Low-Rank Matrix Completion problem
Truncated nuclear norm minimization algorithm

Background

The problem of recovering a matrix from partial observations, called a **matrix completion (MC) problem**, has attracted much attention in the last two decades due to a variety of applications, such as recommendation systems, image processing, localization of IoT networks, etc. Given a partially defined matrix

$$M_o = \begin{pmatrix} 0 & 1 & ? & 3 \\ ? & 2 & 3 & ? \\ 1 & ? & 1 & 1 \\ 2 & 1 & 5 & ? \end{pmatrix},$$

the MC problem is to determine the unknown entries ? such that one of the criteria is met:

1. The rank of the completion M is the lowest possible.
2. The sum of the singular values $\|M\|_*$ of M, called a *nuclear norm*, is the smallest possible.

In recommendation systems, such as Netflix, the rows of M_o represent users, the columns represent movies, and the ij entry represents user i's rating of movie j. Since users tend to share common interests, they will give similar ratings to movies, resulting in a completion M with low rank or low kernel norm. Based on this completion, the system then recommends a list of movies that the user might be interested in based on his previous ratings.

In image restroration when there is dirt present in the two-dimensional image represented by a matrix containing values of pixels, the idea is to use clean pixels as observed entries of the image and restore the image as a low rank or low nuclear norm matrix completion problem.

Since optimizing rank is a very hard problem, while optimizing the nuclear norm is more tractable, many MC algorithms are designed to solve the **nuclear norm minimization (NNM)** problem:

$$\begin{aligned} \min_{X \in \mathbb{R}^{n \times m}} \quad & \|X\|_*, \\ \text{subject to} \quad & P_\Omega(X) = P_\Omega(M_o), \end{aligned} \tag{1}$$

where

$$P_\Omega(A) = \begin{cases} a_{ij}, & (i,j) \in \Omega, \\ 0, & (i,j) \notin \Omega, \end{cases}$$

is a projection of the matrix A to

is a projection of the matrix A to entries from the *observed set* $\Omega \subseteq \{(i,j)\colon 1 \leq i \leq n, 1 \leq j \leq m\}$ and M_o is a given observation matrix. In case of large matrices (e.g., for recommendation systems) the problem (1) requires huge computational burden and thus require further reductions such as the one proposed in [1]:

$$\begin{aligned} \min_{X \in \mathbb{R}^{n\times m}} \quad & \|X\|_r \\ \text{subject to} \quad & P_\Omega(X) = P_\Omega(M_o), \end{aligned} \tag{2}$$

where $\|X\|_r = \sum_{i=r+1}^{n} \sigma_i(X)$ is the sum of all but the first r largest singular values of X.

The algorithms [1, Algorithm 1 and 2, p. 20], called **truncated nuclear norm minimization (TNNR-ADMM)**, which solve the optimization problem in (2), rely on the truncated singular value decomposition of the iterate at each step of the procedure until convergence.

Task

1. Study the algorithm and briefly explain the idea for each of the steps.

   FOOT Here you are not expected to reproduce the proofs from [1], only to understand the idea behind each step of the algorithm

2. Implement the algorithm and try it out on the image reconstruction problem on an image of your choice that is noisy with the text written over it, and then reconstruct the original content as an NNM problem. See [1, Figure 7, p. 2127].

3. Present some experimental results on the ratio of text noise to still get a good reconstruction.

References

[1] Y. Hu, D. Zhan, J. Ye, X. Li, X. He, *Fast and accurate matrix completion via a truncated nuclear norm regularization,* IEEE Trans. Pattern Anal. Mach. Intell. 35, 2013, 2117–2130. `https://arxiv.org/pdf/0810.3286.pdf`

[2] L.T. Nguyen, J. Kim, B. Shim, *Low-Rank Matrix Completion: A Contemporary Survey,* IEEE Access 7, 2019, 94215–94237. `https://arxiv.org/pdf/1907.11705.pdf`