

UC Irvine

UC Irvine Previously Published Works

Title

Convex Clustering via Optimal Mass Transport

Permalink

<https://escholarship.org/uc/item/39v28745>

Authors

Carli, Francesca P

Ning, Lipeng

Georgiou, Tryphon T

Publication Date

2013-07-20

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

CONVEX CLUSTERING VIA OPTIMAL MASS TRANSPORT

FRANCESCA P. CARLI, LIPENG NING, AND TRYPHON T. GEORGIU *

Abstract. We consider approximating distributions within the framework of optimal mass transport and specialize to the problem of clustering data sets. Distances between distributions are measured in the Wasserstein metric. The main problem we consider is that of approximating sample distributions by ones with sparse support. This provides a new viewpoint to clustering. We propose different relaxations of a cardinality function which penalizes the size of the support set. We establish that a certain relaxation provides the tightest convex lower approximation to the cardinality penalty. We compare the performance of alternative relaxations on a numerical study on clustering.

1. Introduction. The analysis of data sets invariably requires approximating observed sample distributions by ones that belong to a particular family. Instances include modeling using sums of Gaussians. While such families can in principle be quite general, the metric used to quantify mismatch ought to reflect appropriate features. A natural geometry is that provided by optimal mass transport endowing the space of distribution with, for instance, the Wasserstein metric d_{W_2} . Throughout we will use this metric to study approximation problems. We specialize to a family \mathcal{S} of distributions with sparse support on discrete spaces with application to clustering.

In more detail, we consider the problem of approximating a given sample distribution p_0 with a distribution p_1 belonging to a class \mathcal{S} by solving the problem

$$\min_{p_1 \in \mathcal{S} \subset \mathcal{P}} d_{W_2}(p_0, p_1). \quad (1.1)$$

Throughout, our spaces are discrete and \mathcal{P} denotes the probability simplex. Membership in \mathcal{S} can typically be relaxed by introducing a suitable penalty function

$$\min_{p_1 \in \mathcal{P}} d_{W_2}(p_0, p_1) + \lambda \mathcal{I}(p_1) \quad (1.2)$$

where the function \mathcal{I} penalizes $p_1 \notin \mathcal{S}$ and $\lambda > 0$. In this paper, we specialize to the case where the set \mathcal{S} represents the family of distributions having a sparse support on a discrete space, which leads to a cardinality penalized optimization problem. The sparse atoms of the support correspond to the representatives of different clusters. Association to those clusters' representatives is provided by the solution to the aforementioned optimal mass transportation problem and is dictated by the optimal transportation plan (see below for details).

A heuristic for cardinality that has attracted a lot of attention in recent years is ℓ_1 -norm regularization [19, 7, 3, 2, 4]. However, this cannot be used to promote sparsity on a probability simplex since the ℓ_1 norm of a probability measure is always one (see e.g., the recent paper by Pilanci et al. [16]). In the present paper, we propose relaxations of the cardinality penalty on the probability simplex. The main idea is to express the cardinality penalty in terms of the optimal transportation plan Π between the sample distribution and a target distribution with required sparsity properties. The transportation plan Π is itself a probability distribution on a larger space. We show that *convex* relaxations of the cardinality penalty can be realized via a sum-of-norms penalty on the transportation plan matrix and via the introduction of a suitable indicator function. This leads to a *convex* optimization scheme to solve the

*Department of Electrical & Computer Eng., University of Minnesota, Minneapolis, MN 55455; Email: {fpcarli, ningx015, tryphon}@umn.edu; Supported in part by the NSF and the AFOSR.

clustering problem. We show that the sum-of-norms relaxation provides the *tightest* convex lower approximation of the original cardinality penalty. Finally, we present numerical examples that underscore the effectiveness and relevance of the proposed approaches with regard to the problem of clustering data sets.

The paper is organized as follows. In Section 2 we introduce the optimal mass transportation problem. Optimal transport with a cardinality penalty is introduced in Section 3 and different relaxations are proposed. In particular a relaxation based on a sum-of-norms penalty is discussed in Section 3.1 while a relaxation based on the introduction of auxiliary boolean variables is discussed in Section 3.2. In Section 3.3 we adapt the approach in [16] to our setting and compare with the proposed techniques. In Section 4 we specialize to the clustering problem of and explain how it fits into the more general framework introduced so far. We conclude by comparing the effectiveness of the different relaxations on numerical experiments on clustering in Section 5.

2. Mass transport and the Wasserstein metric. In this section, we briefly introduce the problem of optimal mass transport. We refer the reader to [20, 17, 18, 21] for a survey of the subject. The original formulation of the problem goes back to G. Monge in 1781 [15], while the modern formulation is due to L. Kantorovich in 1942 [12]. Below we present the Monge-Kantorovich *optimal mass transportation (OMT)* problem restricting our discussion to (finite) discrete spaces.

Optimal mass transport. Let \mathcal{X} and \mathcal{Y} be finite discrete spaces and p_X and p_Y be probability measures on \mathcal{X} and \mathcal{Y} , respectively. Let $c : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R} \cup \{+\infty\}$ be a cost function, i.e. $c(x_i, y_j) \geq 0$ represents the transportation cost of transferring one unit of mass from $x_i \in \mathcal{X}$ to $y_j \in \mathcal{Y}$. Let $\pi : \mathcal{X} \times \mathcal{Y} \mapsto [0, 1]$ be a transference plan (informally, $\pi(x_i, y_j)$ measures the amount of mass transferred from location x_i to location y_j). The OMT problem is to minimize the total transportation cost over the set of (joint) probability measures π with given marginals p_X and p_Y and it is as follows:

$$\begin{aligned} \min_{\pi} \quad & \sum_{i,j} c(x_i, y_j) \pi(x_i, y_j) \\ \text{subject to} \quad & \sum_j \pi(x_i, y_j) = p_X(x_i), \quad \forall i \\ & \sum_i \pi(x_i, y_j) = p_Y(y_j), \quad \forall j \\ & \pi(x_i, y_j) \geq 0, \quad \forall i, j. \end{aligned}$$

The following probabilistic interpretation is standard. If X, Y are random variables taking values on \mathcal{X} and \mathcal{Y} with probability distributions p_X and p_Y , respectively, the OMT problem is to minimize the expectation $\mathbb{E}_{\pi} [c(X, Y)]$ over all admissible joint distributions π of (X, Y) .

We denote by Π the matrix associated to the transference plan π and by C the matrix associated to the transference cost, i.e.,

$$\Pi = [\pi(x_i, y_j)], \quad C = [c(x_i, y_j)],$$

for $i = 1, \dots, |\mathcal{X}|$, $j = 1, \dots, |\mathcal{Y}|$, with $|\mathcal{Z}|$ denoting the cardinality of the set \mathcal{Z} . The OMT problem can be expressed in matrix notation as follows:

$$\min_{\Pi \in \mathcal{M}(p_X, p_Y)} \text{trace}(C^T \Pi) \tag{2.1}$$

where

$$\mathcal{M}(p_X, p_Y) := \{\Pi \mid \Pi \mathbf{1} = p_X, \Pi^T \mathbf{1} = p_Y, \Pi \geq 0\}$$

and $\mathbf{1}$ is a vector of ones of suitable dimension.

Wasserstein metric. Consider the case where \mathcal{Z} is a metric space with metric $d(\cdot, \cdot)$, and $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{Z}$, and consider the problem of optimal transport between two probability measures as before. When the transportation cost $c(\cdot, \cdot)$ is equal to $d(\cdot, \cdot)^q$ with $q > 0$, the OMT induces a metric on the space of probability measures having finite q th-moments [20, Chapter 7]. Herein, we are interested in the case where $q = 2$. More specifically, we view \mathcal{X}, \mathcal{Y} as sets of points in a Euclidean space. The Euclidean metric induces a metric on $\mathcal{Z} := \mathcal{X} \cup \mathcal{Y}$ and thereby a cost $c(\cdot, \cdot)$ so that

$$C_{i,j} = \|z_i - z_j\|^2.$$

The optimal transport cost

$$d_{W_2}(p_1, p_2) := \min_{\Pi \in \mathcal{M}(p_1, p_2)} \text{trace}(C^T \Pi), \quad (2.2)$$

where Π is the joint probability on $\mathcal{Z} \times \mathcal{Z}$ as before, gives rise to the 2-*Wasserstein metric* between p_1 and p_2 :

$$\sqrt{d_{W_2}(p_1, p_2)}.$$

For the case where $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{Z}$ are not necessarily equal and for $p_1 = p_X$ having support on \mathcal{X} and $p_2 = p_Y$ having support on \mathcal{Y} , Π will have support on $\mathcal{X} \times \mathcal{Y}$. Therefore the optimization in (2.2) can be carried out with Π restricted to be a probability distribution on $\mathcal{X} \times \mathcal{Y}$ and C restricted to correspond to distances between points in these two spaces, i.e., $C_{i,j} = \|x_i - y_j\|^2$.

3. The cardinality penalty & relaxations. We now return to considering the optimization problem in (1.2) specializing $\mathcal{I}(\cdot)$ to be the cardinality function $\text{card}(\cdot)$ giving the number of nonzero entries of the argument. We propose alternative convex relaxations that are suitable to the case where the optimization variable is a probability vector.

We first rewrite problem (1.2), namely,

$$\min_{\substack{p_1 \geq 0 \\ \mathbf{1}^T p_1 = 1}} d_{W_2}(p_0, p_1) + \lambda \text{card}(p_1), \quad (3.1)$$

in terms of the transportation plan Π . Indeed, from (2.2) and since

$$p_1 = \Pi^T \mathbf{1},$$

problem (3.1) can be rewritten as

$$\min_{\substack{\Pi \geq 0 \\ \Pi \mathbf{1} = p_0}} \text{trace}(C^T \Pi) + \lambda \text{card}(\Pi^T \mathbf{1}). \quad (3.2)$$

The optimal value will be denoted by J^{opt} .

Cardinality penalized problems are in general NP-hard to solve, being combinatorial in nature. The cardinality function is nonconvex and is usually replaced by the ℓ_1 -norm which is a convex surrogate [1, Chapter 6]. Evidently, such a relaxation is not applicable here since we are dealing with probability vectors. Below, in 3.1 and 3.2, we propose two alternative relaxations that are applicable to our setting and, in 3.3, we discuss an additional relaxation which has recently been proposed in [16].

3.1. Rank regularization and sum-of-norms relaxation. The main idea underlying the relaxation proposed below is to express the cardinality penalty in terms of the rank of a certain linear map of the transportation plan Π . To this end, we denote by $e_i \in \mathbb{R}^N$ the standard unit N -vector with 1 in the i th entry and with E_i the single-entry diagonal matrix with a 1 in position (i, i) . Moreover, we denote by Π_i the i th column of Π . We introduce the map F that associates to each transportation plan Π the rectangular block-diagonal matrix with diagonal block-entries the columns of Π ,

$$F(\Pi) := \sum_{i=1}^N E_i \otimes (\Pi e_i) = \begin{bmatrix} \Pi_1 & 0 & \dots & 0 \\ 0 & \Pi_2 & \dots & 0 \\ \vdots & \vdots & & \\ 0 & 0 & \dots & \Pi_N \end{bmatrix}, \quad (3.3)$$

where \otimes stands for the Kronecker product. Whenever it is clear from the context we simplify the notation and denote $F(\Pi)$ simply by F . The matrix $F^T F$ is diagonal with i th diagonal entry the scalar product $\langle \Pi_i, \Pi_i \rangle$. Since $\langle \Pi_i, \Pi_i \rangle = 0$ if and only if $\Pi_i = 0$, we have

$$\begin{aligned} \text{card}(\Pi^T \mathbf{1}) &= \text{rank}(F^T F) \\ &= \text{rank}(F) \end{aligned}$$

and (3.2) now becomes

$$\min_{\substack{\Pi \geq 0 \\ \Pi \mathbf{1} = p_0}} \text{trace}(C^T \Pi) + \lambda \text{rank}(F(\Pi)). \quad (3.4)$$

Denote by $\|F\|_2$ the spectral norm of F and by $\|F\|_*$ its nuclear norm. Due to the particular structure of F , $F^T F$ is diagonal and the cardinality of $\text{diag}(F^T F)$ coincides with the rank of F . Then, $\|F\|_*$, which is defined as the sum of the singular values, is simply

$$\begin{aligned} \|F\|_* &= \|(\text{diag}(F^T F))^{\frac{1}{2}}\|_1 \\ &= \sum_{i=1}^N \|\Pi_i\|_2. \end{aligned}$$

Utilizing the well-known fact that the nuclear norm represents the convex envelope of the rank function on a (bounded set) of matrices (Fazel *et al.* [10]), we obtain the following result.

PROPOSITION 3.1. The convex envelope of $\text{rank}(F(\Pi))$ on the set $\{\Pi \mid \Pi \mathbf{1} = p_0, \Pi \geq 0\}$ is

$$\phi(\Pi) = \frac{1}{\|p_0\|_2} \sum_{i=1}^N \|\Pi_i\|_2.$$

Proof: Since by [10, Theorem 1], it holds that $\|M\|_*/a$ is the *convex envelope* of $\text{rank}(M)$ over the set $\{M \mid \|M\|_2 \leq a\}$, the only thing that remains to prove is that $\|F\|_2$ is bounded by $\|p_0\|_2$ on the feasible set. Indeed, for every

$\Pi \in \{\Pi \mid \Pi \mathbf{1} = p_0, \Pi \geq 0\}$, $F(\Pi)$ given by (3.3), it holds that

$$\begin{aligned} \|F\|_2 &= \max_i \sqrt{\text{eig}(F^T F)} \\ &= \max_i \text{eig} \left\{ \begin{bmatrix} \|\Pi_1\|_2 & 0 & \dots & 0 \\ 0 & \|\Pi_2\|_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \|\Pi_N\|_2 \end{bmatrix} \right\} \\ &= \max_i \{\|\Pi_i\|_2\} \leq \left\| \sum_i \Pi_i \right\|_2 = \|p_0\|_2, \end{aligned}$$

as claimed. \square

It follows that the rank-penalized problem (3.4) can be relaxed into

$$\min_{\substack{\Pi \geq 0 \\ \Pi \mathbf{1} = p_0}} \text{trace}(C^T \Pi) + \frac{\lambda}{\|p_0\|_2} \sum_{i=1}^N \|\Pi_i\|_2. \quad (3.5)$$

where the sum-of-norms penalty is a Group-Lasso-type penalty [22] with the groups given by the columns of the transference plan Π .

3.2. Integer programming and fractional relaxation. In this section, an alternative relaxation of Problem (3.2) is presented. To this aim, we introduce the indicator function

$$y_i = \begin{cases} 1, & \text{if } \Pi_i \neq 0 \\ 0, & \text{otherwise} \end{cases}$$

whose entries reflect the sparsity pattern of the columns of Π . Then, clearly,

$$\text{card}(\Pi^T \mathbf{1}) = \|y\|_1,$$

Moreover, since the the columns of Π must sum up to p_0 , the following inequality holds

$$\Pi \leq p_0 y^T.$$

This leads to the following equivalent formulation of (3.2)

$$\begin{aligned} &\underset{\Pi, y}{\text{minimize}} && \text{trace}(C^T \Pi) + \lambda \|y\|_1 && (3.6) \\ &\text{subject to} && \Pi \mathbf{1} = p_0, \\ &&& \Pi \geq 0, \\ &&& \Pi \leq p_0 y^T, \\ &&& y_i \in \{0, 1\}. \end{aligned}$$

Problem (3.6) is not convex due to the integer constraints $y_i \in \{0, 1\}$. A standard relaxation of an integer program is

$$\begin{aligned} &\underset{\Pi, y}{\text{minimize}} && \text{trace}(C^T \Pi) + \lambda \|y\|_1 && (3.7) \\ &\text{subject to} && \Pi \mathbf{1} = p_0, \\ &&& \Pi \geq 0, \\ &&& \Pi \leq p_0 y^T, \\ &&& y_i \in [0, 1], \end{aligned}$$

where $y_i \in \{0, 1\}$ is relaxed into $y_i \in [0, 1]$. This formulation is similar to the so called *facility location* problem in Operations Research [8, 9], where typically the matrix Π is Boolean ($(0, 1)$ -entries) and p_0 has integer entries, and a number of works has been devoted to the design of specialized LP-rounding algorithm for reconstructing integer solutions starting from a solution of a relaxed problem (see e.g. [14, 5, 6, 11]).

In section 5, the relaxation (3.7) will be compared with the sum-of-norms relaxation (3.5) on a clustering application.

3.3. Relaxation inverse of the ℓ_∞ -norm . The relaxation discussed below is a special case of a problem that is addressed by Pilanci *et al.* in [16] where they consider sparse minimizers of general convex functions on probability simplices. The key idea in [16] is to utilize the inverse of the ℓ_∞ -norm of the probability vector (here, p_1) as a surrogate for the cardinality. Indeed, in general,

$$\|p_1\|_1 \leq \text{card}(p_1) \|p_1\|_\infty,$$

and since p_1 is a probability vector,

$$\frac{1}{\|p_1\|_\infty} \leq \text{card}(p_1).$$

Problem (3.2) can thus be relaxed to

$$\min_{\substack{\Pi \geq 0 \\ \Pi \mathbf{1} = p_0}} \text{trace}(C^T \Pi) + \frac{\lambda}{\|\Pi^T \mathbf{1}\|_\infty}. \quad (3.8)$$

Note that (3.8) is still not a convex problem. Nevertheless, it can be solved exactly by using the following N convex programs (see [16, Proposition 2.1])

$$\min_{i=1, \dots, N} \left\{ \min_{\substack{\Pi \geq 0 \\ \Pi \mathbf{1} = p_0}} \text{trace}(C^T \Pi) + \frac{\lambda}{[\Pi^T \mathbf{1}]_i} \right\} \quad (3.9)$$

where $[\Pi^T \mathbf{1}]_i$ denotes the i th component of the vector $\Pi^T \mathbf{1}$.

In [16], this relaxation has been applied to estimate the (sparse) coefficients vector of a Gaussian mixture in the exemplar based convex clustering framework of [13].

The relaxation based on the ℓ_∞ -norm will be compared with the sum-of-norm relaxation (3.5) and the relaxation based on the integer programming formulation (3.7) in Section 5.

4. Clustering via OMT. In this section, we describe how ideas from optimal mass transport can be applied for clustering points in a metric space. Consider a data set $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset \mathbb{R}^d$. A common way to address the problem is to select a subset of cluster representatives and associate points to these representatives in such a way to minimize a given cost functional, which is usually the sum of the square distances between points and associated cluster representatives. Here we observe that the problem of optimally selecting a subset of data points as cluster centers, can be seen as an optimal mass transportation problem where a sample distribution p_0 is associated to the data points in \mathcal{X} and a distribution p_1 is associated to the cluster centers. Cluster centers are chosen in such a way to minimize the optimal transportation cost between p_0 and p_1 . In this, p_1 is to be determined based on the requirements that its support consists of a few points (few cluster representatives).

This leads to a *convex* clustering scheme where the optimal transportation plan Π^{opt} can be computed by solving one of the proposed relaxations of (3.2) and clustering is achieved according to the following rules:

- Choice of cluster representatives: the point \mathbf{x}_j is a representative of a cluster if \exists a point \mathbf{x}_i such that $\Pi_{ij}^{\text{opt}} > \Pi_{i\ell}^{\text{opt}}$ for all $\ell \neq j$,
- Association of points to cluster representatives: assign the point \mathbf{x}_i to the cluster representative \mathbf{x}_j if $\Pi_{ij}^{\text{opt}} > \Pi_{i\ell}^{\text{opt}}$ for all $\ell \neq j$.

5. Experimental results. Following the scheme of the previous section, we now compare the relaxations introduced so far on a clustering example. In particular, we consider synthetic data in \mathbb{R}^2 sampled from 4 Gaussian distributions with means

$$\left\{ (0, 5), \left(-5\frac{\sqrt{3}}{2}, -\frac{5}{2}\right), \left(5\frac{\sqrt{3}}{2}, -\frac{5}{2}\right), (8, 2) \right\}$$

respectively, and common covariance given by

$$\begin{bmatrix} 0.8 & 0 \\ 0 & 0.8 \end{bmatrix}.$$

The clustering obtained by solving (3.5), (3.7) and (3.9) for different values of λ are shown in Figures 6.1, 6.2 and 6.3, respectively. The clusters representatives are denoted by a black cross while data points belonging to different clusters are denoted by different shapes and colors.

From the experimental results, we see that both the relaxation based on the sum-of-norms and the one based on the integer programming formulation are able to achieve the correct clustering. In particular, as it can be seen from panels (c) and (d) in Figures 6.1 and 6.2, for values of λ approximately between 20 and 100, both methods correctly partition the data set into 4 clusters. For $\lambda = 220$ both methods partition the data set into 3 clusters. By further increasing the value of the parameter λ , data are finally “grouped” in a unique cluster, which happens approximately for $\lambda = 750$.

The clustering produced by optimizing (3.9) is shown in Figure 6.3. When applied to our particular problem, the penalty based on the ℓ_∞ -norm is less effective than the proposed relaxations in promoting sparsity of the clusters representatives distribution. Clustering results for different values of the parameter λ are shown in Figure 6.3. For $\lambda = 100$, the data points are “grouped” in a unique cluster. By decreasing the value of λ the method is able to isolate the biggest cluster, while all the other points remain unclustered (each point is chosen as the representative of itself). There are no values of λ for which the penalty based on the ℓ_∞ -norm is able to correctly partition the data set in four clusters.

The three relaxations have also been applied to a data set with 10 clusters generated by sampling 10 Gaussian distributions with means

$$\begin{aligned} & \{ (-2.5, -12.5), (5, -10), (0, -5), (-4.5, -5), (-5, 0), \\ & (-6, 5), (-1.5, 2.5), (3.5, -1), (7.5, -2.5), (10, 2.5) \} \end{aligned}$$

and common covariance

$$\begin{bmatrix} 0.2 & 0 \\ 0 & 0.2 \end{bmatrix}.$$

The experimental results are reported in Figure 6.4, 6.5 and 6.6, respectively. Once again, the sum-of-norms relaxation and the relaxation based on integer programming are able to achieve the correct clustering, while there are no values of λ for which the penalty based on the ℓ_∞ -norm is able to correctly partition the data set in ten clusters.

6. Conclusions. In this paper, we considered the problem of approximating distributions within the framework of optimal mass transport. We focused on approximating sample distributions with distributions having sparse support. Standard ℓ_1 regularization cannot be used to promote sparsity on a probability simplex since the ℓ_1 norm of a probability measure is always one. We proposed relaxations of the cardinality penalty which are applicable to probability simplices. One of these relaxation has the property to be the tightest convex lower approximation of the original cardinality penalty. When applied to a clustering problem, the proposed framework leads to *convex* clustering schemes, thus overcoming sensitivity to initialization of classical clustering algorithms such as k-means. The proposed relaxations have been tested on a clustering problem with synthetic generated data. Both the relaxation based on the sum-of-norms and on the integer programming formulation act effectively in promoting sparsity of the clusters representatives distribution and are able to produce the correct clustering.

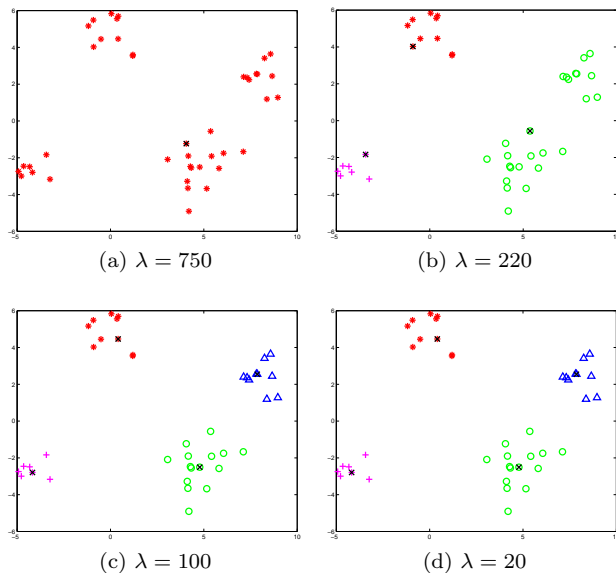


Fig. 6.1: Output clustering via the solution of the sum-of-norms relaxation (3.5) for different values of the parameter λ . The clusters representatives are denoted by a black cross. Different shapes and colors have been used to denote points belonging to different clusters.

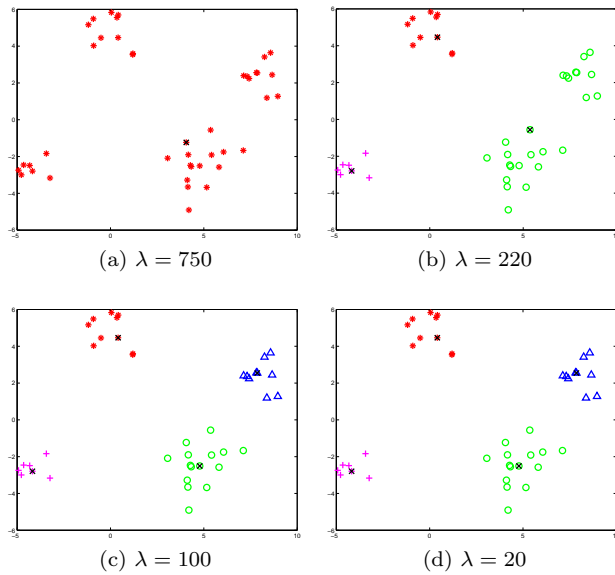


Fig. 6.2: Output clustering via the solution of the LP formulation (3.7) for different values of the parameter λ . The clusters representatives are denoted by a black cross. Different shapes and colors have been used to denote points belonging to different clusters.

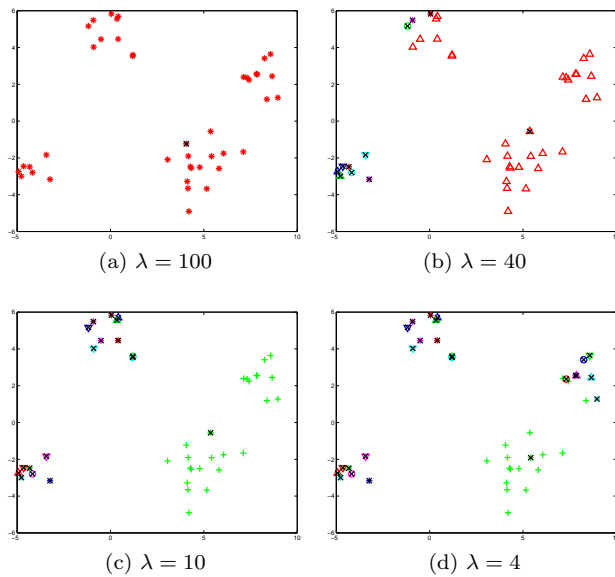


Fig. 6.3: Output clustering via the solution of the ℓ_∞ -norm relaxation (3.9) for different values of the parameter λ . The clusters representatives are denoted by a black cross. Different shapes and colors have been used to denote points belonging to different clusters.

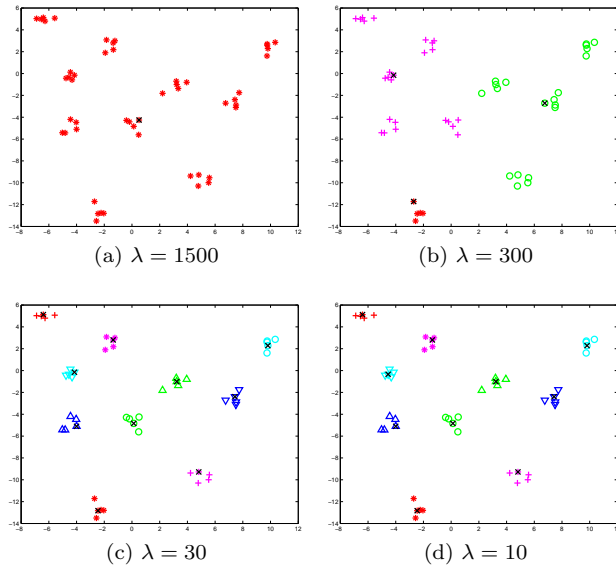


Fig. 6.4: Output clustering via the solution of the sum-of-norms relaxation (3.5) for different values of the parameter λ . The clusters representatives are denoted by a black cross. Different shapes and colors have been used to denote points belonging to different clusters.

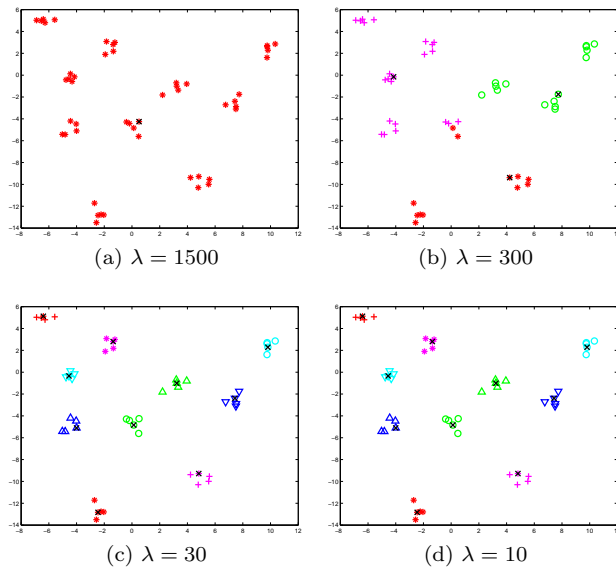


Fig. 6.5: Output clustering via the solution of the LP formulation (3.7) for different values of the parameter λ . The clusters representatives are denoted by a black cross. Different shapes and colors have been used to denote points belonging to different clusters.

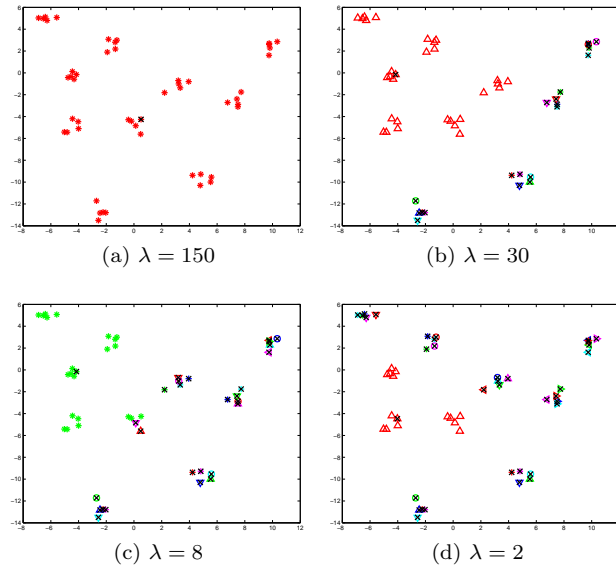


Fig. 6.6: Output clustering via the solution of the ℓ_∞ -norm relaxation (3.9) for different values of the parameter λ . The clusters representatives are denoted by a black cross. Different shapes and colors have been used to denote points belonging to different clusters.

REFERENCES

- [1] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [2] A.M. Bruckstein, D.L. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM review*, 51(1):34–81, 2009.
- [3] E.J. Candes and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.
- [4] V. Chandrasekaran, B. Recht, P.A. Parrilo, and A.S. Willsky. The convex algebraic geometry of linear inverse problems. In *48th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2010*, pages 699–703, 2010.
- [5] M. Charikar, S. Guha, É. Tardos, and D.B. Shmoys. A constant-factor approximation algorithm for the k-median problem. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 1–10. ACM, 1999.
- [6] M. Charikar, S. Guha, É. Tardos, and D.B. Shmoys. A constant-factor approximation algorithm for the k-median problem. *Journal of Computer and System Sciences*, 65(1):129–149, 2002.
- [7] S.S. Chen, D.L. Donoho, and M.A. Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1998.
- [8] M.S. Daskin. *Network and discrete location: models, algorithms, and applications*. Wiley, 1995.
- [9] Z. Drezner and H.W. Hamacher. *Facility location: applications and theory*. Springer-Verlag, 2001.
- [10] M. Fazel, H. Hindi, and S.P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American Control Conference*, volume 6, pages 4734–4739, 2001.
- [11] K. Jain and V.V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *Journal of the ACM (JACM)*, 48(2):274–296, 2001.
- [12] L.V. Kantorovich. On the transfer of masses. In *Dokl. Akad. Nauk. SSSR*, volume 37, pages 227–229, 1942.

- [13] Danial Lashkari and Polina Golland. Convex clustering with exemplar-based models. In *Advances in Neural Information Processing Systems*, pages 825–832, 2007.
- [14] J.H. Lin and J.S. Vitter. ϵ -approximations with minimum packing constraint violation. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 771–782. ACM, 1992.
- [15] G. Monge. *Mémoire sur la théorie des déblais et des remblais*. De l’Imprimerie Royale, 1781.
- [16] M. Pilanci, L. El Ghaoui, and V. Chandrasekaran. Recovery of sparse probability measures via convex programming. In *Advances in Neural Information Processing Systems*, pages 2429–2437, 2012.
- [17] S.T. Rachev and L. Rüschendorf. *Mass Transportation Problems: Volume I: Theory*. Springer, 1998.
- [18] S.T. Rachev and L. Rüschendorf. *Mass Transportation Problems: Volume II: Applications*. Springer, 1998.
- [19] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [20] C. Villani. *Topics in optimal transportation*, volume 58. American Mathematical Society, 2003.
- [21] C. Villani. *Optimal transport: old and new*, volume 338. Springer, 2008.
- [22] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2005.