

What is the Right Notion of Distance between Predict-then-Optimize Tasks?



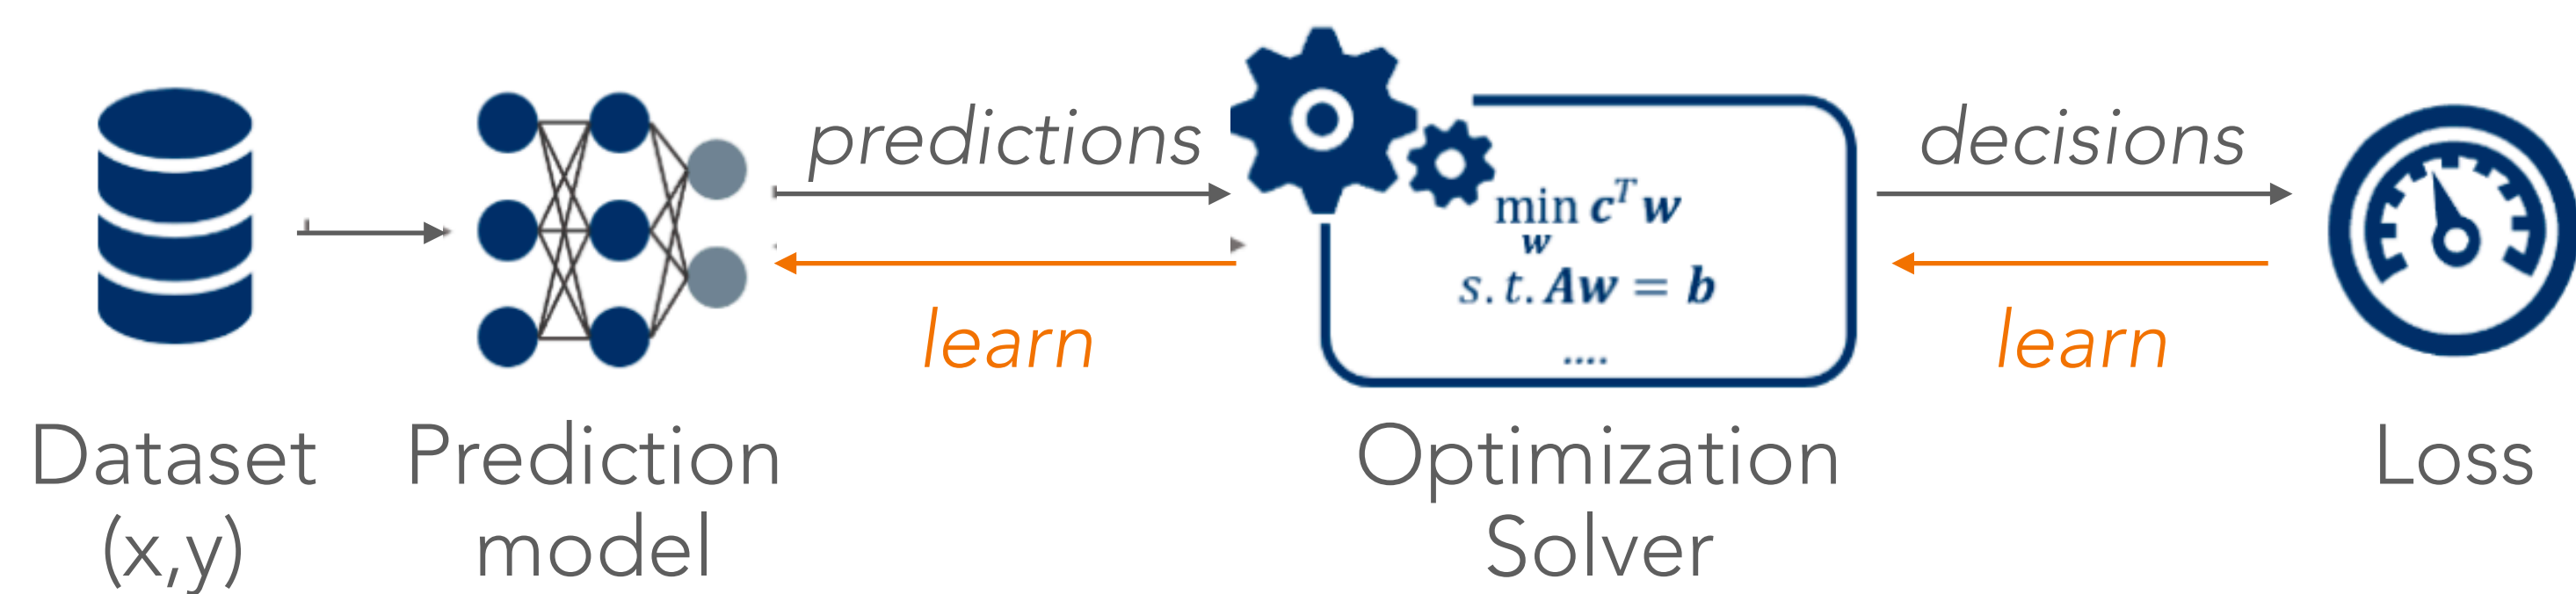
Harvard John A. Paulson School of Engineering and Applied Sciences

Paula Rodriguez-Diaz¹, Lingkai Kong¹, Kai Wang², David Alvarez-Melis^{1,3}, Milind Tambe¹

¹Harvard University ²Georgia Institute of Technology ³Microsoft Research

Predict-then-Optimize (PtO)

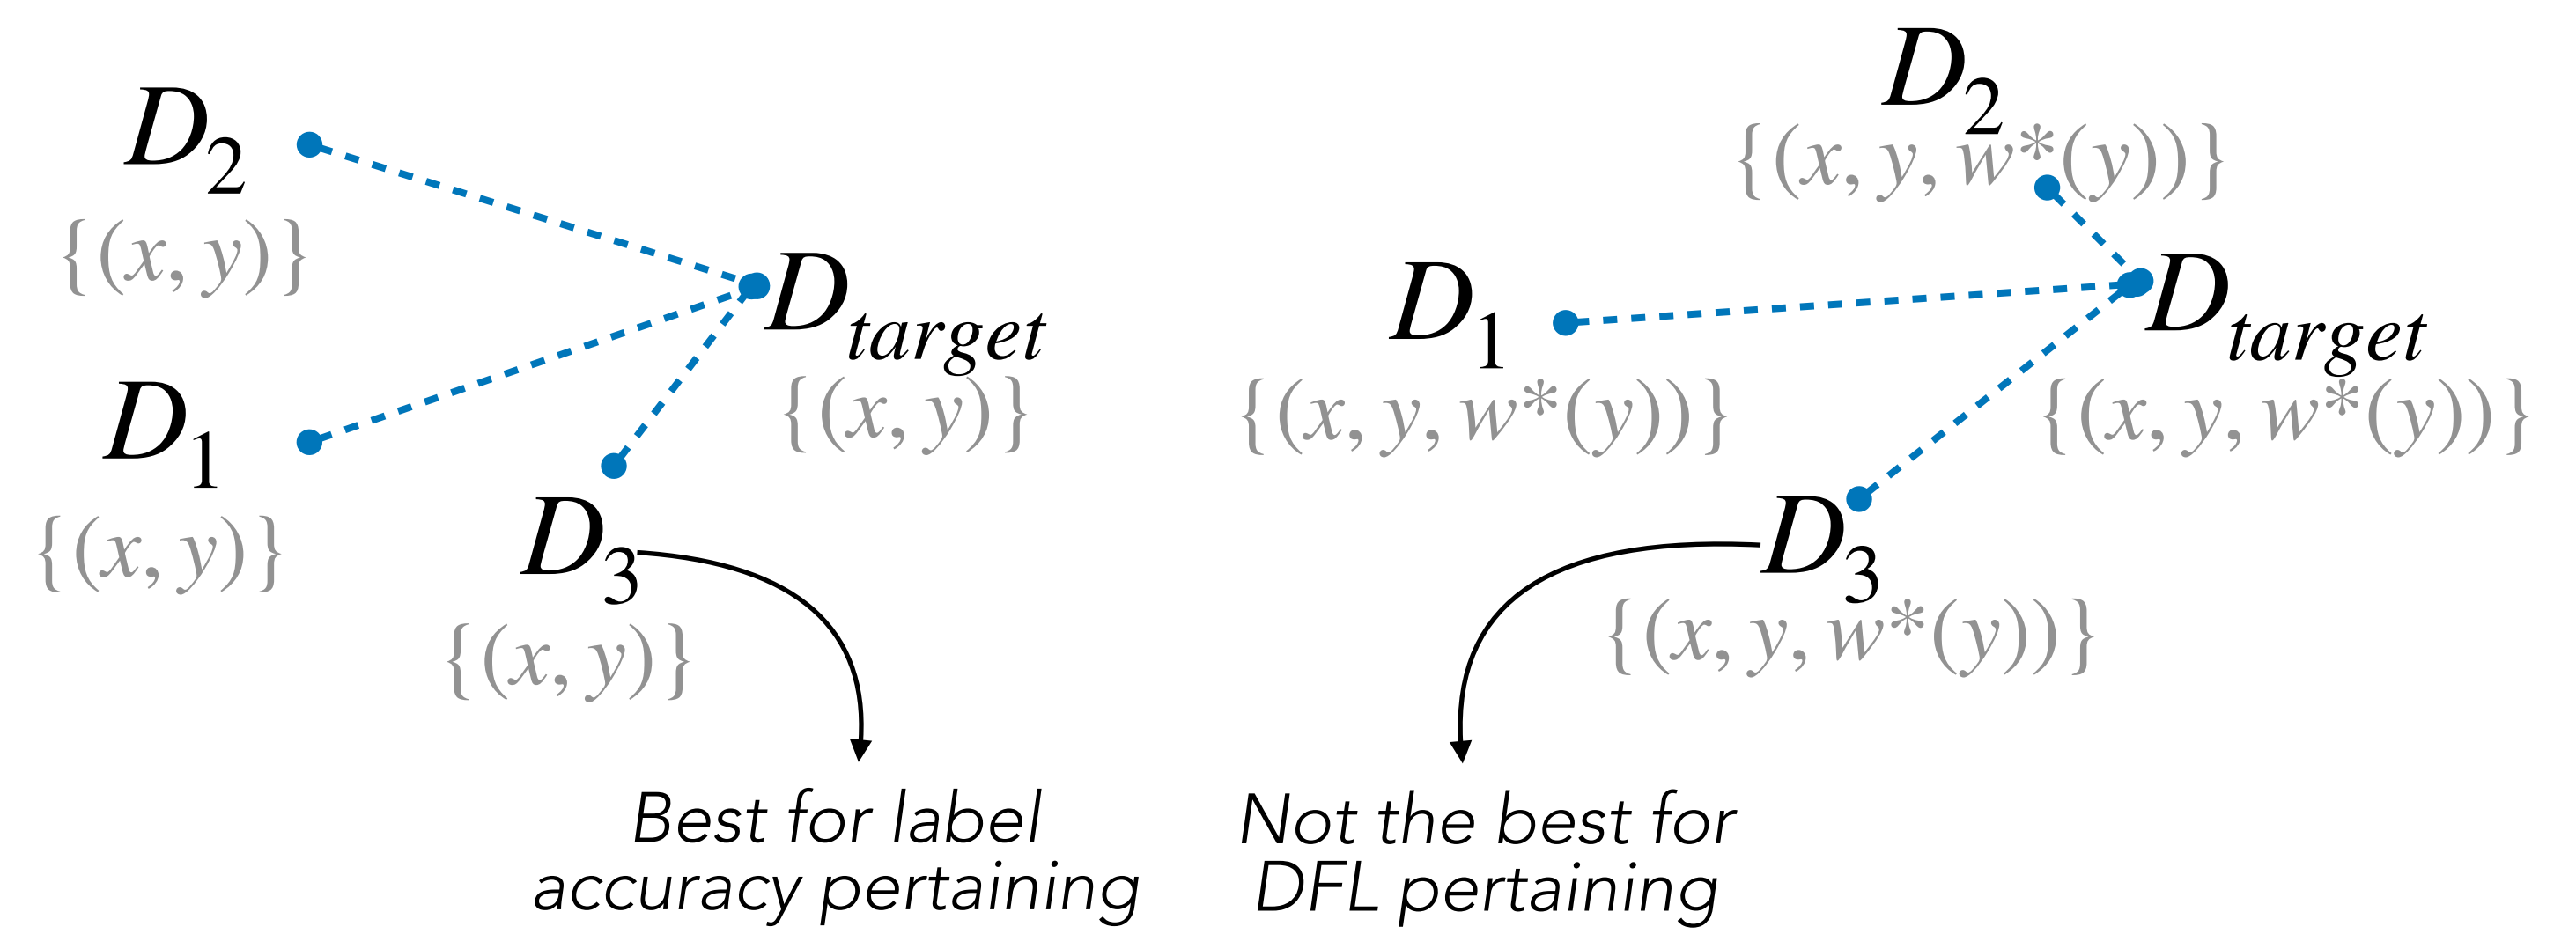
Predictions act as inputs to a downstream optimization, shifting the goal from minimizing prediction error to minimizing **decision regret**



In **decision-focused learning (DFL)**, the loss is decision regret: the difference in objective value between the decision made using the prediction and the optimal decision using the true label.

$$q(\hat{y}, y) = |y \cdot w^*(\hat{y}) - y \cdot w^*(y)|$$

Distances between Datasets



- Enable principled reasoning about task similarity and generalization, especially when data or compute is limited
- Core for transfer learning, multitasks learning, data valuation, etc
- Model agnostic, pre-computed decisions

In decision-focused ML, two tasks that look similar in their feature-label space can lead to very different decisions. **OTD³** is a principled way to measure dataset similarity that accounts for the full feature-label-decision relationship, providing a more reliable signal for transferability when the goal is to learn good decisions—not just good predictions.

How should we compare decisions? We propose *decision quality disparity*: the difference in objective value achieved by two decisions when evaluated on the same true labels.

$$l_q(y, z; y^*, z^*) := |q(y, y^*) - q(z, z^*)|$$

What is the distance between feature-label-decision pairs?

We take a convex combination of the feature-label distance and the disparity in decision quality—used as the cost in the optimal transport formulation.

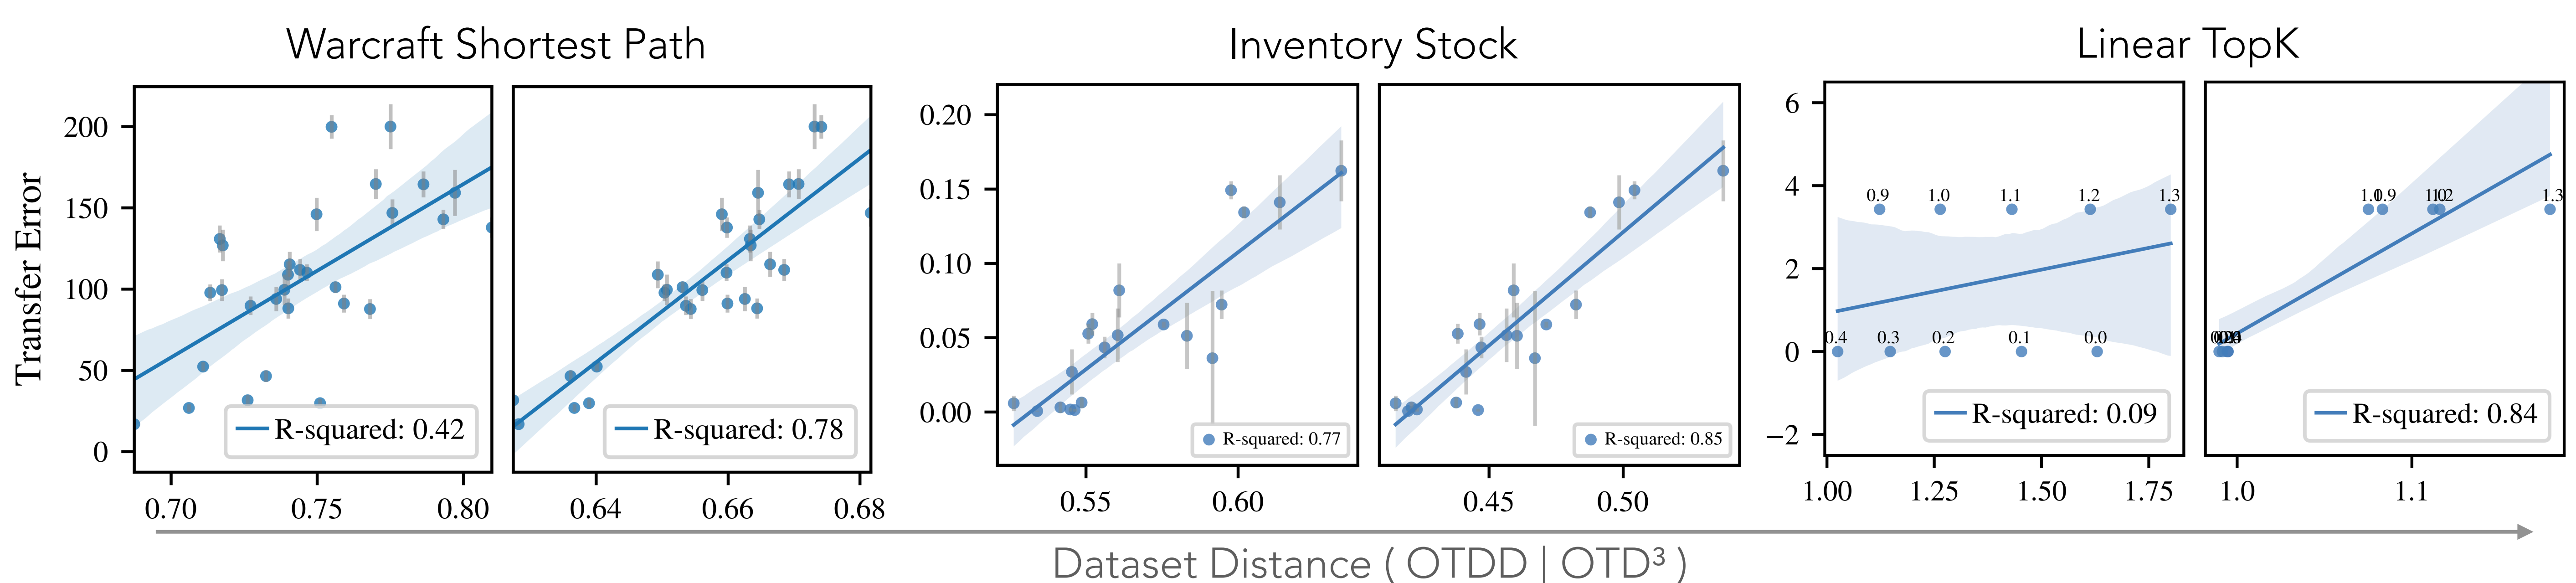
$$c_{PtO}((x, y, z), (x', y', z')) \triangleq \alpha_X \cdot d_X(x, x') + \alpha_Y \cdot d_Y(y, y') + \alpha_W \cdot l_q(z, z'; y', y')$$

$$OTD^3(D, D'; \alpha) \triangleq \min_{\pi \in \Pi(\gamma, \beta)} \int_{W \times W'} c_{PtO}^\alpha(w, w') d\pi$$

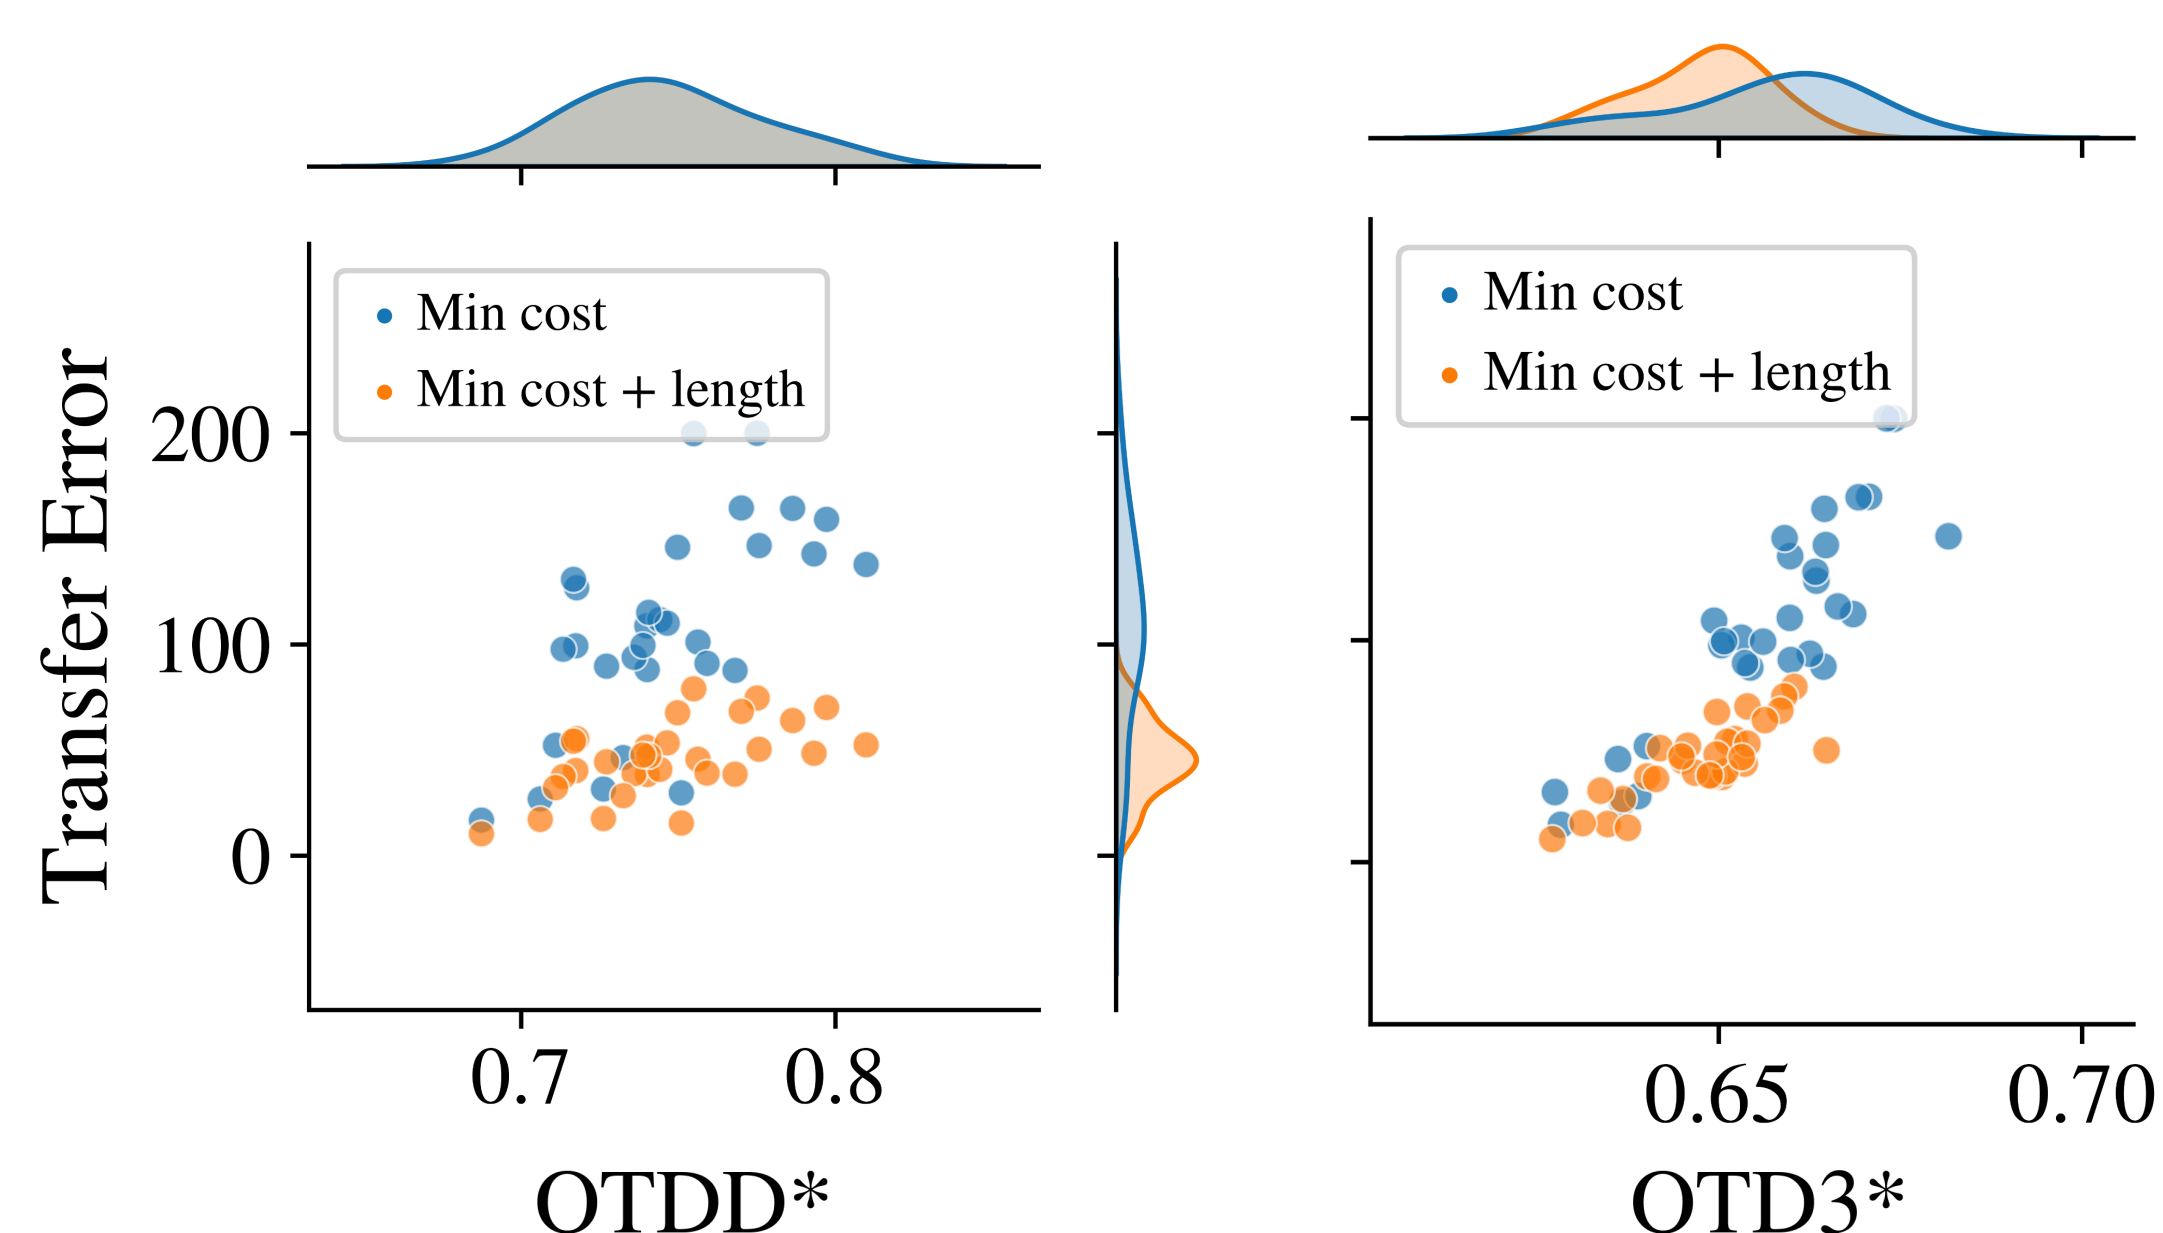
1. A **true metric** in the space of measures of feature-label-decision pairs
2. Provides a **bound on the target regret**

Applications and Experimental Results

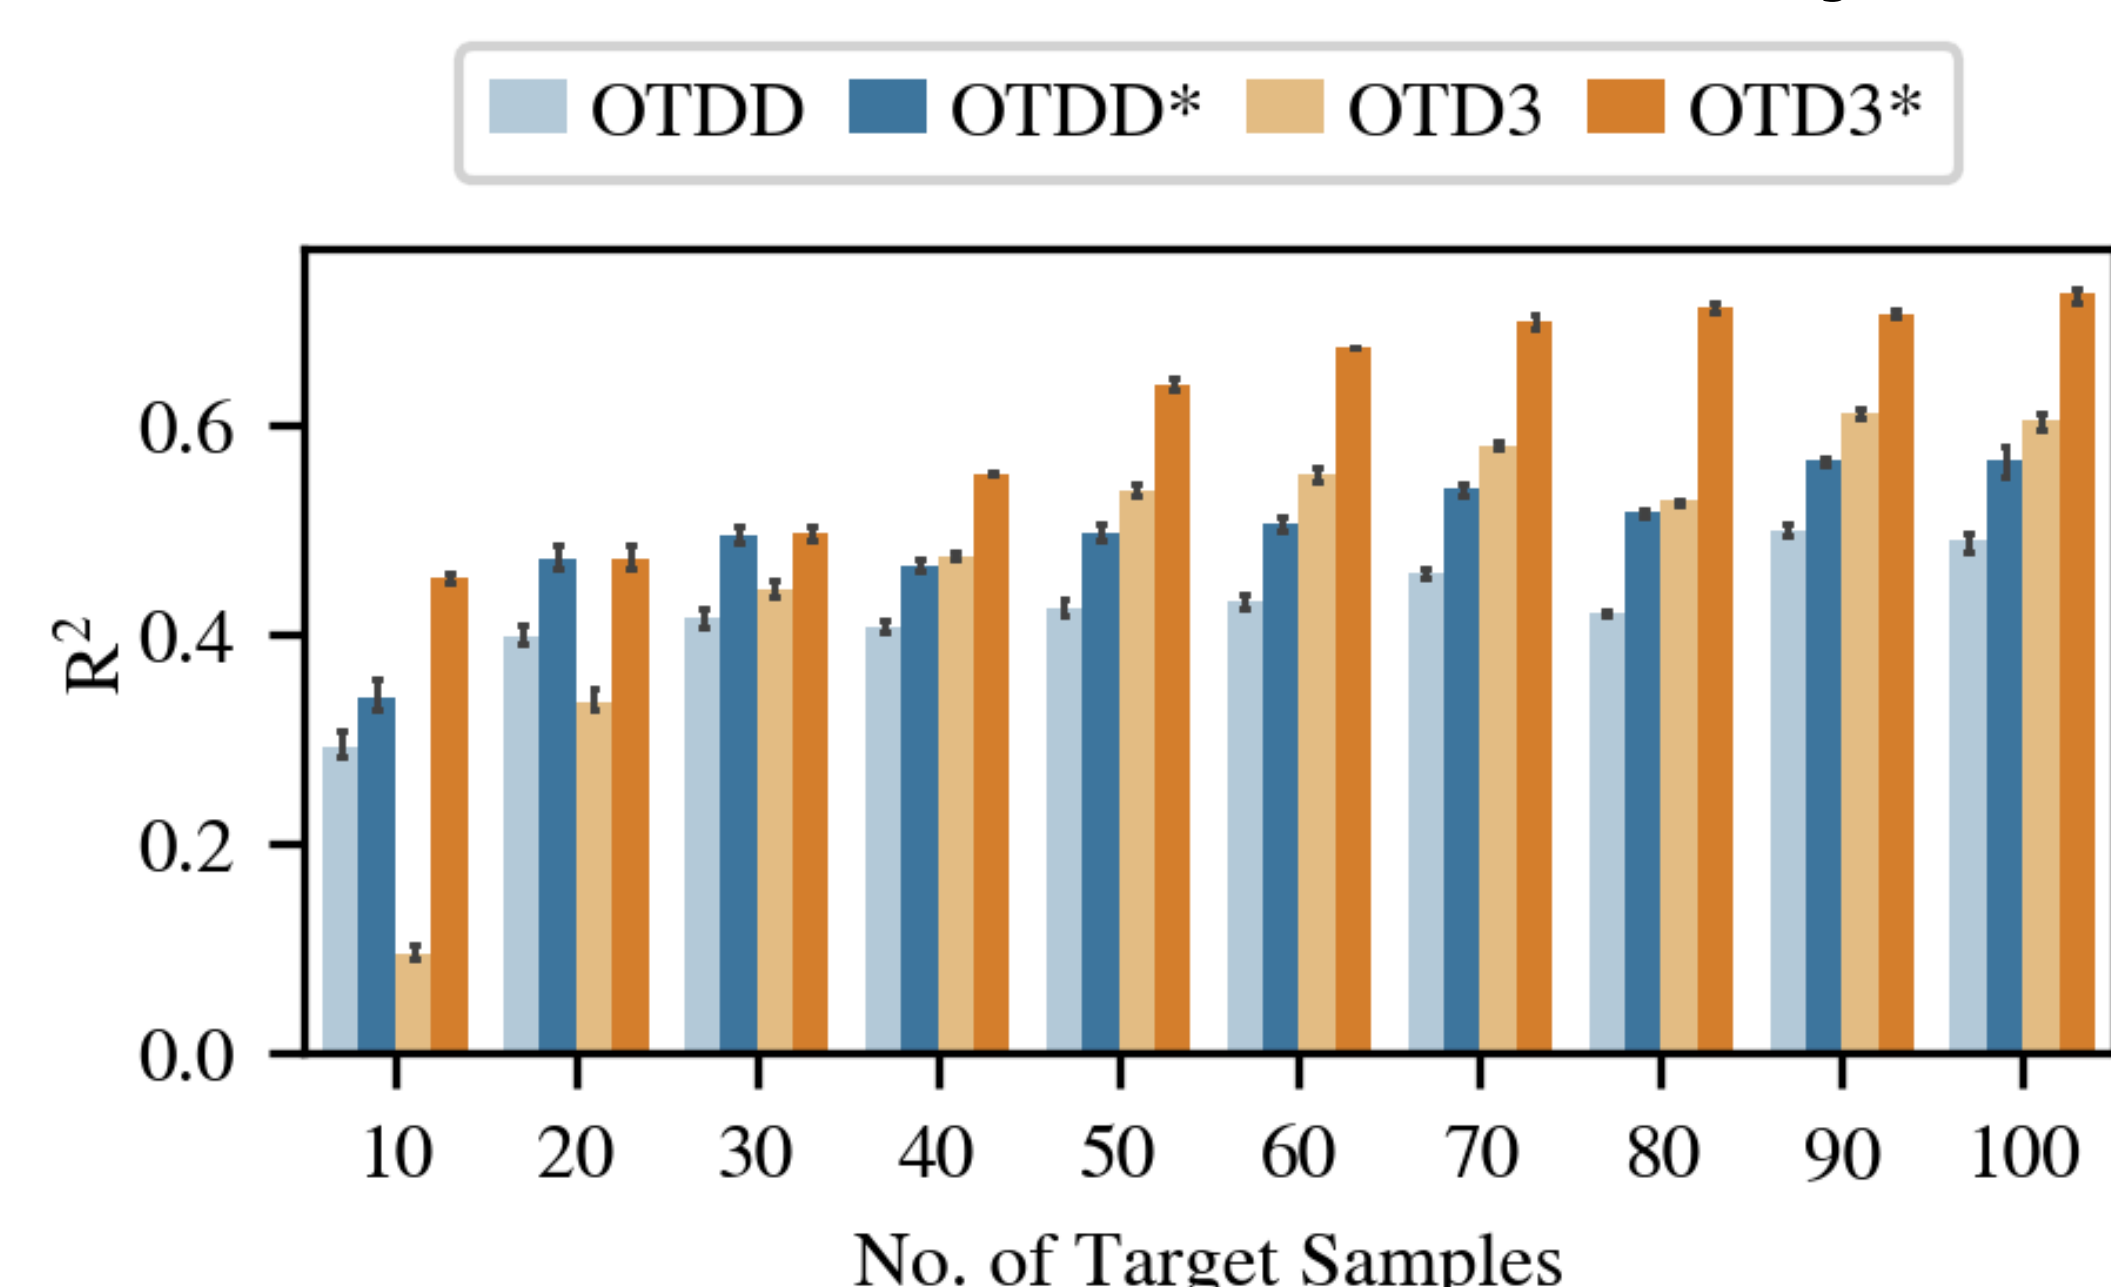
Predict transferability of decision-focused learned models on **shifted domains**



Distinguish transferability success across different downstream tasks



Signal DFL transfer performance under limited data availability



arXiv

