The Computational Performance of Iterated Linear Optimization

Teressa Chambers

Brown University

May 31, 2023

Abstract

Semidefinite programming (SDP) is frequently utilized to relax clustering problems, but its results are heavily dependent on the rounding step used to obtain a clustering. We will assess the viability of combining semidefinite programming with a deterministic rounding approach called iterated linear optimization as an alternative clustering method. This analysis will include examination of the accuracy of the algorithm via experiments on the MNIST dataset.

Background

To make combinatorial optimization problems more computationally feasible, semidefinite relaxations are often used to switch from the combinatorial solution space to the continuous solution space of a problem with semidefinite matrix constraints, such as the method by Goemans and Williamson which gives a 0.878-approximation for the MAX-CUT problem [4]. However, reaching a combinatorial solution from the semidefinite programming solution can prove challenging; while a randomized rounding scheme worked well for MAX-CUT, efforts to extend this method to the MAX k-CUT problem found that in that setting it could produce arbitrarily poor approximations of the solution [3]. Felzenszwalb, Klivans, and Paul have proposed iterated linear optimization (ILO) as a deterministic rounding step for semidefinite relaxations [2]. This approach is a fixed-point method, which can operate over any convex set $\triangle$ and iteratively solves for $T(x)$ until it reaches a point such that $x = T(x)$:

$$T(x) = \underset{y \in \triangle}{\operatorname{argmax}}\, x \cdot y$$

This can be applied to any combinatorial problems with semidefinite relaxations, including clustering problems, which seek a partition of n items into k groups according to some measure of optimality [1]. The relaxed solution space for this problem is the k-way elliptope, $\mathcal{L}_{n,k}$:

$$\mathcal{L}_{n,k} := \{X \in \mathcal{S}(n) | X \succeq 0, X_{ii} = 1, X_{ij} \geq -\frac{1}{k-1}\}$$

Methods & Results

The k-means clustering method was used as a baseline for performance assessment. Clustering was run on random samples of unprocessed images from the MNIST dataset and on lower-dimensional representations preprocessed via TensorFlow [5].

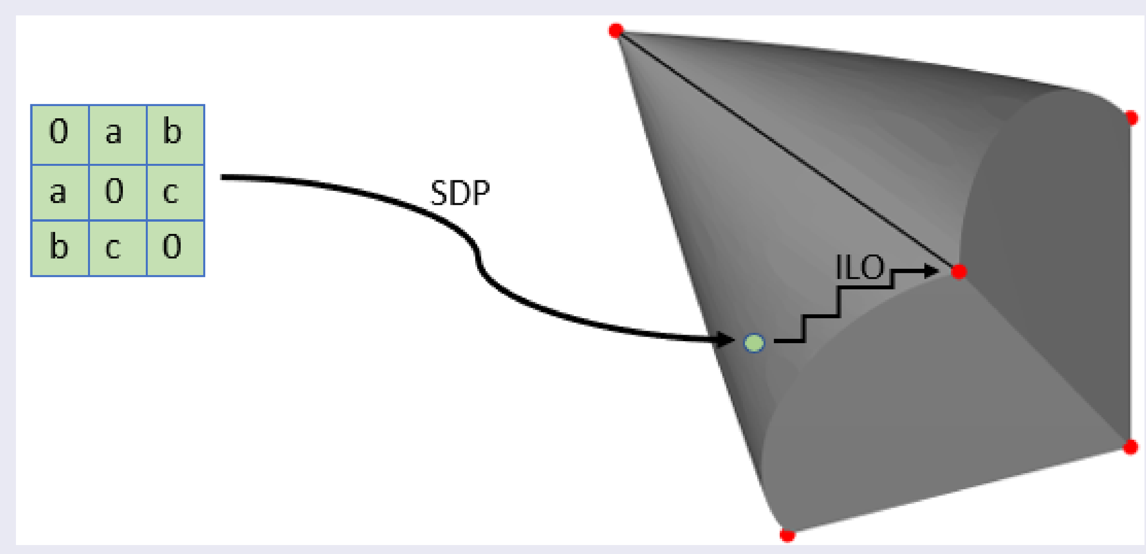

Figure: SDP-ILO strategy: From a n-by-n difference matrix, set up SDP and use ILO to round the result to a vertex of $\mathcal{L}_{n,k}$.

In the unprocessed tests, k-means performed more efficiently but SDP-ILO outscored its clustering outcome according to multiple measures, including adjusted Rand score and Davies-Bouldin score. With pre-processing, the methods had more similar scores, but SDP-ILO showed a stronger preference for the correct number of clusters ($k = 10$).

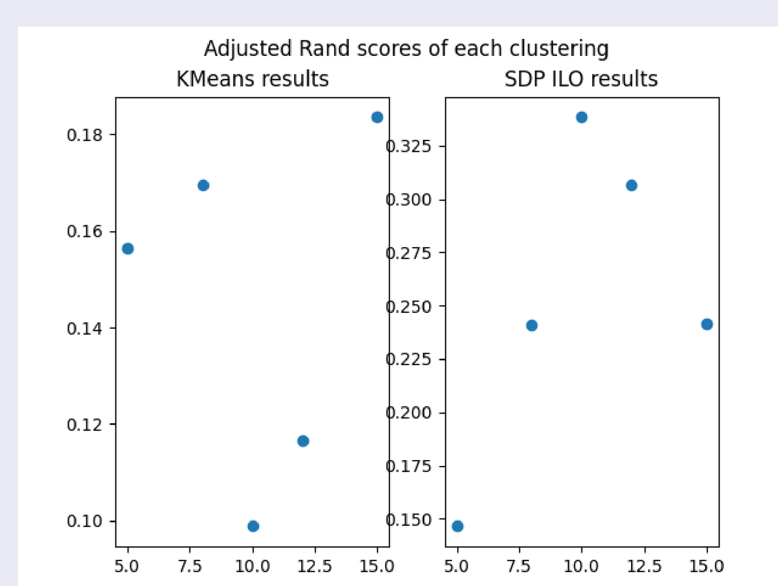

Figure: Sample Rand scores for clustering 50 raw images

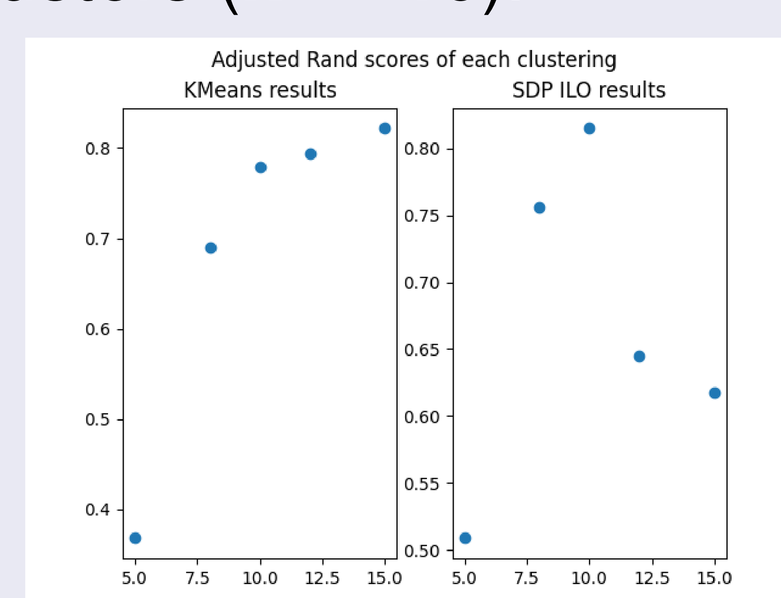

Figure: Rand for clustering 50 preprocessed images

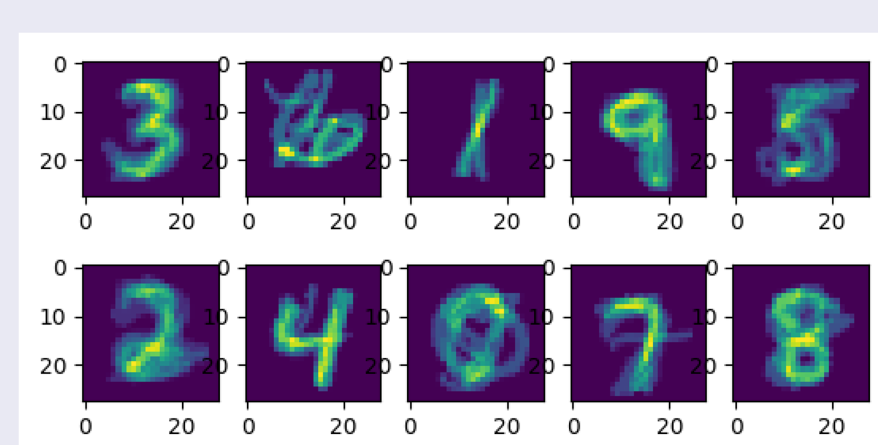

Figure: The image at the center of each SDP-ILO cluster using preprocessed data. All 10 digits are visibly identifiable.

Bibliography

[1] Pedro Felzenszwalb, Caroline Klivans, and Alice Paul. Clustering with Semidefinite Programming and Fixed Point Iteration. *arXiv e-prints*, page arXiv:2012.09202, Dec 2020.

[2] —. Iterated Linear Optimization. *arXiv e-prints*, page arXiv:2012.02213, Dec 2020.

[3] Alan Frieze and M. Jerrum. Improved approximation algorithms for max k-cut and max bisection. *Algorithmica*, 18:67-81, May 1997.

[4] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115-1145, 1995.

[5] Dustin G. Mixon, Soledad Villar, and Rachel Ward. Clustering subgaussian mixtures by semidefinite programming. *arXiv e-prints*, page arXiv:1602.06612, Feb 2016.