
On Coresets for Regularized Regression

Rachit Chhaya¹ Anirban Dasgupta¹ Supratim Shit¹

Abstract

We study the effect of norm based regularization on the size of coresets for regression problems. Specifically, given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with $n \gg d$ and a vector $\mathbf{b} \in \mathbb{R}^n$ and $\lambda > 0$, we analyze the size of coresets for regularized versions of regression of the form $\|\mathbf{Ax} - \mathbf{b}\|_p^r + \lambda \|\mathbf{x}\|_q^s$. Prior work has shown that for ridge regression (where $p, q, r, s = 2$) we can obtain a coreset that is smaller than the coreset for the unregularized counterpart i.e. least squares regression (Avron et al., 2017). We show that when $r \neq s$, no coreset for regularized regression can have size smaller than the optimal coreset of the unregularized version. The well known lasso problem falls under this category and hence does not allow a coreset smaller than the one for least squares regression. We propose a modified version of the lasso problem and obtain for it a coreset of size smaller than the least square regression. We empirically show that the modified version of lasso also induces sparsity in solution, similar to the original lasso. We also obtain smaller coresets for ℓ_p regression with ℓ_p regularization. We extend our methods to multi response regularized regression. Finally, we empirically demonstrate the coreset performance for the modified lasso and the ℓ_1 regression with ℓ_1 regularization.

1. Introduction

Most of the applications of machine learning require huge amounts of data to train models. As the algorithms to train models run in time dependent on the size of the data, the task becomes computationally inefficient. A coreset is a subsample of appropriately reweighted points from the original data which can be used to train models with competitive accuracy and provable guarantees. The size of a coreset is

¹Computer Science and Engineering, Indian Institute of Technology Gandhinagar, Gandhinagar, India . Correspondence to: Rachit Chhaya <rachit.chhaya@iitgn.ac.in>.

usually independent of the size of the original dataset making training on them much quicker. Sampling based coresets also allow better interpretability in contrast to data reduction techniques which project data points in some other space.

Regression is a widely used technique in machine learning and statistics, the most popular variants being the least square regression (ℓ_2 regression) and least absolute deviation (ℓ_1 regression). A coreset construction for ℓ_2 regression based on leverage scores is given in (Drineas et al., 2006) whereas coresets for ℓ_p regression have been created based either on norms of the so called well-conditioned basis (Dasgupta et al., 2009; Sohler & Woodruff, 2011) or based on Lewis weights (Cohen & Peng, 2015).

A common variant of regression is to use regularization, meant to either achieve numerical stability, to prevent overfitting or to induce sparsity in the solution— ridge and lasso being the most commonly used regularizers. Since regularization imposes a constraint on the solution space, we can potentially expect regularized problems to have a smaller size coreset, than the unregularized version. Indeed, this intuition has been formalized in the case of ridge regression (Woodruff et al., 2014) using the ridge leverage scores. Pilanci et al. (Pilanci & Wainwright, 2015) construct small sized coresets for constrained version of lasso using random projections; these coresets are, however, not (scaled) subsample of data points. To the best of our knowledge ours is the first work to study coresets for ℓ_p regularized regression for $p \neq 2$. Our first result is negative— we show that it is not always possible to build smaller coresets for regularized regression. For a specific class of problems we show that smaller coresets are possible and we show how to construct them.

To construct such coresets, we follow the sensitivity framework given in (Feldman & Langberg, 2011). Sensitivities defined in (Langberg & Schulman, 2010) capture the importance of a data point for a particular optimization function. If we sample points using a probability distribution based on sensitivities, the sample size depends on the sum of the sensitivities and the dimension of the solution space. The core idea behind our bounds is that due to regularization, the sensitivities of points change – while the sensitivity of very low sensitivity points might increase slightly (by additive $1/n$), the higher sensitivities are pulled down. The

overall effect is that the sum of sensitivities for a regularized version of regression is less than the sum of sensitivities for its unregularized counterpart by a factor that depends on the value of the regularizer.

1.1. Our Contributions

- We first show that for any regularized problem of the form $\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x}\|_p^r + \lambda \|\mathbf{x}\|_q^s$, where $r \neq s$, a coreset for the problem also works as a coreset for its unregularized version. This implies that when $r \neq s$, we cannot build coresets with size smaller than the ones for the unregularized version.
- We introduce a modified version of lasso regression for which we show that we can construct coresets with size smaller than that of the unregularized linear regression.
- We calculate sensitivity upper bounds for the ℓ_p regression with ℓ_p regularization and construct coresets for the problem. We focus on $p \geq 1, p \neq 2$. Specifically we give smaller coreset for the regularized least deviation problem i.e. for $p = 1$. We extend our result to multiple response regression also.
- We show experimental evidence that the modified lasso problem also preserves sparsity like the lasso problem and hence is a suitable replacement. We also demonstrate the empirical performance of our sampling probabilities for the modified lasso problem and the regularized least deviation problem. We also provide open code for the same [here](#)

1.2. Organization of the Paper

The rest of the paper is organized in the following manner. In section 2 we discuss work in the areas related to coresets in general as well as coresets specifically for regularized problems and distinguish our work from existing work. In section 3 we provide all the notations, definitions and existing results that we use throughout the paper. In section 4 we give our main result relating coresets for regularized regression to the ones for unregularized regression. In section 5 we introduce the modified lasso problem and analyze the size of coreset for it. In section 6 we show smaller coresets for ℓ_p regression with ℓ_p regularization. In section 7 we report the experiments validating our claims and conclude in section 8 with discussion on future scope in this area.

2. Related Work

Coresets are small summaries of data which can be used as a proxy for the original data with provable guarantees. The term was first introduced in (Agarwal et al., 2004) where they used coresets for the shape fitting problem. Coresets for

clustering problem were described in (Har-Peled & Mazumdar, 2004). In (Feldman & Langberg, 2011) authors gave a generalized framework to construct coresets based on importance sampling using sensitivity scores introduced in (Langberg & Schulman, 2010) which was improved in (Braverman et al., 2016) for both offline and streaming settings.

There is a large amount of work to reduce size of the data for the ℓ_2 regression problem (Drineas et al., 2011; Clarkson & Woodruff, 2017). Interested readers may refer to a good survey (Woodruff et al., 2014) and references therein. Coresets for ℓ_2 regression can be obtained using the popular leverage scores (Drineas et al., 2006). Work has also been done to obtain these leverage scores in an efficient manner (Drineas et al., 2012). Practical effectiveness of coresets for least square problems has been recently shown by (Maalouf et al., 2019) compared to available libraries like scikit-learn. Coresets for ℓ_p regression are obtained using the row norms of the well-conditioned basis (Dasgupta et al., 2009) or the Lewis weights (Cohen & Peng, 2015). Other works that construct coresets for ℓ_p regression in general or ℓ_1 regression in particular include (Sohler & Woodruff, 2011; Meng & Mahoney, 2013; Woodruff & Zhang, 2013; Clarkson et al., 2016; Dickens et al., 2018).

However not much has been done for regularized version of the regression problem. Pilanci et al. (Pilanci & Wainwright, 2015) are able to reduce the size of the popular lasso problem using random projections. However they work with the constrained version of the problem and not the regularized one. Also they do not obtain a sample of the original dataset. Reddi et al. (Reddi et al., 2015) obtained additive error coresets for empirical loss minimization. Coresets using local sensitivity scores (Raj et al., 2019) also work for functions with a regularization term. Here they use a quadratic approximation of the function and then use leverage scores to approximate sensitivities locally. Tolochinsky et al. (Tolochinsky & Feldman, 2018) actually add a regularization term to obtain coresets for functions for which otherwise sublinear coresets may not exist.

Coresets for logistic regression and SVM with ℓ_2^2 regularization were obtained by (Curtin et al., 2019) using uniform sampling. They considered empirical loss minimization problems of the form $f(\mathbf{w}) = \sum_{i=1}^n l(-y_i \mathbf{w}^T \mathbf{x}_i) + \lambda r(R\mathbf{w})$ where \mathbf{x}_i is the i^{th} data point with corresponding label $y_i \in \{+1, -1\}$ and \mathbf{w} is solution vector. The parameter R is maximum 2-norm of a row in data matrix. Their coresets do not work for a general response vector as in regression problems. In addition, for uniform sampling to work, they assume a (σ, τ) condition on the loss and the regularization functions l and r respectively which says that $l(-\sigma) > 0$, and if $\|\mathbf{w}\|_2 \geq \sigma$ then $r(\mathbf{w}) \geq \tau l(\|\mathbf{w}\|_2)$. If we consider response vectors consisting of $\{-1, +1\}$, we can

formulate the regression problem to fit their loss function template in the following manner— first create augmented $\mathbf{x}' = \begin{bmatrix} \mathbf{x} \\ y_i \end{bmatrix}$ and $\mathbf{w}' = \begin{bmatrix} \mathbf{w} \\ -1 \end{bmatrix}$. The lasso problem, under the restriction that $y \in \{-1, 1\}^d$, can be represented in their framework using \mathbf{x}' and \mathbf{w}' , setting $l(s) = s^2$ and $r(\mathbf{w}) = \|\mathbf{w}\|_1$. Here $l(-\sigma) > 0$ for all values of σ . However, there is no τ that satisfies $\|\mathbf{w}\|_1 \geq \tau \|\mathbf{w}\|_2^2$ for all \mathbf{w} with $\|\mathbf{w}\|_2 \geq \sigma$. So their framework does not apply to our case.

In spirit, our work most closely relates to the work in (Avron et al., 2017). Here the authors show that for ridge regression, coresets of size smaller than the coreset for the unregularized least square regression problem can be constructed. Their coreset size is a function of the statistical dimension of the matrix for some regularization parameter $\lambda > 0$. It is important to note that a coreset for unregularized regression can be shown to work for regularized regression as well. However the work in (Avron et al., 2017) supports the intuition that coresets for regularized regression may be smaller than coresets for its unregularized counterpart. We generalize this idea from ridge regression to ℓ_p regression with ℓ_p regularization in this paper. We also show that for a broad class of regularized regression problems, it is not possible to construct coreset smaller in size compared to its unregularized form. Another recent related work on ridge regression is (Kacham & Woodruff, 2020). In this paper they construct a deterministic coreset for ridge regression by using the BSS (Batson et al., 2012) technique and improve the dependence on ϵ . They also provide lower bounds showing this is tight. Our lower bound is for the case $r \neq s$ in regularized regression, and is not applicable for ridge regression. Our coresets are randomized, and are for regularized ℓ_p regression for all values of p . The BSS technique only works for $p = 2$. Constructing deterministic coresets for regularized regression for all p is an interesting open question.

3. Notations and Preliminaries

In this section we describe all the notations, definitions and existing results that we will use throughout this paper. A scalar or a function (specified by context) is denoted by a lower case letter, e.g. x while a vector is denoted by a boldface lower case letter, e.g. \mathbf{x} . A matrix or a dataset, as defined by the context, is denoted by a boldface uppercase letter, e.g. \mathbf{X} . A non bold face upper case letter also represents some scalar value unless stated otherwise. We represent a dataset with a matrix where the rows of the matrix represent data points and the columns represent the features. Vectors are considered column vectors unless stated otherwise. \mathbf{a}_i represents the i^{th} row of matrix \mathbf{A} while \mathbf{a}^k represents its k^{th} column and a_{ij} represents the

entry in its i^{th} row and j^{th} column. O represent big-O in the ordered notation. For two quantities x and b and a suitable $\epsilon > 0$ we will often write $x \in (1 \pm \epsilon)b$ to compactly represent that $(1 - \epsilon)b \leq x \leq (1 + \epsilon)b$. All statements where we say “high probability”, hold with probability at least some large constant, e.g. 0.99, unless otherwise stated.

For a vector $\mathbf{x} \in \mathbb{R}^d$, $\|\mathbf{x}\|_p$, $p \geq 1$ represents its p^{th} norm which is defined as $\|\mathbf{x}\|_p = (\sum_{j=1}^d |x_j|^p)^{\frac{1}{p}}$. The infinity norm for a vector is given as $\|\mathbf{x}\|_\infty = \max_j |x_j|$. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\|\mathbf{A}\|_p$ denotes its entry wise p norm which is defined as $\|\mathbf{A}\|_p = (\sum_{i=1}^n \sum_{j=1}^d |a_{ij}|^p)^{\frac{1}{p}}$. For $p = 2$, this is also called the Frobenius norm and is also represented as $\|\mathbf{A}\|_F$. We differentiate the matrix entry wise p norm from the matrix induced operator p norm which is denoted by the symbol $\|\mathbf{A}\|_{(p)}$ and is defined as $\|\mathbf{A}\|_{(p)} = \sup_{\mathbf{x} \neq 0} \frac{\|\mathbf{A}\mathbf{x}\|_p}{\|\mathbf{x}\|_p}$. Specifically for $p = 2$, $\|\mathbf{A}\|_{(2)} = \sigma_{\max}(\mathbf{A})$ i.e. the highest singular value of matrix \mathbf{A} and for $p = 1$, $\|\mathbf{A}\|_{(1)}$ is the highest absolute column sum of the matrix \mathbf{A} . The Singular Value Decomposition (abbreviated as SVD) of a matrix \mathbf{A} is given as $\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^T$ where $\mathbf{U} \in \mathbb{R}^{n \times d}$ and $\mathbf{V} \in \mathbb{R}^{d \times d}$ represent \mathbf{A} 's left and right singular vectors respectively. $\Sigma \in \mathbb{R}^{d \times d}$ is a diagonal matrix containing the singular values of \mathbf{A} in descending order.

3.1. Coresets

A coreset is a small summary of data which can give provable guarantees for a particular optimization problem. Formally, given a dataset $\mathbf{X} \in \mathbb{R}^{n \times d}$ where the rows represent the datapoints, set of queries Q and a non-negative cost function $f_{\mathbf{q}}(\mathbf{x})$ with parameter $\mathbf{q} \in Q$ and data point $\mathbf{x} \in \mathbf{X}$, a dataset consisting of subsampled and appropriately reweighted points \mathbf{C} is an ϵ -coreset if $\forall \mathbf{q} \in Q$,

$$\left| \sum_{\mathbf{x} \in \mathbf{X}} f_{\mathbf{q}}(\mathbf{x}) - \sum_{\tilde{\mathbf{x}} \in \mathbf{C}} f_{\mathbf{q}}(\tilde{\mathbf{x}}) \right| \leq \epsilon \sum_{\mathbf{x} \in \mathbf{X}} f_{\mathbf{q}}(\mathbf{x})$$

for some $\epsilon > 0$. We will denote $\sum_{\mathbf{x} \in \mathbf{X}} f_{\mathbf{q}}(\mathbf{x}) = f_{\mathbf{q}}(\mathbf{X})$. Even if the function is defined over entire dataset \mathbf{X} and not on individual points we can define coreset for it in terms of $f_{\mathbf{q}}(\mathbf{X})$ directly. For e.g. $\|\mathbf{A}\mathbf{x}\|_p^r$ cannot be represented in terms of sum on individual \mathbf{a}_i 's but we can still define coreset property for it. When the above equation is satisfied $\forall \mathbf{q} \in Q$, \mathbf{C} is called a strong coreset. We will refer only to a strong coreset as coreset in this paper. The advantage of a strong coreset, specifically in machine learning problems, is that we can train the model on the coreset, i.e. a smaller data set, and then use the model obtained from it as a surrogate for the model trained on the original data, with comparable accuracy. Formally, via (Bachem et al., 2017), let \mathbf{C} be an ϵ -coreset of \mathbf{X} for $\epsilon \in (0, \frac{1}{3})$ for some function. If $\mathbf{q}_{\mathbf{X}}^*$ and $\mathbf{q}_{\mathbf{C}}^*$ denote the optimal (infimum) solutions for \mathbf{X} and \mathbf{C} respectively, then $f_{\mathbf{q}_{\mathbf{C}}^*}(\mathbf{X}) \leq (1 + 3\epsilon)f_{\mathbf{q}_{\mathbf{X}}^*}(\mathbf{X})$. Given a dataset \mathbf{X} , the query space Q , as well as the cost function

$f_{\mathbf{q}}(\cdot)$, Langberg et al. (Langberg & Schulman, 2010) define a set of scores, termed as sensitivities, that can be used to create coresets via importance sampling. The sensitivity of the i^{th} point \mathbf{x}_i is defined as $s_i = \sup_{\mathbf{q} \in Q} \frac{f_{\mathbf{q}}(\mathbf{x}_i)}{\sum_{\mathbf{x}' \in \mathbf{X}} f_{\mathbf{q}}(\mathbf{x}')}$. Intuitively sensitivity of a point captures its worst case contribution to the value of an objective function. The work by Langberg et al. (Langberg & Schulman, 2010) shows that using any upper bounds to the sensitivity scores, we can create a probability distribution, using which we can then sample a coreset. Throughout the paper we refer to the sensitivity of the i^{th} data point for some objective function as s_i and the function will be clear from the context. Feldman et al. (Feldman & Langberg, 2011) provided a unified framework to build coresets using the sensitivity framework which was improved by Braverman et al. (Braverman et al., 2016). Their theorem on coreset size using sensitivity is the following:

Theorem 3.1. (Braverman et al., 2016) *Let $f_{\mathbf{q} \in Q}(\cdot)$ be the function, \mathbf{X} be the dataset and query space Q be of dimension d . Let the sum of sensitivities be S . Let $(\epsilon, \delta) \in (0, 1)$. Let r be such that*

$$r \geq \frac{10S}{\epsilon^2} \left(d \log S + \log \frac{1}{\delta} \right).$$

\mathbf{C} be a matrix of r rows, each sampled i.i.d from \mathbf{X} such that for every $\mathbf{x}_i \in \mathbf{X}$ and $\tilde{\mathbf{x}}_i \in \mathbf{C}$, $\tilde{\mathbf{x}}_i$ is a scaled version of \mathbf{x}_i with probability $\frac{s_i}{S}$ and scaled by $\frac{S}{rs_i}$, then \mathbf{C} is an ϵ -coreset of \mathbf{X} for function f , with probability at least $1 - \delta$.

We present a slightly better version of this theorem (in terms of dependence on S) obtained by using a tighter tail inequality and use it for our coreset size. The proof of the following theorem is present in Appendix.

Theorem 3.2. *For the same setup as defined in Theorem 3.1, if $r = O\left(\frac{S}{\epsilon^2} \left(d \log \frac{1}{\epsilon} + \log \frac{1}{\delta}\right)\right)$, \mathbf{C} is an ϵ -coreset of \mathbf{X} for function f , with probability at least $1 - \delta$.*

3.2. ℓ_p Regression and Regularization

The ℓ_p regression problem is defined as follows: given \mathbf{A} and \mathbf{b} , find $\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_p^p$. For $p = 1, 2$ the problems are referred to as Least Absolute Deviation and Least Squares Regression respectively. We call a matrix $\mathbf{\Pi}$ to have *subspace preserving* property if, $\forall \mathbf{x} \in \mathbb{R}^d$, $\|\mathbf{\Pi}\mathbf{A}\mathbf{x}\|_p - \|\mathbf{A}\mathbf{x}\|_p \leq \epsilon \|\mathbf{A}\mathbf{x}\|_p$.

Notice that if a sampling and reweighing matrix $\mathbf{\Pi}$ satisfies the ℓ_p subspace preservation property, it will also provide a coreset for the ℓ_p regression problem. If we create matrix \mathbf{A}' by concatenating \mathbf{A} and \mathbf{b} as $\mathbf{A}' = [\mathbf{A} \ \mathbf{b}]$ and consider $\mathbf{x}' = \begin{bmatrix} \mathbf{x} \\ -1 \end{bmatrix}$, then a subspace preserving property of $\mathbf{\Pi}$ in \mathbb{R}^{d+1} implies that $\mathbf{\Pi}\mathbf{A}$ is a coreset for ℓ_p regression also. For $p = 2$, an $O\left(\frac{d \log d}{\epsilon^2}\right)$ sized coreset can be obtained by

sampling using the popular leverage scores which are the squared Euclidean row-norms of any orthogonal column basis of \mathbf{A} (Drineas et al., 2006). A matrix \mathbf{U} is called an (α, β, p) well-conditioned basis for \mathbf{A} if $\|\mathbf{U}\|_p \leq \alpha$ and $\forall \mathbf{x} \in \mathbb{R}^d, \|\mathbf{x}\|_q \leq \beta \|\mathbf{U}\mathbf{x}\|_p$ where $\frac{1}{p} + \frac{1}{q} = 1$. In fact, the \mathbf{U} obtained using the SVD of \mathbf{A} is a $(\sqrt{d}, 1, 2)$ well-conditioned basis of \mathbf{A} . For other values of $p > 1$ (Dasgupta et al., 2009) showed that by sampling using the p^{th} power of the p norm of rows of the (α, β, p) well-conditioned basis of \mathbf{A} , we can obtain a coreset of size $\tilde{O}(\alpha\beta)^p$ with high probability for ℓ_p subspace preservation and hence ℓ_p regression. Well-conditioned basis for p norm can be constructed in various ways like using Cauchy random variables (Sohler & Woodruff, 2011) or exponential random variables (Woodruff & Zhang, 2013). Yang et al. provide a good review of methods (Yang et al., 2015). Our analysis works with any method of constructing the well-conditioned basis. For the norm based regularized regression we consider the following general form for $\lambda > 0$

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_p^r + \lambda \|\mathbf{x}\|_q^s$$

for $p, q \geq 1$ and $r, s > 0$. Notice that not all regularized regression forms falling under above category can be expressed as sum of individual functions. A coreset for this problem is $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}})$ such that $\forall \mathbf{x} \in \mathbb{R}^d$ and $\forall \lambda > 0$,

$$\|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|_p^r + \lambda \|\mathbf{x}\|_q^s \in (1 \pm \epsilon)(\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_p^r + \lambda \|\mathbf{x}\|_q^s)$$

It is not difficult to verify that a coreset for the unregularized version of regression is also a coreset for the regularized one. However, to reiterate, our goal is to analyze whether the size of coreset for the above form of regularized regression can be shown to be provably smaller than the size of the coreset for the unregularized counterpart. We answer this question in the following sections. We first show a negative result where the regularization does not result in a smaller coreset. Next, we show settings where the size of the coreset does decrease inversely with λ . Analogous to the subspace embedding, given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, a matrix \mathbf{A}_c with k rows is a coreset for $(\|\mathbf{A}\mathbf{x}\|_p^r + \lambda \|\mathbf{x}\|_q^s)$ if $\forall \mathbf{x} \in \mathbb{R}^d$ and $\forall \lambda > 0$, the following holds.

$$\begin{aligned} (1 - \epsilon)(\|\mathbf{A}\mathbf{x}\|_p^r + \lambda \|\mathbf{x}\|_q^s) &\leq \|\mathbf{A}_c\mathbf{x}\|_p^r + \lambda \|\mathbf{x}\|_q^s \\ &\leq (1 + \epsilon)(\|\mathbf{A}\mathbf{x}\|_p^r + \lambda \|\mathbf{x}\|_q^s) \end{aligned} \quad (1)$$

It is easy to see that such a coreset will also give a coreset for the ℓ_p regression with ℓ_p regularization by using the concatenated matrix \mathbf{A}' and vector \mathbf{x}' .

4. Coresets For General Form of Regularized Regression

For the case of ridge regression ($p, q, r, s = 2$), Avron et al. (Avron et al., 2017) showed smaller coresets with

size dependent on the statistical dimension of the matrix \mathbf{A} . They constructed their coreset by sampling according to the ridge leverage scores. In this section we show that it is not always possible to get a coreset for a regularized version of regression problem which is strictly smaller in size than the coreset for its unregularized counterpart. Note that, for the purposes of the following theorem, a coreset of \mathbf{A} is any matrix \mathbf{A}_c that satisfies the coreset condition; in particular, \mathbf{A}_c does not need to be a *sampling based coreset* of \mathbf{A} .

Theorem 4.1. *Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\lambda > 0$, any coreset for the problem $\|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s$, where $r \neq s$, $p, q \geq 1$ and $r, s > 0$, is also a coreset for $\|\mathbf{Ax}\|_p^r$.*

Proof. The proof is by contradiction. Let \mathbf{A}_c be a coreset for $\|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s$, where $r \neq s$. Therefore, by definition of coreset, $\forall \mathbf{x} \in \mathbb{R}^d$,

$$\|\mathbf{A}_c \mathbf{x}\|_p^r + \lambda \|\mathbf{x}\|_q^s \in (1 \pm \epsilon) (\|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s)$$

Suppose \mathbf{A}_c is not a coreset for $\|\mathbf{Ax}\|_p^r$. We consider the two cases:

Case 1: $\exists \mathbf{x} \in \mathbb{R}^d$ s.t. $\|\mathbf{A}_c \mathbf{x}\|_p^r > (1 + \epsilon) \|\mathbf{Ax}\|_p^r$. Define ϵ' to be such that $\|\mathbf{A}_c \mathbf{x}\|_p^r = (1 + \epsilon') \|\mathbf{Ax}\|_p^r$. Clearly, $\epsilon' > \epsilon$. Let us define $\mathbf{y} = \alpha \mathbf{x}$ for some suitable α . Consider the ratio

$$\begin{aligned} \frac{\|\mathbf{A}_c \mathbf{y}\|_p^r + \lambda \|\mathbf{y}\|_q^s}{\|\mathbf{Ay}\|_p^r + \lambda \|\mathbf{y}\|_q^s} &= \frac{\alpha^r \|\mathbf{A}_c \mathbf{x}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} \\ &= \frac{\alpha^r (1 + \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} \end{aligned}$$

Suppose $r > s$. Then the ratio $\frac{\alpha^r (1 + \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} = \frac{\alpha^{r-s} (1 + \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s}{\alpha^{r-s} \|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s}$. Here we want the ratio to be greater than $(1 + \epsilon)$. By choosing α appropriately, we can set the above ratio to be greater than $(1 + (\frac{\epsilon + \epsilon'}{2})) > 1 + \epsilon$, since $\epsilon' > \epsilon$. A sufficient condition to choose α for this to happen is: $\alpha^{(r-s)} > \frac{(\epsilon' + \epsilon) \lambda \|\mathbf{x}\|_q^s}{(\epsilon' - \epsilon) \|\mathbf{Ax}\|_p^r}$.

Similarly if $r < s$, the ratio $\frac{\alpha^r (1 + \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} = \frac{(1 + \epsilon') \|\mathbf{Ax}\|_p^r + \alpha^{s-r} \lambda \|\mathbf{x}\|_q^s}{\|\mathbf{Ax}\|_p^r + \alpha^{s-r} \lambda \|\mathbf{x}\|_q^s}$. Here for the ratio to be greater than $(1 + \epsilon)$, we can set α such that $\alpha^{(s-r)} < \frac{(\epsilon' - \epsilon) \|\mathbf{Ax}\|_p^r}{(\epsilon' + \epsilon) \lambda \|\mathbf{x}\|_q^s}$.

Case 2: $\exists \mathbf{x} \in \mathbb{R}^d$ s.t. $\|\mathbf{A}_c \mathbf{x}\|_p^r < (1 - \epsilon) \|\mathbf{Ax}\|_p^r$ and $\|\mathbf{A}_c \mathbf{x}\|_p^r = (1 - \epsilon') \|\mathbf{Ax}\|_p^r$ for some $\epsilon' > \epsilon$. Again define $\mathbf{y} = \alpha \mathbf{x}$ for some suitable α . Consider the ratio

$$\begin{aligned} \frac{\|\mathbf{A}_c \mathbf{y}\|_p^r + \lambda \|\mathbf{y}\|_q^s}{\|\mathbf{Ay}\|_p^r + \lambda \|\mathbf{y}\|_q^s} &= \frac{\alpha^r \|\mathbf{A}_c \mathbf{x}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} \\ &= \frac{\alpha^r (1 - \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} \end{aligned}$$

Suppose $r > s$. Here the ratio $\frac{\alpha^r (1 - \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} = \frac{\alpha^{r-s} (1 - \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s}{\alpha^{r-s} \|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s}$. We want the ratio to be smaller than $(1 - \epsilon)$. Without loss of generality, we can set the ratio to be smaller than $(1 - (\frac{\epsilon + \epsilon'}{2}))$ since $\epsilon' > \epsilon$. For this can set α s.t. $\alpha^{(r-s)} > \frac{(\epsilon' + \epsilon) \lambda \|\mathbf{x}\|_q^s}{(\epsilon' - \epsilon) \|\mathbf{Ax}\|_p^r}$.

Similarly if $r < s$, the ratio $\frac{\alpha^r (1 - \epsilon') \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s}{\alpha^r \|\mathbf{Ax}\|_p^r + \lambda \alpha^s \|\mathbf{x}\|_q^s} = \frac{(1 - \epsilon') \|\mathbf{Ax}\|_p^r + \alpha^{s-r} \lambda \|\mathbf{x}\|_q^s}{\|\mathbf{Ax}\|_p^r + \alpha^{s-r} \lambda \|\mathbf{x}\|_q^s}$. Here for the ratio to be smaller than $(1 - \epsilon)$, we can set α such that $\alpha^{(s-r)} < \frac{(\epsilon' - \epsilon) \|\mathbf{Ax}\|_p^r}{(\epsilon' + \epsilon) \lambda \|\mathbf{x}\|_q^s}$. Hence in both the cases, for both scenarios of r and s we can set an α which gives us a contradiction to the fact that \mathbf{A}_c is a coreset for the regularized function. Hence our assumption is wrong and \mathbf{A}_c is also a coreset for the unregularized function. \square

This theorem implies the following corollary which gives our impossibility result.

Corollary 4.1.1. *Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a corresponding vector $\mathbf{b} \in \mathbb{R}^n$, then for the function $\|\mathbf{Ax} - \mathbf{b}\|_p^r + \lambda \|\mathbf{x}\|_q^s$, $r \neq s$, we cannot get a coreset of size smaller than the size of the optimal sized coreset for $\|\mathbf{Ax} - \mathbf{b}\|_p^r$.*

Proof. Consider the case $\mathbf{b} = 0$, which reduces the above function to $\|\mathbf{Ax}\|_p^r + \lambda \|\mathbf{x}\|_q^s$. Theorem 4.1 then implies that no smaller coreset is possible for the regularized problem. A more generalized version of proof is present in the appendix. \square

5. The Modified Lasso

Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank d and $n \gg d$ and a vector $\mathbf{b} \in \mathbb{R}^n$ and $\lambda > 0$, the lasso problem is stated as

$$\begin{aligned} &\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_1 \\ &= \min_{\mathbf{x} \in \mathbb{R}^d} \sum_{i=1}^n (\mathbf{a}_i^T \mathbf{x} - b_i)^2 + \lambda \|\mathbf{x}\|_1 \end{aligned}$$

However as $r = 2$ and $s = 1$, by corollary 4.1.1 we can not hope to get a coreset smaller than the one for least square regression.

To preserve the sparsity inducing nature of the lasso problem and still obtain a smaller sized coreset, we consider a slightly different version of lasso which we call *modified lasso*. It is stated as follows:

$$\begin{aligned} &\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_1^2 \\ &= \min_{\mathbf{x} \in \mathbb{R}^d} \sum_{i=1}^n (\mathbf{a}_i^T \mathbf{x} - b_i)^2 + \lambda \|\mathbf{x}\|_1^2 \end{aligned}$$

Note that in the constraint based formulation (i.e. least square with constraint $\|\mathbf{x}\|_1 \leq R$), the normal lasso and the

modified one are the same with a change in the constraint radius R . In our experiments, we will also empirically show that just like lasso, the modified version also induces sparsity in the solution vector. However, as we will see, the behaviour of the regularized versions of the two problems are different with respect to the optimal coreset size.

We now show a stronger result which leads us to a smaller coreset for the modified lasso problem.

Theorem 5.1. *Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, corresponding vector $\mathbf{b} \in \mathbb{R}^n$, any coreset for the function $\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_p^p$ is also a coreset of the function $\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_q^p$ where $q \leq p, p, q \geq 1$.*

Proof. Let $(\mathbf{A}_c, \mathbf{b}_c)$ be a coreset for the function $\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_p^p$. Hence,

$$\begin{aligned} & \|\mathbf{A}_c \mathbf{x} - \mathbf{b}_c\|_p^p + \lambda \|\mathbf{x}\|_p^p \\ &= \|\mathbf{A}_c \mathbf{x} - \mathbf{b}_c\|_p^p + \lambda \|\mathbf{x}\|_p^p - \lambda \|\mathbf{x}\|_p^p + \lambda \|\mathbf{x}\|_q^p \\ &\leq (1 + \epsilon) (\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_p^p) - \lambda \|\mathbf{x}\|_p^p + \lambda \|\mathbf{x}\|_q^p \\ &= (1 + \epsilon) \|\mathbf{Ax} - \mathbf{b}\|_p^p + \epsilon \lambda \|\mathbf{x}\|_p^p + \lambda \|\mathbf{x}\|_q^p \\ &\leq (1 + \epsilon) \|\mathbf{Ax} - \mathbf{b}\|_p^p + \epsilon \lambda \|\mathbf{x}\|_q^p + \lambda \|\mathbf{x}\|_q^p \\ &= (1 + \epsilon) (\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_q^p) \end{aligned}$$

The second inequality follows from the fact that for any $\mathbf{x} \in \mathbb{R}^d$, $\|\mathbf{x}\|_q \geq \|\mathbf{x}\|_p$ for $p \geq q$. This proves one direction in the definition of coreset. For the other direction consider

$$\begin{aligned} & \|\mathbf{A}_c \mathbf{x} - \mathbf{b}_c\|_p^p + \lambda \|\mathbf{x}\|_p^p \\ &= \|\mathbf{A}_c \mathbf{x} - \mathbf{b}_c\|_p^p + \lambda \|\mathbf{x}\|_p^p - \lambda \|\mathbf{x}\|_p^p + \lambda \|\mathbf{x}\|_q^p \\ &\geq (1 - \epsilon) (\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_p^p) - \lambda \|\mathbf{x}\|_p^p + \lambda \|\mathbf{x}\|_q^p \\ &= (1 - \epsilon) \|\mathbf{Ax} - \mathbf{b}\|_p^p - \epsilon \lambda \|\mathbf{x}\|_p^p + \lambda \|\mathbf{x}\|_q^p \\ &\geq (1 - \epsilon) \|\mathbf{Ax} - \mathbf{b}\|_p^p - \epsilon \lambda \|\mathbf{x}\|_q^p + \lambda \|\mathbf{x}\|_q^p \\ &= (1 - \epsilon) (\|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_q^p) \end{aligned}$$

Combining both inequalities we get the result. \square

As a result of the above Theorem 5.1, we get the following corollary about the modified lasso problem.

Corollary 5.1.1. *For the modified lasso problem there exists an ϵ -coreset of size $O\left(\frac{sd_\lambda(\mathbf{A}) \log sd_\lambda(\mathbf{A})}{\epsilon^2}\right)$ with a high probability where $sd_\lambda(\mathbf{A}) = \sum_{j \in [d]} \frac{1}{1 + \frac{\lambda}{\sigma_j^2}}$ is the statistical dimension of matrix \mathbf{A} which has singular values σ_j 's.*

Proof. For $p = 2$ and $q = 1$ we get the statement of Theorem 5.1 for the ridge regression and the modified lasso problem. By application of the theorem, a coreset for ridge regression is also a coreset for the modified lasso problem for above values of p and q . Using the results of (Avron et al., 2017) we can show that by sampling points according to the ridge leverage scores we can get a coreset of

size $O\left(\frac{sd_\lambda(\mathbf{A}) \log sd_\lambda(\mathbf{A})}{\epsilon^2}\right)$ for ridge regression with constant probability and hence the corollary. \square

It is important to note that the motivation behind the modified lasso proposal is purely computational. It can have smaller coresets and as the constrained version of both problems are similar, we have proposed it for settings where it is desirable to solve the optimization problem on a coreset for scalability reasons. To illustrate, solving either the modified lasso or lasso on a $100K \times 30$ matrix takes roughly 80-90 seconds using the same algorithm. For the modified version a 200×30 coreset can be constructed and problem can be solved in about 1.6 secs with a relative error of 0.0098. Understanding the statistical properties of modified lasso is an interesting open direction.

6. The ℓ_p Regression with ℓ_p Regularization

The ℓ_p Regression with ℓ_p Regularization is given as

$$\begin{aligned} & \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{Ax} - \mathbf{b}\|_p^p + \lambda \|\mathbf{x}\|_p^p \\ &= \min_{\mathbf{x} \in \mathbb{R}^d} \sum_{i=1}^n |(\mathbf{a}_i^T \mathbf{x} - b_i)|^p + \lambda \|\mathbf{x}\|_p^p \end{aligned}$$

For this problem our main result is as follows

Theorem 6.1. *For ℓ_p regression with ℓ_p regularization, there is a coreset of size $O\left(\frac{(\alpha\beta)^p d \log \frac{1}{\epsilon}}{\left(1 + \frac{\lambda}{\|\mathbf{A}'\|_p^p}\right) \epsilon^2}\right)$ with high probability. Here $\mathbf{A}' = [\mathbf{A} \ \mathbf{b}]$ and has an (α, β, p) well-conditioned basis.*

Proof. We define the sensitivity of the i^{th} point for the ℓ_p regression with ℓ_p regularization problem as follows:

$s_i = \sup_{\mathbf{x}'} \frac{|\mathbf{a}_i^T \mathbf{x}'|^p + \frac{\lambda \|\mathbf{x}'\|_p^p}{n}}{\sum_j |\mathbf{a}_j^T \mathbf{x}'|^p + \lambda \|\mathbf{x}'\|_p^p}$. Again, \mathbf{A}' here is the concatenated matrix and \mathbf{x}' is the concatenated vector. Let $\mathbf{A}' = \mathbf{U}\mathbf{V}$, then $\mathbf{a}_j^T \mathbf{x}' = \mathbf{u}_j^T \mathbf{V}$. Here \mathbf{U} is an (α, β, p) well-conditioned basis for \mathbf{A}' . Now let $\mathbf{a}_j^T \mathbf{x}' = \mathbf{u}_j^T \mathbf{V} \mathbf{x}' = \mathbf{u}_j^T \mathbf{z}$. So

$$\begin{aligned} s_i &= \sup_{\mathbf{z}} \frac{|\mathbf{u}_i^T \mathbf{z}|^p + \frac{\lambda}{n} \|\mathbf{V}^{-1} \mathbf{z}\|_p^p}{\sum_j |\mathbf{u}_j^T \mathbf{z}|^p + \lambda \|\mathbf{V}^{-1} \mathbf{z}\|_p^p} \\ &= \sup_{\mathbf{z}} \left(\frac{|\mathbf{u}_i^T \mathbf{z}|^p}{\sum_j |\mathbf{u}_j^T \mathbf{z}|^p + \lambda \|\mathbf{V}^{-1} \mathbf{z}\|_p^p} \right. \\ &\quad \left. + \frac{\frac{\lambda}{n} \|\mathbf{V}^{-1} \mathbf{z}\|_p^p}{\sum_j |\mathbf{u}_j^T \mathbf{z}|^p + \lambda \|\mathbf{V}^{-1} \mathbf{z}\|_p^p} \right) \\ &\leq \sup_{\mathbf{z}} \frac{|\mathbf{u}_i^T \mathbf{z}|^p}{\sum_j |\mathbf{u}_j^T \mathbf{z}|^p + \lambda \|\mathbf{V}^{-1} \mathbf{z}\|_p^p} + \frac{1}{n} \end{aligned}$$

\mathbf{U} is an (α, β, p) well-conditioned basis for \mathbf{A}' . Hence by definition $\|\mathbf{U}\|_p \leq \alpha$ and $\forall \mathbf{z} \in \mathbb{R}^{d+1}$, $\|\mathbf{z}\|_q \leq \beta \|\mathbf{U}\mathbf{z}\|_p$

where $\frac{1}{p} + \frac{1}{q} = 1$. Now the first term $\frac{|\mathbf{u}_j^T \mathbf{z}|^p}{\sum_j |\mathbf{u}_j^T \mathbf{z}|^{p+\lambda} \|\mathbf{V}^{-1} \mathbf{z}\|_p^p}$ lies in the interval $\left[\frac{|\mathbf{u}_j^T \mathbf{z}|^p}{\alpha^p \|\mathbf{z}\|_q^{p+\lambda} \|\mathbf{V}^{-1} \mathbf{z}\|_p^p}, \frac{|\mathbf{u}_j^T \mathbf{z}|^p}{\beta^p \|\mathbf{z}\|_q^{p+\lambda} \|\mathbf{V}^{-1} \mathbf{z}\|_p^p} \right]$. This is by the application of Holder's inequality and the definition of β . Now instead of calculating supremum over $\frac{|\mathbf{u}_j^T \mathbf{z}|^p}{\beta^p \|\mathbf{z}\|_q^{p+\lambda} \|\mathbf{V}^{-1} \mathbf{z}\|_p^p}$, we calculate the infimum over its reciprocal. Lets call it m

$$\begin{aligned} m &= \inf_{\mathbf{z}} \left(\frac{\|\mathbf{z}\|_q^p}{\beta^p |\mathbf{u}_j^T \mathbf{z}|^p} + \frac{\lambda \|\mathbf{V}^{-1} \mathbf{z}\|_p^p}{|\mathbf{u}_j^T \mathbf{z}|^p} \right) \\ &\geq \frac{1}{\beta^p} \inf_{\mathbf{z}} \frac{\|\mathbf{z}\|_q^p}{\|\mathbf{u}_j\|_p^p \|\mathbf{z}\|_q^p} + \lambda \inf_{\mathbf{z}} \frac{\|\mathbf{V}^{-1} \mathbf{z}\|_p^p}{|\mathbf{u}_j^T \mathbf{z}|^p} \\ &\geq \frac{1}{\beta^p \|\mathbf{u}_j\|_p^p} + \lambda \inf_{\mathbf{z}} \frac{\|\mathbf{V}^{-1} \mathbf{z}\|_p^p}{|\mathbf{u}_j^T \mathbf{z}|^p} \end{aligned}$$

Now sensitivity of i^{th} point is bounded as $s_i \leq \frac{1}{m} + \frac{1}{n}$. Therefore $s_i \leq \frac{1}{\frac{1}{\beta^p \|\mathbf{u}_j\|_p^p} + \lambda \inf_{\mathbf{z}} \frac{\|\mathbf{V}^{-1} \mathbf{z}\|_p^p}{|\mathbf{u}_j^T \mathbf{z}|^p}} + \frac{1}{n}$.

Now let us consider

$$m' = \inf_{\mathbf{z}} \frac{\|\mathbf{A}' \mathbf{V}^{-1} \mathbf{z}\|_p^p}{|\mathbf{u}_j^T \mathbf{z}|^p} \geq \inf_{\mathbf{z}} \frac{\|\mathbf{A}' \mathbf{V}^{-1} \mathbf{z}\|_p^p}{\|\mathbf{u}_j\|_p^p \|\mathbf{z}\|_q^p} \geq \frac{1}{\beta^p \|\mathbf{u}_j\|_p^p}$$

Also $\|\mathbf{A}' \mathbf{V}^{-1} \mathbf{z}\|_p \leq \|\mathbf{A}'\|_{(p)} \|\mathbf{V}^{-1} \mathbf{z}\|_p$. Hence $\|\mathbf{V}^{-1} \mathbf{z}\|_p \geq \frac{\|\mathbf{A}' \mathbf{V}^{-1} \mathbf{z}\|_p}{\|\mathbf{A}'\|_{(p)}}$. Here $\|\mathbf{A}'\|_{(p)}$ represents the induced matrix norm defined as $\|\mathbf{A}'\|_{(p)} = \sup_{\mathbf{x}' \neq 0} \frac{\|\mathbf{A}' \mathbf{x}'\|_p}{\|\mathbf{x}'\|_p}$. Combining this with the inequality for m we get

$$m \geq \frac{1}{\beta^p \|\mathbf{u}_j\|_p^p} \left(1 + \frac{\lambda}{\|\mathbf{A}'\|_{(p)}} \right)$$

This results in $s_i \leq \frac{\beta^p \|\mathbf{u}_j\|_p^p}{1 + \frac{\lambda}{\|\mathbf{A}'\|_{(p)}}} + \frac{1}{n}$. So the sum of sensitivities is bounded by $S \leq \frac{(\alpha\beta)^p}{1 + \frac{\lambda}{\|\mathbf{A}'\|_{(p)}}} + 1$. Using the probability distribution based on sensitivity upper bounds defined in the proof to sample rows of \mathbf{A}' and applying Theorem 3.2, we get the result. \square

As the value is decreasing in λ , this coreset is smaller than the coreset obtained for ℓ_p regression using just the well-conditioned basis. This method works with any well-conditioned basis and sampling complexity is dependent on the quality of the well-conditioned basis. For e.g.: for $p > 2$ we can get $\alpha\beta = d^{1/p+1/q}$. So we can get a coreset of size $O\left(\frac{d^{p+1} \log \frac{1}{\epsilon}}{\left(1 + \frac{\lambda}{\|\mathbf{A}'\|_{(p)}}\right) \epsilon^2}\right)$ with high probability. Let us consider the special case of the Regularized Least Absolute Deviation problem.

6.1. Regularized Least Absolute Deviation (RLAD)

The RLAD problem, given as

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_1 + \lambda \|\mathbf{x}\|_1$$

has the benefits of robustness of ℓ_1 regression and sparsity inducing nature of lasso (Wang et al., 2006). Plugging in $p = 1$ in the above analysis we get the following upper bound for sensitivity of the RLAD problem: $s_i \leq \frac{\beta \|\mathbf{u}_i\|_1}{1 + \frac{\lambda}{\|\mathbf{A}'\|_{(1)}}} + \frac{1}{n} = \frac{\beta \|\mathbf{u}_i\|_1}{1 + \frac{\lambda}{\max_{k \in [d]} \|\mathbf{a}^{*k}\|_1}} + \frac{1}{n}$ where \mathbf{a}^{*k} is the k^{th} column of \mathbf{A}' . So the sum of sensitivities is bounded by $S \leq \frac{\alpha\beta}{1 + \frac{\lambda}{\max_{k \in [d]} \|\mathbf{a}^{*k}\|_1}} + 1$. This implies a coreset size

of $O\left(\frac{(\alpha\beta)d \log \frac{1}{\epsilon}}{\left(1 + \frac{\lambda}{\|\mathbf{A}'\|_{(1)}}\right) \epsilon^2}\right)$ for the RLAD problem with high probability by Theorem 3.2. This is smaller than the size of coreset $O\left(\frac{(\alpha\beta)d}{\epsilon^2}\right)$ obtained for ℓ_1 regression using well-conditioned basis. For $p = 1$, we can get a well-conditioned basis with $\alpha\beta = d^{3/2}$. So we can get a coreset of size $O\left(\frac{d^{5/2} \log \frac{1}{\epsilon}}{\left(1 + \frac{\lambda}{\|\mathbf{A}'\|_{(1)}}\right) \epsilon^2}\right)$.

6.2. Coresets for Multiresponse Regularized Regression

Consider the multiresponse RLAD problem given as

$$\min_{\mathbf{X} \in \mathbb{R}^{d \times k}} \|\mathbf{A} \mathbf{X} - \mathbf{B}\|_1 + \lambda \|\mathbf{X}\|_1$$

Here $\mathbf{B} \in \mathbb{R}^{n \times k}$ i.e. there are k different responses and hence our solution is also a matrix \mathbf{X} . We regularize the problem with entry wise 1 norm of the solution matrix. Let us consider the concatenated matrices $\hat{\mathbf{A}} = [\mathbf{A} \ -\mathbf{B}]$ and $\hat{\mathbf{X}} = \begin{bmatrix} \mathbf{X} \\ \mathbf{I}_k \end{bmatrix}$ where \mathbf{I}_k is k -dimensional identity matrix.

Using this matrices we define the sensitivity of the multiresponse RLAD problem as $s_i = \sup_{\hat{\mathbf{X}}} \frac{\|\hat{\mathbf{a}}_i^T \hat{\mathbf{X}}\|_1 + \frac{\lambda \|\hat{\mathbf{X}}\|_1}{n}}{\sum_j \|\hat{\mathbf{a}}_j^T \hat{\mathbf{X}}\|_1 + \lambda \|\hat{\mathbf{X}}\|_1}$. We can get the sensitivity upper bounds for multiple response RLAD as we obtained for the single response RLAD. This gives us the following result

Corollary 6.1.1. *For multiple response RLAD there exists a coreset of size $O\left(\frac{(\alpha\beta)dk \log \frac{1}{\epsilon}}{\left(1 + \frac{\lambda}{\|\hat{\mathbf{A}}\|_{(1)}}\right) \epsilon^2}\right)$ with high probability*

where we have an $(\alpha, \beta, 1)$ well-conditioned basis of $\hat{\mathbf{A}}$.

The proof is given in the appendix and is similar to the one for the single response case. This can be extended to other values of p .

7. Experiments

In this section we describe the empirical results supporting our claims. We performed experiments for the modified

lasso and the RLAD problem. We generated a matrix \mathbf{A} of size 100000×30 in which there a few rows with high leverage scores. The construction of this matrix is described in (Yang et al., 2015) where they refer to it as an NG matrix. The NG (non uniform leverage scores with good condition number) matrix is generated by the following matlab command:

```
NG=[alpha*randn(n-d/2,d/2);
(10^(-8))*rand(n-d/2,d/2);
zeros(d/2,d/2) eye(d/2)];
```

Here we used $\alpha = 0.00065$ to get a condition number of about 5. A solution vector $\mathbf{x} \in \mathbb{R}^{30}$ was generated randomly and a response vector $\mathbf{b} = \mathbf{A}\mathbf{x} + (10^{-5}) \frac{\|\mathbf{b}\|_2}{\|\mathbf{e}\|_2} \mathbf{e}$ where \mathbf{e} is a vector of noise, was also generated. All the experiments were performed in **MatlabR2017a** on a machine with 16GB memory and 8 cores of 3.40 GHz. In our first experiment we solved the modified lasso problem on the entire data matrix \mathbf{A} and response vector \mathbf{b} to see the effect on sparsity of solution for different values of λ . We also solved the lasso and the ridge problem on the same data and compared the number of zeros in the solution vector for different values of λ . To take numerical precision into account we converted each coordinate of the vector having absolute value less than 10^{-6} to 0 and then plotted the number of zeros against λ for each problem. As seen in Figure 1, just like lasso, the modified lasso also induces sparsity in the solution and hence can be used as an alternate to lasso. As expected, no sparsity was induced by ridge regression and hence is not seen on the graph. In the next

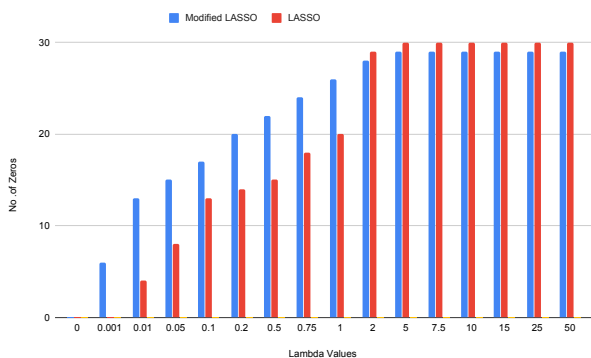


Figure 1. Sparsity Induced by Lasso and Modified Lasso

experiment we compared the performance of a subsample of data sampled using ridge leverage scores with the subsample sampled using uniform sampling. We solved the

Table 1. Relative error of different coreset sizes for Modified Lasso, $\lambda = 0.5$

SAMPLE SIZE	RIDGE LEVERAGE SCORES SAMPLING	UNIFORM SAMPLING
30	0.059	0.8289
50	0.044	0.8289
100	0.031	0.8286
150	0.028	0.8286
200	0.013	0.8287

modified lasso problem on the original data and recorded the objective function value say $V1$. The ridge leverage scores were calculated exactly for the data points. Then we sampled a subsample of points using ridge leverage scores from $[\mathbf{A} \ \mathbf{b}]$. The sampled points were rescaled by the reciprocal of sample size times the probability of sampling the point. We solved the modified lasso on the smaller data and recorded the solution vector. This solution vector was plugged in the modified lasso function with the original large matrix $[\mathbf{A} \ \mathbf{b}]$ and the objective function value was noted as $V2$. We calculated the relative error as $\frac{|V1-V2|}{V1}$. Similar experiment was performed using uniform sampling. In Table 1 we report the values of relative error for different sample sizes. Here $\lambda = 0.5$. The values reported are medians of 5 random experiments for each sample size. It can be seen that even for very small sample sizes, samples obtained using ridge leverage score performed much better than uniform sampling. In fact, uniform sampling showed significant improvement in its own relative error only at much larger sample sizes of the order of 1000's. For e.g. at sample size 1000 error was 0.6816 and at size 2500 error was 0.6554. In the next experiment we fixed a sample size of 200 and solved modified lasso for various values of λ . We report the relative error for both ridge leverage scores based sampling and uniform sampling. We report the medians of 5 random experiments for each value of λ for both schemes of sampling in Table 2. Although both the objective function and the ridge leverage scores depend on the value of λ , still we can observe a decrease in value of relative error for the same sample size with increasing λ except for one case. This effect is more pronounced in case of uniform sampling where the probabilities are independent of λ . In our final experiment we compared the performance of sensitivity based sampling with uniform sampling for the RLAD in similar manner as for modified lasso. Here we used sensitivity upper bounds we calculated for the RLAD problem. The median values of relative error of 5 experiments for each sample size is reported in Table 3. Sensitivity based sampling clearly outperforms uniform sampling in this case.

Table 2. Relative error of different λ values, (sample size = 200) for Modified Lasso

λ	RIDGE LEVERAGE SCORES SAMPLING	UNIFORM SAMPLING
0.1	0.026	2.975
0.5	0.013	0.828
0.75	0.017	0.576
1	0.014	0.443
5	0.007	0.103

Table 3. Relative error of different coreset sizes for RLAD, $\lambda = 0.5$

SAMPLE SIZE	SENSITIVITY BASED SAMPLING	UNIFORM SAMPLING
30	0.69	385.99
50	0.65	112.70
100	0.34	98.53
150	0.19	96.09
200	0.17	27.49

7.1. Experiments on Real Data

We used the *Combined Cycle Power Plant Data Set* (Tüfekci, 2014) available at the UCI Machine learning repository. The data set has 9567 data points and 4 features namely *Temperture(T)*, *Ambient Pressure (AP)*, *Relative Humidity (RH)* and *Exhaust Vacuum (V)*. The task is to predict the *Net hourly electrical energy output (EP)*. We used the modified lasso regression model to fit this data. We normalized each feature by its maximum value so that each feature lies in the same range. We solved the modified lasso problem on the entire data for $\lambda = 1$. We also applied our ridge leverage score based sampling and uniform sampling over the data and created coresets of different sizes. We solved modified lasso problem on the coresets and used the parameter vector obtained, with the original data and compared the errors. We did 5 experiments for each coreset size and obtained the relative error values. The median of relative errors for each coreset size is recorded in Table 4. As can be seen, ridge leverage score sampling performs slightly better than uniform sampling. It is expected, as real data sets tend to have uniform leverage scores. However for data with nonuniform leverage scores our method is useful in practice too.

To the best of our knowledge, this is the first set of experiments evaluating performance of coresets for any regularized form of regression. These clearly verify that coresets constructed using importance scores can give performance comparable to original data at much smaller sizes.

Table 4. Relative error of different coreset sizes for Modified Lasso on Real Data, $\lambda = 1$

SAMPLE SIZE	UNIFORM SAMPLING	RIDGE LEVERAGE SCORES SAMPLING
50	0.0280	0.0267
100	0.0184	0.0161
150	0.0119	0.0082
300	0.0060	0.0048
500	0.0042	0.0028

8. Conclusion and Discussion

In this paper we have studied the coresets for regularized regression problem. We have shown that coresets smaller than those for unregularized regression might not be possible for all forms of regularized regression. We have introduced the modified lasso and shown a smaller coreset for it. We have also shown smaller coresets for ℓ_p regression with ℓ_p regularization by upper bounding the sensitivity scores. Finally we have shown empirical results to support our theoretical claims.

One obvious open question is to see if tighter upper bounds are possible for the sensitivity scores for these problems. A detailed empirical study of coresets for regularized versions of regression under different settings is also required as has been done for ℓ_2 and ℓ_1 regression problems. Coresets for regression with other type of regularization or coresets for other regularized problems is an interesting area for future work. Also study of statistical properties of modified lasso is an interesting question.

Acknowledgements

We would like to thank Jayesh Choudhari and Kshiteej Sheth for helpful discussions. We are also grateful to the anonymous reviewers for their helpful feedback. Anirban Dasgupta acknowledges the kind support of the N. Rama Rao Chair Professorship at IIT Gandhinagar, the Google India AI/ML award (2020), Google Faculty Award (2015), and CISCO University Research Grant (2016). Supratim Shit acknowledges the kind support of Additional Fellowship from IIT Gandhinagar.

References

Agarwal, P. K., Har-Peled, S., and Varadarajan, K. R. Approximating extent measures of points. *Journal of the ACM (JACM)*, 51(4):606–635, 2004.

Avron, H., Clarkson, K. L., and Woodruff, D. P. Sharper bounds for regularized data fitting. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM*

- 2017). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- Bachem, O., Lucic, M., and Krause, A. Practical coreset constructions for machine learning. *arXiv preprint arXiv:1703.06476*, 2017.
- Batson, J., Spielman, D. A., and Srivastava, N. Twice-ramanujan sparsifiers. *SIAM Journal on Computing*, 41(6):1704–1721, 2012.
- Braverman, V., Feldman, D., and Lang, H. New frameworks for offline and streaming coreset constructions. *arXiv preprint arXiv:1612.00889*, 2016.
- Clarkson, K. L. and Woodruff, D. P. Low-rank approximation and regression in input sparsity time. *Journal of the ACM (JACM)*, 63(6):54, 2017.
- Clarkson, K. L., Drineas, P., Magdon-Ismail, M., Mahoney, M. W., Meng, X., and Woodruff, D. P. The fast cauchy transform and faster robust linear regression. *SIAM Journal on Computing*, 45(3):763–810, 2016.
- Cohen, M. B. and Peng, R. ℓ_p row sampling by lewis weights. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 183–192. ACM, 2015.
- Curtin, R., Im, S., Moseley, B., Pruhs, K., and Samadian, A. On coresets for regularized loss minimization. *arXiv preprint arXiv:1905.10845*, 2019.
- Dasgupta, A., Drineas, P., Harb, B., Kumar, R., and Mahoney, M. W. Sampling algorithms and coresets for ℓ_p regression. *SIAM Journal on Computing*, 38(5):2060–2078, 2009.
- Dickens, C., Cormode, G., and Woodruff, D. Leveraging well-conditioned bases: Streaming and distributed summaries in minkowski p -norms. In *International Conference on Machine Learning*, pp. 1243–1251, 2018.
- Drineas, P., Mahoney, M. W., and Muthukrishnan, S. Sampling algorithms for ℓ_2 regression and applications. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 1127–1136. Society for Industrial and Applied Mathematics, 2006.
- Drineas, P., Mahoney, M. W., Muthukrishnan, S., and Sarlós, T. Faster least squares approximation. *Numerische mathematik*, 117(2):219–249, 2011.
- Drineas, P., Magdon-Ismail, M., Mahoney, M. W., and Woodruff, D. P. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research*, 13(Dec):3475–3506, 2012.
- Dubhashi, D. P. and Panconesi, A. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- Feldman, D. and Langberg, M. A unified framework for approximating and clustering data. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 569–578. ACM, 2011.
- Har-Peled, S. and Mazumdar, S. On coresets for k-means and k-median clustering. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pp. 291–300. ACM, 2004.
- Hausler, D. Sphere packing numbers for subsets of the boolean n -cube with bounded vapnik-chervonenkis dimension. *J. Comb. Theory, Ser. A*, 69(2):217–232, 1995.
- Kacham, P. and Woodruff, D. P. Optimal deterministic coresets for ridge regression. *AISTAT*, 2020.
- Langberg, M. and Schulman, L. J. Universal ε -approximators for integrals. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pp. 598–607. SIAM, 2010.
- Maalouf, A., Jubran, I., and Feldman, D. Fast and accurate least-mean-squares solvers. In *Advances in Neural Information Processing Systems*, pp. 8305–8316, 2019.
- Meng, X. and Mahoney, M. W. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 91–100. ACM, 2013.
- Phillips, J. M. Coresets and sketches. *arXiv preprint arXiv:1601.00617*, 2016.
- Pilanci, M. and Wainwright, M. J. Randomized sketches of convex programs with sharp guarantees. *IEEE Transactions on Information Theory*, 61(9):5096–5115, 2015.
- Raj, A., Musco, C., and Mackey, L. Importance sampling via local sensitivity. *arXiv preprint arXiv:1911.01575*, 2019.
- Reddi, S. J., Póczos, B., and Smola, A. J. Communication efficient coresets for empirical loss minimization. In *UAI*, pp. 752–761, 2015.
- Sohler, C. and Woodruff, D. P. Subspace embeddings for the ℓ_1 -norm with applications. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 755–764. ACM, 2011.
- Tolochinsky, E. and Feldman, D. Coresets for monotonic functions with applications to deep learning. *CoRR*, abs/1802.07382, 2018.

- Tüfekci, P. Prediction of full load electrical power output of a base load operated combined cycle power plant using machine learning methods. *International Journal of Electrical Power & Energy Systems*, 60:126–140, 2014.
- Wang, L., Gordon, M. D., and Zhu, J. Regularized least absolute deviations regression and an efficient algorithm for parameter tuning. In *Sixth International Conference on Data Mining (ICDM'06)*, pp. 690–700. IEEE, 2006.
- Woodruff, D. and Zhang, Q. Subspace embeddings and ℓ_p -regression using exponential random variables. In *Conference on Learning Theory*, pp. 546–567, 2013.
- Woodruff, D. P. et al. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- Yang, J., Meng, X., and Mahoney, M. W. Implementing randomized matrix algorithms in parallel and distributed environments. *Proceedings of the IEEE*, 104(1):58–92, 2015.