Published in final edited form as:

J Comput Graph Stat. 2015 ; 24(4): 994–1013. doi:10.1080/10618600.2014.948181.

Splitting Methods for Convex Clustering

Eric C. Chi* and Kenneth Lange†

Eric C. Chi: echi@rice.edu; Kenneth Lange: klange@ucla.edu

*Postdoctoral Research Associate, Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005

†Professor, Departments of Biomathematics, Human Genetics, and Statistics, University of California, Los Angeles, CA 90095-7088

Abstract

Clustering is a fundamental problem in many scientific applications. Standard methods such as k -means, Gaussian mixture models, and hierarchical clustering, however, are beset by local minima, which are sometimes drastically suboptimal. Recently introduced convex relaxations of k -means and hierarchical clustering shrink cluster centroids toward one another and ensure a unique global minimizer. In this work we present two splitting methods for solving the convex clustering problem. The first is an instance of the alternating direction method of multipliers (ADMM); the second is an instance of the alternating minimization algorithm (AMA). In contrast to previously considered algorithms, our ADMM and AMA formulations provide simple and unified frameworks for solving the convex clustering problem under the previously studied norms and open the door to potentially novel norms. We demonstrate the performance of our algorithm on both simulated and real data examples. While the differences between the two algorithms appear to be minor on the surface, complexity analysis and numerical experiments show AMA to be significantly more efficient. This article has supplemental materials available online.

Keywords

Convex optimization; Regularization paths; Alternating minimization algorithm; Alternating direction method of multipliers; Hierarchical clustering; k -means

1 Introduction

In recent years convex relaxations of many fundamental, yet combinatorially hard, optimization problems in engineering, applied mathematics, and statistics have been introduced (Tropp, 2006). Good, and sometimes nearly optimal solutions, can be achieved at affordable computational prices for problems that appear at first blush to be computationally

SUPPLEMENTAL MATERIALS

Proofs, derivations, algorithmic details, and additional numerical results: The Supplemental Materials include proofs for all propositions, details on the more complicated derivations, the stopping criteria for both algorithms, a sketch of algorithm acceleration, derivation of the dual problem, additional discussion on computational complexity, additional figures showing timing comparisons, and application of convex clustering to two additional real data sets. (Supplement.pdf)

Code: An R package, **cvxclu**, which implements the AMA and ADMM algorithms in this paper, is available on the CRAN website. The mammals data set is included in it.

intractable. In this paper, we introduce two new algorithmic frameworks based on variable splitting that generalize and extend recent efforts to convexify the classic unsupervised problem of clustering.

Lindsten et al. (2011) and Hocking et al. (2011) formulate the clustering task as a convex optimization problem. Given n points $\mathbf{x}_1, \dots, \mathbf{x}_n$ in $\mathbb{R}^p$, they suggest minimizing the convex criterion

$$F_\gamma(\mathbf{U}) = \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|_2^2 + \gamma \sum_{i < j} w_{ij} \|\mathbf{u}_i - \mathbf{u}_j\|, \quad (1.1)$$

where γ is a positive tuning constant, $w_{ij} = w_{ji}$ is a nonnegative weight, and the i th column $\mathbf{u}_i$ of the matrix $\mathbf{U}$ is the cluster center (centroid) attached to point $\mathbf{x}_i$. Lindsten et al. (2011) consider an ℓ_q norm penalty on the differences $\mathbf{u}_i - \mathbf{u}_j$ while Hocking et al. (2011) consider ℓ_1 , ℓ_2 , and ℓ_∞ penalties. In the current paper, an arbitrary norm defines the penalty.

The objective function bears some similarity to the fused lasso signal approximator (Tibshirani et al., 2005). When the ℓ_1 penalty is used in definition (1.1), we recover a special case of the General Fused Lasso (Hoeftling, 2010; Tibshirani and Taylor, 2011). In the graphical interpretation of clustering, each point corresponds to a node in a graph, and an edge connects nodes i and j whenever $w_{ij} > 0$. Figure 1 depicts an example. In this case, the objective function $F_\gamma(\mathbf{U})$ separates over the connected components of the underlying graph. Thus, one can solve for the optimal $\mathbf{U}$ component by component. Without loss of generality, we assume the graph is connected.

When $\gamma = 0$, the minimum is attained when $\mathbf{u}_i = \mathbf{x}_i$ and each point occupies a unique cluster. As γ increases, the cluster centers begin to coalesce. Two points $\mathbf{x}_i$ and $\mathbf{x}_j$ with $\mathbf{u}_i = \mathbf{u}_j$ are said to belong to the same cluster. For sufficiently high γ all points coalesce into a single cluster. Because the objective function $F_\gamma(\mathbf{U})$ in equation (1.1) is strictly convex and coercive, it possesses a unique minimum point for each value of γ . If we plot the solution matrix $\mathbf{U}$ as a function of γ , then we can ordinarily identify those values of γ giving k clusters for any integer k between n and 1. In theory, k can decrement by more than 1 as certain critical values of γ are passed. Indeed, when points are not well separated, we observe that many centroids will coalesce abruptly unless care is taken in choosing the weights w_{ij} .

The benefits of this formulation are manifold. As we will show, convex relaxation admits a simple and fast iterative algorithm that is guaranteed to converge to the unique global minimizer. In contrast, the classic k -means problem has been shown to be NP-hard (Aloise et al., 2009; Dasgupta and Freund, 2009). In addition, the classical greedy algorithm for solving k -means clustering often gets trapped in suboptimal local minima (Forgy, 1965; Lloyd, 1982; MacQueen, 1967).

Another vexing issue in clustering is determining the number of clusters. Agglomerative hierarchical clustering (Gower and Ross, 1969; Johnson, 1967; Lance and Williams, 1967; Murtagh, 1983; Ward, 1963) finesses the problem by computing an entire clustering path.

Agglomerative approaches, however, can be computationally demanding and tend to fall into suboptimal local minima since coalescence events are not reversed. The alternative convex relaxation considered here performs continuous clustering just as the lasso (Chen et al., 1998; Tibshirani, 1996) performs continuous variable selection. Figure 2 shows how the solutions to the alternative convex problem traces out an intuitively appealing, globally optimal, and computationally tractable solution path.

1.1 Contributions

Our main contributions are two new methods for solving the convex clustering problem. Relatively little work has been published on algorithms for solving this optimization problem. In fact, the only other paper introducing dedicated algorithms for minimizing criterion (1.1) that we are aware of is Hocking et al. (2011). Lindsten et al. (2011) used the off-the-shelf convex solver CVX (CVX Research, Inc., 2012; Grant and Boyd, 2008) to generate solution paths. Hocking et al. (2011) note that CVX is useful for solving small problems but a dedicated formulation is required for scalability. Thus, they introduced three distinct algorithms for the three most commonly encountered norms. Given the ℓ_1 norm and unit weights w_{ij} , the objective function separates, and they solve the convex clustering problem by the exact path following method designed for the fused lasso (Hoeftling, 2010). For the ℓ_1 and ℓ_2 norms with arbitrary weights w_{ij} , they employ subgradient descent in conjunction with active sets. Finally, they solve the convex clustering problem under the ℓ_∞ norm by viewing it as minimization of a Frobenius norm over a polytope. In this guise, the problem succumbs to the Frank-Wolfe algorithm (Frank and Wolfe, 1956) of quadratic programming.

In contrast to this piecemeal approach, we introduce two similar generic frameworks for minimizing the convex clustering objective function with an arbitrary norm. One approach solves the problem by the alternating direction method of multipliers (ADMM), while the other solves it by the alternating minimization algorithm (AMA). The key step in both cases computes the proximal map of a given norm. Consequently, both of our algorithms apply provided the penalty norm admits efficient computation of its proximal map.

In addition to introducing new algorithms for solving the convex clustering problem, the current paper contributes in other concrete ways: (a) We combine existing results on AMA and ADMM with the special structure of the convex clustering problem to characterize both of the new algorithms theoretically. In particular, the clustering problem formulation gives a minimal set of extra assumptions needed to prove the convergence of the ADMM iterates to the unique global minimum. We also explicitly show how the computational and storage complexity of our algorithms scales with the connectivity of the underlying graph. Examination of the dual problem enables us to identify a fixed step size for AMA that is associated with the Laplacian matrix of the underlying graph. Finally, our complexity analysis enables us to rigorously quantify the efficiency of the two algorithms so the two methods can be compared. (b) We provide new proofs of intuitive properties of the solution path. These results are tied solely to the minimization of the objective function (1.1) and hold regardless of the algorithm used to find the minimum point. (c) We provide guidance on how to choose the weights w_{ij} . Our suggested choices diminish computational complexity

and enhance solution quality. In particular, we show that employing k -nearest neighbor weights allows the storage and computation requirements for our AMA method to grow linearly in problem size.

1.2 Related Work

The literature on clustering is immense; the reader can consult the books (Gordon, 1999; Hartigan, 1975; Kaufman and Rousseeuw, 1990; Mirkin, 1996; Wu and Wunsch, 2009) for a comprehensive review. The clustering function (1.1) can be viewed as a convex relaxation of either k -means clustering (Lindsten et al., 2011) or hierarchical agglomerative clustering (Hocking et al., 2011). Both of these classical clustering methods (Sneath, 1957; Sørensen, 1948; Ward, 1963) come in several varieties. The literature on k -means clustering reports notable improvements in the computation (Elkan, 2003) and quality of solutions (Arthur and Vassilvitskii, 2007; Bradley et al., 1997; Kaufman and Rousseeuw, 1990) delivered by the standard greedy algorithms. Faster methods for agglomerative hierarchical clustering have been developed as well (Fraley, 1998). Many statisticians view the hard cluster assignments of k -means as less desirable than the probabilistic assignments generated by mixture models (McLachlan, 2000; Titterton et al., 1985). Mixture models have the advantage of gracefully assigning points to overlapping clusters. These models are amenable to an EM algorithm and can be extended to infinite mixtures (Ferguson, 1973; Rasmussen, 2000; Neal, 2000).

Alternative approaches to clustering involve identifying components in the associated graph via its Laplacian matrix. Spectral clustering (Luxburg, 2007) can be effective in cases when the clusters are non-convex and linearly inseparable. Although spectral clustering is valuable, it does not conflict with convex relaxation. Indeed, Hocking et al. (2011) demonstrate that convex clustering can be effectively merged with spectral clustering. Although we agree with this point, the solution path revealed by convex clustering is meritorious in its own right because it partially obviates the persistent need for determining the number of clusters.

1.3 Notation and Organization

Throughout, scalars are denoted by lowercase letters (a), vectors by boldface lowercase letters ($\mathbf{u}$), and matrices by boldface capital letters ($\mathbf{U}$). The j th column of a matrix $\mathbf{U}$ is denoted by $\mathbf{u}_j$. At times in our derivations, it will be easier to work with vectorized matrices. We adopt the convention of denoting the vectorization of a matrix ($\mathbf{U}$) by its lower case letter in boldface ($\mathbf{u}$). Finally, we denote sets by upper case letters (B).

The rest of the paper is organized as follows. We first characterize the solution path theoretically. Previous papers take intuitive properties of the path for granted. We then review the ADMM and AMA algorithms and adapt them to solve the convex clustering problem. Once the algorithms are specified, we discuss their convergence and computational and storage complexity. Practical issues such as weight choice and refinements such as accelerating the algorithms are discussed next. We then present some numerical examples of clustering. The paper concludes with a general discussion.

2 Properties of the solution path

The solution path $\mathbf{U}(\gamma)$ enjoys several nice properties as a function of the regularization parameter γ and the weight matrix $\mathbf{W} = (w_{ij})$. These properties expedite its numerical computation. The proof of the following two propositions can be found in the Supplemental Materials.

Proposition 2.1

The solution path $\mathbf{U}(\gamma)$ exists, is unique, and depends continuously on γ . The path also depends continuously on the weight matrix $\mathbf{W}$.

Existence and uniqueness of $\mathbf{U}$ sets the stage for a well-posed optimization problem. Continuity of $\mathbf{U}$ suggests employing homotopy continuation. Indeed, empirically we find great time savings in solving a sequence of problems over a grid of γ values when we use the solution of a previous value of γ as a warm start or initial value for the next larger γ value.

It is plausible that the centroids eventually coalesce to a common point as γ becomes sufficiently large. For the example shown in Figure 1, we intuitively expect for sufficiently large γ that the columns of $\mathbf{U}$ satisfy $\mathbf{u}_3 = \mathbf{u}_4 = \bar{\mathbf{x}}_{34}$ and $\mathbf{u}_1 = \mathbf{u}_2 = \mathbf{u}_5 = \bar{\mathbf{x}}_{125}$, where $\bar{\mathbf{x}}_{34}$ is the mean of $\mathbf{x}_3$ and $\mathbf{x}_4$ and $\bar{\mathbf{x}}_{125}$ is the mean of $\mathbf{x}_1$, $\mathbf{x}_2$, and $\mathbf{x}_5$. The next proposition rigorously confirms our intuition.

Proposition 2.2

Suppose each point corresponds to a node in a graph with an edge between nodes i and j whenever $w_{ij} > 0$. If this graph is connected, then $F_\gamma(\mathbf{U})$ is minimized by $\bar{\mathbf{X}}$ for γ sufficiently large, where each column of $\bar{\mathbf{X}}$ equals the average $\bar{\mathbf{x}}$ of the n vectors $\mathbf{x}_i$.

We close this section by noting that in general the clustering paths are not guaranteed to be agglomerative. In the special case of the ℓ_1 norm with uniform weights $w_{ij} = 1$, Hocking et al. (2011) prove that the path is agglomerative. In the same paper they give an ℓ_2 norm example where the centroids fuse and then unfuse as the regularization parameter increases. This behavior, however, does not seem to occur very frequently in practice. Nonetheless, in the algorithms we describe next, we allow for such fission events to ensure that the computed solution path is truly the global minimizer of the convex criterion (1.1).

3 Algorithms to Compute the Clustering Path

Having characterized the solution path $\mathbf{U}(\gamma)$, we now tackle the task of computing it. We present two closely related optimization approaches: the alternating direction method of multipliers (ADMM) (Boyd et al., 2011; Gabay and Mercier, 1976; Glowinski and Marrocco, 1975) and the alternating minimization algorithm (AMA) (Tseng, 1991). Both approaches employ variable splitting to handle the shrinkage penalties in the convex clustering criterion (1.1).

Let us first recast the convex clustering problem as the equivalent constrained problem

$$\begin{aligned} & \text{minimize } \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|_2^2 + \gamma \sum_{l \in \mathcal{E}} w_l \|\mathbf{v}_l\| \\ & \text{subject to } u_{l_1} - u_{l_2} - v_l = 0. \end{aligned} \quad (3.1)$$

Here we index a centroid pair by $l = (l_1, l_2)$ with $l_1 < l_2$, define the set of edges over the non-zero weights $\mathcal{E} = \{l = (l_1, l_2) : w_l > 0\}$, and introduce a new variable $\mathbf{v}_l = \mathbf{u}_{l_1} - \mathbf{u}_{l_2}$ to account for the difference between the two centroids. The purpose of variable splitting is to simplify optimization with respect to the penalty terms.

Splitting methods such as ADMM and AMA have been successfully used to attack similar problems in image restoration (Goldstein and Osher, 2009). To clarify the similarities and differences between ADMM and AMA, we briefly review how these two methods iteratively solve the following constrained optimization problem

$$\begin{aligned} & \text{minimize } f(\mathbf{u}) + g(\mathbf{v}) \\ & \text{subject to } \mathbf{A}\mathbf{u} + \mathbf{B}\mathbf{v} = \mathbf{c}, \end{aligned} \quad (3.2)$$

which includes the constrained minimization problem (3.1) as a special case. Recall that finding the minimizer to an equality constrained optimization problem is equivalent to the identifying the saddle point of the associated Lagrangian function. Both ADMM and AMA invoke a related function called the augmented Lagrangian,

$$\mathcal{L}_\nu(\mathbf{u}, \mathbf{v}, \boldsymbol{\lambda}) = f(\mathbf{u}) + g(\mathbf{v}) + \langle \boldsymbol{\lambda}, \mathbf{c} - \mathbf{A}\mathbf{u} - \mathbf{B}\mathbf{v} \rangle + \frac{\nu}{2} \|\mathbf{c} - \mathbf{A}\mathbf{u} - \mathbf{B}\mathbf{v}\|_2^2,$$

where the dual variable $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers and ν is a nonnegative tuning parameter. When $\nu = 0$, the augmented Lagrangian coincides with the ordinary Lagrangian.

ADMM minimizes the augmented Lagrangian one block of variables at a time before updating the dual variable $\boldsymbol{\lambda}$. This yields the algorithm

$$\begin{aligned} \mathbf{u}^{m+1} &= \arg \min_{\mathbf{u}} \mathcal{L}_\nu(\mathbf{u}, \mathbf{v}^m, \boldsymbol{\lambda}^m) \\ \mathbf{v}^{m+1} &= \arg \min_{\mathbf{v}} \mathcal{L}_\nu(\mathbf{u}^{m+1}, \mathbf{v}, \boldsymbol{\lambda}^m) \\ \boldsymbol{\lambda}^{m+1} &= \boldsymbol{\lambda}^m + \nu(\mathbf{c} - \mathbf{A}\mathbf{u}^{m+1} - \mathbf{B}\mathbf{v}^{m+1}). \end{aligned} \quad (3.3)$$

AMA takes a slightly different tack and updates the first block $\mathbf{u}$ without augmentation, assuming $f(\mathbf{u})$ is strongly convex. This change is accomplished by setting the positive tuning constant ν to be 0. Thus, we update the first block $\mathbf{u}$ as

$$\mathbf{u}^{m+1} = \arg \min_{\mathbf{u}} \mathcal{L}_0(\mathbf{u}, \mathbf{v}^m, \boldsymbol{\lambda}^m), \quad (3.4)$$

and update $\mathbf{v}$ and $\boldsymbol{\lambda}$ as indicated in (3.3). Later we will see that this seemingly innocuous change pays large dividends in the convex clustering problem. Although block descent appears to complicate matters, it often markedly simplifies optimization in the end. In the case of convex clustering, the updates are either simple linear transformations or evaluations of proximal maps, which we discuss next.

3.1 Proximal Map

For $\sigma > 0$ the function

$$\text{prox}_{\sigma\Omega}(\mathbf{u}) = \arg \min_{\mathbf{v}} \left[\sigma\Omega(\mathbf{v}) + \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 \right]$$

is a well-studied operation called the proximal map of the function $\Omega(\mathbf{v})$. The proximal map exists and is unique whenever the function $\Omega(\mathbf{v})$ is convex and lower semicontinuous. Norms satisfy these conditions, and for many norms of interest the proximal map can be evaluated by either an explicit formula or an efficient algorithm. Table 1 lists some common examples. The proximal maps for the ℓ_1 and ℓ_2 norms have explicit solutions and can be computed in $\mathcal{O}(p)$ operations for a vector $\mathbf{v} \in \mathbb{R}^p$. The proximal map for the ℓ_∞ norm requires projection onto the unit simplex and lacks an explicit solution. However, there are good algorithms for projecting onto the unit simplex (Duchi et al., 2008; Michelot, 1986). In particular, Duchi et al.'s projection algorithm makes it possible to evaluate $\text{prox}_{\sigma\|\cdot\|_\infty}(\mathbf{v})$ in $\mathcal{O}(p \log p)$ operations.

3.2 ADMM updates

The augmented Lagrangian for criterion (3.1) is given by

$$\begin{aligned} \mathcal{L}_\nu(\mathbf{U}, \mathbf{V}, \boldsymbol{\Lambda}) = & \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|_2^2 + \gamma \sum_{l \in \mathcal{E}} w_l \|\mathbf{v}_l\| \\ & + \sum_{l \in \mathcal{E}} \langle \boldsymbol{\lambda}_l, \mathbf{v}_l - \mathbf{u}_{l_1} + \mathbf{u}_{l_2} \rangle + \frac{\nu}{2} \sum_{l \in \mathcal{E}} \|\mathbf{v}_l - \mathbf{u}_{l_1} + \mathbf{u}_{l_2}\|_2^2. \end{aligned} \quad (3.5)$$

To update $\mathbf{U}$, we need to minimize the function

$$f(\mathbf{U}) = \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|_2^2 + \frac{\nu}{2} \sum_{l \in \mathcal{E}} \|\tilde{\mathbf{v}}_l - \mathbf{u}_{l_1} + \mathbf{u}_{l_2}\|_2^2,$$

where $\tilde{\mathbf{v}}_l = \mathbf{v}_l + \nu^{-1} \boldsymbol{\lambda}_l$. The gradient of this function vanishes at $\mathbf{U}$ satisfying the linear system

$$\mathbf{U}\mathbf{M} = \mathbf{X} + \nu \sum_{l \in \mathcal{E}} \tilde{\mathbf{v}}_l (\mathbf{e}_{l_1} - \mathbf{e}_{l_2})^t, \quad (3.6)$$

where $\mathbf{M} = \mathbf{I} + \nu \sum_{l \in \mathcal{E}} (\mathbf{e}_{l_1} - \mathbf{e}_{l_2})(\mathbf{e}_{l_1} - \mathbf{e}_{l_2})^t$. If the edge set $\mathcal{E}$ contains all possible edges, then the update for $\mathbf{U}$ can be computed analytically as

$$\mathbf{u}_i = \frac{1}{1+n\nu} \mathbf{y}_i + \frac{n\nu}{1+n\nu} \bar{\mathbf{x}}, \quad (3.7)$$

where $\bar{\mathbf{x}}$ is the average column of $\mathbf{X}$ and

$$\mathbf{y}_i = \mathbf{x}_i + \sum_{l_1=i} [\boldsymbol{\lambda}_l + \nu \mathbf{v}_l] - \sum_{l_2=i} [\boldsymbol{\lambda}_l + \nu \mathbf{v}_l].$$

Detailed derivations of the equalities in (3.6) and (3.7) can be found in the Supplemental Materials.

Before deriving the updates for $\mathbf{V}$, we remark that although a fully connected weights graph allows one to write explicit updates for $\mathbf{U}$, doing so comes at the cost of increasing the number of variables $\mathbf{v}_l$ and λ_l . Such choices are not immaterial, and we will discuss the tradeoffs later. To update $\mathbf{V}$, we first observe that the augmented Lagrangian $\mathcal{L}_\nu(\mathbf{U}, \mathbf{V}, \Lambda)$ is separable in the vectors $\mathbf{v}_l$. A particular difference vector $\mathbf{v}_l$ is determined by the proximal map

$$\begin{aligned} \mathbf{v}_l = \arg \min_{\mathbf{v}_l} \frac{1}{2} \left[\|\mathbf{v}_l - (\mathbf{u}_{l_1} - \mathbf{u}_{l_2} - \nu^{-1} \lambda_l)\|_2^2 + \frac{\gamma w_l}{\nu} \|\mathbf{v}_l\| \right] \\ = \text{prox}_{\sigma_l \|\cdot\|}(\mathbf{u}_{l_1} - \mathbf{u}_{l_2} - \nu^{-1} \lambda_l), \end{aligned} \quad (3.8)$$

where $\sigma_l = \gamma w_l / \nu$. Finally, the Lagrange multipliers are updated by

$$\lambda_l = \lambda_l + \nu(\mathbf{v}_l - \mathbf{u}_{l_1} - \mathbf{u}_{l_2}).$$

Algorithm 1 summarizes the ADMM algorithm. To track the progress of ADMM, we use standard methods given in (Boyd et al., 2011) based on primal and dual residuals. Details on the stopping rules that we employ are given in the Supplemental Materials.

Algorithm 1 ADMM

Initialize Λ^0 and $\mathbf{V}^0$.

1: **for** $m = 1, 2, 3, \dots$ **do**

2: **for** $i = 1, \dots, n$ **do**

3: $\mathbf{y}_i = \mathbf{x}_i + \sum_{l_1=i} [\lambda_l^{m-1} + \nu \mathbf{v}_l^{m-1}] - \sum_{l_2=i} [\lambda_l^{m-1} + \nu \mathbf{v}_l^{m-1}]$

4: **end for**

5: $\mathbf{U}^m = \frac{1}{1+n\nu} \mathbf{Y} + \frac{n\nu}{1+n\nu} \bar{\mathbf{X}}$

6: **for all** l **do**

7: $\mathbf{v}_l^m = \text{prox}_{\sigma_l \|\cdot\|}(\mathbf{u}_{l_1}^m - \mathbf{u}_{l_2}^m - \nu^{-1} \lambda_l^{m-1})$

8: $\lambda_l^m = \lambda_l^{m-1} + \nu(\mathbf{v}_l^m - \mathbf{u}_{l_1}^m + \mathbf{u}_{l_2}^m)$

9: **end for**

10: **end for**

3.3 AMA updates

Since AMA shares its update rules for $\mathbf{V}$ and Λ with ADMM, consider updating $\mathbf{U}$. Recall that AMA updates $\mathbf{U}$ by minimizing the ordinary Lagrangian. In the $\nu = 0$ case, we have

$$\mathbf{U}^{m+1} = \arg \min_{\mathbf{U}} \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|_2^2 + \sum_{l \in \mathcal{C}} \langle \boldsymbol{\lambda}_l^m, \mathbf{v}_l - \mathbf{u}_{l_1} - \mathbf{u}_{l_2} \rangle.$$

In contrast to ADMM, this minimization separates in each $\mathbf{u}_i$ and gives an update that does not depend on $\mathbf{v}_l$, namely

$$\mathbf{u}_i^{m+1} = \mathbf{x}_i + \sum_{l_1=i} \boldsymbol{\lambda}_{l_1}^m - \sum_{l_2=i} \boldsymbol{\lambda}_{l_2}^m. \quad (3.9)$$

Further scrutiny of the updates for $\mathbf{V}$ and $\boldsymbol{\Lambda}$ reveals that AMA does not even require computing $\mathbf{V}$. Applying standard results from convex calculus, it can be shown that $\boldsymbol{\Lambda}$ has the following update

$$\boldsymbol{\lambda}_l^{m+1} = \mathcal{P}_{C_l} \left(\boldsymbol{\lambda}_l^m - \nu \mathbf{g}_l^{m+1} \right), \quad (3.10)$$

where $\mathbf{g}_l^m = \mathbf{u}_{l_1}^m - \mathbf{u}_{l_2}^m$, $C_l = \{\boldsymbol{\lambda}_l : \|\boldsymbol{\lambda}_l\|_{\dagger} \leq \gamma \omega_l\}$, and $\|\cdot\|_{\dagger}$ is the dual norm of the norm defining the fusion penalty. Algorithm 2 summarizes the AMA algorithm. Detailed derivation of this simplification is in the Supplemental Materials. At the termination of Algorithm 2, the centroids can be recovered from the dual variables using equation (3.9).

Algorithm 2 AMA

Initialize $\boldsymbol{\lambda}^0$.

```

1:  for  $m = 1, 2, 3, \dots$  do
2:    for  $i = 1, \dots, n$  do
3:       $\Delta_i^m = \sum_{l_1=i} \boldsymbol{\lambda}_{l_1}^{m-1} - \sum_{l_2=i} \boldsymbol{\lambda}_{l_2}^{m-1}$ 
4:    end for
5:    for all  $l$  do
6:       $\mathbf{g}_l^m = \mathbf{x}_{l_1} - \mathbf{x}_{l_2} + \Delta_{l_1}^m - \Delta_{l_2}^m$ 
7:       $\boldsymbol{\lambda}_l^m = \mathcal{P}_{C_l} \left( \boldsymbol{\lambda}_l^{m-1} - \nu \mathbf{g}_l^m \right)$ 
8:    end for
9:  end for
```

Note that Algorithm 2 look remarkably like a projected gradient algorithm. Indeed, Tseng (1991) shows that AMA is actually performing proximal gradient ascent to maximize a dual problem. Since the dual of the convex clustering problem (3.1) is essentially a constrained least squares problem, it is hardly surprising that it can be solved numerically by the classic projected gradient algorithm, a special case of the proximal gradient algorithm. In addition to providing a simple interpretation of the AMA method, the dual problem allows us to derive a rigorous stopping criterion for AMA based on the duality gap, a bound on how much the primal objective evaluated at the current iterate exceeds its optimal value. Due to

space limitations, both the dual problem and duality gap computation are covered in the Supplemental Materials. Before proceeding, however, let us emphasize that AMA requires tracking of only as many dual variable λ_l as there are non-zero weights. We will find later that sparse weights often produces better quality clusterings. Thus, when relatively few weights are non-zero, the number of variables introduced by splitting does not become prohibitive under AMA.

4 Convergence

Both ADMM and AMA converge under reasonable conditions. Convergence of ADMM is guaranteed for any $\nu > 0$. Convergence for AMA is guaranteed when ν is not too large. In this section, we show how the specific structure of the convex clustering problem can give stronger convergence results for ADMM and provide guidance on choosing ν in AMA and improve AMA's rate of convergence in practice. All proofs can be found in the Supplemental Materials.

4.1 ADMM

Under minimal assumptions, it can be shown that the limit points of an ADMM iterate sequence coincide with the stationary points of the objective function being minimized (Boyd et al., 2011). Note, however, this does not guarantee that the iterates $\mathbf{U}^m$ converge to $\mathbf{U}^*$. Since the convex clustering criterion $F_\gamma(\mathbf{U})$ defined by equation (1.1) is strictly convex and coercive, we have the stronger result that the ADMM iterate sequence converges to the unique global minimizer $\mathbf{U}^*$ of $F_\gamma(\mathbf{U})$.

Proposition 4.1—The iterates $\mathbf{U}^m$ in Algorithm 1 converge to the unique global minimizer $\mathbf{U}^*$ of the clustering criterion $F_\gamma(\mathbf{U})$.

4.2 AMA

The convergence of Algorithm 2 hinges on the choice of ν , which in turn depends on the connectivity of the associated graph.

Proposition 4.2—Let $\mathbf{u}_i^{m+1} = \mathbf{x}_i + \sum_{l_1=i} \lambda_{l_1}^m - \sum_{l_2=i} \lambda_{l_2}^m$, where λ_l^m , where λ_l^m are the iterates in Algorithm 2. Then the sequence $\mathbf{U}^{m+1}$ converges to the unique global minimizer $\mathbf{U}^*$ of the clustering criterion $F_\gamma(\mathbf{U})$, provided that $\nu < 2/\rho(\mathbf{L})$, where $\rho(\mathbf{L})$ denotes the largest eigenvalue of $\mathbf{L}$, the Laplacian matrix of the associated graph.

In lieu of computing $\rho(\mathbf{L})$ numerically, one can bound it by theoretical arguments. In general $\rho(\mathbf{L}) \leq n$ (Anderson and Morley, 1985), with equality when the graph is fully connected and $w_{ij} > 0$ for all $i < j$. Choosing a fixed step size of $\nu < 2/n$ works in practice when there are fewer than 1000 data points and the graph is dense. For a sparse graph with bounded node degrees, the sharper bound

$$\rho(\mathbf{L}) \leq \max \{d(i) + d(j) : (i, j) \in \mathcal{E}\}$$

applies, where $d(i)$ is the degree of the i th node (Anderson and Morley, 1985). This bound can be computed quickly in $\mathcal{O}(n+\varepsilon)$ operations, where ε denotes the number of edges in $\mathcal{E}$. Section 7.2 demonstrates the overwhelming speed advantage AMA on sparse graphs.

5 Computational Complexity

The computational complexity of an optimization algorithm includes the amount of work per iteration, the number of iterations until convergence, and the overall memory requirements. We discuss the first and third components. For the AMA and ADMM convex clustering algorithms, the number of iterations to achieve a desired level of accuracy can be derived from existing complexity results in the literature. Details can be found in the Supplemental Materials.

5.1 AMA

Inspection of Algorithm 2 shows that computing all $\mathbf{g}_i$ requires $p(2\varepsilon - n)$ total additions and subtractions. Computing all vectors $\mathbf{g}_i$ in Algorithm 2 takes $\mathcal{O}(\varepsilon p)$ operations, and taking the subsequent gradient step costs $\mathcal{O}(\varepsilon p)$ operations. Computing the needed projections costs $\mathcal{O}(\varepsilon p)$ operations for the ℓ_1 and ℓ_2 norms and $\mathcal{O}(\varepsilon p \log p)$ operations for the ℓ_∞ norm. Finally computing the duality gap costs $\mathcal{O}(np + \varepsilon p)$ operations. Details on the duality gap computation appear in the Supplemental Materials. The assumption that n is $\mathcal{O}(\varepsilon)$ entails smaller costs. A single iteration with gap checking then costs just $\mathcal{O}(\varepsilon p)$ operations for the ℓ_1 and ℓ_2 norms and $\mathcal{O}(\varepsilon p \log p)$ operations for the ℓ_∞ norm.

Total storage is $\mathcal{O}(p\varepsilon + np)$. In the worst case ε is $\binom{n}{2}$. However, if we limit a node's connectivity to its k -nearest neighbors, then ε is $\mathcal{O}(kn)$. Thus, the computational complexity of the problem in the worst case is quadratic in the number of points n and linear under the restriction to k -nearest neighbors connectivity. Storage is quadratic in n in the worst case and linear in n under the k -nearest neighbors restriction. Thus, limiting a point's connectivity to its k -nearest neighbors renders both the storage requirements and operation counts linear in the problem size, namely $\mathcal{O}(knp)$.

5.2 ADMM

We have two cases to consider. First consider the explicit updates outlined in Algorithm 1 when the edge set $\mathcal{E}$ is full. By nearly identical arguments as above, the complexity of a single round of ADMM updates with primal and dual residual calculation requires $\mathcal{O}(n^2 p)$ operations for the ℓ_1 and ℓ_2 norms and $\mathcal{O}(n^2 p \log p)$ operations for the ℓ_∞ norm. Thus, the ADMM algorithm based on the explicit updates (3.7) requires the same computational effort as AMA in the worst case. In this setting both ADMM and AMA also have $\mathcal{O}(pn^2)$ storage requirements.

The situation does not improve by much when we consider the more frugal alternative in which $\mathcal{E}$ contains only node pairings corresponding to non-zero weights. In this case, the variables $\mathbf{A}$ and $\mathbf{V}$ have only as many columns as there are non-zero weights. Now the

storage requirements are $\mathcal{O}(p\varepsilon + np)$ like AMA, but the cost of updating $\mathbf{U}$, the most computationally demanding step, remains quadratic in n . Recall we need to solve a linear system of equations (3.6)

$$\mathbf{U}\mathbf{M} = \mathbf{X} + \sum_{l \in \mathcal{E}} \tilde{\mathbf{v}}_l (\mathbf{e}_{l_1} - \mathbf{e}_{l_2})^t,$$

where $\mathbf{M} \in \mathbb{R}^{n \times n}$. Since $\mathbf{M}$ is positive definite and does not change throughout the ADMM iterations, the prudent course of action is to compute and cache its Cholesky factorization. The factorization requires $\mathcal{O}(n^3)$ operations to calculate but that cost can be amortized across the repeated ADMM updates. With the Cholesky factorization in hand, we can update each row of $\mathbf{U}$ by solving two sets of n -by- n triangular systems of equations, which together requires $\mathcal{O}(n^2)$ operations. Since $\mathbf{U}$ has p rows, the total amount of work to update $\mathbf{U}$ is $\mathcal{O}(n^2 p)$. Therefore, the overall amount of work per ADMM iteration is $\mathcal{O}(n^2 p + \varepsilon p)$ operations for the ℓ_1 and ℓ_2 norms and $\mathcal{O}(n^2 p + \varepsilon p \log p)$ operations for the ℓ_∞ norm. Thus, in stark contrast to AMA, both ADMM approaches grow quadratically, either in storage requirements or computational costs, regardless of how we might limit the size of the edge set $\mathcal{E}$.

6 Weights, Cluster Assignment, and Acceleration

We now address some practical issues that arise in applying our algorithms.

Weights

The weight matrix $\mathbf{W}$ can dramatically affect the quality of the clustering path. We set the weight between the i th and j th points equal to $w_{ij} = \iota_{\{i,j\}}^k \exp\left(-\phi \|\mathbf{x}_i - \mathbf{x}_j\|_2^2\right)$, where the indicator function $\iota_{\{i,j\}}^k$ is 1 if j is among i 's k -nearest-neighbors or vice versa and 0 otherwise. The second factor is a Gaussian kernel that slows the coalescence of distant points. The constant ϕ is nonnegative; the value $\phi = 0$ corresponds to uniform weights. As noted earlier, limiting positive weights to nearest neighbors improves both computational efficiency and clustering quality. Although the two factors defining the weights act similarly, their combination increases the sensitivity of the clustering path to the local density of the data.

Cluster Assignment

We determine clustering assignments as a function of γ by reading off which centroids fuse. For both ADMM and AMA, such assignments can be performed in $\mathcal{O}(n)$ operations using the differences variable $\mathbf{V}$. In the case of AMA, where we do not store a running estimate of $\mathbf{V}$, we compute $\mathbf{V}$ via the update (3.8) after the algorithm terminates for a given γ . Once we determine $\mathbf{V}$, we simply apply breadth-first search to identify the connected components of the graph induced by $\mathbf{V}$. This graph identifies a node with every data point and places an edge between the i th pair of points if and only if $\mathbf{v}_i = \mathbf{0}$. Each connected component corresponds to a cluster. Note that the graph described here varies with γ , through the matrix

$\mathbf{V}$, and is unrelated to the graph discussed in the rest of this paper, which depends solely on the weights $\mathbf{W}$ and is invariant to γ .

Acceleration

Both AMA and ADMM admit acceleration at little additional computational cost (Goldstein et al., 2012). In our timing comparisons, accelerated variants of AMA and ADMM are used. Details on our acceleration techniques can be found in the Supplemental Materials.

7 Numerical Experiments

We now report numerical experiments on convex clustering for a synthetic and real data set. Similar experiments on two additional real data sets (iris and Senate) can be found in the Supplemental Materials. In particular, we focus on how the choice of the weights w_{ij} affects the quality of the clustering solution. Prior research on this question is limited. Both Lindsten et al. (2011) and Hocking et al. (2011) suggest weights derived from Gaussian kernels and k -nearest neighbors. Because Hocking et al. try only Gaussian kernels, in this section we follow up on their untested suggestion of combining Gaussian kernels and k -nearest neighbors.

We also compare the run times of our splitting methods to the run times of the subgradient algorithm employed by Hocking et al. for ℓ_2 paths. Our attention here focuses on solving the ℓ_2 path since the rotational invariance of the 2-norm makes it a robust choice in practice. Hocking et al. provide R and C++ code for their algorithms. Our algorithms are implemented in R and C. To make a fair comparison, we run our algorithm until it reaches a primal objective value that is less than or equal to the primal objective value obtained by the subgradient algorithm. Specifically, we first run the Hocking et al. code to generate a clusterpath and record the sequence of γ 's generated by their code. We then run our algorithms over the same sequence of γ 's and stop once our primal objective value falls below their value. We also retain the native stopping rule computations employed by our splitting methods, namely the dual loss calculations for AMA and residual calculations for ADMM. Since AMA already calculates the primal loss, this is not an additional burden. Although convergence monitoring creates additional work for ADMM, the added primal loss calculation at worst changes only the constant in the complexity bound. This follows since computing the primal loss requires $\mathcal{O}(np + \varepsilon p)$ total operations.

7.1 Qualitative Comparisons

The following examples demonstrate how the character of the solution paths can vary drastically with the choice of weights w_{ij} .

Two Half Moons—Consider the standard simulated data of two interlocking half moons in $\mathbb{R}^2$ composed of 100 points each. Figure 3 shows four convex clustering paths computed assuming two different numbers of nearest neighbors (10 and 50) and two different kernel constants φ (0 and 0.5). The upper right panel makes it evident that limiting the number of nearest neighbors ($k = 10$) and imposing non-uniform Gaussian kernel weights ($\varphi = 0.5$) produce the best clustering path. Using too many neighbors and assuming uniform weights results in little agglomerative clustering until late in the clustering path (lower left panel).

The two intermediate cases diverge in interesting ways. The hardest set of points to cluster are the points in the upper half moon's right tip and the lower half moon's left tip. Limiting the number of nearest neighbors and omitting the Gaussian kernel (upper left panel) correctly agglomerates the easier points, but waffles on the harder points, agglomerating them only at the very end when all points coalesce at the grand mean. Conversely, using too many neighbors and the Gaussian kernel (lower right panel) leads to a clustering path that does not hedge but incorrectly assigns the harder points.

Dentition of mammals—Next we consider the problem of clustering mammals based on their dentition (Hartigan, 1975). Eight different kinds of teeth are tallied for each mammal: the number of top incisors, bottom incisors, top canines, bottom canines, top premolars, bottom premolars, top molars, and bottom molars. We removed observations with teeth distributions that were not unique, leaving us with 27 mammals. Figure 4 shows the resulting clustering paths under two different choices of weights. On the left $w_{ij} = 1$ for all i, j , and on the right we use 5-nearest neighbors and $\varphi = 0.5$. Weights sensitive to the local density give superior results. Since there are eight variables, to visualize results we project the data and the fitted clustering paths onto the first two principal components of the data. The cluster path gives a different and perhaps more sensible solution than a two dimensional principal component analysis (PCA). For example, the brown bat is considered more similar to the house bat and red bat, even though it is closer in the first two PCA coordinates to the coyote and opossum.

7.2 Timing Comparisons

We now present results on two batches of experiments, with dense weights in the first batch and sparse ones in the second. For the first set of experiments, we compared the run times of the subgradient descent algorithm of Hocking et al. (2011), accelerated ADMM, and accelerated AMA on 10 replicates of simulated data consisting of 100, 200, 300, 400, and 500 points in $\mathbb{R}^2$ drawn from a multivariate standard normal. We limited our study to at most 500 points because the subgradient algorithm took several hours on a single realization of 500 points. Limiting the number of data points allowed us to use the simpler, but less storage efficient, ADMM formulation given in equation (3.7). For AMA, we fixed the step size at $\nu = 1/n$. For all tests, we assigned full-connectivity weights based on the indicator $\ell_{\{i,j\}}^k \equiv 1$ and tuning constant $\varphi = -2$. The parameter φ was chosen to ensure that the smallest weight was bounded safely away from zero. The full-connectivity assumption illustrates the superiority of AMA even under its least favorable circumstances. To trace out the entire clusterpath, we ran the Hocking subgradient algorithm to completion and invoked its default stopping criterion, namely a gradient with an ℓ_2 norm below 0.001. As noted earlier, we stopped our ADMM and AMA algorithms once their centroid iterates achieved a primal loss less than or equal to that achieved by the subgradient algorithm.

Table 2 shows the resulting mean times in seconds. Box-plots showing how the run time scales with the number of data points n can be found in the Supplemental Materials. All three algorithms scale quadratically in the number of points. This is expected for ADMM and AMA because all weights w_{ij} are positive. Nonetheless, the three algorithms possess different rate constants, with AMA possessing the slowest mean growth. The subgradient

and ADMM algorithms performed commensurately, although the ADMM algorithm was almost twice as fast on average when $n = 500$. Again, to ensure fair comparisons with the subgradient algorithm, we required ADMM to make extra primal loss computations. This change tends to inflate its rate constant. In summary, we see that fast AMA leads to affordable computation times, on the order of minutes for hundreds of data points, in contrast to subgradient descent, which incurs run times on the order of hours for 400 to 500 data points.

In the second batch of experiments, the same set up is retained except for assignments of weights and step length choice for AMA. We used $\varphi = -2$ again, but this time we zeroed out

all weights except those corresponding to the $k = \frac{n}{4}$ nearest neighbors of each point. For AMA we used step sizes based on the bound (4.1). Table 3 shows the resulting mean run times in seconds. As before more detailed box-plot comparisons can be found in the Supplemental Materials. As attested by the shorter run times for all three algorithms, incorporation of sparse weights appears to make the problems easier to solve. In this case ADMM was uniformly better on average than the subgradient method for all n . Even more noteworthy is the pronounced speed advantage of AMA over the other two algorithms for large n . When clustering 500 points, AMA requires on average a mere 7 seconds compared to 5 to 7 minutes for the subgradient and ADMM algorithms.

8 Conclusion & Future Work

In this paper, we introduce two splitting algorithms for solving the convex clustering problem. The splitting perspective encourages path following, one of the chief benefits of convex clustering. The splitting perspective also permits centroid penalties to invoke an arbitrary norm. The only requirement is that the proximal map for the norm be readily computable. Equivalently, projection onto the unit ball of the dual norm should be straightforward. Because proximal maps and projection operators are generally well understood, it is possible for us to quantify the computational complexity and convergence properties of our algorithms.

It is noteworthy that ADMM did not fare as well as AMA. ADMM has become quite popular in machine learning circles in recent years. Applying variable splitting and using ADMM to iteratively solve the convex clustering problem seemed like an obvious and natural initial strategy. Only later during our study did we implement the less favored AMA algorithm. Considering how trivial the differences are between the generic block updates for ADMM (3.3) and AMA (3.4), we were surprised by the performance gap between them. In the convex clustering problem, however, there is a non-trivial difference between minimizing the augmented and unaugmented Lagrangian in the first block update. This task can be accomplished in less time and space by AMA.

Two features of the convex clustering problem make it an especially good candidate for solution by AMA. First, the objective function is strongly convex and therefore has a Lipschitz differentiable dual. Lipschitz differentiability is a standard condition ensuring the convergence of proximal gradient algorithms. Second, a good step size can be readily

computed from the Laplacian matrix generated by the edge set $\mathcal{E}$. Without this prior bound, we would have to employ a more complicated line-search.

Our complexity analysis and simulations show that the accelerated AMA method appears to be the algorithm of choice. Nonetheless, given that alternative variants of ADMM may close the performance gap (Deng and Yin, 2012; Goldfarb et al., 2012), we are reluctant to dismiss ADMM too quickly. Both algorithms deserve further investigation. For instance, in both ADMM and AMA, updates of $\mathbf{A}$ and $\mathbf{V}$ could be parallelized. Hocking et al. also employed an active set approach to reduce computations as the centroids coalesce. A similar strategy could be adopted in our framework, but it incurs additional overhead as checks for fission events have to be introduced. An interesting and practical question brought up by Hocking et al. remains open, namely under what conditions or weights are fusion events guaranteed to be permanent as γ increases? In all our experiments, we did not observe any fission events. Identifying those conditions would eliminate the need to check for fission in such cases and expedite computation.

For AMA, the storage demands and computational complexity of convex clustering depend quadratically on the number of points in the worst case. Limiting a point's connections to its k -nearest neighbors, for example, ensures that the number of edges in the graph is linear in the number of nodes in the graph. Eliminating long-range dependencies is often desirable anyway. Choosing sparse weights can improve both cluster quality and computational efficiency. Moreover, finding the exact k -nearest neighbors is likely not essential, and we conjecture that the quality of solutions would not suffer greatly if approximate nearest neighbors are used and algorithms for fast computation of approximately nearest neighbors are leveraged (Slaney and Casey, 2008). On very large problems, the best strategy might be to exploit the continuity of solution paths in the weights. This suggests starting with even sparser graphs than the desired one and generating a sequence of solutions to increasingly dense problems. A solution with fewer edges can serve as a warm start for the next problem with more edges.

The splitting perspective also invites extensions that impose structured sparsity on the centroids. Witten and Tibshirani (2010) discuss how sparse centroids can improve the quality of a solution, especially when only a relatively few features of data drive clustering. Structured sparsity can be accomplished by adding a sparsity-inducing norm penalty to the $\mathbf{U}$ updates. The update for the centroids for both AMA and ADMM then rely on another proximal map of a gradient step. Introducing a sparsifying norm, however, raises the additional complication of choosing the amount of penalization.

Except for a few hints about weights, our analysis leaves the topic of optimal clustering untouched. Recently, von Luxburg (2010) suggested some principled approaches to assessing the quality of a clustering assignment via data perturbation and resampling. These clues are worthy of further investigation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Jocelyn Chi, Daniel Duckworth, Tom Goldstein, Rob Tibshirani, and Genevera Allen for their helpful comments and suggestions. All plots were made using the open source R package ggplot2 (Wickham, 2009). This research was supported by the NIH United States Public Health Service grants GM53275 and HG006139

References

- Aloise D, Deshpande A, Hansen P, Popat P. NP-hardness of Euclidean sum-of-squares clustering. *Machine Learning*. 2009; 75:245–248.
- Anderson WN, Morley TD. Eigenvalues of the Laplacian of a graph. *Linear and Multilinear Algebra*. 1985; 18:141–145.
- Arthur, D.; Vassilvitskii, S. Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, SODA '07; 2007. k-means++: The advantages of careful seeding; p. 1027-1035.
- Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found Trends Mach Learn*. 2011; 3:1–122.
- Bradley, PS.; Mangasarian, OL.; Street, WN. Advances in Neural Information Processing Systems. MIT Press; 1997. Clustering via Concave Minimization; p. 368-374.
- Chen SS, Donoho DL, Saunders MA. Atomic Decomposition by Basis Pursuit. *SIAM Journal on Scientific Computing*. 1998; 20:33–61.
- CVX Research Inc. CVX: Matlab Software for Disciplined Convex Programming, version 2.0 beta. 2012. <http://cvxr.com/cvx>
- Dasgupta S, Freund Y. Random projection trees for vector quantization. *IEEE Trans Inf Theor*. 2009; 55:3229–3242.
- Deng, W.; Yin, W. Tech. Rep. CAAM Technical Report TR12-14. Rice University; 2012. On the Global and Linear Convergence of the Generalized Alternating Direction Method of Multipliers.
- Duchi, J.; Shalev-Shwartz, S.; Singer, Y.; Chandra, T. Proceedings of the International Conference on Machine Learning. 2008. Efficient Projections onto the ℓ_1 -Ball for Learning in High Dimensions.
- Elkan, C. Proceedings of ICML 2003. 2003. Using the Triangle Inequality to Accelerate k-Means.
- Ferguson TS. A Bayesian Analysis of Some Nonparametric Problems. *Annals of Statistics*. 1973; 1:209–230.
- Forgy E. Cluster analysis of multivariate data: efficiency versus interpretability of classifications. *Biometrics*. 1965; 21:768–780.
- Fraley C. Algorithms for Model-Based Gaussian Hierarchical Clustering. *SIAM Journal on Scientific Computing*. 1998; 20:270–281.
- Frank M, Wolfe P. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*. 1956; 3:95–110.
- Gabay D, Mercier B. A dual algorithm for the solution of nonlinear variational problems via finite-element approximations. *Comp Math Appl*. 1976; 2:17–40.
- Glowinski R, Marrocco A. Sur l'approximation par éléments finis d'ordre un, et la résolution par pénalisation-dualité d'une classe de problèmes de Dirichlet non linéaires. *Rev Française d'Aut Inf Rech Oper*. 1975; 2:41–76.
- Goldfarb D, Ma S, Scheinberg K. Fast alternating linearization methods for minimizing the sum of two convex functions. *Mathematical Programming*. 2012:1–34.
- Goldstein, T.; O'Donoghue, B.; Setzer, S. Tech. Rep. cam12-35. University of California; Los Angeles: 2012. Fast Alternating Direction Optimization Methods.
- Goldstein T, Osher S. The Split Bregman Method for L1-Regularized Problems. *SIAM Journal on Imaging Sciences*. 2009; 2:323–343.
- Gordon, A. Classification. 2. London: Chapman and Hall/CRC Press; 1999.
- Gower JC, Ross GJS. Minimum Spanning Trees and Single Linkage Cluster Analysis. *Journal of the Royal Statistical Society Series C (Applied Statistics)*. 1969; 18:54–64.

- Grant, M.; Boyd, S. Graph implementations for nonsmooth convex programs. In: Blondel, V.; Boyd, S.; Kimura, H., editors. Recent Advances in Learning and Control. Springer-Verlag Limited; 2008. p. 95-110. Lecture Notes in Control and Information Sciences http://stanford.edu/~boyd/graph_dcp.html
- Hartigan, J. Clustering Algorithms. New York: Wiley; 1975.
- Hocking, T.; Vert, JP.; Bach, F.; Joulin, A. Clusterpath: an Algorithm for Clustering using Convex Fusion Penalties. In: Getoor, L.; Scheffer, T., editors. Proceedings of the 28th International Conference on Machine Learning (ICML-11). New York, NY, USA: ACM; 2011. p. 745-752. ICML '11
- Hoefling H. A Path Algorithm for the Fused Lasso Signal Approximator. Journal of Computational and Graphical Statistics. 2010; 19:984–1006.
- Johnson S. Hierarchical clustering schemes. Psychometrika. 1967; 32:241–254. [PubMed: 5234703]
- Kaufman, L.; Rousseeuw, P. Finding Groups in Data: An Introduction to Cluster Analysis. New York: Wiley; 1990.
- Lance GN, Williams WT. A General Theory of Classificatory Sorting Strategies: 1. Hierarchical Systems. The Computer Journal. 1967; 9:373–380.
- Lindsten, F.; Ohlsson, H.; Ljung, L. Tech. rep. Linköpings universitet; 2011. Just Relax and Come Clustering! A Convexification of k-Means Clustering.
- Lloyd S. Least squares quantization in PCM. Information Theory, IEEE Transactions on. 1982; 28:129–137.
- Luxburg U. A tutorial on spectral clustering. Statistics and Computing. 2007; 17:395–416.
- MacQueen, J. Proc Fifth Berkeley Symp on Math Statist and Prob. Vol. 1. Univ. of Calif. Press; 1967. Some methods for classification and analysis of multivariate observations; p. 281-297.
- McLachlan, G. Finite Mixture Models. Hoboken, New Jersey: Wiley; 2000.
- Michelot C. A finite algorithm for finding the projection of a point onto the canonical simplex of $\mathbb{R}^n$. Journal of Optimization Theory and Applications. 1986; 50:195–200.
- Mirkin, B. Mathematical Classification and Clustering. Dordrecht, The Netherlands: Kluwer Academic Publishers; 1996.
- Murtagh F. A Survey of Recent Advances in Hierarchical Clustering Algorithms. The Computer Journal. 1983; 26:354–359.
- Neal RM. Markov Chain Sampling Methods for Dirichlet Process Mixture Models. Journal of Computational and Graphical Statistics. 2000; 9:249–265.
- Rasmussen, CE. In Advances in Neural Information Processing Systems 12. MIT Press; 2000. The Infinite Gaussian Mixture Model; p. 554-560.
- Slaney M, Casey M. Locality-Sensitive Hashing for Finding Nearest Neighbors [Lecture Notes]. Signal Processing Magazine, IEEE. 2008; 25:128–131.
- Sneath PHA. The Application of Computers to Taxonomy. Journal of General Microbiology. 1957; 17:201–226. [PubMed: 13475686]
- Sørensen T. A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons. Biologiske Skrifter. 1948; 5:1–34.
- Tibshirani R. Regression Shrinkage and Selection via the Lasso. Journal of the Royal Statistical Society, Ser B. 1996; 58:267–288.
- Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K. Sparsity and smoothness via the fused lasso. Journal of the Royal Statistical Society: Series B (Statistical Methodology). 2005; 67:91–108.
- Tibshirani RJ, Taylor J. The solution path of the generalized lasso. Annals of Statistics. 2011; 39:1335–1371.
- Titterton, DM.; Smith, AFM.; Makov, UE. Statistical Analysis of Finite Mixture Distributions. Hoboken, New Jersey: John Wiley & Sons; 1985.
- Tropp J. Just relax: Convex programming methods for identifying sparse signals in noise. IEEE Transactions on Information Theory. 2006; 52:1030–1051.
- Tseng P. Applications of a Splitting Algorithm to Decomposition in Convex Programming and Variational Inequalities. SIAM Journal on Control and Optimization. 1991; 29:119–138.

- von Luxburg U. Clustering Stability: An Overview. *Found Trends Mach Learn*. 2010; 2:235–274.
- Ward JH. Hierarchical Grouping to Optimize an Objective function. *Journal of the American Statistical Association*. 1963; 58:236–244.
- Wickham, H. *ggplot2: Elegant Graphics for Data Analysis*. Springer; New York: 2009.
- Witten DM, Tibshirani R. A framework for feature selection in clustering. *J Am Stat Assoc*. 2010; 105:713–726. [PubMed: 20811510]
- Wu, R.; Wunsch, D. *Clustering*. Hoboken: John Wiley & Sons; 2009.

Figure 1.
A graph with positive weights w_{12} , w_{15} , w_{34} and all other weights $w_{jj} = 0$.

Figure 2.

Cluster path assignment: The simulated example shows five well separated clusters and the assigned clustering generated by the convex clustering algorithm under an ℓ_2 -norm. The lines trace the path of the individual cluster centers as the regularization parameter γ increases.

Figure 3. Halfmoons Example: The first and second rows show results using $k = 10$ and 50 nearest neighbors respectively. The first and second columns show results using $\varphi = 0$ and 0.5 respectively.

Figure 4.

Clustering path under the ℓ_2 norm for Mammal Data. Panel on the left (Set A) used $w_{ij} = 1$ for all i, j . Panel on the right (Set B) used $k = 5$ nearest neighbors and $\varphi = 0.5$.

Table 1

Proximal maps for common norms.

Norm	$\Omega(\mathbf{v})$	$\text{prox}_{\sigma\Omega}(\mathbf{v})$	Comment
ℓ_1	$\ \mathbf{v}\ _1$	$\left[1 - \frac{\sigma}{ v_l }\right]_+ v_l$	Element-wise soft-thresholding
ℓ_2	$\ \mathbf{v}\ _2$	$\left[1 - \frac{\sigma}{\ \mathbf{v}\ _2}\right]_+ \mathbf{v}$	Block-wise soft-thresholding
ℓ_∞	$\ \mathbf{v}\ _\infty$	$\mathbf{v} - \mathcal{P}_{\sigma\mathcal{S}}(\mathbf{v})$	$\mathcal{S}$ is the unit simplex

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Timing comparison under the ℓ_2 norm: Dense weights. Mean run times are in seconds. Different methods are listed on each row. Each column reports times for varying number of points.

	100	200	300	400	500
Subgradient	44.40	287.86	2361.84	3231.21	13895.50
AMA	16.09	71.67	295.23	542.45	1109.67
ADMM	57.82	449.68	1430.05	3432.77	6745.82

Table 3

Timing comparison under the ℓ_2 norm: Sparse weights. Mean run times are in seconds. Different methods are listed on each row. Each column reports times for varying number of points.

	100	200	300	400	500
Subgradient	6.52	37.42	161.68	437.32	386.45
AMA	1.50	2.94	4.46	6.02	7.44
ADMM	3.78	21.35	61.23	139.37	297.99