

Robust Regression via Model Based Methods^{*}

Armin Moharrer^[0000-0002-8374-7286], Khashayar Kamran^[0000-0002-4086-1038],
Edmund Yeh^[0000-0002-9544-1567], and Stratis Ioannidis^[0000-0001-8355-4751]

Northeastern University, Boston MA 02115, USA
{amoharrer,kamrank,eyeh,ioannidis}@ece.neu.edu

Abstract. The mean squared error loss is widely used in many applications, including auto-encoders, multi-target regression, and matrix factorization, to name a few. Despite computational advantages due to its differentiability, it is not robust to outliers. In contrast, ℓ_p norms are known to be robust, but cannot be optimized via, e.g., stochastic gradient descent, as they are non-differentiable. We propose an algorithm inspired by so-called model-based optimization (MBO) [36, 37], which replaces a non-convex objective with a convex model function and alternates between optimizing the model function and updating the solution. We apply this to robust regression, proposing SADM, a stochastic variant of the Online Alternating Direction Method of Multipliers (OADM) [51] to solve the inner optimization in MBO. We show that SADM converges with the rate $O(\log T/T)$. Finally, we demonstrate experimentally (a) the robustness of ℓ_p norms to outliers and (b) the efficiency of our proposed model-based algorithms in comparison with gradient methods on autoencoders and multi-target regression.

1 Introduction

Mean Squared Error (MSE) loss problems are ubiquitous in machine learning and data mining. Such problems have the following form:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \|F(\theta; \mathbf{x}_i)\|_2^2 + g(\theta), \quad (1)$$

where function $F : \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}^N$ captures the contribution of a sample $\mathbf{x}_i \in \mathbb{R}^m$, $i = 1, \dots, n$, to the objective under the parameter $\theta \in \mathbb{R}^d$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a regularizer. Example applications include training auto-encoders [18, 28], matrix factorization [16], and multi-target regression [49].

The MSE loss in (1) is computationally convenient, as the resulting problem is smooth and can thus be optimized efficiently via gradient methods, such as stochastic gradient descent (SGD). However, it is well-known that the MSE loss is not robust to *outliers* [8, 15, 19, 20, 35], i.e., samples far from the dataset mean.

^{*} The authors gratefully acknowledge support from the National Science Foundation (Grants CCF-1750539, IIS-1741197, and CNS-1717213), DARPA (Grant HR0011-17-C-0050), and a research grant from American Tower Corp.

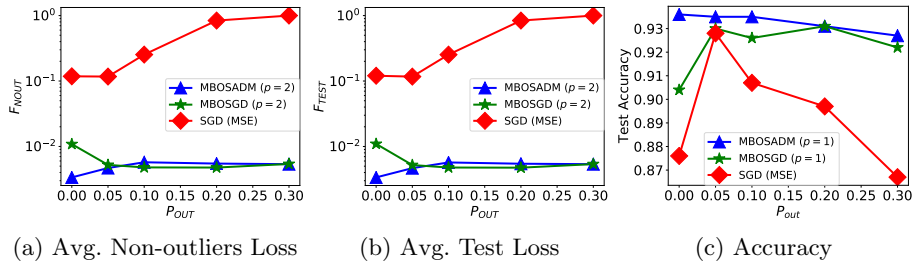


Fig. 1: Robustness of ℓ_p norms vs. MSE to outliers introduced to MNIST when training an autoencoder. Figures 1a and 1b show the average loss over the non-outliers and the test set, respectively; values in each figure are normalized w.r.t. the largest value. The test accuracy of a logistic regression on the latent features is shown in Fig. 1c. We see that, under MSE, both for the loss values and classification accuracy are significantly affected by the fraction of outliers P_{out} . Robust embeddings under $p = 1, 2$ norms optimized via our proposed MBO methods exhibit almost constant behavior w.r.t. P_{out} .

Intuitively, when squaring the error, outliers tend to dominate the objective. To mitigate the effect of outliers, a classic approach is to introduce robustness by replacing the squared error with either the ℓ_2 norm [8, 11, 18, 28, 35, 42] or the ℓ_1 norm [2, 5, 13, 19–21, 24, 41]. This has been applied to several applications, including feature selection [35, 42], PCA [2, 8, 21, 24], K-means clustering [11], training autoencoders [18, 28], matrix factorization [5, 13, 19, 20], and regression [41]. Motivated by this approach, we study the following robust variant of Problem (1):

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \|F(\theta; \mathbf{x}_i)\|_p + g(\theta), \quad (2)$$

where $\|\cdot\|_p$ denotes an ℓ_p norm ($p \geq 1$). We are particularly interested in cases where F is not affine and, in general, Problem (2) is non-convex. This includes, e.g., feature selection [35], matrix factorization [13, 20], auto-encoders [18], and deep multi-target regression [41, 49].

A significant challenge behind solving Prob. (2) is that its objective is not smooth, precisely because the ℓ_p norm is not differentiable at $\mathbf{0} \in \mathbb{R}^N$. For non-convex and non-smooth problems of the form (2), where the objective contains a composite function, *Model-Based Optimization* (MBO) methods [6, 7, 9, 10, 23, 36] come with good experimental performance as well as theoretical guarantees. In particular, these MBO methods define a convex (but non-smooth) approximation of the main objective, called the *model function*. They then iteratively optimize this model function plus a proximal quadratic term. Under certain conditions, MBO converges to a stationary point of the non-convex problem [23].

In this work, we use MBO to solve Problem (2) for arbitrary ℓ_p norms. In particular, each MBO iteration results in a convex optimization problem. We solve these sub-problems using a novel stochastic variant of the Online Alternating Direction Method (OADM) [50], which we call *Stochastic Alternating Direction*

Method (SADM). Using SADM is appealing, as its resulting steps have efficient gradient-free solutions; in particular, we exploit a bisection method [25,30] for finding the proximal operator of ℓ_p norms. We provide theoretical guarantees for SADM. As an additional benefit, SADM comes with a stopping criterion, which is hard to obtain for gradient methods when the objective is non-smooth [27].

Overall, we make the following contributions:

- We study a general outlier-robust optimization that replaces the MSE with ℓ_p norms. We show that such problems can be solved via Model-Based Optimization (MBO) methods.
- We propose SADM, i.e., a stochastic version of OADM, and show that under strong convexity of the regularizer g , it converges with a $O(\log T/T)$ rate when solving the sub-problems arising at each MBO iteration.
- We conduct extensive experiments on training auto-encoders and multi-target regression. We show (a) the higher robustness of ℓ_p norms in comparison with MSE and (b) the superior performance of MBO, against stochastic gradient methods, both in terms of minimizing the objective and performing down-stream classification tasks. In some cases, we see that the MBO variant using SADM obtains objectives that are $29.6\times$ smaller than the ones achieved by the competitors.

The performance of our MBO approach is illustrated in Fig. 1. An autoencoder trained via SGD over the MSE objective is significantly affected by the presence of outliers; in contrast, our MBO methods applied to ℓ_p objectives are robust to outliers. These relative benefits are also evident in a downstream classification task over the latent embeddings. The remainder of this paper is organized as follows. We review related work in Sec. 2. We introduce our robust formulation along with its applications in Sec. 3. We describe the instance of MBO applied to our problem in Sec. 4. We introduce SADM and its convergence analysis in Sec. 5 and present our experiments in Sec. 6. We finally conclude in Sec. 7.

2 Related Work

Robustness of ℓ_p Norms: To improve the sensitivity of MSE to outliers, Ding et al. [8] first suggested replacing the MSE with the ℓ_2 norm in the context of Principal Component Analysis (PCA). This motivated a line of research for developing robust algorithms using the ℓ_2 norm in different applications, e.g., non-negative matrix factorization [20], feature selection [35,42], training autoencoders [18], and k -means clustering [11]. Attaining robustness via the ℓ_1 norm has also been used in matrix factorization [5,13,19], PCA [2,21,24], and regression [41]. Robustness of the ℓ_1 norm can be linked to robustness of median to outliers in comparison to average value (see, e.g., Friedman et al. [15]). Our problem includes robust variations considered in, e.g., [13,18,20,35,41], as special cases. However, these earlier algorithms are tailored to specific ℓ_p norms and/or do not generalize beyond the studied objective or application (some works, e.g., [35,41], only consider convex problems). In contrast, we unify these

variations for different applications as a non-convex and non-smooth problem, and present a general optimization algorithm for arbitrary ℓ_p norms.

Non-smooth/non-convex Optimization: Non-smooth and non-convex optimization problems arise in many applications, such as non-negative matrix factorization [16], compressed sensing with non-convex norms [1], and ℓ_p norm regularized sparse regression problems [3, 33]. A class of non-smooth non-convex optimization problems, known as *weakly convex problems* [47], i.e., problems in which the objective function is the sum of a convex function and a quadratic function, have attracted a lot of attention [6, 10, 12, 22, 23, 27]. Mai and Johansson [27] provided novel theoretical guarantees on the convergence of stochastic gradient descent with momentum for weakly-convex functions. However, in our experiments in Sec. 6, we show that model-based methods considerably outperform these stochastic gradient methods with momentum.

Our approach falls under the class of *prox-linear* methods [6, 9, 10, 12, 22, 23], that solve problems where the objective is a composition of a non-smooth convex function and a smooth function, exactly as in Prob. (2). Such methods iteratively minimize the composition of the non-smooth function with the first-order approximation of the smooth function [6, 10, 12, 23]. Lewis and Wright [23] prove convergence to a stationary point while Drusvyatskiy et al. prove linear convergence [9] and obtain sample complexity guarantees [10]. Ochs et al. [36, 37] generalize prox-linear methods by proposing *Model-Based Optimization* (MBO) for both smooth and non-smooth non-convex problems. MBO reduces to a prox-linear method when the objective has a composite form, as in our case. Ochs et al. further considered non-quadratic proximal penalties in sub-problems and complemented MBO with an Armijo-like line search. We leverage both their line search and theoretical guarantees (c.f. Prop. 1); our main technical departure is in solving sub-problems per iteration via SADM, which we discuss next.

ADMM. The Alternating Direction Method of Multipliers (ADMM) [4] is a convex optimization algorithm that provides efficient methods for non-smooth problems. Applying ADMM often results in sub-problems that can be solved efficiently via proximal operators [4, 40, 46]. To speed up ADMM, stochastic variants [26, 38, 52] have been proposed for minimizing sum-like objectives. These stochastic variants, similar to SGD, update solutions using the gradients of a small batch of terms in the objective, at each iteration. Another group of works proposed online variants of ADMM [17, 45, 50]. In these variants, the goal is to minimize the summation of loss functions that are revealed by an adversary.

Wang and Banerjee [50] proposed the first online variant of ADMM, termed Online Direction Method of Multipliers (OADM). Here, we propose a stochastic version of OADM, Stochastic Alternating Direction Method (SADM), to solve inner-problems in MBO iterations. SADM is similar to OADM with the difference that functions are sampled uniformly at random and are not given by an adversary. We prove that SADM converges with a $O(\log T/T)$ rate when the regularizer is strongly convex. Other existing stochastic or online ADMM variants either require a smooth objective [26, 52] or bounded sub-gradients [38, 45], neither of which apply for the inner problems we solve. In contrast, we show that applying

SADM results in sub-problems that admit gradient-free efficient solutions via a bisection method for finding proximal operators of ℓ_p norms [25, 30].

3 Robust Regression and Applications

Notations. Lowercase boldface letters represent vectors, while capital boldface letters represent matrices. We also use the notation $[n] \triangleq \{1, 2, \dots, n\}$.

Robust Regression. We first extend Prob. (2) to include constraints via:

$$\min_{\boldsymbol{\theta}} \frac{1}{n} \sum_{i \in [n]} \|F(\boldsymbol{\theta}; \mathbf{x}_i)\|_p + g(\boldsymbol{\theta}) + \chi_{\mathcal{C}}(\boldsymbol{\theta}), \quad (3)$$

where, again, $F : \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}^N$ is smooth, $\|\cdot\|_p$ is the ℓ_p norm, $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex regularizer such that $\inf g > -\infty$, while $\chi_{\mathcal{C}} : \mathbb{R}^d \rightarrow \{0, \infty\}$ is the indicator function of the convex set $\mathcal{C} \subseteq \mathbb{R}^d$. In practice, we are often interested in cases where either the regularizer or the constraint is absent.

Applications. For the sake of concreteness, we introduce some applications of Prob. (3). Function g is typically either the lasso (i.e., the ℓ_1 norm $g(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_1$) or ridge regularizer (i.e., the ℓ_2 norm squared $g(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_2^2$). We thus focus on the definition of $F(\cdot; \cdot)$ and constraint set \mathcal{C} in each of these applications.

Auto-encoders [18]. Given n data points $\mathbf{x}_i \in \mathbb{R}^m$, $i \in [n]$, auto-encoders embed them in a m' -dimensional space, $m' \ll m$, as follows. The mapping to $\mathbb{R}^{m'}$ is done by a possibly non-linear function (e.g., a neural network) with d_{enc} parameters $F_{\text{enc}} : \mathbb{R}^{d_{\text{enc}}} \times \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$, called the *encoder*. An inverse mapping, the *decoder* $F_{\text{dec}} : \mathbb{R}^{d_{\text{dec}}} \times \mathbb{R}^{m'} \rightarrow \mathbb{R}^m$ with d_{dec} parameters re-constructs the original points given latent embeddings. Both the encoder and the decoder are trained jointly over a dataset $\{\mathbf{x}_i\}_{i=1}^n$ by minimizing the reconstruction error; cast in our robust setting, this amounts to minimizing (3) with

$$F(\boldsymbol{\theta}; \mathbf{x}_i) = \mathbf{x}_i - F_{\text{dec}}(\boldsymbol{\theta}_{\text{dec}}; F_{\text{enc}}(\boldsymbol{\theta}_{\text{enc}}; \mathbf{x}_i)), \quad (4)$$

where $\boldsymbol{\theta} = [\boldsymbol{\theta}_{\text{dec}}; \boldsymbol{\theta}_{\text{enc}}] \in \mathbb{R}^{d_{\text{enc}} + d_{\text{dec}}}$ comprises the parameters of the encoder and the decoder. Robustness here aims to ameliorate the effect of outliers in the dataset $\{\mathbf{x}_i\}_{i=1}^n$. The constraint set can be \mathbb{R}^d (i.e., the problem is unconstrained) or an ℓ_p -norm ball (i.e., $\{\boldsymbol{\theta} \mid \|\boldsymbol{\theta}\|_p \leq r\}$, for some $r > 0$, $p \geq 1$), when the magnitude of parameters is constrained; this can be used instead of a ℓ_1 or ℓ_2 norm regularizer. In stacked denoising autoencoders [48], the encoder and decoder are shallow and satisfy the additional constraint $\boldsymbol{\theta}_{\text{enc}} = \boldsymbol{\theta}_{\text{dec}}$.

Multi-target Regression [44]. We are given a set of n data points $\mathbf{x}_i \in \mathbb{R}^m$, $i \in [n]$ and the corresponding target labels $\mathbf{y}_i \in \mathbb{R}^{m'}$. The goal is to train a (again possibly non-linear) function $f : \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$, with d parameters, to predict target values for a given vector $\mathbf{x} \in \mathbb{R}^m$. This maps to Prob. (3) via:

$$F(\boldsymbol{\theta}; \mathbf{x}_i, \mathbf{y}_i) = \mathbf{y}_i - f(\boldsymbol{\theta}; \mathbf{x}_i). \quad (5)$$

Robustness in this setting corresponds to ameliorating the effect of outliers in the *label* space, i.e., among labels $\{\mathbf{y}_i\}_{i=1}^n$. The constraint set can again be \mathbb{R}^d or defined through an ℓ_p -norm ball (instead of the corresponding regularizer).

Matrix Factorization [39]. Given a matrix $\mathbf{X} \in \mathbb{R}^{n \times m}$, the goal is to express it as the product of two matrices \mathbf{G}, \mathbf{H} . Cast in our setting, each row $\mathbf{x}_i \in \mathbb{R}^m, i \in [n]$, of \mathbf{X} is mapped to a lower dimensional sub-space as a vector $\mathbf{h}_i \in \mathbb{R}^{m'}$, where the sub-space basis is defined by the rows of the matrix $\mathbf{G} \in \mathbb{R}^{m \times m'}$. Function F is then given by $F(\boldsymbol{\theta}; \mathbf{x}_i) = \mathbf{x}_i - \mathbf{G}\mathbf{h}_i$, where $\boldsymbol{\theta} = (\mathbf{G}, \mathbf{H})$ and the rows of the matrix $\mathbf{H} \in \mathbb{R}^{n \times m'}$ are the low-dimensional embeddings \mathbf{h}_i . Robustness here limits sensitivity to outliers in rows; a similar problem can be defined in terms of robustness to outliers in columns. Beyond usual boundedness constraints, additional constraints are introduced in so-called *non-negative matrix factorization* [14, 39], where matrices \mathbf{G} and \mathbf{H} are constrained to be non-negative.

For all three applications, we assume that F is smooth; this requires, e.g., smooth activation functions in deep models. Moreover, in all three examples, Prob. (3) is non-convex and non-smooth, as $\|\cdot\|_p$ is non-differentiable at $\mathbf{0} \in \mathbb{R}^N$.

4 Robust Regression via MBO

In this section, we outline how non-smooth, non-convex Prob. (3) can be solved via *model-based optimization* (MBO) [36]. MBO relies on the use of a model function, which is a convex approximation of the main objective. In short, the algorithm proceeds iteratively, approximating function $F(\cdot; \cdot)$ by its 1st order Taylor expansion at each iteration. This approximation is affine in $\boldsymbol{\theta}$, and results in a convex optimization problem per iteration.

In more detail, cast into our setting, MBO proceeds as follows. Starting with a feasible solution $\boldsymbol{\theta}^0 \in \mathcal{C}$, it performs the following operations in each step $k \in \mathbb{N}$:

$$\tilde{\boldsymbol{\theta}}^k = \arg \min_{\boldsymbol{\theta}} F_{\boldsymbol{\theta}^k}(\boldsymbol{\theta}) + \frac{h}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}^k\|_2^2, \quad (6a)$$

$$\boldsymbol{\theta}^{k+1} = (1 - \eta^k)\boldsymbol{\theta}^k + \eta^k \tilde{\boldsymbol{\theta}}^k, \quad (6b)$$

where $h > 0$ is a regularization parameter, $\eta^k > 0$ is a step size, and function $F_{\boldsymbol{\theta}^k} : \mathbb{R}^d \rightarrow \mathbb{R}$ is the so-called *model function* at $\boldsymbol{\theta}^k$, defined as:

$$F_{\boldsymbol{\theta}^k}(\boldsymbol{\theta}) \triangleq \frac{1}{n} \sum_{i \in [n]} \|F(\boldsymbol{\theta}^k; \mathbf{x}_i) + \mathbf{D}_{F_i}(\boldsymbol{\theta}^k)(\boldsymbol{\theta} - \boldsymbol{\theta}^k)\|_p + g(\boldsymbol{\theta}) + \chi_{\mathcal{C}}(\boldsymbol{\theta}), \quad (7)$$

where $\mathbf{D}_{F_i}(\boldsymbol{\theta}) \in \mathbb{R}^{N \times d}$ is the Jacobian of $F(\boldsymbol{\theta}; \mathbf{x}_i)$ w.r.t. $\boldsymbol{\theta}$. Thus, in each step, MBO replaces F with its 1st-order Taylor approximation and minimizes the objective plus a proximal penalty; the resulting $\tilde{\boldsymbol{\theta}}^k$ is interpolated with the current solution $\boldsymbol{\theta}^k$.

The above steps are summarized in Alg. 1. The step size η^k is computed via an Armijo-type line search algorithm, which we present in detail in App. A in [31]. Moreover, the inner-step optimization via (6a) can be inexact; the following

Algorithm 1 Model-based Minimization (MBO)

```

1: Input: Initial solution  $\boldsymbol{\theta}^0 \in \text{dom}F$ , iteration number  $K$  set  $\delta, \gamma \in (0, 1)$ ,
   and  $\tilde{\eta} > 0$ 
2: for  $k \in [K]$  do
3:    $\tilde{\boldsymbol{\theta}}^k := \arg \min_{\boldsymbol{\theta}} F_{\boldsymbol{\theta}^k}(\boldsymbol{\theta}) + \frac{h}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}^k\|_2^2$ 
4:   Find  $\gamma^k$  via Armijo search rule
5:    $\boldsymbol{\theta}^{k+1} := (1 - \eta^k)\boldsymbol{\theta}^k + \eta^k \tilde{\boldsymbol{\theta}}^k$ 
6: end for

```

proposition shows asymptotic convergence of MBO to a stationary point using an inexact solver (see also App. A in [31]):

Proposition 1. (Theorem 4.1 of [36]) *Suppose $\boldsymbol{\theta}^*$ is the limit point of the sequence $\boldsymbol{\theta}^k$ generated by Alg. 1. Assume $F_{\boldsymbol{\theta}^k}(\tilde{\boldsymbol{\theta}}^k) + \frac{h}{2} \|\tilde{\boldsymbol{\theta}}^k - \boldsymbol{\theta}^k\|_2^2 - \inf_{\tilde{\boldsymbol{\theta}}} F_{\boldsymbol{\theta}^k}(\tilde{\boldsymbol{\theta}}) + \frac{h}{2} \|\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}^k\|_2^2 \leq \epsilon^k$, for all iterations k , and that $\epsilon^k \rightarrow 0$. Then $\boldsymbol{\theta}^*$ is a stationary point of Prob. (3).*

For completeness, we prove Proposition 1 in Appendix B, by showing that assumptions of Theorem 4.1 of [36] are indeed satisfied. Problem (6a) is convex but still non-smooth; we discuss how it can be solved efficiently via SADM in the next section.

5 Stochastic Alternating Direction Method of Multipliers

After dealing with convexity via MBO, there are still two challenges behind solving the constituent sub-problem (6a). The first is the non-smoothness of $\|\cdot\|_p$; the second is scaling in n , which calls for the use of a stochastic optimization method, akin to SGD (which, however, is not applicable due to the lack of smoothness). We address both through the a novel approach, namely, SADM, which is a stochastic version of the OADM algorithm by Wang and Banerjee [50, 51]. Most importantly, our approach reduces the solution of Prob. (6a) to several gradient-free optimization sub-steps, which can be computed efficiently. In addition, using an SADM/ADMM variant comes with clear stopping criteria, which is challenging for traditional stochastic subgradient methods [27].

5.1 SADM

We first describe how our SADM can be applied to solve Prob. (6a). We introduce the following notation to make our exposition more concise:

$$F^{(k)}(\boldsymbol{\theta}; \mathbf{x}_i) \triangleq \|F(\boldsymbol{\theta}^k; \mathbf{x}_i) + \mathbf{D}_{F_i}(\boldsymbol{\theta}^k)(\boldsymbol{\theta} - \boldsymbol{\theta}^k)\|_p + \frac{h}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}^k\|_2^2, \quad (8a)$$

$$F^{(k)}(\boldsymbol{\theta}) \triangleq \frac{1}{n} \sum_{i \in [n]} F^{(k)}(\boldsymbol{\theta}, \mathbf{x}_i), \quad (8b)$$

$$G(\boldsymbol{\theta}) \triangleq g(\boldsymbol{\theta}) + \chi_C(\boldsymbol{\theta}). \quad (8c)$$

We can then rewrite Prob. (6a) as the following equivalent problem:

$$\text{Minimize } F^{(k)}(\boldsymbol{\theta}_1) + G(\boldsymbol{\theta}_2) \quad (9a)$$

$$\text{subject to: } \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2, \quad (9b)$$

where $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d$ are auxiliary variables.

Note that the objective in (9a) is equivalent to $F^{(k)}(\boldsymbol{\theta}_1) + G(\boldsymbol{\theta}_2)$. SADMM starts with initial solutions, i.e., $\boldsymbol{\theta}_1^0 = \boldsymbol{\theta}_2^0 = \mathbf{u}^0 = 0$. At the t -th iteration, the algorithm performs the following steps:

$$\boldsymbol{\theta}_1^{t+1} := \arg \min_{\boldsymbol{\theta}_1} F^{(k)}(\boldsymbol{\theta}_1; \mathbf{x}_t) + \frac{\rho_t}{2} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2^t + \mathbf{u}^t\|_2^2 + \frac{\gamma_t}{2} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_1^t\|_2^2, \quad (10a)$$

$$\boldsymbol{\theta}_2^{t+1} := \arg \min_{\boldsymbol{\theta}_2} G(\boldsymbol{\theta}_2) + \frac{\rho_t}{2} \|\boldsymbol{\theta}_1^{t+1} - \boldsymbol{\theta}_2 + \mathbf{u}^t\|_2^2, \quad (10b)$$

$$\mathbf{u}^{t+1} := \mathbf{u}^t + \boldsymbol{\theta}_1^{t+1} - \boldsymbol{\theta}_2^{t+1}, \quad (10c)$$

where variables \mathbf{x}_t are sampled uniformly at random from $\{\mathbf{x}_i\}_{i=1}^n$, $\mathbf{u}^t \in \mathbb{R}^d$ is the dual variable, the $\rho_t, \gamma_t > 0$ are scaling coefficients at the t -th iteration. We explain how to set ρ_t, γ_t in Thm. 1.

The solution to Problem (10b) amounts to finding the proximal operator of function G . In general, given that g is smooth and convex, this is a strongly convex optimization problem and can be solved via standard techniques. Nevertheless, for several of the practical cases we described in Sec. 3 this optimization can be done efficiently with gradient-free methods. For example, in the case where the regularizer g is the either a ridge or lasso penalty, and $\mathcal{C} = \mathbb{R}^d$, it is well-known that proximal operators for ℓ_1 and ℓ_2 norms have closed-form solutions [4]. For general ℓ_p norms, an efficient (gradient-free) bi-section method due to Liu and Ye [25] (see App. I in [31]) can be used to compute the proximal operator. Moreover, in the absence of the regularizer, the proximal operator for the indicator function $\chi_{\mathcal{C}}$ is equivalent to projection on the convex set \mathcal{C} . This again has closed-form solution, e.g., when \mathcal{C} is the simplex [29] or an ℓ_p -norm ball [25, 32]. Problem (10a) is harder to solve; we show however that it can also reduced to the (gradient-free) bisection method due to Liu and Ye [25] in the next section.

5.2 Inner ADMM

We solve Problem (10a) using another application of ADMM. In particular, note that (10a) assumes the following general form:

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} + \mathbf{b}\|_p + \lambda \|\mathbf{x} - \mathbf{c}\|_2^2, \quad (11)$$

where $\mathbf{A} = \mathbf{D}_{F_t}(\boldsymbol{\theta}^{(k)})$, the constituent parameter vectors are $\mathbf{c} = \frac{\rho_t}{\rho_t + \gamma_t + h}(\boldsymbol{\theta}_2^t - \mathbf{u}^t) + \frac{\gamma_t}{\rho_t + \gamma_t + h}\boldsymbol{\theta}_1^t + \frac{h}{\rho_t + \gamma_t + h}\boldsymbol{\theta}^{(k)}$, $\mathbf{b} = F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_t) - \mathbf{D}_{F_t}(\boldsymbol{\theta}^{(k)})\boldsymbol{\theta}^{(k)}$, and $\lambda = \frac{\rho_t + \gamma_t + h}{2}$.

We solve (11) via ADMM by reformulating it as the following problem:

$$\min \|\mathbf{y}\|_p + \lambda \|\mathbf{x} - \mathbf{c}\|_2^2 \quad (12a)$$

$$\text{s.t. } \mathbf{A}\mathbf{x} + \mathbf{b} - \mathbf{y} = 0. \quad (12b)$$

The ADMM steps at the k -th iteration for (12) are the following:

$$\mathbf{y}^{k+1} := \arg \min_{\mathbf{y}} \|\mathbf{y}\|_p + \rho'/2 \|\mathbf{y} - \mathbf{A}\mathbf{x}^k - \mathbf{b} + \mathbf{z}^k\|_2^2, \quad (13a)$$

$$\mathbf{x}^{k+1} := \arg \min_{\mathbf{x}} \lambda \|\mathbf{x} - \mathbf{c}\|_2^2 + \rho'/2 \|\mathbf{y}^{k+1} - \mathbf{A}\mathbf{x} - \mathbf{b} + \mathbf{z}^k\|_2^2, \quad (13b)$$

$$\mathbf{z}^{k+1} := \mathbf{z}^k + \mathbf{y}^{k+1} - \mathbf{A}\mathbf{x}^{k+1} - \mathbf{b}, \quad (13c)$$

where $\mathbf{z}^k \in \mathbb{R}^N$ denotes the dual variable at the k -th iteration and $\rho' > 0$ is a hyper-parameter of ADMM.

Problem (13a) is again equivalent to computing the proximal operator of the ℓ_p -norm, which, as mentioned earlier, has closed-form solution for $p = 1, 2$. Moreover, for general ℓ_p -norms the proximal operator can be computed via the bisection algorithm by Liu and Ye [25]. This bisection method yields a solution with an ϵ accuracy in $O(\log_2(1/\epsilon))$ rounds [25, 30] (see App. I in [31]).

5.3 Convergence

To attain the convergence guarantee of MBO given by Proposition 1, we need to solve the inner problem (17) within accuracy ϵ^k at iteration k , where $\epsilon^k \rightarrow 0$. As our major technical contribution, we ensure this by proving the convergence of SADMM when solving Prob. (17).

Consider the sequence $\{\boldsymbol{\theta}_1^t, \boldsymbol{\theta}_2^t, \mathbf{u}^t\}_{t=1}^T$ generated by our SADMM algorithm (10), where $\mathbf{x}_t, t \in [T]$, are sampled u.a.r. from $\{\mathbf{x}_i\}_{i=1}^n$. Let also

$$\bar{\boldsymbol{\theta}}_1^T \triangleq \frac{1}{T} \sum_{t=1}^T \boldsymbol{\theta}_1^t, \quad \bar{\boldsymbol{\theta}}_2^T \triangleq \frac{1}{T} \sum_{t=1}^T \boldsymbol{\theta}_2^{t+1}, \quad (14)$$

denote the time averages of the two solutions. Let also $\boldsymbol{\theta}^* = \boldsymbol{\theta}_1^* = \boldsymbol{\theta}_2^*$ be the optimal solution of Prob. (9). Finally, denote by

$$R^T \triangleq F^{(k)}(\bar{\boldsymbol{\theta}}_1^T) + G(\bar{\boldsymbol{\theta}}_2^T) - F^{(k)}(\boldsymbol{\theta}^*) - G(\boldsymbol{\theta}^*) \quad (15)$$

the residual error of the objective from the optimal. Then, the following holds:

Theorem 1. *Assume that \mathcal{C} is convex, closed, and bounded, while $g(0) = 0$, $g(\boldsymbol{\theta}) \geq 0$, and $g(\cdot)$ is both Lipschitz continuous and β -strongly convex over \mathcal{C} . Moreover, assume that both the function $F(\boldsymbol{\theta}; \mathbf{x}_i)$ and its Jacobian $\mathbf{D}_{F_i}(\boldsymbol{\theta})$ are bounded on the set \mathcal{C} , for all $i \in [n]$. We set $\gamma_t = ht$ and $\rho_t = \beta t$. Then,*

$$\|\bar{\boldsymbol{\theta}}_1^T - \bar{\boldsymbol{\theta}}_2^T\|_2^2 = O\left(\frac{\log T}{T}\right) \quad (16a)$$

$$\mathbb{E}[R^T] = O\left(\frac{\log T}{T}\right) \quad (16b)$$

$$\mathbb{P}\left(R^T \geq k_1 \frac{\log T}{T} + k_2 \frac{M}{\sqrt{T}}\right) \leq e^{-\frac{M^2}{16}} \quad \text{for all } M > 0, T \geq 3, \quad (16c)$$

where $k_1, k_2 > 0$ are constants (see (32) in App. C in [31] for exact definitions).

We prove Theorem 1 in Appendix C. The theorem has the following important consequences. First, (16a) implies that the infeasibility gap between θ_1 and θ_2 decreases as $O(\frac{\log T}{T})$ *deterministically*. Second, by (16b) the residual error R^T decreases as $O(\frac{\log T}{T})$ in expectation. Finally, (16c) shows that the tail of the residual error as iterations increase is exponentially bounded. In particular, given a desirable accuracy ϵ_k , (16c) gives the number of iterations necessary be within ϵ_k of the optimal with any probability $1 - \delta$. Therefore, according to Proposition 1, using SADM will result in convergence of Algorithm 1 with high probability. Finally, we note that, although we write Theorem 1 for updates using only one random sample per iteration, the analysis and guarantees readily extend to the case where a batch selected u.a.r. is used instead. A formal statement and proof can be found in App. E in [31].

6 Experiments

Algorithms. We run two variants of MBO; the first one, which we call MBOSADM, uses SADM (see Sec. 5) for solving the inner problems (6a). The second one, which we call MBOSGD, solves inner problems via a sub-gradient method. We also apply stochastic gradient descent with momentum directly to Prob. (3); we refer to this algorithm as SGD. This corresponds to the algorithm by [27], applied to our setting. We also solve the problem instances with an MSE objective using SGD, as the MSE is smooth and SGD is efficient in this case. Hyperparameters and implementation details are in App. F in [31]. Our code is publicly available.¹

Applications and Datasets. We focus on two applications: training autoencoders and multi-target regression, with a ridge regularizer and $\mathcal{C} = \mathbb{R}^d$. The architectures we use are described in App. F in [31]. For autoencoders, We use MNIST and Fashion-MNIST to train autoencoders and SCM1d [44] for multi target regression. All three datasets, including training and test splits, are described in App. F in [31].

Outliers. We denote the outliers ratio with P_{out} ; each datapoint \mathbf{x}_i , $i \in [n]$, is independently corrupted with outliers with probability P_{out} . The probability P_{out} ranges from 0.0 to 0.3 in our experiments. In particular, we corrupt training samples by replacing them with samples randomly drawn from a Gaussian distribution whose mean is α away from the original data and its standard deviation equals that of the original dataset. For MNIST and FashionMNIST, we set α to 1.5 times the original standard deviation, while for SCM1d, we set α to 2.5 times the standard deviation.

Metrics. We evaluate the solution obtained by different algorithms by using the following three metrics. The first is F_{OBJ} , the regularized objective of Prob. (3) evaluated over the training set. The other two are: $F_{\text{NOU}} \triangleq \frac{\sum_{i \in \mathcal{S}_{\text{OUTL}}} \|F(\theta; \mathbf{x}_i)\|_p}{n - |\mathcal{S}_{\text{OUTL}}|}$, and $F_{\text{TEST}} \triangleq \frac{\sum_{i \in \mathcal{S}_{\text{TEST}}} \|F(\theta; \mathbf{x}_i)\|_p}{|\mathcal{S}_{\text{TEST}}|}$, where $\mathcal{S}_{\text{OUTL}}$, $\mathcal{S}_{\text{TEST}}$ are the outlier and test sets, respectively. Metric F_{NOU} measures the robustness of algorithms w.r.t. outliers;

¹ <https://github.com/neu-spiral/ModelBasedOptimization>

Table 1: Time and Objective Performance. We report objective and time metrics for under different outlier ratios and different p -norms. We observe from the table that MBOSADM significantly outperforms other competitors in terms of objective metrics. In terms of running time, SGD is generally fastest, due to fast gradient updates. However, we see that the time MBO variants take to get to the same or better objective value (i.e., T^*), were comparable to running time of SGD.

P_{out}	p	MBOSADM					MBOSGD					SGD			
		$F_{NOU TL}$	F_{OBJ}	F_{TEST}	$T(h)$	$T^*(h)$	$F_{NOU TL}$	F_{OBJ}	F_{TEST}	$T(h)$	$T^*(h)$	$F_{NOU TL}$	F_{OBJ}	F_{TEST}	$T(h)$
MNIST															
0.0	2.0	2.50	2.51	2.50	5.69	0.14	8.08	8.08	8.12	64.17	6.47	9.21	9.22	9.30	9.83
0.0	1.5	2.63	2.63	2.63	11.67	0.79	20.19	20.20	20.39	65.67	59.98	20.35	20.36	20.57	14.71
0.0	1.0	3.46	3.47	3.44	17.82	3.09	102.79	102.80	104.24	81.53	NA	102.44	102.46	103.89	11.50
0.05	2.0	3.48	5.35	3.46	6.33	0.31	3.89	6.36	3.86	54.92	38.31	8.03	12.52	8.09	13.96
0.05	1.5	4.10	9.74	4.08	45.32	2.08	5.86	11.70	5.82	57.03	25.12	20.34	34.69	20.57	14.60
0.05	1.0	5.23	20.20	5.24	44.03	5.61	27.68	73.53	27.56	32.76	9.67	102.40	236.43	103.90	11.70
0.1	2.0	4.27	7.77	4.23	11.67	1.34	3.56	7.83	3.54	64.20	33.70	7.02	11.64	7.04	13.97
0.1	1.5	4.18	11.84	4.17	68.74	0.29	5.50	13.77	5.45	67.04	9.88	20.34	48.79	20.57	14.04
0.1	1.0	5.90	36.02	5.92	37.73	8.20	30.08	109.79	30.16	39.77	6.72	102.36	368.22	103.90	11.81
0.2	2.0	4.07	8.97	4.04	51.69	4.39	3.54	8.23	3.52	57.08	19.19	7.48	16.44	7.51	14.25
0.2	1.5	3.90	11.58	3.89	195.69	1.56	7.00	20.63	6.95	45.46	6.44	20.36	77.78	20.59	15.06
0.2	1.0	3.85	28.25	3.83	36.98	5.15	40.12	224.47	40.11	19.71	2.13	102.37	639.32	103.90	8.25
0.3	2.0	3.99	14.21	3.98	9.83	4.93	4.02	10.55	3.99	55.60	24.14	7.46	20.63	7.48	13.53
0.3	1.5	20.55	22.56	20.78	159.92	39.24	7.22	24.30	7.16	42.65	15.09	23.90	58.52	23.89	16.25
0.3	1.0	102.70	99.36	104.27	51.48	0.87	56.60	438.89	56.17	20.48	3.32	102.34	910.68	103.90	8.52
Fashion-MNIST															
0.0	2.0	3.51	3.51	3.51	4.33	0.31	5.01	5.01	5.01	42.13	14.80	8.72	8.73	8.70	9.78
0.0	1.5	6.13	6.14	6.14	14.35	2.18	8.87	8.88	8.89	63.39	12.80	22.62	22.63	22.56	14.70
0.0	1.0	10.59	10.61	10.56	29.69	2.63	41.24	41.26	41.35	50.41	3.32	224.26	224.28	224.89	9.72
0.05	2.0	3.80	5.82	3.80	14.70	1.40	4.53	6.75	4.54	67.29	19.15	8.30	11.98	8.27	9.71
0.05	1.5	7.38	14.57	7.40	96.73	2.56	7.91	11.97	7.93	64.25	16.16	20.88	27.22	20.83	10.94
0.05	1.0	16.64	30.51	16.68	43.55	16.04	65.01	109.94	65.31	29.60	9.70	158.65	227.48	158.27	9.97
0.1	2.0	4.05	6.73	4.06	14.66	2.35	4.28	7.90	4.29	65.06	16.46	8.96	14.35	8.94	13.67
0.1	1.5	11.08	32.69	11.10	20.07	NA	8.46	15.23	8.49	65.78	9.92	17.98	31.41	17.95	10.63
0.1	1.0	9.79	27.96	9.81	35.50	2.43	58.70	126.18	58.90	45.06	2.08	235.02	452.45	234.49	13.32
0.2	2.0	6.07	10.19	6.08	14.25	3.67	4.77	9.28	4.77	69.84	30.16	5.71	14.34	5.71	9.31
0.2	1.5	28.51	53.69	28.49	39.42	NA	10.94	27.12	10.97	39.11	16.09	19.36	42.97	19.36	10.87
0.2	1.0	10.50	27.95	10.50	94.72	6.57	140.00	390.02	140.08	17.03	3.33	204.88	644.99	205.13	14.72
0.3	2.0	6.63	23.18	6.63	32.87	NA	5.84	13.04	5.85	50.52	29.95	7.45	20.12	7.46	13.65
0.3	1.5	7.08	22.51	7.10	86.27	30.02	11.09	24.73	11.12	52.41	12.08	19.52	58.26	19.56	11.05
0.3	1.0	14.43	50.91	14.46	95.08	19.48	404.77	893.56	404.52	9.51	NA	410.82	522.50	411.84	10.74
SCMD1d															
0.0	2.0	2.88	2.88	3.02	1.82	0.12	2.85	2.85	2.99	0.36	0.04	3.62	3.63	3.72	1.37
0.0	1.5	4.23	4.24	4.39	7.22	0.43	4.22	4.23	4.44	0.36	0.04	5.47	5.47	5.60	1.58
0.0	1.0	9.78	9.79	10.18	7.13	0.47	9.86	9.86	10.32	0.37	0.04	12.95	12.95	13.25	1.22
0.05	2.0	2.88	3.13	2.99	2.52	0.13	2.86	3.11	3.00	0.54	0.05	3.64	3.89	3.71	1.31
0.05	1.5	4.23	4.61	4.37	10.23	4.61	4.22	4.59	4.46	0.50	0.05	5.50	5.87	5.61	1.23
0.05	1.0	9.69	10.52	10.09	0.59	0.18	9.86	10.66	10.35	0.51	0.05	13.03	13.87	13.29	1.17
0.1	2.0	2.90	3.41	3.01	2.22	0.13	2.84	3.34	3.00	0.46	0.05	3.63	4.12	3.69	1.30
0.1	1.5	4.23	4.99	4.42	9.42	0.68	4.18	4.90	4.40	0.50	0.05	5.52	6.11	5.62	1.15
0.1	1.0	9.56	10.99	10.18	9.92	0.78	9.77	11.11	10.54	0.54	0.07	13.09	13.72	13.32	1.11
0.2	2.0	2.93	3.90	3.03	1.83	0.95	2.86	3.79	3.02	0.5	0.3	3.63	3.97	3.66	1.15
0.2	1.5	4.23	5.60	4.37	8.17	3.60	4.21	5.48	4.47	0.36	0.2	5.50	5.83	5.56	1.17
0.2	1.0	9.46	11.00	10.04	6.55	1.50	9.82	11.30	10.60	0.45	0.20	13.09	13.38	13.28	1.11
0.3	2.0	2.93	4.32	3.03	1.80	NA	2.85	4.10	3.03	0.46	NA	3.61	3.90	3.64	1.18
0.3	1.5	4.25	6.05	4.44	8.19	NA	4.21	5.73	4.49	0.50	NA	5.43	5.69	5.43	1.18
0.3	1.0	9.53	11.07	10.00	6.44	2.72	9.68	11.05	10.32	0.51	0.28	12.95	13.13	12.96	1.11

ideally, $F_{NOU TL}$ should remain unchanged as the fraction of outliers increases. Metric F_{TEST} evaluates the generalization ability of algorithms on unseen (test) data, which also does not contain outliers; ideally, F_{TEST} be similar $F_{NOU TL}$. Moreover, we report total running time (T) of all algorithms. For the two variants of MBO, we additionally report the time (T^*) until they reach the optimal value attained by SGD (N/A if never reached). Finally, for autoencoders, we also use dataset labels to train a logistic regression classifier over latent embeddings, and also report the prediction accuracy on the test set. Classifier hyperparameters are described in App. G in [31].

6.1 Time and Objective Performance Comparison

We evaluate our algorithms w.r.t. both objective and time metrics, which we report for different outlier ratios P_{out} and p -norms in Table 1. By comparing

objective metrics, we see that MBOSADM and MBOSGD significantly outperform SGD. SGD achieves a better F_{OBJ} in only 2 out of 45 cases, i.e., SCM1d dataset for $p = 1.5, 2$ and $P_{\text{out}} = 0.3$; however, even for these two cases, MBOSADM and MBOSGD obtain better F_{NOUTL} and F_{TEST} values. In terms of overall running time T , SGD is generally faster than MBOSADM and MBOSGD; this is expected, as each iteration of SGD only computes the gradient of a mini-batch of terms in the objective, while the other methods need to solve an inner-problem. Nonetheless, by comparing T^* , we see that the MBO variants obtain the same or better objective as SGD in a comparable time. In particular, T^* is less than T for SGD in 33 and 15 cases (out of 45) for MBOSADM and MBOSGD, respectively.

Comparing the performance between MBOSADM and MBOSGD, we first note that MBOSADM has a superior performance w.r.t. all three objective metrics for 25 out of 45 cases. In some cases, MBOSADM obtains considerably smaller objective values; for example, for MNIST and $P_{\text{out}} = 0.0, p = 1$, F_{NOUT} is 0.03 of the value obtained by MBOSGD (also see Figures 2c and 2f). However, it seems that in the high-outlier setting $P_{\text{out}} = 0.3$ the performance of MBOSADM deteriorates; this is mostly due to the fact that the high number of outliers adversely affects the convergence of SADM and it takes more iterations to satisfy the desired accuracy.

6.2 Robustness Analysis

We further study the robustness of different p -norms and MSE to the presence of outliers. For brevity, we only report results for MNIST and for $p = 1, 2$, and MSE. For more results refer to Fig. 4 in App. H in [31]. We show the scaling of F_{NOUT} and F_{TEST} w.r.t. the fraction P_{out} in Fig. 2, for different norms. To make comparisons between different objectives interpretable, we normalize all values in each figure by the largest value in that figure.

By comparing Figures 2a and 2d, corresponding to MSE, with other plots in Fig. 2, we see that the loss values considerably increase by adding outliers. For other p -norms, we see that SGD generally stays unchanged, w.r.t. outliers. However, the loss for SGD is higher than MBO variants. Loss values for MBOSGD also do not increase significantly by adding outliers. Moreover, we see that, when no outliers are present $P_{\text{out}} = 0.0$, MBOSGD obtains higher loss values. MBOSADM generally achieves the lowest loss values and these values again do not increase with increasing P_{out} ; however, for the highest outliers ($P_{\text{out}} = 0.3$), the performance of MBOSADM is considerably worse for $p = 1$. As we emphasize in Sec. 6.1, high number of outliers adversely affects the convergence of SADM, and hence the poor performance of MBOSADM for $P_{\text{out}} = 0.3$.

6.3 Classification Performance

Fig. 3 shows the quality of the latent embeddings obtained by different trained autoencoders on the downstream classification over MNIST and FashionMNIST. Additional results are shown in Tables 2 and 3 in App. G. We see that MBO variants again outperform SGD. For MNIST, (reported in Figures 3a to 3c), we see that MBOSADM for $p = 1$ obtains the highest accuracy. Moreover, for Fashion-MNIST

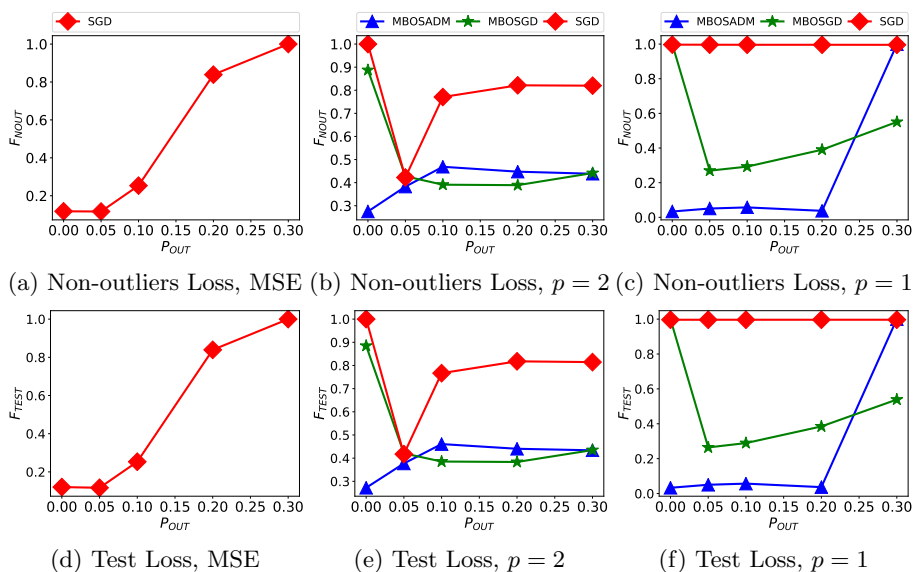


Fig. 2: A comparison of scalability of the non-outliers loss F_{NOUT} and the test loss F_{TEST} for different p -norms, w.r.t., outliers fraction P_{out} . We normalize values in each figure by the largest observed value, to make comparisons between different objectives possible. We see that MSE in Figures 2a and 2 are drastically affected by outliers and scale with outliers fraction P_{out} . Other ℓ_p norms for different methods in Figures 2b, 2c, 2e, and 2f generally stay unchanged w.r.t. P_{out} . However, MBOSADM in the high outlier regime and $p = 1$ performs poorly.

(reported in Fig. 3d to 3f), we observe that again MBOSADM for $p = 1$ outperforms other methods. We also observe that MSE (reported in Figures 3a and 3d) is sensitive to outliers; the corresponding accuracy drastically drops for $P_{\text{out}} \geq 0.1$. An interesting observation is that adding outliers improves the performance of SGD; however, we see that SGD always results in lower accuracy, except in two cases ($P_{\text{out}} = 0.2$ in Fig. 3b and $P_{\text{out}} = 0.3$ in Fig. 3e).

7 Conclusion

We present a generic class of robust formulations that includes many applications, i.e., auto-encoders, multi-target regression, and matrix factorization. We show that SADM, in combination with MBO, provides efficient solutions for our class of robust problems. Studying other proximal measures described by Ochs et al. [36] is an open area. Moreover, characterizing the sample complexity of our proposed method for obtaining a stationary point, as in MBO variants that use gradient methods [6, 10], is an interesting future direction.

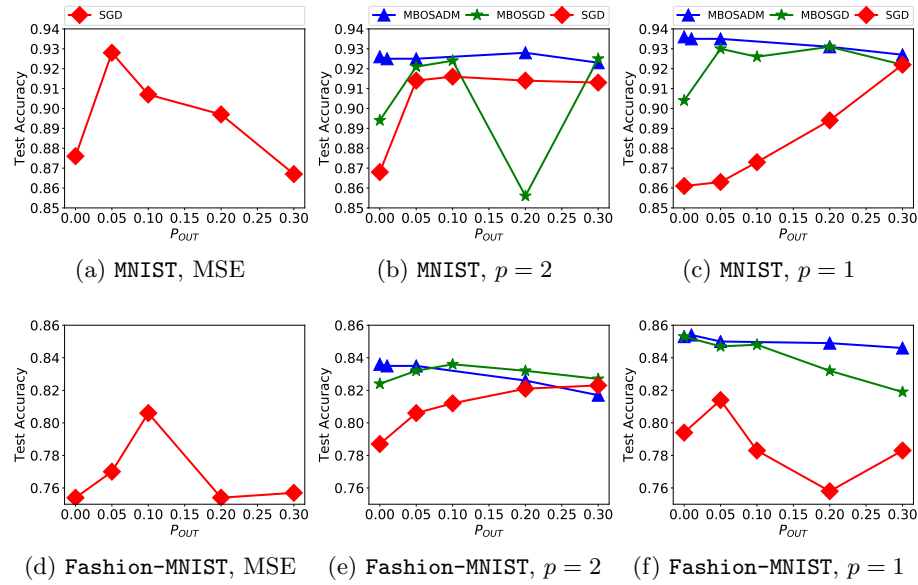


Fig. 3: Classification performance for different methods and datasets. We use the embeddings obtained by auto-encoders trained via different algorithms to train a logistic regression model for classification. We generally observe that MBOSADM results in higher accuracy on the test sets. Moreover, we see that MSE is evidently sensitive to outliers, see Figures 3a and 3d for $P_{out} \geq 0.2$.

References

1. Attouch, H., Bolte, J., Redont, P., Soubeyran, A.: Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Lojasiewicz inequality. *Mathematics of operations research* **35**(2), 438–457 (2010)
2. Baccini, A., Besse, P., Falguerolles, A.: A l_1 -norm PCA and a heuristic approach. *Ordinal and symbolic data analysis* **1**(1), 359–368 (1996)
3. Blumensath, T., Davies, M.E.: Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis* **27**(3), 265–274 (2009)
4. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning* **3**(1), 1–122 (2011)
5. Croux, C., Filzmoser, P.: Robust factorization of a data matrix. In: *COMPSTAT*. pp. 245–250. Springer (1998)
6. Davis, D., Drusvyatskiy, D.: Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization* **29**(1), 207–239 (2019)
7. Davis, D., Grimmer, B.: Proximally guided stochastic subgradient method for nonsmooth, nonconvex problems. *SIAM Journal on Optimization* **29**(3), 1908–1930 (2019)
8. Ding, C., Zhou, D., He, X., Zha, H.: R1-PCA: rotational invariant l_1 -norm principal component analysis for robust subspace factorization. In: *ICML* (2006)

9. Drusvyatskiy, D., Lewis, A.S.: Error bounds, quadratic growth, and linear convergence of proximal methods. *Mathematics of Operations Research* **43**(3), 919–948 (2018)
10. Drusvyatskiy, D., Paquette, C.: Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming* **178**(1), 503–558 (2019)
11. Du, L., Zhou, P., Shi, L., Wang, H., Fan, M., Wang, W., Shen, Y.D.: Robust multiple kernel k-means using l_{21} -norm. In: *IJCAI* (2015)
12. Duchi, J.C., Ruan, F.: Stochastic methods for composite and weakly convex optimization problems. *SIAM Journal on Optimization* **28**(4), 3229–3259 (2018)
13. Eriksson, A., Van Den Hengel, A.: Efficient computation of robust low-rank matrix approximations in the presence of missing data using the l_1 norm. In: *CVPR* (2010)
14. Févotte, C., Idier, J.: Algorithms for nonnegative matrix factorization with the β -divergence. *Neural computation* **23**(9), 2421–2456 (2011)
15. Friedman, J., Hastie, T., Tibshirani, R.: *The Elements of Statistical Learning*, vol. 1. Springer Series in Statistics New York (2001)
16. Gillis, N.: *Nonnegative Matrix Factorization*. SIAM - Society for Industrial and Applied Mathematics, Philadelphia, PA (2020)
17. Hosseini, S., Chapman, A., Mesbahi, M.: Online distributed ADMM via dual averaging. In: *CDC* (2014)
18. Jiang, W., Gao, H., Chung, F.I., Huang, H.: The $l_{2,1}$ -norm stacked robust autoencoders for domain adaptation. In: *AAAI* (2016)
19. Ke, Q., Kanade, T.: Robust l_1 factorization in the presence of outliers and missing data by alternative convex programming. In: *CVPR* (2005)
20. Kong, D., Ding, C., Huang, H.: Robust nonnegative matrix factorization using l_{21} -norm. In: *CIKM* (2011)
21. Kwak, N.: Principal component analysis based on l_1 -norm maximization. *IEEE transactions on pattern analysis and machine intelligence* **30**(9), 1672–1680 (2008)
22. Le, H., Gillis, N., Patrinos, P.: Inertial block proximal methods for non-convex non-smooth optimization. In: *ICML* (2020)
23. Lewis, A.S., Wright, S.J.: A proximal method for composite minimization. *Mathematical Programming* **158**(1), 501–546 (2016)
24. Li, X., Pang, Y., Yuan, Y.: l_1 -norm-based 2DPCA. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **40**(4), 1170–1175 (2010)
25. Liu, J., Ye, J.: Efficient l_1/l_q NormRregularization. arXiv preprint arXiv:1009.4766 (2010)
26. Liu, Y., Shang, F., Cheng, J.: Accelerated variance reduced stochastic admm. In: *AAAI* (2017)
27. Mai, V., Johansson, M.: Convergence of a stochastic gradient method with momentum for non-smooth non-convex optimization. In: *ICML* (2020)
28. Mehta, J., Gupta, K., Gogna, A., Majumdar, A., Anand, S.: Stacked robust autoencoder for classification. In: *NeurIPS* (2016)
29. Michelot, C.: A finite algorithm for finding the projection of a point onto the canonical simplex of n . *Journal of Optimization Theory and Applications* (1986)
30. Moharrer, A., Gao, J., Wang, S., Bento, J., Ioannidis, S.: Massively distributed graph distances. *IEEE Transactions on Signal and Information Processing over Networks* **6**, 667–683 (2020)
31. Moharrer, A., Kamran, K., Yeh, E., Ioannidis, S.: Robust regression via model based methods (2021)
32. Moreau, J.J.: Décomposition orthogonale d’un espace hilbertien selon deux cônes mutuellement polaires. *Comptes rendus hebdomadaires des séances de l’Académie des sciences* **255**, 238–240 (1962)

33. Natarajan, B.K.: Sparse approximate solutions to linear systems. *SIAM journal on computing* **24**(2), 227–234 (1995)
34. Nemirovski, A., Juditsky, A., Lan, G., Shapiro, A.: Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization* **19**(4), 1574–1609 (2009)
35. Nie, F., Huang, H., Cai, X., Ding, C.H.: Efficient and robust feature selection via joint ℓ_2, ℓ_1 -norms minimization. In: *NIPS* (2010)
36. Ochs, P., Fadili, J., Brox, T.: Non-smooth non-convex bregman minimization: Unification and new algorithms. *Journal of Optimization Theory and Applications* **181**(1), 244–278 (2019)
37. Ochs, P., Malitsky, Y.: Model function based conditional gradient method with Armijo-like line search. In: *Proceedings of the 36th International Conference on Machine Learning* (2019)
38. Ouyang, H., He, N., Tran, L., Gray, A.: Stochastic alternating direction method of multipliers. In: *International Conference on Machine Learning*. pp. 80–88. PMLR (2013)
39. Paatero, P., Tapper, U.: Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* **5**(2), 111–126 (1994)
40. Peng, Y., Ganesh, A., Wright, J., Xu, W., Ma, Y.: Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images. *IEEE transactions on pattern analysis and machine intelligence* **34**(11), 2233–2246 (2012)
41. Pesme, S., Flammarion, N.: Online robust regression via sgd on the ℓ_1 loss. In: *NeurIPS* (2020)
42. Qian, M., Zhai, C.: Robust unsupervised feature selection. In: *IJCAI* (2013)
43. Shalev-Shwartz, S., et al.: Online learning and online convex optimization. *Foundations and trends in Machine Learning* **4**(2), 107–194 (2011)
44. Spyromitros-Xioufis, E., Tsoumakas, G., Groves, W., Vlahavas, I.: Multi-target regression via input space expansion: treating targets as inputs. *Machine Learning* **104**(1), 55–98 (2016)
45. Suzuki, T.: Dual averaging and proximal gradient descent for online alternating direction multiplier method. In: *ICML* (2013)
46. Tao, M., Yuan, X.: Recovering low-rank and sparse components of matrices from incomplete and noisy observations. *SIAM Journal on Optimization* **21**(1), 57–81 (2011)
47. Vial, J.P.: Strong and weak convexity of sets and functions. *Mathematics of Operations Research* **8**(2), 231–259 (1983)
48. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.A.: Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* **11**(Dec), 3371–3408 (2010)
49. Waegeman, W., Dembczyński, K., Hüllermeier, E.: Multi-target prediction: a unifying view on problems and methods. *Data Mining and Knowledge Discovery* **33**(2), 293–324 (2019)
50. Wang, H., Banerjee, A.: Online alternating direction method. In: *ICML* (2012)
51. Wang, H., Banerjee, A.: Online alternating direction method (longer version). arXiv preprint arXiv:1306.3721 (2013)
52. Zheng, S., Kwok, J.T.: Fast-and-light stochastic ADMM. In: *IJCAI*. pp. 2407–2613 (2016)

Algorithm 2 Line Search Algorithm (LSA)

-
- 1: **Input:** Solutions $\boldsymbol{\theta}^k$, $\tilde{\boldsymbol{\theta}}^k$, and parameters $\delta, \gamma \in (0, 1), \tilde{\eta} > 0$
 - 2: Initialize $\tilde{\boldsymbol{\theta}} := (1 - \tilde{\eta})\boldsymbol{\theta}^k + \tilde{\eta}\tilde{\boldsymbol{\theta}}^k$
 - 3: **while** $\frac{1}{n} \sum_{i \in [n]} \|F(\tilde{\boldsymbol{\theta}}; \mathbf{x}_i)\|_p + g(\tilde{\boldsymbol{\theta}}) > \frac{1}{n} \sum_{i \in [n]} \|F(\boldsymbol{\theta}^k; \mathbf{x}_i)\|_p + g(\boldsymbol{\theta}^k) + \gamma\tilde{\eta}[F_{\boldsymbol{\theta}^k}(\tilde{\boldsymbol{\theta}}^k) - F_{\boldsymbol{\theta}^k}(\boldsymbol{\theta}^k) + \frac{h}{2}\|\tilde{\boldsymbol{\theta}}^k - \boldsymbol{\theta}^k\|_2^2]$ **do**
 - 4: Set $\tilde{\eta} := \delta\tilde{\eta}$
 - 5: Set $\tilde{\boldsymbol{\theta}} := (1 - \tilde{\eta})\boldsymbol{\theta}^k + \tilde{\eta}\tilde{\boldsymbol{\theta}}^k$
 - 6: **end while**
 - 7: **Return:** $\tilde{\eta}$
-

A Model Based Optimization

Ochs et al. [36] allows to use more general Bregman divergences for the second term (c.f. Sec. 5.2 of [36]). The specific model functions we study fall under Example 5.3 in [36] (see also [6, 23].) Moreover, Ochs et al. allow (6a) to be solved inexactly; at each iteration, the solution $\tilde{\boldsymbol{\theta}}^k$ only needs to improve the model function value by, i.e.,

$$\Delta_k \triangleq F_{\boldsymbol{\theta}^k}(\tilde{\boldsymbol{\theta}}^k) + \frac{h}{2}\|\tilde{\boldsymbol{\theta}}^k - \boldsymbol{\theta}^k\|_2^2 - F_{\boldsymbol{\theta}^k}(\boldsymbol{\theta}^k) < 0. \quad (17)$$

The step-size η^k is found via an Armijo line search algorithm. In particular, the line search algorithm finds a step-size, s.t., $\boldsymbol{\theta}^{k+1}$ improves the current objective comparable with the model improvement Δ^k , i.e.,

$$\frac{1}{n} \sum_{i \in [n]} \|F(\boldsymbol{\theta}^{k+1}; \mathbf{x}_i)\|_p + g(\boldsymbol{\theta}^{k+1}) - \left(\frac{1}{n} \sum_{i \in [n]} \|F(\boldsymbol{\theta}^k; \mathbf{x}_i)\|_p + g(\boldsymbol{\theta}^k) \right) \leq \gamma\eta^k \Delta^k,$$

where $\gamma \in (0, 1)$ is a hyper-parameter of the linear search algorithm.

The exact Line Search Algorithm from Ochs et al. [36] is summarized in Alg. 2. Note that Ochs et al. prove that LSA is guaranteed to finish within finite number of iterations.

B Proof of Proposition 1

We show that all assumptions for Theorem 4.1 of Ochs et al. [36] are satisfied, therefore the result holds. We solve convex problems (6a) via OADM iterations. Since we established the $O(\frac{\log T}{T})$ convergence rate of OADM in Theorem 1, we can solve (6a) with an arbitrary accuracy ϵ , which goes to zero for $T \rightarrow \infty$; therefore, Assumption 4.1 of [36] is satisfied. Moreover, due to the choice of the Euclidean norm as our Bergman distance function, Assumption 4.2 of [36] is satisfied (see Section 5 of [36]). Finally, our model function (7) is precisely Example 5.3 in [36]; it is written as the form $f_0 + h \circ F$, where $f_0(\boldsymbol{\theta}) = \chi c(\boldsymbol{\theta}) + g(\boldsymbol{\theta})$,

$h(\mathbf{F}) = \|\mathbf{F}\|_{p,1}$, and $F : \mathbb{R}^d \rightarrow \mathbb{R}^{n \times N}$. Therefore, Assumption 4.3 is also satisfied. In addition, the domain of Euclidean distance is \mathbb{R}^n , which implies that Condition (ii) in Theorem 4.1 of [36] is satisfied for any limit point. Thus, all the conditions in Theorem 4.1 of Ochs et al. are satisfied, and every limit point of Alg. 1 is a stationary point.

C Proof of Theorem 1

Proof. Since we assume that the constraint set \mathcal{C} is convex, closed, and bounded, we take the diameter of the set \mathcal{C} to be

$$D_{\mathcal{C}} = \max_{\theta_1, \theta_2 \in \mathcal{C}} \|\theta_1 - \theta_2\|_2. \quad (18)$$

In addition, since function F and its Jacobian $\mathbf{D}_{F_i}(\boldsymbol{\theta})$ are bounded on the set \mathcal{X} , for all $i \in [n]$, there exist $M_F, M_D < \infty$, s.t.,

$$\|F(\boldsymbol{\theta}; \mathbf{x})\|_{\infty} \leq M_F \quad \forall \boldsymbol{\theta} \in \mathcal{C}, \mathbf{x} \in \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \quad (19a)$$

$$\|\mathbf{D}_{F_i}(\boldsymbol{\theta})\|_{\infty} \leq M_D \quad \forall \boldsymbol{\theta} \in \mathcal{C}, i \in [n]. \quad (19b)$$

We now paraphrase Theorem 6 of [51] as the following lemma, which we prove in Appendix D.

Lemma 1 (Theorem 6 in [51]). *If the assumptions of Theorem 1 hold, for the sequence $\{\boldsymbol{\theta}_1^t, \boldsymbol{\theta}_2^t, \mathbf{u}^t\}, t \in [T]$ generated by OADM algorithm for any sequence of the variables $\mathbf{x}_t \in \{\mathbf{x}_1, \dots, \mathbf{x}_n\}, t \in [T]$ the following holds*

$$\begin{aligned} & \sum_{t=1}^T \left(F^{(k)}(\boldsymbol{\theta}_1^t; \mathbf{x}_t) + G(\boldsymbol{\theta}_2^{t+1}) \right) - \sum_{t=1}^T \left(F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t) + G(\boldsymbol{\theta}^*) \right) \\ & \leq \frac{(N^{1/p}dM_D + hD_{\mathcal{C}})^2}{2h} \log(T+1) + \frac{\beta + h}{2} D_{\mathcal{C}}^2 \end{aligned} \quad (20a)$$

$$\begin{aligned} & \sum_{t=1}^T \|\boldsymbol{\theta}_1^{t+1} - \boldsymbol{\theta}_2^{t+1}\|_2^2 + \|\boldsymbol{\theta}_2^{t+1} - \boldsymbol{\theta}_2^t\|_2^2 \\ & \leq \frac{2}{\beta} (\sqrt{d}M_G + hD_{\mathcal{C}} + L_G) \log(T+1) + \left(1 + \frac{h}{\beta}\right) D_{\mathcal{C}}^2, \end{aligned} \quad (20b)$$

where $\boldsymbol{\theta}^* = \boldsymbol{\theta}_1^* = \boldsymbol{\theta}_2^*$ is the optimal solution for (9), and L_G is the Lipschitz coefficient for $g(\cdot)$, and $M_G \triangleq N^{1/p} \sqrt{d}M_D$.

From (20b) we obtain the following

$$\sum_{t=1}^T \|\boldsymbol{\theta}_1^t - \boldsymbol{\theta}_2^t\|_2^2 = O(\log T) \quad (21a)$$

$$\sum_{t=1}^T \|\boldsymbol{\theta}_2^{t+1} - \boldsymbol{\theta}_2^t\|_2^2 = O(\log T). \quad (21b)$$

Now we derive the result

$$\begin{aligned}
\|\bar{\boldsymbol{\theta}}_1^T - \bar{\boldsymbol{\theta}}_2^T\|_2^2 &= \left\| \frac{1}{T} \sum_{t=1}^T (\boldsymbol{\theta}_1^t - \boldsymbol{\theta}_2^{t+1}) \right\|_2^2 \\
&= \left\| \frac{1}{T} \sum_{t=1}^T (\boldsymbol{\theta}_1^t - \boldsymbol{\theta}_2^t + \boldsymbol{\theta}_2^t - \boldsymbol{\theta}_2^{t+1}) \right\|_2^2 \\
&\leq \left\| \frac{2}{T} \sum_{t=1}^T (\boldsymbol{\theta}_1^t - \boldsymbol{\theta}_2^t) \right\|_2^2 + \left\| \frac{2}{T} \sum_{t=1}^T (\boldsymbol{\theta}_2^t - \boldsymbol{\theta}_2^{t+1}) \right\|_2^2 \\
&\leq \frac{2}{T} \sum_{t=1}^T \|\boldsymbol{\theta}_1^t - \boldsymbol{\theta}_2^t\|_2^2 + \frac{2}{T} \sum_{t=1}^T \|\boldsymbol{\theta}_2^t - \boldsymbol{\theta}_2^{t+1}\|_2^2 \\
&\stackrel{(21)}{=} O\left(\frac{\log T}{T}\right),
\end{aligned}$$

where in deriving the first inequality we have used the fact that

$$\|\mathbf{x} + \mathbf{y}\|_2^2 \leq 2\|\mathbf{x}\|_2^2 + 2\|\mathbf{y}\|_2^2 \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d.$$

Now we prove the second part of theorem about the optimality of solutions. Using convexity of $F^{(k)}$ and G we have that

$$\begin{aligned}
&F^{(k)}(\bar{\boldsymbol{\theta}}_1^T) + G(\bar{\boldsymbol{\theta}}_2^T) - F^{(k)}(\boldsymbol{\theta}^*) - G(\boldsymbol{\theta}^*) \\
&\leq \frac{1}{T} \sum_{t=1}^T \left(F^{(k)}(\boldsymbol{\theta}_1^t) - F^{(k)}(\boldsymbol{\theta}^*) + G(\boldsymbol{\theta}_2^{t+1}) - G(\boldsymbol{\theta}^*) \right) \\
&= \frac{1}{T} \sum_{t=1}^T \left(F^{(k)}(\boldsymbol{\theta}_1^t; \mathbf{x}_t) - F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t) + G(\boldsymbol{\theta}_2^{t+1}) - G(\boldsymbol{\theta}^*) \right) \\
&+ \frac{1}{T} \sum_{t=1}^T \left(F^{(k)}(\boldsymbol{\theta}_1^t) - F^{(k)}(\boldsymbol{\theta}_1^t; \mathbf{x}_t) - F^{(k)}(\boldsymbol{\theta}^*) + F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t) \right) \\
&\stackrel{(20a)}{\leq} \frac{(N^{1/p}dM_D + hD_C)^2}{2hT} \log(T+1) + \frac{\beta + h}{2T} D_C^2 + \frac{1}{T} \sum_{t=1}^T \delta_t, \tag{22}
\end{aligned}$$

where the first inequality is due to the Jensen's inequality and

$$\delta_t \triangleq F^{(k)}(\boldsymbol{\theta}_1^t) - F^{(k)}(\boldsymbol{\theta}_1^t; \mathbf{x}_t) - F^{(k)}(\boldsymbol{\theta}^*) + F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t).$$

As the variables \mathbf{x}_t and $\boldsymbol{\theta}_1^t$ are independent, we have

$$\begin{aligned}
&\mathbb{E}[\delta_t | \mathbf{x}_1, \dots, \mathbf{x}_{t-1}] \\
&= \mathbb{E} \left[F^{(k)}(\boldsymbol{\theta}_1^t) - F^{(k)}(\boldsymbol{\theta}_1^t; \mathbf{x}_t) - F^{(k)}(\boldsymbol{\theta}^*) + F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t) | \mathbf{x}_1, \dots, \mathbf{x}_{t-1} \right] = 0.
\end{aligned}$$

Therefore, we obtain

$$\mathbb{E}_{\mathbf{x}_t, t \in [T]} [\delta_t] = 0. \tag{23}$$

Now taking expectations of both sides of (22) w.r.t. the sequence $\zeta_t, t \in [T]$ and noting (23) we have that:

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_t, t \in [T]} \left[F^{(k)}(\bar{\boldsymbol{\theta}}_1^T) + G(\bar{\boldsymbol{\theta}}_2^T) - F^{(k)}(\boldsymbol{\theta}^*) - G(\boldsymbol{\theta}^*) \right] \\ & \leq \frac{(N^{1/p} dM_D + hD_C)^2}{2hT} \log(T+1) + \frac{\beta+h}{2T} D_C^2 \\ & = O\left(\frac{\log T}{T}\right) \end{aligned}$$

We first derive the following for all $\boldsymbol{\theta} \in \mathcal{C}, \mathbf{x}_i \in \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$:

$$\begin{aligned} & |F^{(k)}(\boldsymbol{\theta}) - F^{(k)}(\boldsymbol{\theta}; \mathbf{x}_i)| \\ & = \left| F_{\boldsymbol{\theta}^{(k)}}(\boldsymbol{\theta}; \mathbf{x}_i) - \frac{1}{n} \sum_{j=1}^n F_{\boldsymbol{\theta}^{(k)}}(\boldsymbol{\theta}; \mathbf{x}_j) \right| \\ & \leq \frac{1}{n} \sum_{j \neq i} |F_{\boldsymbol{\theta}^{(k)}}(\boldsymbol{\theta}; \mathbf{x}_i) - F_{\boldsymbol{\theta}^{(k)}}(\boldsymbol{\theta}; \mathbf{x}_j)| \\ & = \frac{1}{n} \sum_{j \neq i} \left| \|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_i) + \mathbf{D}_{F_i}(\boldsymbol{\theta}^{(k)})(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)})\|_p \right. \\ & \quad \left. - \|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_j) + \mathbf{D}_{F_j}(\boldsymbol{\theta}^{(k)})(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)})\|_p \right| \\ & \leq \frac{1}{n} \sum_{j \neq i} \|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_i) + \mathbf{D}_{F_i}(\boldsymbol{\theta}^{(k)})(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}) - F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_j) + \mathbf{D}_{F_j}(\boldsymbol{\theta}^{(k)})(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)})\|_p \\ & \leq \frac{1}{n} \sum_{j \neq i} \|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_i) - F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_j)\|_p + \frac{1}{n} \sum_{j \neq i} \left\| \left(\mathbf{D}_{F_i}(\boldsymbol{\theta}^{(k)}) - \mathbf{D}_{F_j}(\boldsymbol{\theta}^{(k)}) \right) (\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}) \right\|_p \\ & \stackrel{(19a)}{\leq} \vartheta \sqrt{N} M_F + \frac{1}{n} \sum_{j \neq i} \left(\sum_{i'=1}^N \left| \left(\mathbf{D}_{F_i}(\boldsymbol{\theta}^{(k)}) - \mathbf{D}_{F_j}(\boldsymbol{\theta}^{(k)}) \right)_{i'}^\top (\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}) \right|^p \right)^{1/p} \\ & \leq \vartheta \sqrt{N} M_F + \frac{1}{n} \sum_{j \neq i} \left(\sum_{i'=1}^N \left(\left\| \left(\mathbf{D}_{F_i}(\boldsymbol{\theta}^{(k)}) - \mathbf{D}_{F_j}(\boldsymbol{\theta}^{(k)}) \right)_{i'} \right\|_2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}\|_2 \right)^p \right)^{1/p} \\ & \stackrel{(19b), (18)}{\leq} \vartheta \sqrt{N} M_F + \vartheta \sqrt{N} \sqrt{d} M_D D_C. \end{aligned} \tag{24}$$

In order to show the rest of the results we first show that the variance of δ_t are bounded for all $t \in [T]$

$$\begin{aligned} \delta_t^2 & = \left(F^{(k)}(\boldsymbol{\theta}_1^t) - F^{(k)}(\boldsymbol{\theta}_1^t, \zeta_t) - F^{(k)}(\boldsymbol{\theta}^*) + F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t) \right)^2 \\ & \leq 2 \left(F^{(k)}(\boldsymbol{\theta}_1^t) - F^{(k)}(\boldsymbol{\theta}_1^t; \mathbf{x}_t) \right)^2 + 2 \left(-F^{(k)}(\boldsymbol{\theta}^*) + F^{(k)}(\boldsymbol{\theta}^*; \mathbf{x}_t) \right)^2 \\ & \stackrel{(24)}{\leq} 4 \left(\vartheta \sqrt{N} M_F + \vartheta \sqrt{N} \sqrt{d} M_D D_C \right)^2 \\ & \equiv \sigma^2. \end{aligned} \tag{25}$$

From (25) it is obvious that

$$\left(\frac{\delta_t^2}{\sigma^2}\right) \leq \exp(1) \quad \forall t \in [T]. \quad (26)$$

Now we show that the following holds

$$\mathbb{E}[\exp(\alpha\delta_t)|\mathbf{x}_1, \dots, \mathbf{x}_{t-1}] \leq \exp(\alpha^2\sigma^2) \quad \forall \alpha > 0. \quad (27)$$

To show (27), we follow results from [34]; similar to them, we use the fact that $\exp(x) \leq x + \exp(x^2)$. Then we have that

$$\begin{aligned} \mathbb{E}[\exp(\alpha\delta_t)|\mathbf{x}_1, \dots, \mathbf{x}_{t-1}] &\leq \mathbb{E}[\exp(\alpha^2\delta_t^2)|\mathbf{x}_1, \dots, \mathbf{x}_{t-1}] \\ &= \mathbb{E}\left[\left(\exp\left(\frac{\delta_t^2}{\sigma^2}\right)\right)^{\alpha^2\sigma^2} \mid \mathbf{x}_1, \dots, \mathbf{x}_{t-1}\right] \\ &\stackrel{(26)}{\leq} \exp(\alpha^2\sigma^2) \end{aligned}$$

Now for the sum $\sum_{t=1}^T \delta_t$ we have that

$$\begin{aligned} \mathbb{E}\left[\exp\left(\alpha \sum_{t=1}^T \delta_t\right)\right] &= \mathbb{E}\left[\exp\left(\alpha \sum_{t=1}^{T-1} \delta_t\right) \exp(\alpha\delta_T)\right] \\ &= \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_{T-1}} \left[\exp\left(\alpha \sum_{t=1}^{T-1} \delta_t\right) \mathbb{E}_{\mathbf{x}_T} [\exp(\alpha\delta_T) | \mathbf{x}_1, \dots, \mathbf{x}_{T-1}] \right] \\ &\stackrel{(27)}{\leq} \exp(\alpha^2\sigma^2) \mathbb{E}\left[\exp\left(\alpha \sum_{t=1}^{T-1} \delta_t\right)\right]. \end{aligned} \quad (28)$$

Having (28) for all T and $\mathbb{E}[\exp(\alpha\delta_1)] \leq \exp(\alpha^2\sigma^2)$ by induction we obtain that:

$$\mathbb{E}\left[\exp\left(\alpha \sum_{t=1}^T \delta_t\right)\right] \leq \exp(T\alpha^2\sigma^2). \quad (29)$$

Now applying the Markov's inequality we have that for all $\alpha > 0, M_\delta$:

$$p\left(\sum_{t=1}^T \delta_t \geq M_\delta\right) \leq \frac{\mathbb{E}\left[\exp\left(\alpha \sum_{t=1}^T \delta_t\right)\right]}{\exp(\alpha M_\delta)} \stackrel{(29)}{\leq} \frac{\exp(T\alpha^2\sigma^2)}{\exp(\alpha M_\delta)}. \quad (30)$$

Now for deriving bounds for our solution we have that for any $M > 0$:

$$\begin{aligned}
& P\left(F^{(k)}(\bar{\boldsymbol{\theta}}_1^T) + G(\bar{\boldsymbol{\theta}}_2^T) - F^{(k)}(\boldsymbol{\theta}^*) - G(\boldsymbol{\theta}^*) \geq \frac{(N^{1/p}dM_D + hD_C)^2}{2hT} \log(T+1) + \frac{\beta+h}{2T}D_C^2 + \frac{M}{\sqrt{T}}\right) \\
& \stackrel{(22)}{\leq} P\left(\frac{1}{T} \sum_{t=1}^T \delta_t \geq \frac{M}{\sqrt{T}}\right) \\
& = P\left(\sum_{t=1}^T \delta_t \geq M\sigma\sqrt{T}\right) \\
& \stackrel{(30)}{\leq} \exp\left(-\frac{M^2}{4}\right), \tag{31}
\end{aligned}$$

where for deriving the last inequality we set $M_\delta = M\sigma\sqrt{T}$ and $\alpha = \frac{M}{2\sigma\sqrt{T}}$ in (30). Eq. (31) is equivalent to (16c), where we set

$$k_1 \triangleq \max\left(\frac{(N^{1/p}dM_D + hD_C)^2 \log(3)}{2h \log(2)}, \frac{\beta+h}{2}D_C^2\right) \tag{32a}$$

$$k_2 \triangleq M. \tag{32b}$$

D Proof of Lemma 1

Proof. We show that Assumption 3 of [51] is satisfied, thus the results follow from Theorem 6 of [51]. For case (a), we need to show that the subgradient of the functions $F^{(k)}(\boldsymbol{\theta}; \mathbf{x}_t)$ are bounded. For any subgradient $\mathbf{g} \in \partial\left(\|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_t) + \mathbf{D}_t^{(k)}(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)})\|_p\right)$ and for all $\boldsymbol{\theta} \in \mathbb{R}^d$, we have

$$\begin{aligned}
\mathbf{g}^\top(\boldsymbol{\theta} - \boldsymbol{\theta}_1) & \leq \|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_t) + \mathbf{D}_t^{(k)}(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)})\|_p - \|F(\boldsymbol{\theta}^{(k)}; \mathbf{x}_t) + \mathbf{D}_t^{(k)}(\boldsymbol{\theta}_1 - \boldsymbol{\theta}^{(k)})\|_p \\
& \leq \|\mathbf{D}_t^{(k)}(\boldsymbol{\theta} - \boldsymbol{\theta}_1)\|_p \\
& = \left(\sum_{i=1}^N |\mathbf{D}_i^\top(\boldsymbol{\theta} - \boldsymbol{\theta}_1)|^p\right)^{1/p} \\
& \stackrel{\text{Cauchy-Schwarz Ineq.}}{\leq} \left(\sum_{i=1}^N (\|\mathbf{D}_i\|_2 \|\boldsymbol{\theta} - \boldsymbol{\theta}_1\|_2)^p\right)^{1/p} \\
& \stackrel{(19b)}{\leq} N^{1/p} \sqrt{d} M_D \|\boldsymbol{\theta} - \boldsymbol{\theta}_1\| = M_G \|\boldsymbol{\theta} - \boldsymbol{\theta}_1\|,
\end{aligned}$$

where \mathbf{D}_i is the i -th row of $\mathbf{D}_t^{(k)}$. Now given that the above holds for all $\boldsymbol{\theta}$, we can show that for every element $i \in [d]$, \mathbf{g}_i is bounded by M_G ; to see this set $\boldsymbol{\theta} = \boldsymbol{\theta}_1 + e_i$ and it follows that $\mathbf{g}_i \leq M_G$. We therefore conclude that the subgradients $\partial F^{(k)}(\boldsymbol{\theta}, \zeta_t) = \{\mathbf{g} + h(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}) | \mathbf{g} \in \partial F_{\boldsymbol{\theta}^{(k)}}(\boldsymbol{\theta}_1, \zeta_t), \boldsymbol{\theta} \in \mathcal{C}\}$ are bounded:

$$\begin{aligned}
\|\mathbf{g}_F\|_2 = \|\mathbf{g} + h(\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)})\|_2 & \leq \|\mathbf{g}\|_2 + h\|\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}\|_2 \\
& \leq \sqrt{d}M_G + hD_C \quad \forall \mathbf{g}_F \in \partial F^{(k)}(\boldsymbol{\theta}; \mathbf{x}_t), \tag{33}
\end{aligned}$$

where the last inequality is due to the fact that $D_{\mathcal{C}} = \max_{\theta_1, \theta_2 \in \mathcal{C}} \|\theta_1 - \theta_2\|_2$.

Case (b) is satisfied with the choice of ℓ_2 norm squared for Bregman distance, i.e., the term $\gamma \|\theta_1 - \theta_1^t\|_2^2$. Since we assumed the constraint set \mathcal{C} is convex, closed and bounded, we take the diameter of the set \mathcal{C} to be

$$D_{\mathcal{C}} = \max_{\theta_1, \theta_2 \in \mathcal{C}} \|\theta_1 - \theta_2\|_2.$$

Thus, case (c) is satisfied as a result of initialization ($\theta_1^1 = \theta_2^1 = \mathbf{u}^1 = 0$), and the fact that $\|\theta_1^1 - \theta^*\|_2 \leq D_{\mathcal{C}}$ and $\|\theta_2^1 - \theta^*\|_2 \leq D_{\mathcal{C}}$. Case (d) is directly included in the assumption of Theorem 1. Finally, for Case (e) we have

$$\begin{aligned} & \left| F^{(k)}(\theta_1^{t+1}; \mathbf{x}_t) + G(\theta_2^{t+1}) - \left(F^{(k)}(\theta^*; \mathbf{x}_t) + G(\theta^*) \right) \right| \\ & \leq \left| F^{(k)}(\theta_1^{t+1}, \zeta_t) - F^{(k)}(\theta^*, \zeta_t) \right| + \left| G(\theta_2^{t+1}) - G(\theta^*) \right| \\ & \stackrel{(a)}{\leq} (\sqrt{d}M_G + hD_{\mathcal{C}})\|\theta_1^{t+1} - \theta^*\|_2 + L_G\|\theta_2^{t+1} - \theta^*\|_2 + \\ & |\chi_{\mathcal{C}}(\theta_2^{t+1}) - \chi_{\mathcal{C}}(\theta^*)| \\ & \stackrel{(b)}{\leq} (\sqrt{d}M_G + hD_{\mathcal{C}} + L_G)D_{\mathcal{C}}, \end{aligned}$$

where we derive (a) using the Lipschitz continuity of $g(\cdot)$ along with Lemma 2.6 from [43], which states that $F^{(k)}$ is Lipschitz continuous if and only if some ℓ_p norm of its subgradients is bounded (that we showed in (33)). Moreover, in deriving (b) we use the fact that the constraint set \mathcal{C} is convex, closed, and bounded, and that $\chi_{\mathcal{C}}(\theta_2^{t+1}) = \chi_{\mathcal{C}}(\theta^*) = 0$, as $\theta_2^{t+1}, \theta^* \in \mathcal{C}$. So far we have shown that all cases in Assumption 3 of [51] are satisfied. Also, both $F^{(k)}$ and G are h and β strongly convex, respectively; the former is due to the quadratic term and the latter is explicitly stated in our assumptions. Therefore, we have shown that all assumptions in Theorem 6 [51] are satisfied and the results in (20) follow from the theorem.

E Batch of Independently Random Samples

Corollary 1. (*Batch Setting*) Assume the assumptions of Theorem 1 and let the sequence $\{\theta_1^t, \theta_2^t, \mathbf{u}^t\}, t \in [T]$ be generated by OADM algorithm, i.e., (10), where the step (10a) is replaced with the following step:

$$\theta_1^{t+1} := \arg \min_{\theta_1} \frac{1}{J} \sum_{j=1}^J F^{(k)}(\theta_1; \mathbf{x}_t^j) + \frac{\rho}{2} \|\theta_1 - \theta_2^t + \mathbf{u}^t\|_2^2 + \frac{\gamma}{2} \|\theta_1 - \theta_1^t\|_2^2, \quad (34)$$

where at each iteration $t \in [T]$, the points $\mathbf{x}_t^j \in \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ are i.i.d. samples drawn uniformly at random. Then all the results of Theorem 1 in (16) hold.

Proof. We only need to adapt the definition in (8a) to this batch setting as follows: (8a) as follows

$$F^{(k)}(\boldsymbol{\theta}; [\mathbf{x}^j]_{j \in [J]}) \triangleq \frac{1}{J} \sum_{j=1}^J F_{\boldsymbol{\theta}^{(k)}}(\boldsymbol{\theta}, \mathbf{x}^j) + \frac{h}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{(k)}\|_2^2. \quad (35)$$

Then all the results follow similar to the proof presented in App. C.

F Experimental Details

Algorithm Hyperparameters and Stopping Criteria. For all algorithms we use a batch size of 8; we avoid using larger batch sizes as the computation time increases. Additional details and stopping criteria used for each algorithm are as follows:

- **MBOSADM:** We run Alg. 1 for 20 iterations, however we stop earlier if we do not see an improvement in the objective. In the k -th iteration of Alg. 1, for solving the sub-problems (Line 3), we run iterations of SADP (10) for a maximum of 200 rounds or until the primal and dual residuals are less than $0.95^k \epsilon$, where we set $\epsilon = 0.1$ and 0.001 for training auto-encoders and multi-target regression, respectively.
- **MBOSGD:** Again, we run Alg. 1 for 20 iterations and stop earlier if the objective does not improve. For solving the sub-problems (Line 3), we run the stochastic gradient descent (SGD) with the learning rate of 10^{-5} and 10^{-3} , respectively, for training auto-encoders and multi-target regression. In all cases we set the momentum parameter to 0.9 and run SGD for 500 iterations.
- **SGD:** We run the stochastic gradient descent (SGD) algorithm with momentum with the learning rate of 10^{-6} and the momentum parameter set to 0.9. This corresponds to the algorithm by Mai and Johansson [27] applied to our setting. We run SGD for 10^3 and 10^4 iterations for training auto-encoders and multi-target regression, respectively; we observe that the algorithm achieves its minimum within this number of iterations. At each iteration, SGD evaluates the objective F_{OBJ} and outputs the solution for the best observed objective.

Implementation. We implement all the algorithms in Python 3.7 and using the PyTorch backend. We run all algorithms on CPU machines that have Intel(R) Xeon(R) CPUs (E5-2680 v4) with 2.4GHz clock speed.

Applications. For both applications, we set the regularizer as $g(\boldsymbol{\theta}) = \frac{0.001}{2} \|\boldsymbol{\theta}\|_2^2$ and do not consider a constraint set, i.e., $\mathcal{C} = \mathbb{R}^d$.

- **Training Autoencoders.** We use a neural network with two convolutional and two de-convolutional layers; the convolutional layers have 8 and 4 output channels, respectively, and 3×3 kernel weights. The de-convolutional layers exactly mirror the convolutional layers. We do not apply zero padding or dilation for any of the layers and use a convolution step size of 1. We apply a soft-plus activation function after each layer.

- **Multi-target Regression.** We use a network with two layers; the first layer is a 1-dimensional convolutional layer with the kernel size of 3 and no zero padding or dilation and the step size of 1. The second layer is a fully-connected layer with 278 hidden units and output size of 16 (the target size). We again apply the soft-plus activation after each layer.

Note that we choose soft-plus activation, i.e., a smooth version of the ReLu, to make sure that the functions $F(\boldsymbol{\theta}; \mathbf{x})$ are smooth.

Datasets. For each of the applications we use two datasets:

- **Multi-Target Regression.** We use `SCM1d`, from the collection of regression data made available by [44]. `SCM1d`² is a supply chain management dataset comprising 9803 samples with 280 predictors and 16 targets. We use 80 percent of data for training, i.e., solving (3), and the rest of it as the test set.
- **Auto-encoders.** We use two well-known datasets `MNIST`³ and `Fashion-MNIST`⁴. They both have greyscale 28×28 images with 60,000 training samples and 10,000 testing samples. `MNIST` contains handwritten digits with 10 classes and `Fashion-MNIST` contains images of clothing items from 10 classes.

G Classification Task

For autoencoder tasks, we train a logistic regression model, where the input features are the outputs of the first two convolutional layers used for encoding data. We set the parameters of the encoder to be the solutions obtained by the respective autoencoder training algorithm. Using the encoder output for training sets and the corresponding target labels, we train the logistic regression, using the ℓ_2 regularizer. We train the latter with different regularizer coefficients (0.01, 0.1, and 1.0) and report the best observed accuracy on the test sets in Fig. 3. Full experimental results for `MNIST` and `Fashion-MNIST` are reported in Tables 2 and 3, respectively.

H Robustness Analysis

Here we report non-outliers loss and test loss for `MNIST` for $p = 1.5$. Similar to results for $p = 1, 2$ in Fig. 2, we see that both `MBOSADM` variants and `SGD` generally stay constant w.r.t P_{out} , for $p = 1.5$ as well. Moreover, `MBO` variants again obtain lower loss values. We also see that in the high outlier regime, $P_{\text{out}} = 0.3$, `MBOSADM` performs poorly, similar to Figures 2c and 2f.

² <http://users.auth.gr/espyromi/mtr/mtr-datasets.zip>

³ <https://pytorch.org/vision/stable/datasets.htm#mnist>

⁴ <https://pytorch.org/vision/stable/datasets.html#fashion-mnist>

P_{out}	p	Acc.		
		MBOSADM	MBOSGD	SGD
0.0	ℓ_2^2	-	-	0.876
0.0	2.0	0.926	0.894	0.868
0.0	1.5	0.930	0.882	0.861
0.0	1.0	0.936	0.904	0.861
<hr/>				
0.05	ℓ_2^2	-	-	0.928
0.05	2.0	0.925	0.921	0.914
0.05	1.5	0.928	0.929	0.867
0.05	1.0	0.935	0.930	0.863
<hr/>				
0.1	ℓ_2^2	-	-	0.907
0.1	2.0	0.925	0.924	0.916
0.1	1.5	0.927	0.930	0.877
0.1	1.0	0.935	0.926	0.873
<hr/>				
0.2	ℓ_2^2	-	-	0.897
0.2	2.0	0.928	0.856	0.914
0.2	1.5	0.928	0.928	0.892
0.2	1.0	0.931	0.931	0.894
<hr/>				
0.3	ℓ_2^2	-	-	0.867
0.3	2.0	0.923	0.925	0.913
0.3	1.5	0.927	0.929	0.914
0.3	1.0	0.927	0.922	0.918

Table 2: Accuracy of the classifiers for MNIST dataset.

P_{out}	p	Acc.		
		MBOSADM	MBOSGD	SGD
0.0	ℓ_2^2	-	-	0.754
0.0	2.0	0.836	0.824	0.787
0.0	1.5	0.848	0.843	0.793
0.0	1.0	0.853	0.853	0.794
<hr/>				
0.05	ℓ_2^2	-	-	0.770
0.05	2.0	0.835	0.832	0.806
0.05	1.5	0.841	0.847	0.803
0.05	1.0	0.850	0.847	0.814
<hr/>				
0.1	ℓ_2^2	-	-	0.806
0.1	2.0	0.835	0.836	0.812
0.1	1.5	0.837	0.846	0.811
0.1	1.0	0.854	0.848	0.783
<hr/>				
0.2	ℓ_2^2	-	-	0.754
0.2	2.0	0.826	0.832	0.821
0.2	1.5	0.815	0.837	0.815
0.2	1.0	0.849	0.832	0.758
<hr/>				
0.3	ℓ_2^2	-	-	0.757
0.3	2.0	0.817	0.827	0.823
0.3	1.5	0.838	0.837	0.826
0.3	1.0	0.846	0.819	0.783

Table 3: Accuracy of the classifiers for Fashion-MNIST dataset.

I Proximal Operator for ℓ_p norms

The proximal operator for ℓ_p norms ($p \geq 1$), is the following:

$$\min_{\mathbf{u} \in \mathbb{R}^d} \|\mathbf{u}\|_p + \frac{\rho}{2} \|\mathbf{u} - \mathbf{w}\|_2^2, \quad (36)$$

for a given $\mathbf{w} \in \mathbb{R}^d$. Liu and Ye [25] define a non-negative vector $\hat{\mathbf{w}}$ via

$$\hat{w}_i = \rho |w_i| \quad \forall i \in [d]. \quad (37)$$

They then consider the following simpler problem:

$$\min_{\mathbf{u} \in \mathbb{R}_+^d} \|\mathbf{u}\|_p + \frac{1}{2} \|\mathbf{u} - \hat{\mathbf{w}}\|_2^2. \quad (38)$$

Note that this differs from Prob. (36) in that (a) $\rho = 1$ and (b) vector $\mathbf{w} \in \mathbb{R}^d$ replaced with non-negative vector $\mathbf{w} \in \mathbb{R}_+^d$, and (c) optimization happens over $\mathbf{u} \in \mathbb{R}_+^d$. Nevertheless, Prob. (36) is equivalent to Prob. (38) [30].

Liu and Ye [24] and Moharrer et al. [30] then solve Prob. (38). To do so, they define first an auxiliary function as follows. Given $\alpha \in (0, \infty)$, they define the function $\alpha \mapsto g(\alpha)$, as the unique solution of the following equation over $x \geq 0$:

$$(x/\alpha)^{p-1} + x - 1 = 0, \quad (39)$$

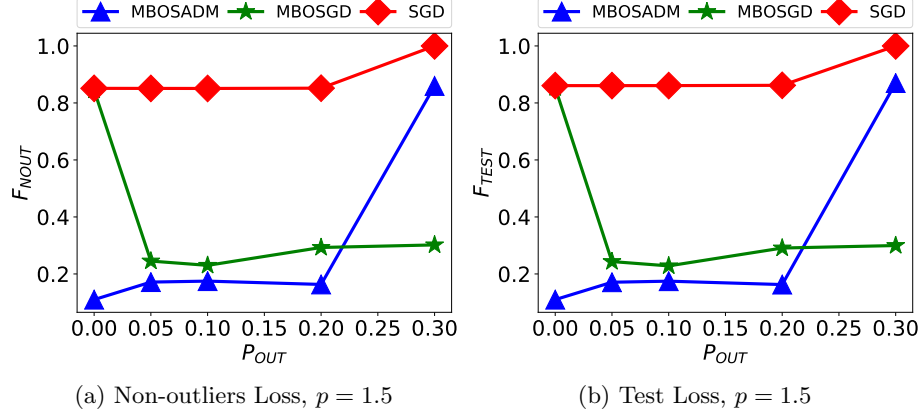


Fig. 4: Scalability of the non-outliers loss F_{NOUT} and the test loss F_{TEST} for $p = 1.5$. Similar to Fig. 2, we see that both MBOSADM variants and SGD generally stay constant w.r.t P_{out} when for $p = 1.5$. We see that MBO variants obtain lower loss values. We also see that in the high outlier regime, $P_{out} = 0.3$, MBOSADM performs poorly, similar to Figures 2c and 2f.

Algorithm 3 p -norm Prox. Operator

- 1: **Input:** $\mathbf{w} \in \mathbb{R}^d$, $p \geq 1$, $\rho > 0$, $\varepsilon > 0$
 - 2: Set $\hat{w}_i \leftarrow \rho|w_i|$ for $i = 1, \dots, d$.
 - 3: **if** $\|\hat{\mathbf{w}}\|_q \leq 1$ **then**
 - 4: **Return** $\mathbf{u}^* \leftarrow \mathbf{0}$
 - 5: **end if**
 - 6: Set $\mathbf{u} \leftarrow \mathbf{0}$, $s_L \leftarrow 0$, and $s_U \leftarrow \|\hat{\mathbf{w}}\|_p$
 - 7: **for** $k = 1, \dots, \log_2 \lceil \frac{1}{\varepsilon} \rceil$ **do**
 - 8: Set $s \leftarrow (s_L + s_U)/2$
 - 9: Compute $u_i \leftarrow \hat{w}_i g\left(s \cdot (\hat{w}_i)^{\frac{2-p}{p-1}}\right)$ for all $i \in [d]$;
 - 10: Compute $\|\mathbf{u}\|_p$;
 - 11: **if** $\|\mathbf{u}\|_p < s$ **then**
 - 12: Set $s_U \leftarrow s$
 - 13: **else**
 - 14: Set $s_L \leftarrow s$
 - 15: **end if**
 - 16: **end for**
 - 17: Set $u_i^* \leftarrow \frac{\text{sign } w_i}{\rho} u_i$ for $i = 1, \dots, d$.
 - 18: **Return** \mathbf{u}^*
-

They extend g to $[0, \infty)$ by setting $g(0) \equiv 0$ for $\alpha = 0$, by definition. Having defined g , given a vector $\hat{\mathbf{w}} \in \mathbb{R}_+^d$, they also define functions $g_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $i \in [d]$ as:

$$g_i(s) = \hat{w}_i \cdot g\left(s \cdot (\hat{w}_i)^{\frac{2-p}{p-1}}\right), \quad (40)$$

as well as function $h : \mathbb{R}_+ \rightarrow \mathbb{R}$ as:

$$h(s) = \left(\sum_{i=1}^d g_i(s)^p \right)^{\frac{1}{p}} - s. \quad (41)$$

The bisection algorithm [25,30] for solving (38), summarized in Alg. 3, proceeds as follows: given $\hat{\mathbf{w}} \in \mathbb{R}_+^d$, it tests whether the condition $\|\hat{\mathbf{w}}\|_q \leq 1$ holds; if so, it returns $\mathbf{u}^* = \mathbf{0}$. Otherwise, Liu and Ye [25] show that the solution is a function of s^* , i.e., the root of the equation $h(s^*) = 0$. Alg. 3 finds a root s^* via bisecting $[0, \|\hat{\mathbf{w}}\|_p]$. That is, at each iteration, it maintains an upper (s_U) and lower (s_L) bound on s^* , initialized at the above values. By construction, function h alternates signs on each of the two bounds: i.e., $h(s_L)h(s_U) \leq 0$; at each iteration, the method (a) computes the average $s = 0.5(s_L + s_U)$, between the two bounds, (b) find the sign of h on this average, and then (c) update the bounds accordingly. For more information on the guarantees of the algorithm see Theorems 4.2 and 4.3 in [30].