

MATRIX COMPLETION: AN OVERVIEW*

Anthony Austin
 Jose Garcia
 Stephen Jong
 Gilberto Hernandez

This work is produced by OpenStax-CNX and licensed under the
 Creative Commons Attribution License 3.0[†]

Abstract

This module is part of a collection of modules for a class project on matrix completion techniques for the sensor network localization problem done for the Fall, 2009 offering of Prof. Baraniuk's ELEC 301 course at Rice University.

1 Overview of Matrix Completion

The fundamental question that the new and emerging field of matrix completion seeks to answer is this: Given a matrix with some of its entries missing, is it possible to determine what those entries should be? Answering this question has an enormous number of potential practical applications. To be more concrete, consider the problem of *collaborative filtering*, of which perhaps the most famous example is the *Netflix* problem [6]. The *Netflix* problem asks how one may be able to predict how an individual would rate movies he or she has not seen based on the ratings that individual has made in the past and on the ratings of other individuals stored in the database. This can be cast as a matrix completion problem in which each row of the matrix corresponds to a particular user, each column to a movie, and each entry a rating that the user of that entry's row has given to the movie in that entry's column. Because there is a large number of users and movies and because each user has probably seen relatively few of the available movies, there are a large number of entries missing. The idea is to somehow fill in the missing entries and thereby determine how every user would rate every movie available. For more examples of potential uses of matrix completion, see the introduction of [2].

In general, matrix recovery is an impossible task because the unknown entries really could be anything; however, if one makes a few reasonable assumptions about the original matrix underlying the one being completed, then the matrix can indeed be reconstructed and often from a surprisingly low number of entries. More precisely, in their May, 2008 paper *Exact Matrix Completion via Convex Optimization*, matrix completion pioneers Emmanuel J. Candès and Benjamin Recht offer the following definitions [3]:

Definition: Let U be a subspace of \mathbb{R}^n of dimension r , and let P_U be the operator that projects orthogonally onto U . The *coherence* $\mu(U)$ of U is defined by

$$\mu(U) = \frac{n}{r} \max_{1 \leq i \leq n} \|P_U e_i\|^2, \quad (1)$$

*Version 1.1: Dec 16, 2009 3:36 pm -0600

[†]<http://creativecommons.org/licenses/by/3.0/>

where e_i is the standard basis vector with a 1 in the i^{th} coordinate and all other coordinates are zero.

Definition: Let A be an m -by- n matrix of rank r with singular value decomposition $\sum_{k=1}^r \sigma_k u_k v_k^*$, and denote its column and row spaces by U and V , respectively. A is said to be (μ_0, μ_1) -incoherent if

1. There exists $\mu_0 > 0$ such that $\max(\mu(U), \mu(V)) < \mu_0$.
2. There exists $\mu_1 > 0$ such that all entries of the m -by- n matrix $\sum_{k=1}^r u_k v_k^*$ are less than or equal to $\mu_1 \sqrt{\frac{r}{mn}}$ in magnitude.

Qualitatively, this definition means that the singular vectors of a (μ_0, μ_1) -incoherent matrix aren't too "spiky" and don't do anything "wild."

In the same paper, Candès and Recht go on to show that if A is an m -by- n (μ_0, μ_1) -incoherent matrix that has rank $r \ll N = \max(m, n)$, then A can be recovered with high probability from a uniform sampling of M of its entries, where $M \geq O(N^{1.2} r \log N)$ [3]. This result was later strengthened to $M \geq O(Nr \max(r, \log N))$ by Keshavan, Montanari, and Oh in [4]. These results, coupled with the fact that many matrices that one encounters in practice both satisfy the incoherence property and are of low rank means that matrix completion has some serious potential for use in practical applications.

Once one knows that matrix completion can be done, the next question is how to go about doing it. There are a variety of different matrix completion algorithms available. Candès et al. have developed a method that they call Singular Value Thresholding (SVT), which attempts to complete the matrix by solving the following optimization problem [1]: Find a matrix X of that minimizes $\|X\|_*$ subject to the condition that the entries of X be equal to those entries of the matrix A to be completed for which we know the value. Here, $\|X\|_*$ is the *nuclear norm* of X , defined to be the sum of the singular values of X . Keshavan, Montanari, and Oh offer an alternative algorithm, dubbed OptSpace, which is based on trimming the incomplete matrix to remove so-called "overrepresented" rows and columns whose values do not help reveal much about the unknown entries and then adjusting the trimmed matrix to minimize the error that is made at the entries whose values are known via a gradient descent procedure [4], [5]. There are other algorithms as well, and which algorithm to choose is really up to the user. For our work, we elected to use the OptSpace algorithm, since it just seems to produce better results.

References

- [1] J. Cai, E. J. Cand [U+FFFD] and Z. Shen. A singular value thresholding algorithm for matrix completion. *arXiv:0810.3286v1*, October 2008.
- [2] E. J. Cand [U+FFFD] and Y. Plan. Matrix completion with noise. *arXiv:0903.3131v1*, March 2009.
- [3] E. J. Cand [U+FFFD] and B. Recht. Exact matrix completion via convex optimization. *arXiv:0805.4471v1*, May 2008.
- [4] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *arXiv:0901.3150v4*, September 2009.
- [5] R. H. Keshavan and S. Oh. Optspace: A gradient descent algorithm on the grassman manifold for matrix completion. *arXiv:0910.5260v2*, November 2009.
- [6] Netflix. Netflix prize. <http://www.netflixprize.com//index>.