

It's All in the Mix:

Wasserstein Classification and Regression with Mixed Features

(Author names blinded for review process)

Abstract

Problem definition: A key challenge in supervised learning is data scarcity, which can cause prediction models to overfit to the training data and perform poorly out of sample. A contemporary approach to combat overfitting is offered by distributionally robust problem formulations that consider all data-generating distributions close to the empirical distribution derived from historical samples, where ‘closeness’ is determined by the Wasserstein distance. While such formulations show significant promise in prediction tasks where all input features are continuous, they scale exponentially when discrete features are present.

Methodology/results: We demonstrate that distributionally robust mixed-feature classification and regression problems can indeed be solved in polynomial time. Our proof relies on classical ellipsoid method-based solution schemes that do not scale well in practice. To overcome this limitation, we develop a practically efficient (yet, in the worst case, exponential time) cutting plane-based algorithm that admits a polynomial time separation oracle, despite the presence of exponentially many constraints. We compare our method against alternative techniques both theoretically and empirically on standard benchmark instances.

Managerial implications: Data-driven operations management problems often involve prediction models with discrete features. We develop and analyze distributionally robust prediction models that faithfully account for the presence of discrete features, and we demonstrate that our models can significantly outperform existing methods that are agnostic to the presence of discrete features, both theoretically and on standard benchmark instances.

1 Introduction

The recent application of machine learning tools across all areas of operations management has led to a plethora of data-driven and end-to-end approaches that blend predictive models from machine learning with optimization frameworks from operations research and operations management. Notable examples include inventory management (Ban and Rudin, 2019; Bertsimas and Kallus, 2020), logistics (Bertsimas et al., 2019; Behrendt et al., 2023) and supply chain management (Glaeser et al., 2019), assortment optimization (Kallus and Udell, 2020; Feldman et al., 2022) and revenue management (Ferreira et al., 2016; Alley et al., 2023) as well as healthcare operations (Bertsimas et al., 2016; Bastani and Bayati, 2020; Bertsimas and Pauphilet, 2023).

The machine learning algorithms used for prediction are prone to overfitting the available data. Overfitted models perform well on the training data used to calibrate the model, but their performance deteriorates when exposed to new, unseen data. This undesirable effect is amplified if the output of a machine learning model is used as input to a downstream optimization model; this phenomenon is known by different communities as the *Optimizer’s Curse* (Smith and Winkler, 2006) or the *Error-Maximization Effect of Optimization* (Michaud, 1989). Traditionally, overfitting is addressed with regularization techniques that penalize complex models characterized by large and/or dense model parameters (Hastie et al., 2009; Murphy, 2022). A contemporary alternative from the robust optimization community frames machine learning problems as Stackelberg leader-follower games where the learner selects a model that performs best against a worst-case data-generating distribution selected by a conceptual adversary (‘nature’) from a predefined ambiguity set (Ben-Tal et al., 2009; Rahimian and Mehrotra, 2022; Bertsimas and den Hertog, 2022). We talk about Wasserstein machine learning problems when the ambiguity set constitutes a Wasserstein ball centered around the empirical distribution of the available historical observations (Mohajerin Esfahani and Kuhn, 2018; Blanchet and Murthy, 2019; Gao and Kleywegt, 2023). Over the last few years, Wasserstein machine learning problems have

attracted enormous attention in the machine learning and optimization communities; we refer to Kuhn et al. (2019, 2024) for recent reviews of the literature. Interestingly, Wasserstein learning problems admit dual characterizations as regularized learning problems (Shafieezadeh-Abadeh et al., 2015, 2019; Gao et al., 2022), and they thus contribute to a deeper understanding of the impact of regularization in machine learning. We note that other classes of ambiguity sets have been explored as well, such as moment ambiguity sets and those based on ϕ -divergences (such as the Kullback-Leibler divergence). We will not delve into the comparative advantages of different ambiguity sets, and we instead refer the interested reader to the existing literature (see, *e.g.*, Van Parys et al., 2021, Kuhn et al., 2019, 2024 and Lam, 2019).

Although Wasserstein formulations of many classical machine learning tasks admit formulations as convex optimization problems, these formulations scale exponentially in the discrete input features. This limitation has, thus far, confined the use of Wasserstein machine learning models primarily to datasets with exclusively continuous features. This constitutes a major restriction in operations management, where estimation problems frequently include discrete features. Recent examples include Qi et al. (2022), who apply Wasserstein-based quantile regression to a bike sharing inventory management problem characterized by numeric and discrete features (*e.g.*, the weather conditions, the hour of the day as well as the day of the week); Samorani et al. (2022), who study appointment scheduling problems where most features are discrete (*e.g.*, the day of the week, the patient’s marital status and her insurance type); Li et al. (2023), who detect human trafficking from user review websites (here, the discrete features describe the presence or absence of indicative words and phrases); Chan et al. (2023), who predict the macronutrient content of human milk donations using discrete features such as the infant status (term vs preterm); and Duchi et al. (2023), who enforce fairness in offender recidivism prediction through the use of discrete features such as the offender’s race, gender and the existence of prior misdemeanour charges. More broadly, at the time of writing, 240 of the 496 classification and 64

of the 159 regression problems in the popular UCI machine learning repository contain discrete input features (Kelly et al., 2017).

We will demonstrate that naïvely replacing discrete features with unbounded continuous features leads to pathological ambiguity sets in the Wasserstein learning problem whose worst-case distributions lack theoretical appeal and whose resulting prediction models underperform in practice. We also show that replacing discrete features with continuous features that are supported on suitably chosen convex hulls of the discrete supports leads to ambiguity sets that are *equivalent* to those resulting from faithfully accounting for discrete features. Unfortunately, however, tractable reformulations of the Wasserstein learning problem over such bounded continuous features are only available for piece-wise affine convex loss functions, where we will demonstrate empirically that their solution times can be prohibitively long.

This paper studies linear Wasserstein classification and regression problems with mixed (continuous and discrete) features from a theoretical, computational and numerical perspective. Our work makes several contributions to the state-of-the-art:

- (i) From a *theoretical perspective*, we demonstrate that while linear Wasserstein classification and regression with mixed features are inherently NP-hard, a wide range of problems can be solved in polynomial time. Also, contrary to problems with exclusively continuous features, we establish that mixed-feature linear Wasserstein classification and regression do not reduce to regularized problems. On the other hand, we demonstrate that mixed-feature linear Wasserstein classification and regression problems are equivalent to bounded continuous-feature problems with suitably chosen supports.
- (ii) From a *computational perspective*, we propose a cutting plane scheme that solves progressively refined relaxations of the linear Wasserstein classification and regression problems as convex optimization problems. While our overall scheme is not guaranteed to terminate in polynomial time, we show that the key step of our algorithm—the identification of the

most violated constraint—can be implemented efficiently for broad classes of problems, despite the presence of exponentially many constraints.

(iii) From a *numerical perspective*, we show that our cutting plane scheme is substantially faster than an equivalent polynomial-size problem reformulation via bounded continuous features. We also show that our model can perform favorably against classical, regularized and alternative robust problem formulations on standard benchmark instances.

Our paper is most closely related to the recent work of Shafieezadeh-Abadeh et al. (2015) and Shafieezadeh-Abadeh et al. (2019). Shafieezadeh-Abadeh et al. (2015) formulate the Wasserstein logistic regression problem as a convex optimization problem, they discuss the out-of-sample guarantees of their model, and they report numerical results on simulated and benchmark instances. Shafieezadeh-Abadeh et al. (2019) extend their previous work to a wider class of Wasserstein classification and regression problems. Both papers focus on problems with exclusively continuous features, and their proposed formulations would scale exponentially in any discrete features. In contrast, our work studies linear Wasserstein classification and regression problems with mixed features: we examine the theoretical properties of such problems, we develop a practically efficient solution scheme, and we report numerical results. Our work also relates closely to a recent stream of literature that characterizes Wasserstein learning problems as regularized learning problems (Shafieezadeh-Abadeh et al., 2015, 2019; Blanchet et al., 2019; Gao et al., 2022). In particular, we demonstrate that our mixed-feature Wasserstein learning problems do not admit an equivalent representation as regularized learning problems, which forms a notable contrast to the existing findings from the literature.

The present work constitutes a completely revised and substantially expanded version of a conference paper (Redacted, 2022). While that work focuses on logistic regression, the present paper studies broad classes of linear Wasserstein classification and regression problems. This expansion necessitates significant adaptations of the proof for the computational complexity

of the Wasserstein learning problem (*cf.* Theorem 3 in Section 4), an entirely new proof for the absence of regularized problem formulations (*cf.* Theorem 4 in Section 4) that applies to any loss function (as opposed to only the log-loss function in Redacted, 2022) and any non-trivial Wasserstein learning instance, as well as a substantially generalized cutting plane scheme (*cf.* Algorithms 1 and 2 as well as Theorem 2 in Section 3). We also show that the mixed-feature Wasserstein learning problem is equivalent to a bounded continuous-feature formulation, and we present a considerably augmented set of numerical results that encompass both classification and regression problems (*cf.* Section 6).

The remainder of this paper is organized as follows. Section 2 introduces the mixed-feature linear Wasserstein classification and regression problems of interest, and it develops a unified representation of both problems as convex programs of exponential size. Section 3 presents and analyzes our cutting plane solution approach for this exponential-size convex program. Section 4 shows that while mixed-feature linear Wasserstein classification and regression problems are generically NP-hard, important special cases can be solved in polynomial time. We also show that mixed-feature linear Wasserstein classification and regression problems do not reduce to regularized problems. Section 5 contrasts our mixed-feature model against formulations that replace discrete features with (bounded or unbounded) continuous ones. We report on numerical experiments in Section 6, and we offer concluding remarks in Section 7. Exponential-size convex reformulations of the mixed-feature linear Wasserstein classification and regression problems, which are unified by our representation in Section 2, as well as all proofs are relegated to the e-companion. All datasets and source codes accompanying this work are available open source.¹

Notation. We denote by $\mathbb{R}$ ($\mathbb{R}_+$, $\mathbb{R}_-$) the set of (non-negative, non-positive) real numbers, by $\mathbb{N}$ the set of positive integers, and we define $\mathbb{B} = \{0, 1\}$ as well as $[N] = \{1, \dots, N\}$ for $N \in \mathbb{N}$. For a proper cone $\mathcal{C} \subseteq \mathbb{R}^n$, we write $\mathbf{x} \leq_{\mathcal{C}} \mathbf{x}'$ and $\mathbf{x} <_{\mathcal{C}} \mathbf{x}'$ to abbreviate $\mathbf{x}' - \mathbf{x} \in \mathcal{C}$

¹Website: <https://anonymous.4open.science/r/Wasserstein-Mixed-Features-088D/>.

and $\mathbf{x}' - \mathbf{x} \in \text{int } \mathcal{C}$, respectively. The dual norm of $\|\cdot\|$ is $\|\mathbf{x}\|_* = \sup_{\mathbf{x}' \in \mathbb{R}^n} \{\mathbf{x}^\top \mathbf{x}' : \|\mathbf{x}'\| \leq 1\}$, and the cone dual to a cone $\mathcal{C}$ is $\mathcal{C}^* = \{\mathbf{x}' : \mathbf{x}^\top \mathbf{x}' \geq 0 \ \forall \mathbf{x} \in \mathcal{C}\}$. The support function of a set $\mathbb{X} \subseteq \mathbb{R}^n$ is $\mathcal{S}_{\mathbb{X}}(\mathbf{x}) = \sup\{\mathbf{x}^\top \mathbf{x}' : \mathbf{x}' \in \mathbb{X}\}$. For a function $L : \mathbb{X} \rightarrow \mathbb{R}$, we define the Lipschitz modulus as $\text{lip}(L) = \sup\{|L(\mathbf{x}) - L(\mathbf{x}')| / \|\mathbf{x} - \mathbf{x}'\| : \mathbf{x}, \mathbf{x}' \in \mathbb{X}, \mathbf{x} \neq \mathbf{x}'\}$. The set $\mathcal{P}_0(\Xi)$ contains all probability distributions supported on Ξ , and the Dirac distribution $\delta_{\mathbf{x}} \in \mathcal{P}_0(\mathbb{R}^n)$ places unit probability mass on $\mathbf{x} \in \mathbb{R}^n$. The indicator function $\mathbb{1}[\mathcal{E}]$ attains the value 1 (0) whenever the logical expression $\mathcal{E}$ is (not) satisfied.

2 Mixed-Feature Wasserstein Classification and Regression

We study learning problems over N data points $\boldsymbol{\xi}^n = (\mathbf{x}^n, \mathbf{z}^n, y^n) \in \Xi = \mathbb{X} \times \mathbb{Z} \times \mathbb{Y}$, $n \in [N]$, where $\mathbf{x}^n$, $\mathbf{z}^n$ and y^n represent the numerical features, the discrete features and the output variable, respectively. We assume that the support $\mathbb{X}$ of the numerical features is a closed and convex subset of $\mathbb{R}^{M_x}$. The support $\mathbb{Z}$ of the K discrete features satisfies $\mathbb{Z} = \mathbb{Z}(k_1) \times \dots \times \mathbb{Z}(k_K)$, where $k_m \in \mathbb{N} \setminus \{1\}$ denotes the number of values that the m -th discrete feature can attain, $m \in [K]$, and $\mathbb{Z}(s) = \{\mathbf{z} \in \mathbb{B}^{s-1} : \sum_{i \in [s-1]} z_i \leq 1\}$ is the one-hot feature encoding. We let $M_z = \sum_{m \in [K]} (k_m - 1)$ denote the number of coefficients associated with the discrete features. The support $\mathbb{Y}$ of the output variable is $\{-1, +1\}$ for classification and a closed and convex subset of $\mathbb{R}$ for regression problems, respectively. We wish to solve the Wasserstein learning problem

$$\begin{aligned} & \underset{\boldsymbol{\beta}}{\text{minimize}} && \sup_{\mathbb{Q} \in \mathfrak{B}_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}} [l_{\boldsymbol{\beta}}(\mathbf{x}, \mathbf{z}, y)] \\ & \text{subject to} && \boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}_x, \boldsymbol{\beta}_z) \in \mathbb{R}^{1+M_x+M_z}, \end{aligned} \tag{1}$$

where the ambiguity set $\mathfrak{B}_\epsilon(\hat{\mathbb{P}}_N) = \{\mathbb{Q} \in \mathcal{P}_0(\Xi) : W(\mathbb{Q}, \hat{\mathbb{P}}_N) \leq \epsilon\}$ represents the Wasserstein ball of radius $\epsilon > 0$ that is centered at the empirical distribution $\hat{\mathbb{P}}_N = \frac{1}{N} \sum_{n \in [N]} \delta_{\boldsymbol{\xi}^n}$ placing equal probability mass on the N data points $\boldsymbol{\xi}^n$, $n \in [N]$, as per the following definition.

Definition 1 (Wasserstein Distance). *The type-1 Wasserstein (Kantorovich-Rubinstein, or earth mover's) distance between two distributions $\mathbb{P} \in \mathcal{P}_0(\Xi)$ and $\mathbb{Q} \in \mathcal{P}_0(\Xi)$ is defined as*

$$W(\mathbb{P}, \mathbb{Q}) := \inf_{\Pi \in \mathcal{P}_0(\Xi^2)} \left\{ \int_{\Xi^2} d(\boldsymbol{\xi}, \boldsymbol{\xi}') \Pi(d\boldsymbol{\xi}, d\boldsymbol{\xi}') : \Pi(d\boldsymbol{\xi}, \Xi) = \mathbb{P}(d\boldsymbol{\xi}), \Pi(\Xi, d\boldsymbol{\xi}') = \mathbb{Q}(d\boldsymbol{\xi}') \right\},$$

where the ground metric d on Ξ satisfies

$$d(\boldsymbol{\xi}, \boldsymbol{\xi}') = \|\mathbf{x} - \mathbf{x}'\| + \kappa_z d_z(\mathbf{z}, \mathbf{z}') + \kappa_y d_y(y, y') \quad \forall \boldsymbol{\xi} = (\mathbf{x}, \mathbf{z}, y) \in \Xi, \boldsymbol{\xi}' = (\mathbf{x}', \mathbf{z}', y') \in \Xi \quad (2a)$$

with $\kappa_z, \kappa_y > 0$ as well as, for some $p \geq 1$,

$$d_z(\mathbf{z}, \mathbf{z}') = \left(\sum_{m \in [K]} \mathbb{1}[\mathbf{z}_m \neq \mathbf{z}'_m] \right)^{1/p} \quad \text{and} \quad d_y(y, y') = \begin{cases} \mathbb{1}[y \neq y'] & \text{if } \mathbb{Y} = \{-1, +1\}, \\ |y - y'| & \text{otherwise.} \end{cases} \quad (2b)$$

An important special case of Definition 1 arises when $\kappa_y = \infty$. In this setting, any transport between samples with different outputs incurs infinite cost, and hence the ambiguity set only contains distributions that preserve the empirical output values. In other words, $\mathfrak{B}_\epsilon(\hat{\mathbb{P}}_N)$ then consists of perturbations of the feature distribution conditional on y , while the marginal distribution of y remains fixed.

The loss function $l_{\boldsymbol{\beta}}(\mathbf{x}, \mathbf{z}, y) : \mathbb{X} \times \mathbb{Z} \times \mathbb{Y} \rightarrow \mathbb{R}_+$ in problem (1) satisfies

$$l_{\boldsymbol{\beta}}(\mathbf{x}, \mathbf{z}, y) = \begin{cases} L(y \cdot [\beta_0 + \boldsymbol{\beta}_x^\top \mathbf{x} + \boldsymbol{\beta}_z^\top \mathbf{z}]) & \text{if } \mathbb{Y} = \{-1, +1\}, \\ L(\beta_0 + \boldsymbol{\beta}_x^\top \mathbf{x} + \boldsymbol{\beta}_z^\top \mathbf{z} - y) & \text{otherwise,} \end{cases}$$

where $L : \mathbb{R} \rightarrow \mathbb{R}_+$ measures the similarity between the prediction $\beta_0 + \boldsymbol{\beta}_x^\top \mathbf{x} + \boldsymbol{\beta}_z^\top \mathbf{z}$ and the output y . With a slight abuse of terminology, we also refer to L as the loss function whenever the context is clear. For both classification and regression problems, we consider two settings:

- (i) L is convex and Lipschitz continuous with Lipschitz modulus $\text{lip}(L)$, $\mathbb{X} = \mathbb{R}^{M_x}$ and $\mathbb{Y} = \{-1, +1\}$ (for classification problems) or $\mathbb{Y} = \mathbb{R}$ (for regression problems);
- (ii) L satisfies $L(e) = \max_{j \in [J]} \{a_j e + b_j\}$, and $\mathbb{X} \subseteq \mathbb{R}^{M_x}$ and $\mathbb{Y} \subseteq \mathbb{R}$ are closed and convex.

In either case, we assume that L is not constant.

The Wasserstein learning problem (1) offers attractive generalization guarantees. While the classical choice of Wasserstein radii suffers from the curse of dimensionality (Mohajerin Esfahani and Kuhn, 2018), recent work has developed asymptotic (Blanchet et al., 2019; Blanchet and Kang, 2021) as well as finite sample guarantees (Shafieezadeh-Abadeh et al., 2019; Gao, 2022) that apply to Wasserstein radii of the order $\mathcal{O}(1/\sqrt{N})$.

We next review a result that expresses the Wasserstein learning problem (1) as a convex optimization problem.

Observation 1. *The Wasserstein learning problem (1) admits the equivalent formulation*

$$\begin{aligned}
& \underset{\boldsymbol{\beta}, \lambda, \mathbf{s}}{\text{minimize}} && \lambda \epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\
& \text{subject to} && \sup_{(\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y}} \{l_{\boldsymbol{\beta}}(\mathbf{x}, \mathbf{z}, y) - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n)\} - \lambda \kappa_z d_z(\mathbf{z}, \mathbf{z}^n) \leq s_n \quad (3) \\
& && \forall n \in [N], \forall \mathbf{z} \in \mathbb{Z} \\
& && \boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}_x, \boldsymbol{\beta}_z) \in \mathbb{R}^{1+M_x+M_z}, \quad \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N.
\end{aligned}$$

When $\kappa_y = \infty$, the semi-infinite constraint in Observation 1 simplifies to

$$\sup_{\mathbf{x} \in \mathbb{X}} \{l_{\boldsymbol{\beta}}(\mathbf{x}, \mathbf{z}, y^n) - \lambda \|\mathbf{x} - \mathbf{x}^n\|\} - \lambda \kappa_z d_z(\mathbf{z}, \mathbf{z}^n) \leq s_n \quad \forall n \in [N], \forall \mathbf{z} \in \mathbb{Z}.$$

Problem (3) contains embedded maximization problems, and it comprises exponentially many constraints. We next prove the existence of equivalent reformulations of (3) for classification

and regression problems, respectively, that do not contain embedded optimization problems. Since the resulting reformulations exhibit an exponential number of constraints, the next section will develop a cutting plane approach to introduce these constraints iteratively.

Our equivalent reformulations of (3) leverage the unified problem representation

$$\begin{aligned}
& \underset{\boldsymbol{\theta}, \boldsymbol{\sigma}}{\text{minimize}} && f_0(\boldsymbol{\theta}, \boldsymbol{\sigma}) \\
& \text{subject to} && f_{ni}(g_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)) \leq \sigma_n \quad \forall n \in [N], \forall i \in \mathcal{I}, \forall \mathbf{z} \in \mathbb{Z} \\
& && \boldsymbol{\theta} \in \Theta, \quad \boldsymbol{\sigma} \in \mathbb{R}_+^N,
\end{aligned} \tag{4}$$

where $\boldsymbol{\xi}_{-\mathbf{z}}^n = (\mathbf{x}^n, y^n)$ and $\mathcal{I}$ is a finite index set. To ensure that (4) is convex, we stipulate that $f_0 : \mathbb{R}^{M_\theta} \times \mathbb{R}_+^N \rightarrow \mathbb{R}$ and $f_{ni} : \mathbb{R} \rightarrow \mathbb{R}$ are convex, $g_{ni} : \mathbb{R}^{M_\theta} \times (\mathbb{X} \times \mathbb{Y}) \times \mathbb{R}^{M_z} \rightarrow \mathbb{R}$ is bi-affine in $\boldsymbol{\theta}$ and $\mathbf{z}$ for every fixed $\boldsymbol{\xi}_{-\mathbf{z}}^n$, $h_{ni} : \mathbb{R}^{M_\theta} \times (\mathbb{X} \times \mathbb{Y}) \times \mathbb{R} \rightarrow \mathbb{R}$ is concave in $\boldsymbol{\theta}$ for every fixed $\boldsymbol{\xi}_{-\mathbf{z}}^n$ and every fixed value of its last component, $n \in [N]$ and $i \in \mathcal{I}$, and $\Theta \subseteq \mathbb{R}^{M_\theta}$ is a convex set. We further assume that f_0 is non-decreasing in $\boldsymbol{\sigma}$ so that (4) is not unbounded.

Theorem 1. *The following problems admit equivalent reformulations in the form of (4):*

- (i) *the mixed-feature Wasserstein classification and regression problems with convex and Lipschitz continuous loss functions L as well as the continuous feature support $\mathbb{X} = \mathbb{R}^{M_x}$.*
- (ii) *the mixed-feature Wasserstein classification and regression problems with piece-wise affine convex loss functions $L(e) = \max_{j \in [J]} \{a_j e + b_j\}$ as well as the continuous feature support*

$$\begin{aligned}
\mathbb{X} &= \{\mathbf{x} \in \mathbb{R}^{M_x} : \mathbf{C}\mathbf{x} \leq_{\mathcal{C}} \mathbf{d}\} \text{ for some } \mathbf{C} \in \mathbb{R}^{r \times M_x}, \mathbf{d} \in \mathbb{R}^r \\
&\text{and proper convex cone } \mathcal{C} \subseteq \mathbb{R}^r
\end{aligned}$$

for classification problems as well as the support

$$\mathbb{X} \times \mathbb{Y} = \{(\mathbf{x}, y) \in \mathbb{R}^{M_x+1} : \mathbf{C}_x \mathbf{x} + \mathbf{c}_y \cdot y \leq_{\mathcal{C}} \mathbf{d}\} \text{ for some } \mathbf{C}_x \in \mathbb{R}^{r \times M_x}, \mathbf{c}_y \in \mathbb{R}^r, \mathbf{d} \in \mathbb{R}^r$$

and proper convex cone $\mathcal{C} \subseteq \mathbb{R}^r$

for regression problems, assuming that the supports admit Slater points.

The Appendices A and B in the e-companion prove Theorem 1 for classification and regression problems, respectively. The proofs follow similar arguments as those of Shafieezadeh-Abadeh et al. (2019), adapted to the presence of discrete features as well as our ground metric.

3 Cutting Plane Solution Scheme

The unified problem representation (4) of the mixed-feature Wasserstein regression and classification problems (*cf.* Theorem 1) is convex, but it comprises constraints whose number scales exponentially in K , the number of discrete features. We therefore develop a cutting plane approach that iteratively introduces only those constraints that are most violated by a sequence of incumbent solutions.

Algorithm 1 outlines our cutting plane scheme, which follows the standard methodology of this class of techniques. The algorithm iteratively solves a sequence of relaxations of problem (4), where each relaxation only includes a subset $\mathcal{W}$ of the constraints indexed by $[N] \times \mathcal{I} \times \mathbb{Z}$. Consequently, any optimal solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ obtained from the relaxation gives rise to a valid lower bound for the optimal value of problem (4). At each iteration, the algorithm identifies and incorporates the most violated constraint from problem (4) with respect to the current solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$. This procedure ensures the generation of a monotonically non-decreasing sequence of lower bounds, which is guaranteed to converge to the optimal value of problem (4) in a finite number of iterations. A common challenge in cutting plane techniques is the derivation of

Algorithm 1 Cutting Plane Scheme for Problem (4)

Input: (Possibly empty) initial constraint set $\mathcal{W} \subseteq [N] \times \mathcal{I} \times \mathbb{Z}$.

Output: Optimal solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ to problem (4).

Initialize $(\text{LB}, \text{UB}) = (-\infty, +\infty)$ as lower and upper bounds for problem (4).

while $\text{LB} < \text{UB}$ **do**

 Let $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ be optimal in the relaxation of (4) involving the constraints $(n, i, \mathbf{z}) \in \mathcal{W}$.

for $n \in [N]$ **do**

 Identify, for each $i \in \mathcal{I}$, a most violated constraint

$$\mathbf{z}(n, i) \in \arg \max_{\mathbf{z} \in \mathbb{Z}} \{f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n)) - \sigma_n^*\}$$

 associated with $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ and (n, i) , and denote the constraint violation by $\vartheta(n, i)$.

 Let $i(n) \in \arg \max\{\vartheta(n, i) : i \in \mathcal{I}\}$ and add $(n, i(n), \mathbf{z}(n, i(n)))$ to $\mathcal{W}$ if $\vartheta(n, i(n)) > 0$.

end for

 Define $\boldsymbol{\vartheta}^* \in \mathbb{R}^N$ via $\vartheta_n^* = \max\{\vartheta(n, i(n)), 0\}$, $n \in [N]$.

 Update $\text{LB} = f_0(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ and $\text{UB} = \min\{\text{UB}, f_0(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^* + \boldsymbol{\vartheta}^*)\}$.

end while

upper bounds, which allow us to quantify the optimality gap, particularly when the algorithm is terminated prematurely, such as when reaching a predefined time limit. In our case, the special structure of problem (4) enables us to obtain such upper bounds efficiently: We adjust the optimal solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ of the relaxation by increasing the slack variables $\boldsymbol{\sigma}^*$ so as to account for the maximum constraint violations $\boldsymbol{\vartheta}^*$.

The next result formalizes our intuition.

Proposition 1. *Algorithm 1 terminates in finite time with an optimal solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ to problem (4). Moreover, LB and UB constitute monotonic sequences of lower and upper bounds on the optimal value of (4) throughout the execution of the algorithm.*

A key step in Algorithm 1 concerns the identification of a constraint $(n, i, \mathbf{z}) \in [N] \times \mathcal{I} \times \mathbb{Z}$ that is most violated by the incumbent solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$. We next argue that the underlying combinatorial problem can be solved efficiently by Algorithm 2.

Algorithm 2 determines a most violated constraint index, denoted as $\mathbf{z}(n, i)$, for a given

Algorithm 2 Identification of a Most Violated Constraint in Problem (4)

Input: Incumbent solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ and constraint group index $(n, i) \in [N] \times \mathcal{I}$.

Output: A most violated constraint index $\mathbf{z}(n, i)$ in constraint group (n, i) .

Initialize the candidate constraint index set $\mathcal{Z} = \emptyset$.

Let $(\mathbf{w}, w_0) \in \mathbb{R}^{M_z} \times \mathbb{R}$ be such that $g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}) = \mathbf{w}^\top \mathbf{z} + w_0$.

for $\mu \in \{\pm 1\}$ **do**

 Compute $\mathbf{z}_m^* \in \arg \max \{\mu \cdot \mathbf{w}_m^\top \mathbf{z}_m : \mathbf{z}_m \in \mathbb{Z}(k_m) \setminus \{\mathbf{z}_m^n\}\}$ for all $m \in [K]$.

 Compute a permutation $\pi : [K] \rightarrow [K]$ such that

$$\mu \cdot \mathbf{w}_{\pi(m)}^\top (\mathbf{z}_{\pi(m)}^* - \mathbf{z}_{\pi(m)}^n) \geq \mu \cdot \mathbf{w}_{\pi(m')}^\top (\mathbf{z}_{\pi(m')}^* - \mathbf{z}_{\pi(m')}^n) \quad \forall 1 \leq m \leq m' \leq K.$$

for $\delta \in [K] \cup \{0\}$ **do**

 Add $\mathbf{z}(\delta)$ to $\mathcal{Z}$, where $\mathbf{z}_m(\delta) = \mathbf{z}_m^*$ if $\pi(m) \leq \delta$; $= \mathbf{z}_m^n$ otherwise, $m \in [K]$.

end for

end for

Select $\mathbf{z}(n, i) \in \arg \max \{f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)) - \sigma_n^* : \mathbf{z} \in \mathcal{Z}\}$.

incumbent solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ and a given constraint group $(n, i) \in [N] \times \mathcal{I}$ in the optimization problem (4). We obtain such a constraint by maximizing the constraint left-hand side in (4),

$$f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)),$$

over $\mathbf{z} \in \mathbb{Z}$. Note that the constraint function comprises two distinct terms. The second term, $h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n))$, depends on $\mathbf{z}$ only through the distance metric $d_z(\mathbf{z}, \mathbf{z}^n)$ within the ground metric d (cf. Definition 1). Consequently, we can decompose the maximization of the difference of both terms into two steps. In a first step, we maximize the first term $f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}))$ over all $\mathbf{z} \in \mathbb{Z}$ that satisfy $d_z(\mathbf{z}, \mathbf{z}^n) = \delta$ for any fixed $\delta \in [K] \cup \{0\}$. This is what the inner for-loop in Algorithm 2 accomplishes, and it results in $K + 1$ candidate solutions $\mathbf{z}(\delta)$. In a second step, we can then select an optimal solution $\mathbf{z}(\delta^*)$ that maximizes

$$f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}(\delta))) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \delta)$$

across all $\delta \in [K] \cup \{0\}$ by evaluating the constraint left-hand side for the $K + 1$ candidate solutions $\mathbf{z}(\delta)$. This is what the final step in Algorithm 2 implements.

To maximize $f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}))$ over all $\mathbf{z} \in \mathbb{Z}$ satisfying $d_z(\mathbf{z}, \mathbf{z}^n) = \delta$, note first that f_{ni} is assumed to be univariate and convex. Thus, its maximum over any compact set $\mathcal{X} \subseteq \mathbb{R}$ is attained either at $\min\{x : x \in \mathcal{X}\}$ or $\max\{x : x \in \mathcal{X}\}$. Hence, to maximize $f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}))$ over all $\mathbf{z} \in \mathbb{Z}$ with $d_z(\mathbf{z}, \mathbf{z}^n) = \delta$, we only need to maximize and minimize the inner function $g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})$ over all $\mathbf{z} \in \mathbb{Z}$ with $d_z(\mathbf{z}, \mathbf{z}^n) = \delta$. The maximization and minimization of g_{ni} is unified by the outer for-loop in Algorithm 2, which maximizes $\mu \cdot g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})$ for $\mu \in \{\pm 1\}$.

We claim that the optimization of $g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})$ over $\mathbf{z} \in \mathbb{Z}$ with $d_z(\mathbf{z}, \mathbf{z}^n) = \delta$ can be decomposed into separate optimizations over $\mathbf{z}_m \in \mathbb{Z}(k_m)$, $m \in [K]$. Indeed, note that g_{ni} is affine in $\mathbf{z}$ for fixed $\boldsymbol{\theta}^*$, which allows us to express the mapping $\mathbf{z} \mapsto g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})$ as $\mathbf{w}^\top \mathbf{z} + w_0$ for some $\mathbf{w} \in \mathbb{R}^{M_z}$ and $w_0 \in \mathbb{R}$. Moreover, the distance metric $d_z(\mathbf{z}, \mathbf{z}^n)$ only counts the number of subvectors $m \in [K]$ where $\mathbf{z}_m$ differs from $\mathbf{z}_m^n$, without considering the magnitude of the difference (as both vectors are one-hot encoded). Finally, we recall that $\mathbb{Z}$ is rectangular, that is, $\mathbb{Z} = \mathbb{Z}(k_1) \times \dots \times \mathbb{Z}(k_K)$. Taken together, the aforementioned three observations imply that the optimization of $g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})$ over $\mathbf{z} \in \mathbb{Z}$ with $d_z(\mathbf{z}, \mathbf{z}^n) = \delta$ can be broken into two interconnected steps: We need to select the δ subvectors $m_1 < \dots < m_\delta$ where $\mathbf{z}$ differs from $\mathbf{z}^n$, and for each differing subvector m , we need to select a solution $\mathbf{z}_m^*$ that optimizes $\mathbf{w}_m^\top (\mathbf{z}_m - \mathbf{z}_m^n)$ across all $\mathbf{z}_m \neq \mathbf{z}_m^n$. To this end, Algorithm 2 first computes for every $m \in [K]$ a solution $\mathbf{z}_m^*$ that optimizes $\mathbf{w}_m^\top (\mathbf{z}_m - \mathbf{z}_m^n)$ over $\mathbf{z}_m \neq \mathbf{z}_m^n$ (first step inside the outer for-loop), it subsequently sorts the contributions $\mathbf{w}_m^\top (\mathbf{z}_m^* - \mathbf{z}_m^n)$, $m \in [K]$ (second step inside the outer for-loop), and it finally selects the δ largest (smallest) contributions to maximize (minimize) g_{ni} (inner for-loop).

We formalize the above intuition in the next result.

Theorem 2. *For a given incumbent solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ and a constraint group index $(n, i) \in [N] \times \mathcal{I}$ in problem (4), Algorithm 2 identifies a most violated constraint index $\mathbf{z}(n, i)$ in time*

$\mathcal{O}(M_{\mathbf{z}} + KT + K \log K)$, where T is the time required to compute f_{ni} , g_{ni} and h_{ni} .

4 Complexity Analysis

Despite its exponential size, our unified problem representation (4) admits a polynomial time solution for the classes of loss functions that we consider in this paper. In fact, it follows from Theorem 2 that Algorithm 2 provides a polynomial time separation oracle for problem (4).

Theorem 3 (Complexity of the Unified Problem Representation (4)). *Assume that the subgradients of the functions f_{ni} and $-h_{ni}$ in problem (4) can be computed in polynomial time and that Θ is compact, full-dimensional and admits a polynomial time weak separation oracle. Then problem (4) can be solved to δ -accuracy in polynomial time.*

Recall that an optimization problem is solved to δ -accuracy if a δ -suboptimal solution is identified that satisfies all constraints modulo a violation of at most δ . The consideration of δ -accurate solutions is standard in the numerical solution of nonlinear programs where an optimal solution may be irrational. We show in the Appendices A and B that the subgradients of f_{ni} and $-h_{ni}$ can be computed in polynomial time and that Θ can be assumed to be compact in the Wasserstein classification and regression problems that we consider. The regularity assumptions imposed on problem (4) are crucial in this regard. In fact, it follows from Theorem 2 of Redacted (2022) that if the functions f_{ni} in problem (4) were allowed to be non-convex, then the problem would be strongly NP-hard even if $M_{\boldsymbol{\theta}} = 0$, $N = |\mathcal{I}| = 1$ and h_{11} vanished.

While Theorem 3 principally allows us to solve problem (4) in polynomial time, the available solution algorithms (see, *e.g.*, Lee et al. 2015) rely on variants of the ellipsoid method and are not competitive on practical problem instances. For this reason, we rely on our cutting plane algorithm throughout our numerical experiments.

A by now classical result shows that when $K = 0$ (absence of discrete features) and $\mathbb{X} = \mathbb{R}^{M_{\mathbf{x}}}$,

the Wasserstein learning problem (1) reduces to a classical learning problem with an additional regularization term in the objective function whenever the output weight κ_y in Definition 1 approaches ∞ (Shafieezadeh-Abadeh et al., 2015, 2019; Blanchet et al., 2019; Gao et al., 2022). It turns out that this reduction no longer holds when discrete features are present.

Theorem 4 (Absence of Regularizers). *Fix any convex Lipschitz continuous or piece-wise affine convex loss function L that is not constant, any Wasserstein classification or regression instance (cf. Theorem 1 as well as Appendices A and B) with any number $N \geq 1$ training samples, non-zero numbers M_x of continuous and K of discrete features, $\mathbb{X} = \mathbb{R}^{M_x}$, as well as any $\kappa_z > 0$ and any $p \geq 1$ in the definition of the Wasserstein ground metric. For sufficiently large radii ϵ of the ambiguity set, the objective function*

$$\sup_{\mathbb{Q} \in \mathfrak{B}_\epsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{Q}} [l_{\beta}(\mathbf{x}, \mathbf{z}, y)],$$

of the Wasserstein learning problem does not admit an equivalent reformulation as a classical regularized learning problem,

$$\mathbb{E}_{\hat{\mathbb{P}}_N} [l_{\beta}(\mathbf{x}, \mathbf{z}, y)] + \mathfrak{R}(\beta) \quad \text{for any } \mathfrak{R} : \mathbb{R}^{1+M_x+M_z} \rightarrow \mathbb{R},$$

even when the weight κ_y of the output distance d_y approaches ∞ .

We emphasize that Theorem 4 applies to *any* loss function L and *any* regularizer $\mathfrak{R}$. We are not aware of any prior results of this form in the literature.

5 Comparison with Continuous-Feature Formulations

We next contrast the mixed-feature Wasserstein learning problem (3) with an *unbounded* continuous-feature formulation that replaces the support $\mathbb{Z}$ of the discrete features $\mathbf{z}$ with $\mathbb{R}^{M_z}$

(Section 5.1), as well as with a *bounded* continuous-feature formulation that replaces $\mathbb{Z}$ with its convex hull $\text{conv}(\mathbb{Z})$ (Section 5.2). This section focuses on a qualitative comparison of the three formulations; a quantitative comparison in terms of runtimes and generalization errors on benchmark instances is relegated to our numerical experiments (Section 6).

5.1 Comparison with Unbounded Continuous-Feature Formulation

Our reformulation (3) of the mixed-feature Wasserstein learning problem scales exponentially in the discrete features. It may thus be tempting to replace the support $\mathbb{Z}$ with $\mathbb{R}^{M_z}$, which would allow us to solve problem (3) in polynomial time using the reformulations proposed by Shafieezadeh-Abadeh et al. (2015) and Shafieezadeh-Abadeh et al. (2019). We next present a stylized example to illustrate that disregarding the discrete nature of the features $\mathbf{z}$ can result in pathological worst-case distributions that lack relevance to the problem.

Consider a classification problem where the single feature is binary and follows the Bernoulli distribution $z \sim \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$, and where the output variable $y \in \{-1, +1\}$ is related to the feature z via $\text{Prob}(y = z \mid z) = 0.8$. Note that in slight deviation from our earlier convention, the binary feature is supported on $\{-1, +1\}$ instead of $\{0, 1\}$; this allows us to present the example without the use of an intercept β_0 , which in turn will simplify our exposition. We set the label mismatch cost to $\kappa_y = 1$, and we use a non-smooth Hinge loss function. We generate random datasets comprising $N = 10$ samples, and we compare the following three formulations:

- (i) *Empirical risk model.* This model replaces the worst-case expectation in the Wasserstein learning problem (1) with an expectation over the empirical distribution $\hat{\mathbb{P}}_N$.
- (ii) *Mixed-feature Wasserstein model.* We solve the Wasserstein learning problem (3) with the discrete support $\mathbb{Z} = \{-1, +1\}$.
- (iii) *Unbounded continuous-feature Wasserstein model.* We solve the Wasserstein learning




Figure 1: (Worst-case) distributions for our three formulations. The bars represent the mean probabilities, and the whiskers indicate the corresponding standard deviations, computed over 10,000 statistically independent simulations.

problem (3) with the unbounded continuous support $\mathbb{Z} = \mathbb{R}$.

For the fixed hypothesis $\beta_z = 1$ and across 10,000 randomly generated datasets, the mean expected Hinge loss of the empirical risk model evaluates to 0.41. The worst-case expected Hinge loss of the mixed-feature Wasserstein learning problem varies from 0.41 to 2; the largest worst-case expected loss is attained for large Wasserstein radii ϵ , where the worst-case distributions approach any mixture of $\delta_{(+1,-1)}$ and $\delta_{(-1,+1)}$. The worst-case expected Hinge loss of the unbounded continuous-feature Wasserstein learning problem, on the other hand, diverges as the Wasserstein radii ϵ increase. In fact, for sufficiently large ϵ the worst-case distributions place a probability mass of 0.1 on a non-sensical atom $(\pm c, \pm 1)$, where c increases with ϵ .

Figure 1 illustrates the (worst-case) distributions of the three formulations for an intermediate Wasserstein radius of $\epsilon = 0.85$. As can be seen from the figure, the bounded continuous-feature Wasserstein learning problem places a mean probability mass of 0.063 (std. dev. 0.048) on an atom $(\pm c, \pm 1)$, where c is instance-specific and where $|c|$ lies in $[1.5, 7.5]$ (mean 2.29).

One might argue that an illogical worst-case distribution is less concerning in practical applications, where the primary role of Wasserstein learning may be viewed as providing a regularizing effect through the worst-case formulation. However, our numerical results will demonstrate that the unbounded continuous-feature formulation is consistently outperformed by our mixed-feature formulation in terms of generalization error, even when the Wasserstein radii ϵ are chosen via cross-validation.

5.2 Equivalence to Bounded Continuous-Feature Formulation

It is natural to ask whether the mixed-feature Wasserstein learning problem (1) is equivalent to its *bounded* continuous-feature counterpart that replaces the support $\mathbb{Z}$ of the discrete features with its convex hull, $\text{conv}(\mathbb{Z})$. We show in the following that this is indeed the case, but we argue that this equivalence bears little computational consequence.

To construct the bounded continuous-feature model, fix any $s \geq 2$ and define $\mathbb{U}(s)$ as the convex hull of the one-hot discrete feature encoding $\mathbb{Z}(s)$. One readily observes that

$$\mathbb{U}(s) = \left\{ \mathbf{u} \in [0, 1]^{s-1} : \sum_{i \in [s-1]} u_i \leq 1 \right\}.$$

Indeed, we have $\mathbb{Z}(s) \subseteq \mathbb{U}(s)$ by construction. To see that $\mathbb{U}(s) \subseteq \text{conv}(\mathbb{Z}(s))$, on the other hand, we note that any $\mathbf{u} \in \mathbb{U}(s)$ can be represented as $\mathbf{u} = \lambda_0 \cdot \mathbf{0} + \sum_{i \in [s-1]} \lambda_i \cdot \mathbf{e}_i$, where $\mathbf{0} \in \mathbb{Z}(s)$ is the vector of all zeros and where $\mathbf{e}_i \in \mathbb{Z}(s)$ is the i -th canonic basis vector in $\mathbb{R}^{s-1}$, with $\lambda_i = u_i$ and $\lambda_0 = 1 - \sum_{i \in [s-1]} \lambda_i$. Defining $\mathbb{U}$ as the convex hull of $\mathbb{Z}$, we observe that

$$\mathbb{U} = \text{conv}(\mathbb{Z}) = \bigtimes_{m \in [K]} \text{conv}(\mathbb{Z}(k_m)) = \bigtimes_{m \in [K]} \text{conv}(\mathbb{U}(k_m)),$$

where the second identity follows from Exercise 1.19 of Bertsekas (2009). Since the ground metric d_z from Definition 1 is only suitable for pairs of discrete feature vectors, we replace it

with the ground metric

$$d_u(\mathbf{u}, \mathbf{u}') = \left(\sum_{m \in [K]} \frac{1}{2} \|\mathbf{u}_m - \mathbf{u}'_m\|_1 + \frac{1}{2} |\mathbf{1}^\top (\mathbf{u}_m - \mathbf{u}'_m)| \right)^{1/p}$$

that extends to pairs of continuous feature vectors. One readily verifies that d_u is a metric, that is, it satisfies the identity of indiscernibles, positivity, symmetry and the triangle inequality. The proof of Proposition 2 below will also show that d_z and d_u coincide over all pairs of discrete features. We finally define the bounded continuous-feature model as

$$\begin{aligned} & \underset{\boldsymbol{\beta}, \lambda, \mathbf{s}}{\text{minimize}} && \lambda \epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\ & \text{subject to} && \sup_{(\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y}} \{l_{\boldsymbol{\beta}}(\mathbf{x}, \mathbf{u}, y) - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n)\} - \lambda \kappa_z d_u(\mathbf{u}, \mathbf{z}^n) \leq s_n \quad (3') \\ & && \forall n \in [N], \forall \mathbf{u} \in \mathbb{U} \end{aligned}$$

$$\boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}_x, \boldsymbol{\beta}_z) \in \mathbb{R}^{1+M_x+M_z}, \quad \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N,$$

where $\mathbf{u}$ and $\mathbb{U}$ serve as the continuous representations of the discrete features $\mathbf{z}$ and their support $\mathbb{Z}$, respectively. Note that in this formulation, the features $\mathbf{u} \in \mathbb{U}$ can be readily absorbed in the continuous feature set $\mathbf{x} \in \mathbb{X}$; we keep the distinction for improved clarity.

We next state the main result of this section.

Proposition 2. *The mixed-feature Wasserstein learning problem (3) and its corresponding bounded continuous-feature counterpart (3') share the same optimal value and the same set of optimal solutions.*

The proof of Proposition 2 crucially relies on the concavity of the family of mappings $\mathbf{u} \mapsto d_u(\mathbf{u}, \mathbf{u}')$, $\mathbf{u}' \in \mathbb{Z}$. Indeed, the equivalence stated in the proposition would *not* hold if we were to replace d_u with $d'_u(\mathbf{u}, \mathbf{u}') = \left(\sum_{m \in [K]} \|\mathbf{u}_m - \mathbf{u}'_m\|_\infty \right)^{1/p}$, despite d'_u being a metric on $\mathbb{U}$.

Unfortunately, the equivalence of Proposition 2 does not lead to more efficient solution

schemes for the mixed-feature Wasserstein learning problem (3). Indeed, we are not aware of any finite-dimensional convex reformulations of the bounded continuous-feature counterpart (3') when the loss function L is convex and Lipschitz continuous. When L is piece-wise affine and convex, on the other hand, problem (3') indeed admits a finite-dimensional convex reformulation of polynomial size. We will see in our numerical experiments, however, that our cutting plane scheme from Section 3 is much faster than the monolithic solution of this reformulation. This is due to the fact that the reformulation comprises a large number of decision variables as well as a constraint set that lacks desirable sparsity patterns.

6 Numerical Results

We compare empirically the performance of our mixed-feature Wasserstein learning problem (1) with classical and regularized learning methods as well as continuous-feature formulations of problem (1). To this end, Section 6.1 compares the out-of-sample losses of our mixed-feature Wasserstein learning problem (1) with those of the unbounded continuous-feature approximation when the Wasserstein radius is selected via cross-validation. Section 6.2 compares the runtime of our cutting plane solution scheme for the mixed-feature Wasserstein learning problem with that of solving the equivalent bounded continuous-feature formulation monolithically. Finally, Section 6.3 compares the out-of-sample performance of our formulation (1) with that of alternative methods on standard benchmark instances.

All algorithms were implemented in Julia v1.9.2 using the JuMP package (Lubin et al., 2023) and MOSEK v10.0, and all experiments were run on Intel Xeon 2.66GHz cluster nodes with 8GB memory in single-core and single-thread mode (unless otherwise specified). All implementations, datasets and experimental results are available on the GitHub repository accompanying this work.




Figure 2: Mean out-of-sample losses for various classification (left) and regression (right) tasks when the number of discrete features that are treated as such varies. All losses are scaled to $[0, 1]$ and shifted so that the curves do not overlap.

6.1 Comparison with Unbounded Continuous-Feature Formulation

Section 5.1 demonstrated that modeling discrete features in the Wasserstein learning problem (1) as continuous and subsequently solving an unbounded continuous-feature formulation may inadvertently hedge against pathological worst-case distributions.

We now explore whether the findings from Section 5.1 have implications on the out-of-sample performance of the mixed-feature and unbounded continuous-feature formulations when the Wasserstein radius ϵ is selected via cross-validation. We generate synthetic problem instances to be in full control of the problem dimensions. In particular, we generate 100 synthetic instances for each of the 6 loss functions considered in Appendices A and B: log-loss (LR), non-smooth Hinge Loss (SVM), smooth Hinge loss (SSVM), Huber loss (RR), τ -insensitive loss (SVR) and pinball loss (QR), assuming that the loss functions explain a large part (but not all) of the variability in the data. All instances comprise $N = 20$ data points, no numerical features, and $K = 20$ binary features. The small number of data points ensures that distributional robustness is required to obtain small out-of-sample losses. For each instance, we solve a mixed-feature Wasserstein learning problem that treats some of the features as binary, whereas the other features are

treated as continuous and unbounded. In particular, the extreme cases of modeling all and none of the features as binary correspond to our mixed-feature Wasserstein learning problem (1) and its unbounded continuous-feature approximation, respectively. We cross-validate the Wasserstein radius from the set $\epsilon \in \{0.005, 0.01, 0.015, \dots, 0.1\}$, and we choose $(\kappa_z, \kappa_y, p) = (1, 1, 1)$ as well as $\|\cdot\| = \|\cdot\|_1$ in our ground metric (2). We refer to the GitHub repository for further details of the instance generation procedure. Figure 2 reports the average out-of-sample losses across 100 randomly generated problem instances for the different loss functions. The figure reveals an overall trend of improved results when the number of binary features that are treated as such increases. Qualitatively, we observe that this conclusion is robust to different choices of the ϵ -grid used for cross-validation; we refer to the GitHub repository for further details.

6.2 Comparison with Bounded Continuous-Feature Formulation

In Section 5.2, we showed that our mixed-feature Wasserstein learning problem (1) is equivalent to a bounded continuous-feature model that replaces the support $\mathbb{Z}$ of the discrete features with its convex hull $\text{conv}(\mathbb{Z})$. In this section, we compare the runtimes of solving the mixed-feature formulation via our cutting plane method with those of solving the bounded continuous-feature formulation via an off-the-shelf solver.

When the loss function L in the mixed-feature Wasserstein learning problem (1) is piece-wise affine and convex, the equivalent bounded continuous-feature model admits a reformulation as a polynomial-size convex optimization problem. However, solving this reformulation is computationally expensive since it involves $\mathcal{O}(N \cdot J \cdot K)$ decision variables, where N , J and K denote the numbers of data points, affine pieces in the loss function and discrete features, respectively. Moreover, the reformulation lacks desirable sparsity features. The consequences are illustrated in Table 1, which compares the monolithic solution of the bounded continuous-feature reformulation with the solution of the mixed-feature model using our cutting plane approach.

The table reports median runtimes over 100 datasets with $N \in \{500, 1,000, 2,500\}$ data points. All instances contain 30 features, out of which 25%, 50%, 75% are discrete with 4 possible values, whereas the rest are continuous. We set the Wasserstein radius $\epsilon = 10^{-2}$ and use $(\kappa_z, \kappa_y, p) = (1, K, 1)$ for our ground metric. Further details of the implementation are available on the GitHub repository. We observe that our cutting plane scheme is significantly faster than the solution of the equivalent bounded continuous-feature model.

Method	N = 500			N = 1,000			N = 2,500		
	%25	%50	%75	%25	%50	%75	%25	%50	%75
SVM-cut	0.05	0.20	0.25	0.13	0.21	0.45	0.29	0.61	1.23
SVM-piece	6.35	11.80	15.25	12.48	22.99	40.97	30.4	109.78	146.62
QR-cut	0.15	0.24	0.36	0.28	0.57	1.18	1.08	1.64	3.08
QR-piece	3.35	5.59	9.24	8.71	19.95	26.1	22.83	48.34	114.19
SVR-cut	0.12	0.22	0.24	0.26	0.44	0.55	0.76	0.92	2.73
SVR-piece	4.35	9.20	17.29	10.2	16.78	35.46	27.13	52.02	95.08

Table 1: Median runtimes in secs to solve the mixed-feature model (-cut; using our cutting plane method) and the bounded continuous-feature model (-piece; solving a polynomial-size convex reformulation) for the hinge (SVM), pinball (QR; with $\tau = 0.5$) and τ -insensitive (SVR; with $\tau = 10^{-2}$) loss functions.

Although no polynomial-size convex reformulations are known for the equivalent bounded continuous-feature model when the loss function L in the mixed-feature Wasserstein learning problem (1) is Lipschitz continuous, we can approximate such loss functions to any desired accuracy via piece-wise affine convex functions. To this end, we take the pointwise maximum of tangent lines at equally spaced breakpoints and subsequently solve the bounded continuous-feature formulation for piece-wise affine convex loss functions. This provides a lower bound on the original problem, whose value approaches the optimal value of the original problem as the number of affine pieces in the approximated loss function increases. Table 2 reports the median runtimes for 100 datasets; we use the same problem parameters as in our previous experiment, and we report results for loss function approximations with $J \in \{5, 10, 25\}$ affine pieces. Similar

to our previous experiments, the table shows that our cutting plane scheme is significantly faster than the solution of the equivalent bounded continuous-feature model. We note that the approximated loss functions would require more than 25 pieces to guarantee a suboptimality of 0.1% or less; details are relegated to the GitHub repository.

Method	N = 500			N = 1,000			N = 2,500		
	%25	%50	%75	%25	%50	%75	%25	%50	%75
LR-cut	0.16	0.61	1.32	0.22	0.60	1.48	0.56	0.67	1.87
LR-piece [5]	42.08	108.0	157.06	127.97	382.6	608.76	743.01	2,371.29*	NaN
LR-piece [10]	50.32	133.0	186.56	482.51	236.94*	339.82*	3,107.19*	NaN	2,595.02*
LR-piece [25]	526.72	280.11	499.72	2,178.79*	1,634.96*	1,137.18*	NaN	NaN	NaN
SSVM-cut	0.14	0.31	0.54	0.22	0.46	0.63	0.6	0.61	1.57
SSVM-piece [5]	20.54	59.6	84.11	44.43	125.33	201.43	460.29	349.74	612.18*
SSVM-piece [10]	34.85	80.19	114.63	82.52	166.36	277.17*	1,924.21*	439.99*	730.46*
SSVM-piece [25]	90.25	198.32*	332.86*	221.74*	406.68	763.53*	NaN	1,161.48*	1,996.73*
RR-cut	0.16	0.19	0.3	0.27	0.3	0.52	0.7	0.78	1.31
RR-piece [5]	9.87	23.23	38.57	23.13	55.95	87.75	73.02	171.16	269.98
RR-piece [10]	16.52	38.92	60.38	37.69	87.57	143.39	105.4	281.6	464.49
RR-piece [25]	43.16	95.53	155.26	93.21	222.79	373.44*	296.18*	824.93*	1,584.85*

Table 2: Median runtimes in secs to solve the mixed-feature model (-cut; using our cutting plane method) and the bounded continuous-feature model (-piece [J]; solving a polynomial-size convex reformulation) for the logistic (LR), smooth hinge (SSVM) and Huber (RR; with $\delta = 0.05$) loss functions. A **NaN** indicates that more than half of the instances could not be solved within an hour of runtime. An asterisk (*) indicates that at least one instance could not be solved within an hour of runtime.

6.3 Performance on Benchmark Instances

In Section 6.1, we observed that the mixed-feature Wasserstein learning problem (1) can outperform its unbounded continuous-feature approximation in terms of out-of-sample losses. In practice, however, loss functions merely serve as surrogates for the misclassification rate (in classification problems) or mean squared error (in regression problems). Moreover, Section 6.1 considered synthetically generated problem instances, which may lack some of the intricate structure of real-life datasets. This section therefore compares the out-of-sample misclassification rates and mean squared errors of problem (1) with those of alternative methods on standard

benchmark instances. In particular, we selected 8 of the most popular classification and 7 of the most popular regression datasets from the UCI machine learning repository (Kelly et al., 2017). We processed the datasets to (i) handle missing values and inconsistencies in the data, (ii) employ a one-hot encoding for discrete variables, and (iii) convert the output variable (to a binary value for classification tasks and a $[-1, 1]$ -interval for regression tasks). All datasets, processing scripts as well as further results on other UCI datasets can be found in the GitHub repository.

Tables 3 and 4 report averaged results for logistic regression (classification) and support vector regression (regression) over 100 random splits into 80% training set and 20% test set for both the original datasets as well as parsimonious variants where only half of the data points are available. In the tables, the column groups report (from left to right) the problem instances’ names and dimensions; the results for the unregularized implementations of the nominal problem as well as the mixed-feature Wasserstein learning problem (1) using $(\kappa_z, p) = (1, 1)$ and $\|\cdot\| = \|\cdot\|_1$ in our ground metric (2); the results for the corresponding l_2 -regularized versions; and the results for two unbounded continuous-feature approximations of problem (1). For the unregularized mixed-feature and continuous-feature Wasserstein learning problems, we jointly cross-validate the hyperparameters (ϵ, κ_y) where $\epsilon \in \{0, 10^{-5}, 10^{-3}, 10^{-1}\}$ and $\kappa_y \in \{1, K, \infty\}$. For the regularized nominal model, we cross-validate the regularization penalty α from the set $\{0\} \cup \{c \cdot 10^{-p} : c \in \{1, 5\}, p = 1, \dots, 6\}$. For the regularized mixed-feature Wasserstein learning problem, finally, we cross-validate the hyperparameters $(\epsilon, \kappa_y, \alpha)$ from the Cartesian product of the previous three sets. The method with the smallest error in each column group is highlighted in bold, and the method with the smallest error across all column groups has a grey background. For instances where a version of the mixed-feature Wasserstein learning problem attains the smallest error across all column groups, the symbols $\dagger$ and $\ddagger$ indicate a statistically significant improvement of that model over the nominal and regularized nominal learning problem as well

as over the continuous-feature approximation, respectively, at a p -value of 0.05. Most nominal models were solved within seconds, most continuous-feature models were solved within minutes, and most of the mixed-feature models were solved within 10 minutes. We relegate the details of the statistical significance tests and runtimes to the GitHub repository, which also contains comparisons across alternative loss functions. The results are qualitatively similar to those of Tables 3 and 4.

Overall, we observe from the tables that the mixed-feature Wasserstein learning problems outperform both the (regularized) nominal problems and the unbounded continuous-feature approximations in the classification and regression tasks. The outperformance of the mixed-feature Wasserstein learning problems over the continuous-feature approximations tends to be more substantial for problem instances with many discrete features, which further confirms our findings from Sections 5.1 and 6.1. Also, the mixed-feature Wasserstein learning problems tend to perform better on the parsimonious versions of the problem instances. This is intuitive as we expect robust optimization to be particularly effective in data sparse environments. Regularizing the mixed-feature Wasserstein learning problem does not yield significant advantages in the classification tasks, but it does help in the regression tasks. Finally, while we do not observe any significant differences among the classification loss functions, the Huber loss function performs best on the regression problem instances.

7 Conclusions

Wasserstein learning offers a principled approach to mitigating overfitting in supervised learning. While Wasserstein learning has been studied extensively for datasets with exclusively continuous features, the existing solution techniques cannot handle discrete features faithfully without sacrificing performance in real-world applications.

In this paper, we propose a cutting plane algorithm for Wasserstein classification and

regression problems that involve both continuous and discrete features. While the worst-case complexity of our method is exponential, it performs well on both synthetic and benchmark datasets. The core of our algorithm is an efficient identification of the most violated constraints with respect to the current solution. Despite the discrete nature of the problem, we show that it can be solved in polynomial time. We also study the complexity of mixed-feature Wasserstein learning more broadly, and we investigate its connections to regularized as well as different continuous-feature formulations.

Our work opens several promising directions for future research. Most notably, it would be instructive to explore whether similar algorithmic techniques can be developed for other supervised learning models, including decision trees, ensemble methods, and neural networks. Furthermore, it would be valuable to investigate the development of polynomial-time solution schemes that avoid reliance on the ellipsoid method and that overcome the scalability limitations of monolithic reformulations. For example, the fact that our algorithm for the mixed-feature problem relies on a sorting argument to identify the most violated constraint could potentially pave the way to developing robust counterparts using standard convex reformulations of sorting-type constraints.

Acknowledgments The authors gratefully acknowledge the detailed and constructive feedback provided by the department editor, the anonymous associate editor, and the two anonymous referees, which has substantially improved the quality of this manuscript. We are especially indebted to the reviewers who suggested the possibility that our mixed-feature formulation might be equivalent to a bounded continuous-feature model, which directly inspired Section 5.2 of the manuscript. The authors are also indebted to Oscar Dowson for his invaluable guidance on efficient implementations of our Julia source codes.

Dataset	N	M_x	M_z	K	Reduced?	Nom	MixF	r-Nom	r-MixF	ConF
<u>balance-scale</u>	625	0	16	4	✖ ✓	0.56% 1.65%	0.44% † 1.47% †‡	0.52% 1.92%	0.48% 1.77%	0.44% 1.60%
<u>breast-cancer</u>	277	0	42	9	✖ ✓	30.10% 31.87%	28.46% †‡ 29.88%	29.46% 29.58%	29.36% 28.31% †‡	29.45% 30.48%
<u>credit-approval</u>	690	6	36	9	✖ ✓	13.59% 16.30%	13.12% †‡ 14.94%	13.70% 14.58%	13.55% 14.94%	13.88% 15.11%
<u>cylinder-bands</u>	539	19	43	14	✖ ✓	22.85% 27.37%	22.71% 26.36%	22.90% 27.18%	22.62% ‡ 26.02% †	23.13% 26.41%
<u>lymphography</u>	148	0	42	18	✖ ✓	17.24% 23.98%	14.66% †‡ 17.73% ‡	15.86% 17.96%	16.55% 18.07%	17.41% 19.14%
<u>primacy</u>	339	0	25	17	✖ ✓	13.51% 15.89%	13.96% 14.70%	14.25% 14.70%	13.58% 14.48% †	14.33% 14.68%
<u>spect</u>	267	0	22	22	✖ ✓	20.28% 23.22%	18.40% 19.44%	20.09% 23.34%	18.30% 19.13% †‡	17.93% 19.66%
<u>tic-tac-toe</u>	958	0	18	9	✖ ✓	1.94% 2.75%	1.70% 1.65% †	1.70% 1.81%	1.70% 1.68%	1.70% 1.65%

Table 3: Mean classification error for the **logistic regression loss function**. N , M_x , M_z and K refer to the problem parameters from Section 2. The column ‘Reduced?’ indicates whether the full or sparse version of the dataset is being considered. Nom, MixF and ConF refer to the nominal problem formulation, the mixed-feature Wasserstein learning problem (1) and its unbounded continuous-feature approximation, respectively. The prefix ‘r-’ indicates the use of an l_2 -regularization.

Dataset	N	M_x	M_z	K	Reduced?	Nom	MixF	r-Nom	r-MixF	ConF
<u>bike</u>	17,379	4	31	5	✖ ✓	217.87 218.18	217.61 217.96	217.88 218.15	217.62 217.95	217.61 217.99
<u>fire</u>	517	10	31	4	✖ ✓	133.10 185.04	122.58 108.54	136.26 125.23	121.13 †‡ 108.28 †	122.58 108.54
<u>flare</u>	1,066	0	21	9	✖ ✓	212.93 1,562.88	213.55 209.84	215.49 232.67	211.32 †‡ 207.79 †‡	214.00 210.22
<u>garments</u>	1,197	7	22	4	✖ ✓	369.10 871.46	368.69 371.12	370.08 370.72	369.04 371.45	368.89 371.79
<u>imports</u>	193	14	45	10	✖ ✓	621.38 761.05	293.31 †‡ 378.83	331.36 370.97	297.68 338.95 †‡	314.97 382.32
<u>student</u>	395	13	26	17	✖ ✓	399.65 451.93	397.63 397.62 †‡	383.57 410.40	383.18 ‡ 399.46	404.89 412.98
<u>vegas</u>	504	5	106	14	✖ ✓	232,332.44 287,597.53	502.89 507.90	503.37 507.74	492.94 †‡ 502.73 †‡	506.67 510.15

Table 4: Mean squared errors for the **τ -insensitive loss function** (with $\tau = 0.1$). We use the same abbreviations as in Table 3.

References

- Alley, M., M. Biggs, R. Hariss, C. Herrmann, M. L. Li, and G. Perakis (2023). Pricing for heterogeneous products: Analytics for ticket reselling. *Manufacturing & Service Operations Management* 25(2), 409–426.
- Ban, G.-Y. and C. Rudin (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research* 67(1), 90–108.
- Bastani, H. and M. Bayati (2020). Online decision making with high-dimensional covariates. *Operations Research* 68(1), 276–294.
- Behrendt, A., M. Savelsbergh, and H. Wang (2023). A prescriptive machine learning method for courier scheduling on crowdsourced delivery platforms. *Transportation Science* 57(4), 889–907.
- Ben-Tal, A., L. E. Ghaoui, and A. Nemirovski (2009). *Robust Optimization*. Princeton University Press.
- Bertsekas, D. (2009). *Convex optimization theory*, Volume 1. Athena Scientific.
- Bertsimas, D., A. Delarue, P. Jaillet, and S. Martin (2019). Travel time estimation in the age of big data. *Operations Research* 67(2), 498–515.
- Bertsimas, D. and D. den Hertog (2022). *Robust and Adaptive Optimization*. Dynamic Ideas.
- Bertsimas, D. and N. Kallus (2020). From predictive to prescriptive analytics. *Management Science* 66(3), 1025–1044.
- Bertsimas, D., A. O’Hair, S. Relyea, and J. Silberholz (2016). An analytics approach to designing combination chemotherapy regimens for cancer. *Management Science* 62(5), 1511–1531.
- Bertsimas, D. and J. Pauphilet (2023). Hospital-wide inpatient flow optimization. *Management Science*, Available ahead of print.
- Blanchet, J. and Y. Kang (2021). Sample out-of-sample inference based on Wasserstein distance. *Operations Research* 69(3), 985–1013.
- Blanchet, J., Y. Kang, and K. Murthy (2019). Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability* 56(3), 830–857.
- Blanchet, J. and K. Murthy (2019). Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research* 44(2), 565–600.
- Boyd, S. and L. Vandenberghe (2004). *Convex optimization*. Cambridge university press.
- Chan, T. C. Y., R. Mahmood, D. L. O’Connor, D. Stone, S. Unger, R. K. Wong, and I. Y. Zhu (2023). Got (optimal) milk? Pooling donations in human milk banks with machine learning and optimization. *Manufacturing & Service Operations Management*, Available ahead of print.
- Duchi, J., T. Hashimoto, and H. Namkoong (2023). Distributionally robust losses for latent covariate mixtures. *Operations Research* 71(2), 649–664.

- Feldman, J., D. J. Zhang, X. Liu, and N. Zhang (2022). Customer choice models vs. machine learning: Finding optimal product displays on Alibaba. *Operations Research* 70(1), 309–328.
- Ferreira, K. J., B. H. A. Lee, and D. Simchi-Levi (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1), 69–88.
- Gao, R. (2022). Finite-sample guarantees for Wasserstein distributionally robust optimization: Breaking the curse of dimensionality. *Operations Research*, Available ahead of print.
- Gao, R., X. Chen, and A. J. Kleywegt (2022). Wasserstein distributionally robust optimization and variation regularization. *Operations Research*, Available ahead of print.
- Gao, R. and A. Kleywegt (2023). Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research* 48(2), 603–655.
- Glaeser, C. K., M. Fisher, and X. Su (2019). Optimal retail location: Empirical methodology and application to practice. *Manufacturing & Service Operations Management* 21(1), 86–102.
- Grötschel, M., L. Lovász, and A. Schrijver (1988). *Geometric Algorithms and Combinatorial Optimization*. Springer.
- Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- Kallus, N. and M. Udell (2020). Dynamic assortment personalization in high dimensions. *Operations Research* 68(4), 1020–1037.
- Kelly, M., R. Longjohn, and K. Nottingham (2017). The UCI machine learning repository. <https://archive.ics.uci.edu>.
- Kuhn, D., P. Mohajerin Esfahani, V. Nguyen, and S. Shafieezadeh-Abadeh (2019). Wasserstein distributionally robust optimization: Theory and applications in machine learning. *INFORMS TutORials in Operations Research*, 130–169.
- Kuhn, D., S. Shafiee, and W. Wiesemann (2024). Distributionally robust optimization. Working Paper.
- Lam, H. (2019). Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *Operations Research* 6(4), 1090–1105.
- Lee, Y. T., A. Sidford, and S. C.-W. Wong (2015). A faster cutting plane method and its implications for combinatorial and convex optimization. pp. 1049–1065.
- Li, R., M. Tobey, M. E. Mayorga, S. Caltagirone, and O. Y. Özaltın (2023). Detecting human trafficking: Automated classification of online customer reviews of massage businesses. *Manufacturing & Service Operations Management* 25(3), 1051–1065.
- Lubin, M., O. Dowson, J. Dias Garcia, J. Huchette, B. Legat, and J. P. Vielma (2023). JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation* 15(3), 581–589.

- Michaud, R. O. (1989). The Markowitz optimization enigma: Is ‘optimized’ optimal? *Financial Analysts Journal* 45(1), 31–42.
- Mohajerin Esfahani, P. and D. Kuhn (2018). Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* 171(1-2), 1–52.
- Murphy, K. P. (2022). *Probabilistic Machine Learning: An Introduction*. MIT Press.
- Qi, M., Y. Cao, and Z.-J. Shen (2022). Distributionally robust conditional quantile prediction with fixed design. *Management Science* 68(3), 1639–1658.
- Rahimian, H. and S. Mehrotra (2022). Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization* 3, 1–85.
- Redacted (2022). Redacted for the purpose of the review process.
- Samorani, M., S. L. Harris, L. G. Blount, H. Lu, and M. A. Santoro (2022). Overbooked and overlooked: Machine learning and racial bias in medical appointment scheduling. *Manufacturing & Service Operations Management* 24(6), 2825–2842.
- Shafieezadeh-Abadeh, S., D. Kuhn, and P. Mohajerin Esfahani (2019). Regularization via mass transportation. *Journal of Machine Learning Research* 20(103), 1–68.
- Shafieezadeh-Abadeh, S., P. Mohajerin Esfahani, and D. Kuhn (2015). Distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, Volume 28, pp. 1576–1584.
- Smith, J. E. and R. L. Winkler (2006). The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science* 52(3), 311–322.
- Van Parys, B. P. G., P. Mohajerin Esfahani, and D. Kuhn (2021). From data to decisions: Distributionally robust optimization is optimal. *Management Science* 67(6), 3387–3402.

E-Companion to “It’s All in the Mix:

Wasserstein Classification and Regression with Mixed Features”

Appendix A: Wasserstein Classification with Mixed Features

This section derives exponential-size convex reformulations of the mixed-feature Wasserstein learning problem (3) for classification problems with convex and Lipschitz continuous as well as piece-wise affine and convex loss functions L . One readily confirms that the resulting reformulations are special cases of our unified representation (4), which implies that our results in this appendix prove Theorem 1 for the case of classification problems.

Proposition 3. *Consider a classification problem with a convex and Lipschitz continuous loss function as well as $\mathbb{X} = \mathbb{R}^{M_x}$. In this case, the Wasserstein learning problem (3) is equivalent to*

$$\begin{aligned} & \underset{\beta, \lambda, \mathbf{s}}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\ & \text{subject to} && \left. \begin{aligned} l_{\beta}(\mathbf{x}^n, \mathbf{z}, y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) &\leq s_n \\ l_{\beta}(\mathbf{x}^n, \mathbf{z}, -y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) - \lambda\kappa_y &\leq s_n \end{aligned} \right\} \forall n \in [N], \forall \mathbf{z} \in \mathbb{Z} \quad (\text{E-1}) \\ & && \text{lip}(L) \cdot \|\beta_x\|_* \leq \lambda \\ & && \beta = (\beta_0, \beta_x, \beta_z) \in \mathbb{R}^{1+M_x+M_z}, \quad \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N. \end{aligned}$$

When $\kappa_y = \infty$, the second constraint set in Proposition 3 is redundant for any $\lambda > 0$. This is the case whenever λ does not vanish at optimality (*i.e.*, the Wasserstein radius ϵ actually affects the optimal solution).

Formulation (E-1) emerges as a special case of our unified representation (4) if we set $\mathcal{I} = \{\pm 1\}$, $\boldsymbol{\theta} = (\beta, \lambda)$, $\boldsymbol{\sigma} = \mathbf{s}$, $f_0(\boldsymbol{\theta}, \boldsymbol{\sigma}) = \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n$, $f_{ni}(e) = L(e)$, $g_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}) = iy^n \cdot [\beta_0 + \beta_x^\top \mathbf{x}^n + \beta_z^\top \mathbf{z}]$, $h_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)) = \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) + \lambda\kappa_y \cdot \mathbb{1}[i = -1]$ and $\Theta = \{\boldsymbol{\theta} : \text{lip}(L) \cdot \|\beta_x\|_* \leq \lambda, \lambda \in \mathbb{R}_+\}$. One verifies that this choice of $\mathcal{I}$, $\boldsymbol{\theta}$, $\boldsymbol{\sigma}$, f_0 , f_{ni} , g_{ni} , h_{ni} and Θ

satisfies the regularity conditions imposed on problem (4) in the main text.

Proposition 4. *If β in problem (E-1) is restricted to a bounded hypothesis set $\mathcal{H}$, then λ can be restricted to a bounded set as well.*

Proposition 3 covers support vector machines with a smooth Hinge loss, and logistic regression with a log-loss; we provide the corresponding reformulations in our GitHub repository.

We now provide a reformulation of problem (3) without embedded maximizations when the loss function L is piece-wise affine and convex.

Proposition 5. *Consider a classification problem with a piece-wise affine and convex loss function $L(e) = \max_{j \in [J]} \{a_j e + b_j\}$, and assume that $\mathbb{X} = \{\mathbf{x} \in \mathbb{R}^{M_x} : \mathbf{C}\mathbf{x} \leq_{\mathcal{C}} \mathbf{d}\}$ for some $\mathbf{C} \in \mathbb{R}^{r \times M_x}$, $\mathbf{d} \in \mathbb{R}^r$ and proper convex cone $\mathcal{C} \subseteq \mathbb{R}^r$. If $\mathbb{X}$ admits a Slater point $\mathbf{x}^s \in \mathbb{R}^{M_x}$ such that $\mathbf{C}\mathbf{x}^s <_{\mathcal{C}} \mathbf{d}$, then the Wasserstein learning problem (3) is equivalent to*

$$\begin{aligned}
& \underset{\beta, \lambda, \mathbf{s}, \mathbf{q}_{nj}^+, \mathbf{q}_{nj}^-}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\
& \text{subject to} && \left. \begin{aligned} & \mathbf{q}_{nj}^{+ \top} (\mathbf{d} - \mathbf{C}\mathbf{x}^n) + a_j y^n \cdot (\beta_0 + \beta_x^\top \mathbf{x}^n + \beta_z^\top \mathbf{z}) - \lambda \kappa_z d_z(\mathbf{z}, \mathbf{z}^n) + b_j \leq s_n \\ & \mathbf{q}_{nj}^{- \top} (\mathbf{d} - \mathbf{C}\mathbf{x}^n) - a_j y^n \cdot (\beta_0 + \beta_x^\top \mathbf{x}^n + \beta_z^\top \mathbf{z}) - \lambda \kappa_z d_z(\mathbf{z}, \mathbf{z}^n) - \lambda \kappa_y + b_j \leq s_n \end{aligned} \right\} \\
& && \forall n \in [N], \forall j \in [J], \forall \mathbf{z} \in \mathbb{Z} \\
& && \|a_j y^n \cdot \beta_x - \mathbf{C}^\top \mathbf{q}_{nj}^+\|_* \leq \lambda, \quad \|a_j y^n \cdot \beta_x + \mathbf{C}^\top \mathbf{q}_{nj}^-\|_* \leq \lambda \quad \forall n \in [N], \forall j \in [J] \\
& && \beta = (\beta_0, \beta_x, \beta_z) \in \mathbb{R}^{1+M_x+M_z}, \quad \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N \\
& && \mathbf{q}_{nj}^+, \mathbf{q}_{nj}^- \in \mathcal{C}^*, \quad n \in [N] \text{ and } j \in [J].
\end{aligned} \tag{E-2}$$

When $\kappa_y = \infty$, the constraints involving $\mathbf{q}_{nj}^-$ are redundant whenever $\lambda > 0$. In this case, the variables $\mathbf{q}_{nj}^-$ and all constraints involving $\mathbf{q}_{nj}^-$ can be removed from Proposition 5.

Formulation (E-2) emerges as a special case of our unified representation (4) if we set $i = (j, t)$, $\mathcal{I} = [J] \times \{\pm 1\}$, $\boldsymbol{\theta} = (\beta, \lambda, \mathbf{q}_{ni})$, $\boldsymbol{\sigma} = \mathbf{s}$, $f_0(\boldsymbol{\theta}, \boldsymbol{\sigma}) = \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n$, $f_{ni}(e) = a_j e + b_j$, $g_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}) = t y^n \cdot [\beta_0 + \beta_x^\top \mathbf{x}^n + \beta_z^\top \mathbf{z}]$, $h_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)) = -\mathbf{q}_{ni}^\top (\mathbf{d} - \mathbf{C}\mathbf{x}^n) +$

$\lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) + \lambda\kappa_y \cdot \mathbb{1}[t = -1]$ and $\Theta = \{\boldsymbol{\theta} : \|a_j y^n \cdot \boldsymbol{\beta}_x - t \cdot \mathbf{C}^\top \mathbf{q}_{ni}\|_* \leq \lambda \forall (n, i) \in [N] \times \mathcal{I}, \mathbf{q}_{ni} \in \mathcal{C}^* \forall (n, i) \in [N] \times \mathcal{I}, \lambda \in \mathbb{R}_+\}$. One verifies that this choice of $\mathcal{I}, \boldsymbol{\theta}, \boldsymbol{\sigma}, f_0, f_{ni}, g_{ni}, h_{ni}$ and Θ satisfies the regularity conditions imposed on problem (4) in the main text.

Proposition 6. *If $\boldsymbol{\beta}$ in problem (E-2) is restricted to a bounded hypothesis set $\mathcal{H}$, then $(\mathbf{q}_{ni}^+, \mathbf{q}_{ni}^-)_{n,i}$ and λ can be restricted to a bounded set as well.*

Proposition 5 covers support vector machines with a non-smooth Hinge loss; we provide the corresponding reformulation in our GitHub repository.

Appendix B: Wasserstein Regression with Mixed Features

This section derives exponential-size convex reformulations of the mixed-feature Wasserstein learning problem (3) for regression problems with convex and Lipschitz continuous as well as piece-wise affine and convex loss functions L . One readily confirms that the resulting reformulations are special cases of our unified representation (4), which implies that our results in this appendix prove Theorem 1 for the case of regression problems.

Proposition 7. *Consider a regression problem with a convex and Lipschitz continuous loss function L and $(\mathbb{X}, \mathbb{Y}) = (\mathbb{R}^{M_x}, \mathbb{R})$. In this case, the Wasserstein learning problem (3) equals*

$$\begin{aligned}
& \underset{\boldsymbol{\beta}, \lambda, \mathbf{s}}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\
& \text{subject to} && l_{\boldsymbol{\beta}}(\mathbf{x}^n, \mathbf{z}, y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) \leq s_n \quad \forall n \in [N], \forall \mathbf{z} \in \mathbb{Z} \\
& && \text{lip}(L) \cdot \|\boldsymbol{\beta}_x\|_* \leq \lambda \\
& && \text{lip}(L) \leq \lambda\kappa_y \\
& && \boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}_x, \boldsymbol{\beta}_z) \in \mathbb{R}^{1+M_x+M_z}, \quad \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N.
\end{aligned} \tag{E-3}$$

When $\kappa_y = \infty$ and $\lambda > 0$, the third constraint in Proposition 7 is redundant.

Formulation (E-3) emerges as a special case of our unified representation (4) if we set $\boldsymbol{\theta} = (\boldsymbol{\beta}, \lambda)$, $\boldsymbol{\sigma} = \mathbf{s}$, $f_0(\boldsymbol{\theta}, \boldsymbol{\sigma}) = \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n$, $f_{ni}(e) = L(e)$, $g_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}) = \beta_0 + \boldsymbol{\beta}_x^\top \mathbf{x}^n + \boldsymbol{\beta}_z^\top \mathbf{z} - y^n$, $h_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)) = \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n)$ and $\Theta = \{\boldsymbol{\theta} : \text{lip}(L) \cdot \|\boldsymbol{\beta}_x\|_* \leq \lambda, \text{lip}(L) \leq \lambda\kappa_y, \lambda \in \mathbb{R}_+\}$. One verifies that this choice of $\boldsymbol{\theta}$, $\boldsymbol{\sigma}$, f_0 , f_{ni} , g_{ni} , h_{ni} and Θ satisfies the regularity conditions imposed on problem (4) in the main text.

Corollary 1. *If $\boldsymbol{\beta}$ in problem (E-3) is restricted to a bounded hypothesis set $\mathcal{H}$, then λ can be restricted to a bounded set as well.*

Proposition 7 covers regression with a Huber loss; we provide the corresponding reformulation in our GitHub repository. We now provide a reformulation of problem (3) without embedded maximizations when the loss function L is piece-wise affine and convex.

Proposition 8. *Consider a regression problem with a piece-wise affine and convex loss function $L(e) = \max_{j \in [J]} \{a_j e + b_j\}$, and assume that $\mathbb{X} \times \mathbb{Y} = \{(\mathbf{x}, y) \in \mathbb{R}^{M_x+1} : \mathbf{C}_x \mathbf{x} + \mathbf{c}_y \cdot y \leq_{\mathcal{C}} \mathbf{d}\}$ for some $\mathbf{C}_x \in \mathbb{R}^{r \times M_x}$, $\mathbf{c}_y \in \mathbb{R}^r$, $\mathbf{d} \in \mathbb{R}^r$ and proper convex cone $\mathcal{C} \subseteq \mathbb{R}^r$. If this set admits a Slater point $(\mathbf{x}^s, y^s) \in \mathbb{R}^{M_x+1}$ such that $\mathbf{C}_x \mathbf{x}^s + \mathbf{c}_y \cdot y^s <_{\mathcal{C}} \mathbf{d}$, then problem (3) is equivalent to*

$$\begin{aligned}
& \underset{\boldsymbol{\beta}, \lambda, \mathbf{s}, \mathbf{q}_{nj}}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\
& \text{subject to} && \mathbf{q}_{nj}^\top (\mathbf{d} - \mathbf{C}_x \mathbf{x}^n - \mathbf{c}_y \cdot y^n) + a_j \cdot (\beta_0 + \boldsymbol{\beta}_x^\top \mathbf{x}^n + \boldsymbol{\beta}_z^\top \mathbf{z} - y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) + b_j \leq s_n \\
& && \forall n \in [N], \forall j \in [J], \forall \mathbf{z} \in \mathbb{Z} \\
& && \|a_j \cdot \boldsymbol{\beta}_x - \mathbf{C}_x^\top \mathbf{q}_{nj}\|_* \leq \lambda, \quad |-a_j - \mathbf{c}_y^\top \mathbf{q}_{nj}| \leq \lambda\kappa_y \quad \forall n \in [N], \forall j \in [J] \\
& && \boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}_x, \boldsymbol{\beta}_z) \in \mathbb{R}^{1+M_x+M_z}, \quad \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N \\
& && \mathbf{q}_{nj} \in \mathcal{C}^\star, \quad n \in [N] \text{ and } j \in [J].
\end{aligned} \tag{E-4}$$

When $\kappa_y = \infty$ and $\lambda > 0$, the constraints $|-a_j - \mathbf{c}_y^\top \mathbf{q}_{nj}| \leq \lambda\kappa_y$ in Proposition 8 are redundant.

Formulation (E-4) emerges as a special case of our unified representation (4) if we set $\mathcal{I} = [J]$, $\boldsymbol{\theta} = (\boldsymbol{\beta}, \lambda, \mathbf{q}_{ni})$, $\boldsymbol{\sigma} = \mathbf{s}$, $f_0(\boldsymbol{\theta}, \boldsymbol{\sigma}) = \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n$, $f_{ni}(e) = a_i e + b_i$, $g_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}) = \beta_0 + \boldsymbol{\beta}_x^\top \mathbf{x}^n + \boldsymbol{\beta}_z^\top \mathbf{z} - y^n$, $h_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_z(\mathbf{z}, \mathbf{z}^n)) = -\mathbf{q}_{ni}^\top (\mathbf{d} - \mathbf{C}_x \mathbf{x}^n - \mathbf{c}_y \cdot y^n) + \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n)$ and

$\Theta = \{\boldsymbol{\theta} : \|a_i \cdot \boldsymbol{\beta}_x - \mathbf{C}^\top \mathbf{q}_{ni}\|_* \leq \lambda \forall (n, i) \in [N] \times \mathcal{I}, |-a_i - \mathbf{c}_y^\top \mathbf{q}_{ni}| \leq \lambda \kappa_y \forall (n, i) \in [N] \times \mathcal{I}, \mathbf{q}_{ni} \in \mathcal{C}^* \forall (n, i) \in [N] \times \mathcal{I}, \lambda \in \mathbb{R}_+\}$. One verifies that this choice of $\mathcal{I}$, $\boldsymbol{\theta}$, $\boldsymbol{\sigma}$, f_0 , f_{ni} , g_{ni} , h_{ni} and Θ satisfies the regularity conditions imposed on problem (4) in the main text.

Corollary 2. *If β in problem (E-4) is restricted to a bounded hypothesis set $\mathcal{H}$, then $(\mathbf{q}_{ni})_{n,i}$ and λ can be restricted to a bounded set as well.*

Proposition 8 covers support vector regression with a τ -insensitive loss, and quantile regression with a pinball loss; we provide the corresponding reformulations in our GitHub repository.

Appendix C: Proofs

In this section, we present the proofs of all the results in the main body of this paper. The proofs of Appendices A and B are relegated to our GitHub repository.

Proof of Observation 1. The statement follows from similar arguments as in the proof of Theorem 1 by Shafieezadeh-Abadeh et al. (2015). Details are omitted for the sake of brevity. $\square$

Proof of Proposition 1. We first show that LB and UB constitute lower and upper bounds on the optimal value of (4) throughout the execution of the algorithm, and then we conclude that Algorithm 1 terminates in finite time with an optimal solution $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ to problem (4).

Algorithm 1 updates LB to $f_0(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ in each iteration of the while-loop. Since $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$ is an optimal solution to the relaxation of problem (4) that only contains the constraints $\mathcal{W} \subseteq [N] \times \mathcal{I} \times \mathbb{Z}$, LB indeed constitutes a lower bound on the optimal value of (4). Moreover, since no elements are ever removed from $\mathcal{W}$, the sequence of lower bounds LB is monotonic.

To see that UB constitutes an upper bound on the optimal value of (4) throughout the execution of Algorithm 1, we claim that in each iteration, $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^* + \boldsymbol{\vartheta}^*)$ constitutes a feasible solution to problem (4). Indeed, we have $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*) \in \Theta \times \mathbb{R}_+^N$ and $\boldsymbol{\vartheta}^* \in \mathbb{R}_+^N$ by construction, while

for all $n \in [N]$, $i \in \mathcal{I}$ and $\mathbf{z} \in \mathbb{Z}$, we have that

$$f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n)) \leq \sigma_n^* + \vartheta(n, i) \leq \sigma_n^* + \vartheta(n, i(n)) \leq \sigma_n^* + \vartheta_n^*,$$

that is, $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^* + \boldsymbol{\vartheta}^*)$ is indeed feasible in problem (4). Moreover, the sequence of upper bounds UB is monotonic by construction of the updates.

To see that Algorithm 1 terminates in finite time, finally, note that each iteration of the while-loop either adds a new constraint index $(n, i(n), \mathbf{z}(n, i(n)))$ to $\mathcal{W}$, or we have $\vartheta(n, i(n)) \leq 0$ for all $n \in [N]$. In the latter case, however, we have $\boldsymbol{\vartheta}^* = \mathbf{0}$ and thus LB = UB at the end of the iteration. The claim now follows from the fact that the index set $[N] \times \mathcal{I} \times \mathbb{Z}$ is finite. $\square$

Proof of Theorem 2. For fixed $(\boldsymbol{\theta}^*, \boldsymbol{\sigma}^*)$, finding the most violated constraint index $\mathbf{z}(n, i)$ in constraint group (n, i) amounts to solving the combinatorial optimization problem

$$\begin{aligned} & \underset{\mathbf{z}}{\text{maximize}} && f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n)) - \sigma_n^* \\ & \text{subject to} && \mathbf{z} \in \mathbb{Z}. \end{aligned}$$

We can solve this problem by solving the $K + 1$ problems

$$\left[\begin{array}{ll} \underset{\mathbf{z}}{\text{maximize}} & f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \delta) - \sigma_n^* \\ \text{subject to} & d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n) = \delta \\ & \mathbf{z} \in \mathbb{Z} \end{array} \right] \quad \forall \delta \in [K] \cup \{0\},$$

where each problem conditions on a fixed number δ of discrepancies between $\mathbf{z}$ and $\mathbf{z}^n$, and subsequently choosing any solution $\mathbf{z}(\delta)$ that attains the maximum optimal objective value among those $K + 1$ problems. Removing constant terms from those $K + 1$ problems, we observe

that the δ -th problem shares its set of optimal solutions with the problem

$$\begin{aligned} & \underset{\mathbf{z}}{\text{maximize}} && f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) \\ & \text{subject to} && d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n) = \delta \\ & && \mathbf{z} \in \mathbb{Z}. \end{aligned}$$

Since the outer function f_{ni} in the objective function is convex, the objective function is maximized whenever the inner function g_{ni} in the objective function is either maximized or minimized. We thus conclude that the δ -th problem is solved by solving the two problems

$$\left[\begin{array}{ll} \underset{\mathbf{z}}{\text{maximize}} & \mu \cdot g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}) \\ \text{subject to} & d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n) = \delta \\ & \mathbf{z} \in \mathbb{Z} \end{array} \right] \quad \forall \mu \in \{\pm 1\}$$

and subsequently choosing the solution $\mathbf{z}(\mu, \delta)$ that attains the larger value $f_{ni}(g_{ni}(\boldsymbol{\theta}^*, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z}(\mu, \delta)))$ among those two solutions (with ties broken arbitrarily). Fixing μ to either value, adopting the notation for $\mathbf{w}$ and w_0 of Algorithm 2 and ignoring constant terms, we can write the problem as

$$\begin{aligned} & \underset{\mathbf{z}}{\text{maximize}} && \mu \cdot \left[\sum_{m \in [K]} \mathbf{w}_m^\top \mathbf{z}_m \right] \\ & \text{subject to} && d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n) = \delta \\ & && \mathbf{z} \in \mathbb{Z}. \end{aligned}$$

The rectangularity of $\mathbb{Z}$ implies that the decisions of this problem admit a decomposition into the selection $\mathcal{M} \subseteq [K]$ of δ discrete features $m \in [K]$ along which $\mathbf{z}_m$ differs from $\mathbf{z}_m^n$ and, for those features $m \in [K]$ where $\mathbf{z}_m \neq \mathbf{z}_m^n$, the choice of $\mathbf{z}_m \in \mathbb{Z}(k_m) \setminus \{\mathbf{z}_m^n\}$:

$$\begin{aligned} & \underset{\mathcal{M}}{\text{maximize}} && \max_{\mathbf{z}} \left\{ \mu \cdot \left[\sum_{m \in [K]} \mathbf{w}_m^\top \mathbf{z}_m \right] : \left[\begin{array}{ll} \mathbf{z}_m \in \mathbb{Z}(k_m) \setminus \{\mathbf{z}_m^n\} & \forall m \in \mathcal{M} \\ \mathbf{z}_m = \mathbf{z}_m^n & \forall m \in [K] \setminus \mathcal{M} \end{array} \right] \right\} \\ & \text{subject to} && \mathcal{M} \subseteq [K], \quad |\mathcal{M}| = \delta. \end{aligned}$$

Noticing that the embedded maximization problem decomposes along the discrete features $m \in [K]$, we can adopt the notation for $\mathbf{z}_m^\star$ of Algorithm 2 to obtain the equivalent formulation

$$\begin{aligned} & \underset{\mathcal{M}}{\text{maximize}} \quad \left[\sum_{m \in \mathcal{M}} \mu \cdot \mathbf{w}_m^\top \mathbf{z}_m^\star \right] + \left[\sum_{m \in [K] \setminus \mathcal{M}} \mu \cdot \mathbf{w}_m^\top \mathbf{z}_m^n \right] \\ & \text{subject to} \quad \mathcal{M} \subseteq [K], \quad |\mathcal{M}| = \delta. \end{aligned}$$

The two summations in the objective function of this problem admit the reformulation

$$\left[\sum_{m \in \mathcal{M}} \mu \cdot \mathbf{w}_m^\top \mathbf{z}_m^\star \right] + \left[\sum_{m \in [K] \setminus \mathcal{M}} \mu \cdot \mathbf{w}_m^\top \mathbf{z}_m^n \right] = \left[\sum_{m \in [K]} \mu \cdot \mathbf{w}_m^\top \mathbf{z}_m^n \right] + \left[\sum_{m \in \mathcal{M}} \mu \cdot \mathbf{w}_m^\top (\mathbf{z}_m^\star - \mathbf{z}_m^n) \right],$$

and ignoring constant terms once more simplifies our optimization problem to

$$\begin{aligned} & \underset{\mathcal{M}}{\text{maximize}} \quad \sum_{m \in \mathcal{M}} \mu \cdot \mathbf{w}_m^\top (\mathbf{z}_m^\star - \mathbf{z}_m^n) \\ & \text{subject to} \quad \mathcal{M} \subseteq [K], \quad |\mathcal{M}| = \delta. \end{aligned}$$

This problem is solved by identifying $\mathcal{M}$ with the indices of the δ largest elements of the sequence $\mu \cdot \mathbf{w}_1^\top (\mathbf{z}_1^\star - \mathbf{z}_1^n), \dots, \mu \cdot \mathbf{w}_K^\top (\mathbf{z}_K^\star - \mathbf{z}_K^n)$, and this problem can be solved by a simple sorting algorithm. One readily verifies that Algorithm 2 adopts the solution approach just described to determine a maximally violated constraint index $\mathbf{z}(n, i)$.

The runtime of Algorithm 2, finally, is dominated by determining the $2K$ maximizers $\mathbf{z}_m^\star$, $m \in [K]$ and $\mu \in \{\pm 1\}$, which takes time $\mathcal{O}(M_{\mathbf{z}})$ due to the one-hot encoding employed by $\mathbb{Z}$, sorting the $2K$ values $\mu \cdot \mathbf{w}_m^\top (\mathbf{z}_m^\star - \mathbf{z}_m^n)$, which takes time $\mathcal{O}(K \log K)$, as well as determining a maximally violated constraint among the $2K + 2$ candidates $\mathbf{z} \in \mathcal{Z}$, which takes time $\mathcal{O}(KT)$. $\square$

Proof of Theorem 3. Recall from Grötschel et al. (1988, Corollary 4.2.7) that problem (4) can be solved to δ -accuracy in polynomial time if (i) the problem admits a polynomial time weak separation oracle and if (ii) the feasible region of the problem is a circumscribed convex body. We prove both of these properties next.

Fix a convex and compact set $\mathcal{K} \subseteq \mathbb{R}^n$ and a rational number $\delta > 0$. Grötschel et al. (1988, Definition 2.1.13) define a *weak* separation oracle for $\mathcal{K}$ as an algorithm which for any vector $\mathbf{q} \in \mathbb{R}^n$ either confirms that $\mathbf{q} \in \mathcal{S}(\mathcal{K}, \delta)$, where $\mathcal{S}(\mathcal{K}, \delta) = \{\mathbf{r} \in \mathbb{R}^n : \|\mathbf{r} - \mathbf{r}'\|_2 \leq \delta \text{ for some } \mathbf{r}' \in \mathcal{K}\}$ is the δ -enclosure around $\mathcal{K}$ (*i.e.*, $\mathbf{q}$ is *almost* in $\mathcal{K}$), or finds a vector $\mathbf{c} \in \mathbb{R}^n$ with $\|\mathbf{c}\|_\infty = 1$ such that $\mathbf{c}^\top \mathbf{p} \leq \mathbf{c}^\top \mathbf{q} + \delta$ for all $\mathbf{p} \in \mathcal{S}(\mathcal{K}, -\delta)$, where $\mathcal{S}(\mathcal{K}, -\delta) = \{\mathbf{r} \in \mathcal{K} : \mathcal{S}(\{\mathbf{r}\}, \delta) \subseteq \mathcal{K}\}$ is the δ -interior of $\mathcal{K}$ (*i.e.*, $\mathbf{c}$ is an *almost* separating hyperplane). We construct a weak separation oracle for problem (4) as follows. Denote the feasible region of problem (4) by $\mathcal{K} \subseteq \Theta \times \mathbb{R}_+^N$, and fix any $\mathbf{q}' = (\boldsymbol{\theta}', \boldsymbol{\sigma}') \in \mathbb{R}^{M_\Theta} \times \mathbb{R}^N$. If $\sigma'_n < 0$ for any $n \in [N]$, then our algorithm returns the separating hyperplane $\mathbf{c}^\top = (\mathbf{0}^\top - \mathbf{e}_n^\top)$. If $\boldsymbol{\theta}' \notin \Theta$, then we can augment the weakly separating hyperplane $\mathbf{c} \in \mathbb{R}^{M_\Theta}$ for the set Θ to a weakly separating hyperplane $(\mathbf{c}^\top \mathbf{0}^\top)$ for problem (4). Otherwise, we leverage Algorithm 2 to obtain a *strong* separation oracle for $\mathcal{K}$, which Grötschel et al. (1988, Definition 2.1.4) define as an algorithm which for any vector $\mathbf{q} \in \mathbb{R}^n$ either confirms that $\mathbf{q} \in \mathcal{K}$ or finds a vector $\mathbf{c} \in \mathbb{R}^n$ such that $\mathbf{c}^\top \mathbf{p} < \mathbf{c}^\top \mathbf{q}$ for all $\mathbf{p} \in \mathcal{K}$. To this end, we execute Algorithm 2 for all constraint group indices $(n, i) \in [N] \times \mathcal{I}$ to identify most violated constraint indices $\mathbf{z}(n, i)$ along with the constraint violations $\vartheta(n, i)$. If $\vartheta(n, i) \leq 0$ for all $(n, i) \in [N] \times \mathcal{I}$, then $\mathbf{q}'$ is feasible in problem (4) and our oracle terminates. If $\vartheta(n^*, i^*) > 0$ for some $(n^*, i^*) \in [N] \times \mathcal{I}$, then our oracle computes a subgradient $\mathbf{c}$ of the function

$$r(\boldsymbol{\theta}, \boldsymbol{\sigma}) = f_{n^* i^*}(g_{n^* i^*}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^{n^*}; \mathbf{z}(n^*, i^*))) - h_{n^* i^*}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^{n^*}; d_{\mathbf{z}}(\mathbf{z}(n^*, i^*), \mathbf{z}^{n^*})) - \sigma_{n^*}$$

at $(\boldsymbol{\theta}, \boldsymbol{\sigma}) = (\boldsymbol{\theta}', \boldsymbol{\sigma}')$. We have

$$\mathbf{c}^\top(\boldsymbol{\theta}, \boldsymbol{\sigma}) \leq \mathbf{c}^\top(\boldsymbol{\theta}', \boldsymbol{\sigma}') + r(\boldsymbol{\theta}, \boldsymbol{\sigma}) - r(\boldsymbol{\theta}', \boldsymbol{\sigma}') < \mathbf{c}^\top(\boldsymbol{\theta}', \boldsymbol{\sigma}') \quad \forall (\boldsymbol{\theta}, \boldsymbol{\sigma}) \in \mathcal{K},$$

where the first inequality is due to the definition of subgradients and the fact that r is convex, and the second inequality holds since $r(\boldsymbol{\theta}', \boldsymbol{\sigma}') > 0$ while $r(\boldsymbol{\theta}, \boldsymbol{\sigma}) \leq 0$ for all $(\boldsymbol{\theta}, \boldsymbol{\sigma}) \in \mathcal{K}$. In conclusion, $\mathbf{c}$ is a strong separating hyperplane for $(\boldsymbol{\theta}', \boldsymbol{\sigma}')$ as desired. Note that by construction, a strong separation oracle is also a weak separation oracle. Finally, the assumptions in the statement of

our theorem, together with Theorem 2, imply that our oracle runs in polynomial time.

We next turn to the claim that the feasible region $\mathcal{K} \subseteq \Theta \times \mathbb{R}_+^N$ of problem (4) is a circumscribed convex body. According to Grötschel et al. (1988, Definition 2.1.16), this is the case whenever $\mathcal{K}$ is a finite-dimensional, full-dimensional, closed and convex subset of a ball whose finite radius we can specify. By construction, $\mathcal{K}$ is finite-dimensional, closed and convex. Moreover, $\mathcal{K}$ is full-dimensional since Θ is full-dimensional by construction and the constraints have an epigraphical structure. To see that $\mathcal{K}$ can be circumscribed by a ball whose finite radius we can specify, we note that Θ is bounded by assumption and $\boldsymbol{\sigma}$ is non-negative by construction. Since the objective function f_0 in (4) is non-decreasing in $\boldsymbol{\sigma}$, we can without loss of generality include in (4) the additional constraints $\sigma_n \leq \bar{\sigma}_n$ for

$$\bar{\sigma}_n = \max_{\boldsymbol{\theta} \in \Theta} \max_{i \in \mathcal{I}} \max_{\mathbf{z} \in \mathbb{Z}} \{f_{ni}(g_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; \mathbf{z})) - h_{ni}(\boldsymbol{\theta}, \boldsymbol{\xi}_{-\mathbf{z}}^n; d_{\mathbf{z}}(\mathbf{z}, \mathbf{z}^n))\}, \quad n \in [N].$$

Note that all $\bar{\sigma}_n$ are finite since Θ is compact, $\mathcal{I}$ and $\mathbb{Z}$ are finite sets and the objective function is continuous. We thus conclude that $\boldsymbol{\sigma}$ can be bounded as well, and thus $\mathcal{K}$ can indeed be circumscribed by a ball whose finite radius we can specify. $\square$

Proof of Theorem 4. Fix any Wasserstein classification or regression instance as described in the statement of the theorem, fix any ambiguity radius $\epsilon > \kappa_{\mathbf{z}} K^{1/p}$, any non-zero numbers $M_{\mathbf{x}}$ of continuous and K of discrete features, and denote the training set as $\{\boldsymbol{\xi}^n\}_{n \in [N]}$ with $\boldsymbol{\xi}^n = (\mathbf{x}^n, \mathbf{z}^n, y^n)$, $n \in [N]$. Our proof proceeds in three steps. We first derive a closed-form expression for the objective function of the Wasserstein learning problem (1) at a judiciously chosen model $\hat{\boldsymbol{\beta}}$. Our derivation will show that this objective function constitutes the sum of the empirical loss $N^{-1} \cdot \sum_{n \in [N]} l_{\hat{\boldsymbol{\beta}}}(\boldsymbol{\xi}^n)$ and a function $h_{\hat{\boldsymbol{\beta}}}(\{\boldsymbol{\xi}^n\}_{n \in [N]})$. We then construct two sets of data points $\{\hat{\boldsymbol{\xi}}^n\}_{n \in [N]}$ and $\{\check{\boldsymbol{\xi}}^n\}_{n \in [N]}$ at which $h_{\hat{\boldsymbol{\beta}}}(\{\hat{\boldsymbol{\xi}}^n\}_{n \in [N]}) \neq h_{\hat{\boldsymbol{\beta}}}(\{\check{\boldsymbol{\xi}}^n\}_{n \in [N]})$, showing that $h_{\hat{\boldsymbol{\beta}}}(\{\boldsymbol{\xi}^n\}_{n \in [N]})$ exhibits a dependence on the dataset $\{\boldsymbol{\xi}^n\}_{n \in [N]}$ that cannot be replicated by any data-agnostic regularizer $\mathfrak{R}(\hat{\boldsymbol{\beta}})$.

Fix the model $\hat{\boldsymbol{\beta}} = (\hat{\beta}_0, \hat{\boldsymbol{\beta}}_{\mathbf{x}}, \hat{\boldsymbol{\beta}}_{\mathbf{z}})$ with $\hat{\beta}_0 \in \mathbb{R}$ and $\hat{\boldsymbol{\beta}}_{\mathbf{x}} \neq \mathbf{0}$ selected arbitrarily. Fix any discrete

feature realization $\mathbf{z}^* \in \mathbb{Z} \setminus \{\mathbf{z}^1\}$ as well as the discrete feature coefficients

$$\hat{\beta}_{\mathbf{z},mi} = \begin{cases} -\frac{2}{d'} \cdot \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1) & \text{if } z_{mi}^1 = 1, \\ \frac{2}{d'} \cdot \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1) & \text{otherwise} \end{cases} \quad \text{for all } m \in [K] \text{ and } i \in [k_m],$$

where $d' = (d/de)L(e) \big|_{e=x'}$ is the derivative of the loss function L at any point $x' \in \mathbb{R}$ where the derivative does not vanish. Such points x' exist due to Rademacher's theorem, which ensures that a Lipschitz continuous function is differentiable almost everywhere, as well as the assumption that the loss function L is non-constant. We make two observations that we will leverage later on in this proof:

- (i) We have $d' \cdot \hat{\beta}_{\mathbf{z},mi} \cdot (z_{mi} - z_{mi}^1) \geq 0$ for all $\mathbf{z} \in \mathbb{Z}$, $m \in [K]$ and $i \in [k_m]$. Indeed, fix any $\mathbf{z} \in \mathbb{Z}$, $m \in [K]$ and $i \in [k_m]$. If $z_{mi}^1 = 0$, then $d' \cdot \hat{\beta}_{\mathbf{z},mi} = 2 \cdot \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1) \geq 0$ and $z_{mi} - z_{mi}^1 \geq 0$. If $z_{mi}^1 = 1$, on the other hand, then $d' \cdot \hat{\beta}_{\mathbf{z},mi} = -2 \cdot \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1) \leq 0$ and $z_{mi} - z_{mi}^1 \leq 0$. In particular, we have $d' \cdot \hat{\beta}_{\mathbf{z}}^\top (\mathbf{z} - \mathbf{z}^1) \geq 0$ for all $\mathbf{z} \in \mathbb{Z}$.
- (ii) We have $d' \cdot \hat{\beta}_{\mathbf{z}}^\top (\mathbf{z}^* - \mathbf{z}^1) > \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1)$. Indeed, since $\mathbf{z}^* \neq \mathbf{z}^1$, there is at least one $m^* \in [K]$ and $i^* \in [k_{m^*}]$ where $z_{m^*i^*}^* - z_{m^*i^*}^1 \neq 0$, and thus $d' \cdot \hat{\beta}_{\mathbf{z},m^*i^*} \cdot (z_{m^*i^*}^* - z_{m^*i^*}^1) = 2 \cdot \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1) > \kappa_{\mathbf{z}} \text{lip}(L) \cdot \|\hat{\beta}_{\mathbf{x}}\|_* \cdot d_{\mathbf{z}}(\mathbf{z}^*, \mathbf{z}^1)$. The claim now follows from the fact that $d' \cdot \hat{\beta}_{\mathbf{z},mi} \cdot (z_{mi} - z_{mi}^1) \geq 0$ for all other $m \in [K]$ and $i \in [k_m]$ thanks to our previous observation (i).

Note that piece-wise affine loss functions are Lipschitz continuous. For our selected problem instance, Proposition 3 from Appendix A therefore implies that the objective function of the

Wasserstein classification problem becomes

$$\begin{aligned}
& \underset{\lambda, \mathbf{s}}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\
& \text{subject to} && \left. \begin{aligned} & l_{\hat{\beta}}(\mathbf{x}^n, \mathbf{z}, y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) \leq s_n \\ & l_{\hat{\beta}}(\mathbf{x}^n, \mathbf{z}, -y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) - \lambda\kappa_y \leq s_n \end{aligned} \right\} \forall n \in [N], \forall \mathbf{z} \in \mathbb{Z} \\
& && \text{lip}(L) \cdot \|\hat{\beta}_x\|_* \leq \lambda \\
& && \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N,
\end{aligned}$$

and Proposition 7 from Appendix B implies that the objective function of the Wasserstein regression problem becomes

$$\begin{aligned}
& \underset{\lambda, \mathbf{s}}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} s_n \\
& \text{subject to} && l_{\hat{\beta}}(\mathbf{x}^n, \mathbf{z}, y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) \leq s_n \quad \forall n \in [N], \forall \mathbf{z} \in \mathbb{Z} \\
& && \text{lip}(L) \cdot \|\hat{\beta}_x\|_* \leq \lambda \\
& && \text{lip}(L) \leq \lambda\kappa_y \\
& && \lambda \in \mathbb{R}_+, \quad \mathbf{s} \in \mathbb{R}_+^N.
\end{aligned}$$

When $\kappa_y \rightarrow \infty$, the second set of constraints in the classification problem and the third constraint in the regression problem become redundant since $\lambda \geq \text{lip}(L) \cdot \|\hat{\beta}_x\|_* > 0$ because $\hat{\beta}_x \neq \mathbf{0}$. For all $n \in [N]$, we can then replace each s_n with a maximum of the left-hand side of the first constraint set over all $\mathbf{z} \in \mathbb{Z}$ in either problem to obtain the unified formulation

$$\begin{aligned}
& \underset{\lambda}{\text{minimize}} && \lambda\epsilon + \frac{1}{N} \sum_{n \in [N]} \max_{\mathbf{z} \in \mathbb{Z}} \left\{ l_{\hat{\beta}}(\mathbf{x}^n, \mathbf{z}, y^n) - \lambda\kappa_z d_z(\mathbf{z}, \mathbf{z}^n) \right\} \\
& \text{subject to} && \text{lip}(L) \cdot \|\hat{\beta}_x\|_* \leq \lambda \\
& && \lambda \in \mathbb{R}_+
\end{aligned}$$

of the objective function of both the classification and the regression problem. Note that the non-negativity of s_n for $n \in [N]$ is preserved in the unified formulation since $d_z(\mathbf{z}^n, \mathbf{z}^n) = 0$

and $l_{\hat{\beta}}(\mathbf{x}^n, \mathbf{z}^n, y^n) \geq 0$ for all $n \in [N]$ by definition of the loss function L that underlies $l_{\hat{\beta}}$. We claim that $\lambda^* = \text{lip}(L) \cdot \|\hat{\beta}_x\|_*$ at optimality. Indeed, any increment $\Delta\lambda > 0$ in λ will cause an increase of $\Delta\lambda \cdot \epsilon$ and a maximum decrease of $\Delta\lambda \cdot \kappa_z K^{1/p}$ in the objective function, and we have $\epsilon > \kappa_z K^{1/p}$ by assumption. Hence, the objective function of the Wasserstein classification and regression problem simplifies to

$$f_1(\hat{\beta}, \{\xi^n\}_{n \in [N]}) = \lambda^* \epsilon + \frac{1}{N} \sum_{n \in [N]} \max_{z \in \mathbb{Z}} \left\{ l_{\hat{\beta}}(\mathbf{x}^n, z, y^n) - \lambda^* \kappa_z d_z(z, \mathbf{z}^n) \right\} = \frac{1}{N} \sum_{n \in [N]} l_{\hat{\beta}}(\xi^n) + h_{\hat{\beta}}(\{\xi^n\}_{n \in [N]}),$$

$$\text{where } h_{\hat{\beta}}(\{\xi^n\}_{n \in [N]}) = \lambda^* \epsilon + \frac{1}{N} \sum_{n \in [N]} \max_{z \in \mathbb{Z}} \left\{ l_{\hat{\beta}}(\mathbf{x}^n, z, y^n) - l_{\hat{\beta}}(\mathbf{x}^n, \mathbf{z}^n, y^n) - \lambda^* \kappa_z d_z(z, \mathbf{z}^n) \right\}.$$

In contrast, the objective function of a generic regularized learning problem has the form

$$f_2(\hat{\beta}, \{\xi^n\}_{n \in [N]}) = \frac{1}{N} \sum_{n \in [N]} l_{\hat{\beta}}(\xi^n) + \mathfrak{R}(\hat{\beta}).$$

By construction, $\mathfrak{R}(\hat{\beta})$ does not vary with $\{\xi^n\}_{n \in [N]}$. In contrast, we claim that $h_{\hat{\beta}}(\{\xi^n\}_{n \in [N]})$ varies with $\{\xi^n\}_{n \in [N]}$. To this end, we will construct two sets of data points $\{\hat{\xi}^n\}_{n \in [N]}$ and $\{\check{\xi}^n\}_{n \in [N]}$ at which $h_{\hat{\beta}}(\{\hat{\xi}^n\}_{n \in [N]}) \neq h_{\hat{\beta}}(\{\check{\xi}^n\}_{n \in [N]})$. The two datasets $\{\hat{\xi}^n\}_{n \in [N]}$ and $\{\check{\xi}^n\}_{n \in [N]}$ are identical to $\{\xi^n\}_{n \in [N]}$ except for the realization $\mathbf{x}^1$ of the continuous features in sample 1. In other words, we have $(\hat{\mathbf{x}}^n, \hat{\mathbf{z}}^n, \hat{y}^n) = (\check{\mathbf{x}}^n, \check{\mathbf{z}}^n, \check{y}^n) = (\mathbf{x}^n, \mathbf{z}^n, y^n)$ for all $n \in [N] \setminus 1$ as well as $(\hat{\mathbf{z}}^1, \hat{y}^1) = (\check{\mathbf{z}}^1, \check{y}^1) = (\mathbf{z}^1, y^1)$.

We select $\hat{\mathbf{x}}^1$ such that the maximum in the definition of $h_{\hat{\beta}}$ is strictly positive at sample $n = 1$. To achieve this, we choose $\hat{\mathbf{x}}^1$ such that $\hat{y}^1 \cdot [\hat{\beta}_0 + \hat{\beta}_x^\top \hat{\mathbf{x}}^1 + \hat{\beta}_z^\top \hat{\mathbf{z}}^1] = x'$ for classification problems and $\hat{\beta}_0 + \hat{\beta}_x^\top \hat{\mathbf{x}}^1 + \hat{\beta}_z^\top \hat{\mathbf{z}}^1 - \hat{y}^1 = x'$ for regression problems, respectively, where x' was defined earlier as a point where the derivative of the loss function does not vanish. Note that this is always possible since $\hat{\beta}_x \neq \mathbf{0}$. We then have

$$\max_{z \in \mathbb{Z}} l_{\hat{\beta}}(\hat{\mathbf{x}}^1, z, \hat{y}^1) - l_{\hat{\beta}}(\hat{\mathbf{x}}^1, \hat{\mathbf{z}}^1, \hat{y}^1) - \lambda^* \kappa_z d_z(z, \hat{\mathbf{z}}^1) \geq l_{\hat{\beta}}(\hat{\mathbf{x}}^1, \mathbf{z}^*, \hat{y}^1) - l_{\hat{\beta}}(\hat{\mathbf{x}}^1, \hat{\mathbf{z}}^1, \hat{y}^1) - \lambda^* \kappa_z d_z(\mathbf{z}^*, \hat{\mathbf{z}}^1)$$

$$\begin{aligned}
&= L(x' + \hat{\beta}_z^\top(z^\star - \hat{z}^1)) - L(x') - \lambda^\star \kappa_z d_z(z^\star, \hat{z}^1) \\
&\geq \hat{\beta}_z^\top(z^\star - \hat{z}^1) \cdot \frac{d}{de} L(e) \big|_{e=x'} - \lambda^\star \kappa_z d_z(z^\star, \hat{z}^1) \\
&= d' \cdot \hat{\beta}_z^\top(z^\star - \hat{z}^1) - \kappa_z \text{lip}(L) \cdot \|\hat{\beta}_x\|_* \cdot d_z(z^\star, \hat{z}^1) > 0,
\end{aligned}$$

where the first inequality holds since $z^\star \in \mathbb{Z}$, the first identity uses the definitions of $l_{\hat{\beta}}$ and $\hat{x}^1$, the second inequality exploits the convexity of L , the second identity uses the definition of d' as well as $\lambda^\star$, and the third inequality follows from our earlier observation (ii). Since the term inside the maximum in the definition of $h_{\hat{\beta}}$ is strictly positive for $z = z^\star$ at sample $n = 1$, the maximum is guaranteed to be positive at sample $n = 1$ as well.

We choose $\tilde{x}^1$ such that the term inside the maximum in the definition of $h_{\hat{\beta}}$ is strictly negative at sample $n = 1$ for all $z \in \mathbb{Z} \setminus \{\tilde{z}^1\}$. Consider the case where $d' > 0$; the alternative case where $d' < 0$ follows from analogous arguments. Since L is non-negative and convex, we have

$$\lim_{x \rightarrow -\infty} \frac{d}{de} L(e) \big|_{e=x} \leq 0, \quad (\text{E-5})$$

which in turn implies that for all $z \in \mathbb{Z} \setminus \{\tilde{z}^1\}$, we have

$$\lim_{x \rightarrow -\infty} L(x + \hat{\beta}_z^\top(z - \tilde{z}^1)) - L(x) \leq \lim_{x \rightarrow -\infty} \hat{\beta}_z^\top(z - \tilde{z}^1) \cdot \frac{d}{de} L(e) \big|_{e=x+\hat{\beta}_z^\top(z-\tilde{z}^1)} \leq 0, \quad (\text{E-6})$$

where the first inequality exploits the fact that the derivative of a convex function is non-decreasing and the second inequality holds because $d' \cdot \hat{\beta}_z^\top(z - \tilde{z}^1) \geq 0$ due to our earlier observation (i) and the fact that $\lim_{x \rightarrow -\infty} \frac{d}{de} L(e) \big|_{e=x+\hat{\beta}_z^\top(z-\tilde{z}^1)} = \lim_{x \rightarrow -\infty} \frac{d}{de} L(e) \big|_{e=x} \leq 0$ by equation (E-5). Moreover, for $n = 1$ and for all $z \in \mathbb{Z} \setminus \{\tilde{z}^1\}$, we have

$$\begin{aligned}
&\lim_{\hat{\beta}_x^\top \tilde{x}^1 \rightarrow -\infty} l_{\hat{\beta}}(\tilde{x}^1, z, \tilde{y}^1) - l_{\hat{\beta}}(\tilde{x}^1, \tilde{z}^1, \tilde{y}^1) - \lambda^\star \kappa_z d_z(z, \tilde{z}^1) \\
&= \lim_{x'' \rightarrow -\infty} L(x'' + \hat{\beta}_z^\top(z - \tilde{z}^1)) - L(x'') - \lambda^\star \kappa_z d_z(z, \tilde{z}^1) \leq -\lambda^\star \kappa_z d_z(z, \tilde{z}^1) < 0,
\end{aligned}$$

where the identity applies the change of variables $\tilde{y}^1 \cdot [\hat{\beta}_0 + \hat{\beta}_x^\top \tilde{x}^1 + \hat{\beta}_z^\top \tilde{z}^1] = x''$ for classification problems and $\hat{\beta}_0 + \hat{\beta}_x^\top \tilde{x}^1 + \hat{\beta}_z^\top \tilde{z}^1 - \tilde{y}^1 = x''$ for regression problems, respectively, the first inequality uses (E-6), and the last inequality is due to the fact that λ^* , κ_z and $d_z(z, \tilde{z}^1)$ are all strictly positive. Since the term inside the maximum in the definition of $h_{\hat{\beta}}$ is strictly negative at sample $n = 1$ for all $z \in \mathbb{Z} \setminus \{\tilde{z}^1\}$, the maximum is guaranteed to evaluate to 0 for $n = 1$. Given that the datasets $\{\hat{\xi}^n\}_{n \in [N]}$ and $\{\check{\xi}^n\}_{n \in [N]}$ are identical for all other samples $n \in [N] \setminus 1$, this implies that $h_{\hat{\beta}}(\{\hat{\xi}^n\}_{n \in [N]}) \neq h_{\hat{\beta}}(\{\check{\xi}^n\}_{n \in [N]})$ and confirms that $h_{\hat{\beta}}(\{\xi^n\}_{n \in [N]})$ has a dependency on the dataset $\{\xi^n\}_{n \in [N]}$ that cannot be replicated by any data-agnostic regularizer $\mathfrak{R}(\hat{\beta})$. $\square$

Proof of Proposition 2. We prove the statement of the proposition in three steps. We first show that our revised ground metric d_u coincides with the original ground metric d_z over all pairs of discrete feature vectors $z, z' \in \mathbb{Z}$. This implies that the left-hand sides of the constraints indexed by $u \in \mathbb{Z}$ in (3') reduce to their counterparts in (3). In other words, problem (3') is a restriction of (3) since it contains all of the constraints of (3). We then prove that $d_u(u, z')$ is concave in u for every fixed $z' \in \mathbb{Z}$. Our third step leverages this intermediate result to show that every feasible solution to our mixed-feature Wasserstein learning problem (3) is also feasible in the bounded continuous-feature formulation (3').

In view of the first step, fix any $z, z' \in \mathbb{Z}$ and observe that for any $m \in [K]$, we have

$$\frac{1}{2} \|z_m - z'_m\|_1 + \frac{1}{2} |\mathbf{1}^\top (z_m - z'_m)| = \frac{1}{2} \left\| \begin{pmatrix} z_m \\ 1 - \mathbf{1}^\top z_m \end{pmatrix} - \begin{pmatrix} z'_m \\ 1 - \mathbf{1}^\top z'_m \end{pmatrix} \right\|_1 = \begin{cases} 0 & \text{if } z_m = z'_m, \\ 1 & \text{otherwise.} \end{cases}$$

Here, the first identity applies basic algebraic manipulations, and the second identity holds since $(z_m^\top, 1 - \mathbf{1}^\top z_m)$ and $(z'_m{}^\top, 1 - \mathbf{1}^\top z'_m)$ are canonic basis vectors in $\mathbb{R}^{k_m}$. We thus conclude the claim:

$$d_u(z, z') = \left(\sum_{m \in [K]} \frac{1}{2} \|z_m - z'_m\|_1 + \frac{1}{2} |\mathbf{1}^\top (z_m - z'_m)| \right)^{1/p} = \sum_{m \in [K]} (\mathbf{1}[z_m \neq z'_m])^{1/p} = d_z(z, z').$$

As for the second statement, fix any $z' \in \mathbb{Z}$. We then find that

$$d_u(\mathbf{u}, \mathbf{z}') = \left(\frac{1}{2} \sum_{m \in [K]} \sum_{j \in [k_m-1]} [\mathbb{1}[z'_{mj} = 0] \cdot u_{mj} + \mathbb{1}[z'_{mj} = 1] \cdot (1 - u_{mj})] + \frac{1}{2} \sum_{m \in [K]} [\mathbb{1}[\mathbf{z}'_m = \mathbf{0}] \cdot \mathbf{1}^\top \mathbf{u}_m + \mathbb{1}[\mathbf{z}'_m \neq \mathbf{0}] \cdot (1 - \mathbf{1}^\top \mathbf{u}_m)] \right)^{1/p},$$

where the first summation evaluates the 1-norm differences between $\mathbf{u}_m$ and $\mathbf{z}'_m$ and the second summation computes the absolute values in d_u , respectively. The above reformulation shows that the mapping $\mathbf{u} \mapsto d_u(\mathbf{u}, \mathbf{z}')$ can be represented as $f(g(\mathbf{u}))$, where $f(x) = x^{1/p}$ and $g : \mathbb{U} \rightarrow \mathbb{R}$ is affine. It follows from §3.2.2 of Boyd and Vandenberghe (2004) that $\mathbf{u} \mapsto d_u(\mathbf{u}, \mathbf{z}')$ is concave.

In view of the third statement, finally, fix any feasible solution $(\beta, \lambda, \mathbf{s})$ to problem (3). Since the objective functions of (3) and (3') coincide, we only need to show that $(\beta, \lambda, \mathbf{s})$ satisfies the constraints of problem (3'). To this end, we observe that for all $n \in [N]$, we have

$$\begin{aligned} & \sup_{(\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y}} \{l_\beta(\mathbf{x}, \mathbf{u}, y) - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n)\} - \lambda \kappa_z d_u(\mathbf{u}, \mathbf{z}^n) \leq s_n & \forall \mathbf{u} \in \mathbb{U} \\ \iff & \sup_{\mathbf{u} \in \mathbb{U}} \{l_\beta(\mathbf{x}, \mathbf{u}, y) - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n)\} - \lambda \kappa_z d_u(\mathbf{u}, \mathbf{z}^n) \leq s_n & \forall (\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y} \\ \iff & \sup_{\mathbf{u} \in \mathbb{U}} \{l_\beta(\mathbf{x}, \mathbf{u}, y) - \lambda \kappa_z d_u(\mathbf{u}, \mathbf{z}^n)\} - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n) \leq s_n & \forall (\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y} \\ \iff & \sup_{\mathbf{z} \in \mathbb{Z}} \{l_\beta(\mathbf{x}, \mathbf{z}, y) - \lambda \kappa_z d_u(\mathbf{z}, \mathbf{z}^n)\} - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n) \leq s_n & \forall (\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y} \\ \iff & \sup_{\mathbf{z} \in \mathbb{Z}} \{l_\beta(\mathbf{x}, \mathbf{z}, y) - \lambda \kappa_z d_z(\mathbf{z}, \mathbf{z}^n)\} - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n) \leq s_n & \forall (\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y} \\ \iff & \sup_{(\mathbf{x}, y) \in \mathbb{X} \times \mathbb{Y}} \{l_\beta(\mathbf{x}, \mathbf{z}, y) - \lambda \kappa_z d_z(\mathbf{z}, \mathbf{z}^n)\} - \lambda \|\mathbf{x} - \mathbf{x}^n\| - \lambda \kappa_y d_y(y, y^n) \leq s_n & \forall \mathbf{z} \in \mathbb{Z}. \end{aligned}$$

Here, the first, the second and the last equivalences reorder terms. As for the third equivalence, recall that $d_u(\mathbf{u}, \mathbf{z}^n)$ is concave in $\mathbf{u}$ for every fixed $\mathbf{z}^n \in \mathbb{Z}$. Thus, the expression inside the supremum in the third row is convex in $\mathbf{u}$, which implies that it attains its maximum at an extreme point of $\mathbb{U}$. The equivalence then follows from the fact that $\mathbb{Z}$ coincides with the extreme points of $\mathbb{U}$. The fourth equivalence, finally, leverages the first step of this proof, which showed that d_u and d_z agree over all pairs of discrete feature vectors $\mathbf{z}, \mathbf{z}' \in \mathbb{Z}$. $\square$