# Maximum Margin Matrix Factorization using Smooth Semidefinite Optimization 

Alexandre d'Aspremont, Nathan Srebro
ORFE, Princeton University \& CS, University of Toronto

Thanks to Yurii Nesterov for numerous suggestions!

## Introduction

- Users assign ratings to a certain number of movies:

|  |  | 2 |  | 1 |  |  | 4 |  |  |  | 5 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | 5 |  | 4 |  |  |  | ? |  | 1 |  | 3 |
|  |  |  | 3 |  | 5 |  |  | 2 |  |  |  |  |
|  | 4 |  |  | ? |  |  | 5 |  | 3 |  | ? |  |
|  |  |  | 4 |  | 1 | 3 |  |  |  | 5 |  |  |
|  |  |  |  | 2 |  |  |  | 1 | ? |  |  | 4 |
|  |  | 1 |  |  |  |  | 5 | 5 | 5 |  | 4 |  |
| $\stackrel{N}{\omega}$ |  |  | 2 |  | ? | 5 |  | ? |  | 4 |  |  |
| $\stackrel{N}{0}$ |  | 3 |  | 3 |  | 1 |  | 5 |  | 2 |  | 1 |
|  |  | 3 |  |  |  | 1 |  |  | 2 |  | 3 |  |
|  |  | 4 |  |  | 5 | 1 |  |  | 3 |  |  |  |
|  |  |  | 3 |  |  |  | 3 | 3? |  |  | 5 |  |
|  | 2 | ? |  | 1 |  | 1 |  |  |  |  |  |  |
|  |  |  | 5 |  |  | 2 | ? |  | 4 |  | 4 |  |
|  |  | 1 |  | 3 |  | 1 | 5 | 5 | 4 |  | 5 |  |
|  | 1 |  | 2 |  |  | 4 |  |  |  | 5 | ? |  |
| Movies |  |  |  |  |  |  |  |  |  |  |  |  |

- Objective: make recommendations for other movies. . .


## Collaborative prediction

- Infer user preferences and movie features from user ratings.
- We use a linear prediction model:

$$
\operatorname{rating}_{i j}=u_{i}^{T} v_{j}
$$

where $u_{i}$ represents user characteristics and $v_{j}$ movie features.

- This makes collaborative prediction a matrix factorization problem
- Overcomplete representation. . .


## Collaborative prediction

- Inputs: a matrix of ratings $M_{i j}=\{-1,+1\}$ for $(i, j) \in S$, where $S$ is a subset of all possible user/movies combinations.
- We look for a linear model by factorizing $M \in \mathbf{R}^{n \times m}$ as:

$$
M=U^{T} V
$$

where $U \in \mathbf{R}^{n \times k}$ represents user characteristics and $V \in \mathbf{R}^{k \times m}$ movie features.

- Parsimony. . We want $k$ to be as small as possible.
- Output: a matrix $X \in \mathbf{R}^{n \times m}$ which is a low-rank approximation of the ratings matrix $M$.


## Least-Squares

- Choose Means Squared Error as measure of discrepancy.
- Suppose $S$ is the full set, our problem becomes:

$$
\min _{\{X: \operatorname{Rank}(X)=k\}}\|X-M\|^{2}
$$

- This is just a singular value decomposition (SVD). . .

Problem: Not true when $S$ is not the full set (partial observations). Also, MSE not a good measure of prediction performance. . .

## Soft Margin

$$
\operatorname{minimize} \quad \operatorname{Rank}(X)+c \sum_{(i, j) \in S} \max \left(0,1-X_{i j} M_{i j}\right)
$$

non-convex and numerically hard. . .

- Relaxation result in Fazel, Hindi \& Boyd (2001): replace Rank $(X)$ by its convex envelope on the spectahedron to solve:

$$
\operatorname{minimize}\|X\|_{*}+c \sum_{(i, j) \in S} \max \left(0,1-X_{i j} M_{i j}\right)
$$

where $\|X\|_{*}$ is the nuclear norm, i.e. sum of the singular values of $X$.

- Srebro (2004): This relaxation also corresponds to multiple large margin SVM classifications.


## Soft Margin

- The dual of this program:

$$
\begin{array}{cl}
\underset{\operatorname{maximize}}{ } & \sum_{i j} Y_{i j} \\
\text { subject to } & \|Y \odot M\|_{2} \leq 1 \\
& 0 \leq Y_{i j} \leq c
\end{array}
$$

in the variable $Y \in \mathbf{R}^{n \times m}$, where $Y \odot M$ is the Schur (componentwise) product of $Y$ and $M$ and $\|Y\|_{2}$ the largest singular value of $Y$.

- This problem is sparse: $Y_{i j}^{*}=c$ for $(i, j) \in S^{c}$


## Semidefinite Program

- How do we solve it?
- Rewrite the dual

$$
\begin{array}{ll}
\operatorname{maximize} & \sum_{i j} Y_{i j} \\
\text { subject to } & \|Y \odot M\|_{2} \leq 1 \\
& 0 \leq Y_{i j} \leq c
\end{array}
$$

as:

$$
\begin{array}{llc}
\operatorname{maximize} & \sum_{i j} Y_{i j} & \\
\text { subject to } & {\left[\begin{array}{cc}
I & -(Y \odot M) \\
-(Y \odot M)^{T} & I
\end{array}\right] \succeq 0} \\
& 0 \leq Y_{i j} \leq c &
\end{array}
$$

which is a sparse semidefinite program in $Y \in \mathbf{R}^{n \times m}$.

## Complexity

Complexity?

- Small subset $S$ : the dual in $Y$ is sparse, primal (in ratings $X$ ) is dense.
- Interior point solvers work fine for problem sizes up to $400 \ldots$
- We need to solve much larger instances.
- High precision is not necessary. . .


## Smoothing Technique

- Solution, formulate this as a saddle problem using binary search:

$$
\begin{array}{ll}
\operatorname{minimize} & \lambda^{\max }\left(\left[\begin{array}{cc}
I & -(Y \odot M) \\
-(Y \odot M)^{T} & I
\end{array}\right]\right) \\
\text { subject to } & \sum_{i j} Y_{i j}=t \\
& 0 \leq Y_{i j} \leq c
\end{array}
$$

for some $t>0$.

- Use the smoothing technique in Nesterov (2005): first-order algorithm with optimal complexity of $O(1 / \epsilon)$.
- Homogeneity means we also get a solution to:

$$
\begin{array}{ll}
\operatorname{maximize} & \sum_{i j} Y_{i j} \\
\text { subject to } & \|Y \odot M\|_{2} \leq 1 \\
& 0 \leq Y_{i j} \leq c^{*}
\end{array}
$$

## Nesterov's method

Assuming problem has a particular min-max structure:

- Regularization. Add strongly convex penalty inside the min-max representation to produce an $\epsilon$-approximation of $f$ with Lipschitz continuous gradient (generalized Moreau-Yosida regularization step, see Lemaréchal \& Sagastizábal (1997) for example).
- Optimal first order minimization. Use optimal first order scheme for Lipschitz continuous functions detailed in Nesterov (1983) to the solve the regularized problem.

Caveat: Only efficient if the subproblems involved in these steps can be solved explicitly or very efficiently. . . Change of granularity: larger number of cheaper iterations.

## Regularization

Replace $\lambda^{\max }(X)$ by

$$
f_{\mu}(X)=\mu \log \left(\sum_{i=1}^{k} e^{\frac{\lambda_{i}}{\mu}}\right)
$$

For a good choice of $\mu$ :

- $f_{\mu}(X)$ is an $\epsilon$-approximation of $f$.
- $f_{\mu}(X)$ has a Lipschitz continuous gradient with constant $L=O(1 / \epsilon)$.


## First-Order Minimization

The minimization algorithm in Nesterov (1983) then involves the following steps:

Choose $\epsilon>0$ and set $X_{0}=\beta I_{n}$, For $k=0, \ldots, N(\epsilon)$ do

1. Compute $f_{\mu}$ and $\nabla f_{\mu}$
2. Find

$$
Y_{k}=\arg \min _{Y}\left\{\operatorname{Tr}\left(\nabla f_{\epsilon}\left(X_{k}\right)\left(Y-X_{k}\right)\right)+\frac{1}{2} L_{\epsilon}\left\|Y-X_{k}\right\|_{F}^{2}: Y \in \mathcal{Q}_{1}\right\} .
$$

3. Find $Z_{k}=$
$\arg \min _{X}\left\{L_{\epsilon} \beta^{2}\|X\|+\sum_{i=0}^{k} \frac{i+1}{2} \operatorname{Tr}\left(\nabla f_{\epsilon}\left(X_{i}\right)\left(X-X_{i}\right)\right): X \in \mathcal{Q}_{1}\right\}$.
4. Update $X_{k}=\frac{2}{k+3} Z_{k}+\frac{k+1}{k+3} Y_{k}$.

## Numerical Cost

At each iteration:

- Step 1: computes $f$ and $\nabla f$ and is a (full) eigenvalue decomposition (in fact SVD here, because of structure)
- Step 2 \& 3: involve projections on a the set:

$$
\mathcal{Q}_{1}=\left\{Y: \sum_{i j} Y_{i j}=t, 0 \leq Y_{i j} \leq c\right\}
$$

and are numerically easy.
Complexity, i.e. maximum number of iterations to reach absolute precision $\epsilon$

$$
\frac{4 \sqrt{m+n+m n c^{2}}}{\epsilon}
$$

with each iteration (roughly) costing $O\left(m n^{2}+n^{3}\right)$.

## Numerical Results

- No movies to recommend but. . .
- Compare CPU time and memory usage for CSDP and smooth optimization code.
- Both codes are C/MEX with calls to (dense) LAPACK/BLAS.


## Numerical Results



Figure 1: CPU time and memory usage versus $n$.

## Numerical Results

Large scale tests on a 3,06 Ghz CPU with 2Gb RAM:

| n | $1 \%$ observed | $10 \%$ observed | $50 \%$ observed |
| ---: | :---: | :---: | :---: |
| 100 | 2 sec | 3 sec | 10 sec |
| 178 | 2 sec | 18 sec | 35 sec |
| 316 | 19 sec | $2: 34 \mathrm{~min}$ | $2: 41 \mathrm{~min}$ |
| 562 | $3: 27 \mathrm{~min}$ | $3: 37 \mathrm{~min}$ | $19: 11 \mathrm{~min}$ |
| 1000 | $34: 35 \mathrm{~min}$ | $41: 15 \mathrm{~min}$ | $1: 35: 28$ hours |
| 1778 | $5: 44: 07$ hours | $6: 40: 06$ hours | $19: 09: 49$ hours |
| 3162 | $57: 23: 09$ hours | $67: 35: 34$ hours | $62: 12: 21$ hours |

## References

Fazel, M., Hindi, H. \& Boyd, S. (2001), 'A rank minimization heuristic with application to minimum order system approximation', Proceedings American Control Conference 6, 4734-4739.

Lemaréchal, C. \& Sagastizábal, C. (1997), 'Practical aspects of the Moreau-Yosida regularization: theoretical preliminaries', SIAM Journal on Optimization 7(2), 367-385.

Nesterov, Y. (1983), 'A method of solving a convex programming problem with convergence rate $O\left(1 / k^{2}\right)^{\prime}$, Soviet Mathematics Doklady 27(2), 372-376.

Nesterov, Y. (2005), 'Smooth minimization of nonsmooth functions', Mathematical Programming, Series A 103, 127-152.

Srebro, N. (2004), Learning with Matrix Factorization, PhD thesis, Massachusetts Institute of Technology.

