

Project and Forget: Solving Large-Scale Metric Constrained Problems

Rishi Sonthalia

*Department of Mathematics
University of Michigan
Ann Arbor, Michigan, 48104*

RSONTAL@UMICH.EDU

Anna C. Gilbert

*Departments of Mathematics and Statistics & Data Science
Yale University
New Haven, Connecticut, 06510*

ANNA.GILBERT@YALE.EDU

Editor: Suvrit Sra

Abstract

Many important machine learning problems can be formulated as highly constrained convex optimization problems. One important example is metric constrained problems. In this paper, we show that standard optimization techniques can not be used to solve metric constrained problem.

To solve such problems, we provide a general active set framework, called PROJECT AND FORGET, and several variants thereof that use Bregman projections. PROJECT AND FORGET is a general purpose method that can be used to solve highly constrained convex problems with many (possibly exponentially) constraints. We provide a theoretical analysis of PROJECT AND FORGET and prove that our algorithms converge to the global optimal solution and have a linear rate of convergence. We demonstrate that using our method, we can solve large problem instances of general weighted correlation clustering, metric nearness, information theoretic metric learning and quadratically regularized optimal transport; in each case, out-performing the state of the art methods with respect to CPU times and problem sizes.

Keywords: Large Scale Convex Optimization, Metric Constrained Optimization, Metric Learning, Correlation Clustering

1. Introduction

Given a set of dissimilarity measures amongst data points, many machine learning problems are considerably “easier” if these dissimilarity measures adhere to a metric. Furthermore, learning the metric that is most “consistent” with the input dissimilarities or the metric that best captures the relevant geometric features of the data (e.g., the correlation structure in the data) is a key step in efficient, approximation algorithms for classification, clustering, regression, and feature selection. In practice, these metric learning problems are formulated as convex optimization problems subject to metric constraints, such as the triangle inequality, on all the output variables. Because of the large number of constraints, researchers have been forced to restrict either the kinds of metrics learned or the size of the problem that can be solved. In many cases, researchers have restricted themselves to learning (weighted)

Euclidean or Mahalanobis metrics. This approach is, however, far from ideal as the inherent geometry of many data sets necessitates different types of metrics. Therefore, we need to develop optimization techniques that can optimize over the space of all metrics on a data set.

Many of the existing standard optimization methods suffer from significant drawbacks when solving metric constrained problems that hamper performance and restrict the instance size. Gradient based algorithms such as projected gradient descent (e.g., Beck and Teboulle (2009); Nesterov (1983)) or Riemannian gradient descent require a projection onto the space of all metrics which, in general, is an intractable problem. One modification of this approach is to sub-sample the constraints and then project onto the sampled set (see Nedić (2011); Polyak (2001); Wang and Bertsekas (2013); Wang et al. (2015)). For metric constrained problems, however, there are many more constraints than data points, so the condition numbers of the problems are quite high and, as a result, these algorithms tend to require a large number of iterations.

Another standard general optimization technique is based on the Lagrangian method. These algorithms augment the objective (or introduce a barrier function) by adding a term for each constraint. Examples of such methods include the interior point method, the barrier method, and the Alternating Direction Method of Multipliers. These methods run into two different kinds of problems. First, computing the gradient becomes an intractable problem because of the large number of constraints. One fix could be to sub-sample the constraints and compute only those gradients, but this approach runs into the same drawbacks as we discussed before. The other option is to incrementally update the Lagrangian, looking at one constraint at a time. Traditionally, these methods require us to cycle through all the constraints. One such method is Bregman cyclic method and we note that the requirement to examine the constraints cyclically is an aspect that is highlighted in various previous works Censor and Reich (1998); Bauschke and Lewis (2000); Censor and Zenios (1997). This is time consuming with metric constraints as we have cycle through at least $O(n^3)$ many constraints. Many applications that use this method sidestep the issue either by restricting the number of constraints and solving a heuristic problem (Davis et al., 2007), by solving smaller sized problems (Dhillon and Tropp, 2007), or by trying to parallelize the projections (Ruggles et al., 2020).

A third approach is to use traditional active set methods such as Sequential Linear (Quadratic) Programming (Palacios-Gomez et al., 1982; Boggs and Tolle, 1995) or the algorithm from Schittkowski (2009). In general, these methods maintain a set of constraints C that is assumed to be the true set of active constraints. During each iteration, they *completely* solve the sub-problem defined by the constraints in C . They then check which of the constraints in C are inactive and remove those constraints. They also search for and add new violated constraints. These methods run into the problem that each iteration is computationally expensive and, in some cases, may not be tractable because of the large number violated constraints. If the size of C is reduced, then the number of iterations becomes too large, again making the problem intractable.

One final approach is to use cutting planes. The performance of this method is heavily dependent on the cut selection process (see Dey and Molinaro (2018); Poirrier and Yu (2019) for deep discussions). The discovery of Gomory cuts (Gomory, 1960) and other subsequent methods such as branch and bound, has led to the viability of the cutting plane method for solving mixed integer linear programs. This success has, however, not transferred to other

problems. In general, if the cuts are not selected appropriately, the algorithm could take an exponential number of iterations; i.e., it might add an exponential number of constraints. To use this method, we must show for each problem that the specific cutting plane selection method results in a feasible algorithm (see Chandrasekaran et al. (2012) for an example).

We note that the only method that has found any success for metric constrained problems is the Lagrangian based method known as the cyclic Bregman method (Bregman, 1967). This was first used by Brickell et al. (2008) to solve metric constrained problems. Follow up work such as Davis et al. (2007); Veldt et al. (2019); Ruggles et al. (2020) also used this methods or variants of this methods to solve metric constrained problems. Although, we note the significant drawbacks to this approach above.

In this paper, we provide an active set algorithm, PROJECT AND FORGET, that uses Bregman projections to solve **convex optimization problems with a large number of (possibly exponentially linear inequality) constraints**. Since our algorithm is based on the cyclic Bregman method, it has the rapid rate of convergence of the Bregman method, along with all the benefits of being an active set method. This method overcomes the weaknesses of both the traditional active set methods and Bregman cyclic method. First, we overcome the drawbacks of the active set methods by allowing the introduction of new and the removal of old constraints *without having to completely solve convex programs as intermediate steps*. In particular, our algorithm examines each constraint once, before it introduces new constraints and forgets old constraints, thus allowing us to converge rapidly to the true active constraint set. We overcome the drawback of the traditional Bregman method by cycling through the current active set only. Thus, making each iteration much faster. Our new algorithm PROJECT AND FORGET, is the first iterative Bregman projection based algorithm for convex programs that does not require us to cyclically examine all the constraints.

The major contributions of our paper are as follows:

1. For the case when we have linear inequality constraints, we provide a Bregman projection based algorithm that does not need to look at the constraints cyclically. We prove that our algorithm converges to the global optimal solution and that the optimality error (L_2 distance of the current iterate to the optimal) asymptotically decays at an exponential rate. We also show that because of the FORGET step, when the algorithm terminates, the set of constraints remembered are exactly the active constraints.
2. For the case when we have general convex constraints, we provide a Bregman projection based algorithm that does not need to look at the constraints cyclically. We prove that our algorithm converges to the global optimal solution and that when we have quadratic objective function, the optimality error (L_2 distance of the current iterate to the optimal) asymptotically decays at an exponential rate. We also show that because of the FORGET step, when the algorithm terminates, the set of constraints remembered are exactly the active constraints.
3. We solve the weighted correlation clustering problem Bansal et al. (2002) on a graph with over 130,000 nodes. To solve this problem with previous methods, we would need to solve a linear program with over 10^{15} constraints. Furthermore, we demonstrate our algorithms superiority by outperforming the current state of the art in terms of CPU times.

4. We use our algorithm to develop a new one that solves the metric nearness problem Brickell et al. (2008). We show that our algorithm outperforms the current state of the art with respect to CPU time and can be used to solve the problem for non-complete graphs.
5. We also show the generality of our algorithm by using it to solve the quadratically regularized optimal transport problem. We show that, using our algorithm, we can solve this problem faster and using less memory than many standard solvers.

2. Preliminaries

We start by presenting the general version of the problem before focusing on metric constrained problems in later sections.

2.1 Convex Programming

Given a strictly convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, and a finite family of convex sets $\mathcal{F} = \{C_i\}$ we want to find the unique point $x^* \in \bigcap_i C_i =: C$ that solves the following problem.

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && \forall i, x \in C_i. \end{aligned} \tag{2.1}$$

We refer to each C_i as a *constraint set* and $C := \bigcap_i C_i$ as the *feasible region*. We shall assume that C is not empty; i.e., there is at least one feasible point. Since we have a large number of constraint sets, we access the constraint sets only through an oracle that has one of the two following separation properties.

Property 1 \mathcal{Q} is a deterministic separation oracle for a family of convex sets $\mathcal{F} = \{C_i\}$, if there exists a non-decreasing, continuous function ϕ , with $\phi(y) = 0 \iff y = 0$, such that on input $x \in \mathbb{R}^d$, \mathcal{Q} either certifies $x \in C$ or returns a list $\mathcal{L} \subset \mathcal{F}$ such that

$$\max_{\tilde{C} \in \mathcal{L}} \text{dist}(x, \tilde{C}) \geq \phi(\text{dist}(x, C)),$$

where for a point x and set B , $\text{dist}(x, B) = \inf_{w \in B} \|w - x\|$.

There are a few things that we would like to highlight about this definition. First, the list \mathcal{L} need not contain all violated constraints. That is, given x , there can be some $C_i \in \mathcal{F}$ such that $x \notin C_i$ and $C_i \notin \mathcal{L}$. In fact, there could be many such C_i . However, what we do require is that the maximum distance from x to the constraints in \mathcal{L} is at least some non-decreasing function ϕ of the distance from x to C .

Property 2 \mathcal{Q} is a random separation oracle for a family of convex sets \mathcal{F} , if there exists a lower bound $\tau > 0$, such that on input $x \in \mathbb{R}^d$, \mathcal{Q} returns a list $\mathcal{L} \subset \mathcal{F}$ such that

$$\forall \tilde{C} \in \mathcal{F}, \Pr[\tilde{C} \in \mathcal{L}] \geq \tau.$$

Remark 1 The random separation oracle need not decide whether $x \in C$.

2.1.1 LINEAR INEQUALITY CONSTRAINTS

In practice, most problems do not have the most general of convex constraints. Indeed, linear inequality constraints are common. In such a case, our constraints sets C_i are half spaces. In particular, we denote each half space by H_i instead of C_i . Additionally, for each H_i , we know that there exists $a_i \in \mathbb{R}^d$ and $b_i \in \mathbb{R}$ such that

$$H_i = \{x \in \mathbb{R}^d : \langle a_i, x \rangle \leq b_i\}.$$

Thus, if A is the matrix whose rows are given by a_i and b is the vector whose coordinates are b_i , then our feasible region C can be represented as follows:

$$C = \{x \in \mathbb{R}^d : Ax \leq b\}.$$

In this case, we can reformulate Problem 2.1 as follows.

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && Ax \leq b \end{aligned} \tag{2.2}$$

Solving this problem for general convex functions f is too monumental a task. We restrict ourselves to a rich class of functions known as Bregman functions. First, we define the generalized Bregman distance.

Definition 1 *Given a convex function $f(x) : S \rightarrow \mathbb{R}$ whose gradient is defined on S , we define its generalized Bregman distance $D_f : S \times S \rightarrow \mathbb{R}$ as $D_f(x, y) = f(x) - f(y) - \langle \nabla f(y), x - y \rangle$.*

Definition 2 *A function $f : \Lambda \rightarrow \mathbb{R}$ is called a Bregman function if there exists a non-empty convex set S such that $\bar{S} \subset \Lambda$ and the following hold:*

- (i) $f(x)$ is continuous, strictly convex on \bar{S} , and has continuous partial derivatives in S .
- (ii) For every $\alpha \in \mathbb{R}$, the partial level sets $L_1^f(y, \alpha) := \{x \in \bar{S} : D_f(x, y) \leq \alpha\}$ and $L_2^f(x, \alpha) := \{y \in S : D_f(x, y) \leq \alpha\}$ are bounded for all $x \in \bar{S}, y \in S$.
- (iii) If $y_n \in S$ and $\lim_{n \rightarrow \infty} y_n = y^*$, then $\lim_{n \rightarrow \infty} D_f(y^*, y_n) = 0$.
- (iv) If $y_n \in S, x_n \in \bar{S}, \lim_{n \rightarrow \infty} D_f(x_n, y_n) = 0, y_n \rightarrow y^*$, and x_n is bounded, then $x_n \rightarrow y^*$.

We denote the family of Bregman functions by $\mathcal{B}(S)$. We refer to S as the zone of the function and we take the closure of the S to be the domain of f . Here \bar{S} is the closure of S .

This class of function includes many natural objective functions, including entropy $f(x) = -\sum_{i=1}^n x_i \log(x_i)$ with zone $S = \mathbb{R}_+^n$ (here f is defined on the boundary of S by taking the limit) and $f(x) = \frac{1}{p} \|x\|_p^p$ for $p \in (1, \infty)$. The ℓ_p norms for $p = 1, \infty$ are not Bregman functions but can be made Bregman functions by adding a quadratic term. That is, $f(x) = c^T x$ is a not Bregman function, but $c^T x + x^T Q x$ for any positive definite Q is a Bregman function.

Definition 3 *We say that a hyper-plane H_i is strongly zone consistent with respect to a Bregman function f and its zone S , if for all $y \in S$ and for all hyper-planes H , parallel to H_i that lie in between y and H_i , the Bregman projection of y onto H lies in S instead of in \bar{S} .*

In addition to the restrictions on the functions for which we can solve Problem 2.2, we will need the assumption that all hyper-planes in \mathcal{H} (our family of half spaces) are strongly zone consistent with respect to $f(x)$. This assumption is used to guarantee that when we do a projection the point we project onto lies within our domain. This is also not too restrictive. For example, all hyper-planes are strongly zone consistent with respect to the objective functions $f(x) = 0.5\|x\|^2$ and $f(x) = -\sum_i x_i \log(x_i)$. The final assumption, that we mentioned earlier, is that C is non-empty. This is needed to ensure the algorithm converges.

2.1.2 GENERAL CONVEX CONSTRAINTS

While general convex constraints do not often appear in practical problems, such a formulation is of theoretical interest. When our constraints are simply convex sets rather than linear ones, we must adjust our algorithmic approach and, as a result, our assumptions about the objective function $f(x)$ and the constraints will be slightly different as compared to the linear case. The problem that we are interested in is

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && x \in C_i \quad \forall C_i. \end{aligned} \tag{2.3}$$

Before we state our assumptions, we need the following definitions.

Definition 4 *A function f is Legendre, if f is a closed proper map, $\text{int}(\text{dom}f) \neq \emptyset$, and $\lim_{t \downarrow 0} \langle \nabla f(x + t(y - x)), y - x \rangle = -\infty$ for all $x \in \partial(\text{dom}f)$ and for all $y \in \text{int}(\text{dom}f)$.*

Definition 5 *A function f is strictly convex, if f is twice differentiable everywhere and its Hessian is positive definite everywhere.*

Definition 6 *A function f is co-finite if $\lim_{r \rightarrow \infty} f(rx)/r = \infty$ for all $x \in \text{dom}f$.*

For this version of the problem we will assume that f is a strictly convex, co-finite, Legendre function. Note that these are strong conditions and rule out some objective functions such as Burg's entropy $f(x) = \sum_i \log(x_i)$. They do still, however, allow a rich class of functions. Some examples can be seen below:

1. $f(x) = 0.5\|x\|^2$ on \mathbb{R}^n .
2. $f(x) = \sum_i x_i \log(x_i) - x_i$ (Boltzman/Shannon entropy) on \mathbb{R}_+^n .
3. $f(x) = -\sum_i \sqrt{1 - x_i^2}$ (Hellinger distance) in $[-1, 1]^n$.
4. $f(x) = \sum_i x_i \log(x_i) + (1 - x_i) \log(1 - x_i)$ (Fermi/Dirac entropy) on $[0, 1]^n$.
5. $f(x) = \begin{cases} \frac{1}{2}x^2 + 2x + \frac{1}{2} & x \leq -1 \\ -1 - \log(-x) & -1 \leq x < 0 \end{cases}$.

Note that the last example Legendre function is not a Bregman function. In addition to the Legendre function assumptions, we also assume that $\text{int}(\text{dom}f) \cap C \neq \emptyset$. See Bauschke and Lewis (2000) for a discussion comparing the two different classes functions presented.

2.2 Metric Constrained Problems

The primary motivation of this paper is to solve metric constrained problems and in this section we set up such problems. Prior work such as Veldt et al. (2019); Brickell et al. (2008) define a metric constrained problem as follows. Let $x \in \mathbb{R}^{\binom{n}{2}}$ so that the entries are indexed as $x(i, j)$ for $i < j$. Let $f : \mathbb{R}^{\binom{n}{2}} \rightarrow \mathbb{R}$ be the function to optimize. Then a metric constrained problem, as defined in prior work, is the following optimization problem.

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && \forall i < j < k, x(i, j) \leq x(i, k) + x(j, k) \end{aligned} \tag{2.4}$$

In this paper, however, we will generalize this problem.

To define general metric constrained problems, we first define the metric polytope.

Definition 7 Let $MET_n \subset \mathbb{R}^{\binom{n}{2}}$ be the space of all pseudo-metrics on n points. Given a graph G the metric polytope $MET_n(G)$ is the projection of MET_n onto the coordinates given by the edges of G (i.e., we consider distances only between pairs of points that are adjacent in G).

It can be easily seen that for any $x \in \mathbb{R}^{\binom{n}{2}}$, $x \in MET_n(G)$ if and only if $\forall e \in G, x(e) \geq 0$ and for every cycle \mathcal{C} in G and $\forall e \in \mathcal{C}$, we have that

$$x(e) \leq \sum_{\tilde{e} \in \mathcal{C}, \tilde{e} \neq e} x(\tilde{e}).$$

Therefore, $MET_n(G)$ can be described as the intersection of exponentially many half-spaces. Further each constraint in $MET_n(G)$ can be indexed by a path $P(i, j)$ that connects two adjacent nodes i, j .

Remark 2 When we optimize over MET_n (or over $MET_n(G)$), we are optimizing over a much larger and more complex space than the space of Euclidean metrics or the space of all Mahalanobis metrics.

Now that we have the set over which we want to optimize, we give a general formulation for metric constrained optimization problems.

Definition 8 Given a strictly convex function f , a graph G , and a finite family of half-spaces $\mathcal{H} = \{H_i\}$ such that $H_i = \{x : \langle a_i, x \rangle \leq b_i\}$, we seek the unique point $x^* \in \bigcap_i H_i \cap MET(G) =: \mathcal{C}$ that minimizes f . That is, if we set A to be the matrix whose rows are a_i and b be the vector whose coordinates are b_i we seek

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && Ax \leq b \\ & && x \in MET(G). \end{aligned} \tag{2.5}$$

The constraints encoded in the matrix A let us impose additional constraints, beyond the metric constraints. For example, in correlation clustering, the matrix A encodes $x_{ij} \in [0, 1]$. In general, we will assume that the number of additional constraints encoded in A (beyond

the metric constraints) is relatively small so that the predominant difficulty in solving these optimization problems comes from the metric constraints.

It is clear if $G = K_n$, the complete graph on n points, then $\text{MET}_n = \text{MET}(K_n)$. Therefore, the metric constrained problems from Veldt et al. (2019); Brickell et al. (2008) are simply special cases of our problem. There is, however, an equivalence beyond this as well when f does not depend on some of the coordinates of x . For example, suppose $n = 3$, $x = (x(1, 2), x(1, 3), x(2, 3))$, and $f(x) = x(1, 2) + x(2, 3)$. Then, if we optimize f over MET_3 or $\text{MET}(G)$ where G is the path $1 - 2 - 3$, then the solutions are the same. This is formalized by the following proposition.

Proposition 1 *Let $f(x)$ be a function whose values only depends on the values $x(i, j)$ for $e = (i, j) \in G$ and consider the following constrained optimization problem.*

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && x \in \text{MET}(K_n). \end{aligned} \tag{2.6}$$

Let π be the projection from $\text{MET}(K_n)$ to $\text{MET}(G)$ and let $\tilde{f}(x) := f(\pi^{-1}(x))$. Then for any optimal solution x^* to the following problem

$$\begin{aligned} & \text{minimize} && \tilde{f}(x) \\ & \text{subject to} && x \in \text{MET}(G). \end{aligned} \tag{2.7}$$

we have that for all $\hat{x} \in \pi^{-1}(x^*)$, \hat{x} is an optimal solution to 2.6.

Proof Here, we see that if \tilde{x} is a minimizer of Problem 2.6 and x^* is the minimizer of Problem 2.7 then

$$f(\tilde{x}) = \tilde{f}(\pi(\tilde{x})) \geq \tilde{f}(x^*) = f(\pi^{-1}(x^*)),$$

where the middle inequality is true, since $\pi(\tilde{x})$ need not be an optimal solution to Problem 2.7. Thus, any element in $\pi^{-1}(x^*)$ is an optimal solution to Problem 2.6. ■

In these cases, we see that while we are solving the same problem, there are some practical differences. When $\text{MET}_n = \text{MET}(K_n)$, we are looking at two different representations of the same polytope; one with $O(n^3)$ constraints and the other with significantly more. In the case when f is constant for some coordinates of x , we have another practical difference; that is, we compute x^* the optimal point in $\text{MET}(G)$. It is, however, unclear how to pick the point $\hat{x} \in \pi^{-1}(x^*)$.

2.3 Projections

All of our algorithms will be based on iteratively computing Bregman projections.

Definition 9 *Given a strictly convex function f , a closed convex set Y , and a point y , the Bregman projection of y onto Y with respect to D_f is a point $x^* \in \text{dom}(f)$ such that*

$$x^* = \arg \min_{x \in Y \cap \text{dom}(f)} D_f(x, y).$$

In the case we have linear inequality constraints, we project onto the boundary of the half space ∂H . In this case, the Bregman projection has some additional special properties.

Lemma 1 *Let x be the point that we project onto $\partial H_i = \{y \in \mathbb{R}^d : \langle y, a_i \rangle = b_i\}$, then there exists a unique x^*, θ such that $\nabla f(x^*) = \nabla f(x) + \theta a_i$ and $\langle x^*, a_i \rangle = b_i$. This unique x^* is also the Bregman projection of x on ∂H_i . Furthermore,*

1. $\theta > 0$ if and only if $\langle x, a_i \rangle > b_i$;
2. $\theta < 0$ if and only if $\langle x, a_i \rangle < b_i$;
3. $\theta = 0$ if and only if $\langle x, a_i \rangle = b_i$.

3. Project and Forget: Linear Inequalities

To set the stage for subsequent discussions, we present the general structure of our algorithm first and then detail adjustments we make for the different kinds of constraints (linear inequalities versus the more general convex constraints). It is iterative and, in general, will be run until some convergence criterion has been met. The convergence criterion depends largely on the specific application for which the algorithm is tailored. For this reason, we postpone the discussion of the convergence criterion until the applications section.

The PROJECT AND FORGET algorithm keeps track of three quantities; $x^{(\nu)}$, the vector of variables over which we optimize, $L^{(\nu)}$ a list of the constraints that the algorithm deems active, and $z^{(\nu)}$ a vector of dual variables. Each iteration of the PROJECT AND FORGET algorithm consists of three phases. In the first phase, we query our oracle \mathcal{Q} to obtain a list of constraints L . In the second phase, we merge $L^{(\nu)}$ with L to form $\tilde{L}^{(\nu+1)}$ and project onto each of the constraints in $\tilde{L}^{(\nu+1)}$ one at a time. When we do these projections, we update $x^{(\nu)}$ and $z^{(\nu)}$. Finally, in the third phase, we forget some constraints from $\tilde{L}^{(\nu+1)}$ to yield $L^{(\nu+1)}$.

Algorithm 1 General Algorithm.

```

1: function PROJECT AND FORGET( $f$  convex function)
2:    $L^{(0)} = \emptyset, z^{(0)} = 0$ . Initialize  $x^{(0)}$  so that  $\nabla f(x^{(0)}) = 0$ .
3:   while Not Converged do
4:      $L = \mathcal{Q}(x^{(\nu)})$ 
5:      $\tilde{L}^{(\nu+1)} = L^{(\nu)} \cup L$ 
6:      $x^{(\nu+1)}, z^{(\nu+1)} = \text{Project}(x^{(\nu)}, z^{(\nu)}, \tilde{L}^{(\nu+1)})$ 
7:      $L^{(\nu+1)} = \text{Forget}(z^{(\nu+1)}, \tilde{L}^{(\nu+1)})$ 
   return  $x$ 

```

3.1 Finding Violated (Metric) Constraints

The first step of the method is to find violated constraints and in this subsection we detail how to find violated metric constraints in particular (which are a special case of linear inequality constraints). In many applications, we could do this by searching through the list of constraints until we found a violated constraint. However, in our case, since $\text{MET}_n(G)$ has

exponentially many faces, we cannot list all of them, so we seek an efficient separation oracle \mathcal{Q} . That is, given a point x , the oracle efficiently return a list L of violated constraints, such that the constraints in L satisfy some properties. We will assume that \mathcal{Q} satisfies either the Property 1 or Property 2.

Algorithm 2 Finding Metric Violations.

```

1: function METRIC VIOLATIONS( $d$ )
2:    $L = \emptyset$ 
3:   Let  $d(i, j)$  be the weight of shortest path between nodes  $i$  and  $j$  or  $\infty$  if none exists.
4:   for Edge  $e = (i, j) \in E$  do
5:     if  $x(i, j) > d(i, j)$  then
6:       Let  $P$  be the shortest path between  $i$  and  $j$ 
7:       Add  $C = P \cup \{(i, j)\}$  to  $L$ 
   return  $L$ 

```

For metric constrained problems, Algorithm 2 finds violated constraints. If the metric constrained problem has additional constraints (i.e $Ax \leq b$), then we augment our oracle accordingly.

Proposition 2 METRIC VIOLATION runs $\Theta(n^2 \log(n) + n|E|)$ time and satisfies Property 1 with $\phi(y) = \frac{y}{n^{1.5}}$.

Proof The first step in METRIC VIOLATION is to calculate the shortest distance between all pairs of nodes. This can be done using Dijkstra’s algorithm in $\Theta(n^2 \log(n) + n|E|)$ time. Then, if the shortest path between any adjacent pair of vertices is not the edge connecting them, then the algorithm has found a violated cycle inequality. Note that if no such path exists, then all cycle inequalities have been satisfied and the input point x (representing distances) is within the metric polytope. Thus, we have an oracle that separates the polytope.

However, we want an oracle that also satisfies property 1. To that end, let us define the deficit of a constraint. Given a point x and a hyper-plane $H_{\mathcal{C},e}$, defined by some cycle \mathcal{C} and an edge e , the deficit of this constraint is given by

$$d(\mathcal{C}, e) = x(e) - \sum_{\tilde{e} \in \mathcal{C}, \tilde{e} \neq e} x(\tilde{e}).$$

If this quantity is positive, then x violates this constraint. In this case, the squared distance from x to this constraint is $\frac{d(\mathcal{C},e)^2}{|\mathcal{C}|}$ (i.e., we add/subtract $\frac{d(\mathcal{C},e)}{|\mathcal{C}|}$ to each edge weight of \mathcal{C}).

Now let x_{apsp} be the all pair shortest path metric obtained from x and \mathcal{L} be the list returned by the oracle. Then

$$\|x - x_{\text{apsp}}\|_2^2 = \sum_{\mathcal{C}, e \in \mathcal{L}} d(\mathcal{C}, e)^2.$$

Thus, if $H_{\tilde{\mathcal{C}}, \tilde{e}}$ is the constraint that maximizes $d(\tilde{\mathcal{C}}, \tilde{e})$, then we have that

$$\text{dist}(x, \tilde{\mathcal{C}}, \tilde{e})^2 = \frac{d(\tilde{\mathcal{C}}, \tilde{e})^2}{|\tilde{\mathcal{C}}|} \geq \frac{\|x - x_{\text{apsp}}\|_2^2}{|\tilde{\mathcal{C}}||\mathcal{L}|}.$$

Since our oracle returns at most 1 constraint per edge, we have that $|\mathcal{L}| \leq |E| \leq n^2$. This along with the fact that $|\tilde{\mathcal{C}}| \leq n$, gives us that

$$\text{dist}(x, H_{\tilde{\mathcal{C}}, \tilde{\epsilon}})^2 \geq \frac{\|x - x_{\text{apsp}}\|_2^2}{n^3}.$$

Finally, we know that $x_{\text{apsp}} \in \text{MET}(G)$. Thus, we see that

$$\|x_{\text{apsp}} - x\|_2^2 \geq \text{dist}(x, \text{MET}(G))^2 = \text{dist}(x, C)^2.$$

Putting it all together, we have that

$$\begin{aligned} \max_{\hat{C} \in \mathcal{L}} \text{dist}(x, \hat{C})^2 &\geq \text{dist}(x, H_{\tilde{\mathcal{C}}, \tilde{\epsilon}})^2 \\ &\geq \frac{\|x - x_{\text{apsp}}\|_2^2}{n^3} \\ &\geq \frac{\text{dist}(x, C)^2}{n^3}. \end{aligned}$$

Taking the square root of both sides gives the needed result. ■

This oracle is the reason we can use the version of the metric polytope with exponentially many constraints rather than the one with cubically many constraints. This oracle runs in sub-cubic time in many cases; hence, we find violated constraints quickly. Note the oracle also returns at most $O(n^2)$ constraints. Thus, if we have sub-cubically many active constraints, this version does many fewer projections.

3.2 Project and Forget Steps

The Project and Forget steps for the algorithm are presented in Algorithm 3. Let us step through the code to obtain an intuitive understanding of its behavior. Let $H_i = \{x : \langle a_i, x \rangle \leq b_i\}$ be a constraint and x the current iterate. The first step is to calculate x^* and θ . Here x^* is the projection of x onto the boundary of H_i and θ is a “measure” of how far x is from x^* . In general, θ can be any real number and so we examine two cases: θ positive or negative.

It can be easily seen from Lemma 1 that θ is negative if and only if the constraint is violated. In this case, we have $c = \theta$ because (as we will see in proof) the algorithm always maintains $z_i \geq 0$. Then on line 5, we compute the projection of x onto H_i . Finally, since we corrected x for this constraint, we add $|\theta|$ to z_i . Since each time we correct for H_i , we add to z_i , we see that z_i stores the total corrections made for H_i . On the other hand, if θ is positive, this constraint is satisfied. In this case, if we also have that z_i is positive; i.e., we have corrected for H_i before and we have over compensated for this constraint. Thus, we must undo some of the corrections. If $c = z_i$, then we undo all of the corrections and z_i is set to 0. Otherwise, if $c = \theta$ we only undo part of the correction.

For the Forget step, given a constraint $H_i \in \tilde{L}^{(\nu+1)}$, we check if $z_i^{(\nu+1)} = 0$. If so, then we have not done any net corrections for this constraint and we can forget it; i.e., delete it from $\tilde{L}^{(\nu+1)}$.

If we think of $L^{(\nu)}$ as matrix, with each constraint being a row, we see that at each iteration $L^{(\nu)}$ is a sketch of the matrix of active constraints. Hence, during each iteration we

update this sketch by adding new constraints (rows). During the Forget step, we determine which parts of our sketch are superfluous and we erase (forget) these parts (rows) of the sketch.

Algorithm 3 Project and Forget algorithms.

```

1: function PROJECT( $x, z, L$ )
2:   for  $H_i = \{y : \langle a_i, y \rangle = b_i\} \in L$  do
3:     Find  $x^*, \theta$  by solving  $\nabla f(x^*) - \nabla f(x) = \theta a_i$  and  $x^* \in H_i$ 
4:      $c_i = \min(z_i, \theta)$ 
5:      $x \leftarrow x_{new}, x_{new} \leftarrow$  such that  $\nabla f(x_{new}) - \nabla f(x) = c_i a_i$ 
6:      $z_i \leftarrow z_i - c_i$ 
   return  $x, z$ 
7: function FORGET( $z, L$ )
8:   for  $H_i = \{x : \langle a_i, x \rangle = b_i\} \in L$  do
9:     if  $z_i == 0$  then Forget  $H_i$ 
   return  $L$ 

```

In general, calculating the Bregman projection (line 3) may not have a closed form formula. See Dhillon and Tropp (2007) for a general method to perform the calculation on line 3. For example, if $f(x) = x^T Q x + r^T x + s$ where Q positive definite, then for a given hyper-plane $\langle a, x \rangle = b$ and a point x we have that

$$\theta = \frac{\langle a, x \rangle - b}{a^T Q^{-1} a}. \quad (3.1)$$

3.3 Truly Stochastic Variant

In some problems, we have constraints defined using only subsets of the data points and we may not have an oracle that satisfies Property 1. For such cases, we present a stochastic version of our algorithm. Instead of calling METRIC VIOLATION or an oracle with Property 1, we want an oracle with Property 2. This version of our algorithm is very similar to the algorithms presented in Nedić (2011); Wang et al. (2015). The major difference being that we do not need to perform a gradient descent step. Instead, we maintain the KKT conditions by keeping track of the dual variables and doing dual corrections. In practice, using PROJECT AND FORGET with the random oracle tends to produce better results than Nedić (2011); Wang et al. (2015) because we remember the active constraints that we have seen, instead of hoping that we sample them.

In some cases, we may want a more stochastic variant. With the algorithm as specified, we have to keep track of the constraints that we have seen and carefully pick which constraints to forget. We can, nevertheless, modify the Forget step to forget all constraints and obtain a truly stochastic version of the algorithm. In this version, at each iteration, we choose a random set of constraints and project onto these constraints only, independently of what constraints were used in previous iterations. We cannot, however, forget the values of the dual variables. This version is similar to that in Bauschke and Borwein (1997). However, Bauschke and Borwein (1997) only looks at the problem when we have linear equality constraints. To employ such an approach, we could modify our problem to add slack variables and change all of constraints into equality constraints, however these modifications will not yield an

equivalent problem. One of the major assumptions of Bauschke and Borwein (1997) is that the objective function is strictly convex. Thus, if we add slack variables, then we would need to modify our objective function to be strictly convex on these variables as well. This changes the problem.

3.4 Convergence Analysis: Linear Inequality Constraints

Before we can use PROJECT AND FORGET, it is crucial to establish a few theoretical properties. The previous work on the convergence of the Bregman method relies on the fact that the algorithm cyclically visits all of the constraints (Censor and Zenios, 1997). For our method, however, this is not the case and so it is not clear that the convergence results for the traditional Bregman method still apply. Fortunately, the proofs for the traditional Bregman method can be adapted in subtle ways, so that we can establish crucial theoretical properties of the PROJECT AND FORGET algorithm.

Theorem 1 *If $f \in \mathcal{B}(S)$, H_i are strongly zone consistent with respect to f , and $\exists x^0 \in S$ such that $\nabla f(x^0) = 0$, then*

1. *If the oracle \mathcal{Q} satisfies property 1 (property 2), then any sequence x^n produced by the above algorithm converges (with probability 1) to the optimal solution of problem 2.2.*
2. *If x^* is the optimal solution, f is twice differentiable at x^* , and the Hessian $H := Hf(x^*)$ is positive definite, then there exists $\rho \in (0, 1)$ such that*

$$\lim_{\nu \rightarrow \infty} \frac{\|x^* - x^{\nu+1}\|_H}{\|x^* - x^\nu\|_H} \leq \rho \quad (3.2)$$

where $\|y\|_H^2 = y^T H y$. In the case when we have an oracle that satisfies property 2, the limit in 3.2 holds with probability 1.

The proof of Theorem 1 also establishes another important theoretical property.

Proposition 3 *If a_i is an inactive constraint, then there exists an N , such that for all $n \geq N$, we have that $z_i^n = 0$. That is, after some finite time, we never project onto inactive constraints ever again.*

Corollary 1 *Under the assumptions for part 2 of Theorem 1, the sequence $z^n \rightarrow z^*$ also converges.*

These properties are important as they permit the following interpretation of our algorithm. The algorithm spends the initial few iterations identifying the active constraints from amongst a large number of constraints. This is the active set part of the algorithm. The algorithm then spends the remainder of the iterations finding the optimal solution with respect to these constraints. Empirically, we notice this phenomenon as well. At first, the error metrics decreases very slowly, while the number of constraints that are being considered grows rapidly. Eventually, we reach a point when the number of constraints that we are currently considering stabilizes, at this point the error metrics start decreasing very rapidly. An example of this phenomenon can be seen in Figure 3(b). This behavior is one of the major advantages of our method. Additionally, the ability to find the set of active constraints without having to solve the problem is another advantage of our algorithm.

Remark 3 We note that while these results show that the algorithm converges linearly, ρ is close one. Indeed, the proof bounds $\rho \leq \frac{F}{F+1}$, where F is the number of hyperplanes that the optimal solution lies on. From a heuristic perspective, it is beneficial to use only those hyperplanes that define the facets of the constraint polytope.

For the truly stochastic case, we have the following theorem instead.

Theorem 2 If $f \in \mathcal{B}(S)$, H_i are strongly zone consistent with respect to f , and $\exists x^0 \in S$ such that $\nabla f(x^0) = 0$, then with probability 1 any sequence x^n produced by the above truly stochastic algorithm converges to the optimal solution of problem 2.2. Furthermore, if x^* is the optimal solution, f is twice differentiable at x^* , and the Hessian $H := Hf(x^*)$ is positive semi-definite, then there exists $\rho \in (0, 1)$ such that with probability 1,

$$\liminf_{\nu \rightarrow \infty} \frac{\|x^* - x^{\nu+1}\|_H}{\|x^* - x^\nu\|_H} \leq \rho. \quad (3.3)$$

Because the proofs of Theorem 1 and 2 are quite technical and involve two different types of separation oracles, we split them into several parts. In Subsections 8.1 and 8.2, we prove the first part of Theorem 1 for separation oracles with property 1 and 2, respectively. In Subsection 8.3, we prove the second part of Theorem 1 (also subdividing this proof into several cases). Finally, in Subsection 8.4 we prove Theorem 2.

4. Project and Forget: General Convex Constraints

In the previous section, we detailed the PROJECT AND FORGET algorithm for half-space (or linear inequality) constraints. In this section, we use similar idea to develop the appropriate variations for general convex constraints. We draw inspiration from Dijkstra’s method and use a slightly different set of assumptions on our objective functions and constraints. These assumptions are detailed in Section 2.1.2.

4.1 Algorithm

Before we can present the algorithm we are going to need a few key definitions.

Definition 10 Given a closed convex set C , a strictly convex Legendre function, the Bregman Projection Operator P_C is the map that takes x to its Bregman projection of x onto C with respect to f .

Definition 11 Given a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, the convex conjugate $h^* : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as follows

$$h^*(x^*) = \sup_{x \in \mathbb{R}^n} x^T x^* - h(x).$$

Definition 12 The control sequence for an optimization method, is the order in which the constraints are seen. That is, on the k th iteration $i(k)$ is the index of constraint that is looked at.

Let x^n be the primal sequence of iterates and q^n an auxiliary (secondary) sequence. Let P_k denote the Bregman projection operator onto the k th constraint set. Let f^* be the convex conjugate of f . Let $i(k)$ denote the control sequence, and let $p(c, k) = \arg \max_{k' < k} i(k') = c$. We will abbreviate $p(i(k), k)$ as $p(k)$. With this notation established, the PROJECT AND FORGET algorithm for the case of general convex constraints is shown in Algorithm 4.

Algorithm 4 Project and Forget algorithms

```

1: Initialize  $q^0 = 0, L = \emptyset$ 
2: function PROJECT( $x_0 = x, q, L$ )
3:   for  $C_i \in L$  do
4:      $x^n := (P_i \circ \nabla f^*)(\nabla f(x^{n-1}) + q^{p(n)})$ 
5:      $q^n = \nabla f(x^{n-1}) + q^{p(n)} - \nabla f(x^n)$ 
   return  $x^{|L|}$ 
6: function FORGET( $x, q, L$ )
7:   for  $C_i \in L$  do
8:     if  $q^{p(i,n)} == 0$  then Forget  $C_i$ 
   return  $L$ 

```

Lines 4 and 5 of the above algorithm come from Dijkstra's method. To understand their role, note that since f is co-finite (see Definition 6) Legendre, we have that $\nabla f^* = (\nabla f)^{-1}$. Thus, on line 4, we are perturbing x^{n-1} by modifying its gradient with the auxiliary variable $q^{p(n)}$. For example, if $f(x) = x^T Q x$ for some positive definite matrix Q , then line 4 would be

$$x^n = P_i(Q^{-1}(Qx^{n-1} + q^{p(n)})) = P_i(x^{n-1} + Q^{-1}q^{p(n)}).$$

4.2 Convergence Analysis

Now that we have defined the above algorithm, the second major theoretical result of this paper is the following.

Theorem 3 *If f is a closed very strictly convex, co-finite, Legendre function and we are given a point $x^0 \in \text{dom} f$, then*

1. *if the oracle \mathcal{Q} satisfies either property 1 or 2, then any sequence x^n produced by the above algorithm converges to Bregman projection of x^0 on C with respect to f ; and,*
2. *if f is also a quadratic function, then for large enough ν there exists a $\rho \in (0, 1)$ such that*

$$\|x^* - x^{\nu+1}\| \leq \rho \|x^* - x^\nu\|.$$

5. Applications: Metric Constrained Problems

To demonstrate the effectiveness of our method in solving metric constrained problems, we solve large instances of two different types of such problems: metric nearness and correlation clustering. We focus on these types of problems first as they were the original motivation behind our work. ¹

1. All implementations and experiments can be found at <https://github.com/rsenthal/ProjectAndForget>.

5.1 Metric Nearness

The first and simplest form of a metric constrained problem is the metric nearness problem. Following Brickell et al. (2008), the metric nearness problem is:

given a point $x \in \mathbb{R}^{\binom{n}{2}}$, find the closest (in some ℓ_p norm) point $x^* \in \text{MET}_n$ to x .

This problem is a form of metric learning; see Brickell et al. (2008) for an application to clustering and see Gilbert and Sonthalia (2018) for an application to unsupervised metric learning. Additionally, if we further restrict to finding the closest Euclidean metric, then this problem is a well studied one. See Qi and Yuan (2014); Cayton and Dasgupta (2006), and Glunt et al. (1990) for examples of this problem. Recently this problem has also been looked at from a discrete setting. Gilbert and Jain (2017) and Fan et al. (2018) looked to solve the problem with the fewest possible changes. Following this, Fan et al. (2020) generalized the problem to finding the closest point in $\text{MET}(G)$ instead of MET_n .

We use ℓ_2 version of this problem to demonstrate that standard solvers have significant drawbacks on large scale metric constrained problems while PROJECT AND FORGET handles these problems easily. In particular, in addition to comparing against the cyclic Bregman used in Brickell et al. (2008), we compare against commercial solvers CPLEX (ILOG, 2021) and Mosek (ApS, 2021), ADMM based solvers OSQP (Stellato et al., 2020), SCS (O’Donoghue et al., 2019) and COSMO (Garstka et al., 2019), operator splitting and interior point solvers Ipopt (Nocedal et al., 2009), ProxSDP (Souto et al., 2020) and ECOS (Domahidi et al., 2013), and active set solver SLSQP (Kraft, 1988). We also use Mosek as an active set solver (MASS) as follows. We go through all $O(n^3)$ constraints and find the subset S of the violated constraints. We then use Mosek to solve the problem on S . We then remove the inactive constraints and add in the new violated constraints and use Mosek to solve the problem again and iterate.

5.1.1 EXPERIMENTAL SET UP:

Before we see the experimental results, let us look at the experimental set up more closely.

Data. For this experiment, we will generate three different types of synthetic data. We will refer to these as Type I, Type II, and Type III data.

For Type I data, we generate random weighted complete graphs with Gaussian weights. For Type II data, for each edge e we set $w(e) = 1$ with probability 0.8 and $w(e) = 0$ with probability 0.2. For Type III data, we let u_{ij} be sampled from the uniform distribution on $[0, 1]$ and v_{ij} from a standard normal, then the weight for an edge $e = ij$ is given by

$$w_{ij} = \lceil 1000 \cdot u_{ij} \cdot v_{ij}^2 \rceil.$$

Implementation Details. We implemented the algorithm from Brickell et al. (2008). We made a small modification that improves the running time. In Brickell et al. (2008), it is recommended that we store the dual variable z as a sparse vector. However, as we do not want the overhead of handling sparse vectors, we store z as a dense vector.

PROJECT AND FORGET, as presented Algorithm 5, was implemented with two modifications. As we can see from Algorithm 2, when the oracle finds violated constraints, it looks at each edge in G and then decides whether there is a violated inequality with that edge. It

is cleaner, in theory, to find all such violated constraints at once and then do the project and forget steps. It is, however, much more efficient in practice to do the project and forget steps for a single constraint as we find it. This approach also helps cut down on memory usage. Once our oracle returns a list of constraints (note we have already projected onto these once), we project onto our whole list of constraints again. Thus, for the constraints returned by the oracle, we project onto these constraints twice per iteration. Note this does not affect the convergence results for the algorithm.

Algorithm 5 Pseudo-code for the implementation for Metric Nearness.

```

1:  $L^0 = \emptyset, z^0 = 0$ . Initialize  $x^0$  so that  $\nabla f(x^0) = 0$ .
2: while Not Converged do
3:   Let  $d(i, j)$  be the weight of shortest path between nodes  $i$  and  $j$  or  $\infty$  if none exists.
4:    $L = \emptyset$ 
5:   for Edge  $e = i, j \in E$  do
6:     if  $x(i, j) > d(i, j)$  then
7:       Let  $P$  be the shortest path between  $i$  and  $j$ .
8:       Let  $C = P \cup \{(i, j)\}$ .
9:       Project onto  $C$  and update  $x, z$ .
10:      if  $z_C \neq 0$  then
11:        Add  $C$  to  $L$ .
12:       $\tilde{L}^{\nu+1} = L^\nu \cup L$ 
13:       $x^{\nu+1}, z^{\nu+1} = \text{Project}(x^\nu, z^\nu, \tilde{L}^{\nu+1})$ 
14:       $L^{\nu+1} = \text{Forget}(\tilde{L}^{\nu+1})$ 
return  $x$ 

```

The solvers CPLEX, Mosek, OSQP, SCS, COSMO, Ipopt, ProxSDP, ECOS, and MASS all were accessed via Julia’s JuMP library (Dunning et al., 2017). SLSQP was interfaced with using Scipy. We do not compare against Schittkowsky (2009) as their code is not public.

Convergence criterion. When we presented the algorithm in Section 3, we did not specify a convergence criterion because we wanted the criterion to be application dependent. The convergence criterion for PROJECT AND FORGET and Brickell et al. (2008) is as follows. One variant of the metric nearness problem is the decrease only variant, in which we are not allowed to increase the distances and must only decrease them. This problem can be solved in $O(n^3)$ time by calculating the all pairs shortest path metric (Gilbert and Jain, 2017). Given x^n as input, let \hat{x}^n be the optimal decrease only metric. For PROJECT AND FORGET, and Brickell et al. (2008), we ran these experiments until $\|\hat{x}^n - x^n\|_2 \leq 10^{-10}$. For convenience, given x^n let $D(x^n) = \|\hat{x}^n - x^n\|_2$. Note that $D(x^n)$ is an upper bound for how far the current iterate is from the polytope. We can use the information returned by our oracle to compute $D(x^n)$ in quadratic time, as our oracle computes \hat{x}^n . For these experiments, however, at the end of each iteration, we rerun the computation. We do this for fairness, as this extra computation has to be added to Bregman’s cyclic method.

For CPLEX, Mosek, OSQP, SCS, COSMO, Ipopt, ProxSDP, ECOS, SLSQP, we let the convergence criterion be the default criterion. For MASS, we ran it until the constraint set converged.

Comparison Statistics We compare the solvers on two different kinds of measures. The first is the convergence details of each solver. To determine convergence, we look at two statistics. First, we look at the relative objective difference (ROD), which is given by

$$ROD = \frac{\text{OtherSolverObjective} - \text{ProjectForgetObjective}}{\text{ProjectForgetObjective}}.$$

Second, if x is the solution returned by PROJECT AND FORGET and \tilde{x} is the solution returned by one of the other solvers, then the feasibility difference (FD) is given by

$$FD = D(\tilde{x}) - D(x).$$

For both metrics, if the quantities are positive, then PROJECT AND FORGET does better.

The second measure we compare is the time taken to solve the optimization problem. Here we use the solve time reported by the solvers. Note that this does not include the interface time and in many cases the interface time could be significant.

5.1.2 RESULTS

As we can see from Table 1, standard solvers run out memory or start taking too long extremely rapidly as a function of instance size. In fact, all times reported for the standard solvers is the solve time returned by the optimizer and does not include the interface time. In many cases, such as ProxSDP, the interface time could be multiple hours. On the other hand, while initially Brickell et al. (2008) is faster than PROJECT AND FORGET as n gets larger, PROJECT AND FORGET starts to dominate. Thus, showing that PROJECT AND FORGET is the only viable algorithm to solve large metric constrained problems.

Algorithm	Number of Nodes									
	100	200	300	400	500	600	700	800	900	1000
Ours	13.5	32.7	85.1	170	271	458	720	983	1356	1649
Cyclic Bregman	1.77	10.5	47.1	141	322	558	910	1472	2251	3167
Mosek	11.7	542	Out of Memory							
SCS	1632	19466	Timed Out							
OSQP	64.5	3383	Timed Out							
ProxSDP	353	684	Timed Out							
Ipopt	2792	Timed Out								
ECOS	597	Timed Out								
CPLEX	Out of Memory									
SLSQP	Timed Out									
COSMO	Timed Out									

Table 1: Table comparing PROJECT AND FORGET against a variety of different solvers to solve the Metric Nearness problem for Type 1 graphs in terms of time taken in seconds. All experiments were run on a Computer with 52 GB of memory. All times reported are averaged over 5 instances.

Solver		n									
		100	200	300	400	500	600	700	800	900	1000
Cyclic Bregman	<i>ROD</i>	-3e-13	-4e-14	4e-15	-2e-14	6e-15	-8e-16	1e-15	-4e-15	5e-15	2e-17
	<i>FD</i>	-6e-12	5e-13	4e-13	-3e-12	1e-12	3e-12	5e-12	3e-12	-2e-12	7e-12
		n									
		10	20	30	40	50	60	70	80	90	100
Mosek	<i>ROD</i>	-4e-4	2e-9	1e-9	6e-9	7e-9	1e-8	2e-8	2e-8	2e-8	2e-8
	<i>FD</i>	2e-11	2e-11	1e-9	7e-10	3e-10	9e-10	1e-9	9e-10	8e-10	7e-10
OSQP	<i>ROD</i>	-4e-4	5e-7	-9e-7	-2e-7	-7e-8	1e-7	-1e-8	7e-8	4e-8	5e-8
	<i>FD</i>	1e-3	2e-3	1e-3	9e-4	8e-4	5e-4	9e-4	1e-3	8e-4	1e-3
SCS	<i>ROD</i>	-4e-4	-3e-8	8e-7	-8e-6	-5e-7	-1e-8	-2e-8	3e-9	-7e-6	1e-9
	<i>FD</i>	7e-5	1e-4	5e-3	6e-3	5e-4	9e-5	4e-5	4e-5	0.09	4e-5

Table 2: Convergence statistics for the metric nearness problem for the different solvers. Each value is the average over 10 trials.


To make sure that we are running all of the algorithms to the same level of convergence, we take the five best solvers and check their convergence statistics. That is, we check the convergence statistics for commercial solver Mosek, for ADMM based solvers OSQP and SCS, and for operator splitting solver ProxSDP.

Each data point in Table 2 is average over 10 trials. When comparing against Brickell et al. (2008), we let n range from a 100 to 1000, i.e., the same values as in Table 1. Table 2 then shows that both method have similar levels of convergence. Thus, since PROJECT AND FORGET is faster; it is the superior algorithm.

To compare against the rest of the solvers, as they simply cannot solve the problem for larger values of n , we let n range from 10 to 100. Here, we see that PROJECT AND FORGET has much smaller feasibility error than the other solvers. Solvers such as SCS and QSQP have a feasibility error of about 10^{-3} to 10^{-5} , however, we run PROJECT AND FORGET until the feasibility error is smaller than 10^{-10} , which is many orders of magnitude smaller. For ProxSDP, *ROD* was consistently around 1 and *FD* was consistently greater than 2. We conclude that these solvers have not converged.

The only solver, other than PROJECT AND FORGET, that consistently converges is Mosek, which has higher objective values than PROJECT AND FORGET in all but one case. The active set method MASS was run until the the feasibility error was at most 1e-10 or the constraint set was stable, in practice we found that the constraint set stabilized first and hence the solver has similar convergence statistics to Mosek.

As we can see, the cyclic Bregman method is the one that has the closest performance to PROJECT AND FORGET. Therefore, we want to take a closer look at how *ROD* and *FD* evolve for the two algorithms. Figure 1, shows a plot for both of these statistics as a function of number of projections. Note we did not compute these quantities after every projection but only after each iteration of the algorithm. As we can see from Figure 1,



(a) Plot for Relative Objective Difference (ROD) versus number of projections for PROJECT AND FORGET and Bregman's cyclic method (b) Plot for Feasibility Difference (FD) versus number of projections for PROJECT AND FORGET and Bregman's cyclic method

Figure 1: Plot for ROD and FD versus number of projections for PROJECT AND FORGET and Bregman's cyclic method.

PROJECT AND FORGET uses far fewer projections. In fact, PROJECT AND FORGET uses fewer projections than the cyclic Bregman method does in one iteration. Thus, highlighting the great advantage obtained by the forget step.

Figure 1 also highlights another aspect of the algorithm from its analysis in Proposition 3. That is, the algorithm spends the first few iterations finding the correct set of constraints. Once it has found the correct set of constraints, it converges to the correct solution very quickly (with an exponential decay in the error). This is seen in Figure 1, where sometime between 10^6 and 10^7 projections, the algorithm seems to have found the correct active set of solutions.

We further tested our method against the cyclic Bregman method on Type II and Type III data. In this case, we got the running times shown in Figure 2. For this we relaxed the convergence criteria to being within 1 of the closest decrease only metric solution. One thing we learn from relaxing the convergence criteria is that cyclic Bregman method outperforms PROJECT AND FORGET for larger values of n . This is because, that PROJECT AND FORGET only focuses on the set of the active constraints. Thus, the more stringent the convergence criterion, the better our method does compared to the standard cyclic Bregman method used in Brickell et al. (2008).

Finally, we must note that when we use the cycle constraints to represent $MET(K_n)$, there is a significant amount of dependency amongst the constraints. That is, many constraints are not independent from other constraints. We may suspect that our oracle finds a lot of active constraints that have linear dependencies and so $L^{(\nu)}$ becomes saturated with dependent constraints. It is surprising, however, that after we solve the problem, $L^{(\nu)}$, more often than not, only contains 3-cycle constraints. We also see that the size of $L^{(\nu)}$ is about $O(n^2)$.

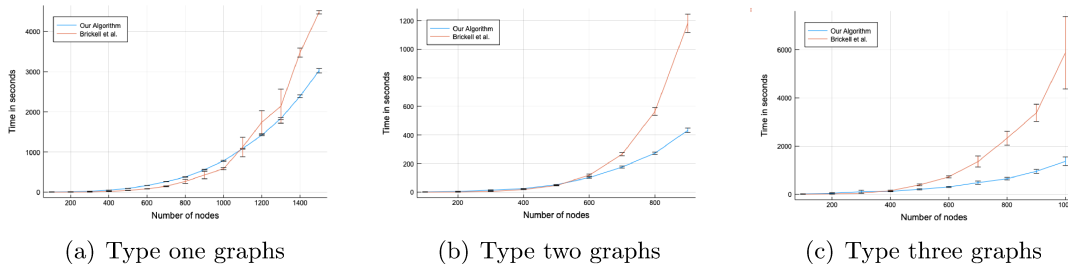


Figure 2: The red line is the mean running time for the algorithm from Brickell et al. (2008). The blue line is the running mean time for our algorithm. All computations were done on a machine with 4 physical cores, each with 13 GB of RAM.

5.2 Weighted Correlation Clustering on General Graphs.

For the correlation clustering problem, we are given a graph $G = (V, E)$ (not necessarily complete) and two non-negative numbers $w^+(e)$ and $w^-(e)$ for each edge e . These numbers indicate the level of similarity and dissimilarity between the end points of e . The goal of correlation clustering is to partition the nodes into clusters so as to minimize some objective function. The most common objective is $\sum_{e \in E} w^+(e)x(e) + w^-(e)(1 - x(e))$, where $x(e) \in \{0, 1\}$ indicates whether the end points of the edge e belong to different clusters. This problem is NP-hard and many different approximation algorithms and heuristics have been developed to solve it. The best approximation results (Charikar et al., 2005; Emanuel and Fiat, 2003), however, are obtained by rounding the solution to the following relaxed linear problem

$$\begin{aligned} & \text{minimize} && \sum_{e \in E} w^+(e)x(e) + w^-(e)(1 - x(e)) \\ & \text{subject to} && x(i, j) \leq x(i, k) + x(k, j) \quad i, j, k = 1, \dots, n \\ & && x(i, j) \in [0, 1] \quad i, j = 1, \dots, n. \end{aligned} \quad (5.1)$$

Special cases, such as when the weights are ± 1 and $G = K_n$, have faster algorithms for the same approximation ratio (Ailon et al., 2005).

The LP formulation for correlation clustering in Equation 5.1 has $\Theta(n^3)$ constraints. Hence, solving the LP for large n becomes infeasible quickly in terms of both memory and time. Veldt et al. (2019) showed that for instances with $n \approx 4000$, standard solvers such as Gurobi ran out of memory on a 100 GB machine. On the other hand, Veldt et al. (2019) developed a method using which they can feasibly solve the problem for $n \approx 11000$. To do so, they transform Problem 5.1 into Problem 5.2. To do this transformation, for $e \in E$, we define $\tilde{w}(e) = |w^+(e) - w^-(e)|$. For $e \notin E$, we let $\tilde{w}(e) = 0$. W is defined to be a diagonal matrix whose entries are given by \tilde{w} . Then, we define d as follows

$$d_e = \begin{cases} 1 & w^-(e) > w^+(e) \\ 0 & \text{otherwise} \end{cases}.$$

And the transformed problem is as follows.

$$\begin{aligned} & \text{minimize} && \tilde{w}^T |x - d| + \frac{1}{\gamma} |x - d|^T W |x - d|. \\ & \text{subject to} && x \in \text{MET}(K_n). \end{aligned} \quad (5.2)$$

For general γ , the solution to Problem 5.2 approximates the optimal solution to 5.1. However, for large enough γ it has been shown that the two problems are equivalent (Veldt et al., 2019). As we will see for our instances, Problem 5.2 is at most a 2 approximation to problem 5.1, which is an $O(\log(n))$ approximation of the NP-hard correlation clustering problem.

Finally, we also relax the condition that $x \in \text{MET}_n$ to $x \in \text{MET}(G)$. The formulation of the LP that we solve is as follows:

$$\begin{aligned} \text{minimize} \quad & \tilde{w}^T f + \frac{1}{\gamma} f^T \cdot W \cdot f \\ \text{subject to} \quad & x \in \text{MET}(G) \\ & f(i, j) = |x(i, j) - d(i, j)|, \quad (i, j) \in E. \end{aligned} \tag{5.3}$$

While it seems like we have relaxed the problem, however, since our objective function does not depend on the value of x for the edges not in G we see that the two problems are equivalent due to Proposition 1.

Remark 4 *When Veldt et al. (2019) tested their solver against Gurobi, they did so in an active set manner. That is, they found a set of violated constraints, fed this set into Gurobi, and solved the problem with this subset of the constraints. They then took the solution from Gurobi and found constraints that the current solution violated and added these constraints and repeated. They did this until Gurobi solved the problem. We view this convoluted process as yet more evidence for standard active set methods not being feasible at large scales.*

5.2.1 EXPERIMENTAL SET UP

Before looking at the results, let us see the experimental set up.

Data. We solve the problem for two different types of graphs; dense and sparse.

For dense graphs, we use four graphs from the Stanford sparse network repository. Then, following Veldt et al. (2019), we use the method from Wang et al. (2013) to convert these graphs into instances of weighted correlation clustering on the complete graph. We compare our method against Ruggles et al. (2020), a parallel version of Veldt et al. (2019), in terms of running time, quality of the solutions, and memory usage.

For much real-world data, the graph G is larger than our previous experiments but is also sparse. Since the weighting of the edges does not affect the size of the linear program that needs to be solved, we tested our algorithm on sparse signed graphs to get an estimate of the running time for the algorithm. The two graphs used for this experiment are much bigger instances than our previous experiments and have 82,140 nodes and 131,828 nodes, respectively.

Implementation Details. For the case when $G = K_n$, in addition to the modifications that were done for metric nearness experiment, we made two more modifications to PROJECT AND FORGET. First, we did the project step and the forget step one additional time per iteration. Second, we parallelized the oracle by running Dijkstra’s algorithm in parallel. The pseudo-code for this version of Algorithm 1 can be seen in Algorithm 6. For the sparse version, we also used the parallel version of the oracle and did the project and the forget step 75 times per iteration.

Note that for both experiments, the additional constraints that were introduced due to the transformation were all projected onto once per iteration and never forgotten. The

Algorithm 6 Pseudo-code for the implementation for CC for the dense case.

```

1:  $L^0 = \emptyset, z^0 = 0$ . Initialize  $x^0$  so that  $\nabla f(x^0) = 0$ .
2:  $L_a$  is the list of additional constraints.  $z_a^0 = 0$  (dual for additional constraints)
3: while Not Converged do
4:   Let  $d(i, j)$  be the weight of shortest path between nodes  $i$  and  $j$  or  $\infty$  if none exists.
   This is found using a parallel algorithm.
5:    $L = \emptyset$ 
6:   for Edge  $e = i, j \in E$  do
7:     if  $x(i, j) > d(i, j)$  then
8:       Let  $C = P \cup \{(i, j)\}$ . Where  $P$  be the shortest path between  $i$  and  $j$ .
9:       Project onto  $C$  and update  $x, z$ .
10:      if  $z_C \neq 0$  then
11:         $C$  to  $L$ .
12:    $L^\nu \leftarrow L^\nu \cup L$ 
13:   for  $i = 1, 2$  do
14:      $x^\nu, z^\nu \leftarrow \text{Project}(x^\nu, z^\nu, L^\nu)$ 
15:      $L^\nu \leftarrow \text{Forget}(L^\nu)$ 
16:    $x^\nu, z_a^\nu \leftarrow \text{Project}(x^\nu, z_a^\nu, L_a)$ 
17:    $x^{\nu+1} = x^\nu, L^{\nu+1} = L^\nu, z^{\nu+1} = z^\nu, z_a^{\nu+1} = z_a^\nu$ ,
return  $x$ 

```

pseudo-code for this version can be seen in Algorithm 7. We used the implementation provided by the authors of Veldt et al. (2019) to run the experiments for their algorithm.

Algorithm 7 Pseudo-code for the implementation for CC for the sparse case.

```

1:  $L^0 = \emptyset, z^0 = 0$ . Initialize  $x^0$  so that  $\nabla f(x^0) = 0$ .
2:  $L_a$  is the list of additional constraints.  $z_a^0 = 0$  (dual for additional constraints)
3: while Not Converged do
4:   Let  $d(i, j)$  be the weight of shortest path between nodes  $i$  and  $j$  or  $\infty$  if none exists.
   This is found using a parallel algorithm.
5:    $L = \emptyset$ 
6:   for Edge  $e = i, j \in E$  do
7:     if  $x(i, j) > d(i, j)$  then
8:       Let  $C = P \cup \{(i, j)\}$ . Where  $P$  be the shortest path between  $i$  and  $j$ .
9:       Add  $C$  to  $L$ .
10:   $L^\nu \leftarrow L^\nu \cup L$ 
11:  for  $i = 1, \dots, 75$  do
12:     $x^\nu, z^\nu \leftarrow \text{Project}(x^\nu, z^\nu, L^\nu)$ 
13:     $L^\nu \leftarrow \text{Forget}(L^\nu)$ 
14:   $x^\nu, z_a^\nu \leftarrow \text{Project}(x^\nu, z_a^\nu, L_a)$ 
15:   $x^{\nu+1} = x^\nu, L^{\nu+1} = L^\nu, z^{\nu+1} = z^\nu, z_a^{\nu+1} = z_a^\nu$ ,
return  $x$ 

```

Calculating the approximation ratio. Let \hat{x} be the optimal solution to 5.2, then if we let $R = \frac{\hat{x}^T W \hat{x}}{2\gamma \tilde{w}^T \hat{x}}$, by Veldt et al. (2019), we have that \hat{x} is an $\frac{1+\gamma}{1+R}$ approximation to the optimal solution of 5.1. This is the formula we used to calculate the approximation ratios reported in Tables 3 and 4. For our experiments we used $\gamma = 1$. Note that since the entries of \hat{x} are non negative, and the entries of W and \tilde{w} are non-negative, we have that $R > 0$. Thus, our approximation ratio is at most 2.

Convergence criterion. We ran the experiment until the maximum violation of a metric inequality was at most 0.01. However, the two algorithms, Ruggles et al. (2020) and ours, have different metric constraints. Specifically Ruggles et al. (2020) only uses all constraints that come from 3 cycles, whereas we use all cycle constraints. Theoretically both sets of constraints define the same polytope, but practically there is a difference. Thus, in practice our algorithm was run to a slightly greater level of convergence than the one from Ruggles et al. (2020).

5.2.2 RESULTS

Graph	Time (s)		Opt Ratio		Avg. mem. / iter. (GiB)		
	n	Ours	Ruggles et al.	Ours	Ruggles et al.	Ours	Ruggles et al.
CAGrQc	4158	2098	5577	1.33	1.38	4.4	1.3
Power	4941	1393	6082	1.33	1.37	5.9	2
CAHepTh	8638	9660	35021	1.33	1.36	24	8
CAHepPh	11204	71071	135568	1.33	1.46	27.5	15

Table 3: Table comparing PROJECT AND FORGET against Ruggles et al. (2020) in terms of time taken, quality of solution, and average memory usage when solving the weighted correlation clustering problem on dense graphs.

For the case of dense graphs, we see from Table 3 that our algorithm takes less time to obtain a better approximation ratio, but requires more memory per iteration. Thus, demonstrating the superiority of our method in terms of CPU time. Our algorithm requires more memory because the initial few iterations find a large number of constraints. Later, the algorithm forgets these constraints until the number of constraints stabilizes at a reasonable level. Hence, our initial memory usage is much larger than our later memory usage. To see how the number of constraints found by the oracle evolves, we plot the number of constraints found by the oracle and the number of constraints after the forget step in Figure 3(a). As we can see, after the initial few iterations, the number constraints found sharply reduces and the method has found the true set of active constraints by the 15th iteration. Figure 3(b) also shows us, as expected, *the exponential decay of the maximum violation of a metric constraint*.

Figure 3(b), also highlights another aspect our algorithm. That is, for the initial few iterations the error statistics do not decrease. In fact, we have experimentally seen that the error statistics may actually increase for the first few iterations. Proposition 3, which tells us that our algorithm spends the initial few iterations finding the active constraint set, explains

this phenomenon. During this time, the algorithm makes minimal progress towards reducing the error statistics. However, once we have found the active constraints, Theorem 4 is now applicable and we have an exponential decay of the error statistics.

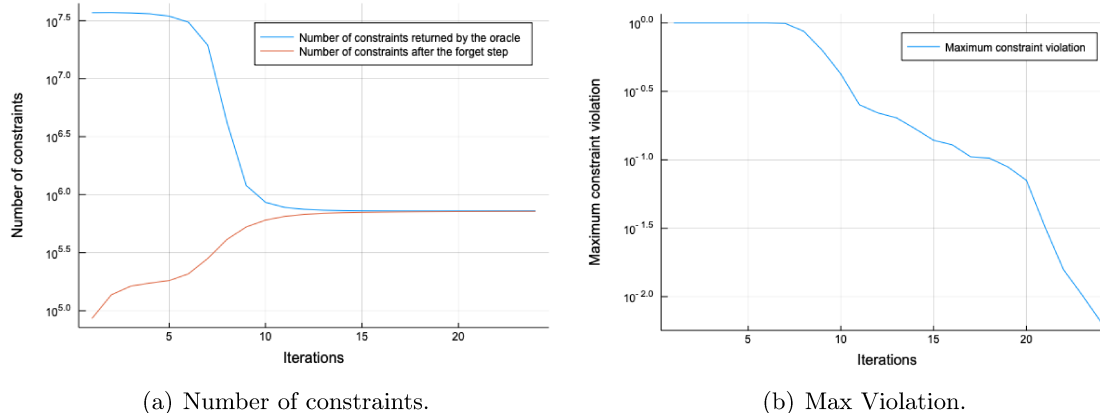


Figure 3: Plots showing the number of constraints returned by the oracle, the number of constraints after the forget step, and the maximum violation of a metric constraint when solving correlation clustering on the Ca-HepTh graph

Remark 5 *Figure 3(a) highlights the crucial difference between our method and standard active set methods. Standard active set methods would have to initially solve the convex optimization problem with 10^8 constraints. This would take a very long time. However, using PROJECT AND FORGET, we only need to compute projections onto each constraint once before we can forget constraints. Thus, we forget constraints more frequently and much earlier.*

Graph	n	# Constraints	Time	Opt Ratio	# Active Constraints	Iters.
Slashdot	82140	5.54×10^{14}	46.7 hours	1.78	384227	145
Epinions	131,828	2.29×10^{15}	121.2 hours	1.77	579926	193

Table 4: Time taken and quality of solution returned by PROJECT AND FORGET when solving the weighted correlation clustering problem for sparse graphs. The table also displays the number of constraints the traditional LP formulation would have.

For sparse graphs, even if we use Ruggles et al. (2020), based on the average time it took for a single iteration for the CA-HepPh graph, it would take Ruggles et al. (2020) an estimated two days to complete a single iteration, for a graph with $n \approx 80,000$. This is because each of their iterations takes $\Theta(n^3)$ time. Since most graphs require at least 100 iterations, Veldt et al. (2019); Ruggles et al. (2020) cannot be used to solve problems of this magnitude. Other methods of solving the LP are also not feasible as they run out of memory on much smaller instances.

We can see the performance of our method PROJECT AND FORGET in Table 4. Here, we see that our method has solved these problems in a reasonable amount of time. As we can see from Table 4, these instances have over 500 trillion constraints, but the number of active constraints is only a tiny fraction of the total number of constraints. Thus, using our approach, we can solve the weighted correlation clustering problem on much larger graphs than ever before.

There are two reasons as to why PROJECT AND FORGET can be used to solve these problems at these scales. First, while these instances have over 500 trillion constraints, the number of active constraints is only a tiny fraction of the total number of constraints. Second, due to the sparsity of the graph, our oracle finds violated cycle inequalities relatively quickly (sub-cubic time) and, since we forget inactive constraints, we project onto this relatively small number of constraints in each iteration. Thus, using our approach, we can solve the weighted correlation clustering problem on much larger (10 times bigger) graphs than ever before.

6. Applications: General algorithm

We shall now also provide experimental evidence that supports the generality of the algorithm. To do so we will solve three different problems, each with general linear constraints (not simply metric constraints) and a variety of objective functions.

First, we solve the dual of the quadratically regularized optimal transport problem from Blondel et al. (2018). Solving this problem helps highlight various different aspects of our algorithm that are not highlighted by the metric constrained problem. Second, we solve the information theoretic metric learning from Davis et al. (2007). Solving this problem highlights the wide variety of objective functions that our method can handle. Finally, we use our algorithm to train L_2 SVMs to highlight cases when the truly stochastic variant of our algorithm would be useful.

6.1 Sparse Optimal Transport

Optimal transport is a ubiquitous problem that appears in many different areas of machine learning, data science, statistics, mathematics, physics, and finance.

Problem 1 *Given two measure spaces $(\mathcal{X}, \Sigma_X, \mu)$ and $(\mathcal{Y}, \Sigma_Y, \nu)$, and a cost function $c : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$, the Monge-Kanterovich Optimal Transport seeks joint probability π on $\mathcal{X} \times \mathcal{Y}$ that minimizes*

$$\int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y),$$

subject to constraint that the marginals of π are equal to μ and ν .

In the discrete setting this problem can be formulated as follows:

$$\begin{aligned} OT(\mathbf{a}, \mathbf{b}) &= \min \langle \mathbf{C}, \mathbf{P} \rangle \\ \text{Subject to: } \mathbf{a} &= \mathbf{P} \mathbf{1}_m, \mathbf{b} = \mathbf{P}^T \mathbf{1}_n \end{aligned} \tag{6.1}$$

Where \mathbf{C} is the cost matrix and \mathbf{P} is the transport matrix. It has been shown that optimal solutions to the discrete Monge-Kanterovich problem are sparse (Brualdi, 2006). Specifically, at most $n + m - 1$ entries of \mathbf{P} can be non-zero.

While Problem 1 can be formulated as a linear program, solving it is fairly challenging as it has a quadratic number of variables and so many alternative formulations have been proposed. One such alternative, called ROT, is presented in Blondel et al. (2018); it uses a different regularizer, most importantly an L_2 regularizer. In this case, they showed experimentally that the solutions are still sparse. We show that PROJECT AND FORGET can also be used to solve ROT.

$$\begin{aligned} \text{Minimize: } & \langle \mathbf{C}, \mathbf{P} \rangle + \gamma \|\mathbf{a} - \mathbf{P}\mathbf{1}_m\|^2 + \gamma \|\mathbf{b} - \mathbf{P}^T\mathbf{1}_n\|^2 \\ \text{Subject to: } & \forall i \in [n], \forall j \in [m], \mathbf{P}_{ij} \geq 0 \end{aligned} \quad (6.2)$$

6.1.1 DUAL PROBLEM

In particular, we will not solve ROT directly; instead, we will solve the dual formulation of ROT.

$$\begin{aligned} \text{Min: } & \frac{1}{\gamma} \|\mathbf{f}\|^2 + \frac{1}{\gamma} \|\mathbf{g}\|^2 - \langle \mathbf{f}, \mathbf{a} \rangle - \langle \mathbf{g}, \mathbf{b} \rangle \\ \text{Subject to: } & \mathbf{f}_i + \mathbf{g}_j \leq \mathbf{C}_{ij} \end{aligned} \quad (6.3)$$

Generally, solving the dual problem, while it gives us the optimal value of the primal objective function, does not give you the optimal value of the primal variables directly. Thus, going from the solution to Problem 6.3 to that of Problem 6.2 is non trivial. This brings us to one aspect of PROJECT AND FORGET that is not highlighted by the metric constrained problems. That is,

if we use PROJECT AND FORGET to solve **either** the primal problem or the dual problem, then it finds the optimal solution to **both** problems.

This result holds as a result of Proposition 3. This proposition states that the variables z converge. Since we maintain the KKT conditions throughout the algorithm, the value that z converges to is the optimal solution of the dual problem and, thus, we solve the dual problem.

6.1.2 SPARSITY

We want to solve the dual problem instead of the primal because PROJECT AND FORGET exploits sparsity to its great advantage. Specifically, the dual constraints are extremely sparse; each dual constraint involves two variables so we have rapid computation of the projection.

6.1.3 EXPERIMENTAL RESULTS

In this subsection, we detail the experimental set up and present the results.

Data. The experimental set up is as follows. We take two shifted Gaussian distributions with means ± 15 and variance 10. Then, we split the interval $[-20, 20]$ into n points and create two discrete distributions by sampling the Gaussians on those n points. We use the squared Euclidean distances between the points as the cost function. We set the regularization parameter $1/\gamma = 10^{-3}$. This set up is a very basic example of the optimal transport problem and is an important test example for algorithmic comparisons.

Convergence Details. The feasibility error for Mosek and CPLEX are those reported by the solvers. For PROJECT AND FORGET, we calculate the feasibility error by

$$\max_{i,j} \mathbf{f}_i + \mathbf{g}_j - \mathbf{C}_{ij}.$$

Solvers. For the oracle for the PROJECT AND FORGET, we simply look through all the constraints and the return all of the violated constraints. This oracle clearly satisfies Property 1. We solve the dual version of the problem using PROJECT AND FORGET (PF), Mosek, and CPLEX and the primal version of the problem using Mosek, CPLEX, LFBGSB, and projected gradient descent (PGD).

Results. Table 5 shows that PROJECT AND FORGET can be used to solve the problem for much larger values of n . Additionally, for the smaller values of n , we see that PROJECT AND FORGET is competitive in terms of solve time. One thing to note from Table 6 is that the difference between the Primal objective and the Dual objective for PROJECT AND FORGET is smaller than 10^{-7} ; however, for Mosek, and CPLEX this is not the case. Thus, while both have similar levels of feasibility error, we see that our solver is still more “converged” since the primal dual gap is smaller.

Algorithm	501	1001	5001	10001	20001
Poject and Forget	12	151	1972	5909	21665
Cyclic Bregman	2681	-	-	-	-
LFBGSB	24	162	4080	Out of memory.	
Mosek dual	56	328	1927	Out of memory.	
Mosek primal	5	Out of memory.			
CPLEX primal	105	Out of memory.			
CPLEX dual	Out of memory.				
PGD	Did not converge.				

Table 5: Time taken in seconds to solve the quadratic regularized optimal transport problem. All experiments were run on a machine with 52 GB of RAM.


Again, we look at how ROD and FD evolve. Figure 4 shows that PROJECT AND FORGET does fewer projections overall than cyclic Bregman does in one iteration. This is the same as the result that we saw for metric nearness. Here, FD is measured as the true feasibility error.

These results hold for synthetic data. We also test our method on real world data. For these tests, we consider the application of optimal transport known as color transfer (as done in Blondel et al. (2018)). Here, we only compare against the best solver from the previous experiment. The results can be seen in Table 7. As we can see from the table, as we scale the size of the problem, PROJECT AND FORGET takes less time than Mosek and we can scale to larger problems before running out of memory. We also have better convergence statistics than Mosek.

PROJECT AND FORGET

Solver	Objective		Feasibility Error	
	Dual	Primal	Dual	Primal
$n = 501$				
Project and Forget	3.8416077	3.8416077	1.7e-09	0
Mosek Dual	3.8414023	3.8414023	3.8e-08	2.8e-10
LBFGBS	n/a	3.8416114	n/a	0
Mosek Primal	3.8303160	3.8303203	3.5e-8	2.2e-11
CPLEX Primal	3.8416376	3.8416076	8.44e-07	1.13e-04
CPLEX Dual		Ran out of memory		
$n = 1001$				
Project and Forget	1.947532046	1.947532046	2.0e-8	0
Mosek Dual	1.947091229	1.947091203	2.7e-08	8.4e-11
LBFGBS	n/a	1.947548404	n/a	0
Mosek Primal		Ran out of memory		
CPLEX Primal		Ran out of memory		
CPLEX Dual		Ran out of memory		
$n = 5001$				
Project and Forget	3.946556152e-01	3.946556154e-01	2e-08	0
Mosek Dual	3.880175376e-01	3.880175255e-01	1.2e-08	7.4e-11
LBFGBS	n/a	3.947709104e-01	n/a	0
Mosek Primal		Ran out of memory		
CPLEX Primal		Ran out of memory		
CPLEX Dual		Ran out of memory		
$n = 10001$				
Project and Forget	0.197687348	0.197687348	2.0e-8	0
Mosek Dual		Ran out of memory		
LBFGBS		Ran out of memory		
Mosek Primal		Ran out of memory		
CPLEX Primal		Ran out of memory		
CPLEX Dual		Ran out of memory		
$n = 20001$				
Project and Forget	0.0989355070	0.0989355070	2.0e-8	0
Mosek Dual		Ran out of memory		
LBFGBS		Ran out of memory		
Mosek Primal		Ran out of memory		
CPLEX Primal		Ran out of memory		
CPLEX Dual		Ran out of memory		

Table 6: Table showing the convergence details for the various solvers.



(a) Plot for ROD versus number of projections for PROJECT AND FORGET and Bregman’s cyclic method
 (b) Plot for ROD versus number of projections for PROJECT AND FORGET and Bregman’s cyclic method

Figure 4: Plot for ROD and FD versus number of projections for PROJECT AND FORGET and Bregman’s cyclic method

Method	n	Time	Objective Gap	Feasibility Error	
				Primal	Dual
Project and Forget Mosek Dual	512	611	8e-11	0	9.8e-11
		6	2e-10	7e-8	7e-9
Project and Forget Mosek Dual	1024	724	1.3e-10	0	5.2e-11
		31	9e-10	3e-8	2.9e-9
Project and Forget Mosek Dual	4096	825	2.8e-9	0	6.9e-9
		995	2e-7	5.5e-7	2e-8
Project and Forget Mosek Dual	8192	2197	3e-9	0	9.7e-8
		Out of Memory			

Table 7: Table showing the time taken in seconds, the duality gap (difference between primal and dual objectives), as well as the primal and dual feasibility error for PROJECT AND FORGET and Mosek for solving the dual formulation for color transfer. Here we set $\gamma = 1e2$

6.2 Information Theoretic Metric Learning

The next problem we consider is metric learning. There are many different versions of this problem (see Bellet et al. (2013); Suárez Díaz et al. (2018) for surveys on the topic). We focus on a specific instantiation, information theoretic metric learning (ITML) from Davis et al. (2007). This formulation allows us to demonstrate the wide range of objective functions that our method can handle.

Given two sets S, D that represent the sets of similar and dissimilar points, we learn the co-variance matrix M of a Gaussian distribution $p(x; M)$. The metric learned is then the Mahalanobis metric associated with M . Concretely, the problem is as follows:

$$\begin{aligned} & \text{minimize} && KL(p(x; M) \| p(x; I)) \\ & \text{subject to} && d_A(x_i, x_j) \leq u && (i, j) \in S \\ & && d_A(x_i, x_j) \geq l && (i, j) \in D. \end{aligned} \tag{6.4}$$

Davis et al. (2007) suggest solving this problem by sampling a small subset of the constraints and then using Bregman’s cyclic method. However, since our method solves the problem with the large number of constraints, when we compare against the method from Davis et al. (2007), we can no longer compare in terms of convergence statistics or time taken to solve the problem. Hence, as the method from Davis et al. (2007) is also based on Bregman projections, we limit both algorithms to the same number of projections. Then, we compare the algorithms on the quality of the solution; that is, the testing accuracy based on the metrics learned.

For each data set, we uniformly at random choose 80% of the data points to be the training set and the remaining to be the test set. We then let the similar pairs S be those pairs that have the same label and the dissimilar pairs D be all of the other pairs. For the algorithm from Davis et al. (2007), as suggested in the paper, we randomly sampled $20c^2$ constraints, where c is the number of different classes and run the algorithm so that it performs approximately 10^6 projections. We implement the algorithm in Julia. PROJECT AND FORGET returns $50c^2$ uniformly sampled similarity constraints and $50c^2$ uniformly random dissimilarity constraints and runs until we perform approximately 10^6 projections.

The hyper-parameters were set as follows: $\gamma = 1, u = 1, l = 10$. The pseudo-code for our algorithm can be seen in Algorithm 8. The classification is done using the k nearest neighbor classifier, with $k = 5$.

As shown in Table 8, our algorithm results in, for the most part, better test accuracy. The test accuracy is significantly better for the Ionosphere data set. This is the smallest of the data sets, from which we conclude that if PROJECT AND FORGET can solve the entire problem, instead of the approximation from Davis et al. (2007), we get a performance boost. As the size of the problem gets larger, since we have limited the number of projections, our algorithm has not converged as much. On the other hand, since the size of the problem for Davis et al. (2007) does not depend on the size of S and D , its convergence does not change.

We are also interested in how the test accuracy evolves as we do more projections. This can be seen in Figure 5. As we can see since ITML solves a smaller problem; it very quickly stabilizes. However, since PROJECT AND FORGET solves the complete problem using the stochastic variant of the method, as we do more projections, we see more data points and the accuracy evolves. In some cases, such as for Ionosphere, this results in increased accuracy. However, in other cases, such as Letter, this results in decreased accuracy.

Algorithm 8 Pseudo-code for the Project and Forget algorithm for ITML.

```

1: function PFITML( $X, C, \gamma, u, l, S, D$ )
2:    $\lambda^0 = 0, \Xi_{ij} = u$  for  $(i, j) \in S$  and  $\Xi_{ij} = l$  for  $(i, j) \in D$ . Initialize  $C = I$ .
3:   while Not done 1e7 projections do
4:     Randomly sample  $50c^2$   $(i, j)$  from  $S$  and  $50c^2$  from  $D$ 
5:     Do projection for these constraints.
6:     Project onto all remembered constraints.
7:   return  $C$ 
8: function PROJECTION( $X, i, j, S, D, u, l, \Xi, \lambda, C$ )
9:    $p = \text{dist}_C(X_i, X_j)$ 
10:   $\delta = 1$  if  $(i, j) \in S$  and  $\delta = -1$  if  $(i, j) \in D$ 
11:   $\alpha = \min\left(\lambda_{ij}, \frac{\delta}{2} \left(\frac{1}{p} - \frac{\gamma}{\Xi_{ij}}\right)\right)$ 
12:   $\beta = \frac{\delta\alpha}{1-\delta\alpha p}$ 
13:   $\Xi_{ij} = \gamma \frac{\Xi_{ij}}{\gamma + \delta\alpha\Xi_{ij}}$ 
14:   $\lambda_{ij} = \lambda_{ij} - \alpha$ 
15:   $C = C + \beta C(x_i - x_j)(x_i - x_j)^T C^T$ 

```

Dataset	S	D	PROJECT AND FORGET		ITML	
			Test Accuracy	#Projections	Test Accuracy	#Projections
Banana	4.5e6	4.5e6	89.6%	1.003e6	89.8%	1e6
Ionsphere	2.1e4	1.8e4	88.3%	1.007e6	83.7%	1e6
Coil2000	2.7e7	3.5e6	93.3%	1.02e6	93.2%	1e6
Letter	4.9e6	1.2e8	90.0%	1.08e6	92.6%	1.0004e6
Penbased	3.9e6	3.5e7	98.5%	1.08e6	98.2%	1e6
Spambase	3.5e6	3.2e6	92.5%	1.01e6	90.7%	1e6
Texture	8.8e5	8.8e6	99.8%	1e6	98.7%	1.002e6

Table 8: Table comparing the testing accuracy of PROJECT AND FORGET and ITML. Numbers are averaged over 5 trials.

6.3 Support Vector Machines

Finally, to show the usefulness of the truly stochastic variant, we use our algorithm to train L_2 SVMs. Given training data x_1, \dots, x_n , a feature map ϕ , and labels $y_1, \dots, y_n \in \{-1, 1\}$, we want to learn a vector w and slack variables ξ_i such that

$$\begin{aligned}
& \text{minimize} && \frac{1}{2}w^T w + \frac{C}{2} \sum_i \xi_i^2 \\
& \text{subject to} && y_i(w^T \phi(x_i)) \geq 1 - \xi_i \quad \forall i = 1, \dots, n \\
& && \xi_i \geq 0, \quad i = 1, \dots, n.
\end{aligned}$$

Here we use an L_2 penalty for the slack variables instead of the usual L_1 penalty. Tang (2013) provides insights into when using an L_2 penalty would provide better results.

Joachims et al. (2009) is a cutting plane based method to train L_2 SVMs. However, in contrast with Joachims et al. (2009), we do not need to use an oracle that finds the maximum




Figure 5: Plots showing the test accuracy after each projection for PROJECT AND FORGET (to solve the whole problem) and cyclic Bregman to solve the restricted problem (ITML) for the Ionosphere and Letter dataset

violated constraint. Instead, we used the truly stochastic variant of our algorithm. Hence, instead of searching through our data points to find large violations, we pick a subset of data points uniformly at random and project onto the constraints defined by these data points.

Implementation Details. We used the LIBLINEAR implementation that is available online. We interfaced with LIBLINEAR using an interface written in Julia. The implementation for our algorithm follows the pseudo-code presented in algorithm 9.

Algorithm 9 Pseudo-code for the implementation training an SVM.

```

1:  $z^0, w^0 = 0$ .  $n$  is the number of training samples
2: for  $iter = 1, \dots, MaxIters$  do
3:   for  $i = 1, \dots, n$  do
4:      $x_j =$  random data point
5:   Project  $w$  onto the constraint defined by  $x_j$ .
return  $w$ 

```

Data. We generated the data matrix X as follows. Each entry X_{ij} was sampled from $\mathcal{N}(0, K^2)$. We then sampled the coordinates of H from $\mathcal{N}(0, 1)$ and used H to label all the points in X . We then added noise so that H did not correctly separate the data. This noise was also sampled from $\mathcal{N}(0, 1)$. We generate data in this manner because a standard pre-processing step is to normalize features (i.e., transform the data so that each feature has mean zero and variance one).

We used $K = 1, 2$ and 5 to generate our data set. The different values for K gave us the different values of the percentage of data points misclassified by the original classifier s that we report in Table 9.

Convergence criterion. We considered a few different convergence criteria, but settled upon running the algorithm for 100 projections since whenever we train an SVM, we have a validation set. This number of projections becomes a hyperparameter in the model and,

we found experimentally that increasing the number of projections (unless increasing it by several orders of magnitude) beyond 100 did not increase the validation accuracy.

6.3.1 RESULTS

Table 9 gives us the results for the experiment. In this experiment, we are not checking for convergence, but demonstrating a possible use case of the truly stochastic variant of our algorithm. Because of the way in which we have generated data, it is easy to see that we have good test accuracy only if the current separating hyper-plane is roughly similar to the optimal one. We also see that this method has roughly similar performance to the LIBLINEAR dual solver. However, the primal solver from LIBLINEAR has the best testing accuracy. Thus, we see that the truly stochastic version of the algorithm can potentially be used to get reasonable results quickly.

Parameters				Time			Test Acc.		
n	d	s	C	Ours	Liblinear Dual	Liblinear Primal	Ours	Liblinear Dual	Liblinear Primal
1,000,000	100	6.3%	10^3	1.33 s	547 s	9.87 s	94.7%	94.7%	96.8%
1,000,000	100	12.6%	10^3	1.38 s	1032 s	10.1 s	91.4%	90.2%	93.6%
1,000,000	100	29.5%	10^3	1.58 s	1532 s	8.32 s	78.4%	77.1%	85.2%

Table 9: Table comparing the testing accuracy and running times for the truly stochastic variant of our algorithm against LIBLINEAR for binary classification using an L_2 SVM.

7. Conclusion and Future Work

In conclusion, in this paper, we present a new algorithm PROJECT AND FORGET that can be used to solve highly constrained convex optimization problems. We show that under some general assumptions, our method has a linear rate of convergence. Additionally, we demonstrate that our method returns both the dual and primal solutions.

The main type of problems we are interested in solving are metric constrained problems. We show that standard solvers cannot solve metric constrained problems at even moderate scales. We then demonstrate that PROJECT AND FORGET can solve the problem at large scales. We do this by solving larger instances of the metric nearness and correlation clustering problems.

We also demonstrate the generality of the method. First, we show that a variety of objective functions can be used with our methods, as demonstrated by the information theoretic metric learning experiment. We also demonstrate the role of sparsity and how it helps our method using the optimal transport experiment.

For future work, we hope that our method inspires a new way of learning metrics on data sets. Earlier researchers were either restricted to learning embeddings and then extracting a metric from those embeddings. However, using our method, we can now learn a much

wider variety of metrics. We also hope that once we can learn optimal metrics on data, these optimal metrics then suggest the types of spaces the data should be embedded into, rather than users picking the space in which to embed in an ad hoc fashion. We also hope to adapt the general convex constraint version to solve highly constrained semi-definite programs.

8. Proofs

8.1 Proof of part 1 of Theorem 1 for oracles that satisfy property 1

We remind the reader of the notation established in Section 2. The vector of variables over which we optimize is x , f is the objective function, $H_i = \{y : \langle y, a_i \rangle = b_i\}$ are the hyperplanes that lie on the boundaries of the half-space constraints, $L(x, z)$ is the Lagrangian, z is the dual variable, A is the matrix with rows given by a_i , and b is the vector with rows b_i .

Next, we clarify the indexing of the variables. Algorithm 1 has three steps per iteration and during the PROJECT step there are multiple projections. When we want to refer to a variable after the ν th iteration, it will have a superscript with a ν . When we refer to a variable after the n th, i th, k th projection, we use the superscript n, i, k . Finally, before the n th projection, $i(n)$ will represent the index of the hyper-plane onto which we project.

Finally, let R be the maximum number of constraints that our oracle \mathcal{Q} returns. This is clearly upper bounded by the total number of constraints, which we have assumed is finite. We are now ready to prove the first part of Theorem 1.

Theorem 1. (Part 1) *If $f \in \mathcal{B}(S)$, H_i are strongly zone consistent with respect to f , $\exists x^0 \in S$, such that $\nabla f(x^0) = 0$, and the oracle \mathcal{Q} satisfies property 1, then any sequence x^n produced by Algorithm 1 converges to optimal solution of problem 2.2.*

Proof The proof of this theorem is an adaptation of the proof of convergence for the traditional Bregman method that is presented in Censor and Zenios (1997) whose proof entails the following four steps. The main difference between Censor and Zenios' proof and ours is that of the last two steps. We present the entire proof, however, for completeness. To that end, we show the following.

- Step 1. The KKT condition, $\nabla f(x) = \nabla f(x^0) - A^T z$, is always maintained.
- Step 2. The sequence x^n is bounded and it has at least one accumulation point.
- Step 3. Any accumulation point of x^n is feasible (i.e., is in C).
- Step 4. Any accumulation point is the optimal solution.

Step 1. The KKT condition, $\nabla f(x) = \nabla f(x^0) - A^T z$, is always maintained.

We show by induction that for all n , $\nabla f(x^n) = -A^T z^n$. In the base case, $z = 0$, thus, $\nabla f(x^0) = 0 = -A^T z^0$. Assume the result holds for iteration n , then

$$\begin{aligned} \nabla f(x^{n+1}) &= \nabla f(x^n) + c^n a_{i(n)} \\ &= -A^T z^n + A^T c^n e_{i(n)} \\ &= -A^T (z^n - c^n e_{i(n)}) \\ &= -A^T z^{n+1}. \end{aligned}$$

We know that $c^n \leq z_{i(n)}^n$; therefore, we maintain $z^{n+1} \geq 0$ as well.

Step 2. The sequence x^n is bounded and has an accumulation point.

To show that x^n is a bounded, we first show that $(L(x^n, z^n))_n$ is a monotonically increasing sequence bounded from above. This observation results from the following string of equalities.

$$\begin{aligned}
 L(x^{n+1}, z^{n+1}) - L(x^n, z^n) &= f(x^{n+1}) - f(x^n) + \langle z^{n+1}, Ax^{n+1} - b \rangle - \langle z^n, Ax^n - b \rangle \\
 &= f(x^{n+1}) - f(x^n) + \langle A^T z^{n+1}, x^{n+1} \rangle - \langle A^T z^n, x^n \rangle - \langle z^{n+1} - z^n, b \rangle \\
 &= f(x^{n+1}) - f(x^n) - \langle \nabla f(x^{n+1}), x^{n+1} \rangle \langle \nabla f(x^n), x^n \rangle + \langle c^n e_{i(n)}, b \rangle \\
 &= f(x^{n+1}) - f(x^n) - \langle \nabla f(x^n) + c^n a_{i(n)}, x^{n+1} \rangle + \langle \nabla f(x^n), x^n \rangle + c^n b_{i(n)} \\
 &= f(x^{n+1}) - f(x^n) - \langle \nabla f(x^n), x^{n+1} - x^n \rangle - \langle c^n a_{i(n)}, x^{n+1} \rangle + c^n b_{i(n)} \\
 &= \underbrace{D_f(x^{n+1}, x^n)}_{(1)} + \underbrace{c^n (b_{i(n)} - \langle a_{i(n)}, x^{n+1} \rangle)}_{(2)}.
 \end{aligned}$$

Next, we show that both terms (1) and (2) are non-negative. We know that D_f is always non-negative so we only need to consider term (2). There are two cases: (i) if $c^n = \theta^n$, then $x^{n+1} \in H_{i(n)}$ and $b_{i(n)} - \langle a_{i(n)}, x^{n+1} \rangle = 0$. On the other hand, (ii) if $c^n = z_{i(n)}^n$, then $b_{i(n)} - \langle a_{i(n)}, x^{n+1} \rangle \geq 0$ and $c^n \geq 0$. We can conclude that the difference between successive terms of $L(x^n, z^n)$ is always non-negative and, hence, it is an increasing sequence.

To bound the sequence, let y be a feasible point (i.e., $Ay \leq b$). (Note that this is the only place we use the assumption that the feasible set is not empty.) Then

$$\begin{aligned}
 D_f(y, x^n) &= f(y) - f(x^n) - \langle \nabla f(x^n), y - x^n \rangle \\
 &= f(y) - f(x^n) + \langle z^n, Ay - Ax^n \rangle \\
 &\leq f(y) - f(x^n) + \langle z^n, b - Ax^n \rangle.
 \end{aligned}$$

Rearranging terms in the inequality, we obtain a bound on the sequence $L(x^n, z^n)$ from above:

$$\begin{aligned}
 L(x^n, z^n) &= f(x^n) + \langle z^n, Ax^n - b \rangle \\
 &\leq f(y) - D_f(y, x^n) \\
 &\leq f(y).
 \end{aligned}$$

Since the sequence $(L(x^n, z^n))_{n \in \mathbb{N}}$ is increasing and bounded, it is a convergent sequence and the difference between successive terms of the sequence goes to 0. Therefore,

$$\lim_{n \rightarrow \infty} D_f(x^{n+1}, x^n) = 0.$$

From the previous inequality we also have that

$$D_f(y, x^n) \leq f(y) - L(x^n, z^n) \leq f(y) - L(x^0, z^0) =: \alpha.$$

Using part (ii) of the definition of a Bregman function, we see that $L_2^f(y, \alpha)$ is bounded and since $(x^n)_{n \in \mathbb{N}} \in L_2^f(y, \alpha)$, x^n is an infinite bounded sequence. Thus, has an accumulation point.

Step 3. Any accumulation point x^* of x^n is feasible (i.e., is in C).

This is the only step in which we use the fact that our oracle satisfies property 1. Let x^* be some accumulation point for x^n and assume for the sake of contradiction that $Ax^* \not\leq b$. Let \tilde{A}, \tilde{b} be the maximal set of constraints that x^* does satisfy; i.e.,

$$\tilde{A}x^* \leq \tilde{b}.$$

Let (x^{n_k}) be a sub-sequence such that $x^{n_k} \rightarrow x^*$ and H be a constraint that x^* violates. Define ϵ as

$$\epsilon := \phi(d(x^*, H)) > 0. \quad (8.1)$$

Because x^n is bounded, x^{n_k} is a convergent sub-sequence $x^{n_k} \rightarrow x^*$, and $D_f(x^{n+1}, x^n) \rightarrow 0$, by Equation 6.48 from Censor and Zenios (1997) we see that for any t ,

$$x^{n_k+t} \rightarrow x^*.$$

In particular, the proposition holds for all $t \leq 2|\tilde{A}| + 2 =: T$.

Let us consider an augmented sub-sequence $x^{n_k}, x^{n_k+1}, \dots, x^{n_k+T}$, i.e., add in extra terms. Note that if $n_{k+1} - n_k \rightarrow \infty$, then this augmented sequence is not the entire sequence. We want to show that infinitely many of the terms in our augmented sequence satisfy a constraint *not* in \tilde{A} . Should this hold, then because we have only finitely many constraints, there exists at least one single constraint \tilde{a} that is *not* in \tilde{A} , such that infinitely many terms of the augmented sequence all satisfy the single constraint \tilde{a} . Finally, because our augmented sequence converges to x^* and we are only looking at closed constraints, we must have that x^* also satisfies the constraint \tilde{a} . Thus, we would arrive at a contradiction of the maximality of \tilde{A} and x^* would have to be in the feasible region.

To see that infinitely many of the terms in our augmented sequence satisfy a constraint *not* in \tilde{A} , let ν_k be the iteration in which the n_k th projection takes place. Note that we can assume without loss of generality that in any iteration, we project onto any constraint at most once. If this were not the case and we projected onto constraints more often, we would simply change the value of T to reflect this larger number of projections. Therefore, we have two possibilities for which iteration the $n_k + |\tilde{A}| + 1$ st projection takes place and we consider each case below.

Case 1: The $n_k + |\tilde{A}| + 1$ st projection, infinitely often, takes places in ν_k th iteration. Since we project onto each constraint at most once, one of the projections between the n_k and $n_k + |\tilde{A}| + 1$ st projection must be onto a hyper-plane defined by a constraint *not* represented in \tilde{A} and amongst the terms $x^{n_k}, x^{n_k+1}, \dots, x^{n_k+|\tilde{A}|+1}$, we must have a term that satisfies a constraint *not* in \tilde{A} infinitely often.

Case 2: The $n_k + |\tilde{A}| + 1$ st projection, infinitely often, takes place in $\nu_k + 1$ st iteration or later.

If this projection happens in $\nu_k + 1$ st iteration, consider the iteration in which we do the $n_k + T$ th projection. If this projection also takes place in the $\nu_k + 1$ st iteration, then we have done at least $|\tilde{A}| + 1$ projections in the $\nu_k + 1$ st iteration. Hence, amongst $x^{n_k+|\tilde{A}|+1}, \dots, x^{n_k+T}$, we must have a term that satisfies a constraint *not* in \tilde{A} .

If the $n_k + |\tilde{A}| + 1$ st or the $n_k + T$ th projection happens in the $\nu_k + 2$ nd iteration or later, then between the n_k th and the $n_k + T$ th projection, we must have projected onto all

constraints returned by oracle in the $\nu_k + 1$ st iteration. Therefore, we must have projected onto some hyper-plane defined by \hat{a}^{n_k} (for some constraint \hat{C}^{n_k}) such that

$$\hat{d}^{n_k} := d(x^{\nu_k+1}, \hat{C}^{n_k}) \geq \phi(d(x^{\nu_k+1}, C)).$$

Then there exists a sufficiently small $\delta > 0$, depending on $\tilde{A}, \tilde{b}, x^*$, such that if $\|y - x^*\| \leq \delta$, then

$$\tilde{A}y \leq \tilde{b} + \frac{\epsilon}{2} \mathbb{1},$$

where $\mathbb{1}$ is vector of all ones.

Since our augmented sequence converges to x^* , we know that there exists a K , such that for all $k \geq K$ and $t \leq T$, $\|x^{n_k+t} - x^*\| < \delta$. That is, for all $k \geq K$ and $t \leq T$,

$$\tilde{A}x^{n_k+t} \leq \tilde{b} + \frac{\epsilon}{2} \mathbb{1}. \tag{8.2}$$

Note x^{ν_k+1} is within our augmented sequence so if \hat{a} is infinitely often in \tilde{A} , by equation 8.2, we have that infinitely often

$$\frac{\epsilon}{2} \geq \hat{d}^{n_k}.$$

Finally, because the augmented sequence converges to x^* ,

$$\begin{aligned} \epsilon &= \phi(d(x^*, H)) \\ &\leq \phi(d(x^*, C)) \\ &= \lim_{k \rightarrow \infty} \phi(d(x^{\nu_k+1}, C)) \\ &\leq \lim_{k \rightarrow \infty} \hat{d}^{n_k} \\ &\leq \frac{\epsilon}{2}. \end{aligned}$$

The first inequality follows from the fact that ϕ is non-decreasing. Thus, we have a contradiction. Therefore, \hat{a}^{n_k} is *not* in \tilde{A} infinitely often and amongst $x^{n_k}, x^{n_k+1}, \dots, x^{n_k+T}$, we must have a term that satisfies a constraint *not* in \tilde{A} infinitely often. Thus, there is a constraint \tilde{a} not in \tilde{A} that is satisfied by infinitely terms of our augmented sequence and we have a contradiction.

Step 4. Optimality of accumulation point.

Because we have established the feasibility of all accumulation points, we show next that any accumulation point $x^{n_k} \rightarrow x^*$ is optimal.

First, we show that there exists an N , such that for any $k \geq N$, and for any a_i such that

$$\langle a_i, x^* \rangle < b_i,$$

we have $z_i^{n_k} = 0$. To do so, we assume for the sake of contradiction that for some a_i , our sequence $z_i^{n_k}$ is infinitely often not 0. The algorithm then projects onto this constraint infinitely often. Therefore, the point x^{n_k} lies on the hyper-plane defined by a_i, b_i infinitely often. Thus, the limit point x^* must lie on this hyper-plane as well and we have a contradiction.

Now we know that for any constraint a_i , we either have that $\langle a_i, x^* \rangle = b_i$ or we have that $z_i^{n_k} = 0$ for the tail of the sequence. Thus, for sufficiently large k ,

$$\begin{aligned} \langle z^{n_k}, Ax^{n_k} - b \rangle &= \langle A^T z^{n_k}, x^{n_k} - x^* \rangle \\ &= \langle -\nabla f(x^{n_k}), x^{n_k} - x^* \rangle \\ &= D_f(x^*, x^{n_k}) - f(x^*) + f(x^{n_k}). \end{aligned}$$

Next, by part (iii) of the definition of a Bregman function,

$$\lim_{k \rightarrow \infty} D_f(x^*, x^{n_k}) = 0.$$

Finally,

$$\lim_{k \rightarrow \infty} L(x^k, z^k) = \lim_{k \rightarrow \infty} f(x^{n_k}) + \langle z^{n_k}, Ax^{n_k} - b \rangle = f(x^*).$$

We also know that $L(x^k, z^k) \leq f(y)$ for any feasible y . Thus, $f(x^*) \leq f(y)$. Hence x^* is an optimal solution. Now since f is strictly convex, this optimal point is unique. Therefore, we have that $(x^n)_{n \in \mathbb{N}}$ has only one accumulation point and $x^n \rightarrow x^*$. \blacksquare

An important consequence of this proof is the following proposition:

Proposition 3. *If a_i is an inactive constraint, then there exists a N , such that for all $n \geq N$, we have that $z_i^n = 0$. That is, after some finite time, we never project onto inactive constraints ever again.*

8.2 Proof of part 1 of Theorem 1 for oracles that satisfy property 2

In this subsection, we prove part 1 of Theorem 1 for oracles that satisfy property 2. We make note of the key ideas in this proof as they are useful in the proof of the truly stochastic variant. To be precise, we prove:

Theorem 1. *(Part 1) If $f \in \mathcal{B}(S)$, H_i are strongly zone consistent with respect to f , $\exists x^0 \in S$, such that $\nabla f(x^0) = 0$, and the oracle \mathcal{Q} satisfies property 2, then with probability 1, any sequence x^n produced by Algorithm 1 converges to optimal solution of problem 2.2.*

Proof Assume that we have an oracle that satisfies property 2. A careful reading of the previous proof shows that if we switch out an oracle with property 1 for an oracle with property 2, then we only need to adjust step 3 of our proof. The crucial part of that step was showing that for our augmented sequence, we had infinitely many terms that satisfied a constraint *not* in \tilde{A} . We make the following adjustments to our analysis.

Let ν_k be the iteration in which the n_k th projection takes place. In the previous proof, we used the property of the oracle only when the $n_k + T$ th projection took place in the $\nu_k + 2$ nd iteration or a later iteration. In this case, the augmented sequence encompasses all of the $\nu_k + 1$ st iteration infinitely often.

Let us choose a constraint \hat{a} that is not satisfied by x^* . Because the oracle satisfies property 2, for each iteration $\nu_k + 1$, our oracle returns \hat{a} with probability at least $\tau > 0$. By the Borel Cantelli Lemma, we know that during the selected iterations, the constraint \hat{a} , with probability one, is returned infinitely often by our oracle. Thus, our augmented sequence satisfies this constraint with probability 1 and x^* lies in the feasible region with probability

1. ■

A direct consequence of this proof is the proof of the probabilistic version of Proposition 3.

Proposition 3. *With probability 1, we project onto inactive constraints a finite number of times.*

8.3 Proof of part 2 of Theorem 1

The discussions in Iusem and De Pierro (1990); Iusem (1991) almost directly apply to that for our algorithm. For completeness, we present it along with the necessary modifications. As with the traditional Bregman algorithm, we first present the case when $f(x)$ is quadratic. That is,

$$f(x) = r + s^T \cdot x + \frac{1}{2}x^T Hx,$$

where H is a positive definite matrix. In this case, it is easy to see that

$$D_f(x, y) = \|x - y\|_H^2 := (x - y)^T H(x - y).$$

8.3.1 PROOF OF PART 2 OF THEOREM 1—QUADRATIC CASE

In this section, we will prove the following variation of Theorem 1.

Theorem 1. *If f is a strictly convex quadratic function, H_i are strongly zone consistent with respect to f , $x^0 = H^{-1}s \in S$, and the oracle \mathcal{Q} satisfies either property 1 or 2, then there exists $\rho \in (0, 1)$ such that*

$$\lim_{\nu \rightarrow \infty} \frac{\|x^* - x^{\nu+1}\|_H}{\|x^* - x^\nu\|_H} \leq \rho, \tag{3.2}$$

where $\|y\|_H^2 = y^T H y$. In the case when we have an oracle that satisfies property 2, the limit in 3.2 holds with probability 1.

We establish some notation ahead of our lemmas. Let I be the set of all active constraints. That is, if x^* is the optimal solution then

$$I = \{i : \langle a_i, x^* \rangle = b_i\}.$$

Let S be the set of all x that satisfy these constraints (namely $S = \{x : \forall i \in I, \langle a_i, x \rangle = b_i\}$). Let H_x be the hyper-plane, such that H_x represents the constraint in I that is furthest from x . Define

$$\mu = \inf_{x \notin S} \frac{d(x, H_x)}{d(x, S)}.$$

By Iusem and De Pierro (1990), we know that $\mu > 0$. Let U be the set of all optimal dual variables z ; i.e., $U = \{z : \nabla f(x^*) = -A^T z\}$ and let $I_\nu = \{i : z_i^{\nu+1} \neq 0\}$.

Next, we present a few preliminary lemmas. These lemmas exist in some form or another in Iusem (1991); Iusem and De Pierro (1990) and we present them suitably modified for our purpose. These lemmas require the following set of assumptions about an iteration ν :

1. $\forall i \notin I, z_i^\nu = 0$;
2. for all $i \notin I$, we do not project onto this constraint in the ν th iteration; and,
3. there exists $z \in U$, such that for all $i \notin I_\nu, z_i = 0$.

Lemma 2 *Let x^* be the optimal solution for an instance of problem 2.2. For any sequence $x^n \rightarrow x^*$ such that x^n, z^n maintain the KKT conditions, there exists an M , such that for all $\nu \geq M$, there exists a $z \in U$, such that for all $i \notin I_\nu$, we have that $z_i = 0$.*

Proof Let $V_\nu = \{z : \forall i \notin I_\nu, z_i = 0\}$. Then assume, for the sake of contradiction, that the result is false. That is, there is a sequence ν_k such that $V_{\nu_k} \cap U = \emptyset$. Then since there finitely many different I_ν (hence finitely many V_ν), we have that one of these must occur infinitely often. Thus, by taking an appropriate sub-sequence, we assume, without loss of generality, that I_{ν_k} are all equal. Let $V = V_{\nu_k}$ and obtain $V \cap U = \emptyset$.

Since V is a closed subspace, U is a closed set, and $V \cap U = \emptyset$, we must have that $d(V, U) > 0$. But $z^{\nu_k+1} \in V$ and so

$$d(z^{\nu_k+1}, U) \geq d(V, U) > 0.$$

Since $x^\nu \rightarrow x^*$ and we maintain the KKT conditions, we have that for any $z \in U$,

$$A^T z^\nu = -\nabla f(x^\nu) \rightarrow -\nabla f(x^*) = A^T z.$$

Thus $d(z^\nu, U) \rightarrow 0$ which is a contradiction. ■

Lemma 3 *For any sequence $x^n \rightarrow x^*$, if for a given ν , we have that the sequence satisfies assumptions (1) and (2), then*

$$\|x^{\nu+1} - x^*\|_Q^2 \leq \|x^\nu - x^*\|_Q^2 - \sum_{n=k}^K \|x^{n+1} - x^n\|_Q^2$$

where k and K are the indices of the first and last projection that take place in the ν th iteration.

Proof This Lemma is simply a statement about Bregman projections and so its proof requires no modification. ■

Before we proceed, we introduce additional notation. Let A_{I_ν}, b_{I_ν} be the sub-matrix of A, b with rows from I_ν and

$$S_\nu = \{x : A_{I_\nu} x = b_{I_\nu}\}.$$

Lemma 4 *For any sequence such that $x^n \rightarrow x^*$, if for a given ν , we have that it satisfies assumptions (1), (2), and (3), then we have that $\|x^{\nu+1} - x^*\|_Q = d(x^{\nu+1}, S_\nu)$.*

Proof Consider the constrained problem

$$\min_{x \in S_\nu} \|x^{\nu+1} - x\|_Q^2. \quad (8.3)$$

Then sufficient conditions for a pair (x, z_{I_ν}) to be optimal for this problem are

$$A_{I_\nu} x = b_{I_\nu} \text{ and } x = x^{\nu+1} - Q^{-1} A_{I_\nu}^T z_{I_\nu}.$$

By Proposition 3, we see that since x^* is solution to problem 2.2, we have that $A_{I_\nu} x^* = b_{I_\nu}$. Then by assumptions and the manner in which we do projections, we have that there exists $z \in U$, such that for all $i \notin I_\nu$, $z_i = 0$ and

$$x^* = x^{\nu+1} - Q^{-1} A^T (z^{\nu+1} - z).$$

Then since $z_i^{\nu+1} = 0$ for all $i \notin I_\nu$, we have that

$$x^* = x^{\nu+1} - Q^{-1} A_{I_\nu}^T (z_{I_\nu}^{\nu+1} - z_{I_\nu}).$$

Thus, x^* is the optimal solution to 8.3. ■

Next for $x \notin S_\nu$, let H_x^ν be the hyper-plane of that is furthest from x and define

$$\mu_\nu = \inf_{x \notin S_\nu} \frac{d(x, H_x^\nu)}{d(x, S_\nu)}.$$

Now we are ready to prove the following theorem.

Theorem 4 *Let x^* is the optimal solution to problem 2.2. Then given ν that satisfies assumptions (1), (2), and (3), we have that*

$$\|x^{\nu+1} - x^*\|_Q^2 \leq \frac{L}{L + \mu^2} \|x^\nu - x^*\|_Q^2$$

where L is the number of projections that happened in ν th iteration.

Proof By Lemma 4, for any such ν we have that $x^{\nu+1} \notin S_\nu$ (or we have converged already). Suppose constraint $j \in I_\nu$ defines the hyper-plane $H_{x^{\nu+1}}^\nu$. Then by Lemma 4 and definitions of μ_ν, μ we have the following inequality.

$$\begin{aligned} \|x^{\nu+1} - x^*\| &= d(x^{\nu+1}, S_\nu) \\ &\leq \frac{1}{\mu_\nu} d(x^{\nu+1}, H_{x^{\nu+1}}^\nu) \\ &\leq \frac{1}{\mu} d(x^{\nu+1}, H_{x^{\nu+1}}^\nu). \end{aligned}$$

Now since $I_\nu = \{i : z_i^{\nu+1} \neq 0\}$, we know that during the ν th iteration we must have projected onto $H_{x^{\nu+1}}^\nu$. Note that this is the only place in the proof where we need the fact that we remember old constraints. Let us say that this happens during the r th projection of the ν th iteration.

Note by assumption, we satisfy the assumptions of Lemma 3. Let $y^r, y^{\nu+1}$ be the projections of $x^r, x^{\nu+1}$ onto $H_{x^{\nu+1}}^\nu$. Then we see that

$$\begin{aligned}
 d(x^{\nu+1}, H_{x^{\nu+1}}^\nu)^2 &= \|y^{\nu+1} - x^{\nu+1}\|_Q^2 \\
 &\leq \|y^r - x^{\nu+1}\|_Q^2 \\
 &\leq \left(\|y^r - x^{r+1}\|_Q + \sum_{i=r+1}^L \|x^i - x^{i+1}\|_Q \right)^2 \\
 &= \left(\sum_{i=r+1}^L \|x^i - x^{i+1}\|_Q \right)^2 \\
 &\leq \left(\sum_{i=0}^L \|x^i - x^{i+1}\|_Q \right)^2 \\
 &\leq L \sum_{i=0}^L \|x^i - x^{i+1}\|_Q^2 \\
 &\leq L (\|x^\nu - x^*\|_Q^2 - \|x^{\nu+1} - x^*\|_Q^2).
 \end{aligned}$$

Thus, we get that

$$\begin{aligned}
 \mu^2 \|x^{\nu+1} - x^*\|^2 &\leq d(x^{\nu+1}, H_{x^{\nu+1}}^\nu)^2 \\
 &\leq L (\|x^\nu - x^*\|^2 - \|x^{\nu+1} - x^*\|^2).
 \end{aligned}$$

Rearranging, we get that

$$\|x^{\nu+1} - x^*\|_Q^2 \leq \frac{L}{L + \mu^2} \|x^\nu - x^*\|_Q^2.$$

■

As a corollary to the above theorem, we have that algorithm 1 converges linearly.

Theorem 1. 2) If f is a strictly convex quadratic function, H_i are strongly zone consistent with respect to f , $x^0 = H^{-1}s \in S$, and the oracle \mathcal{Q} satisfies either property 1 or 2, then there exists $\rho \in (0, 1)$ such that

$$\lim_{\nu \rightarrow \infty} \frac{\|x^* - x^{\nu+1}\|_H}{\|x^* - x^\nu\|_H} \leq \rho \tag{3.2}$$

where $\|y\|_H^2 = y^T H y$. In the case when we have an oracle that satisfies property 2, the limit in 3.2 holds with probability 1.

Proof Using Proposition 3, Lemma 2, and that we have finitely many constraints, we see that if ν is large enough, the assumptions for Theorem 4 are satisfied. Taking the limit gives us the needed result.

In the case when we have an oracle that satisfies property 2, consider the product space of all possible sequences of hyper-planes returned by our oracle. In this product space, we see that with probability 1, we generate a sequence of hyper-planes, such that algorithm 1 converges. For any such sequence of hyper-planes, we have that 3.2 holds. Thus, the limit in 3.2 holds with probability 1 for random separation oracles. \blacksquare

8.3.2 PROOF OF PART 2 OF THEOREM 1—GENERAL

The rate of convergence for the general Bregman method was established in Iusem (1991). To show this, let \tilde{f} be the 2nd degree Taylor polynomial of f centered at the optimal solution x^* .

$$\tilde{f}(x) = f(x^*) + \nabla f(x^*)^T \cdot x + \frac{1}{2} x^T \cdot \nabla^2 f(x^*) \cdot x.$$

For notational convenience, let H be the Hessian of f at x^* . Then we can see that if replace f with \tilde{f} in 2.2 then the optimal solution does not change. Thus, if had access to \tilde{f} and could use this function to do our projections, then from the quadratic case we have our result.

Thus, to get the general result, if x^ν is our standard iterate and \tilde{x}^ν is the iterate produced by using \tilde{f} instead of f , then Iusem (1991) shows that $\|x^\nu - \tilde{x}^\nu\|$ is $o(\|x^\nu - x^*\|_H)$. Specifically, we can extract the following theorem from Iusem (1991).

Theorem 5 *Iusem (1991) Let x^* is the optimal solution for problem 2.2 and \tilde{x}^n is the sequence produced by using the same sequence of hyper-planes but with \tilde{f} instead of f . Given a sequence x^n produced by Bregman projections, such that $x^n \rightarrow x^*$, and for large enough ν we satisfy assumptions (1), (2), and (3), then $\|x^\nu - \tilde{x}^\nu\|$ is $o(\|x^\nu - x^*\|_H)$*

Using this we can get the general result as follows

$$\begin{aligned} \|x^{\nu+1} - x^*\|_H &\leq \|x^{\nu+1} - \tilde{x}^{\nu+1}\|_H + \|\tilde{x}^{\nu+1} - x^*\|_H \\ &\leq \|x^{\nu+1} - \tilde{x}^{\nu+1}\|_H + \rho \|\tilde{x}^\nu - x^*\|_H \\ &\leq \|x^{\nu+1} - \tilde{x}^{\nu+1}\|_H + \rho \|\tilde{x}^\nu - x^\nu\|_H \\ &\quad + \rho \|x^\nu - x^*\|_H. \end{aligned}$$

Then diving by $\|x^{\nu+1} - x^*\|_H$, and using Theorem 5 to take the limit, we get that there exists $\rho \in (0, 1)$ such that

$$\lim_{\nu \rightarrow \infty} \frac{\|x^* - x^{\nu+1}\|_H}{\|x^* - x^\nu\|_H} \leq \rho. \tag{3.2}$$

As with the quadratic case, we see that is an oracle satisfies property 2, then 3.2 holds with probability 1. Thus, we have proved Theorem 1 in its complete generality.

8.4 Proof of Theorem 2

In this section we prove Theorem 2 in essentially the same manner as we did for Theorem 1 and so we outline only what changes are necessary.

Theorem 2. *If $f \in \mathcal{B}(S)$, the hyper-planes H_i are strongly zone consistent with respect to f , and $\exists x^0 \in S$ such that $\nabla f(x^0) = 0$, then with probability 1 any sequence x^n produced by the algorithm converges to the optimal solution of problem 2.2. Furthermore, if x^* is the optimal solution, f is twice differentiable at x^* , and the Hessian $H := Hf(x^*)$ is positive semi-definite, then there exists $\rho \in (0, 1)$ such that with probability 1,*

$$\liminf_{\nu \rightarrow \infty} \frac{\|x^* - x^{\nu+1}\|_H}{\|x^* - x^\nu\|_H} \leq \rho. \quad (3.3)$$

To prove Theorem 2, we need to analyze only what goes wrong if the algorithm “forgets” all of the old constraints. First, consider the proof in the case that we converge to the optimal solution. Then, steps 1,2, and 3 are completely unaffected by forgetting old constraints. The only step that is affected is step 4. In a previous proof, we argued that if for some inactive constraint a_i , z_i is non-zero infinitely often, then we projected onto this constraint infinitely often. In our present setting, we cannot conclude this directly as $z_i^\nu > 0$ does not imply that we remember a_i on the ν th iteration. However, due to property 2, we know that \mathcal{Q} returns a_i with probability at least τ . Thus, again using the Borel Cantelli Lemmas, we see that we have a_i infinitely often and this iteration converges to the optimal.

To prove the second part of the theorem, we recall from Theorem 5 that we only need to analyze the case when f is a quadratic function. Indeed, the only place where we used the fact that we remembered old constraints was in the proof of Theorem 4 in which we needed to remember old constraints to guarantee that during the ν th iteration we project onto the constraint a_i that is furthest from x^ν among those constraints for which $z_i^{\nu+1} > 0$. We cannot guarantee that this happens always but we can guarantee that it happens infinitely often.

Therefore, the conclusion of Theorem 4 holds infinitely often instead of for the tail of the sequence and we replace the limit with a limit infimum to obtain the desired result.

8.5 General Convex Proof

We will need the following facts from Bauschke and Lewis (2000).

1. A convex function is Legendre if and only if its convex conjugate is Legendre. In this case the gradient mapping

$$\nabla f : \text{int}(\text{dom}(f)) \rightarrow \text{int}(\text{dom}(f^*)) : x \mapsto \nabla f(x)$$

is a topological isomorphism with inverse mapping $(\nabla f)^{-1} = \nabla f^*$.

2. Suppose f is Legendre on E and a S is a closed convex set in E with $S \cap \text{int}(\text{dom}(f)) \neq \emptyset$. Suppose further $y \in \text{int}(\text{dom}(f))$. Then the Bregman projection $P_S^f y$ of y onto S with respect to f is characterized by

$$P_S^f y \in S \cap \text{int}(\text{dom}(f)) \text{ and } \langle \nabla f(y) - \nabla f(P_S^f y), S - P_S^f y \rangle \leq 0.$$

In addition, $D_f(P_S^f y, y) \leq D_f(s, y) - D_f(s, P_S^f y)$ for all $s \in S \cap \text{dom}(f)$. Here when we write $\langle \nabla f(y) - \nabla f(P_S^f y), S - P_S^f y \rangle \leq 0$, we mean that $\langle \nabla f(y) - \nabla f(P_S^f y), s - P_S^f y \rangle \leq 0$, is true for all $s \in S$.

3. Three point identity: Suppose f is Legendre on E . If $x, y \in \text{int}(\text{dom}f)$ and $b \in \text{dom}f$, then

$$D_f(b, x) + D_f(x, y) - D_f(b, y) = \langle \nabla f(x) - \nabla f(y), x - b \rangle.$$

4. Suppose f is a very strictly convex function on E . Then for every compact convex subset K of $\text{int}(\text{dom}f)$, there exists reals $0 < \theta$ and $\Theta < +\infty$ such that for every $x, y \in K$ we have that

$$D_f(x, y) \geq \theta \|x - y\|^2 \text{ and } \|\nabla f(x) - \nabla f(y)\| \leq \Theta \|x - y\|.$$

On the n th iteration we projection onto the constraint given by $i(n)$. Thus, we have that

$$x^n \in C_{i(n)} \cap \text{dom}(f). \tag{8.4}$$

Then since f is Legendre, using Facts 1 and 2 with $y = \nabla f^*(\nabla f(x^{n-1}) + q^{p(n)})$, $S = C_{i(n)}$ and $P_S^f y = x^n$, we have that

$$\langle \nabla f(x^{n-1}) + q^{p(n)} - \nabla f(x^n), C_{i(n)} - x^n \rangle \leq 0.$$

Then by definition of q^n we have that

$$\langle q^n, x^n - C_{i(n)} \rangle \geq 0. \tag{8.5}$$

Then we can also see that

$$\nabla f(x^{n-1}) - \nabla f(x^n) = q^n - q^{p(n)}. \tag{8.6}$$

Then summing over this telescoping sum and the fact that the q are initialized to be 0 we get that

$$\nabla f(x^0) - \nabla f(x^n) = \sum_{k \in \mathcal{C}} q^{p(k,n)}. \tag{8.7}$$

Here \mathcal{C} stores the indices of all the constraints. To see why this is true, consider what happens if we sum the terms for just one constraint set. If we see the right hand side of Equation 8.6, then we see that only the q from the last time we projected onto that set remains.

Let $c \in \text{int}(\text{dom}f)$. Then we need to prove the following crucial equality.

$$\begin{aligned} D_f(c, x^n) &= D_f(c, x^{n+1}) + D_f(x^{n+1}, x^n) - \langle \nabla f(x^{n+1}) - \nabla f(x^n), x^{n+1} - c \rangle && \text{[Fact 3]} \\ &= D_f(c, x^{n+1}) + D_f(x^{n+1}, x^n) + \langle q^{(n+1)} - q^{p(n+1)}, x^{n+1} - c \rangle && \text{[Eq 8.6]} \\ &= D_f(c, x^{n+1}) + D_f(x^{n+1}, x^n) + \langle q^{n+1}, x^{n+1} - c \rangle - \langle q^{p(n+1)}, x^{n+1} - c \rangle \\ &= D_f(c, x^{n+1}) + D_f(x^{n+1}, x^n) + \langle q^{n+1}, x^{n+1} - c \rangle \\ &\quad - \langle q^{p(n+1)}, x^{n+1} + x^{p(n+1)} - x^{p(n+1)} - c \rangle \\ &= D_f(c, x^{n+1}) + D_f(x^{n+1}, x^n) + \langle q^{n+1}, x^{n+1} - c \rangle \\ &\quad - \langle q^{p(n+1)}, x^{n+1} - x^{p(n+1)} \rangle + \langle q^{p(n+1)}, x^{p(n+1)} - c \rangle. \end{aligned}$$

Using induction we can now prove that for $c \in \text{int}(\text{dom}f)$ and for every $n \geq 0$ we have that

$$D_f(c, x^0) = D_f(c, x^n) + \sum_{k=1}^n \left(D_f(x^k, x^{k-1}) + \langle q^{p(k)}, x^{p(k)} - x^k \rangle \right) + \sum_{k \in \mathcal{C}} \langle x^{p(k,n)} - c, q^{p(k,n)} \rangle.$$

Using Equation 8.5, we have that all the inner products on the right hand of the above equation are non-negative. Since the Bregman distances are always non-negative, every term on the right hand side in non-negative. Thus, if we take the limit as $n \rightarrow \infty$ of the equation we get that $D_f(c, x^n)$ is a convergent sequence. Thus, is bounded. We also get that $S_n = \sum_{k=1}^n D_f(x^k, x^{k-1})$ is also a convergent sequence and hence is bounded. Thus, we get that

$$\lim_{n \rightarrow \infty} D_f(x^n, x^{n+1}) = 0.$$

Then using the previous reasoning we see that x^n is a bounded sequence and has accumulation points. Furthermore by Bauschke and Lewis (2000) we have that all the accumulation points lie inside $\text{dom}f$ and we have that $x^k - x^{k-1} \rightarrow 0$.

Now using Equation 8.7 and adding and subtracting $x^{p(k,n)}$ from $c - x^n$, we get that

$$\langle c - x^n, \nabla f(x^0) - \nabla f(x^n) \rangle = \underbrace{\sum_{k \in \mathcal{C}} \langle c - x^{p(k,n)}, q^{p(k,n)} \rangle}_{S_1(n)} + \underbrace{\sum_{k \in \mathcal{C}} \langle x^{p(k,n)} - x^n, q^{p(k,n)} \rangle}_{S_2(n)}. \quad (8.8)$$

Using Equation 8.5, we see that $S_1(n)$ is non positive. Now we want to prove that

$$\lim_n \sum_{k \in \mathcal{C}} |\langle x^{p(k,n)} - x^n, q^{p(k,n)} \rangle| \stackrel{?}{=} 0. \quad (8.9)$$

Let $K := \text{cl}(\text{conv}(\{x^n\}))$. That is, K is the smallest closed convex set containing the sequence $\{x^n\}$. Then we have that

$$K = \text{conv}(\text{cl}(\{x^n\})) = \text{conv}(\{x^n\} \cup \{\text{cluster points of } x^n\}) \subseteq \text{conv}(\text{int}(\text{dom}f)) = \text{int}(\text{dom}f).$$

The first equality comes from the fact that the cluster points of x^n are exactly the points needed to make the set $\{x^n\}$ closed. The second comes from the fact that the clusters points x^n are in the interior of the domain of f , as we previously showed. Finally, since we assumed that the domain of f is convex, we get the last equality.

Then, since f is very strictly convex and K is compact (we showed that the sequence $\{x^n\}$ is bounded), we can use Fact 4 to see that there exists $0 < \theta$ and $\Theta < +\infty$, such that

$$D_f(x^k, x^{k-1}) \geq \theta \|x^k - x^{k-1}\|^2 \text{ and } \|\nabla f(x^k) - \nabla f(x^{k-1})\| \leq \Theta \|x^k - x^{k-1}\|.$$

Then using the fact that $\sum_{k=1}^{\infty} D_f(x^k, x^{k-1}) < \infty$ we have that

$$\sum_{k=1}^{\infty} \|x^k - x^{k-1}\|^2 \leq \frac{1}{\theta} \sum_{k=1}^{\infty} D_f(x^k, x^{k-1}) < \infty.$$

Now consider the following telescoping identity

$$q^n = (q^n - q^{p(n)}) - (q^{p(n)} - q^{p(p(n))}) - \dots - (q^{p^k(n)} - q^{p^{k+1}(n)}) - \dots 0.$$

Then, we get that

$$\sum_{k \in \mathcal{C}} \|q^{p(k,n+1)}\| \leq \sum_{k=1}^n \|q^k - q^{p(k)}\| = \sum_{k=1}^n \|\nabla f(x^k) - \nabla f(x^{k-1})\| \leq \Theta \sum_{k=1}^n \|x^k - x^{k-1}\|.$$

Now note that we only forget a constraint if $q^k = 0$ and that in each iteration we look at most $2M$ constraints (where M is number of constraints). Thus, if $q^{p(k,n)} \neq 0$ then, we must have that $p(k,n) \geq n - 4M$. That is, we either have projected onto that constraint in the current iteration or if we haven't projected onto that constraint yet in the current iteration, we must have projected onto it in the previous iteration. Thus, using the above we get that

$$\sum_{k \in \mathcal{C}} |\langle x^{p(k,n)} - x^n, q^{p(k,n)} \rangle| \leq \sum_{k \in \mathcal{C}} \|x^{p(k,n)} - x^n\| \|q^{p(k,n)}\| \quad (8.10)$$

$$\leq \sum_{k \in \mathcal{C}} \|q^{p(k,n+1)}\| \sum_{k=n-8M}^n \|x^k - x^{k-1}\| \quad (8.11)$$

$$\leq \Theta \sum_{k=1}^n \|x^k - x^{k-1}\| \sum_{k=n-4M}^n \|x^k - x^{k-1}\|. \quad (8.12)$$

Then using proposition 3.1 from Bauschke and Lewis (2000) with the fact that $\sum_{k=1}^{\infty} \|x^k - x^{k-1}\|^2 < \infty$ we get that the limit of the right hand side is 0. Thus, Equation 8.9 is true. Now looking at Equation 8.9 again we see that there exists a subsequence k_n such that

$$\sum_{k \in \mathcal{C}} |\langle x^{p(k,k_n)} - x^{k_n}, q^{p(k,k_n)} \rangle| \rightarrow 0 \text{ and } 0 \geq \limsup_n \langle c - x^{k_n}, \nabla f(x^0) - \nabla f(x^{k_n}) \rangle.$$

Let x^* be the accumulation point for this convergent subsequence (we may need to pass to another subsequence). Now we want to show that x^* is feasible.

Assume for the sake of contradiction that x^* is not feasible. Let \tilde{C} be the set of constraint sets that x^* belongs to, i.e., the set of constraints x^* satisfies. Recall that C is the set of all constraints, and let

$$\epsilon = \phi \left(\inf_{\hat{C} \in C \cap \tilde{C}^c} \text{dist}(x^*, \hat{C}) \right).$$

Let N be such that for all $k \geq N$ we have that $\|x^{n_k} - x^*\| < \epsilon/2$. Then by Equation 6.48 from Censor and Zenios (1997), since x^n is bounded, $x^{n_k} \rightarrow x^*$, and $D(x^{n+1}, x^n) \rightarrow 0$ we have that for any t ,

$$x^{n_k+t} \rightarrow x^*.$$

Thus, in particular, we see that for all $t \leq 2|\tilde{C}| + 2M + 1 =: T$ this is true.

Consider our augmented subsequence $x^{n_k}, x^{n_k+1}, \dots, x^{n_k+T}$ (note if $n_{k+1} - n_k \rightarrow \infty$ then this is not all points in the sequence). We want to show that infinitely often one of these values satisfies a constraint not in \tilde{C} . Then since we have finitely many constraints, at least one constraint \tilde{c} not in \tilde{C} must be satisfied infinitely often by our augmented sequence. Then since our augmented sequence converges to x^* and we are only looking at closed constraints, we must have x^* satisfies this constraint. Which is a contradiction. Thus, x^* must be feasible.

Let $k > N$. Then we know that x^{n_k} is the variable that we get after the n_k th projection. Let's assume that this happens in iteration ν_k . Note that in any iteration we project onto any constraint at most twice, once if its in our list C^{ν_k} and once if its in the list L returned by the oracle. Thus we have two possibilities for which iteration the $n_k + 2|\tilde{C}| + 1$ projection takes place. Let us case on when that happens.

Case 1: If $n_k + 2|\tilde{C}| + 1$ takes places in ν_k th iteration infinitely often, then in case since we project onto each constraint at most twice, by pigeonhole principle one of the projections between n_k and $n_k + 2|\tilde{C}| + 1$ must be onto a constraint not represented in \tilde{C} . Thus, if we look amongst $x^{n_k}, x^{n_k+1}, \dots, x^{n_k+2|\tilde{C}|+1}$, we must have projected onto a constraint not in \tilde{C} infinitely often.

Case 2: If $n_k + 2|\tilde{C}| + 1$ takes places in $\nu_k + 1$ th iteration infinitely often. Then since our oracle returns at most M constraints, we must have that amongst $x^{n_k}, x^{n_k+1}, \dots, x^{n_k+T}$ we must have projected onto all of the constraints returned by the oracle during step 1 of the $\nu_k + 1$ th iteration. Thus, we must have projected onto some constraint \hat{c}^{ν_k+1} such that distance of $x^{\nu+1}$ to \hat{c}^{ν_k+1} is at least $\phi(d_{\nu_k+1})$ where d_{ν_k+1} is the distance from x^{ν_k+1} to our feasible region C . Then, we have that

$$\text{dist}(x^{\nu_k+1}, \hat{c}^{\nu_k+1}) \geq \phi(d_{\nu_k+1}) \geq \phi\left(\inf_{\hat{c} \in C \cap \tilde{C}^c} \text{dist}(x^{\nu_k+1}, \hat{c})\right).$$

Noting that $x^{\nu+1}$ is within our augmented sequence and taking the limit, we get that

$$\lim_{k \rightarrow \infty} \text{dist}(x^{\nu_k+1}, \hat{c}^{\nu_k+1}) \geq \epsilon.$$

Since, we have finitely many constraints, we may assume (by passing to a subsequence) that all \hat{c}^{ν_k+1} are equal to some constraint \hat{c} . Then, we have that

$$\text{dist}(x^*, \hat{c}) \geq \epsilon.$$

Thus, $\hat{c} \notin \tilde{C}$. However, we now project onto \hat{c} infinitely often. Hence x^* must satisfy the constraint \hat{c} . Hence we have a contradiction. Thus, x^* is feasible.

Finally, we have that since $D_f(\cdot, \cdot)$ is separately continuous, we have that $D_f(c, x^{k_n}) \rightarrow D_f(c, x^*)$ and $D_f(x^{k_n}, x^0) \rightarrow D_f(x^*, x^0)$. Thus for an arbitrary feasible c we have that

$$\begin{aligned} \langle c - x^*, \nabla f(x^0) - \nabla f(x^*) \rangle &= D_f(c, x^*) + D_f(x^*, x^0) - D_f(c, x^0) \\ &= \lim_n D_f(c, x^{k_n}) + D_f(x^{k_n}, x^0) - D_f(c, x^0) \\ &= \lim_n \langle c - x^{k_n}, \nabla f(x^0) - \nabla f(x^{k_n}) \rangle \\ &\leq 0. \end{aligned}$$

Then by Fact 3, we have that $x^* = P_C^f(x^0)$. Thus, all accumulation points are optimal points. Thus, by strict convexity of f , we see that this is the unique solution.

Now to get the unique solution to constrained minimization problem we need to initialize x^0 such that $\nabla f(x^0) = 0$. This is because x^* is the point that minimized

$$D_f(x, x^0) = f(x) - f(x^0) - \langle \nabla f(x^0), x - x^0 \rangle.$$

Thus, if the last term is 0, then second term is a constant and we minimize the first term which is what we want.

8.6 Convergence Rate for Quadratic Objective Function

Lemma 1 *For any C_i such that $x^* \in \text{int}(C_i)$ we have that we only look at this constraint finitely often.*

Proof Every time we look at a constraint we project onto its boundary. Thus, if we project onto C_i infinitely often x^* cannot be in the interior of C_i . \blacksquare

Let I be the set of all active constraints. That is if x^* is the optimal solution then

$$I = \{i : x^* \in \partial C_i\}$$

Let S be the set of all x that satisfy these constraints (namely $S = \{x : \forall i \in I, x \in C_i\}$). Let B_x be the boundary for a constraint in I furthest from x and define

$$\mu = \inf_{x \notin S} \frac{d(x, B_x)}{d(x, S)} > 0$$

Let U be the set of all optimal duals q 's that is $U = \{(q_1, \dots, q_N) : \nabla f(x^0) - \nabla f(x^*) = \sum_i q_i\}$ and let $I_\nu = \{i : q^{p(i, \nu+1)} \neq 0\}$.

Lemma 2 *For any sequence $x^n \rightarrow x^*$ where x^* is the optimal for an instance of problem 2.3, and x^n, q^n maintain the KKT conditions then there exists an M' such that for all $\nu \geq M'$ there exists a $z \in U$ such that for all $i \notin I_\nu$ we have that $q_i = 0$*

Proof Let $V_\nu = \{(q_1, \dots, q_N) : \forall i \notin I_\nu, q_i = 0\}$. Then assume for the sake of contradiction that the result is false. Thus, there is a sequence ν_k such that $V_{\nu_k} \cap U = \emptyset$. Then since there finitely many different I_ν (hence finitely many V_ν) we have that one of these must occur infinitely often. Thus, by taking an appropriate subsequence we assume that without loss of generality that all I_{ν_k} are all equal. Thus, let $V = V_{\nu_k}$. Thus, we have that $V \cap U = \emptyset$.

Then since V is a subspace (hence closed), U is a closed set, and they are disjoint, we must have that $d(V, U) > 0$. But now $(q^{p(1, \nu_k+1)}, \dots, q^{p(N, \nu_k+1)}) \in V$. Thus,

$$d((q^{p(1, \nu_k+1)}, \dots, q^{p(N, \nu_k+1)}), U) \geq d(V, U) > 0$$

By Theorem 3 we have that $x^\nu \rightarrow x^*$. Thus, for any $(q_1, \dots, q_N) \in U$,

$$\sum_i q^{p(i, \nu_k+1)} = \nabla f(x^0) - \nabla f(x^\nu) \rightarrow \nabla f(x^0) - \nabla f(x^*) = \sum_i q_i$$

Thus $d((q^{p(1, \nu_k+1)}, \dots, q^{p(N, \nu_k+1)}), U) \rightarrow 0$ which is a contradiction \blacksquare

The rest of the proof is same as for Theorem 1. As before we get the following corollaries.

Corollary 2 *If we further have that there exists an $x^0 \in \text{dom} f$ such that $\nabla f(x^0) = 0$. Then we have that the Bregman projection of x^0 onto C is the solution to Problem 2.3.*

Remark 6 *Similar to before while this gives linear convergence but ρ is very close to one. As we will see later $\rho \leq \frac{F}{F+1}$ where F is number of convex sets whose boundary the optimal solution lies on that are seen by the algorithm.*

Remark 7 *Unlike in the linear case, our proof does not guarantee that a truly stochastic version will converge to the optimal solution. Equation 8.11 does not hold in this case.*

Acknowledgements

Rishi Sonthalia was partially supported by NSF grant DMS-2027277.

References

- Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating Inconsistent Information: Ranking and Clustering. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, 2005.
- MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 9.2.*, 2021.
- Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. In *43rd Symposium on Foundations of Computer Science (FOCS 2002)*, 2002.
- Heinz H. Bauschke and Jonathan M. Borwein. Legendre Functions and the Method of Random Bregman Projections. *Journal of Convex Analysis*, 1997.
- Heinz H. Bauschke and Adrian S. Lewis. Dykstra’s Algorithm with Bregman Projections: a Convergence Proof. *Optimization. A Journal of Mathematical Programming and Operations Research*, 2000.
- Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2009.
- Aurélien Bellet, Amaury Habrard, and Marc Sebban. A Survey on Metric Learning for Feature Vectors and Structured Data. *ArXiv*, 2013.
- Mathieu Blondel, Vivien Seguy, and Antoine Rolet. Smooth and Sparse Optimal Transport. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics (AISTATS 2018)*, 2018.
- Paul Boggs and Jon Tolle. Sequential Quadratic Programming. *Acta Numerica*, 1995.
- Lev M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 1967.
- Justin Brickell, Inderjit S. Dhillon, Suvrit Sra, and Joel A. Tropp. The Metric Nearness Problem. *SIAM Journal on Matrix Analysis and Applications*, 2008.
- Richard A. Brualdi. Combinatorial Matrix Classes. *Encyclopedia of Mathematics and its Applications*, 2006.
- Lawrence Cayton and Sanjoy Dasgupta. Robust Euclidean Embedding. In *Proceedings of the 23rd International Conference on Machine Learning (ICML 2006)*, 2006.
- Yair Censor and Simeon Reich. The Dykstra Algorithm with Bregman Projections. *Communications in Applied Analysis*, 1998.

- Yair Censor and Stavros Zenios. *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford University Press, 1997.
- Karthekeyan Chandrasekaran, László A. Végh, and Santosh Vempala. The Cutting Plane method is Polynomial for Perfect Matchings. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS 2012)*. 2012.
- Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with Qualitative Information. *Journal of Computer and System Sciences*, 71, 2005.
- Jason V. Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S. Dhillon. Information-Theoretic Metric Learning. In *Proceedings of the 24th International Conference on Machine Learning (ICML 2007)*, 2007.
- Santanu S. Dey and Marco Molinaro. Theoretical Challenges Towards Cutting-Plane Selection. *Mathematical Programming*, 2018.
- Inderjit S. Dhillon and Joel A. Tropp. Matrix Nearness Problems with Bregman Divergences. *SIAM Journal on Matrix Analysis and Applications*, 2007.
- Alexander Domahidi, Eric Chu, and Stephen Boyd. ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC 2013)*, 2013.
- Iain Dunning, Joey Huchette, and Miles Lubin. Jump: A modeling language for mathematical optimization. *SIAM Review*, 2017.
- Dotan Emanuel and Amos Fiat. Correlation Clustering—Minimizing Disagreements on Arbitrary Weighted Graphs. In *European Symposium on Algorithms (ESA 2003)*. 2003.
- Chenglin Fan, Benjamin Raichel, and Gregory Van Buskirk. Metric Violation Distance: Hardness and Approximation. In *Symposium on Discrete Algorithms (SODA 2018)*, 2018.
- Chenglin Fan, Anna C. Gilbert, Benjamin Raichel, Rishi Sonthalia, and Gregory Van Buskirk. Generalized Metric Repair on Graphs. In *17th Scandinavian Symposium and Workshops on Algorithm Theory, (SWAT 2020)*, 2020.
- Michael Garstka, Mark Cannon, and Paul Goulart. COSMO: A conic operator splitting method for large convex problems. In *European Control Conference (ECC 2019)*, 2019.
- A. C. Gilbert and L. Jain. If it Ain't Broke, Don't Fix it: Sparse Metric Repair. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton 2017)*, 2017.
- Anna C. Gilbert and Rishi Sonthalia. Unsupervised Metric Learning in Presence of Missing Data. *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton 2018)*, 2018.
- W. Glunt, T. L. Hayden, S. Hong, and J. Wells. An Alternating Projection Algorithm for Computing the Nearest Euclidean Distance Matrix. *SIAM Journal on Matrix Analysis and Applications*, 1990.

- Ralph Edward Gomory. An Algorithm for the Mixed Integer Problem, 1960.
- IBM ILOG. v10.10. cplex. *International Business Machines Corporation*, 2021.
- Alfredo N. Iusem. On Dual Convergence and the Rate of Primal Convergence of Bregman’s Convex Programming Method. *SIAM Journal on Optimization*, 1991.
- Alfredo N. Iusem and Alvaro R. De Pierro. On the Convergence Properties of Hildreth’s Quadratic Programming Algorithm. *Mathematical Programming*, 1990.
- Thorsten Joachims, Thomas Finley, and Chun-Nam John Yu. Cutting-Plane Training of Structural SVMs. *Machine Learning*, 77, 2009.
- Donald H. Kraft. A software package for sequential quadratic programming. 1988.
- Angelia Nedić. Random Algorithms for Convex Minimization Problems. *Mathematical Programming*, 2011.
- Yu. E. Nesterov. A Method for Solving the Convex Programming Problem with Convergence Rate $O(1/k^2)$. *Doklady Akademii Nauk SSSR*, 1983.
- Jorge Nocedal, Andreas Wächter, and Richard A. Waltz. Adaptive barrier update strategies for nonlinear interior methods. *SIAM Journal on Optimization*, 2009.
- B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. SCS: Splitting Conic Solver, version 2.1.2. <https://github.com/cvxgrp/scs>, 2019.
- Fernando Palacios-Gomez, Leon S. Lasdon, and Michael Engquist. Nonlinear Optimization by Successive Linear Programming. *Management Science*, 1982.
- Laurent Poirrier and James Yu. On the Depth of Cutting Planes. *arXiv e-prints*, art. arXiv:1903.05304, Mar 2019.
- Boris T. Polyak. Random Algorithms for Solving Convex Inequalities. In *Inherently parallel algorithms in feasibility and optimization and their applications (Haifa, 2000)*, Studies in Computational Mathematics. 2001.
- Houduo Qi and Xiaoming Yuan. Computing the nearest Euclidean distance matrix with low embedding dimensions. *Mathematical Programming*, 2014.
- Cameron Ruggles, Nate Veldt, and David F. Gleich. A parallel projection method for metric constrained optimization. In *Proceedings of the SIAM Workshop on Combinatorial Scientific Computing, CSC 2020, Seattle, USA, February 11-13, 2020*, 2020.
- Klaus Schittkowski. An active set strategy for solving optimization problems with up to 200,000,000 nonlinear constraints. *Applied Numerical Mathematics*, 2009.
- Mario Souto, Joaquim Dias Garcia, and Álvaro Veiga. Exploiting low-rank structure in semidefinite programming by approximate operator splitting. *Optimization*, 2020.
- B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: An operator splitting solver for quadratic programs. *Mathematical Programming Computation*, 2020.

- Juan Luis Suárez Díaz, Salvador García, and Francisco Herrera. A Tutorial on Distance Metric Learning: Mathematical Foundations, Algorithms and Software. *ArXiv eprint*, 2018.
- Yichuan Tang. Deep Learning Using Support Vector Machines. *ICML 2013 Challenges in Representation Learning*, abs/1306.0239, 2013.
- Nate Veldt, David Gleich, Anthony Wirth, and James Saunderson. Metric-Constrained Optimization for Graph Clustering Algorithms. *SIAM Journal on Mathematics of Data Science*, 2019.
- Mengdi Wang and Dimitri P. Bertsekas. Incremental Constraint Projection-Proximal Methods for Nonsmooth Convex Optimization. <http://www.mit.edu/~dimitrib/>, 2013.
- Mengdi Wang, Yichen Chen, Jialin Liu, and Yuantao Gu. Random Multi-Constraint Projection: Stochastic Gradient Methods for Convex Optimization with Many Constraints. *ArXiv*, 2015.
- Yubo Wang, Linli Xu, Yucheng Chen, and Hao Wang. A Scalable Approach for General Correlation Clustering. In *Advanced Data Mining and Applications*, 2013.