Majorization Minimization (MM) and Block Coordinate Descent (BCD)

Wing-Kin (Ken) Ma

Department of Electronic Engineering,
The Chinese University Hong Kong, Hong Kong

ELEG5481, Lecture 15

Acknowledgment: Qiang Li for helping prepare the slides.

Outline

- Majorization Minimization (MM)
  - Convergence
  - Applications
- Block Coordinate Descent (BCD)
  - Applications
  - Convergence
- Summary

Majorization Minimization

Consider the following problem

$$\min_{\boldsymbol{x}} f(\boldsymbol{x}) \quad \text{s.t. } \boldsymbol{x} \in \mathcal{X} \tag{1}$$

where $\mathcal{X}$ is a closed convex set; $f(\cdot)$ may be non-convex and/or nonsmooth.

- **Challenge**: For a general $f(\cdot)$, problem (1) can be difficult to solve.

- **Majorization Minimization**: Iteratively generate $\{\boldsymbol{x}^r\}$ as follows

$$\boldsymbol{x}^r \in \min_{\boldsymbol{x}} u(\boldsymbol{x}, \boldsymbol{x}^{r-1}) \quad \text{s.t. } \boldsymbol{x} \in \mathcal{X} \tag{2}$$

where $u(\boldsymbol{x}, \boldsymbol{x}^{r-1})$ is a surrogate function of $f(\boldsymbol{x})$, satisfying

1. $u(\boldsymbol{x}, \boldsymbol{x}^r) \geq f(\boldsymbol{x}), \quad \forall \boldsymbol{x}^r, \boldsymbol{x} \in \mathcal{X}$;
2. $u(\boldsymbol{x}^r, \boldsymbol{x}^r) = f(\boldsymbol{x}^r)$;

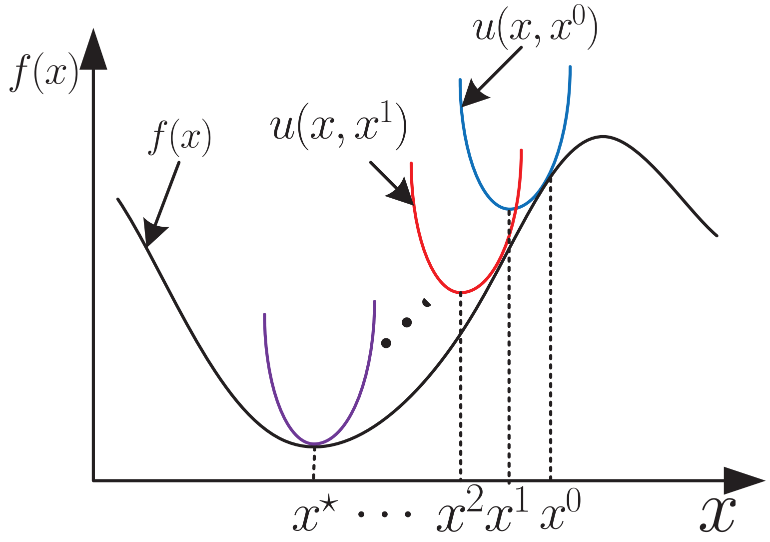

Figure 1: An pictorial illustration of MM algorithm.

Property 1. $\{f(\boldsymbol{x}^r)\}$ *is nonincreasing, i.e.,* $f(\boldsymbol{x}^r) \leq f(\boldsymbol{x}^{r-1}),\ \forall r = 1, 2, \ldots$.

Proof. $f(\boldsymbol{x}^r) \leq u(\boldsymbol{x}^r, \boldsymbol{x}^{r-1}) \leq u(\boldsymbol{x}^{r-1}, \boldsymbol{x}^{r-1}) = f(\boldsymbol{x}^{r-1})$ ■

- The nonincreasing property of $\{f(\boldsymbol{x}^r)\}$ implies that $f(\boldsymbol{x}^r) \to \bar{f}$. But how about the convergence of the iterates $\{\boldsymbol{x}^r\}$?

Technical Preliminaries

- **Limit point**: $\bar{x}$ is a limit point of $\{x_k\}$ if there exists a subsequence of $\{x_k\}$ that converges to $\bar{x}$. Note that every bounded sequence in $\mathbb{R}^n$ has a limit point (or convergent subsequence);

- **Directional derivative**: Let $f : \mathcal{D} \to \mathbb{R}$ be a function where $\mathcal{D} \subseteq \mathbb{R}^m$ is a convex set. The directional derivative of f at point x in direction $\boldsymbol{d}$ is defined by

$$f'(\boldsymbol{x}; \boldsymbol{d}) \triangleq \liminf_{\lambda \downarrow 0} \frac{f(\boldsymbol{x} + \lambda \boldsymbol{d}) - f(\boldsymbol{x})}{\lambda}.$$

  – If f is differentiable, then $f'(\boldsymbol{x}; \boldsymbol{d}) = \boldsymbol{d}^T \nabla f(\boldsymbol{x})$.

- **Stationary point**: $\boldsymbol{x} \in \mathcal{X}$ is a stationary point of $f(\cdot)$ if

$$f'(\boldsymbol{x}; \boldsymbol{d}) \geq 0, \ \forall \boldsymbol{d} \text{ such that } \boldsymbol{x} + \boldsymbol{d} \in \mathcal{D}. \quad (3)$$

  – A stationary point may be a local min., a local max. or a saddle point;
  – If $\mathcal{D} = \mathbb{R}^n$ and f is differentiable, then (3) $\Longleftrightarrow \nabla f(\boldsymbol{x}) = \boldsymbol{0}$.

Convergence of MM

- **Assumption 1** $u(\cdot,\cdot)$ satisfies the following conditions

$$
\begin{cases}
u(\boldsymbol{y},\boldsymbol{y}) = f(\boldsymbol{y}), \quad \forall \boldsymbol{y} \in \mathcal{X}, & \text{(4a)}\\
u(\boldsymbol{x},\boldsymbol{y}) \geq f(\boldsymbol{x}), \quad \forall \boldsymbol{x},\boldsymbol{y} \in \mathcal{X}, & \text{(4b)}\\
u'(\boldsymbol{x},\boldsymbol{y};\boldsymbol{d})|_{\boldsymbol{x}=\boldsymbol{y}} = f'(\boldsymbol{y};\boldsymbol{d}), \quad \forall \boldsymbol{d} \text{ with } \boldsymbol{y}+\boldsymbol{d} \in \mathcal{X}, & \text{(4c)}\\
u(\boldsymbol{x},\boldsymbol{y}) \text{ is continuous in } (\boldsymbol{x},\boldsymbol{y}), & \text{(4d)}
\end{cases}
$$

- (4c) means the 1st order local behavior of $u(\cdot,\boldsymbol{x}^{r-1})$ is the same as $f(\cdot)$.

Convergence of MM

Theorem 1. [Razaviyayn-Hong-Luo] *Assume that Assumption 1 is satisfied. Then every limit point of the iterates generated by MM algorithm is a stationary point of problem* (1)*.*

Proof. From **Property 1**, we know that $f(\boldsymbol{x}^{r+1}) \leq u(\boldsymbol{x}^{r+1}, \boldsymbol{x}^r) \leq u(\boldsymbol{x}, \boldsymbol{x}^r), \ \forall \boldsymbol{x} \in \mathcal{X}$. Now assume that there exists a subsequence $\{\boldsymbol{x}^{r_j}\}$ of $\{\boldsymbol{x}^r\}$ converging to a limit point $\boldsymbol{z}$, i.e., $\lim_{j\to\infty} \boldsymbol{x}^{r_j} = \boldsymbol{z}$. Then

$$u(\boldsymbol{x}^{r_{j+1}}, \boldsymbol{x}^{r_{j+1}}) = f(\boldsymbol{x}^{r_{j+1}}) \leq f(\boldsymbol{x}^{r_j+1}) \leq u(\boldsymbol{x}^{r_j+1}, \boldsymbol{x}^{r_j}) \leq u(\boldsymbol{x}, \boldsymbol{x}^{r_j}), \ \forall \boldsymbol{x} \in \mathcal{X}.$$

Letting $j \to \infty$, we obtain $u(\boldsymbol{z}, \boldsymbol{z}) \leq u(\boldsymbol{x}, \boldsymbol{z}), \quad \forall \boldsymbol{x} \in \mathcal{X}$, which implies that

$$u'(\boldsymbol{x}, \boldsymbol{z}; \boldsymbol{d})|_{\boldsymbol{x}=\boldsymbol{z}} \geq 0, \quad \forall \boldsymbol{z} + \boldsymbol{d} \in \mathcal{X}.$$

Combining the above inequality with (4c) (i.e., $u'(\boldsymbol{x}, \boldsymbol{y}; \boldsymbol{d})|_{\boldsymbol{x}=\boldsymbol{y}} = f'(\boldsymbol{y}; \boldsymbol{d}), \quad \forall \boldsymbol{d} \text{ with } \boldsymbol{y} + \boldsymbol{d} \in \mathcal{X}$), we have

$$f'(\boldsymbol{z}; \boldsymbol{d}) \geq 0, \quad \forall \boldsymbol{z} + \boldsymbol{d} \in \mathcal{X}.$$

■

Applications — Nonnegative Least Squares

In many engineering applications, we encounter the following problem

$$\text{(NLS)} \quad \min_{\boldsymbol{x} \geq \mathbf{0}} \ \|\mathbf{A}\boldsymbol{x} - \boldsymbol{b}\|_2^2 \tag{5}$$

where $\boldsymbol{b} \in \mathbb{R}_+^m$, $\boldsymbol{b} \neq \mathbf{0}$, and $\mathbf{A} \in \mathbb{R}_{++}^{m \times n}$.

- It's an LS problem with nonnegative constraints, so the conventional LS solution may not be feasible for (5).

- A simple multiplicative updating algorithm:

$$\boldsymbol{x}_l^r = c_l^r \boldsymbol{x}_l^{r-1}, \quad l = 1, \ldots, n \tag{6}$$

  where $\boldsymbol{x}_l^r$ is the lth component of $\boldsymbol{x}^r$, and $c_l^r = \frac{[\mathbf{A}^T \boldsymbol{b}]_l}{[\mathbf{A}^T \mathbf{A} \boldsymbol{x}^{r-1}]_l}$.

- Starting with $\boldsymbol{x}^0 > \mathbf{0}$, then all $\boldsymbol{x}^r$ generated by (6) are nonnegative.

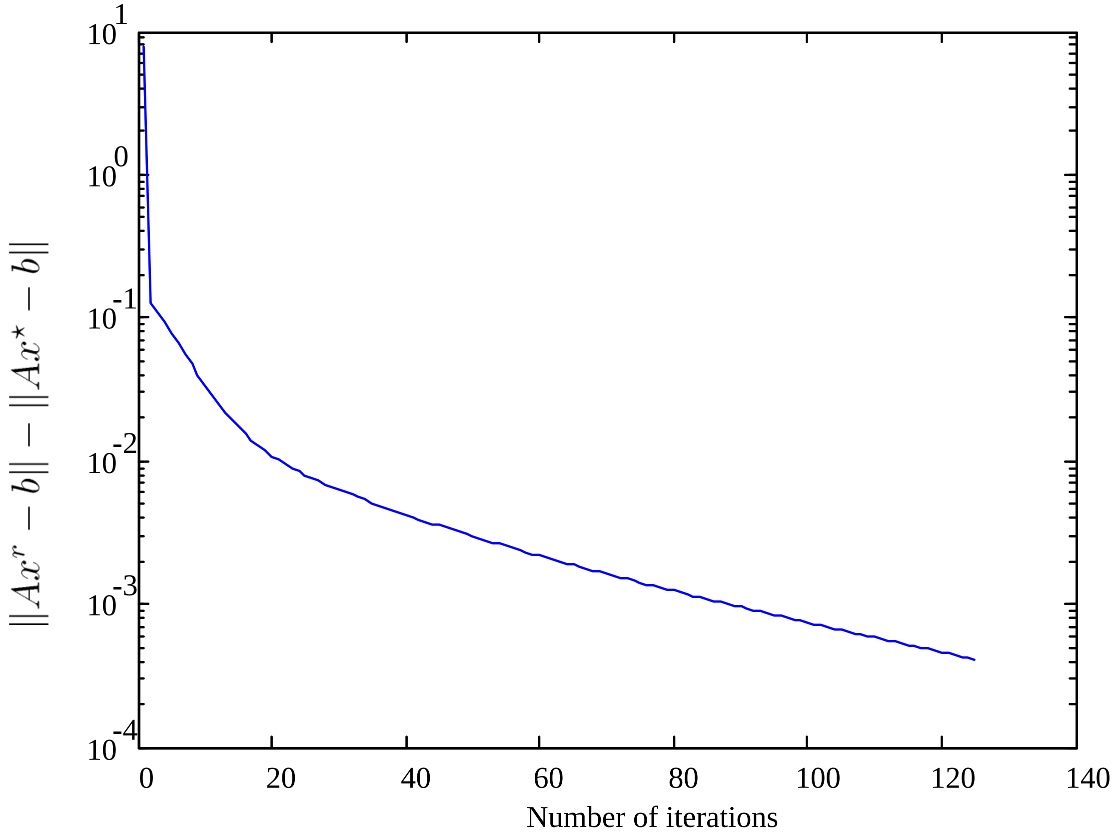

Figure 2: $\|\mathbf{A}\boldsymbol{x}^r - \boldsymbol{b}\|_2$ vs. the number of iterations.

- Usually the multiplicative update converges within a few tens of iterations.

- **MM interpretation**: Let $f(\boldsymbol{x}) \triangleq \|\mathbf{A}\boldsymbol{x} - \boldsymbol{b}\|_2^2$. The multiplicative update essentially solves the following problem

$$\min_{\boldsymbol{x} \geq \mathbf{0}} \; u(\boldsymbol{x}, \boldsymbol{x}^{r-1})$$

  where

$$u(\boldsymbol{x}, \boldsymbol{x}^{r-1}) \triangleq f(\boldsymbol{x}^{r-1}) + (\boldsymbol{x} - \boldsymbol{x}^{r-1})^T \nabla f(\boldsymbol{x}^{r-1}) + \frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}^{r-1})^T \boldsymbol{\Phi}(\boldsymbol{x}^{r-1})(\boldsymbol{x} - \boldsymbol{x}^{r-1}),$$

$$\boldsymbol{\Phi}(\boldsymbol{x}^{r-1}) = \mathrm{Diag}\left(\frac{[\mathbf{A}^T\mathbf{A}\boldsymbol{x}^{r-1}]_1}{\boldsymbol{x}_1^{r-1}}, \ldots, \frac{[\mathbf{A}^T\mathbf{A}\boldsymbol{x}^{r-1}]_n}{\boldsymbol{x}_n^{r-1}}\right).$$

  – Observations:

$$\begin{cases} u(\boldsymbol{x}, \boldsymbol{x}^{r-1}) \;\; \text{is quadratic approx. of } f(\boldsymbol{x}), \\ \boldsymbol{\Phi}(\boldsymbol{x}^{r-1}) \succeq \mathbf{A}^T\mathbf{A}, \end{cases} \Longrightarrow \begin{cases} u(\boldsymbol{x}, \boldsymbol{x}^{r-1}) \geq f(\boldsymbol{x}), \;\; \forall \boldsymbol{x} \in \mathbb{R}^n, \\ u(\boldsymbol{x}^{r-1}, \boldsymbol{x}^{r-1}) = f(\boldsymbol{x}^{r-1}). \end{cases}$$

- The multiplicative update converges to an optimal solution of NLS (by the MM convergence in **Theorem 1** and convexity of NLS).

Applications — Convex-Concave Procedure/ DC Programming

- Suppose that $f(\boldsymbol{x})$ has the following form

$$f(\boldsymbol{x}) = g(\boldsymbol{x}) - h(\boldsymbol{x}),$$

where $g(\boldsymbol{x})$ and $h(\boldsymbol{x})$ are convex and differentiable. Thus, $f(\boldsymbol{x})$ is in general nonconvex.

- **DC Programming:** Construct $u(\cdot, \cdot)$ as

$$u(\boldsymbol{x}, \boldsymbol{x}^r) = g(\boldsymbol{x}) - \underbrace{\left(h(\boldsymbol{x}^r) + \nabla_{\boldsymbol{x}} h(\boldsymbol{x}^r)^T (\boldsymbol{x} - \boldsymbol{x}^r)\right)}_{\text{linearization of } h \text{ at } x^r}$$

- By the 1st order condition of $h(\boldsymbol{x})$, it's easy to show that

$$u(\boldsymbol{x}, \boldsymbol{x}^r) \geq f(\boldsymbol{x}), \ \forall \boldsymbol{x} \in \mathcal{X}, \qquad u(\boldsymbol{x}^r, \boldsymbol{x}^r) = f(\boldsymbol{x}^r).$$

- **Sparse Signal Recovery by DC Programming**

$$\min_{\boldsymbol{x}} \ \|\boldsymbol{x}\|_0 \quad \text{s.t. } \boldsymbol{y} = \mathbf{A}\boldsymbol{x} \tag{7}$$

  – Apart from the popular ℓ_1 approximation, consider the following concave approximation:

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \sum_{i=1}^{n} \log(1 + |x_i|/\epsilon) \quad \text{s.t. } \boldsymbol{y} = \mathbf{A}\boldsymbol{x},$$

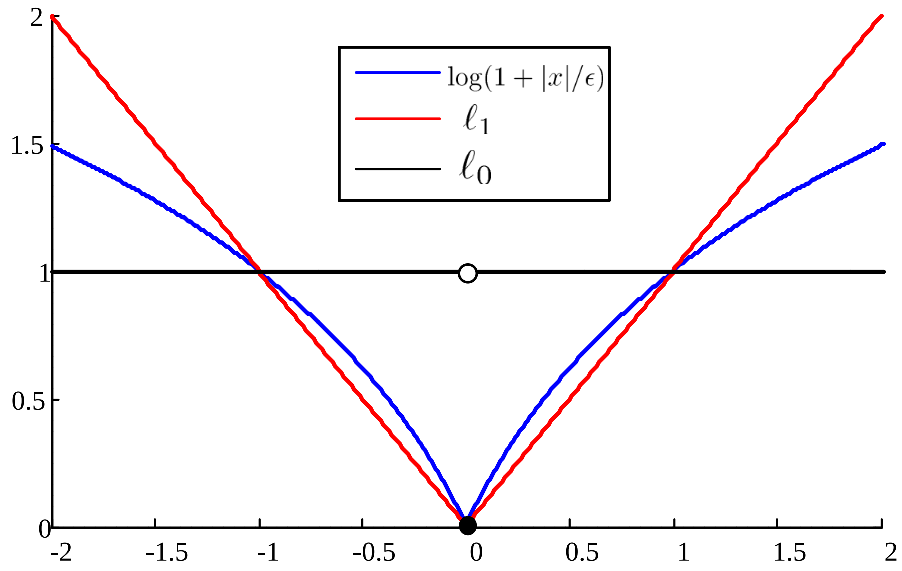

Figure 3: $\log(1 + |x|/\epsilon)$ promotes more sparsity than ℓ_1

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \sum_{i=1}^{n} \log(1+|x_i|/\epsilon) \quad \text{s.t. } \boldsymbol{y}=\mathbf{A}\boldsymbol{x},$$

which can be equivalently written as

$$\min_{\boldsymbol{x},\boldsymbol{z}\in\mathbb{R}^n} \sum_{i=1}^{n} \log(z_i+\epsilon) \quad \text{s.t. } \boldsymbol{y}=\mathbf{A}\boldsymbol{x},\ |x_i|\le z_i,\ i=1,\ldots,n \tag{8}$$

– Problem (8) minimizes a concave objective, so it's a special case of DC programming ($g(\boldsymbol{x})=0$). Linearizing the concave function at $(\boldsymbol{x}^r,\boldsymbol{z}^r)$ yields

$$(\boldsymbol{x}^{r+1},\boldsymbol{z}^{r+1}) = \arg\min \sum_{i=1}^{n} \frac{z_i}{z_i^r+\epsilon} \quad \text{s.t. } \boldsymbol{y}=\mathbf{A}\boldsymbol{x},\ |x_i|\le z_i,\ i=1,\ldots,n$$

– We solve a sequence of reweighted ℓ_1 problems.

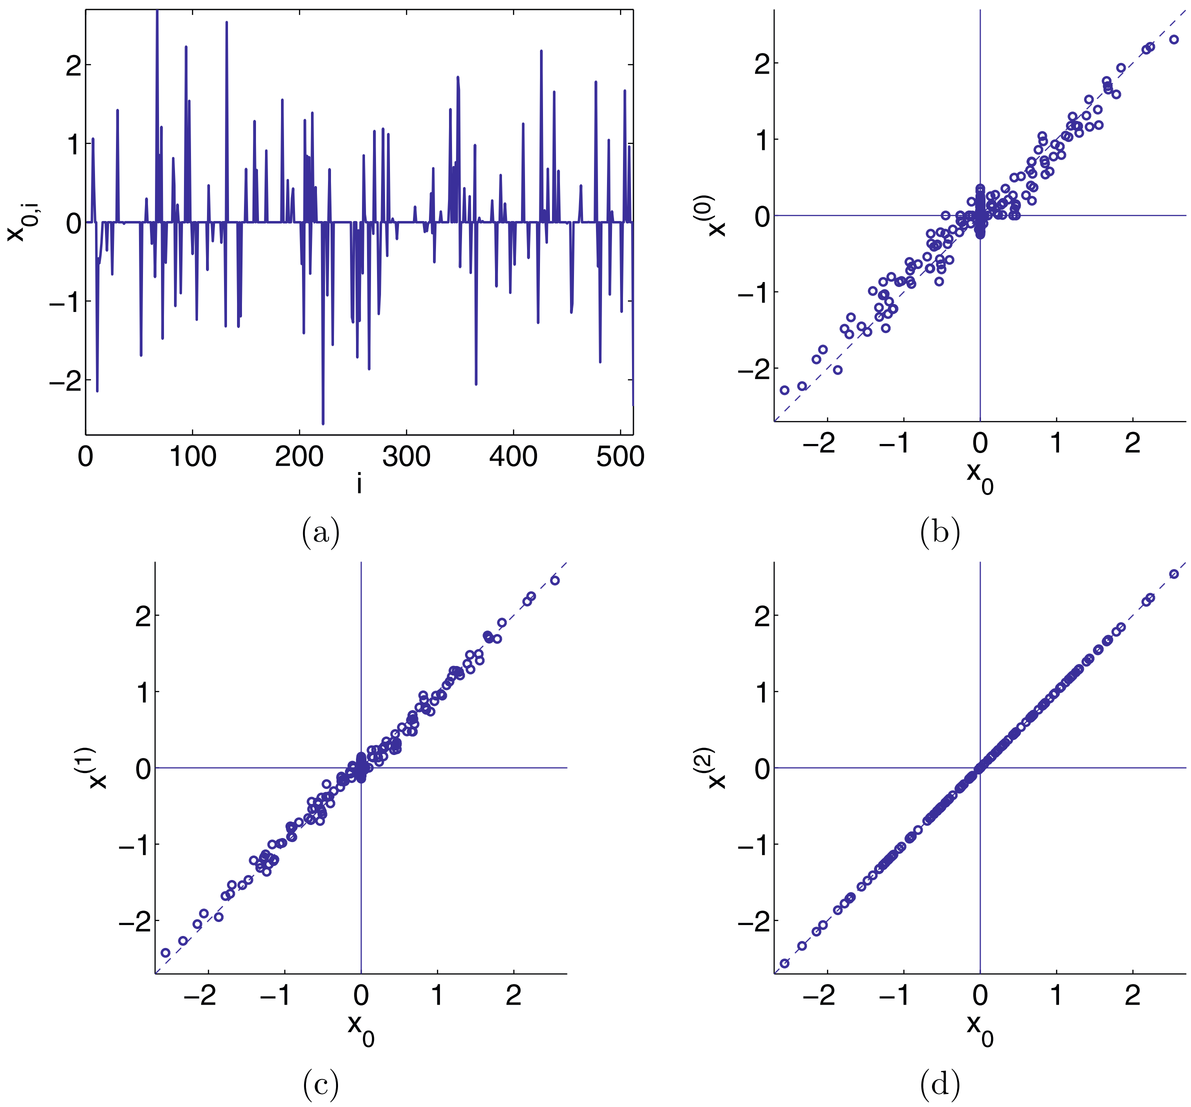

Fig. 2 Sparse signal recovery through reweighted ℓ_1 iterations. (**a**) Original length $n = 512$ signal x_0 with 130 spikes. (**b**) Scatter plot, coefficient-by-coefficient, of x_0 versus its reconstruction $x^{(0)}$ using unweighted ℓ_1 minimization. (**c**) Reconstruction $x^{(1)}$ after the first reweighted iteration. (**d**) Reconstruction $x^{(2)}$ after the second reweighted iteration

Applications — $\ell_2 - \ell_p$ Optimization

- Many problems involve solving the following problem (e.g., basis-pursuit denoising)

$$\min_{\boldsymbol{x}} \; f(\boldsymbol{x}) \triangleq \frac{1}{2}\|\boldsymbol{y} - \mathbf{A}\boldsymbol{x}\|_2^2 + \mu\|\boldsymbol{x}\|_p \tag{9}$$

where $p \geq 1$.

- If $\mathbf{A} = \mathbf{I}$ or $\mathbf{A}$ is unitary, optimal $\boldsymbol{x}^\star$ is computed in closed-form as

$$\boldsymbol{x}^\star = \mathbf{A}^T\boldsymbol{y} - \mathrm{Proj}_C(\mathbf{A}^T\boldsymbol{y})$$

where $C \triangleq \{\boldsymbol{x} : \|\boldsymbol{x}\|_{p*} \leq \mu\}$, $\|\cdot\|_{p*}$ is the dual norm of $\|\cdot\|_p$ and Proj_C denotes the projection operator. In particular, for $p = 1$

$$x_i^\star = \mathrm{soft}(y_i, \mu), \quad i = 1, \ldots, n$$

where $\mathrm{soft}(u, a) \triangleq \mathrm{sign}(u)\max\{|u| - a, 0\}$ denotes a *soft-thresholding* operation.

- For general $\mathbf{A}$, there is no simple closed-form solution for (9).

- **MM for $\ell_2 - \ell_p$ Problem**: Consider a modified $\ell_2 - \ell_p$ problem

$$\min_{\boldsymbol{x}} \; u(\boldsymbol{x}, \boldsymbol{x}^r) \triangleq f(\boldsymbol{x}) + \mathrm{dist}(\boldsymbol{x}, \boldsymbol{x}^r) \tag{10}$$

  where $\mathrm{dist}(\boldsymbol{x}, \boldsymbol{x}^r) \triangleq \frac{c}{2}\|\boldsymbol{x} - \boldsymbol{x}^r\|_2^2 - \frac{1}{2}\|\mathbf{A}\boldsymbol{x} - \mathbf{A}\boldsymbol{x}^r\|_2^2$ and $c > \lambda_{\max}(\mathbf{A}^T\mathbf{A})$.

  - $\mathrm{dist}(\boldsymbol{x}, \boldsymbol{x}^r) \geq 0 \; \forall \boldsymbol{x} \Longrightarrow u(\boldsymbol{x}, \boldsymbol{x}^r)$ majorizes $f(\boldsymbol{x})$.
  - $u(\boldsymbol{x}, \boldsymbol{x}^r)$ can be reexpressed as

$$u(\boldsymbol{x}, \boldsymbol{x}^r) = \frac{c}{2}\|\boldsymbol{x} - \bar{\boldsymbol{x}}^r\|_2^2 + \mu\|\boldsymbol{x}\|_p + \mathrm{const.},$$

    where

$$\bar{\boldsymbol{x}}^r = \frac{1}{c}\mathbf{A}^T(\boldsymbol{y} - \mathbf{A}\boldsymbol{x}^r) + \boldsymbol{x}^r.$$

  - The modified $\ell_2 - \ell_p$ problem (10) has a simple soft-thresholding solution.
  - Repeatedly solving problem (10) leads to an optimal solution of the $\ell_2 - \ell_p$ problem (by the MM convergence in **Theorem 1**)

Applications — Expectation Maximization (EM)

- Consider an ML estimate of θ, given the random observation w

$$\hat{\theta}_{\mathrm{ML}} = \arg\min_{\theta} - \ln p(w|\theta)$$

- Suppose that there are some missing data or hidden variables z in the model. Then, EM algorithm iteratively compute an ML estimate $\hat{\theta}$ as follows:
  - E-step:
  $$g(\theta, \theta^r) \triangleq \mathbb{E}_{z|w,\theta^r}\{\ln p(w, z|\theta)\}$$
  - M-step:
  $$\theta^{r+1} = \arg\max_{\theta} g(\theta, \theta^r)$$
  - repeat the above two steps until convergence.

- EM algorithm generates a nonincreasing sequence of $\{-\ln p(w|\theta^r)\}$.

- EM algorithm can be interpreted by MM.

- MM interpretation of EM algorithm:

$$
\begin{aligned}
& -\ln p(w|\theta) \\
=& -\ln \mathbb{E}_{z|\theta} p(w|z,\theta) \\
=& -\ln \mathbb{E}_{z|\theta}\left[\frac{p(z|w,\theta^r)p(w|z,\theta)}{p(z|w,\theta^r)}\right] \\
=& -\ln \mathbb{E}_{z|w,\theta^r}\left[\frac{p(z|\theta)p(w|z,\theta)}{p(z|w,\theta^r)}\right] \quad \text{(interchange the integrations)} \\
\leq& -\mathbb{E}_{z|w,\theta^r} \ln \left[\frac{p(z|\theta)p(w|z,\theta)}{p(z|w,\theta^r)}\right] \quad \text{(Jensen's inequality)} \\
=& -\mathbb{E}_{z|w,\theta^r} \ln p(w,z|\theta) + \mathbb{E}_{z|w,\theta^r} \ln p(z|w,\theta^r) \qquad \text{(11a)} \\
\triangleq& \, u(\theta,\theta^r)
\end{aligned}
$$

  - $u(\theta,\theta^r)$ majorizes $-\ln p(w|\theta)$, and $-\ln p(w|\theta^r) = u(\theta^r,\theta^r)$;
  - E-step essentially constructs $u(\theta,\theta^r)$;
  - M-step minimizes $u(\theta,\theta^r)$ (note θ appears in the 1st term of (11a) only).

Outline

- Majorization Minimization (MM)
  - Convergence
  - Applications
- Block Coordinate Descent (BCD)
  - Applications
  - Convergence
- Summary

Block Coordinate Descent

- Consider the following problem

$$\min_{\boldsymbol{x}} \; f(\boldsymbol{x}) \quad \text{s.t. } \boldsymbol{x} \in \mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \ldots \times \mathcal{X}_m \subseteq \mathbb{R}^n \tag{12}$$

where each $\mathcal{X}_i \subseteq \mathbb{R}^{n_i}$ is closed, nonempty and convex.

- **BCD Algorithm**:

```
Find a feasible point x^0 ∈ X and set r = 0
repeat
    r = r + 1, i = (r − 1 mod m) + 1
    Let x_i^⋆ ∈ arg min_{x∈X_i} f(x_1^{r−1}, ..., x_{i−1}^{r−1}, x, x_{i+1}^{r−1}, ..., x_m^{r−1})
    Set x_i^r = x_i^⋆ and x_k^r = x_k^{r−1}, ∀k ≠ i
until some convergence criterion is met
```

- Merits of BCD

1. each subproblem is much easier to solve, or even has a closed-form solution;
2. The objective value is nonincreasing along the BCD updates;
3. it allows parallel or distributed implementations.

Applications — $\ell_2 - \ell_1$ Optimization Problem

- Let us revisit the $\ell_2 - \ell_1$ problem

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \; f(\boldsymbol{x}) \triangleq \frac{1}{2}\|\boldsymbol{y} - \mathbf{A}\boldsymbol{x}\|_2^2 + \mu\|\boldsymbol{x}\|_1 \tag{13}$$

- Apart from MM, BCD is another efficient approach to solve (13):
  - Optimize x_k while fixing $x_j = x_j^r, \; \forall j \neq k$:

$$\min_{x_k} \; f_k(x_k) \triangleq \frac{1}{2}\| \underbrace{\boldsymbol{y} - \sum_{j \neq k} \boldsymbol{a}_j x_j^r}_{\triangleq \bar{\boldsymbol{y}}} - \boldsymbol{a}_k x_k\|_2^2 + \mu|x_k|$$

  - The optimal x_k has a closed form:

$$x_k^\star = \operatorname{soft}\left(\boldsymbol{a}_k^T \bar{\boldsymbol{y}}/\|\boldsymbol{a}_k\|^2, \mu/\|\boldsymbol{a}_k\|^2\right)$$

  - Cyclically update $x_k, \; k = 1, \ldots, n$ until convergence.

Applications — Iterative Water-filling for MIMO MAC Sum Capacity Maximization

- **MIMO Channel Capacity Maximization**
  - MIMO received signal model:

$$\boldsymbol{y}(t) = \mathbf{H}\boldsymbol{x}(t) + \boldsymbol{n}(t)$$

where

$\boldsymbol{x}(t) \in \mathbb{C}^N$	Tx signal
$\mathbf{H} \in \mathbb{C}^{N \times N}$	MIMO channel matrix
$\boldsymbol{n}(t) \in \mathbb{C}^N$	standard additive Gaussian noise, i.e., $\boldsymbol{n}(t) \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$.

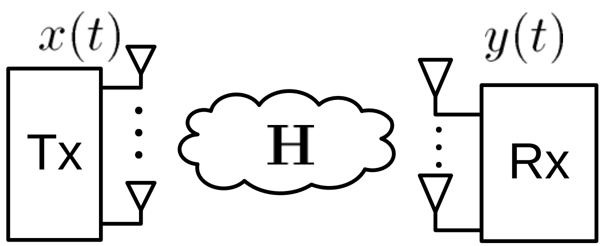

Figure 4: MIMO system model.

– MIMO channel capacity:

$$C(\mathbf{Q}) = \log\det\left(\mathbf{I} + \mathbf{H}\mathbf{Q}\mathbf{H}^H\right)$$

where $\mathbf{Q} = \mathrm{E}\{\boldsymbol{x}(t)\boldsymbol{x}(t)^H\}$ is the covariance of the tx signal.

– MIMO channel capacity maximization:

$$\max_{\mathbf{Q}\succeq\mathbf{0}} \;\log\det\left(\mathbf{I} + \mathbf{H}\mathbf{Q}\mathbf{H}^H\right) \quad \text{s.t. } \mathrm{Tr}(\mathbf{Q}) \leq P$$

where $P > 0$ is the transmit power budget.

– The optimal $\mathbf{Q}^\star$ is given by the well-known *water-filling* solution, i.e.,

$$\mathbf{Q}^\star = \mathbf{V}\mathrm{Diag}(\boldsymbol{p}^\star)\mathbf{V}^H$$

where $\mathbf{H} = \mathbf{U}\mathrm{Diag}(\sigma_1, \ldots, \sigma_N)\mathbf{V}^H$ is the SVD of $\mathbf{H}$, and $\boldsymbol{p}^\star = [p_1^\star, \ldots, p_N^\star]$ is the power allocation with $p_i^\star = \max(0, \mu - 1/\sigma_i^2)$ and $\mu \geq 0$ being the water-level such that $\sum_i p_i^\star = P$.

- **MIMO Multiple-Access Channel (MAC) Sum-Capacity Maximization**
  - Multiple transmitters simultaneously communicate with one receiver:

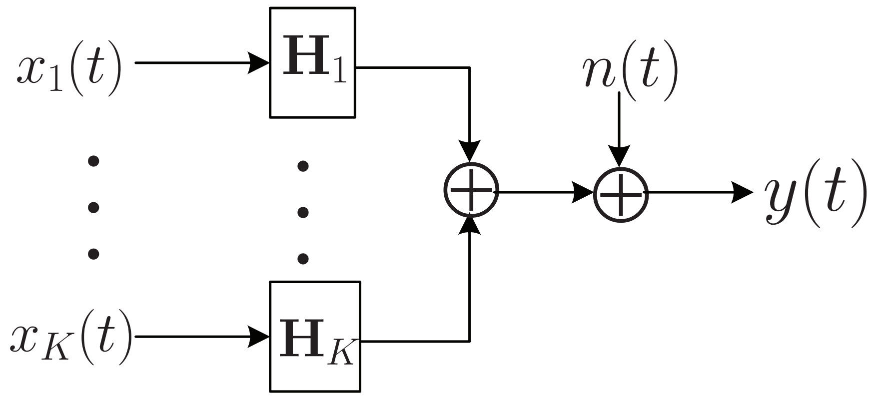

Figure 5: MIMO multiple-access channel (MAC).

  - Received signal model:

$$\boldsymbol{y}(t) = \sum_{k=1}^{K} \mathbf{H}_k \boldsymbol{x}_k(t) + \boldsymbol{n}(t)$$

  - MAC sum capacity:

$$C_{\text{MAC}}(\{\mathbf{Q}_k\}_{k=1}^{K}) = \log\det\left(\sum_{k=1}^{K} \mathbf{H}_k \mathbf{Q}_k \mathbf{H}_k^H + \mathbf{I}\right)$$

– MAC sum capacity maximization:

$$\max_{\{\mathbf{Q}_k\}_{k=1}^K} \quad \log\det\left(\sum_{k=1}^K \mathbf{H}_k\mathbf{Q}_k\mathbf{H}_k^H + \mathbf{I}\right)$$
$$\text{s.t. } \mathrm{Tr}(\mathbf{Q}_k) \le P_k,\ \mathbf{Q}_k \succeq \mathbf{0},\ k = 1,\ldots,K \tag{14}$$

– Problem (14) is convex w.r.t. $\{\mathbf{Q}_k\}$, but it has no simple closed-form solution.
– Alternatively, we can apply BCD to (14) and cyclically update $\mathbf{Q}_k$ while fixing $\mathbf{Q}_j$ for $j \neq k$

$$(\triangle) \quad \max_{\mathbf{Q}_k} \ \log\det\left(\mathbf{H}_k\mathbf{Q}_k\mathbf{H}_k^H + \mathbf{\Phi}\right)$$
$$\text{s.t. } \mathrm{Tr}(\mathbf{Q}_k) \le P_k, \quad \mathbf{Q}_k \succeq \mathbf{0},$$

where $\mathbf{\Phi} = \sum_{j\neq k} \mathbf{H}_j\mathbf{Q}_j\mathbf{H}_j^H + \mathbf{I}$
– $(\triangle)$ has a closed-form water-filling solution, just like the previous single-user MIMO case.

Applications — Low-Rank Matrix Completion

- In a previous lecture, we have introduced the low-rank matrix completion problem, which has huge potential in sales recommendation.

- For example, we would like to predict how much someone is going to like a movie based on its movie preferences:

$$M = \overset{\text{movies}}{\begin{bmatrix} 2 & 3 & 1 & ? & ? & 5 & 5 \\ 1 & ? & 4 & 2 & ? & ? & ? \\ ? & 3 & 1 & ? & 2 & 2 & 2 \\ ? & ? & ? & 3 & ? & 1 & 5 \\ 2 & ? & 4 & ? & ? & 5 & 3 \end{bmatrix}} \text{ users}$$

- M is assumed to be of low rank, as only a few factors affect users' preferences.

$$\min_{\mathbf{W}\in\mathbb{R}^{m\times n}} \ \text{rank}(\mathbf{W}) \to \|\mathbf{W}\|_* \quad \text{s.t. } W_{ij} = M_{ij}, \ \forall (i,j) \in \mathbf{\Omega}$$

- An alternative low-rank matrix completion formulation [Wen-Yin-Zhang]:

$$(\triangle) \quad \min_{\mathbf{X},\mathbf{Y},\mathbf{Z}} \frac{1}{2}\|\mathbf{X}\mathbf{Y}-\mathbf{Z}\|_F^2 \quad \text{s.t. } Z_{ij} = M_{ij},\ \forall (i,j) \in \mathbf{\Omega}$$

where $\mathbf{X} \in \mathbb{R}^{M\times L}$, $\mathbf{Y} \in \mathbb{R}^{L\times N}$, $\mathbf{Z} \in \mathbb{R}^{M\times N}$, and L is an estimate of min. rank.

- Advantage of adopting $(\triangle)$: When BCD is applied, each subproblem of $(\triangle)$ has a closed-form solution:

$$\mathbf{X}^{r+1} = \mathbf{Z}^r \mathbf{Y}^{rT}(\mathbf{Y}^r\mathbf{Y}^{rT})^\dagger,$$

$$\mathbf{Y}^{r+1} = (\mathbf{X}^{r+1^T}\mathbf{X}^{r+1})^\dagger(\mathbf{X}^{r+1^T}\mathbf{Z}^r),$$

$$[\mathbf{Z}^{r+1}]_{i,j} = \begin{cases} [\mathbf{X}^{r+1}\mathbf{Y}^{r+1}]_{i,j}, & \text{for } (i,j) \notin \mathbf{\Omega} \\ M_{i,j}, & \text{for } (i,j) \in \mathbf{\Omega} \end{cases}$$

Applications — Maximizing A Convex Quadratic Function

- Consider maximizing a convex quadratic problem:

$$(\square) \quad \max_{\boldsymbol{x}} \ \frac{1}{2}\boldsymbol{x}^T \mathbf{Q} \boldsymbol{x} + \boldsymbol{c}^T \boldsymbol{x} \quad \text{s.t. } \boldsymbol{x} \in \mathcal{X}$$

  where $\mathcal{X}$ is a polyhedral set, and $\mathbf{Q} \succeq \mathbf{0}$.

- $(\square)$ is equivalent to the following problem[1]

$$(\triangle) \quad \max_{\boldsymbol{x}_1, \boldsymbol{x}_2} \ \frac{1}{2}\boldsymbol{x}_1^T \mathbf{Q} \boldsymbol{x}_2 + \frac{1}{2}\boldsymbol{c}^T \boldsymbol{x}_1 + \frac{1}{2}\boldsymbol{c}^T \boldsymbol{x}_2 \quad \text{s.t. } (\boldsymbol{x}_1, \boldsymbol{x}_2) \in \mathcal{X} \times \mathcal{X}$$

- When fixing either $\boldsymbol{x}_1$ or $\boldsymbol{x}_2$, problem $(\triangle)$ is an LP, thereby efficiently solvable.

[1] The equivalence is in the following sense: If $\boldsymbol{x}^\star$ is an optimal solution of $(\square)$, then $(\boldsymbol{x}^\star, \boldsymbol{x}^\star)$ is optimal for $(\triangle)$; Conversely, if $(\boldsymbol{x}_1^\star, \boldsymbol{x}_2^\star)$ is an optimal solution of $(\triangle)$, then both $\boldsymbol{x}_1^\star, \boldsymbol{x}_2^\star$ are optimal for $(\square)$.

Applications — Nonnegative Matrix Factorization (NMF)

- NMF is concerned with the following problem [Lee-Seung]:

$$\min_{\mathbf{U}\in\mathbb{R}^{m\times k},\mathbf{V}\in\mathbb{R}^{k\times n}} \|\mathbf{M}-\mathbf{U}\mathbf{V}\|_F^2 \qquad \text{s.t. } \mathbf{U}\geq\mathbf{0},\ \mathbf{V}\geq\mathbf{0} \tag{15}$$

where $\mathbf{M}\geq\mathbf{0}$.

- Usually $k \ll \min(m,n)$ or $mk+nk \ll mn$, so NMF can be seen as a linear dimensionality reduction technique for nonnegative data.

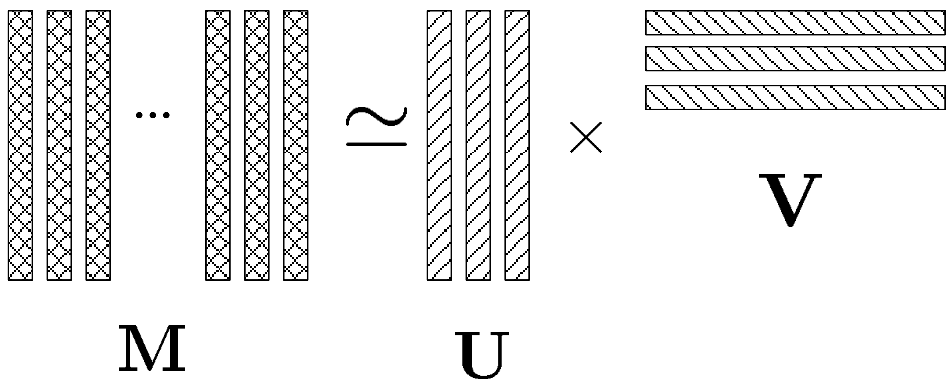

NMF Examples

- **Image Processing:**
  - $\mathbf{U} \geq \mathbf{0}$ constraints the basis elements to be nonnegative.
  - $\mathbf{V} \geq \mathbf{0}$ imposes an additive reconstruction.

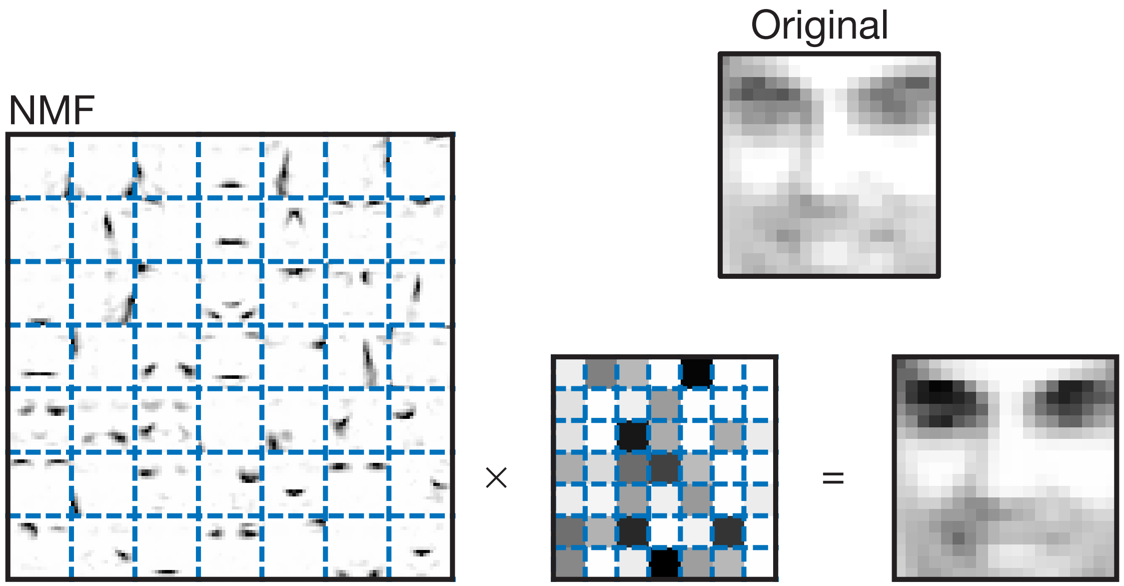

The basis elements extract facial features such as eyes, nose and lips.

- **Text Mining**

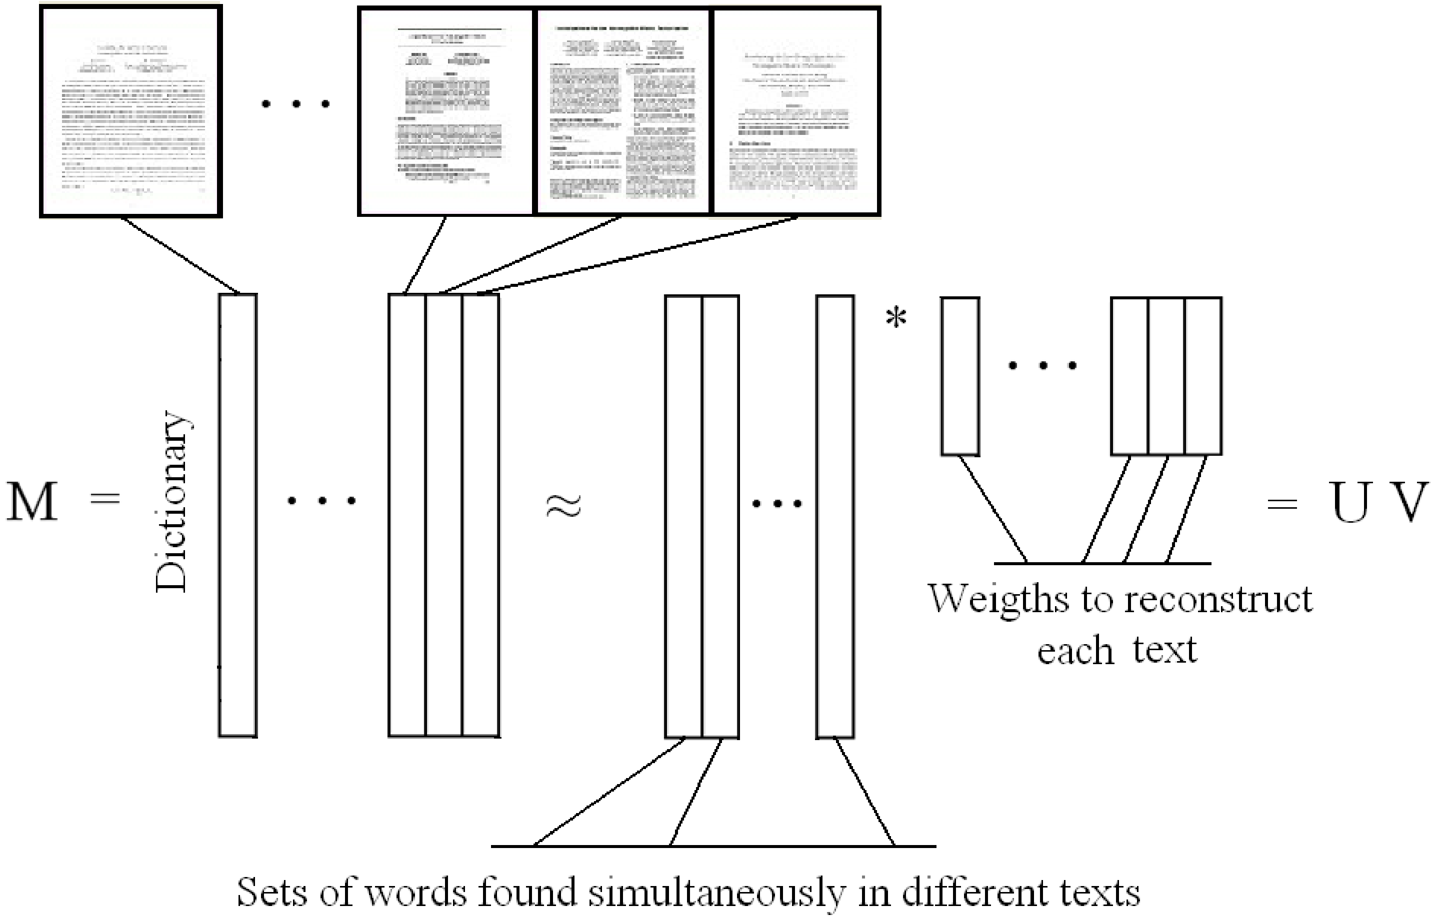

  - Basis elements allow to recover different topics;
  - Weights allow to assign each text to its corresponding topics.

- **Hyperspectral Unmixing**

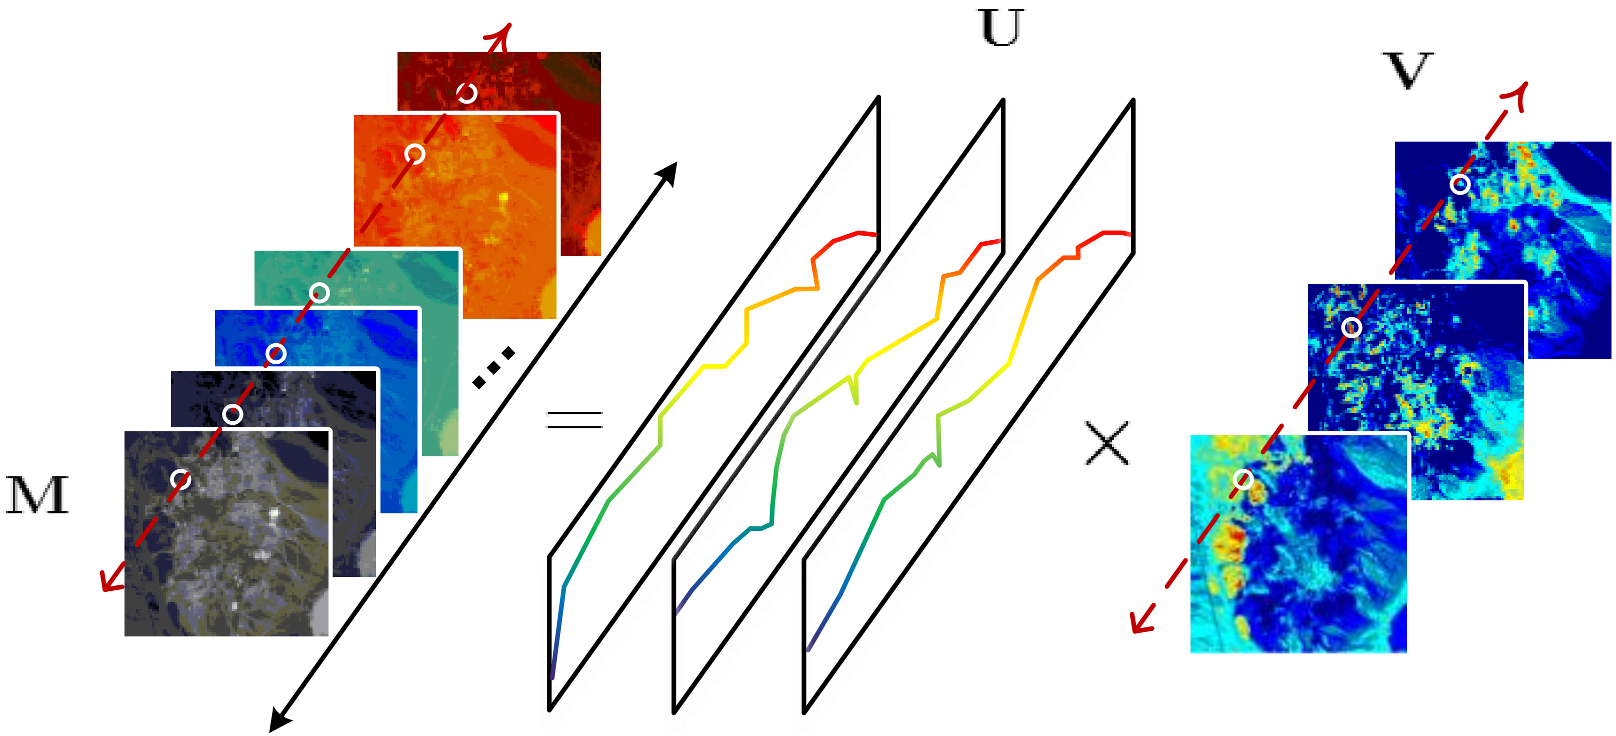

  - Basis elements $\mathbf{U}$ represent different materials;
  - Weights $\mathbf{V}$ allow to know which pixel contains which material.

- Let's turn back to the NMF problem:

$$\min_{\mathbf{U}\in\mathbb{R}^{m\times k},\mathbf{V}\in\mathbb{R}^{k\times n}} \|\mathbf{M}-\mathbf{U}\mathbf{V}\|_F^2 \qquad \text{s.t. } \mathbf{U}\geq\mathbf{0},\ \mathbf{V}\geq\mathbf{0} \tag{16}$$

- Without "$\geq \mathbf{0}$" constraints, the optimal $\mathbf{U}^\star$ and $\mathbf{V}^\star$ can be obtained by SVD.

- With "$\geq \mathbf{0}$" constraints, problem (16) is generally NP-hard.

- When fixing $\mathbf{U}$ (resp. $\mathbf{V}$), problem (16) is convex w.r.t. $\mathbf{V}$ (resp. $\mathbf{U}$).

- For example, for a given $\mathbf{U}$, the ith column of $\mathbf{V}$ is updated by solving the following NLS problem:

$$\min_{\mathbf{V}(:,i)\in\mathbb{R}^k} \|\mathbf{M}(:,i)-\mathbf{U}\mathbf{V}(:,i)\|_2^2, \quad \text{s.t. } \mathbf{V}(:,i)\geq\mathbf{0}, \tag{17}$$

BCD Algorithm for NMF:

Initialize $\mathbf{U} = \mathbf{U}^0$, $\mathbf{V} = \mathbf{V}^0$ and $r = 0$;

repeat

solve the NLS problem

$$\mathbf{V}^\star \in \arg \min_{\mathbf{V} \in \mathbb{R}^{k \times n}} \|\mathbf{M} - \mathbf{U}^r \mathbf{V}\|_F^2, \quad \text{s.t. } \mathbf{V} \geq \mathbf{0}$$

$\mathbf{V}^{r+1} = \mathbf{V}^\star$;

solve the NLS problem

$$\mathbf{U}^\star \in \arg \min_{\mathbf{U} \in \mathbb{R}^{m \times k}} \|\mathbf{M} - \mathbf{U} \mathbf{V}^{r+1}\|_F^2, \quad \text{s.t. } \mathbf{U} \geq \mathbf{0}$$

$\mathbf{U}^{r+1} = \mathbf{U}^\star$;

$r = r + 1$;

until some convergence criterion is met

Outline

- Majorization Minimization (MM)
  - Convergence
  - Applications
- Block Coordinate Descent (BCD)
  - Applications
  - Convergence
- Summary

BCD Convergence

- The idea of BCD is to divide and conquer. However, there is no free lunch; BCD may get stuck or converge to some point of no interest.

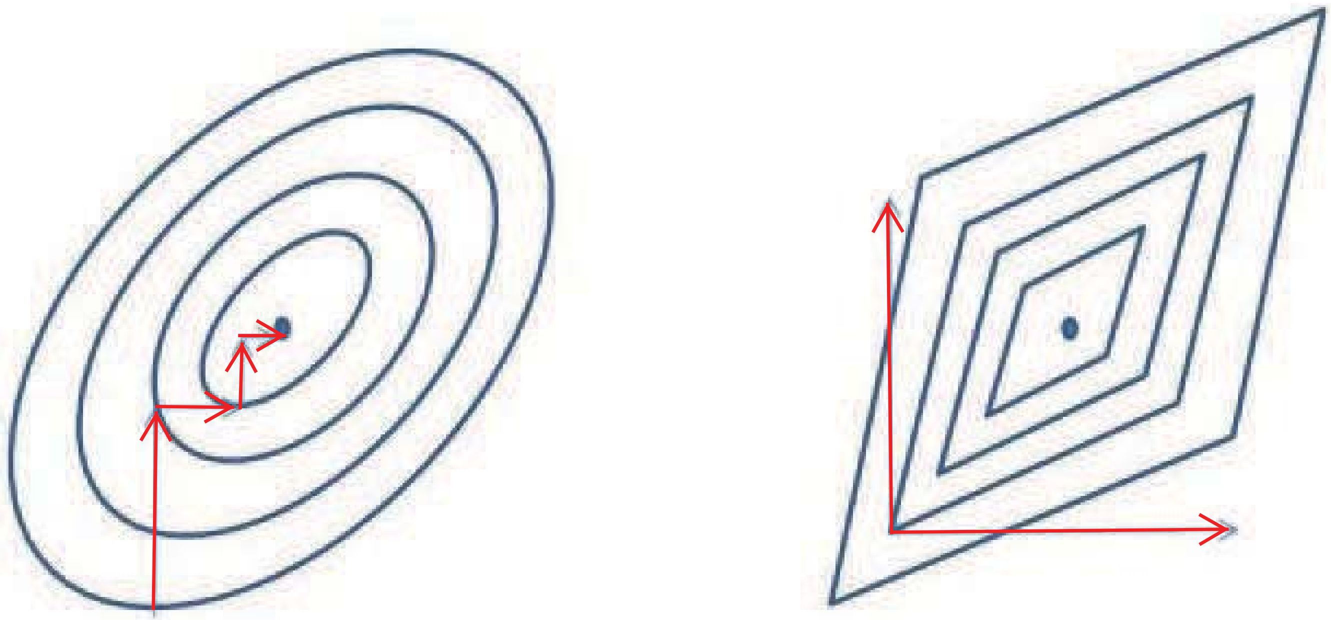

Figure 6: BCD for smooth/non-smooth minimization.

BCD Convergence

$$\min_{\boldsymbol{x}} \ f(\boldsymbol{x}) \quad \text{s.t. } \boldsymbol{x} \in \mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \ldots \times \mathcal{X}_m \subseteq \mathbb{R}^n \tag{18}$$

- A well-known BCD convergence result due to Bertsekas:

Theorem 2. *([Bertsekas]) Suppose that f is continuously differentiable over the convex closed set $\mathcal{X}$. Furthermore, suppose that for each i*

$$g_i(\boldsymbol{\xi}) \triangleq f(\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_{i-1}, \boldsymbol{\xi}, \boldsymbol{x}_{i+1}, \ldots, \boldsymbol{x}_m)$$

is strictly convex. Let $\{\boldsymbol{x}^r\}$ be the sequence generated by BCD method. Then every limit point of $\{\boldsymbol{x}^r\}$ is a stationary point of problem (18).

- If $\mathcal{X}$ is (convex) compact, i.e., closed and bounded, then strict convexity of $g_i(\boldsymbol{\xi})$ can be relaxed to having a unique optimal solution.

- Application: Iterative water-filling for MIMO MAC sum capacity max.:

$$(\triangle) \quad \max_{\{\mathbf{Q}_k\}_{k=1}^K} \quad \log\det\big(\textstyle\sum_{k=1}^K \mathbf{H}_k\mathbf{Q}_k\mathbf{H}_k^H + \mathbf{I}\big), \quad \text{s.t. } \mathrm{Tr}(\mathbf{Q}_k) \le P_k,\ \mathbf{Q}_k \succeq \mathbf{0},\ \forall k$$

- Iterative water-filling converges to a global optimal solution of $(\triangle)$, because
  - BCD subproblem is strictly convex (assuming full column rankness of $\mathbf{H}_k$);
  - $\mathcal{X}_k$ is a convex closed subset;
  - $(\triangle)$ is a convex problem, so stationary point $\Longrightarrow$ global optimal solution

Generalization of Bertsekas' Convergence Result

- Generalization 1: Relax Strict Convexity to Strict Quasiconvexity[2] [Grippo-Sciandrone]

  Theorem 3. *Suppose that the function f is continuously differentiable and strictly quasiconvex with respect to $\boldsymbol{x}_i$ on $\mathcal{X}$, for each $i = 1, \ldots, m-2$ and that the sequence $\{\boldsymbol{x}^r\}$ generated by the BCD method has limit points. Then, every limit point is a stationary point of problem* (18).

- Application: Low-Rank Matrix Completion

$$(\triangle) \quad \min_{\mathbf{X},\mathbf{Y},\mathbf{Z}} \frac{1}{2}\|\mathbf{X}\mathbf{Y} - \mathbf{Z}\|_F^2 \quad \text{s.t. } Z_{ij} = M_{ij}, \ \forall (i,j) \in \mathbf{\Omega}$$

  – $m = 3$ and $(\triangle)$ is strictly convex w.r.t. $\mathbf{Z} \Longrightarrow$ BCD converges to a stationary point.

[2] f is strictly quasiconvex w.r.t. $\boldsymbol{x}_i \in \mathcal{X}_i$ on $\mathcal{X}$ if for every $\boldsymbol{x} \in \mathcal{X}$ and $\boldsymbol{y}_i \in \mathcal{X}_i$ with $\boldsymbol{y}_i \neq \boldsymbol{x}_i$ we have

$$f(\boldsymbol{x}_1, \ldots, t\boldsymbol{x}_i + (1-t)\boldsymbol{y}_i, \ldots, \boldsymbol{x}_m) < \max\{f(\boldsymbol{x}), f(\boldsymbol{x}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{x}_m)\}, \ \forall t \in (0,\ 1).$$

- Generalization 2: Without Solution Uniqueness

Theorem 4. *Suppose that f is pseudoconvex[3] on $\mathcal{X}$ and that $\mathcal{L}_{\mathcal{X}}^0 := \{\boldsymbol{x} \in \mathcal{X} : f(\boldsymbol{x}) \leq f(\boldsymbol{x}^0)\}$ is compact. Then, the sequence generated by BCD method has limit points and every limit point is a global minimizer of f.*

- Application: Iterative water-filling for MIMO-MAC sum capacity max.

$$\max_{\{\mathbf{Q}_k\}_{k=1}^K} \quad \log\det\left(\sum_{k=1}^K \mathbf{H}_k\mathbf{Q}_k\mathbf{H}_k^H + \mathbf{I}\right)$$
$$\text{s.t. } \mathrm{Tr}(\mathbf{Q}_k) \leq P_k,\ \mathbf{Q}_k \succeq \mathbf{0}, \quad k = 1, \ldots, K$$

- f is convex, thus pseudoconvex;
- $\{\mathbf{Q}_k \mid \mathrm{Tr}(\mathbf{Q}_k) \leq P_k,\ \mathbf{Q}_k \succeq \mathbf{0}\}$ is compact;
- iterative water-filling converges to a globally optimal solution.

[3] f is pseudoconvex if for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$ such that $\nabla f(\boldsymbol{x})^T(\boldsymbol{y} - \boldsymbol{x}) \geq 0$, we have $f(\boldsymbol{y}) \geq f(\boldsymbol{x})$. Notice that "convex $\subset$ pseudoconvex $\subset$ quasiconvex".

- Generalization 3: Without Solution Uniqueness, Pseudoconvexity and Compactness

Theorem 5. *Suppose that f is continuously differentiable, and that $\mathcal{X}$ is convex and closed. Moreover, if there are only two blocks, i.e., $m = 2$, then every limit point generated by BCD is a stationary point of f.*

- Application: NMF

$$\min_{\mathbf{U}\in\mathbb{R}^{m\times k},\mathbf{V}\in\mathbb{R}^{k\times n}} \|\mathbf{M}-\mathbf{UV}\|_F^2 \qquad \text{s.t. } \mathbf{U}\geq\mathbf{0},\ \mathbf{V}\geq\mathbf{0}$$

- Alternating NLS converges to a stationary point of the NMF problem, since
  - the objective is continuously differentiable;
  - the feasible set is convex and closed;
  - $m = 2$.

Summary

- MM and BCD have great potential in handling nonconvex problems and realizing fast/distributed implementations for large-scale convex problems;
- Many well-known algorithms can be interpreted as special cases of MM and BCD;
- Under some conditions, convergence to stationary point can be guaranteed by MM and BCD.

References

M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," submitted to *SIAM Journal on Optimization*, available online at http://arxiv.org/abs/ 1209.2385.

L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear Gauss-Seidel method under convex constraints," *Operation research letter* vol. 26, pp. 127-136, 2000

E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, 14 (2008), pp. 877-905.

M. Zibulevsky and M. Elad, "$\ell_1 - \ell_2$ optimization in signal and image processing," *IEEE Signal Process. Magazine*, May 2010, pp.76-88.

D. P. Bertsekas, "Nonlinear Programming," Athena Scientific, 1st Ed., 1995

W. Yu and J. M. Cioffi, "Sum capacity of a Gaussian vector broadcast channel", *IEEE Trans. Inf. Theory*, vol. 50, no. 1, pp. 145-152, Jan. 2004

Z. Wen, W. Yin, and Y. Zhang, "Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm," Rice CAAM Tech Report 10-07.

Daniel D. Lee and H. Sebastian Seung, "Algorithms for Non-negative Matrix Factorization". *Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference*. MIT Press. pp. 556-562, 2001.