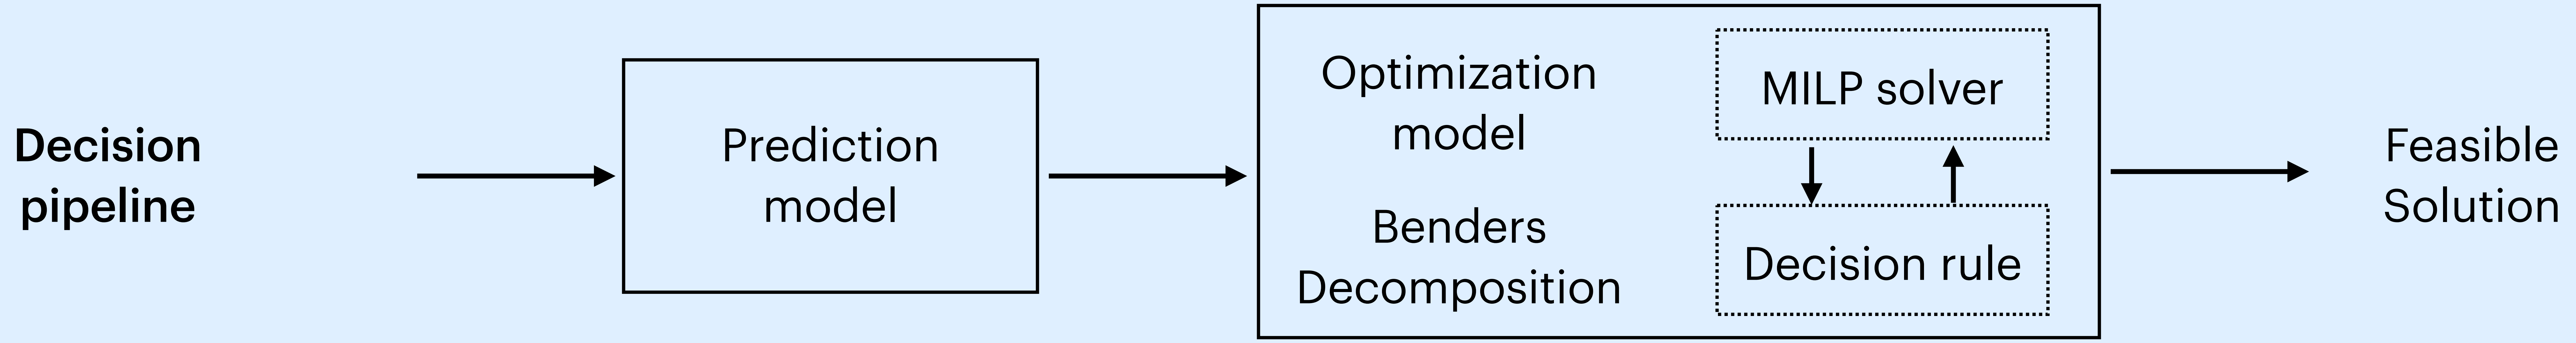

Larsen, Lachapelle, Bengio, Frejinger, Lacoste Julien, Lodi, Predicting Tactical Solutions to Operational Planning Problems Under Imperfect Information, INFORMS Journal on Computing 34(1):227-242, 2022.

Larsen, Frejinger, Gendron, Lodi, Fast Continuous and Integer L-Shaped Heuristics Through Supervised Learning, INFORMS Journal on Computing 36(1):203-223, 2024.

SLO AND "DECISION RULES" SUCCESSFULLY USED IN A MATHEURISTIC TO SOLVE TWO-STAGE STOCHASTIC PROGRAMS (MIXED INTEGER)

2c

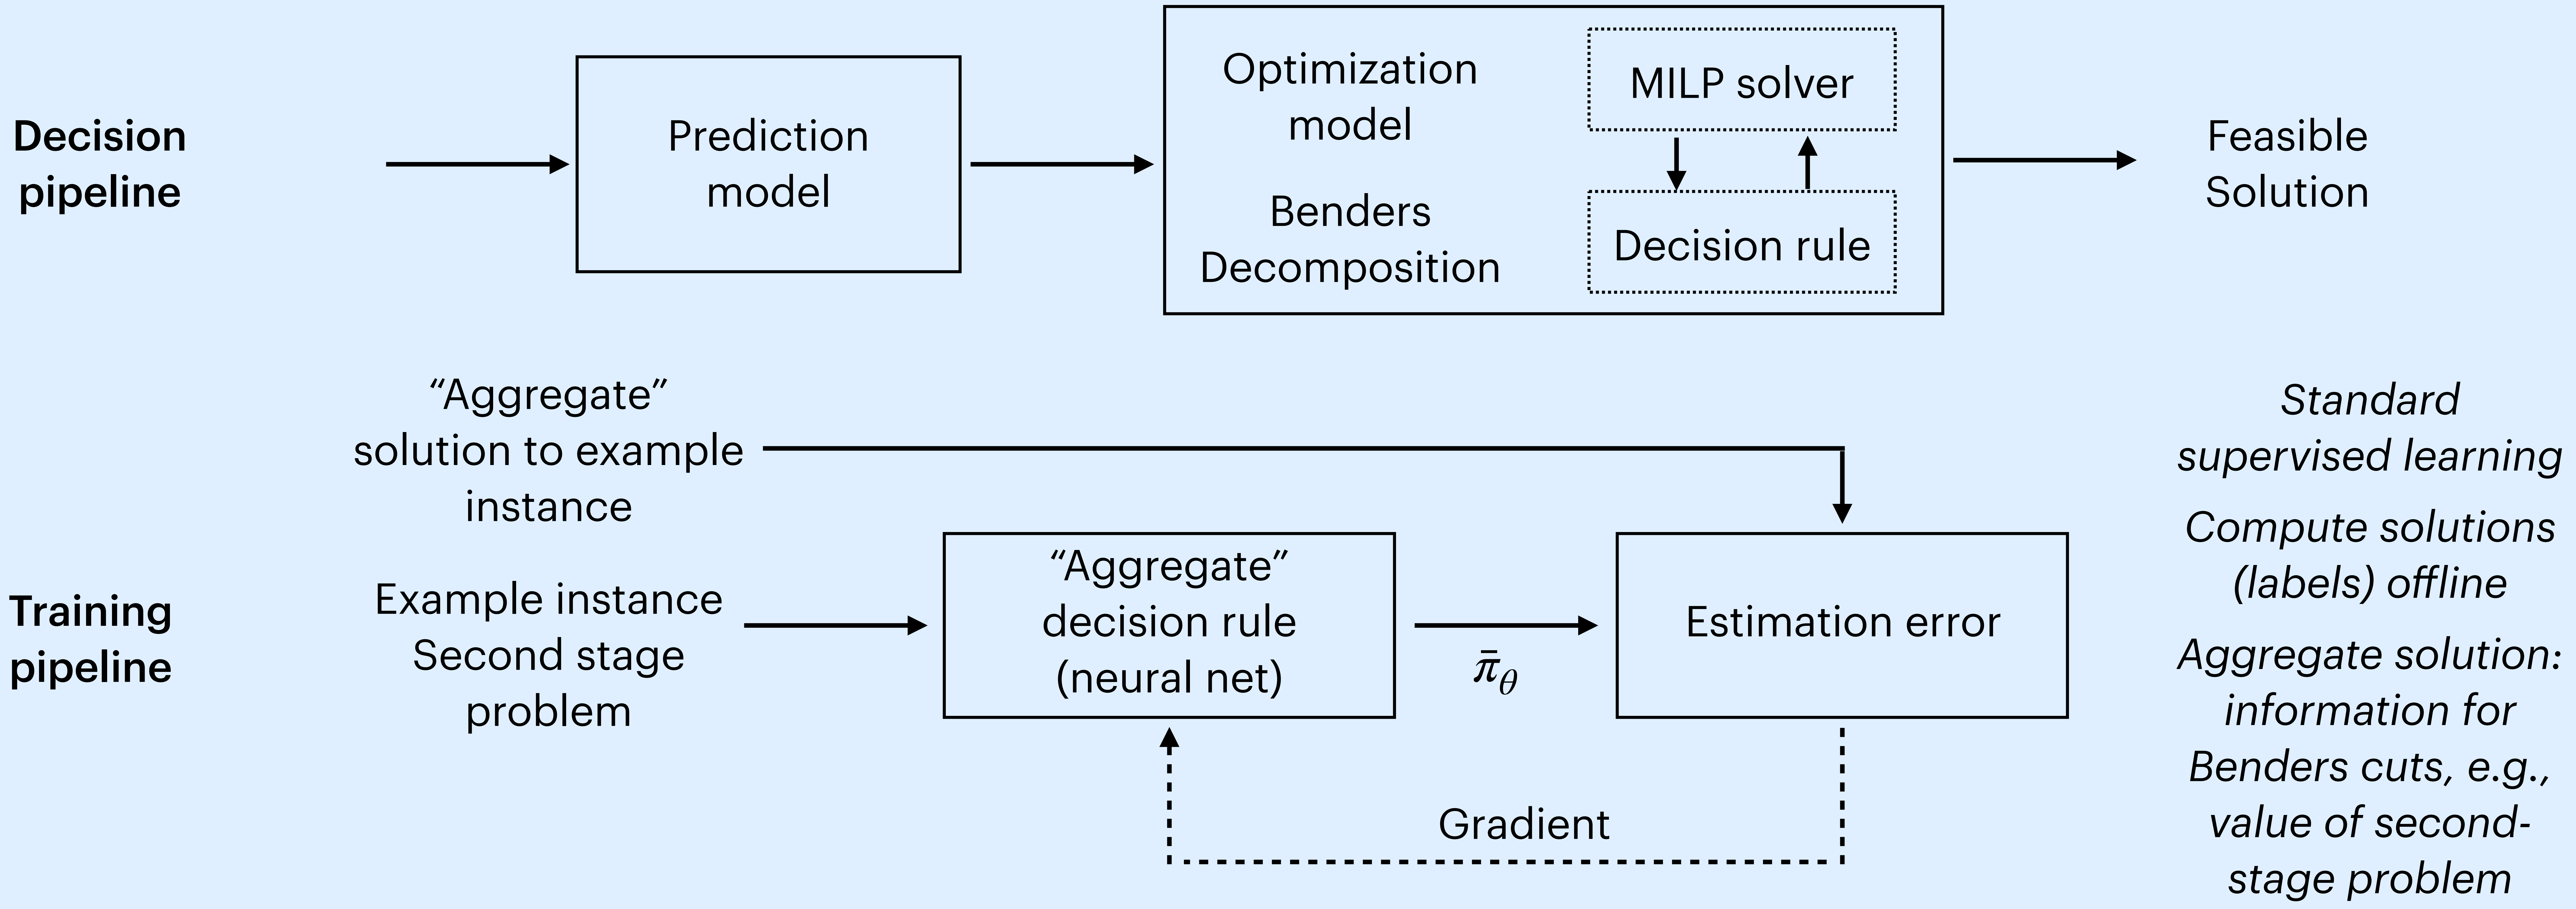

Larsen, Lachapelle, Bengio, Frejinger, Lacoste Julien, Lodi, Predicting Tactical Solutions to Operational Planning Problems Under Imperfect Information, INFORMS Journal on Computing 34(1):227-242, 2022.

Larsen, Frejinger, Gendron, Lodi, Fast Continuous and Integer L-Shaped Heuristics Through Supervised Learning, INFORMS Journal on Computing 36(1):203-223, 2024.

SLO AND "DECISION RULES" SUCCESSFULLY USED IN A MATHEURISTIC TO SOLVE TWO-STAGE STOCHASTIC PROGRAMS (MIXED INTEGER)

2c

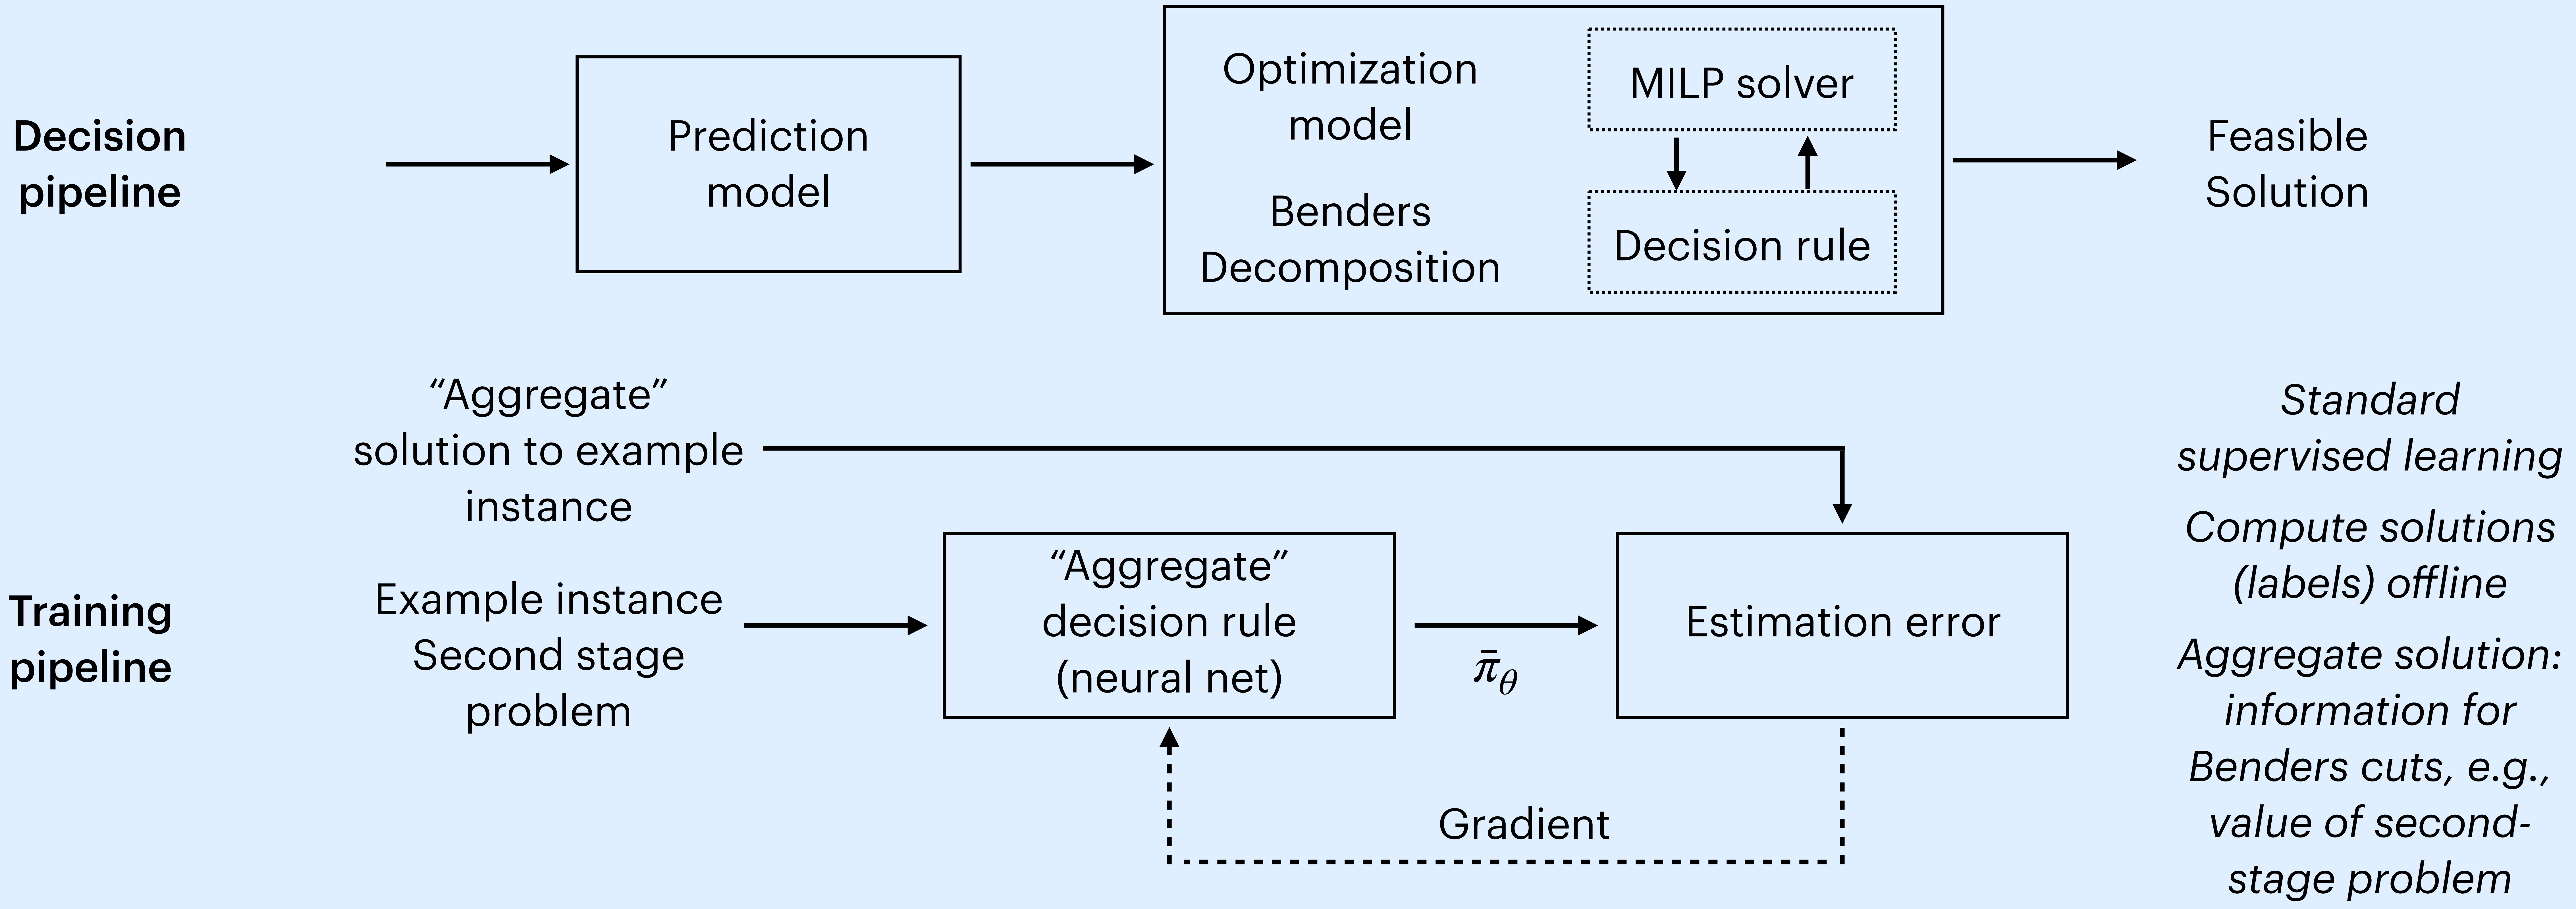

Feasibility guaranteed because general purpose solver is used at deployment.
Speed ups of one to two orders of magnitude and optimality gaps close to zero.

Larsen, Lachapelle, Bengio, Frejinger, Lacoste Julien, Lodi, Predicting Tactical Solutions to Operational Planning Problems Under Imperfect Information, INFORMS Journal on Computing 34(1):227-242, 2022.

Larsen, Frejinger, Gendron, Lodi, Fast Continuous and Integer L-Shaped Heuristics Through Supervised Learning, INFORMS Journal on Computing 36(1):203-223, 2024.

Climbing to the moon, Chris Rowney www.devianart.com

My moonshot: From data to real-world impact

Solve large-scale discrete optimization problems subject to uncertainty (exogenous or endogenous) while leveraging recent relevant data

Outline

1. Contextual optimization lens: Foundational concepts and motivation
2. From data to decision
   - Two decision pipelines: from context to decision
   - Three training pipelines: how to use data to improve prescriptive performance
3. From data to real-world impact
   - What do we know and what topics require more research

NO REAL-WORLD IMPACT WITHOUT ADOPTION

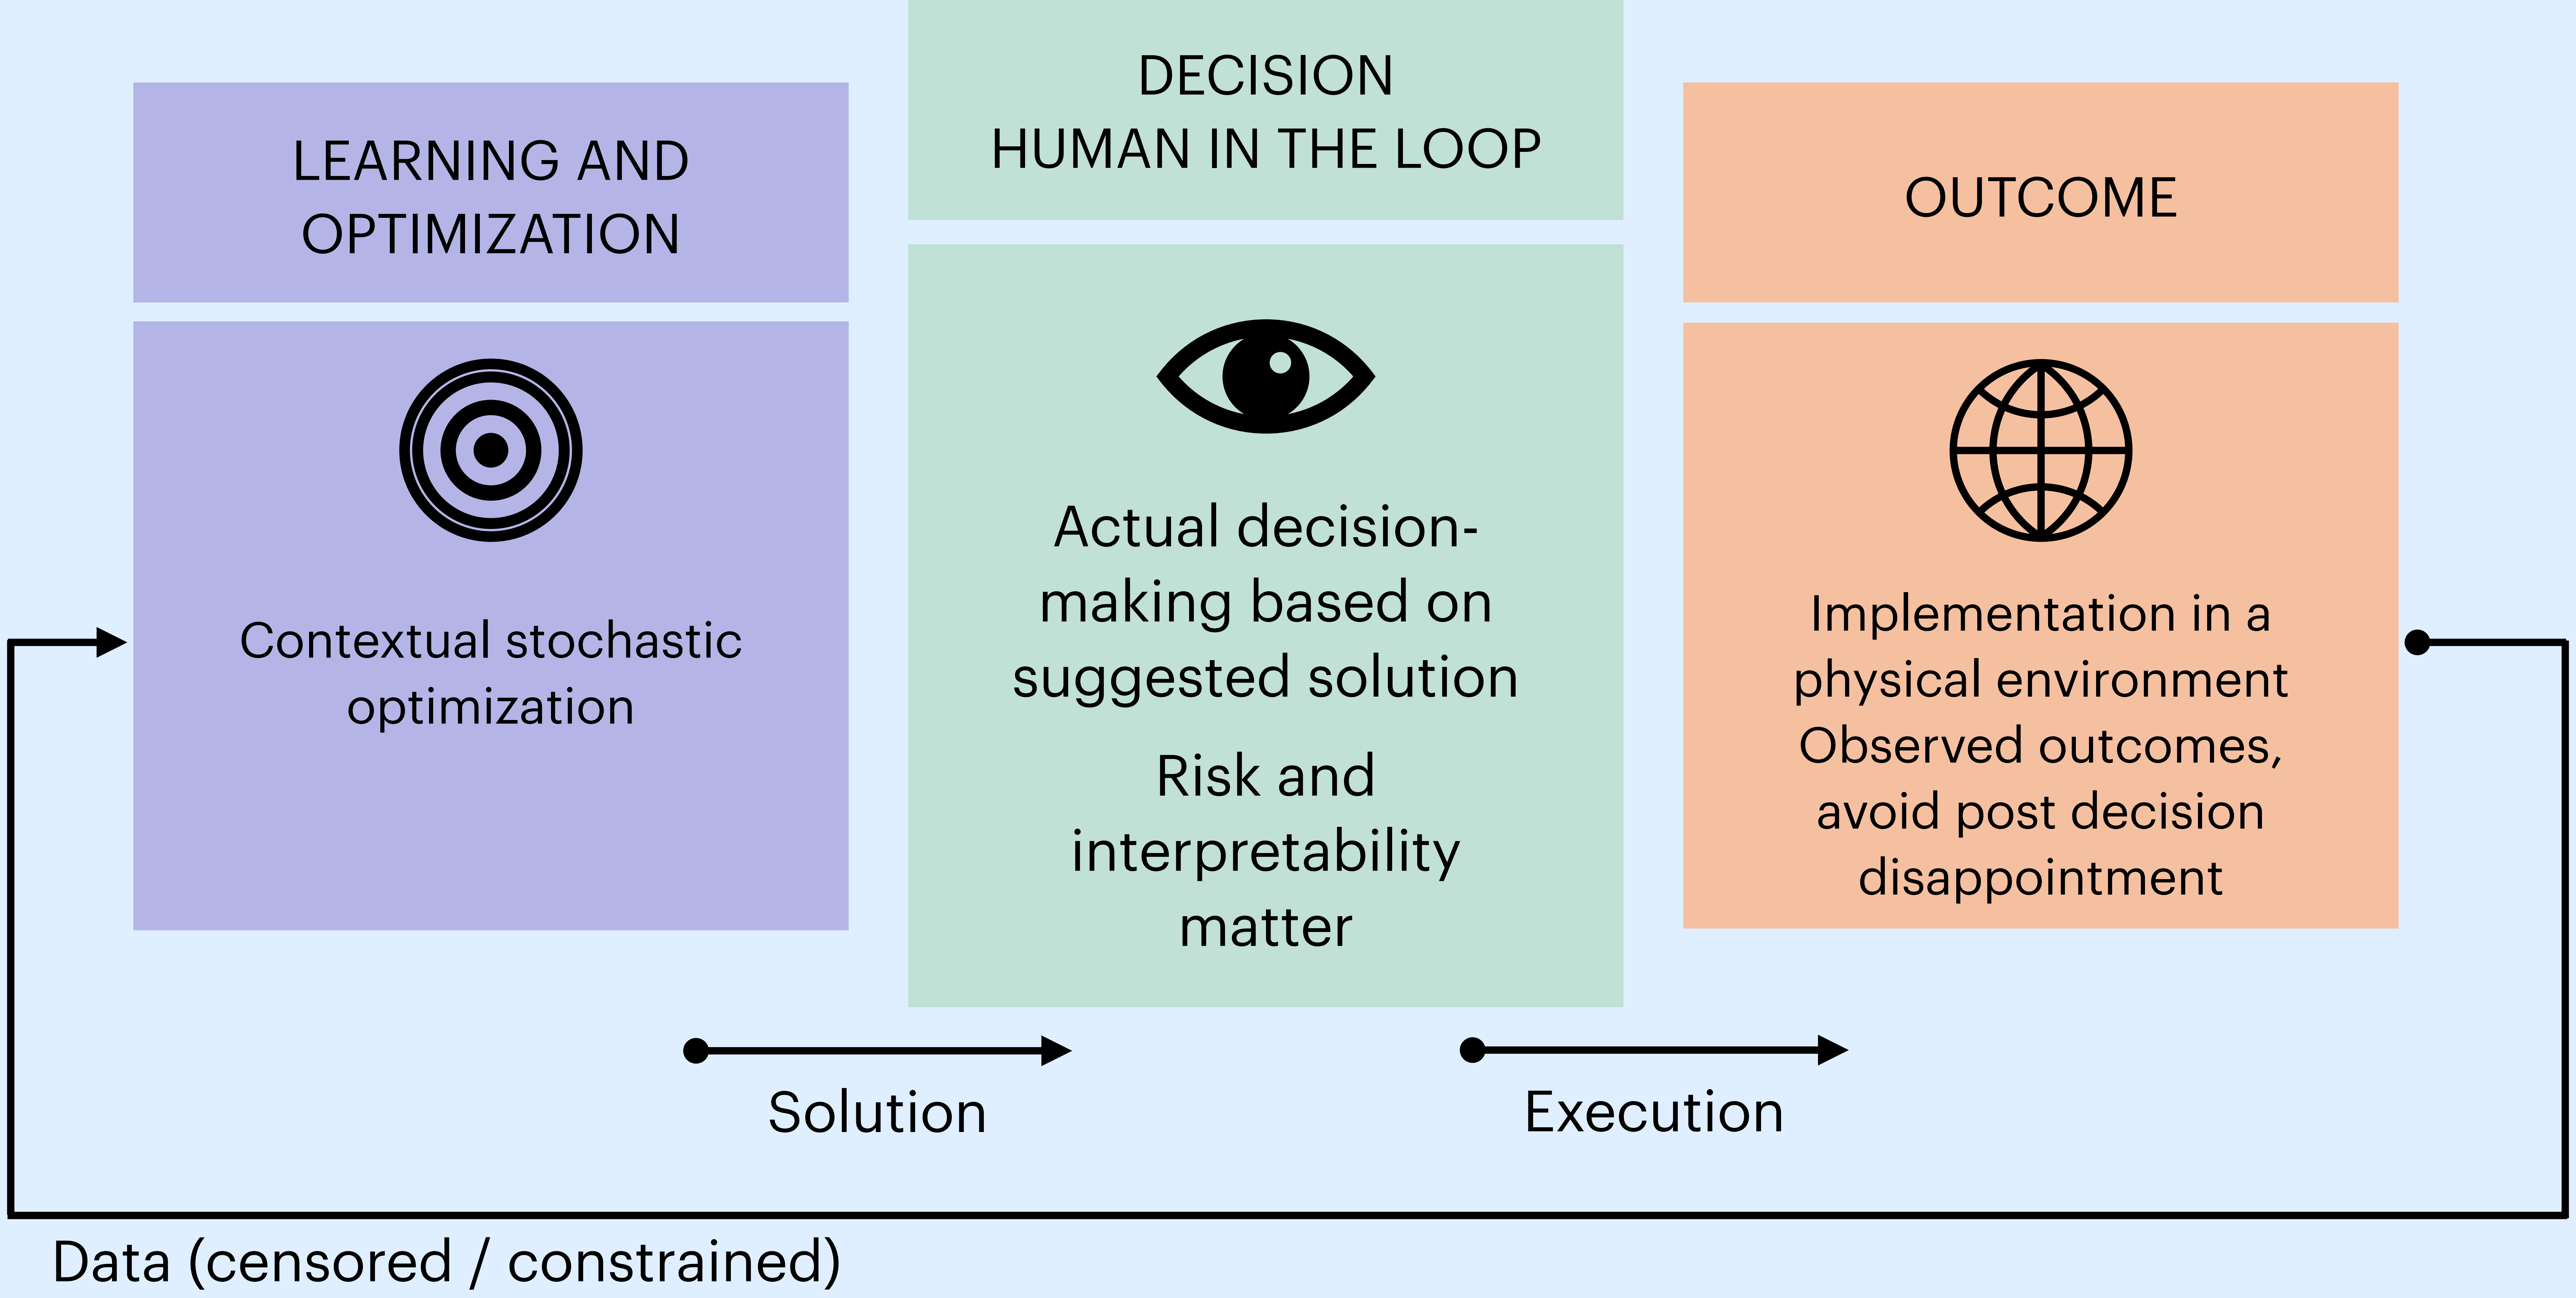

Not only a “practical issue”. Today we did not have time to cover risk models, counterfactual explanations, interpretability (and related regulation), effective exploration, …

WHAT WE KNOW – TAKEAWAYS

Contextual optimization problems subject to known (non-trivial) constraints, exogenous uncertainty in the objective only (convex problems, mixed-integer linear programs)

Select decision pipeline based on application-specific requirements

Predict/estimate, then optimize

or

Decision rule

Methodologies for training machine learning models to *maximize prescriptive performance*

SLO

Computational advantages and strong performance for well-specified models

Estimate distributions

ILO

Important advantages, but has scalability issues (computationally challenging)

Most works limited to expected value-based models. Exceptions include contextual uncertainty sets (Chenreddy & Delage, 2024, arXiv:2304.04670v1)

Decision rule

Computational advantage at inference/ deployment

WHAT WE DO NOT KNOW – TOPICS THAT NEED MORE RESEARCH

3

ILO for problems subject to endogenous uncertainty

Model **misspecification likely**

- Often **small data setting**
- Demand distributions are endogenous in many problems and hard to estimate

Computational challenges! Optimization problems subject to endogenous uncertainty are often computationally costly to solve (e.g. use of bilevel programming)

ILO for problem formulations with uncertainty in the constraints

Challenging to guarantee feasibility

Thank you

for your attention and for the invitation!

Eric Larsen, Utsav Sadana, Kim Yu helped me improve this presentation. I also benefited from discussions with students and postdocs in my group.

emmafrejinger.org