## Big Data Analytics: Optimization and Randomization

Tianbao Yang $^{\dagger}$, Qihang Lin ${ }^{\natural}$, Rong Jin* $\ddagger$

Tutorial@SIGKDD 2015
Sydney, Australia
${ }^{\dagger}$ Department of Computer Science, The University of Iowa, IA, USA
${ }^{4}$ Department of Management Sciences, The University of lowa, IA, USA
*Department of Computer Science and Engineering, Michigan State University, MI, USA
$\ddagger$ Institute of Data Science and Technologies at Alibaba Group, Seattle, USA

August 10, 2015

## URL

http://www.cs.uiowa.edu/~tyng/kdd15-tutorial.pdf

## Some Claims

No

- This tutorial is not an exhaustive literature survey
- It is not a survey on different machine learning/data mining algorithms
Yes
- It is about how to efficiently solve machine learning/data mining (formulated as optimization) problems for big data


## Outline

- Part I: Basics
- Part II: Optimization
- Part III: Randomization


# Big Data Analytics: Optimization and Randomization Part I: Basics 

## Outline

## (1) Basics

- Introduction
- Notations and Definitions


## Three Steps for Machine Learning



Data


Model


## Optimization

## Big Data Challenge



## Big Data Challenge

## Big Model



60 million parameters

## Learning as Optimization

Ridge Regression Problem:

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-\mathbf{w}^{\top} \mathbf{x}_{i}\right)^{2}+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$



- $\mathbf{x}_{i} \in \mathbb{R}^{d}: d$-dimensional feature vector
- $y_{i} \in \mathbb{R}$ : target variable
- $\mathbf{w} \in \mathbb{R}^{d}$ : model parameters
- $n$ : number of data points


## Learning as Optimization

Ridge Regression Problem:

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \underbrace{\frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-\mathbf{w}^{\top} \mathbf{x}_{i}\right)^{2}}_{\text {Empirical Loss }}+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$



- $\mathbf{x}_{i} \in \mathbb{R}^{d}: d$-dimensional feature vector
- $y_{i} \in \mathbb{R}$ : target variable
- $\mathbf{w} \in \mathbb{R}^{d}$ : model parameters
- $n$ : number of data points


## Learning as Optimization

Ridge Regression Problem:

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-\mathbf{w}^{\top} \mathbf{x}_{i}\right)^{2}+\underbrace{\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}}_{\text {Regularization }}
$$



- $\mathbf{x}_{i} \in \mathbb{R}^{d}: d$-dimensional feature vector
- $y_{i} \in \mathbb{R}$ : target variable
- $\mathbf{w} \in \mathbb{R}^{d}$ : model parameters
- $n$ : number of data points


## Learning as Optimization

Classification Problems:

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i} \mathbf{w}^{\top} \mathbf{x}_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$



- $y_{i} \in\{+1,-1\}$ : label
- Loss function $\ell(z): z=y \mathbf{w}^{\top} \mathbf{x}$

1. SVMs: (squared) hinge loss $\ell(z)=\max (0,1-z)^{p}$, where $p=1,2$
2. Logistic Regression: $\ell(z)=\log (1+\exp (-z))$

## Learning as Optimization

Feature Selection:

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda\|\mathbf{w}\|_{1}
$$


(a)

(b)

- $\ell_{1}$ regularization $\|\mathbf{w}\|_{1}=\sum_{i=1}^{d}\left|w_{i}\right|$
- $\lambda$ controls sparsity level


## Learning as Optimization

Feature Selection using Elastic Net:
(

- Elastic net regularizer, more robust than $\ell_{1}$ regularizer


## Learning as Optimization

Multi-class/Multi-task Learning:

$$
\min _{\mathbf{W}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{W} \mathbf{x}_{i}, y_{i}\right)+\lambda r(\mathbf{W})
$$

- $\mathbf{W} \in \mathbb{R}^{K \times d}$
- $r(\mathbf{W})=\|\mathbf{W}\|_{F}^{2}=\sum_{k=1}^{K} \sum_{j=1}^{d} W_{k j}^{2}$ : Frobenius Norm
- $r(\mathbf{W})=\|\mathbf{W}\|_{*}=\sum_{i} \sigma_{i}$ : Nuclear Norm (sum of singular values)
- $r(\mathbf{W})=\|\mathbf{W}\|_{1, \infty}=\sum_{j=1}^{d}\left\|W_{: j}\right\|_{\infty}: \ell_{1, \infty}$ mixed norm


## Learning as Optimization

Regularized Empirical Loss Minimization

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})
$$

- Both $\ell$ and $R$ are convex functions
- Extensions to Matrix Cases are possible (sometimes straightforward)
- Extensions to Kernel methods can be combined with randomized approaches
- Extensions to Non-convex (e.g., deep learning) are in progress


## Data Matrices and Machine Learning

The Instance-feature Matrix: $X \in \mathbb{R}^{n \times d}$



## Data Matrices and Machine Learning

The output vector: $\mathbf{y}=\left(\begin{array}{c}y_{1} \\ y_{2} \\ \cdot \\ \cdot \\ \cdot \\ y_{n}\end{array}\right) \in \mathbb{R}^{n \times 1}$

- continuous $y_{i} \in \mathbb{R}$ : regression (e.g., house price)
- discrete, e.g., $y_{i} \in\{1,2,3\}$ : classification (e.g., species of iris)


Setosa


Versicolour


Virginica

## Data Matrices and Machine Learning

The Instance-Instance Matrix: $K \in \mathbb{R}^{n \times n}$

- Similarity Matrix
- Kernel Matrix



## Data Matrices and Machine Learning

Some machine learning tasks are formulated on the kernel matrix

- Clustering
- Kernel Methods



## Data Matrices and Machine Learning

The Feature-Feature Matrix: $C \in \mathbb{R}^{d \times d}$

- Covariance Matrix
- Distance Metric Matrix



## Data Matrices and Machine Learning

Some machine learning tasks requires the covariance matrix

- Principal Component Analysis
- Top-k Singular Value (Eigen-Value) Decomposition of the Covariance Matrix



## Why Learning from Big Data is Challenging?

- High per-iteration cost
- High memory cost
- High communication cost
- Large iteration complexity


## Outline

## (1) Basics

- Introduction
- Notations and Definitions


## Norms

## Vector $\mathbf{x} \in \mathbb{R}^{d}$

- Euclidean vector norm: $\|\mathbf{x}\|_{2}=\sqrt{\mathbf{x}^{\top} \mathbf{x}}=\sqrt{\sum_{i=1}^{d} x_{i}^{2}}$
- $\ell_{p}$-norm of a vector: $\|\mathbf{x}\|_{p}=\left(\sum_{i=1}^{d}\left|x_{i}\right|^{p}\right)^{1 / p}$ where $p \geq 1$
(a) $\mathrm{O}_{2}$ norm $\| \boldsymbol{x}_{2}=\sqrt{\sum_{i=1}^{d} x_{i}^{2}}$
(2) $\ell_{1}$ norm $\|\mathbf{x}\|_{1}=\sum_{i=1}^{d}\left|x_{i}\right|$
(3) $\ell \infty \operatorname{norm}\|x\|_{\infty}=\max \left|x_{i}\right|$


## Norms

## Vector $\mathbf{x} \in \mathbb{R}^{d}$

- Euclidean vector norm: $\|\mathbf{x}\|_{2}=\sqrt{\mathbf{x}^{\top} \mathbf{x}}=\sqrt{\sum_{i=1}^{d} x_{i}^{2}}$
- $\ell_{p}$-norm of a vector: $\|\mathbf{x}\|_{p}=\left(\sum_{i=1}^{d}\left|x_{i}\right|^{p}\right)^{1 / p}$ where $p \geq 1$
(1) $\ell_{2}$ norm $\|\mathbf{x}\|_{2}=\sqrt{\sum_{i=1}^{d} x_{i}^{2}}$
(2) $\ell_{1}$ norm $\|\mathbf{x}\|_{1}=\sum_{i=1}^{d}\left|x_{i}\right|$
(3) $\ell_{\infty}$ norm $\|\mathbf{x}\|_{\infty}=\max _{i}\left|x_{i}\right|$


## Norms

## Vector $\mathbf{x} \in \mathbb{R}^{d}$

- Euclidean vector norm: $\|\mathbf{x}\|_{2}=\sqrt{\mathbf{x}^{\top} \mathbf{x}}=\sqrt{\sum_{i=1}^{d} x_{i}^{2}}$
- $\ell_{p}$-norm of a vector: $\|\mathbf{x}\|_{p}=\left(\sum_{i=1}^{d}\left|x_{i}\right|^{p}\right)^{1 / p}$ where $p \geq 1$
(1) $\ell_{2}$ norm $\|\mathbf{x}\|_{2}=\sqrt{\sum_{i=1}^{d} x_{i}^{2}}$
(2) $\ell_{1}$ norm $\|\mathbf{x}\|_{1}=\sum_{i=1}^{d}\left|x_{i}\right|$
(3) $\ell_{\infty}$ norm $\|\mathbf{x}\|_{\infty}=\max _{i}\left|x_{i}\right|$


## Matrix Factorization

Matrix $X \in \mathbb{R}^{n \times d}$

- Singular Value Decomposition $X=U \Sigma V^{\top}$
(1) $U \in \mathbb{R}^{n \times r}$ : orthonormal columns $\left(U^{\top} U=I\right)$ : span column space
(2) $\Sigma \in \mathbb{R}^{r \times r}$ : diagonal matrix $\Sigma_{i i}=\sigma_{i}>0, \sigma_{1} \geq \sigma_{2} \ldots \geq \sigma_{r}$
(3) $V \in \mathbb{R}^{d \times r}$ : orthonormal columns $\left(V^{\top} V=I\right)$ : span row space
(9) $r \leq \min (n, d)$ : max value such that $\sigma_{r}>0$ : rank of $X$
(3) $U_{k} \Sigma_{k} V_{k}^{\top}$ : top- $k$ approximation
- Pseudo inverse:

- QR factorization: $X=Q R(n \geq d)$
- $Q \in \mathbb{R}^{n \times d}$ : orthonormal columns
- $R \in \mathbb{R}^{d \times d}$ : upper triangular matrix


## Matrix Factorization

Matrix $X \in \mathbb{R}^{n \times d}$

- Singular Value Decomposition $X=U \Sigma V^{\top}$
(1) $U \in \mathbb{R}^{n \times r}$ : orthonormal columns $\left(U^{\top} U=I\right)$ : span column space
(2) $\Sigma \in \mathbb{R}^{r \times r}$ : diagonal matrix $\Sigma_{i i}=\sigma_{i}>0, \sigma_{1} \geq \sigma_{2} \ldots \geq \sigma_{r}$
(3) $V \in \mathbb{R}^{d \times r}$ : orthonormal columns $\left(V^{\top} V=I\right)$ : span row space
(9) $r \leq \min (n, d)$ : max value such that $\sigma_{r}>0$ : rank of $X$
(5) $U_{k} \Sigma_{k} V_{k}^{\top}$ : top- $k$ approximation
- Pseudo inverse: $X^{\dagger}=V \Sigma^{-1} U^{\top}$
- QR factorization: $X=Q R(n \geq d)$ - $Q \in \mathbb{R}^{n \times d}$ : orthonormal columns


## Matrix Factorization

Matrix $X \in \mathbb{R}^{n \times d}$

- Singular Value Decomposition $X=U \Sigma V^{\top}$
(1) $U \in \mathbb{R}^{n \times r}$ : orthonormal columns $\left(U^{\top} U=I\right)$ : span column space
(2) $\Sigma \in \mathbb{R}^{r \times r}$ : diagonal matrix $\Sigma_{i i}=\sigma_{i}>0, \sigma_{1} \geq \sigma_{2} \ldots \geq \sigma_{r}$
(3) $V \in \mathbb{R}^{d \times r}$ : orthonormal columns $\left(V^{\top} V=I\right)$ : span row space
(9) $r \leq \min (n, d)$ : $\max$ value such that $\sigma_{r}>0$ : rank of $X$
(5) $U_{k} \Sigma_{k} V_{k}^{\top}$ : top- $k$ approximation
- Pseudo inverse: $X^{\dagger}=V \Sigma^{-1} U^{\top}$
- QR factorization: $X=Q R(n \geq d)$
- $Q \in \mathbb{R}^{n \times d}$ : orthonormal columns
- $R \in \mathbb{R}^{d \times d}$ : upper triangular matrix


## Norms

Matrix $X \in \mathbb{R}^{n \times d}$

- Frobenius norm: $\|X\|_{F}=\sqrt{\operatorname{tr}\left(X^{\top} X\right)}=\sqrt{\sum_{i=1}^{n} \sum_{j=1}^{d} X_{i j}^{2}}$
- Spectral (induced norm) of a matrix: $\|X\|_{2}=\max _{\|u\|_{2}=1}\|X \mathbf{u}\|_{2}$ - $\|A\|_{2}=\sigma_{1}$ (maximum singular value)


## Norms

Matrix $X \in \mathbb{R}^{n \times d}$

- Frobenius norm: $\|X\|_{F}=\sqrt{\operatorname{tr}\left(X^{\top} X\right)}=\sqrt{\sum_{i=1}^{n} \sum_{j=1}^{d} X_{i j}^{2}}$
- Spectral (induced norm) of a matrix: $\|X\|_{2}=\max _{\|\mathbf{u}\|_{2}=1}\|X \mathbf{u}\|_{2}$
- $\|A\|_{2}=\sigma_{1}$ (maximum singular value)


## Convex Optimization

$$
\min _{x \in \mathcal{X}} f(x)
$$



- $\mathcal{X}$ is a convex domain
- for any $x, y \in \mathcal{X}$, their convex combination $\alpha x+(1-\alpha) y \in \mathcal{X}$
- $f(x)$ is a convex function


## Convex Function

## Characterization of Convex Function



$$
\begin{gathered}
f(\alpha x+(1-\alpha) y) \leq \alpha f(x)+(1-\alpha) f(y), \\
\forall x, y \in \mathcal{X}, \alpha \in[0,1]
\end{gathered}
$$



$$
f(x) \geq f(y)+\nabla f(y)^{\top}(x-y) \forall x, y \in \mathcal{X}
$$

## Convex Function

## Characterization of Convex Function



$$
\begin{gathered}
f(\alpha x+(1-\alpha) y) \leq \alpha f(x)+(1-\alpha) f(y), \\
\forall x, y \in \mathcal{X}, \alpha \in[0,1]
\end{gathered}
$$



$$
f(x) \geq f(y)+\nabla f(y)^{\top}(x-y) \forall x, y \in \mathcal{X}
$$

## Convex vs Strongly Convex

Convex function:

$$
f(x) \geq f(y)+\nabla f(y)^{\top}(x-y) \forall x, y \in \mathcal{X}
$$

Strongly Convex function:

$$
f(x) \geq f(y)+\nabla f(y)^{\top}(x-y)+\frac{\lambda}{2}\|x-y\|_{2}^{2} \forall x, y \in \mathcal{X}
$$

Global optimum is unique

## Convex vs Strongly Convex

Convex function:

$$
\begin{aligned}
& \qquad f(x) \geq f(y)+\nabla f(y)^{\top}\left(\begin{array}{c}
\text { strong convexity } \\
\text { constant }
\end{array}\right. \\
& \text { Strongly Convex function: } \\
& \qquad f(x) \geq f(y)+\nabla f(y)^{\top}(x-y)+\frac{\lambda}{2}\|x-y\|_{2}^{2} \forall x, y \in \mathcal{X}
\end{aligned}
$$

Global optimum is unique

## Non-smooth function vs Smooth function

Non-smooth function

- Lipschitz continuous: e.g. absolute loss $f(x)=|x|$
- $|f(x)-f(y)| \leq G\|x-y\|_{2}$
- Subgradient: $f(x) \geq f(y)+\partial f(y)^{\top}(x-y)$



## Smooth function

## Non-smooth function vs Smooth function

Non-smooth function

- Lipschitz continuous: e.g. absolute loss $f(x)=|x|$
- $|f(x)-f(y)| \leq G\|x-y\|_{2}$
- Subgradient: $f(x) \geq f(y)+\partial f(y)^{\top}(x-y)$


Smooth function

- e.g. logistic loss $f(x)=\log (1+\exp (-x))$
- $\|\nabla f(x)-\nabla f(y)\|_{2} \leq L\|x-y\|_{2}$



## Non-smooth function vs Smooth function

Non-smooth function

- Lipschitz continuouconstant bsolyte loss
$f(x)=|x|$
- $|f(x)-f(y)| \leq G\|x-y\|_{2}$
- Subgradient: $f(x) \geq f(y)+\partial f(y)^{\top}(x-y)$


Smooth function

## Lipschitz

Smooth function | smoothness |
| :---: |
| constant |

- e.g. logistic loss $f(x)=1 / g(1+\exp (-x))$
- $\|\nabla f(x)-\nabla f(y)\|_{2} \leq L\|x-y\|_{2}$



## Next

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})
$$

Part II: Optimization

- stochastic optimization
- distributed optimization

Reduce Iteration Complexity: utilizing properties of functions

## Next ...



Part III: Randomization

- Classification, Regression
- SVD, K-means, Kernel methods

Reduce Data Size: utilizing properties of data

Please stay tuned!

## Big Data Analytics: Optimization and Randomization Part II: Optimization

## Outline

(2) Optimization

- (Sub)Gradient Methods
- Stochastic Optimization Algorithms for Big Data
- Stochastic Optimization
- Distributed Optimization


## Learning as Optimization

Regularized Empirical Loss Minimization

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \underbrace{\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})}_{F(\mathbf{w})}
$$

## Convergence Measure

- Most optimization algorithms are iterative

$$
\mathbf{w}_{t+1}=\mathbf{w}_{t}+\Delta \mathbf{w}_{t}
$$



## Convergence Measure

- Most optimization algorithms are iterative

$$
\mathbf{w}_{t+1}=\mathbf{w}_{t}+\Delta \mathbf{w}_{t}
$$

- Iteration Complexity: the number of iterations $T(\epsilon)$ needed to have

$$
F\left(\widehat{\mathbf{w}}_{T}\right)-\min _{\mathbf{w}} F(\mathbf{w}) \leq \epsilon \quad(\epsilon \ll 1)
$$



## Convergence Measure

- Most optimization algorithms are iterative

$$
\mathbf{w}_{t+1}=\mathbf{w}_{t}+\Delta \mathbf{w}_{t}
$$

- Iteration Complexity: the number of iterations $T(\epsilon)$ needed to have

$$
F\left(\widehat{\mathbf{w}}_{T}\right)-\min _{\mathbf{w}} F(\mathbf{w}) \leq \epsilon \quad(\epsilon \ll 1)
$$

- Convergence Rate: after $T$ iterations, how good is the solution


$$
F\left(\widehat{\mathbf{w}}_{T}\right)-\min _{\mathbf{w}} F(\mathbf{w}) \leq \epsilon(T)
$$

## Convergence Measure

- Most optimization algorithms are iterative

$$
\mathbf{w}_{t+1}=\mathbf{w}_{t}+\Delta \mathbf{w}_{t}
$$

- Iteration Complexity: the number of iterations $T(\epsilon)$ needed to have

$$
F\left(\widehat{\mathbf{w}}_{T}\right)-\min _{\mathbf{w}} F(\mathbf{w}) \leq \epsilon \quad(\epsilon \ll 1)
$$

- Convergence Rate: after $T$ iterations, how good is the solution


$$
F\left(\widehat{\mathbf{w}}_{T}\right)-\min _{\mathbf{w}} F(\mathbf{w}) \leq \epsilon(T)
$$

- Total Runtime $=$ Per-iteration Cost $\times$ Iteration Complexity


## More on Convergence Measure

- Big $O(\cdot)$ notation: explicit dependence on $T$ or $\epsilon$

|  | Convergence Rate |  | Iteration Complexity |
| :---: | :---: | :---: | :---: |
| linear | $O\left(\mu^{T}\right)$ | $(\mu<1)$ | $O\left(\log \left(\frac{1}{\epsilon}\right)\right)$ |
| sub-linear | $O\left(\frac{1}{T^{\alpha}}\right)$ | $\alpha>0$ | $O\left(\frac{1}{\epsilon^{1 / \alpha}}\right)$ |

## More on Convergence Measure

- Big $O(\cdot)$ notation: explicit dependence on $T$ or $\epsilon$

|  | Convergence Rate |  |
| :---: | :---: | :---: |
| Iteration Complexity |  |  |
| linear | $O\left(\mu^{T}\right)$ | $(\mu<1)$ |$) O\left(\log \left(\frac{1}{\epsilon}\right)\right)$.

Why are we interested in Bounds?

## More on Convergence Measure

|  | Convergence Rate | Iteration Complexity |
| :---: | :---: | :---: |
| linear | $O\left(\mu^{T}\right) \quad(\mu<1)$ | $O\left(\log \left(\frac{1}{\epsilon}\right)\right)$ |
| sub-linear | $O\left(\frac{1}{T^{\alpha}}\right) \quad \alpha>0$ | $O\left(\frac{1}{\epsilon^{1 / \alpha}}\right)$ |



## More on Convergence Measure

|  | Convergence Rate | Iteration Complexity |
| :---: | :---: | :---: |
| linear | $O\left(\mu^{T}\right) \quad(\mu<1)$ | $O\left(\log \left(\frac{1}{\epsilon}\right)\right)$ |
| sub-linear | $O\left(\frac{1}{T^{\alpha}}\right) \quad \alpha>0$ | $O\left(\frac{1}{\epsilon^{1 / \alpha}}\right)$ |



## More on Convergence Measure

|  | Convergence Rate | Iteration Complexity |
| :---: | :---: | :---: |
| linear | $O\left(\mu^{T}\right) \quad(\mu<1)$ | $O\left(\log \left(\frac{1}{\epsilon}\right)\right)$ |
| sub-linear | $O\left(\frac{1}{T^{\alpha}}\right) \quad \alpha>0$ | $O\left(\frac{1}{\epsilon^{1 / \alpha}}\right)$ |



## More on Convergence Measure

|  | Convergence Rate |  |
| :---: | :---: | :---: |
| Iteration Complexity |  |  |
| linear | $O\left(\mu^{T}\right) \quad(\mu<1)$ | $O\left(\log \left(\frac{1}{\epsilon}\right)\right)$ |
| sub-linear | $O\left(\frac{1}{T^{\alpha}}\right) \quad \alpha>0$ | $O\left(\frac{1}{\epsilon^{1 / \alpha}}\right)$ |



## Non-smooth V.S. Smooth

- Non-smooth $\ell(z)$
- hinge loss: $\ell\left(\mathbf{w}^{\top} \mathbf{x}, y\right)=\max \left(0,1-y \mathbf{w}^{\top} \mathbf{x}\right)$
- absolute loss: $\ell\left(\mathbf{w}^{\top} \mathbf{x}, y\right)=\left|\mathbf{w}^{\top} \mathbf{x}-y\right|$
- Smooth $\ell(z)$
- squared hinge loss: $\ell\left(\mathbf{w}^{\top} \mathbf{x}, y\right)=\max \left(0,1-y \mathbf{w}^{\top} \mathbf{x}\right)^{2}$
- logistic loss: $\ell\left(\mathbf{w}^{\top} \mathbf{x}, y\right)=\log \left(1+\exp \left(-y \mathbf{w}^{\top} \mathbf{x}\right)\right)$
- square loss: $\ell\left(\mathbf{w}^{\top} \mathbf{x}, y\right)=\left(\mathbf{w}^{\top} \mathbf{x}-y\right)^{2}$


## Strong convex V.S. Non-strongly convex

- $\lambda$-strongly convex $R(\mathbf{w})$
- $\ell_{2}$ regularizer: $\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}$
- Elastic net regularizer: $\tau\|\mathbf{w}\|_{1}+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}$
- Non-strongly convex $R(\mathbf{w})$
- unregularized problem: $R(\mathbf{w}) \equiv 0$
- $\ell_{1}$ regularizer: $\tau\|\mathbf{w}\|_{1}$


## Gradient Method in Machine Learning

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Suppose $\ell(z)$ is smooth
- Full gradient: $\nabla F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}$
- Per-iteration cost: $O(n d)$


## Gradient Method in Machine Learning

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Suppose $\ell(z)$ is smooth
- Full gradient: $\nabla F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}$
- Per-iteration cost: $O(n d)$
Gradient Descent $\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \nabla F\left(\mathbf{w}_{t-1}\right)$


## Gradient Method in Machine Learning

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- If $\lambda=0: R(\mathbf{w})$ is non-strongly convex
- Iteration complexity $O\left(\frac{1}{\epsilon}\right)$
- If $\lambda>0: R(\mathbf{w})$ is $\lambda$-strongly convex
- Iteration complexity $O\left(\frac{1}{\lambda} \log \left(\frac{1}{\epsilon}\right)\right)$


## Accelerated Gradient Method

## Accelerated Gradient Descent

$$
\begin{aligned}
& \mathbf{w}_{t}=\mathbf{v}_{t-1}-\gamma_{t} \nabla F\left(\mathbf{v}_{t-1}\right) \\
& \mathbf{v}_{t}=\mathbf{w}_{t}+\eta_{t}\left(\mathbf{w}_{t}-\mathbf{w}_{t-1}\right)
\end{aligned}
$$

- $\mathbf{w}_{t}$ is the output and $\mathbf{v}_{t}$ is an auxiliary sequence.


## Accelerated Gradient Method

## Accelerated Gradient Descent

$$
\begin{aligned}
& \mathbf{w}_{t}=\mathbf{v}_{t-1}-\gamma_{t} \nabla F\left(\mathbf{v}_{t-1} \begin{array}{c}
\text { Momentum } \\
\text { Step }
\end{array}\right. \\
& \mathbf{v}_{t}=\mathbf{w}_{t}+\eta_{t}\left(\mathbf{w}_{t}-\mathbf{w}_{t-1}\right)
\end{aligned}
$$

- $\mathbf{w}_{t}$ is the output and $\mathbf{v}_{t}$ is an auxiliary sequence.


## Accelerated Gradient Method

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- If $\lambda=0: R(\mathbf{w})$ is non-strongly convex
- Iteration complexity $O\left(\frac{1}{\sqrt{\epsilon}}\right)$, better than $O\left(\frac{1}{\epsilon}\right)$
- If $\lambda>0: R(\mathbf{w})$ is $\lambda$-strongly convex
- Iteration complexity $O\left(\frac{1}{\sqrt{\lambda}} \log \left(\frac{1}{\epsilon}\right)\right)$, better than $O\left(\frac{1}{\lambda} \log \left(\frac{1}{\epsilon}\right)\right)$ for small $\lambda$


## Deal with $\ell_{1}$ regularizer

Consider a more general case

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\underbrace{R^{\prime}(\mathbf{w})+\tau\|\mathbf{w}\|_{1}}_{R(\mathbf{w})}
$$

- $R(\mathbf{w})=R^{\prime}(\mathbf{w})+\tau\|\mathbf{w}\|_{1}$
- $R^{\prime}(\mathbf{w}): \lambda$-strongly convex and smooth


## Deal with $\ell_{1}$ regularizer

Consider a more general case

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\underbrace{\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R^{\prime}(\mathbf{w})}_{F^{\prime}(\mathbf{w})}+\tau\|\mathbf{w}\|_{1}
$$

- $R(\mathbf{w})=R^{\prime}(\mathbf{w})+\tau\|\mathbf{w}\|_{1}$
- $R^{\prime}(\mathbf{w}): \lambda$-strongly convex and smooth


## Deal with $\ell_{1}$ regularizer

## Accelerated Gradient Descent

$$
\begin{gathered}
\mathbf{w}_{t}=\underset{\mathbf{w} \in \mathbb{R}^{d}}{\arg \min } \nabla F^{\prime}\left(\mathbf{v}_{t-1}\right)^{\top} \mathbf{w}+\frac{1}{2 \gamma_{t}}\left\|\mathbf{w}-\mathbf{v}_{t-1}\right\|_{2}^{2}+\tau\|\mathbf{w}\|_{1} \\
\mathbf{v}_{t}=\mathbf{w}_{t}+\eta_{t}\left(\mathbf{w}_{t}-\mathbf{w}_{t-1}\right)
\end{gathered}
$$

- Proximal mapping has close-form solution: Soft-thresholding
- Iteration complexity and runtime remain unchanged.


## Deal with $\ell_{1}$ regularizer



- Proximal mapping has close-form solution: Soft-thresholding
- Iteration complexity and runtime remain unchanged.


## Sub-Gradient Method in Machine Learning

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Suppose $\ell(z)$ is non-smooth
- Full sub-gradient: $\partial F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \partial \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}$



## Sub-Gradient Method in Machine Learning

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Suppose $\ell(z)$ is non-smooth
- Full sub-gradient: $\partial F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \partial \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}$


## Sub-Gradient Descent

$$
\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \partial F\left(\mathbf{w}_{t-1}\right)
$$

## Sub-Gradient Method

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- If $\lambda=0: R(\mathbf{w})$ is non-strongly convex
- Iteration complexity $O\left(\frac{1}{\epsilon^{2}}\right)$
- If $\lambda>0: R(\mathbf{w})$ is $\lambda$-strongly convex
- Iteration complexity $O\left(\frac{1}{\lambda \epsilon}\right)$
- No efficient acceleration scheme in general


## Problem Classes and Iteration Complexity

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})
$$

- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-strongly convex | $O\left(\frac{1}{\epsilon^{2}}\right)$ |
|  | $\lambda$-strongly convex | $O\left(\frac{1}{\lambda \epsilon}\right)$ | $O\left(\frac{1}{\sqrt{\epsilon}}\right)$ |

- Per-iteration cost: $O(n d)$, too high if $n$ or $d$ are large.


## Outline

(2) Optimization

- (Sub)Gradient Methods
- Stochastic Optimization Algorithms for Big Data
- Stochastic Optimization
- Distributed Optimization


## Stochastic First-Order Method by Data Sampling

- Stochastic Gradient Descent (SGD)
- Stochastic Variance Reduced Gradient (SVRG)
- Stochastic Average Gradient Algorithm (SAGA)
- Stochastic Dual Coordinate Ascent (SDCA)
- Accelerated Proximal Coordinate Gradient (APCG)

$$
\text { Assumption: }\left\|\mathbf{x}_{i}\right\| \leq 1 \text { for any } i
$$

## Basic SGD (Nemirovski \& Yudin (1978))

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Full sub-gradient: $\partial F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \partial \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}$
- Randomly sample $i \in\{1, \ldots, n\}$
- Stochastic sub-gradient: $\partial \ell\left(\mathbf{w}^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}$

$$
\mathbb{E}_{i}\left[\partial \ell\left(\mathbf{w}^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}\right]=\partial F(\mathbf{w})
$$

## Basic SGD (Nemirovski \& Yudin (1978))

Applicable in all settings!

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$


output:


## Basic SGD (Nemirovski \& Yudin (1978))

Applicable in all settings!

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

sample: $\quad i_{t} \in\{1, \ldots, n\}$
update: $\quad \mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t}\left(\partial \ell\left(\mathbf{w}_{t-1}^{T} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}_{t-1}\right)$
output: $\quad \overline{\mathbf{w}}_{T}=\frac{1}{T} \sum_{t=1}^{T} \mathbf{w}_{t}$

## Basic SGD (Nemirovski \& Yudin (1978))

$$
F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- If $\lambda=0: R(\mathbf{w})$ is non-strongly convex
- Iteration complexity $O\left(\frac{1}{\epsilon^{2}}\right)$
- If $\lambda>0: R(\mathbf{w})$ is $\lambda$-strongly convex
- Iteration complexity $O\left(\frac{1}{\lambda \epsilon}\right)$
- Exactly the same as sub-gradient descent!


## Total Runtime

- Per-iteration cost: $O(d)$
- Much lower than full gradient method
- e.g. hinge loss (SVM)
stochastic gradient: $\partial \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)=\left\{\begin{array}{cc}-y_{i_{t}} \mathbf{x}_{i_{t}}, & 1-y_{i_{t}} \mathbf{w}^{\top} \mathbf{x}_{i_{t}}>0 \\ 0, & \text { otherwise }\end{array}\right.$


## Total Runtime

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})
$$

- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-strongly convex | $O\left(\frac{1}{\epsilon^{2}}\right)$ |
|  | $\lambda$-strongly convex | $O\left(\frac{1}{\epsilon^{2}}\right)$ |  |
|  |  | $O\left(\frac{1}{\lambda \epsilon}\right)$ | $O\left(\frac{1}{\lambda \epsilon}\right)$ |

- For SGD, only strongly convexity helps but the smoothness does not make any difference!


## Full Gradient V.S. Stochastic Gradient

- Full gradient method needs fewer iterations
- Stochastic gradient method has lower cost per iteration
- For small $\epsilon$, use full gradient
- Satisfied with large $\epsilon$, use stochastic gradient
- Full gradient can be accelerated
- Stochastic gradient cannot
- Full gradient's iterations complexity depends on smoothness and strong convexity
- Stochastic gradient's iteration complexity only depends on strong convexity

SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Applicable when $\ell(z)$ is smooth and $R(\mathbf{w})$ is $\lambda$-strongly convex
- Stochastic gradient:

- $\mathrm{E}_{i_{t}}\left[g_{i_{t}}(\mathbf{w})\right]=\nabla F(\mathbf{w})$ but...
- $\operatorname{Var}\left[g_{i t}(\mathbf{w})\right] \neq 0$ even if $\mathbf{w}=\mathbf{w}$


## SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- Applicable when $\ell(z)$ is smooth and $R(\mathbf{w})$ is $\lambda$-strongly convex
- Stochastic gradient:

$$
g_{i_{t}}(\mathbf{w})=\nabla \ell\left(\mathbf{w}^{T} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}
$$

- $\mathrm{E}_{i_{t}}\left[g_{i_{t}}(\mathbf{w})\right]=\nabla F(\mathbf{w})$ but...
- $\operatorname{Var}\left[g_{i_{t}}(\mathbf{w})\right] \neq 0$ even if $\mathbf{w}=\mathbf{w}^{\star}$


## SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

- Compute the full gradient at a reference point $\tilde{\mathbf{w}}$

$$
\nabla F(\tilde{\mathbf{w}})=\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\tilde{\mathbf{w}}^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \tilde{\mathbf{w}}
$$

- Stochastic variance reduced gradient:

$$
\tilde{g}_{i_{t}}(\mathbf{w})=\nabla F(\tilde{\mathbf{w}})-g_{i_{t}}(\tilde{\mathbf{w}})+g_{i_{t}}(\mathbf{w})
$$

- $\mathrm{E}_{i_{t}}\left[\tilde{g}_{i_{t}}(\mathbf{w})\right]=\nabla F(\mathbf{w})$
- $\operatorname{Var}\left[\tilde{g}_{i_{t}}(\mathbf{w})\right] \longrightarrow 0$ as $\tilde{\mathbf{w}}, \mathbf{w} \rightarrow \mathbf{w}^{\star}$

SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

- At optimal solution $\mathbf{w}^{\star}: \nabla F\left(\mathbf{w}^{\star}\right)=\mathbf{0}$
- It does not mean

$$
g_{i_{t}}(\mathbf{w}) \longrightarrow \mathbf{0}
$$

as $\mathbf{w} \rightarrow \mathbf{w}^{\star}$

- However, we have

$$
\tilde{g}_{i_{t}}(\mathbf{w})=\nabla F(\tilde{\mathbf{w}})-g_{i_{t}}(\tilde{\mathbf{w}})+g_{i_{t}}(\mathbf{w}) \longrightarrow \mathbf{0}
$$

as $\tilde{\mathbf{w}}, \mathbf{w} \rightarrow \mathbf{w}^{\star}$

SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

Iterate $s=1, \ldots, T-1$
Let $\mathbf{w}_{0}=\tilde{\mathbf{w}}_{s}$ and compute $\nabla F\left(\tilde{\mathbf{w}}_{s}\right)$
Iterate $t=1, \ldots, K$
$\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right)$
$\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)$
$\tilde{\mathbf{w}}_{s+1}=\frac{1}{K} \sum_{t=1}^{K} \mathbf{w}_{t}$
output: $\tilde{\mathbf{w}}_{T}$

- $K=O\left(\frac{1}{\lambda}\right)$

SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

- Per-iteration cost: $O\left(d\left(n+\frac{1}{\lambda}\right)\right)$
- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-smooth | Smooth |
|  | $\lambda$-strongly convex | N.A. | N.A. |

- Use proximal mapping for $\ell_{1}$ regularizer

SVRG (Johnson \& Zhang, 2013; Zhang et al., 2013; Xiao \& Zhang, 2014)

- Per-iteration cost: $O\left(d\left(n+\frac{1}{\lambda}\right)\right)$
- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-smooth | Smooth |
|  | $\lambda$-strongly convex | N.A. | N.A. |

- Total Runtime: $O\left(d\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$
- Use proximal mapping for $\ell_{1}$ regularizer


## SAGA (Defazio et al. (2014))

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

- A new version of SAG (Roux et al. (2012))
- Applicable when $\ell(z)$ is smooth
- Strong convexity is not required.


## SAGA (Defazio et al. (2014))

- SAGA also reduces the variance of stochastic gradient but with a different technique
- SVRG uses gradients at the same point $\tilde{w}$

$$
\begin{aligned}
\tilde{g}_{i_{t}}(\mathbf{w}) & =\nabla F(\tilde{\mathbf{w}})-g_{i_{t}}(\tilde{\mathbf{w}})+g_{i_{t}}(\mathbf{w}) \\
\nabla F(\tilde{\mathbf{w}}) & =\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\tilde{\mathbf{w}}^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \tilde{\mathbf{w}}
\end{aligned}
$$

- SAGA uses gradients at different point $\left\{\tilde{\mathbf{w}}_{1}, \tilde{\mathbf{w}}_{2}, \cdots, \tilde{\mathbf{w}}_{n}\right\}$

$$
\begin{aligned}
\tilde{g}_{i_{t}}(\mathbf{w}) & =G-g_{i_{t}}\left(\tilde{\mathbf{w}}_{i_{t}}\right)+g_{i_{t}}(\mathbf{w}) \\
G & =\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\tilde{\mathbf{w}}_{i}^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \tilde{\mathbf{w}}_{i_{t}}
\end{aligned}
$$

## SAGA (Defazio et al. (2014))

- SAGA also reduces the variance of stochastic gradient but with a different technique
- SVRG uses gradients at the same point $\tilde{w}$

$$
\begin{aligned}
\tilde{g}_{i_{t}}(\mathbf{w}) & =\nabla F(\tilde{\mathbf{w}})-g_{i_{t}}(\tilde{\mathbf{w}})+g_{i_{t}}(\mathbf{w}) \\
\nabla F(\tilde{\mathbf{w}}) & =\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\tilde{\mathbf{w}}^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \tilde{\mathbf{w}}
\end{aligned}
$$

- SAGA uses gradients at different point $\left\{\tilde{\mathbf{w}}_{1}, \tilde{\mathbf{w}}_{2}, \cdots, \tilde{\mathbf{w}}_{n}\right\}$

$$
\tilde{g}_{i_{t}} \text { Average Gradient } \boldsymbol{w}^{\text {ow }}=\frac{1}{n} \sum_{i=1}^{n} \nabla \ell\left(\tilde{\mathbf{w}}_{i}{ }^{T} \mathbf{x}_{i}, y_{i}\right)+\lambda \tilde{\mathbf{w}}_{i_{t}}
$$

## SAGA (Defazio et al. (2014))

- Initialize average gradient $G_{0}$ :

$$
G_{0}=\frac{1}{n} \sum_{i=1}^{n} g_{i}, \quad g_{i}=\nabla \ell\left(\mathbf{w}_{0}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}_{0}
$$

- stochastic variance reduced gradient:

- Update average gradient



## SAGA (Defazio et al. (2014))

- Initialize average gradient $G_{0}$ :

$$
G_{0}=\frac{1}{n} \sum_{i=1}^{n} g_{i}, \quad g_{i}=\nabla \ell\left(\mathbf{w}_{0}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}_{0}
$$

- stochastic variance reduced gradient:

$$
\begin{aligned}
\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right) & =G_{t-1}-g_{i_{t}}+\left(\nabla \ell\left(\mathbf{w}_{t-1}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}_{t-1}\right) \\
\mathbf{w}_{t} & =\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)
\end{aligned}
$$

- Update average gradient



## SAGA (Defazio et al. (2014))

- Initialize average gradient $G_{0}$ :

$$
G_{0}=\frac{1}{n} \sum_{i=1}^{n} g_{i}, \quad g_{i}=\nabla \ell\left(\mathbf{w}_{0}^{\top} \mathbf{x}_{i}, y_{i}\right)+\lambda \mathbf{w}_{0}
$$

- stochastic variance reduced gradient:

$$
\begin{aligned}
\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right) & =G_{t-1}-g_{i_{t}}+\left(\nabla \ell\left(\mathbf{w}_{t-1}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}_{t-1}\right) \\
\mathbf{w}_{t} & =\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)
\end{aligned}
$$

- Update average gradient

$$
G_{t}=\frac{1}{n} \sum_{i=1}^{n} g_{i}, \quad g_{i_{t}}=\nabla \ell\left(\mathbf{w}_{t-1}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}_{t-1}
$$

## SAGA: efficient update of averaged gradient

- $G_{t}$ and $G_{t-1}$ only differs in $g_{i}$ for $i=i_{t}$
- Before we update $g_{i}$, we update

$$
G_{t}=\frac{1}{n} \sum_{i=1}^{n} g_{i}=G_{t-1}-\frac{1}{n} g_{i_{t}}+\frac{1}{n}\left(\nabla \ell\left(\mathbf{w}_{t-1}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}_{t-1}\right)
$$

- computation cost: $O(d)$
- Require extra memory space $O$ (nd)


## SAGA: efficient update of averaged gradient

- $G_{t}$ and $G_{t-1}$ only differs in $g_{i}$ for $i=i_{t}$
- Before we update $g_{i}$, we update

$$
G_{t}=\frac{1}{n} \sum_{i=1}^{n} g_{i}=G_{t-1}-\frac{1}{n} g_{i_{t}}+\frac{1}{n}\left(\nabla \ell\left(\mathbf{w}_{t-1}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\lambda \mathbf{w}_{t-1}\right)
$$

- computation cost: $O(d)$
- To implemente SAGA, we have to store and update all: $g_{1}, g_{2}, \ldots, g_{n}$
- Require extra memory space $O(n d)$


## SAGA (Defazio et al. (2014))

- Per-iteration cost: $O(d)$
- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
|  |  | Non-smooth | Smooth |
| $R(\mathbf{w})$ | Non-strongly convex | N.A. | $O\left(\frac{n}{\epsilon}\right)$ |
|  | $\lambda$-strongly convex | N.A. | $O\left(\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$ |

- Use proximal mapping for $\ell_{1}$ regularizer


## SAGA (Defazio et al. (2014))

- Per-iteration cost: $O(d)$
- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-strongly convex | N.A. |
|  | $\lambda$-strongly convex | N.A. | $O\left(\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$ |

- Total Runtime (strongly convex): $O\left(d\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$. Same as SVRG!
- Use proximal mapping for $\ell_{1}$ regularizer


## Compare the Runtime of SGD and SVRG/SAGA

- Smooth but non-strongly convex:
- SGD: $O\left(\frac{d}{\epsilon^{2}}\right)$
- SAGA: $O\left(\frac{d n}{\epsilon}\right)$
- Smooth and strongly convex:
- SGD: $O\left(\frac{d}{\lambda \epsilon}\right)$
- SVRG/SAGA: $O\left(d\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$
- For small $\epsilon$, use SVRG/SAGA
- Satisfied with large $\epsilon$, use SGD


## Conjugate Duality

- Define $\ell_{i}(z) \equiv \ell\left(z, y_{i}\right)$
- Conjugate function: $\ell_{i}^{*}(\alpha) \Longleftrightarrow \ell_{i}(z)$

$$
\ell_{i}(z)=\max _{\alpha \in \mathbb{R}}\left[\alpha z-\ell^{*}(\alpha)\right], \quad \ell_{i}^{*}(\alpha)=\max _{z \in \mathbb{R}}[\alpha z-\ell(z)]
$$

- E.g. hinge loss: $\ell_{i}(z)=\max \left(0,1-y_{i} z\right)$
- E.g. square hinge loss: $\ell_{i}(z)=\max \left(0,1-y_{i} z\right)^{2}$


## Conjugate Duality

- Define $\ell_{i}(z) \equiv \ell\left(z, y_{i}\right)$
- Conjugate function: $\ell_{i}^{*}(\alpha) \Longleftrightarrow \ell_{i}(z)$

$$
\ell_{i}(z)=\max _{\alpha \in \mathbb{R}}\left[\alpha z-\ell^{*}(\alpha)\right], \quad \ell_{i}^{*}(\alpha)=\max _{z \in \mathbb{R}}[\alpha z-\ell(z)]
$$

- E.g. hinge loss: $\ell_{i}(z)=\max \left(0,1-y_{i} z\right)$

$$
\ell_{i}^{*}(\alpha)= \begin{cases}\alpha y_{i} & \text { if }-1 \leq \alpha y_{i} \leq 0 \\ +\infty & \text { otherwise }\end{cases}
$$

- E.g. square hinge loss: $\ell_{i}(z)=\max \left(0,1-y_{i} z\right)^{2}$

$$
\ell_{i}^{*}(\alpha)= \begin{cases}\frac{\alpha^{2}}{4}+\alpha y_{i} & \text { if } \alpha y_{i} \leq 0 \\ +\infty & \text { otherwise }\end{cases}
$$

## SDCA (Shalev-Shwartz \& Zhang (2013))

- Stochastic Dual Coordinate Ascent (liblinear (Hsieh et al., 2008))
- Applicable when $R(\mathbf{w})$ is $\lambda$-strongly convex
- Smoothness is not required
- From Primal problem to Dual problem:



## SDCA (Shalev-Shwartz \& Zhang (2013))

- Stochastic Dual Coordinate Ascent (liblinear (Hsieh et al., 2008))
- Applicable when $R(\mathbf{w})$ is $\lambda$-strongly convex
- Smoothness is not required
- From Primal problem to Dual problem:

$$
\begin{aligned}
& \min _{\mathbf{w}} \frac{1}{n} \sum_{i=1}^{n} \ell(\underbrace{\mathbf{w}^{\top} \mathbf{x}_{i}}_{z}, y_{i})+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2} \\
= & \min _{\mathbf{w}} \frac{1}{n} \sum_{i=1}^{n} \max _{\alpha_{i} \in \mathbb{R}}\left[\alpha_{i}\left(\mathbf{w}^{\top} \mathbf{x}_{i}\right)-\ell_{i}^{*}\left(\alpha_{i}\right)\right]+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2} \\
= & \max _{\alpha \in \mathbb{R}^{n}} \frac{1}{n} \sum_{i=1}^{n}-\ell_{i}^{*}\left(\alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda n} \sum_{i=1}^{n} \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
\end{aligned}
$$

## SDCA (Shalev-Shwartz \& Zhang (2013))

- Solve Dual Problem:

$$
\max _{\alpha \in \mathbb{R}^{n}} \frac{1}{n} \sum_{i=1}^{n}-\ell_{i}^{*}\left(\alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda n} \sum_{i=1}^{n} \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
$$

- Sample $i_{t} \in\{1, \ldots, n\}$. Optimize $\alpha_{i_{t}}$ while fixing others


## SDCA (Shalev-Shwartz \& Zhang (2013))

- Maintain a primal solution: $\mathbf{w}_{t}=\frac{1}{\lambda n} \sum_{i=1}^{n} \alpha_{i}^{t} \mathbf{x}_{i}$
- Change variable $\alpha_{i} \longrightarrow \Delta \alpha_{i}$

$$
\begin{aligned}
& \max _{\Delta \alpha \in \mathbb{R}^{n}} \frac{1}{n} \sum_{i=1}^{n}-\ell_{i}^{*}\left(\alpha_{i}^{t}+\Delta \alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda n}\left(\sum_{i=1}^{n} \alpha_{i}^{t} \mathbf{x}_{i}+\sum_{i=1}^{n} \Delta \alpha_{i} \mathbf{x}_{i}\right)\right\|_{2}^{2} \\
\Longleftrightarrow & \max _{\Delta \alpha \in \mathbb{R}^{n}} \frac{1}{n} \sum_{i=1}^{n}-\ell_{i}^{*}\left(\alpha_{i}^{t}+\Delta \alpha_{i}\right)-\frac{\lambda}{2}\left\|\mathbf{w}_{t}+\frac{1}{\lambda n} \sum_{i=1}^{n} \Delta \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
\end{aligned}
$$

## SDCA (Shalev-Shwartz \& Zhang (2013))

## Dual Coordinate Updates

$$
\begin{aligned}
\Delta \alpha_{i_{t}} & =\max _{\Delta \alpha_{i_{t}} \in \mathbb{R}^{n}}-\frac{1}{n} \ell_{i_{t}}^{*}\left(-\alpha_{i_{t}}^{t}-\Delta \alpha_{i_{t}}\right)-\frac{\lambda}{2}\left\|\mathbf{w}_{t}+\frac{1}{\lambda n} \Delta \alpha_{i_{t}} \mathbf{x}_{i_{t}}\right\|_{2}^{2} \\
\alpha_{i_{t}}^{t+1} & =\alpha_{i_{t}}^{t}+\Delta \alpha_{i_{t}} \\
\mathbf{w}_{t+1} & =\mathbf{w}_{t}+\frac{1}{\lambda n} \Delta \alpha_{i_{t}} \mathbf{x}_{i}
\end{aligned}
$$

## SDCA updates

- Close-form solution for $\Delta \alpha_{i}$ : hinge loss, squared hinge loss, absolute loss and square loss (Shalev-Shwartz \& Zhang (2013))
- e.g. square loss

$$
\Delta \alpha_{i}=\frac{y_{i}-\mathbf{w}_{t}^{\top} \mathbf{x}_{i}-\alpha_{i}^{t}}{1+\left\|\mathbf{x}_{i}\right\|_{2}^{2} /(\lambda n)}
$$

- Per-iteration cost: $O(d)$
- Approximate solution: logistic loss (Shalev-Shwartz \& Zhang (2013))


## SDCA

- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
|  |  | Non-smooth | Smooth |
| $R(\mathbf{w})$ | Non-strongly convex | N.A. | N.A. |
|  | $\lambda$-strongly convex | $O\left(n+\frac{1}{\lambda \epsilon}\right)$ | $O\left(\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$ |

## SDCA

- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-smooth | Smooth |
|  | Non-strongly convex | N.A. | N.A. |

- Total Runtime (smooth loss): $O\left(d\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)\right)$. The same as SVRG and SAGA!


## SVRG V.S. SDCA V.S. SGD

- $\ell_{2}$-regularized logistic regression with $\lambda=10^{-4}$
- MNIST data
- Johnson \& Zhang (2013)




## APCG (Lin et al. (2014))

- Recall the acceleration scheme for full gradient method
- Auxiliary sequence ( $\beta^{t}$ )
- Momentum step
- Maintain a primal solution: $\mathbf{w}_{t}=\frac{1}{\lambda n} \sum_{i=1}^{n} \beta_{i}^{t} \mathbf{x}_{i}$


## Dual Coordinate Updates

$$
\begin{aligned}
\Delta \beta_{i_{t}} & =\max _{\Delta \beta_{i_{t}} \in \mathbb{R}^{n}}-\frac{1}{n} \ell_{i_{t}}^{*}\left(-\beta_{i_{t}}^{t}-\Delta \beta_{i_{t}}\right)-\frac{\lambda}{2}\left\|\mathbf{w}_{t}+\frac{1}{\lambda n} \Delta \beta_{i_{t}} \mathbf{x}_{i_{t}}\right\|_{2}^{2} \\
\alpha_{i_{t}}^{t+1} & =\beta_{i_{t}}^{t}+\Delta \beta_{i_{t}} \\
\beta^{t+1} & =\alpha^{t+1}+\eta_{t}\left(\alpha^{t+1}-\alpha^{t}\right)
\end{aligned}
$$

## APCG (Lin et al. (2014))

- Recall the acceleration scheme for full gradient method
- Auxiliary sequence ( $\beta^{t}$ )
- Momentum step
- Maintain a primal solution: $\mathbf{w}_{t}=\frac{1}{\lambda n} \sum_{i=1}^{n} \beta_{i}^{t} \mathbf{x}_{i}$


## Dual Coordinate Updates

$$
\begin{aligned}
\Delta \beta_{i_{t}} & =\max _{\Delta \beta_{i_{t} \in \mathbb{R}^{n}}-\frac{1}{n} \ell_{i_{t}}^{*}\left(-\beta_{i_{i}}^{t}-\Delta \beta_{i}\right)-\frac{\lambda}{\text { Momentum }}\left\|_{\mathbf{w}_{t}}+\frac{1}{\lambda n} \Delta \beta_{i_{t}} \mathbf{x}_{i_{t}}\right\|_{2}^{2}}^{2} \\
\alpha_{i_{t}}^{t+1} & =\beta_{i_{t}}^{t}+\Delta \beta_{i_{t}} \\
\beta^{t+1} & =\alpha^{t+1}+\eta_{t}\left(\alpha^{t+1}-\alpha^{t}\right)
\end{aligned}
$$

## APCG (Lin et al. (2014))

- Per-iteration cost: $O(d)$
- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-smooth | Smooth |
|  | $\lambda$-strongly convex | $O\left(n+\sqrt{\frac{n}{\lambda \epsilon}}\right)$ | $O\left(\left(n+\sqrt{\frac{n}{\lambda}}\right) \log \left(\frac{1}{\epsilon}\right)\right)$ |

- Compared to SDCA, APCG has shorter runtime when $\lambda$ is very small.


## APCG V.S. SDCA

- squared hinge loss SVM
- real data

| datasets | number of samples $n$ | number of features $d$ | sparsity |
| :---: | :---: | :---: | :---: |
| rcv1 | 20,242 | 47,236 | $0.16 \%$ |
| covtype | 581,012 | 54 | $22 \%$ |
| news20 | 19,996 | $1,355,191$ | $0.04 \%$ |

- $F\left(\mathbf{w}_{t}\right)-F\left(\mathbf{w}^{\star}\right)$ V.S. the number of passes of data


## APCG V.S. SDCA

Lin et al. (2014)


## For general $R(\mathbf{w})$

- Dual Problem:

$$
\max _{\alpha \in \mathbb{R}^{n}} \frac{1}{n} \sum_{i=1}^{n}-\ell_{i}^{*}\left(\alpha_{i}\right)-R^{*}\left(\frac{1}{\lambda n} \sum_{i=1}^{n} \alpha_{i} \mathbf{x}_{i}\right)
$$

- $R^{*}$ is the conjugate of $R$
- Sample $i_{t} \in\{1, \ldots, n\}$. Optimize $\alpha_{i_{t}}$ while fixing others
- Can be still updated in $O(d)$ in many cases (Shalev-Shwartz \& Zhang (2013))
- Iteration complexity and runtime of SDCA and APCG remain unchanged.


## APCG for primal problem

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}+\tau\|\mathbf{w}\|_{1}
$$

- Suppose $d \gg n$. Per-iteration cost $O(d)$ is too high
- Apply APCG to the primal instead of dual problem
- Sample over features instead of data
- Per-iteration cost becomes $O(n)$


## APCG for primal problem

$$
\begin{gathered}
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{2}\|X \mathbf{w}-\mathbf{y}\|_{2}^{2}+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}+\tau\|\mathbf{w}\|_{1} \\
X=\left[\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{n}\right]
\end{gathered}
$$

- Full gradient: $\nabla F(\mathbf{w})=X^{T}(X \mathbf{w}-\mathbf{y})+\lambda \mathbf{w}$
- Partial gradient: $\nabla_{i} F(\mathbf{w})=x_{i}^{T}(X \mathbf{w}-\mathbf{y})+\lambda w_{i}$


## APCG for primal problem

$$
\begin{gathered}
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{2}\|X \mathbf{w}-\mathbf{y}\|_{2}^{2}+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}+\tau\|\mathbf{w}\|_{1} \\
X=\left[\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{n}\right]
\end{gathered}
$$

- Full gradient: $\nabla F(\mathbf{w})=X^{T}(X \mathbf{w}-\mathbf{y})+\lambda \mathbf{w}$
- Partial gradient: $\nabla_{i} F(\mathbf{w})=x_{i}^{T}(X \mathbf{w}-\mathbf{y})+\lambda w_{i}$


## Proximal Coordinate Gradient (PCG) (Nesterov (2012))

$$
\mathbf{w}_{i}^{t}= \begin{cases}\arg \min _{w \in \mathbb{R}} \nabla_{i} F\left(\mathbf{w}^{t-1}\right) w_{i}+\frac{1}{2 \gamma_{t}}\left(w_{i}-\mathbf{w}_{i}^{t-1}\right)^{2}+\tau\left|w_{i}\right| & \text { if } i=t_{i} \\ \mathbf{w}_{i}^{t-1} & \text { otherwise }\end{cases}
$$

- $\nabla_{i} F\left(\mathbf{w}^{t}\right)$ can be updated in $O(n)$


## APCG for primal problem

$$
\begin{gathered}
\min _{\mathbf{w} \in \mathbb{R}^{d}} F(\mathbf{w})=\frac{1}{2}\|X \mathbf{w}-\mathbf{y}\|_{2}^{2}+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}+\tau\|\mathbf{w}\|_{1} \\
X=\left[\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{n}\right]
\end{gathered}
$$

- Full gradient: $\nabla F(\mathbf{w})=X^{T}(X \mathbf{w}-\mathbf{y})+\lambda \mathbf{w}$
- Partial gradient: $\nabla_{i} F(\mathbf{w})=x_{i}^{T}(X \mathbf{w}-\mathbf{y})+\lambda w_{i}$

Proximal mapping
if $i=t_{i}$
otherwise

- $\nabla_{i} F\left(\mathbf{w}^{t}\right)$ can be updated in $O(n)$


## APCG for primal problem

- APCG accelerates PCG using
- Auxiliary sequence ( $\mathbf{v}^{t}$ )
- Momentum step


## APCG

$$
\begin{aligned}
\mathbf{w}_{i}^{t} & = \begin{cases}\arg \min _{w_{i} \in \mathbb{R}} \nabla_{i} F\left(\mathbf{v}^{t-1}\right) w_{i}+\frac{1}{2 \gamma_{t}}\left(w_{i}-\mathbf{v}_{i}^{t-1}\right)^{2}+\tau\left|w_{i}\right| & \text { if } i=t_{i} \\
\mathbf{w}_{i}^{t-1} & \text { otherwise }\end{cases} \\
\mathbf{v}^{t} & =\mathbf{w}^{t}+\eta_{t}\left(\mathbf{w}^{t}-\mathbf{w}^{t-1}\right)
\end{aligned}
$$

## APCG for primal problem

- Per-iteration cost: $O(n)$
- Iteration complexity

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-strongly convex | N.A. |
|  | $\lambda$-strongly convex | N.A. | $O\left(\left(\frac{d}{\sqrt{\lambda}}\right) \log \left(\frac{d}{\epsilon}\right)\right)$ |

- $n \gg d$ : Apply APCG to the dual problem.
- $d \gg n$ : Apply APCG to the primal problem.


## Which Algorithm to Use

- Satisfied with large $\epsilon$ : SGD
- For small $\epsilon$ :

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-strongly convex |  |
|  | $\lambda$-strongly convex | SGD |  |
| APCG | SAGA |  |  |

## Summary

| smooth | str-cvx | SGD | SAGA | SDCA | APCG |
| :---: | :---: | :---: | :---: | :---: | :---: |
| No | No | $\frac{1}{\epsilon^{2}}$ | N.A. | N.A. | N.A. |
| Yes | No | $\frac{1}{\epsilon^{2}}$ | $\frac{n}{\epsilon}$ | N.A. | N.A. |
| No | Yes | $\frac{1}{\lambda \epsilon}$ | N.A. | $n+\frac{1}{\epsilon}$ | $n+\sqrt{\frac{n}{\lambda \epsilon}}$ |
| Yes | Yes | $\frac{1}{\lambda \epsilon}$ | $\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)$ | $\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)$ | $\left(n+\frac{1}{\sqrt{\lambda}}\right) \log \left(\frac{1}{\epsilon}\right)$ |

Table: Per-iteration cost: $O(d)$

## Summary

| smooth | str-cvx | SGD | SAGA | SDCA | APCG |
| :---: | :---: | :---: | :---: | :---: | :---: |
| No | No | $\frac{1}{\epsilon^{2}}$ | N.A. | N.A. | N.A. |
| Yes | No | $\frac{1}{\epsilon^{2}}$ | $\frac{n}{\epsilon}$ | N.A. | N.A. |
| No | Yes | $\frac{1}{\lambda \epsilon}$ | N.A. | $n+\frac{1}{\epsilon}$ | $n+\sqrt{\frac{n}{\lambda \epsilon}}$ |
| Yes | Yes | $\frac{1}{\lambda \epsilon}$ | $\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)$ | $\left(n+\frac{1}{\lambda}\right) \log \left(\frac{1}{\epsilon}\right)$ | $\left(n+\frac{1}{\sqrt{\lambda}}\right) \log \left(\frac{1}{\epsilon}\right)$ |

Table: Per-iteration cost: $O(d)$

| smooth | str-cvx | SVRG |
| :---: | :---: | :---: |
| No | No | N.A. |
| Yes | No | N.A. |
| No | Yes | N.A. |
| Yes | Yes | $\log \left(\frac{1}{\epsilon}\right)$ |

Table: Per-iteration cost: $O\left(d\left(n+1 \frac{1}{2}\right)\right)$

## Summary

|  | SGD | SVRG | SAGA | SDCA | APCG |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | $O(d)$ | $O(d)$ | $O(d n)$ | $O(d)$ | $O(d)$ |
| Parameters | $\gamma_{t}$ | $\gamma_{t}, K$ | $\gamma_{t}$ | None | $\eta_{t}$ |
| absolute |  |  |  |  |  |
| hinge |  |  |  |  |  |
| square |  |  |  |  |  |
| squared hinge |  |  |  |  |  |
| logistic |  |  |  |  |  |
| $\lambda>0$ |  |  |  |  |  |
| $\lambda=0$ |  |  |  |  |  |
| Primal |  |  |  |  |  |
| Dual |  |  |  |  |  |

## Summary

|  | SGD | SVRG | SAGA | SDCA | APCG |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | $O(d)$ | $O(d)$ | $O(d n)$ | $O(d)$ | $O(d)$ |
| Parameters | $\gamma_{t}$ | $\gamma_{t}, K$ | $\gamma_{t}$ | None | $\eta_{t}$ |
| absolute | $\checkmark$ | $\boldsymbol{X}$ | $\boldsymbol{X}$ | $\checkmark$ | $\checkmark$ |
| hinge | $\checkmark$ | $X$ | $X$ | $\checkmark$ | $\checkmark$ |
| square | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| squared hinge | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| logistic | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| $\lambda>0$ |  |  |  |  |  |
| $\lambda=0$ |  |  |  |  |  |
| Primal |  |  |  |  |  |
| Dual |  |  |  |  |  |

## Summary

|  | SGD | SVRG | SAGA | SDCA | APCG |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | $O(d)$ | $O(d)$ | $O(d n)$ | $O(d)$ | $O(d)$ |
| Parameters | $\gamma_{t}$ | $\gamma_{t}, K$ | $\gamma_{t}$ | None | $\eta_{t}$ |
| absolute | $\checkmark$ | $\boldsymbol{X}$ | $\boldsymbol{X}$ | $\checkmark$ | $\checkmark$ |
| hinge | $\checkmark$ | $\boldsymbol{X}$ | $\boldsymbol{X}$ | $\checkmark$ | $\checkmark$ |
| square | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| squared hinge | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| logistic | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| $\lambda>0$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| $\lambda=0$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\mathbf{X}$ | $\mathbf{X}$ |
| Primal |  |  |  |  |  |
| Dual |  |  |  |  |  |

## Summary

|  | SGD | SVRG | SAGA | SDCA | APCG |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | $O(d)$ | $O(d)$ | $O(d n)$ | $O(d)$ | $O(d)$ |
| Parameters | $\gamma_{t}$ | $\gamma_{t}, K$ | $\gamma_{t}$ | None | $\eta_{t}$ |
| absolute | $\checkmark$ | $X$ | $\boldsymbol{X}$ | $\checkmark$ | $\checkmark$ |
| hinge | $\checkmark$ | $\boldsymbol{X}$ | $\boldsymbol{X}$ | $\checkmark$ | $\checkmark$ |
| square | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| squared hinge | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| logistic | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| $\lambda>0$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| $\lambda=0$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\boldsymbol{X}$ | $\mathbf{X}$ |
| Primal | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\mathbf{X}$ | $\checkmark$ |
| Dual | $X$ | $X$ | $\mathbf{X}$ | $\checkmark$ | $\checkmark$ |

## Outline

(2) Optimization

- (Sub)Gradient Methods
- Stochastic Optimization Algorithms for Big Data
- Stochastic Optimization
- Distributed Optimization


## Big Data and Distributed Optimization

## Distributed Optimization

 1010100010101010010010101001001000101100 10101010001001101110101010010111000110101 1010101010001010010010100100101000101010 $10100100010101100 \equiv 101 F \left\lvert\, \begin{aligned} & 010=1010 \\ & 1010\end{aligned}\right.$ 10101000101010010 ص10 10101001010010010 : $001:$ : 11 c : 010011


- data distributed over a cluster of multiple machines
- moving to single machine suffers
- low network bandwidth
- limited disk or memory
- communication V.S. computation
- RAM 100 nanoseconds
- standard network connection 250,000 nanoseconds


## Distributed Data

- $N$ data points are partitioned and distributed to $m$ machines
- $\left[\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots, \mathbf{x}_{N}\right]=S_{1} \cup S_{2} \cup \cdots \cup S_{m}$
- Machine $j$ only has access to $S_{j}$.
- W.L.O.G: $\left|S_{j}\right|=n=\frac{N}{m}$



## A simple solution: Average Solution

- Global problem

$$
\mathbf{w}^{\star}=\underset{\mathbf{w} \in \mathbb{R}^{d}}{\arg \min }\left\{F(\mathbf{w})=\frac{1}{N} \sum_{i=1}^{N} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})\right\}
$$

- Machine $j$ solves a local problem

$$
\mathbf{w}_{j}=\underset{\mathbf{w} \in \mathbb{R}^{d}}{\arg \min }\left\{f_{j}(\mathbf{w})=\frac{1}{n} \sum_{i \in S_{j}} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})\right\}
$$


$W_{1}$

$W_{2}$

$W_{3}$

$W_{4}$

$W_{5}$

$W_{6}$

## A simple solution: Average Solution

- Global problem

$$
\mathbf{w}^{\star}=\underset{\mathbf{w} \in \mathbb{R}^{d}}{\arg \min }\left\{F(\mathbf{w})=\frac{1}{N} \sum_{i=1}^{N} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})\right\}
$$

- Machine $j$ solves a local problem

$$
\mathbf{w}_{j}=\underset{\mathbf{w} \in \mathbb{R}^{d}}{\arg \min }\left\{f_{j}(\mathbf{w})=\frac{1}{n} \sum_{i \in S_{j}} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})\right\}
$$


$\mathbf{w}_{1}$

$\mathbf{w}_{2}$

$\mathbf{w}_{3}$

$\mathbf{W}_{4}$

$\mathbf{w}_{5}$

$\mathbf{w}_{6}$

Center computes: $\widehat{\mathbf{w}}=\frac{1}{m} \sum_{j=1}^{m} \mathbf{w}_{j}$, Issue: Will not converge to $\mathbf{w}^{\star}$

## Mini-Batch SGD: Average Stochastic Gradient

- Machine $j$ sample $i_{t} \in S_{j}$ and construct a stochastic gradient

$$
g_{j}\left(\mathbf{w}_{t-1}\right)=\partial \ell\left(\mathbf{w}_{t-1}^{\top} \mathbf{x}_{i_{t}}, y_{i_{t}}\right)+\partial R\left(\mathbf{w}_{t-1}\right)
$$


$g_{1}\left(\mathbf{w}_{t-1}\right) \quad g_{2}\left(\mathbf{w}_{t-1}\right) \quad g_{3}\left(\mathbf{w}_{t-1}\right) \quad g_{4}\left(\mathbf{w}_{t-1}\right) \quad g_{5}\left(\mathbf{w}_{t-1}\right) \quad g_{6}\left(\mathbf{w}_{t-1}\right)$
Center computes: $\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \underbrace{\frac{1}{m} \sum_{j=1}^{m} g_{j}\left(\mathbf{w}_{t-1}\right)}_{\text {Mini-batch SG }}$

## Total Runtime

Single machine

- Total Runtime
$=$ Per-iteration Cost $\times$ Iteration Complexity
Distributed optimization
- Total Runtime
$=($ Communication Time Per-round + Local Runtime Per-round $)$ $\times$ Rounds of Communication


## Mini-Batch SGD

Applicable in all settings!

- Communication Time Per-round: increase in $m$ in a complicated way.
- Local Runtime Per-round: $O(1)$

Suppose $R(\mathbf{w})$ is $\lambda$-strongly convex

- Rounds of Communication: $O\left(\frac{1}{m} \lambda_{c}\right)$

Suppose $R(\mathbf{w})$ is non-strongly convex

- Rounds of Communication: $O\left(\frac{1}{m^{2}}\right)$

More machines reduce the rounds of communication but increase communication time per-round.

## Mini-Batch SGD

Applicable in all settings!

- Communication Time Per-round: increase in $m$ in a complicated way.
- Local Runtime Per-round: $O(1)$

Suppose $R(\mathbf{w})$ is $\lambda$-strongly convex

- Rounds of Communication: $O\left(\frac{1}{m \lambda \epsilon}\right)$

Suppose $R(\mathbf{w})$ is non-strongly convex

- Rounds of Communication: $O\left(\frac{1}{m \epsilon^{2}}\right)$

More machines reduce the rounds of communication but increase communication time per-round.

## Distributed SDCA (Yang, 2013; Ma et al., 2015)

- Only works when $R(\mathbf{w})=\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}$ with $\lambda>0$. (No $\left.\ell_{1}\right)$
- Global dual problem

$$
\max _{\alpha \in \mathbb{R}^{N}} \frac{1}{N} \sum_{i=1}^{N}-\ell_{i}^{*}\left(\alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda N} \sum_{i=1}^{N} \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
$$

- $\alpha=\left[\alpha_{S_{1}}, \alpha_{S_{2}}, \cdots, \alpha_{S_{m}}\right]$
- Machine $j$ solves a local dual problem only over asj



## Distributed SDCA (Yang, 2013; Ma et al., 2015)

- Only works when $R(\mathbf{w})=\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}$ with $\lambda>0$. (No $\left.\ell_{1}\right)$
- Global dual problem

$$
\max _{\alpha \in \mathbb{R}^{N}} \frac{1}{N} \sum_{i=1}^{N}-\ell_{i}^{*}\left(\alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda N} \sum_{i=1}^{N} \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
$$

- $\alpha=\left[\alpha_{S_{1}}, \alpha_{S_{2}}, \cdots, \alpha_{S_{m}}\right]$
- Machine $j$ solves a local dual problem only over $\alpha_{S_{j}}$

$$
\max _{\alpha s_{j} \in \mathbb{R}^{n}} \frac{1}{N} \sum_{i \in S_{j}}-\ell_{i}^{*}\left(\alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda N} \sum_{i=1}^{N} \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
$$

## DisDCA (Yang, 2013), CoCoA+ (Ma et al., 2015)

- Center maintains a primal solution: $\mathbf{w}^{t}=\frac{1}{\lambda N} \sum_{i=1}^{N} \alpha_{i}^{t} \mathbf{x}_{i}$
- Change variable $\alpha_{i} \longrightarrow \Delta \alpha_{i}$

$$
\begin{aligned}
& \max _{\Delta \alpha S_{j} \in \mathbb{R}^{n}} \frac{1}{N} \sum_{i \in S_{j}}-\ell_{i}^{*}\left(\alpha_{i}^{t}+\Delta \alpha_{i}\right)-\frac{\lambda}{2}\left\|\frac{1}{\lambda N}\left(\sum_{i=1}^{N} \alpha_{i}^{t} \mathbf{x}_{i}+\sum_{i \in S_{j}} \Delta \alpha_{i} \mathbf{x}_{i}\right)\right\|_{2}^{2} \\
\Longleftrightarrow & \max _{\Delta \alpha S_{j} \in \mathbb{R}^{n}} \frac{1}{N} \sum_{i \in S_{j}}-\ell_{i}^{*}\left(\alpha_{i}^{t}+\Delta \alpha_{i}\right)-\frac{\lambda}{2}\left\|\mathbf{w}^{t}+\frac{1}{\lambda N} \sum_{i \in S_{j}} \Delta \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
\end{aligned}
$$

## DisDCA (Yang, 2013), CoCoA+ (Ma et al., 2015)

- Machine $j$ approximately solves

$$
\begin{gathered}
\Delta \alpha_{S_{j}}^{t} \approx \underset{\Delta \alpha_{S_{j}} \in \mathbb{R}^{n}}{\arg \max } \frac{1}{N} \sum_{i \in S_{j}}-\ell_{i}^{*}\left(\alpha_{i}^{t}+\Delta \alpha_{i}\right)-\frac{\lambda}{2}\left\|\mathbf{w}^{t}+\frac{1}{\lambda N} \sum_{i \in S_{j}} \Delta \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2} \\
\alpha_{S_{j}}^{t+1}=\alpha_{S_{j}}^{t}+\Delta \alpha_{S_{j}}^{t}, \quad \Delta \mathbf{w}_{j}^{t}=\frac{1}{\lambda N} \sum_{i \in S_{j}} \Delta \alpha_{S_{j}}^{t} \mathbf{x}_{i}
\end{gathered}
$$



## DisDCA (Yang, 2013), CoCoA+ (Ma et al., 2015)

- Machine $j$ approximately solves



## CoCoA + (Ma et al., 2015)

- Local objective value

$$
\mathcal{G}_{j}\left(\Delta \alpha_{S_{j}}, \mathbf{w}^{t}\right)=\frac{1}{N} \sum_{i \in S_{j}}-\ell_{i}^{*}\left(\alpha_{i}^{t}+\Delta \alpha_{i}\right)-\frac{\lambda}{2}\left\|\mathbf{w}^{t}+\frac{1}{\lambda N} \sum_{i \in S_{j}} \Delta \alpha_{i} \mathbf{x}_{i}\right\|_{2}^{2}
$$

- Solve $\Delta \alpha_{S_{j}}^{t}$ by any local solver as long as

$$
\begin{gathered}
\left(\max _{\Delta \alpha_{S_{j}}} \mathcal{G}_{j}\left(\Delta \alpha_{S_{j}}, \mathbf{w}^{t}\right)-\mathcal{G}_{j}\left(\Delta \alpha_{S_{j}}^{t}, \mathbf{w}^{t}\right)\right) \leq \Theta\left(\max _{\Delta \alpha_{S_{j}}} \mathcal{G}_{j}\left(\Delta \alpha_{S_{j}}, \mathbf{w}^{t}\right)-\mathcal{G}_{j}\left(0, \mathbf{w}^{t}\right)\right) \\
0<\Theta<1
\end{gathered}
$$

## CoCoA+ (Ma et al., 2015)

Suppose $\ell(z)$ is smooth, $R(\mathbf{w})$ is $\lambda$-strongly convex and SDCA is the local solver

- Local Runtime Per-round: $O\left(\left(\frac{1}{\lambda}+\frac{N}{m}\right) \log \left(\frac{1}{\Theta}\right)\right)$
- Rounds of Communication: $O\left(\frac{1}{1-\Theta} \frac{1}{\lambda} \log \left(\frac{1}{\epsilon}\right)\right)$ Suppose $\ell(z)$ is non-smooth, $R(\mathbf{w})$ is $\lambda$-strongly convex and SDCA is the local solver
- Rounds of Communication: $O\left(\frac{1}{1-\Theta} \frac{1}{\lambda \epsilon}\right)$


## CoCoA $+(M a$ et al., 2015)

Suppose $\ell(z)$ is smooth, $R(\mathbf{w})$ is $\lambda$-strongly convex and SDCA is the local solver

- Local Runtime Per-round: $O\left(\left(\frac{1}{\lambda}+\frac{N}{m}\right) \log \left(\frac{1}{\Theta}\right)\right)$
- Rounds of Communication: $O\left(\frac{1}{1-\Theta} \frac{1}{\lambda} \log \left(\frac{1}{\epsilon}\right)\right)$

Suppose $\ell(z)$ is non-smooth, $R(\mathbf{w})$ is $\lambda$-strongly convex and SDCA is the local solver

- Local Runtime Per-round: $O\left(\left(\frac{1}{\lambda}+\frac{N}{m}\right) \frac{1}{\Theta}\right)$
- Rounds of Communication: $O\left(\frac{1}{1-\Theta} \frac{1}{\lambda \epsilon}\right)$


## Distributed SDCA in Practice

- Choice of $\Theta$ (how long we run the local solver?)
- Choice of $m$ (how many machines to use?)
- Fast machines but slow network: Use small $\Theta$ and small $m$
- Fast network but slow machines: Use large $\Theta$ and large $m$


## Distributed SDCA in Practice

- Choice of $\Theta$ (how long we run the local solver?)
- Choice of $m$ (how many machines to use?)
- Fast machines but slow network: Use small $\Theta$ and small $m$
- Fast network but slow machines: Use large $\Theta$ and large $m$


## DiSCO (Zhang \& Xiao, 2015)

- The rounds of communication of Dual SDCA does not depends on $m$.
- DiSCO: Distributed Second-Order method (Zhang \& Xiao, 2015)
- The rounds of communication of DiSCO depends on $m$.


## DiSCO (Zhang \& Xiao, 2015)

- Global problem

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}}\left\{F(\mathbf{w})=\frac{1}{N} \sum_{i=1}^{N} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})\right\}
$$

- Local problem

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}}\left\{f_{j}(\mathbf{w})=\frac{1}{n} \sum_{i \in S_{j}} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+R(\mathbf{w})\right\}
$$

- Global problem can be written as

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}}\left\{F(\mathbf{w})=\frac{1}{m} \sum_{j=1}^{m} f_{j}(\mathbf{w})\right\}
$$

- Applicable when $f_{j}(\mathbf{w})$ is smooth, $\lambda$-strongly convex and self-concordant


## Newton Direction

- At optimal solution $w^{\star}$ :

$$
\nabla F\left(\mathbf{w}^{\star}\right)=\mathbf{0}
$$

- We hope moving $\mathbf{w}^{t}$ along $-\mathbf{v}^{t}$ leads to $\nabla F\left(\mathbf{w}^{t}-\mathbf{v}^{t}\right)=\mathbf{0}$
- Taylor expansion:

$$
F\left(\mathbf{w}^{t}-\mathbf{v}^{t}\right) \approx \nabla F\left(\mathbf{w}^{t}\right)-\nabla^{2} F\left(\mathbf{w}^{t}\right) \mathbf{v}^{t}=\mathbf{0}
$$

- Such $\mathbf{v}^{t}$ is called a Newton's direction


## Newton Method

## Newton Method

Find a Netwon direction $\mathbf{v}_{t}$ by solving

$$
\nabla^{2} F\left(\mathbf{w}_{t}\right) \mathbf{v}_{t}=\nabla F\left(\mathbf{w}_{t}\right)
$$

Then update

$$
\mathbf{w}_{t+1}=\mathbf{w}_{t}-\gamma_{t} \mathbf{v}_{t}
$$

Require solving a linear $d \times d$ equation system. Costly when $d>1000$.

## DiSCO (Zhang \& Xiao, 2015)

## Inexact Newton Method

Find an inexact Netwon direction $v_{t}$ using Preconditioned Conjugate Gradient (PCG) (Golub \& Ye, 1997)

$$
\left\|\nabla^{2} F\left(\mathbf{w}_{t}\right) \mathbf{v}_{t}-\nabla F\left(\mathbf{w}_{t}\right)\right\|_{2} \leq \epsilon_{t}
$$

Then update

$$
\mathbf{w}_{t+1}=\mathbf{w}_{t}-\gamma_{t} \mathbf{v}_{t}
$$

## DiSCO (Zhang \& Xiao, 2015)

## PCG <br> Keep computing

$$
\mathbf{v} \longleftarrow P^{-1} \times \nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v}
$$

iteratively until v becomes an inexact Newton direction.

- Preconditioner: $P=\nabla^{2} f_{1}\left(\mathbf{w}_{t}\right)+\mu l$
- $\mu$ : a tuning parameter such that

$$
\left\|\nabla^{2} f_{1}\left(\mathbf{w}_{t}\right)-\nabla^{2} F\left(\mathbf{w}_{t}\right)\right\|_{2} \leq \mu
$$

- $f_{1}$ is "similar" to $F$ so that $P$ is a good local preconditioner.



## DiSCO (Zhang \& Xiao, 2015)

## PCG

Keep computing

$$
\mathbf{v} \longleftarrow P^{-1} \times \nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v}
$$

iteratively until v becomes an inexact Newton direction.

- Preconditioner: $P=\nabla^{2} f_{1}\left(\mathbf{w}_{t}\right)+\mu l$
- $\mu$ : a tuning parameter such that

$$
\left\|\nabla^{2} f_{1}\left(\mathbf{w}_{t}\right)-\nabla^{2} F\left(\mathbf{w}_{t}\right)\right\|_{2} \leq \mu
$$

- $f_{1}$ is "similar" to $F$ so that $P$ is a good local preconditioner.
- $\mu=O\left(\frac{1}{\sqrt{n}}\right)=O\left(\sqrt{\frac{m}{N}}\right)$


## DiSCO (Zhang \& Xiao, 2015)

DiSCO compute $\nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v}$ distributedly:

$$
\begin{aligned}
& \begin{array}{l}
\nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v} \\
= \\
\nabla^{2} f_{1}\left(\mathbf{w}_{t}\right) \times \mathbf{v}
\end{array} \text { machine } 1+\underbrace{\nabla^{2} f_{2}\left(\mathbf{w}_{t}\right) \times \mathbf{v}} \text { machine } 2+\cdots+\underbrace{\nabla^{2} f_{m}\left(\mathbf{w}_{t}\right)}
\end{aligned}
$$

Then, compute $P^{-1} \times \nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v}$ only in machine 1

## DiSCO (Zhang \& Xiao, 2015)



## DiSCO (Zhang \& Xiao, 2015)



## DiSCO (Zhang \& Xiao, 2015)



## DiSCO (Zhang \& Xiao, 2015)



Machine 1: $\quad \mathrm{P}^{-1}\left(\nabla^{2} F\left(w_{t}\right) v_{t}\right)=P^{-1}\left(\frac{1}{\mathrm{~m}} \sum_{j=1} \nabla^{2} f_{j}\left(w_{t}\right) v_{t}\right)$
Run PCG until $\left\|\nabla^{2} F\left(w_{t}\right) v_{t}-\nabla F\left(w_{t}\right)\right\|_{2} \leq \epsilon_{t}$

## DiSCO (Zhang \& Xiao, 2015)

- For high-dimensional data (e.g. $d \geq 1000$ ), computing $P^{-1} \times \nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v}$ is time costly.
- Instead, use SDCA in machine 1 to solve

$$
P^{-1} \times \nabla^{2} F\left(\mathbf{w}_{t}\right) \times \mathbf{v} \approx \underset{\mathbf{u} \in \mathbb{R}^{d}}{\arg \min } \frac{1}{2} \mathbf{u}^{\top} P \mathbf{u}-\mathbf{u}^{\top} F\left(\mathbf{w}_{t}\right) \mathbf{v}
$$

- Local runtime: $O\left(\frac{N}{m}+\frac{1+\mu}{\lambda+\mu}\right)$


## DiSCO (Zhang \& Xiao, 2015)

Suppose SDCA is the local solver

- Local Runtime Per-round: $O\left(\frac{N}{m}+\frac{1+\mu}{\lambda+\mu}\right)$
- Rounds of Communication: $O\left(\sqrt{\frac{\mu}{\lambda}} \log \left(\frac{1}{\epsilon}\right)\right)$

Choice of $m$ (how many machines to use?). $\left(\mu=O\left(\sqrt{ } \frac{m}{N}\right)\right)$

- Fast machines but slow network: Use small $m$
- Fast network but slow machine: Use large $m$


## DiSCO (Zhang \& Xiao, 2015)

Suppose SDCA is the local solver

- Local Runtime Per-round: $O\left(\frac{N}{m}+\frac{1+\mu}{\lambda+\mu}\right)$
- Rounds of Communication: $O\left(\sqrt{\frac{\mu}{\lambda}} \log \left(\frac{1}{\epsilon}\right)\right)$

Choice of $m$ (how many machines to use?). $\left(\mu=O\left(\sqrt{\frac{m}{N}}\right)\right)$

- Fast machines but slow network: Use small $m$
- Fast network but slow machine: Use large $m$


## DSVRG (Lee et al., 2015)

- A distributed version of SVRG using a "round-robin" scheme
- Assume the user can control the distribution of data before algorithm.
- Applicable when $f_{j}(\mathbf{w})$ is smooth and $\lambda$-strongly convex


## DSVRG (Lee et al., 2015)

```
Iterate \(s=1, \ldots, T-1\)
    Let \(\mathbf{w}_{0}=\tilde{\mathbf{w}}_{s}\) and compute \(\nabla F\left(\tilde{\mathbf{w}}_{s}\right)\)
    Iterate \(t=1, \ldots, K\)
        \(\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right)\)
        \(\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{\underline{w}}_{i_{t}}\left(\mathbf{w}_{t-1}\right)\)
    \(\tilde{\mathbf{w}}_{s+1}=\frac{1}{K} \sum_{t=1}^{K} \mathbf{w}_{t}\)
output: \(\tilde{\mathbf{w}}_{T}\)
```

- Each machine can only sample from its own data. However,
- $\mathbb{E}_{i_{t} \in S_{j}}\left[\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)\right] \neq \nabla F\left(\mathbf{w}_{t-1}\right)$


## DSVRG (Lee et al., 2015)

$$
\begin{aligned}
& \text { Iterate } s=1, \ldots, T-1 \\
& \text { Let } \mathbf{w}_{0}=\tilde{\mathbf{w}}_{s} \text { and compute } \nabla F\left(\tilde{\mathbf{w}}_{s}\right) \\
& \quad \text { Iterate } t=1, \ldots, K \\
& \quad \tilde{g}_{i t}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right) \\
& \quad \mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right) \\
& \text { Easy to } \\
& \text { distribute }
\end{aligned}
$$

- Each machine can only sample from its own data. However,
- $\mathbb{E}_{i_{t} \in S_{j}}\left[\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)\right] \neq \nabla F\left(\mathbf{w}_{t-1}\right)$


## DSVRG (Lee et al., 2015)

```
Iterate \(s=1, \ldots, T-1\)
    Let \(\mathbf{w}_{0}=\tilde{\mathbf{w}}_{s}\) and compute \(\nabla F\left(\tilde{\mathbf{w}}_{s}\right)\)
    Iterate \(t=1, \ldots, K\)
        \(\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right)\)
        \(\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{\underline{w}}_{i_{t}}\left(\mathbf{w}_{t-1}\right)\)
    \(\tilde{\mathbf{w}}_{s+1}=\frac{1}{K} \sum_{t=1}^{K} \mathbf{w}_{t}\)
output: \(\tilde{\mathbf{w}}_{T}\)
```

- Each machine can only sample from its own data. However,
$\square$


## DSVRG (Lee et al., 2015)

```
Iterate \(s=1, \ldots, T-1\)
    Let \(\mathbf{w}_{0}=\tilde{\mathbf{w}}_{s}\) and compute \(\nabla F\left(\tilde{\mathbf{w}}_{s}\right)\)
    Iterate \(t=1, \ldots, K\)
        \(\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right)\)
        \(\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)\)
    \(\tilde{\mathbf{w}}_{s+1}=\frac{1}{K} \sum_{t=1}^{K} \mathbf{w}_{t}\)
output: \(\tilde{\mathbf{w}}_{T}\)
```

- Each machine can only sample from its own data. However,
- $\mathbb{E}_{i_{t} \in S_{j}}\left[\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)\right] \neq \nabla F\left(\mathbf{w}_{t-1}\right)$


## DSVRG (Lee et al., 2015)

Solution:

- Store a second set of data $R_{j}$ in machine $j$, which are sampled with replacement from $\left\{x_{1}, x_{2}, \ldots, x_{n}\right\}$ before the algorithm starts.
- Construct the stochastic gradient $\tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)$ by sampling $i_{t} \in R_{j}$ and removing $i_{t}$ from $R_{j}$ after.

$$
\mathbb{E}_{i_{t} \in R_{j}}\left[\tilde{g}_{i t}\left(\mathbf{w}_{t-1}\right)\right]=\nabla F\left(\mathbf{w}_{t-1}\right)
$$

- When $R_{j}=\emptyset$, pass $\mathbf{w}_{t}$ to next machine.


## Full Gradient Step



Center: $\quad \nabla F\left(\widetilde{w}_{s}\right)=\frac{1}{\mathrm{~m}} \sum_{j=1} \nabla f_{j}\left(\widetilde{w}_{s}\right)$

## Full Gradient Step



## Stochastic Gradient Step

Iterate $t=1, \ldots, m$

$$
\begin{aligned}
& \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right) \\
& \mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)
\end{aligned}
$$



## Stochastic Gradient Step

$$
\begin{aligned}
& \text { Iterate } t=1, \ldots, m \\
& \qquad \begin{array}{l}
\tilde{g}_{j_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right) \\
\mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)
\end{array} \quad \begin{array}{l}
\quad \mathbf{w}_{t}
\end{array}
\end{aligned}
$$



## Stochastic Gradient Step

$$
\begin{aligned}
& \text { Iterate } t=1, \ldots, m \\
& \quad \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)=\nabla F\left(\tilde{\mathbf{w}}_{s}\right)-g_{i_{t}}\left(\tilde{\mathbf{w}}_{s}\right)+g_{i_{t}}\left(\mathbf{w}_{t-1}\right) \\
& \mathbf{w}_{t}=\mathbf{w}_{t-1}-\gamma_{t} \tilde{g}_{i_{t}}\left(\mathbf{w}_{t-1}\right)
\end{aligned}
$$



## DSVRG (Lee et al., 2015)

Suppose $\left|R_{j}\right|=r$ for all $j$.

- Local Runtime Per-round: $O\left(\frac{N}{m}+r\right)$
- Rounds of Communication: $O\left(\frac{1}{r \lambda} \log \left(\frac{1}{\epsilon}\right)\right)$


## DSVRG (Lee et al., 2015)

Choice of $m$ (how many machines to use?).

- Fast machines but slow network: Use small $m$
- Fast network but slow machine: Use large $m$

Choice of $r$ (how many data points to pre-sample in $R_{j}$ ?).

- The larger, the better
- Required machine memory space: $\left|S_{j}\right|+\left|R_{j}\right|=\frac{N}{m}+r$


## Other Distributed Optimization Methods

- ADMM (Boyd et al., 2011; Ozdaglar, 2015)
- Rounds of Communication: $O\left(\right.$ Network Graph Dependency Term $\left.\times \frac{1}{\sqrt{\lambda}} \log \left(\frac{1}{\epsilon}\right)\right)$
- DANE (Shamir et al., 2014)
- Approximate Newton direction with a difference approach from DISCO


## Summary

$\ell(z)$ is smooth and $R(\mathbf{w})$ is $\lambda$-strongly convex.

| alg. | Mini-SGD | Dist-SDCA | DiSCO |
| :--- | :---: | :---: | :---: |
| Runtime Per-round | $O(1)$ | $O\left(\frac{1}{\lambda}+\frac{N}{m}\right)$ | $O\left(\frac{N}{m}\right)$ |
| Round of Comm. | $O\left(\frac{1}{\lambda m \epsilon}\right)$ | $O\left(\frac{1}{\lambda} \log \left(\frac{1}{\epsilon}\right)\right)$ | $O\left(\frac{m^{1 / 4}}{\sqrt{\lambda} N^{1 / 4}} \log \left(\frac{1}{\epsilon}\right)\right)$ |

Table: Assume $\mu=O\left(\sqrt{\frac{m}{N}}\right)$

| alg. | DSVRG |  |
| :--- | :---: | :---: |
|  | Full Grad. | Stoch. Grad. |
| Runtime Per-round | $O\left(\frac{N}{m}\right)$ | $O(r)$ |
| Round of Comm. | $O\left(\log \left(\frac{1}{\epsilon}\right)\right)$ | $O\left(\frac{1}{r \lambda} \log \left(\frac{1}{\epsilon}\right)\right)$ |

## Summary

$\ell(z)$ is non-smooth and $g(\mathbf{w})$ is $\lambda$-strongly convex.

| alg. | Mini-SGD | Dist-SDCA | DiSCO | DSVRG |
| :--- | :---: | :---: | :---: | :---: |
| Runtime Per-round | $O(1)$ | $O\left(\frac{1}{\lambda}+\frac{N}{m}\right)$ | N.A. | N.A. |
| Round of Comm. | $O\left(\frac{1}{\lambda m \epsilon}\right)$ | $O\left(\frac{1}{\lambda \epsilon}\right)$ | N.A. | N.A. |

## Summary

$g(\mathbf{w})$ is non-strongly convex.

| alg. | Mini-SGD | Dist-SDCA | DiSCO | DSVRG |
| :--- | :---: | :---: | :---: | :---: |
| Runtime Per-round | $O(1)$ | N.A. | N.A. | N.A. |
| Round of Comm. | $O\left(\frac{1}{m \epsilon^{2}}\right)$ | N.A. | N.A. | N.A. |

## Which Algorithm to Use

- Algorithm to use

|  |  | $\ell(z) \equiv \ell(z, y)$ |  |
| :---: | :---: | :---: | :---: |
| $R(\mathbf{w})$ |  | Non-strongly convex | Mini-SGD |
|  | $\lambda$-strongly convex | Dist-SDCA | DiSCO/DSVRG |

- Between DiSCO/DSVRG
- Use DiSCO for small $\lambda$, e.g., $\lambda<10^{-5}$
- Use DSVRG for large $\lambda$, e.g., $\lambda>10^{-5}$


## Distributed Machine Learning Systems and Library

- Petuum: http://petuum.github.io
- Apache Spark: http://spark.apache.org/
- Parameter Server: http://parameterserver.org/
- Birds: http://cs.uiowa.edu/~tyng/software.html


## Thank You! Questions?

# Big Data Analytics: Optimization and Randomization Part III: Randomization 

## Outline

(1) Basics
(2) Optimization
(3) Randomized Dimension Reduction
(4) Randomized Algorithms
(5) Concluding Remarks

## Random Sketch

Approximate a large data matrix
by a much smaller sketch


## The Framework of Randomized Algorithms



## The Framework of Randomized Algorithms



## The Framework of Randomized Algorithms



## The Framework of Randomized Algorithms



## Why randomized dimension reduction?

- Efficient
- Robust (e.g., dropout)
- Formal Guarantees
- Can explore parallel algorithms


## Randomized Dimension Reduction

- Johnson-Lindenstauss (JL) transforms
- Subspace embeddings
- Column sampling


## JL Lemma

JL Lemma (Johnson \& Lindenstrauss, 1984)
For any $0<\epsilon, \delta<1 / 2$, there exists a probability distribution on $m \times d$ real matrices $A$ such that there exists a small universal constant $c>0$ and for any fixed $\mathrm{x} \in \mathbb{R}^{d}$ with a probability at least $1-\delta$, we have

$$
\left|\|A \mathbf{x}\|_{2}^{2}-\|\mathbf{x}\|_{2}^{2}\right| \leq c \sqrt{\frac{\log (1 / \delta)}{m}}\|\mathbf{x}\|_{2}^{2}
$$

or for $m=\Theta\left(\epsilon^{-2} \log (1 / \delta)\right)$, then with a probability at least $1-\delta$

$$
\left|\|A \mathbf{x}\|_{2}^{2}-\|\mathbf{x}\|_{2}^{2}\right| \leq \epsilon\|\mathbf{x}\|_{2}^{2}
$$

## Embedding a set of points into low dimensional space

Given a set of points $\mathbf{x}_{1}, \ldots, \mathbf{x}_{n} \in \mathbb{R}^{d}$, we can embed them into a low dimensional space $A \mathbf{x}_{1}, \ldots, A \mathbf{x}_{n} \in \mathbb{R}^{m}$ such that the pairwise distance between any two points are well preserved in the low dimensional space

$$
\begin{aligned}
\left\|A \mathbf{x}_{i}-A \mathbf{x}_{j}\right\|_{2}^{2} & =\left\|A\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)\right\|_{2}^{2} \leq(1+\epsilon)\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\|_{2}^{2} \\
\left\|A \mathbf{x}_{i}-A \mathbf{x}_{j}\right\|_{2}^{2} & =\left\|A\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)\right\|_{2}^{2} \geq(1-\epsilon)\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\|_{2}^{2}
\end{aligned}
$$

In other words, in order to have all pairwise Euclidean distances preserved up to $1 \pm \epsilon$, only $m=\Theta\left(\epsilon^{-2} \log \left(n^{2} / \delta\right)\right)$ dimensions are necessary

## JL transforms: Gaussian Random Projection

Gaussian Random Projection (Dasgupta \& Gupta, 2003): $A \in \mathbb{R}^{m \times d}$

- $A_{i j} \sim \mathcal{N}(0,1 / m)$
- $m=\Theta\left(\epsilon^{-2} \log (1 / \delta)\right)$
- Computational cost of $A X$ : where $X \in \mathbb{R}^{d \times n}$
- mnd for dense matrices
- $\mathrm{nnz}(X) m$ for sparse matrices

Computational Cost is very High (could be as high as solving many problems)

## Accelerate JL transforms: using discrete distributions

Using Discrete Distributions (Achlioptas, 2003):

- $\operatorname{Pr}\left(A_{i j}= \pm \frac{1}{\sqrt{m}}\right)=0.5$
- $\operatorname{Pr}\left(A_{i j}= \pm \sqrt{\frac{3}{m}}\right)=\frac{1}{6}, \operatorname{Pr}\left(A_{i j}=0\right)=\frac{2}{3}$
- Database friendly
- Replace multiplications by additions and subtractions


## Accelerate JL transforms: using Hadmard transform (I)

Fast JL transform based on randomized Hadmard transform:

Motivation: Can we simply use random sampling matrix $P \in \mathbb{R}^{m \times d}$ that randomly selects $m$ coordinates out of $d$ coordinates (scaled by $\sqrt{d / m}$ )?


Unless $\frac{\sqrt{d}\|\mathrm{x}\|_{\infty}}{\|\mathrm{x}\|_{2}} \leq c$, the random sampling doest not work

Remedy is given by randomized Hadmard transform

## Accelerate JL transforms: using Hadmard transform (I)

Fast JL transform based on randomized Hadmard transform:

Motivation: Can we simply use random sampling matrix $P \in \mathbb{R}^{m \times d}$ that randomly selects $m$ coordinates out of $d$ coordinates (scaled by $\sqrt{d / m}$ )?

Unfortunately: by Chernoff bound

$$
\left|\|P \mathbf{x}\|_{2}^{2}-\|\mathbf{x}\|_{2}^{2}\right| \leq \frac{\sqrt{d}\|\mathbf{x}\|_{\infty}}{\|\mathbf{x}\|_{2}} \sqrt{\frac{3 \log (2 / \delta)}{m}}\|\mathbf{x}\|_{2}^{2}
$$

Unless $\frac{\sqrt{d}\|\mathbf{x}\|_{\infty}}{\|\mathbf{x}\|_{2}} \leq c$, the random sampling doest not work

## Accelerate JL transforms: using Hadmard transform (I)

Fast JL transform based on randomized Hadmard transform:

Motivation: Can we simply use random sampling matrix $P \in \mathbb{R}^{m \times d}$ that randomly selects $m$ coordinates out of $d$ coordinates (scaled by $\sqrt{d / m}$ )?

Unfortunately: by Chernoff bound

$$
\left|\|P \mathbf{x}\|_{2}^{2}-\|\mathbf{x}\|_{2}^{2}\right| \leq \frac{\sqrt{d}\|\mathbf{x}\|_{\infty}}{\|\mathbf{x}\|_{2}} \sqrt{\frac{3 \log (2 / \delta)}{m}}\|\mathbf{x}\|_{2}^{2}
$$

Unless $\frac{\sqrt{d}\|\mathbf{x}\|_{\infty}}{\|\mathbf{x}\|_{2}} \leq c$, the random sampling doest not work
Remedy is given by randomized Hadmard transform

## Randomized Hadmard transform

Hadmard transform:

- $H \in \mathbb{R}^{d \times d}: H=\sqrt{\frac{1}{d}} H_{2^{k}}$

$$
H_{1}=[1], \quad H_{2}=\left[\begin{array}{cc}
1 & 1 \\
1 & -1
\end{array}\right], \quad H_{2^{k}}=\left[\begin{array}{cc}
H_{2^{k-1}} & H_{2^{k-1}} \\
H_{2^{k-1}} & -H_{2^{k-1}}
\end{array}\right]
$$

- $\|H \mathbf{x}\|_{2}=\|\mathbf{x}\|_{2}$ and $H$ is orthogonal
- Computational costs of $H x: d \log (d)$
randomized Hadmard transform: HD
- $D \in \mathbb{R}^{d \times d}$ : a diagonal matrix $\operatorname{Pr}\left(D_{i i}= \pm 1\right)=0.5$
- $H D$ is orthogonal and $\|H D \mathbf{x}\|_{2}=\|\mathbf{x}\|_{2}$


## Randomized Hadmard transform

Hadmard transform:

- $H \in \mathbb{R}^{d \times d}: H=\sqrt{\frac{1}{d}} H_{2^{k}}$

$$
H_{1}=[1], \quad H_{2}=\left[\begin{array}{cc}
1 & 1 \\
1 & -1
\end{array}\right], \quad H_{2^{k}}=\left[\begin{array}{cc}
H_{2^{k-1}} & H_{2^{k-1}} \\
H_{2^{k-1}} & -H_{2^{k-1}}
\end{array}\right]
$$

- $\|H \mathbf{x}\|_{2}=\|\mathbf{x}\|_{2}$ and $H$ is orthogonal
- Computational costs of $H x: d \log (d)$ randomized Hadmard transform: $H D$
- $D \in \mathbb{R}^{d \times d}$ : a diagonal matrix $\operatorname{Pr}\left(D_{i i}= \pm 1\right)=0.5$
- $H D$ is orthogonal and $\|H D \mathbf{x}\|_{2}=\|\mathbf{x}\|_{2}$


## Randomized Hadmard transform

Hadmard transform:

- $H \in \mathbb{R}^{d \times d}: H=\sqrt{\frac{1}{d}} H_{2^{k}}$

$$
H_{1}=[1], \quad H_{2}=\left[\begin{array}{cc}
1 & 1 \\
1 & -1
\end{array}\right], \quad H_{2^{k}}=\left[\begin{array}{cc}
H_{2^{k-1}} & H_{2^{k-1}} \\
H_{2^{k-1}} & -H_{2^{k-1}}
\end{array}\right]
$$

- $\|H \mathbf{x}\|_{2}=\|\mathbf{x}\|_{2}$ and $H$ is orthogonal
- Computational costs of $H x: d \log (d)$ randomized Hadmard transform: $H D$
- $D \in \mathbb{R}^{d \times d}$ : a diagonal matrix $\operatorname{Pr}\left(D_{i i}= \pm 1\right)=0.5$
- $H D$ is orthogonal and $\|H D \mathbf{x}\|_{2}=\|\mathbf{x}\|_{2}$

$$
\text { Key property: } \frac{\sqrt{d}\|H D \mathbf{x}\|_{\infty}}{\|H D \mathbf{x}\|_{2}} \leq \sqrt{\log (d / \delta)} \text { w.h.p } 1-\delta
$$

## Accelerate JL transforms: using Hadmard transform (I)

Fast JL transform based on randomized Hadmard transform (Tropp, 2011):

$$
A=\sqrt{\frac{d}{m}} P H D
$$

yields

$$
\left|\|A x\|_{2}^{2}-\|\mathbf{x}\|_{2}^{2}\right| \leq \sqrt{\frac{3 \log (2 / \delta) \log (d / \delta)}{m}}\|\mathbf{x}\|_{2}^{2}
$$

- $m=\Theta\left(\epsilon^{-2} \log (1 / \delta) \log (d / \delta)\right)$ suffice for $1 \pm \epsilon$
- additional factor $\log (d / \delta)$ can be removed
- Computational cost of $A X: O(n d \log (m))$


## Accelerate JL transforms: using a sparse matrix (I)

Random hashing (Dasgupta et al., 2010)

$$
A=H D
$$

where $D \in \mathbb{R}^{d \times d}$ and $H \in \mathbb{R}^{m \times d}$

- random hashing: $h(j):\{1, \ldots, d\} \rightarrow\{1, \ldots, m\}$
- $H_{i j}=1$ if $h(j)=i$ : sparse matrix (each column has only one non-zero entry)
- $D \in \mathbb{R}^{d \times d}$ : a diagonal matrix $\operatorname{Pr}\left(D_{i i}= \pm 1\right)=0.5$
- $[A \mathbf{x}]_{j}=\sum_{i: h(i)=j} x_{i} D_{i i}$


## Accelerate JL transforms: using a sparse matrix (I)

Random hashing (Dasgupta et al., 2010)

$$
A=H D
$$

where $D \in \mathbb{R}^{d \times d}$ and $H \in \mathbb{R}^{m \times d}$

- random hashing: $h(j):\{1, \ldots, d\} \rightarrow\{1, \ldots, m\}$
- $H_{i j}=1$ if $h(j)=i$ : sparse matrix (each column has only one non-zero entry)
- $D \in \mathbb{R}^{d \times d}$ : a diagonal matrix $\operatorname{Pr}\left(D_{i i}= \pm 1\right)=0.5$
- $[A \mathbf{x}]_{j}=\sum_{i: h(i)=j} x_{i} D_{i i}$

Technically speaking, random hashing does not satisfy JL lemma

## Accelerate JL transforms: using a sparse matrix (I)

- key properties:
- $\mathrm{E}\left[\left\langle H D \mathrm{x}_{1}, H D \mathrm{x}_{2}\right\rangle\right]=\left\langle\mathbf{x}_{1}, \mathbf{x}_{2}\right\rangle$
- and norm perserving $\left|\|H D \mathbf{x}\|_{2}^{2}-\|\mathbf{x}\|_{2}^{2}\right| \leq \epsilon\|\mathbf{x}\|_{2}^{2}$, only when

$$
\frac{\|\mathbf{x}\|_{\infty}}{\|\mathbf{x}\|_{2}} \leq \frac{1}{\sqrt{c}}
$$

Apply randomized Hadmard transform $P$ first: $\Theta(c \log (c / \delta))$ blocks of randomized Hadmard transform

$$
\frac{\|P \mathbf{x}\|_{\infty}}{\|P \mathbf{x}\|_{2}} \leq \frac{1}{\sqrt{c}}
$$

## Accelerate JL transforms: using a sparse matrix (II)

Sparse JL transform based on block random hashing (Kane \& Nelson, 2014)

$$
A=\left[\begin{array}{c}
\frac{1}{\sqrt{s}} Q_{1} \\
\cdots \\
\frac{1}{\sqrt{s}} Q_{s}
\end{array}\right]
$$

- Each $Q_{s} \in \mathbb{R}^{v \times d}$ is an independent random hashing (HD) matrix
- Set $v=\Theta\left(\epsilon^{-1}\right)$ and $s=\Theta\left(\epsilon^{-1} \log (1 / \delta)\right)$
- Computational Cost of $A X: O\left(\frac{n n z(X)}{\epsilon} \log \left[\frac{1}{\delta}\right]\right)$


## Randomized Dimension Reduction

- Johnson-Lindenstauss (JL) transforms
- Subspace embeddings
- Column sampling


## Subspace Embeddings

Definition: a subspace embedding given some parameters $0<\epsilon, \delta<1, k \leq d$ is a distribution $\mathcal{D}$ over matrices $A \in \mathbb{R}^{m \times d}$ such that for any fixed linear subspace $W \in \mathbb{R}^{d}$ with $\operatorname{dim}(W)=k$ it holds that

$$
\operatorname{Pr}_{A \sim \mathcal{D}}\left(\forall \mathbf{x} \in W,\|A \mathbf{x}\|_{2} \in(1 \pm \epsilon)\|\mathbf{x}\|_{2}\right) \geq 1-\delta
$$

- These are key properties in the theoretical analysis of many algorithms (e.g., low-rank matrix approximation, randomized least-squares regression, randomized classification)


## Subspace Embeddings

Definition: a subspace embedding given some parameters $0<\epsilon, \delta<1, k \leq d$ is a distribution $\mathcal{D}$ over matrices $A \in \mathbb{R}^{m \times d}$ such that for any fixed linear subspace $W \in \mathbb{R}^{d}$ with $\operatorname{dim}(W)=k$ it holds that

$$
\operatorname{Pr}_{A \sim \mathcal{D}}\left(\forall \mathbf{x} \in W,\|A \mathbf{x}\|_{2} \in(1 \pm \epsilon)\|\mathbf{x}\|_{2}\right) \geq 1-\delta
$$

It implies

- If $U \in \mathbb{R}^{d \times k}$ is orthogonal matrix (contains the orthonormal bases)
- $A U \in \mathbb{R}^{m \times k}$ is of full column rank
- $\|A U\|_{2} \in(1 \pm \epsilon)$
- $(1-\epsilon)^{2} \leq\left\|U^{\top} A^{\top} A U\right\|_{2} \leq(1+\epsilon)^{2}$
$\square$
- These are key properties in the theoretical analysis of many algorithms (e.g., low-rank matrix approximation, randomized least-squares regression, randomized classification)


## Subspace Embeddings

Definition: a subspace embedding given some parameters $0<\epsilon, \delta<1, k \leq d$ is a distribution $\mathcal{D}$ over matrices $A \in \mathbb{R}^{m \times d}$ such that for any fixed linear subspace $W \in \mathbb{R}^{d}$ with $\operatorname{dim}(W)=k$ it holds that

$$
\operatorname{Pr}_{A \sim \mathcal{D}}\left(\forall \mathbf{x} \in W,\|A \mathbf{x}\|_{2} \in(1 \pm \epsilon)\|\mathbf{x}\|_{2}\right) \geq 1-\delta
$$

It implies

- If $U \in \mathbb{R}^{d \times k}$ is orthogonal matrix (contains the orthonormal bases)
- $A U \in \mathbb{R}^{m \times k}$ is of full column rank
- $\|A U\|_{2} \in(1 \pm \epsilon)$
- $(1-\epsilon)^{2} \leq\left\|U^{\top} A^{\top} A U\right\|_{2} \leq(1+\epsilon)^{2}$
- These are key properties in the theoretical analysis of many algorithms (e.g., low-rank matrix approximation, randomized least-squares regression, randomized classification)


## Subspace Embeddings

From a JL transform to a Subspace Embedding (Sarlós, 2006). Let $A \in \mathbb{R}^{m \times d}$ be a JL transform. If

$$
m=O\left(\frac{k \log \left[\frac{k}{\delta \epsilon}\right]}{\epsilon^{2}}\right)
$$

Then w.h.p $1-\delta^{k}, A \in \mathbb{R}^{m \times d}$ is a subspace embedding w.r.t a $k$-dimensional space in $\mathbb{R}^{d}$

## Subspace Embeddings

Making block random hashing a Subspace Embedding (Nelson \& Nguyen, 2013).

$$
A=\left[\begin{array}{c}
\frac{1}{\sqrt{s}} Q_{1} \\
\cdots \\
\frac{1}{\sqrt{s}} Q_{s}
\end{array}\right]
$$

- Each $Q_{s} \in \mathbb{R}^{v \times d}$ is an independent random hashing (HD) matrix
- Set $v=\Theta\left(k \epsilon^{-1} \log ^{5}(k / \delta)\right)$ and $s=\Theta\left(\epsilon^{-1} \log ^{3}(k / \delta)\right)$
- w.h.p $1-\delta, A \in \mathbb{R}^{m \times d}$ with $m=\Theta\left(\frac{k \log ^{8}(k / \delta)}{\epsilon^{2}}\right)$ is a subspace embedding w.r.t a $k$-dimensional space in $\mathbb{R}^{d}$
- Computational Cost of $A X: O\left(\frac{n n z(X)}{\epsilon} \log ^{3}\left[\frac{k}{\delta}\right]\right)$


## Sparse Subspace Embedding (SSE)

Random hashing is SSE with a Constant Probability (Nelson \& Nguyen, 2013)

$$
A=H D
$$

where $D \in \mathbb{R}^{d \times d}$ and $H \in \mathbb{R}^{m \times d}$

- $m=\Omega\left(k^{2} / \epsilon^{2}\right)$ suffice for a subspace embedding with a probability $2 / 3$
- Computational Cost $A X: O(n n z(X))$


## Randomized Dimensionality Reduction

- Johnson-Lindenstauss (JL) transforms
- Subspace embeddings
- Column (Row) sampling


## Column sampling

- Column subset selection (feature selection)
- More interpretable
- Uniform sampling usually does not work (not a JL transform)
- Non-oblivious sampling (data-dependent sampling)
- leverage-score sampling


## Column sampling

- Column subset selection (feature selection)
- More interpretable
- Uniform sampling usually does not work (not a JL transform)
- leverage-score sampling


## Column sampling

- Column subset selection (feature selection)
- More interpretable
- Uniform sampling usually does not work (not a JL transform)
- Non-oblivious sampling (data-dependent sampling)
- leverage-score sampling


## Leverage-score sampling (Drineas et al., 2006)

Let $X \in \mathbb{R}^{d \times n}$ be a rank- $k$ matrix

- $X=U \Sigma V^{\top}: U \in \mathbb{R}^{d \times k}, \Sigma \in \mathbb{R}^{k \times k}$
- Leverage scores $\left\|U_{i *}\right\|_{2}^{2}, i=1, \ldots, d$

- Let $i_{1}, \ldots, i_{m} \in\{1, \ldots, d\}$ denote $m$ indices selected by following $p_{i}$ - Let $A \in \mathbb{R}^{m \times d}$ be sampling-and-rescaling matrix:

- $A X \in \mathbb{R}^{m \times n}$ is a small sketch of $X$


## Leverage-score sampling (Drineas et al., 2006)

Let $X \in \mathbb{R}^{d \times n}$ be a rank- $k$ matrix

- $X=U \Sigma V^{\top}: U \in \mathbb{R}^{d \times k}, \Sigma \in \mathbb{R}^{k \times k}$
- Leverage scores $\left\|U_{i *}\right\|_{2}^{2}, i=1, \ldots, d$

- Let $i_{1}, \ldots, i_{m} \in\{1, \ldots, d\}$ denote $m$ indices selected by following $p_{i}$ - Let $A \in \mathbb{R}^{m \times d}$ be sampling-and-rescaling matrix:

- $A X \in \mathbb{R}^{m \times n}$ is a small sketch of $X$


## Leverage-score sampling (Drineas et al., 2006)

Let $X \in \mathbb{R}^{d \times n}$ be a rank- $k$ matrix

- $X=U \Sigma V^{\top}: U \in \mathbb{R}^{d \times k}, \Sigma \in \mathbb{R}^{k \times k}$
- Leverage scores $\left\|U_{i *}\right\|_{2}^{2}, i=1, \ldots, d$
- Let $p_{i}=\frac{\left\|U_{i *}\right\|_{2}^{2}}{\sum_{i=1}^{d}\left\|U_{i *}\right\|_{2}^{2}}, i=1, \ldots, d$
- Let $i_{1}, \ldots, i_{m} \in\{1, \ldots, d\}$ denote $m$ indices selected by following $p_{i}$
- Let $A \in \mathbb{R}^{m \times d}$ be sampling-and-rescaling matrix:



## Leverage-score sampling (Drineas et al., 2006)

Let $X \in \mathbb{R}^{d \times n}$ be a rank- $k$ matrix

- $X=U \Sigma V^{\top}: U \in \mathbb{R}^{d \times k}, \Sigma \in \mathbb{R}^{k \times k}$
- Leverage scores $\left\|U_{i *}\right\|_{2}^{2}, i=1, \ldots, d$
- Let $p_{i}=\frac{\left\|U_{i *}\right\|_{2}^{2}}{\sum_{i=1}^{d}\left\|U_{i *}\right\|_{2}^{2}}, i=1, \ldots, d$
- Let $i_{1}, \ldots, i_{m} \in\{1, \ldots, d\}$ denote $m$ indices selected by following $p_{i}$
- Let $A \in \mathbb{R}^{m \times d}$ be sampling-and-rescaling matrix:



## Leverage-score sampling (Drineas et al., 2006)

Let $X \in \mathbb{R}^{d \times n}$ be a rank- $k$ matrix

- $X=U \Sigma V^{\top}: U \in \mathbb{R}^{d \times k}, \Sigma \in \mathbb{R}^{k \times k}$
- Leverage scores $\left\|U_{i *}\right\|_{2}^{2}, i=1, \ldots, d$
- Let $p_{i}=\frac{\left\|U_{i *}\right\|_{2}^{2}}{\sum_{i=1}^{d}\left\|U_{i *}\right\|_{2}^{2}}, i=1, \ldots, d$
- Let $i_{1}, \ldots, i_{m} \in\{1, \ldots, d\}$ denote $m$ indices selected by following $p_{i}$
- Let $A \in \mathbb{R}^{m \times d}$ be sampling-and-rescaling matrix:

$$
A_{i j}=\left\{\begin{array}{cc}
\frac{1}{\sqrt{m p_{j}}} & \text { if } j=i_{j} \\
0 & \text { otherwise }
\end{array}\right.
$$

- $A X \in \mathbb{R}^{m \times n}$ is a small sketch of $X$


## Properties of Leverage-score sampling

When $m=\Theta\left(\frac{k}{\epsilon^{2}} \log \left[\frac{2 k}{\delta}\right]\right)$, w.h.p $1-\delta$,

- $A U \in \mathbb{R}^{m \times k}$ is full column rank
- $\sigma_{i}^{2}(A U) \geq(1-\epsilon) \geq(1-\epsilon)^{2}$
- $\sigma_{i}^{2}(A U) \leq 1+\epsilon \leq(1+\epsilon)^{2}$
- Computational cost: compute top-k SVD of $X$, expensive
- Randomized algoritms to compute approximate leverage scores


## Properties of Leverage-score sampling

When $m=\Theta\left(\frac{k}{\epsilon^{2}} \log \left[\frac{2 k}{\delta}\right]\right)$, w.h.p $1-\delta$,

- $A U \in \mathbb{R}^{m \times k}$ is full column rank
- $\sigma_{i}^{2}(A U) \geq(1-\epsilon) \geq(1-\epsilon)^{2}$
- $\sigma_{i}^{2}(A U) \leq 1+\epsilon \leq(1+\epsilon)^{2}$
- Leverage-score sampling performs like a subspace embedding (only for $U$, the top singular vector matrix of $X$ )
- Computational cost: compute top- $k$ SVD of $X$, expensive
- Randomized algoritms to compute approximate leverage scores


## When uniform sampling makes sense?

Coherence measure

$$
\mu_{k}=\frac{d}{k} \max _{1 \leq i \leq d}\left\|U_{i *}\right\|_{2}^{2}
$$

- Valid when the coherence measure is small (some real data mining datasets have small coherence measures)
- The Nyström method usually uses uniform sampling (Gittens, 2011)


## Outline

(4) Randomized Algorithms

- Randomized Classification (Regression)
- Randomized Least-Squares Regression
- Randomized K-means Clustering
- Randomized Kernel methods
- Randomized Low-rank Matrix Approximation


## Classification

Classification problems:

$$
\min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i} \mathbf{w}^{\top} \mathbf{x}_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$



- $y_{i} \in\{+1,-1\}$ : label
- Loss function $\ell(z): z=y \mathbf{w}^{\top} \mathbf{x}$

1. SVMs: (squared) hinge loss $\ell(z)=\max (0,1-z)^{p}$, where $p=1,2$
2. Logistic Regression: $\ell(z)=\log (1+\exp (-z))$

## Randomized Classification

For large-scale high-dimensional problems, the computational cost of optimization is $O((n d+d \kappa) \log (1 / \epsilon))$.

Use random reduction $A \in \mathbb{R}^{d \times m}(m \ll d)$, we reduce $X \in \mathbb{R}^{n \times d}$ to $\widehat{X}=X A \in \mathbb{R}^{n \times m}$. Then solve

$$
\min _{\mathbf{u} \in \mathbb{R}^{m}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i} \mathbf{u}^{\top} \widehat{\mathbf{x}}_{i}\right)+\frac{\lambda}{2}\|\mathbf{u}\|_{2}^{2}
$$

- JL transforms
- Sparse subspace embeddings


## Randomized Classification

Two questions:

- Is there any performance guarantee?

- How to recover an accurate model in the original high-dimensional space?


## Randomized Classification

## Two questions:

- Is there any performance guarantee?
- margin is preserved: if data is linearly separable (Balcan et al., 2006) as long as $m \geq \frac{12}{\epsilon^{2}} \log \left(\frac{6 m}{\delta}\right)$
- generalization performance is preserved:

- How to recover an accurate model in the original high-dimensional space?
Dual Recovery (Zhang et al., 2014) and Dual Sparse Recovery (Yang et al., 2015)


## Randomized Classification

## Two questions:

- Is there any performance guarantee?
- margin is preserved: if data is linearly separable (Balcan et al., 2006) as long as $m \geq \frac{12}{\epsilon^{2}} \log \left(\frac{6 m}{\delta}\right)$
- generalization performance is preserved: if the data matrix if of low rank and $m=\Omega\left(\frac{k p l o y}{}(\log (k / \delta \epsilon)), \epsilon^{2}\right)$ (Paul et al., 2013)
- How to recover an accurate model in the original high-dimensional space?
et al., 2015)


## Randomized Classification

## Two questions:

- Is there any performance guarantee?
- margin is preserved: if data is linearly separable (Balcan et al., 2006) as long as $m \geq \frac{12}{\epsilon^{2}} \log \left(\frac{6 m}{\delta}\right)$
- generalization performance is preserved: if the data matrix if of low rank and $m=\Omega\left(\frac{k p l o y}{}(\log (k / \delta \epsilon)), \epsilon^{2}\right)$ (Paul et al., 2013)
- How to recover an accurate model in the original high-dimensional space?
et al., 2015)


## Randomized Classification

Two questions:

- Is there any performance guarantee?

- How to recover an accurate model in the original high-dimensional space?
Dual Recovery (Zhang et al., 2014) and Dual Sparse Recovery (Yang et al., 2015)


## The Dual probelm

Using Fenchel conjugate

$$
\ell_{i}^{*}\left(\alpha_{i}\right)=\max _{\alpha_{i}} \alpha_{i} z-\ell\left(z, y_{i}\right)
$$

Primal:

$$
\mathbf{w}_{*}=\arg \min _{\mathbf{w} \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{w}^{\top} \mathbf{x}_{i}, y_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}
$$

Dual:

$$
\alpha_{*}=\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} X X^{\top} \alpha
$$

From dual to primal:

$$
\mathbf{w}_{*}=-\frac{1}{\lambda n} X^{\top} \alpha_{*}
$$

## Dual Recovery for Randomized Reduction

From dual formulation: $\mathbf{w}_{*}$ lies in the row space of the data matrix $X \in \mathbb{R}^{n \times d}$

- Dual Recovery: $\widetilde{\mathbf{w}}_{*}=-\frac{1}{\lambda n} X^{\top} \widehat{\alpha}_{*}$, where $\widehat{\alpha}_{*}=\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} \widehat{X} \widehat{X}^{\top} \alpha$
and $\hat{X}=X A \in \mathbb{R}^{n \times m}$
- Subspace Embedding $A$ with $m=\Theta\left(r \log (r / \delta) \epsilon^{-2}\right)$
- Guarantee: under low-rank assumption of the data matrix $X$ (e.g., $\operatorname{rank}(X)=r$ ), with a high probability $1-\delta$,


## Dual Recovery for Randomized Reduction

From dual formulation: $\mathbf{w}_{*}$ lies in the row space of the data matrix $X \in \mathbb{R}^{n \times d}$

- Dual Recovery: $\widetilde{\mathbf{w}}_{*}=-\frac{1}{\lambda n} X^{\top} \widehat{\alpha}_{*}$, where

$$
\widehat{\alpha}_{*}=\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} \widehat{X} \widehat{X}^{\top} \alpha
$$

and $\widehat{X}=X A \in \mathbb{R}^{n \times m}$

- Guarantee: under low-rank assumption of the data matrix $X$ (e.g., $\operatorname{rank}(X)=r)$, with a high probability $1-\delta$,


## Dual Recovery for Randomized Reduction

From dual formulation: $\mathbf{w}_{*}$ lies in the row space of the data matrix $X \in \mathbb{R}^{n \times d}$

- Dual Recovery: $\widetilde{\mathbf{w}}_{*}=-\frac{1}{\lambda n} X^{\top} \widehat{\alpha}_{*}$, where

$$
\widehat{\alpha}_{*}=\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} \widehat{X} \hat{X}^{\top} \alpha
$$

and $\widehat{X}=X A \in \mathbb{R}^{n \times m}$

- Subspace Embedding $A$ with $m=\Theta\left(r \log (r / \delta) \epsilon^{-2}\right)$
- Guarantee: under low-rank assumption of the data matrix $X$ (e.g., $\operatorname{rank}(X)=r)$, with a high probability $1-\delta$,

$$
\left\|\widetilde{\mathbf{w}}_{*}-\mathbf{w}_{*}\right\|_{2} \leq \frac{\epsilon}{1-\epsilon}\left\|\mathbf{w}_{*}\right\|_{2}
$$

## Dual Sparse Recovery for Randomized Reduction

Assume the optimal dual solution $\alpha_{*}$ is sparse (i.e., the number of support vectors is small)

- Dual Sparse Recovery: $\widetilde{\mathbf{w}}_{*}=-\frac{1}{\lambda n} X^{\top} \widehat{\alpha}_{*}$, where

where $\widehat{X}=X A \in \mathbb{R}^{n \times m}$
- JL transform $A$ with $m=\Theta\left(s \log (n / \delta) \epsilon^{-2}\right)$
- Guarantee: if $\alpha_{*}$ is $s$-sparse, with a high probability $1-\delta$

$$
\left\|\widetilde{\mathbf{w}}_{*}-\mathbf{w}_{*}\right\|_{2} \leq \epsilon\left\|\mathbf{w}_{*}\right\|_{2}
$$

## Dual Sparse Recovery for Randomized Reduction

Assume the optimal dual solution $\alpha_{*}$ is sparse (i.e., the number of support vectors is small)

- Dual Sparse Recovery: $\widetilde{\mathbf{w}}_{*}=-\frac{1}{\lambda n} X^{\top} \widehat{\alpha}_{*}$, where

$$
\widehat{\alpha}_{*}=\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} \widehat{X} \widehat{X}^{\top} \alpha-\frac{\tau}{n}\|\alpha\|_{1}
$$

where $\widehat{X}=X A \in \mathbb{R}^{n \times m}$

- Guarantee: if $\alpha_{*}$ is $s$-sparse, with a high probability $1-\delta$,



## Dual Sparse Recovery for Randomized Reduction

Assume the optimal dual solution $\alpha_{*}$ is sparse (i.e., the number of support vectors is small)

- Dual Sparse Recovery: $\widetilde{\mathbf{w}}_{*}=-\frac{1}{\lambda n} X^{\top} \widehat{\alpha}_{*}$, where

$$
\widehat{\alpha}_{*}=\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} \widehat{X} \widehat{X}^{\top} \alpha-\frac{\tau}{n}\|\alpha\|_{1}
$$

where $\widehat{X}=X A \in \mathbb{R}^{n \times m}$

- JL transform $A$ with $m=\Theta\left(s \log (n / \delta) \epsilon^{-2}\right)$
- Guarantee: if $\alpha_{*}$ is $s$-sparse, with a high probability $1-\delta$,

$$
\left\|\widetilde{\mathbf{w}}_{*}-\mathbf{w}_{*}\right\|_{2} \leq \epsilon\left\|\mathbf{w}_{*}\right\|_{2}
$$

## Dual Sparse Recovery

RCV1 text data, $n=677,399$, and $d=47,236$


Primal Error


## Outline

(4) Randomized Algorithms

- Randomized Classification (Regression)
- Randomized Least-Squares Regression
- Randomized K-means Clustering
- Randomized Kernel methods
- Randomized Low-rank Matrix Approximation


## Least-squares regression

Let $X \in \mathbb{R}^{n \times d}$ with $d \ll n$ and $b \in \mathbb{R}^{n}$. The least-squares regression problem is to find $\mathbf{w}_{*}$ such that

$$
\mathbf{w}_{*}=\arg \min _{\mathbf{w} \in \mathbb{R}^{d}}\|X \mathbf{w}-b\|_{2}
$$

- Computational Cost: $O\left(n d^{2}\right)$
- Goal of RA: o(nd $\left.{ }^{2}\right)$


## Randomized Least-squares regression

Let $A \in \mathbb{R}^{m \times n}$ be a random reduction matrix. Solve

$$
\widehat{\mathbf{w}}_{*}=\arg \min _{\mathbf{w} \in \mathbb{R}^{d}}\|A(X \mathbf{w}-b)\|_{2}=\|A X \mathbf{w}-A b\|_{2}
$$

- Computational Cost: $O\left(m d^{2}\right)+$ reduction time


## Randomized Least-squares regression

Theoretical Guarantees (Sarlós, 2006; Drineas et al., 2011; Nelson \& Nguyen, 2012):

$$
\left\|X \widehat{\mathbf{w}}_{*}-b\right\|_{2} \leq(1+\epsilon)\left\|X \mathbf{w}_{*}-b\right\|_{2}
$$

Total Time $O\left(n n z(X)+d^{3} \log (d / \epsilon) \epsilon^{-2}\right)$

## Outline

(4) Randomized Algorithms

- Randomized Classification (Regression)
- Randomized Least-Squares Regression
- Randomized K-means Clustering
- Randomized Kernel methods
- Randomized Low-rank Matrix Approximation


## K-means Clustering

Let $\mathbf{x}_{1}, \ldots, \mathbf{x}_{n} \in \mathbb{R}^{d}$ be a set of data points.
K-means clustering aims to solve

$$
\min _{C_{1}, \ldots, C_{k}} \sum_{j=1}^{k} \sum_{\mathbf{x}_{i} \in C_{j}}\left\|\mathbf{x}_{i}-\mu_{j}\right\|_{2}^{2}
$$

Computational Cost: $O(n d k t)$, where $t$ is number of iterations.

## Randomized Algorithms for K-means Clustering

Let $X=\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right)^{\top} \in \mathbb{R}^{n \times d}$ be the data matrix.
High-dimensional data: Random Sketch: $\widehat{X}=X A \in \mathbb{R}^{n \times m}, \ell \ll d$

Approximate K-means:


## Randomized Algorithms for K-means Clustering

Let $X=\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right)^{\top} \in \mathbb{R}^{n \times d}$ be the data matrix.
High-dimensional data: Random Sketch: $\widehat{X}=X A \in \mathbb{R}^{n \times m}, \ell \ll d$

Approximate K-means:

$$
\min _{C_{1}, \ldots, C_{k}} \sum_{j=1}^{k} \sum_{\widehat{\mathbf{x}}_{i} \in C_{j}}\left\|\widehat{\mathbf{x}}_{i}-\widehat{\mu}_{j}\right\|_{2}^{2}
$$

## Randomized Algorithms for K-means Clustering

For random sketch: JL transforms, sparse subspace embedding all work

- JL transform: $m=O\left(\frac{k \log (k /(\epsilon \delta))}{\epsilon^{2}}\right)$
- Sparse subspace embedding: $m=O\left(\frac{k^{2}}{\epsilon^{2} \delta}\right)$
- $\epsilon$ relates to the approximation accuracy
- Analysis of approximation error for K-means can be formulates as Constrained Low-rank Approximation (Cohen et al., 2015)

$$
\min _{Q^{\top} Q=1}\left\|X-Q Q^{\top} X\right\|_{F}^{2}
$$

where $Q$ is orthonormal.

## Outline

(4) Randomized Algorithms

- Randomized Classification (Regression)
- Randomized Least-Squares Regression
- Randomized K-means Clustering
- Randomized Kernel methods
- Randomized Low-rank Matrix Approximation


## Kernel methods

- Kernel function: $\kappa(\cdot, \cdot)$
- a set of examples $\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}$
- Kernel matrix: $K \in \mathbb{R}^{n \times n}$ with $K_{i j}=\kappa\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)$
- $K$ is a PSD matrix
- Computational and memory costs: $\Omega\left(n^{2}\right)$
- Approximation methos
- The Nyström method
- Random Fourier features


## Kernel methods

- Kernel function: $\kappa(\cdot, \cdot)$
- a set of examples $\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}$
- Kernel matrix: $K \in \mathbb{R}^{n \times n}$ with $K_{i j}=\kappa\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)$
- $K$ is a PSD matrix
- Computational and memory costs: $\Omega\left(n^{2}\right)$
- Approximation methos
- The Nyström method
- Random Fourier features


## The Nyström method



Let $A \in \mathbb{R}^{n \times \ell}$ be uniform sampling matrix.

$$
\begin{gathered}
B=K A \in \mathbb{R}^{n \times \ell} \\
C=A^{\top} B=A^{\top} K A
\end{gathered}
$$

The Nyström approximation (Drineas \& Mahoney, 2005)


Computational Cost: $O\left(\ell^{3}+n \ell^{2}\right)$

## The Nyström method



Let $A \in \mathbb{R}^{n \times \ell}$ be uniform sampling matrix.

$$
\begin{gathered}
B=K A \in \mathbb{R}^{n \times \ell} \\
C=A^{\top} B=A^{\top} K A
\end{gathered}
$$

The Nyström approximation (Drineas \& Mahoney, 2005)

$$
\widehat{K}=B C^{\dagger} B^{\top}
$$

Computational Cost: $O\left(\ell^{3}+n \ell^{2}\right)$

## The Nyström based kernel machine

The dual problem:

$$
\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} B C^{\dagger} B^{\top} \alpha
$$

Solve it like solving a linear method: $\widehat{X}=B C^{-1 / 2} \in \mathbb{R}^{n \times \ell}$

$$
\arg \max _{\alpha \in \mathbb{R}^{n}}-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 \lambda n^{2}} \alpha^{\top} \hat{X} \hat{X}^{\top} \alpha
$$

## The Nyström based kernel machine


webspam



ijenn1

covtype


## Random Fourier Features (RFF)

## Bochner's theorem

A shift-invariant kernel $\kappa(\mathbf{x}, \mathbf{y})=\kappa(\mathbf{x}-\mathbf{y})$ is a valid kernel if only if $\kappa(\delta)$ is the Fourier transform of a non-negative measure, i.e.,

$$
\kappa(\mathbf{x}-\mathbf{y})=\int p(\omega) e^{-j \omega^{\top}(\mathbf{x}-\mathbf{y})} d \omega
$$

RFF (Rahimi \& Recht, 2008): generate a set of $\omega_{1}, \ldots, \omega_{m} \in \mathbb{R}^{d}$ following


## Random Fourier Features (RFF)

Bochner's theorem
A shift-invariant kernel $\kappa(\mathbf{x}, \mathbf{y})=\kappa(\mathbf{x}-\mathbf{y})$ is a valid kernel if only if $\kappa(\delta)$ is the Fourier transform of a non-negative measure, i.e.,

$$
\kappa(\mathbf{x}-\mathbf{y})=\int p(\omega) e^{-j \omega^{\top}(\mathbf{x}-\mathbf{y})} d \omega
$$

RFF (Rahimi \& Recht, 2008): generate a set of $\omega_{1}, \ldots, \omega_{m} \in \mathbb{R}^{d}$ following $p(\omega)$. For an example $\mathbf{x} \in \mathbb{R}^{d}$, construct

$$
\widehat{\mathbf{x}}=\left(\cos \left(\omega_{1}^{\top} \mathbf{x}\right), \sin \left(\omega_{1}^{\top} \mathbf{x}\right), \ldots, \cos \left(\omega_{m}^{\top} \mathbf{x}\right), \sin \left(\omega_{m}^{\top} \mathbf{x}\right)\right)^{\top} \in \mathbb{R}^{2 m}
$$

## Random Fourier Features (RFF)

Bochner's theorem
A shift-invariant kernel $\kappa(\mathbf{x}, \mathbf{y})=\kappa(\mathbf{x}-\mathbf{y})$ is a valid kernel if only if $\kappa(\delta)$ is the Fourier transform of a non-negative measure, i.e.,

$$
\kappa(\mathbf{x}-\mathbf{y})=\int p(\omega) e^{-j \omega^{\top}(\mathbf{x}-\mathbf{y})} d \omega
$$

RFF (Rahimi \& Recht, 2008): generate a set of $\omega_{1}, \ldots, \omega_{m} \in \mathbb{R}^{d}$ following $p(\omega)$. For an example $\mathbf{x} \in \mathbb{R}^{d}$, construct

$$
\widehat{\mathbf{x}}=\left(\cos \left(\omega_{1}^{\top} \mathbf{x}\right), \sin \left(\omega_{1}^{\top} \mathbf{x}\right), \ldots, \cos \left(\omega_{m}^{\top} \mathbf{x}\right), \sin \left(\omega_{m}^{\top} \mathbf{x}\right)\right)^{\top} \in \mathbb{R}^{2 m}
$$

RBF kernel $\exp \left(-\frac{\|\mathbf{x}-\mathbf{y}\|_{2}^{2}}{2 \gamma^{2}}\right): p(\omega)=\mathcal{N}\left(0, \gamma^{2}\right)$

## The Nyström method vs RFF (Yang et al., 2012)

- functional approximation framework
- The Nyström method: data-dependent bases
- RFF: data independent bases
- In certain cases (e.g., large eigen-gap, skewed eigen-value distribution): the generalization performance of the Nyström method is better than RFF


## The Nyström method vs RFF



## Outline

(4) Randomized Algorithms

- Randomized Classification (Regression)
- Randomized Least-Squares Regression
- Randomized K-means Clustering
- Randomized Kernel methods
- Randomized Low-rank Matrix Approximation


## Randomized low-rank matrix approximation

Let $X \in \mathbb{R}^{n \times d}$. The goal is to obtain

$$
\widehat{U} \widehat{\Sigma} \widehat{V}^{\top} \approx X
$$

where $\widehat{U} \in \mathbb{R}^{n \times k}, \widehat{V} \in \mathbb{R}^{d \times k}$ have orthonormal columns, $\widehat{\Sigma} \in \mathbb{R}^{k \times k}$ is a diagonal matrix with nonegative entries

- $k$ is target rank
- The best rank-k approximation $X_{k}=U_{k} \Sigma_{k} V_{k}^{\top}$
- Approximation error

$$
\left\|\hat{U} \widehat{\Sigma} \widehat{V}^{\top}-X\right\|_{\xi} \leq(1+\epsilon)\left\|U_{k} \Sigma_{k} V_{k}^{\top}-X\right\|_{\xi}
$$

where $\xi=F$ or $\xi=2$

## Why low-rank approximation?

Applications in Data mining and Machine learning

- PCA
- Spectral clustering


## Why randomized algorithms?

Deterministic Algorithms

- Truncated SVD $O(n d \min (n, d))$
- Rank-Revealing QR factorization $O(n d k)$
- Krylov subspace method (e.g. Lanczos algorithm):
$O\left(k T_{\text {mult }}+(n+d) k^{2}\right)$, where $T_{\text {mult }}$ denotes the cost of matrix-vector product.
- Speed can be faster (e.g., $O(n d \log (k)))$
- Output more robust (e.g. Lanczos requires sophisticated modifications)
- Can be pass efficient
- Can exploit parallel algorithms


## Why randomized algorithms?

Deterministic Algorithms

- Truncated SVD $O(n d \min (n, d))$
- Rank-Revealing QR factorization $O(n d k)$
- Krylov subspace method (e.g. Lanczos algorithm):
$O\left(k T_{\text {mult }}+(n+d) k^{2}\right)$, where $T_{\text {mult }}$ denotes the cost of matrix-vector product.
Randomized Algorithms
- Speed can be faster (e.g., $O(n d \log (k)))$
- Output more robust (e.g. Lanczos requires sophisticated modifications)
- Can be pass efficient
- Can exploit parallel algorithms


## Randomized algorithms for low-rank matrix approximation

The Basic Randomized Algorithms for Approximating $X \in \mathbb{R}^{n \times d}$ (Halko et al., 2011)
(1) Obtain a small sketch by $Y=X A \in \mathbb{R}^{n \times m}$
(2) Compute $Q \in \mathbb{R}^{n \times m}$ that contains the orthonormal basis of $\operatorname{col}(Y)$
(3) Compute SVD of $Q^{\top} X=U \Sigma V^{\top}$
(9) Approximation $X \approx \widetilde{U} \Sigma V^{\top}$, where $\widetilde{U}=Q U$ Explanation: If $\operatorname{col}(X A)$ captures the top- $k$ column space of $X$ well, i.e.,

$$
\left\|X-Q Q^{\top} X\right\| \leq \varepsilon
$$

then

$$
\left\|X-\widetilde{U} \Sigma V^{\top}\right\| \leq \varepsilon
$$

## Randomized algorithms for low-rank matrix approximation

Three questions:
(1) What is the value of $m$ ?
gives superb results
(2) What is the computational cost?

- Subsampled Randomized Hadmard Transform: can be as fast as
(3) What is the quality?
- Theoretical Guarantee:
- Practically, very accurate


## Randomized algorithms for low-rank matrix approximation

Three questions:
(1) What is the value of $m$ ?

- $m=k+p, p$ is the oversampling parameter. In practice $p=5$ or 10 gives superb results
(2) What is the computational cost?
- Subsampled Randomized Hadmard Transform: can be as fast as
(3) What is the quality?
- Theoretical Guarantee:
- Practically, very accurate


## Randomized algorithms for low-rank matrix approximation

Three questions:
(1) What is the value of $m$ ?

- $m=k+p, p$ is the oversampling parameter. In practice $p=5$ or 10 gives superb results
(2) What is the computational cost?
- Subsampled Randomized Hadmard Transform: can be as fast as $O\left(n d \log (k)+k^{2}(n+d)\right)$
(3) What is the quality?
- Theoretical Guarantee:
- Practically, very accurate


## Randomized algorithms for low-rank matrix approximation

Three questions:
(1) What is the value of $m$ ?

- $m=k+p, p$ is the oversampling parameter. In practice $p=5$ or 10 gives superb results
(2) What is the computational cost?
- Subsampled Randomized Hadmard Transform: can be as fast as $O\left(n d \log (k)+k^{2}(n+d)\right)$
(3) What is the quality?
- Theoretical Guarantee:
- Practically, very accurate


## Randomized algorithms for low-rank matrix approximation



## Randomized algorithms for low-rank matrix approximation

Other things

- Use power iteration to reduce the error: use $\left(X X^{\top}\right)^{q} X$
- Can use sparse JL transform/subspace embedding matrices (Frobenius norm guarantee only)


## Outline

(1) Basics
(2) Optimization
(3) Randomized Dimension Reduction
(4) Randomized Algorithms
(5) Concluding Remarks

## How to address big data challenge?

- Optimization perspective: improve convergence rates, exploring properties of functions
- stochastic optimization (e.g., SDCA, SVRG, SAGA)
- distributed optimization (e.g., DisDCA)
- Randomization perspective: reduce data size, exploring properties of data
- randomized feature reduction (e.g., reduce the number of features)
- randomized instance reduction (e.g., reduce the number of instances)


## How can we address big data challenge?

- Optimization perspective: improve convergence rates, exploring properties of functions
- Pro: can obtain the optimal solution
- Con: high computational/communication costs
- Randomization perspective: reduce data size, exploring properties of data
- Pro: fast
- Con: still exists recovery error

Can we combine the benefits of two techniques?

## Combine Randomization and Optimization (Yang et al., 2015)

- Use randomization (Dual Spare Recovery) to obtain a good initial solution
- Initialize distributed optimization (DisDCA) to reduce cost of computation/communication
- Observe 1 or 2 epochs of computations (1 or 2 communications) suffice to obtain the same performance of pure optimization


## Big Data Experiments

KDDcup Data: $n=8,407,752, d=29,890,095,10$ machines, $m=1024$



## Thank You! Questions?

## References I

Achlioptas, Dimitris. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. Journal of Computer and System Sciences, 66(4):671-687, 2003.
Balcan, Maria-Florina, Blum, Avrim, and Vempala, Santosh. Kernels as features: on kernels, margins, and low-dimensional mappings. Machine Learning, 65(1):79-94, 2006.
Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. Distributed optimization and statistical learning via the alternating direction methods of multiplies. Foundations and Trends in Machine Learning, 3(1):1-122, 2011.

## References II

Cohen, Michael B., Elder, Sam, Musco, Cameron, Musco, Christopher, and Persu, Madalina. Dimensionality reduction for k-means clustering and low rank approximation. In Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC), pp. 163-172, 2015.

Dasgupta, Anirban, Kumar, Ravi, and Sarlos, Tamás. A sparse johnson: Lindenstrauss transform. In Proceedings of the 42nd ACM symposium on Theory of computing, STOC '10, pp. 341-350, 2010.
Dasgupta, Sanjoy and Gupta, Anupam. An elementary proof of a theorem of Johnson and Lindenstrauss. Random Structures \& Algorithms, 22(1): 60-65, 2003.
Defazio, Aaron, Bach, Francis, and Lacoste-Julien, Simon. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In NIPS, 2014.

## References III

Drineas, Petros and Mahoney, Michael W. On the nystrom method for approximating a gram matrix for improved kernel-based learning. Journal of Machine Learning Research, 6:2005, 2005.
Drineas, Petros, Mahoney, Michael W., and Muthukrishnan, S. Sampling algorithms for 12 regression and applications. In ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1127-1136, 2006.
Drineas, Petros, Mahoney, Michael W., Muthukrishnan, S., and Sarlós, Tamàs. Faster least squares approximation. Numerische Mathematik, 117(2):219-249, February 2011.
Gittens, Alex. The spectral norm error of the naive nystrom extension. CoRR, 2011.
Golub, Gene H. and Ye, Qiang. Inexact preconditioned conjugate gradient method with inner-outer iteration. SIAM J. Sci. Comput, 21:1305-1320, 1997.

## References IV

Halko, Nathan, Martinsson, Per Gunnar., and Tropp, Joel A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM Review, 53(2):217-288, May 2011.

Hsieh, Cho-Jui, Chang, Kai-Wei, Lin, Chih-Jen, Keerthi, S. Sathiya, and Sundararajan, S. A dual coordinate descent method for large-scale linear svm. In ICML, pp. 408-415, 2008.
Johnson, Rie and Zhang, Tong. Accelerating stochastic gradient descent using predictive variance reduction. In NIPS, pp. 315-323, 2013.
Johnson, William and Lindenstrauss, Joram. Extensions of Lipschitz mappings into a Hilbert space. In Conference in modern analysis and probability (New Haven, Conn., 1982), volume 26, pp. 189-206. 1984.
Kane, Daniel M. and Nelson, Jelani. Sparser johnson-lindenstrauss transforms. Journal of the ACM, 61:4:1-4:23, 2014.

## References $V$

Lee, Jason, Ma, Tengyu, and Lin, Qihang. Distributed stochastic variance reduced gradient methods. Technical report, UC Berkeley, 2015.
Lin, Qihang, Lu, Zhaosong, and Xiao, Lin. An accelerated proximal
coordinate gradient method and its application to regularized empirical risk minimization. In NIPS, 2014.
Ma, Chenxin, Smith, Virginia, Jaggi, Martin, Jordan, Michael I., Richtárik, Peter, and Takác, Martin. Adding vs. averaging in distributed primal-dual optimization. In ICML, 2015.
Nelson, Jelani and Nguyen, Huy L. OSNAP: faster numerical linear algebra algorithms via sparser subspace embeddings. CoRR, abs/1211.1002, 2012.

## References VI

Nelson, Jelani and Nguyen, Huy L. OSNAP: faster numerical linear algebra algorithms via sparser subspace embeddings. In 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 117-126, 2013.

Nemirovski, A. and Yudin, D. On cezari's convergence of the steepest descent method for approximating saddle point of convex-concave functons. Soviet Math Dkl., 19:341-362, 1978.
Nesterov, Yurii. Efficiency of coordinate descent methods on huge-scale optimization problems. SIAM Journal on Optimization, 22:341-362, 2012.

Ozdaglar, Asu. Distributed multiagent optimization linear convergence rate of admm. Technical report, MIT, 2015.

## References VII

Paul, Saurabh, Boutsidis, Christos, Magdon-Ismail, Malik, and Drineas, Petros. Random projections for support vector machines. In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 498-506, 2013.
Rahimi, Ali and Recht, Benjamin. Random features for large-scale kernel machines. In Advances in Neural Information Processing Systems 20, pp. 1177-1184, 2008.
Recht, Benjamin. A simpler approach to matrix completion. Journal Machine Learning Research (JMLR), pp. 3413-3430, 2011.
Roux, Nicolas Le, Schmidt, Mark, and Bach, Francis. A stochastic gradient method with an exponential convergence rate for strongly-convex optimization with finite training sets. CoRR, 2012.

## References VIII

Sarlós, Tamás. Improved approximation algorithms for large matrices via random projections. In 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 143-152, 2006.
Shalev-Shwartz, Shai and Zhang, Tong. Stochastic dual coordinate ascent methods for regularized loss. Journal of Machine Learning Research, 14: 567-599, 2013.
Shamir, Ohad, Srebro, Nathan, and Zhang, Tong.
Communication-efficient distributed optimiztion using an approximate newton-type method. In ICML, 2014.
Tropp, Joel A. Improved analysis of the subsampled randomized hadamard transform. Advances in Adaptive Data Analysis, 3(1-2):115-126, 2011.
Tropp, Joel A. User-friendly tail bounds for sums of random matrices.
Found. Comput. Math., 12(4):389-434, August 2012. ISSN 1615-3375.

## References IX

Xiao, L. and Zhang, T. A proximal stochastic gradient method with progressive variance reduction. SIAM Journal on Optimization, 24(4): 2057-2075, 2014.
Yang, Tianbao. Trading computation for communication: Distributed stochastic dual coordinate ascent. NIPS'13, pp. -, 2013.
Yang, Tianbao, Li, Yu-Feng, Mahdavi, Mehrdad, Jin, Rong, and Zhou, Zhi-Hua. Nyström method vs random fourier features: A theoretical and empirical comparison". In Advances in Neural Information Processing Systems (NIPS), pp. 485-493, 2012.
Yang, Tianbao, Zhang, Lijun, Jin, Rong, and Zhu, Shenghuo. Theory of dual-sparse regularized randomized reduction. In Proceedings of the 32nd International Conference on Machine Learning, (ICML), pp. 305-314, 2015.

## References $X$

Zhang, Lijun, Mahdavi, Mehrdad, and Jin, Rong. Linear convergence with condition number independent access of full gradients. In NIPS, pp. 980-988. 2013.
Zhang, Lijun, Mahdavi, Mehrdad, Jin, Rong, Yang, Tianbao, and Zhu, Shenghuo. Random projections for classification: A recovery approach. IEEE Transactions on Information Theory (IEEE TIT), 60(11): 7300-7316, 2014.

Zhang, Yuchen and Xiao, Lin. Communication-efficient distributed optimization of self-concordant empirical loss. In ICML, 2015.

## Examples of Convex functions

- $a x+b, A \mathbf{x}+b$
- $x^{2},\|\mathbf{x}\|_{2}^{2}$
- $\exp (a x), \exp \left(\mathbf{w}^{\top} \mathbf{x}\right)$
- $\log (1+\exp (a x)), \log \left(1+\exp \left(\mathbf{w}^{\top} \mathbf{x}\right)\right)$
- $x \log (x), \sum_{i} x_{i} \log \left(x_{i}\right)$
- $\|\mathbf{x}\|_{p}, p \geq 1,\|\mathbf{x}\|_{p}^{2}$
- $\max _{i}\left(x_{i}\right)$


## Operations that preserve convexity

- Nonnegative scale: $a \cdot f(\mathbf{x})$ where $a \geq 0$
- Sum: $f(\mathbf{x})+g(\mathbf{x})$
- Composition with affine function $f(A \mathbf{x}+b)$
- Point-wise maximum: $\max _{i} f_{i}(\mathbf{x})$


## Examples:

- Least-squares regression: $\|A \mathbf{x}-b\|_{2}$
- SVM: $\frac{1}{n} \sum_{i=1}^{n} \max \left(0,1-y_{i} \mathbf{w}^{\top} \mathbf{x}_{i}\right)+\frac{\lambda}{2}\|\mathbf{w}\|_{2}^{2}$


## Smooth Convex function

- smooth: e.g. logistic loss $f(x)=\log (1+\exp (-x))$
$\|\nabla f(x)-\nabla f(y)\|_{2} \leq L\|x-y\|_{2}$ where $L>0$

Second Order Derivative is upper bounded $\left\|\nabla^{2} f(x)\right\|_{2} \leq L$


## Smooth Convex function

- smooth: e.g. $\log$ sticsmoothness $\mathrm{O}_{\text {p }}(1+\exp (-x))$
$\|\nabla f(x)-\nabla f(y)\|_{2} \leq L\|x-y\|_{2}$ where $L>0$

Second Order Derivative is upper bounded $\left\|\nabla^{2} f(x)\right\|_{2} \leq L$


## Strongly Convex function

- strongly convex: e.g. Euclidean norm $f(x)=\frac{1}{2}\|x\|_{2}^{2}$
$\|\nabla f(x)-\nabla f(y)\|_{2} \geq \lambda\|x-y\|_{2}$ where $\lambda>0$

Second Order Derivative is lower bounded $\left\|\nabla^{2} f(x)\right\|_{2} \geq \lambda$


## Strongly Convex function



## Smooth and Strongly Convex function

- smooth and strongly convex: e.g. quadratic function:

$$
\begin{aligned}
& f(z)=\frac{1}{2}(z-1)^{2} \\
& \quad \lambda\|x-y\|_{2} \leq\|\nabla f(x)-\nabla f(y)\|_{2} \leq L\|x-y\|_{2}, \quad L \geq \lambda>0
\end{aligned}
$$

## Chernoff bound

Let $X_{1}, \ldots, X_{n}$ be independent random variables. Assume $0 \leq X_{i} \leq 1$. Let $X=X_{1}+\ldots+X_{n} . \mu=\mathrm{E}[X]$. Then

$$
\begin{aligned}
& \operatorname{Pr}(X \geq(1+\epsilon) \mu) \leq \exp \left(-\frac{\epsilon^{2}}{2+\epsilon} \mu\right) \\
& \operatorname{Pr}(X \leq(1-\epsilon) \mu) \leq \exp \left(-\frac{\epsilon^{2}}{2} \mu\right)
\end{aligned}
$$

or

$$
\operatorname{Pr}(|X-\mu| \geq \epsilon \mu) \leq 2 \exp \left(-\frac{\epsilon^{2}}{2+\epsilon} \mu\right) \leq 2 \exp \left(-\frac{\epsilon^{2}}{3} \mu\right)
$$

the last inequality holds when $0<\epsilon \leq 1$

## Theoretical Guarantee of RA for low-rank approximation

$$
X=U\left[\begin{array}{ll}
\Sigma_{1} & \\
& \Sigma_{2}
\end{array}\right]\left[\begin{array}{l}
V_{1}^{\top} \\
V_{2}^{\top}
\end{array}\right]
$$

- $X \in \mathbb{R}^{m \times n}$ : the target matrix
- $\Sigma_{1} \in \mathbb{R}^{k \times k}, V_{1} \in \mathbb{R}^{n \times k}$
- $A \in \mathbb{R}^{n \times \ell}$ : random reduction matrix
- $Y=X A \in \mathbb{R}^{m \times \ell}$ : the small sketch

Key inequality:

$$
\left\|\left(I-P_{Y}\right) X\right\|^{2} \leq\left\|\Sigma_{2}\right\|^{2}+\left\|\Sigma_{2} \Omega_{2} \Omega_{1}^{\dagger}\right\|^{2}
$$

## Gaussian Matrices

- $G$ is a standard Gaussian matrix
- $U$ and $V$ are orthonormal matrices
- $U^{T} G V$ follows the standard Gaussian distribution
- $\mathrm{E}\left[\|S G T\|_{F}^{2}\right]=\|S\|_{F}^{2}\|T\|_{F}^{2}$
- $\mathrm{E}[\|S G T\|] \leq\|S\|\|T\|_{F}+\|S\|_{F}\|T\|$
- Concentration for function of a Gaussian matrix. Suppose $h$ is a Lipschitz function on matrices

$$
h(X)-h(Y) \leq L\|X-Y\|_{F}
$$

Then

$$
\operatorname{Pr}(h(G) \geq \mathrm{E}[h(G)]+L t) \leq e^{-t^{2} / 2}
$$

## Analysis for Randomized Least-square regression

Let $X=U \Sigma V^{\top}$

$$
\mathbf{w}_{*}=\arg \min _{\mathbf{w} \in \mathbb{R}^{d}}\|X \mathbf{w}-b\|_{2}
$$

Let $Z=\left\|X \mathbf{w}_{*}-b\right\|_{2}, \omega=b-X \mathbf{w}_{*}$, and $X \mathbf{w}_{*}=U \alpha$

$$
\widehat{\mathbf{w}}_{*}=\arg \min _{\mathbf{w} \in \mathbb{R}^{d}}\|A(X \mathbf{w}-b)\|_{2}
$$

Since $b-X \mathbf{w}_{*}=b-X\left(X^{\top} X\right)^{\dagger} X^{\top} b=\left(I-U U^{\top}\right) b, X \widehat{\mathbf{w}}_{*}-X \mathbf{w}_{*}=U \beta$.
Then

$$
\left\|X \widehat{\mathbf{w}}_{*}-b\right\|_{2}=\left\|X \mathbf{w}_{*}-b\right\|_{2}+\left\|X \widehat{\mathbf{w}}_{*}-X \mathbf{w}\right\|_{2}=Z+\|\beta\|_{2}
$$

## Analysis for Randomized Least-square regression

$$
\begin{gathered}
A U(\alpha+\beta)=A X \widehat{\mathbf{w}}_{*}=A X(A X)^{\dagger} A b=P_{A X}(A b)=P_{A U}(A b) \\
P_{A U}(A b)=P_{A U}(A(\omega+U \alpha))=A U \alpha+P_{A U}(A \omega)
\end{gathered}
$$

Hence
$U^{\top} A^{\top} A U \beta=(A U)^{\top}(A U)(A U)^{\dagger} A \omega=(A U)^{\top}(A U)\left((A U)^{\top} A U\right)^{-1}(A U)^{\top} A \omega$ where we use $A U$ is full column matrix. Then

$$
\begin{gathered}
U^{\top} A^{\top} A U \beta=U^{\top} A^{\top} A \omega \\
\|\beta\|_{2}^{2} / 2 \leq\left\|U^{\top} A^{\top} A U \beta\right\|_{2}^{2}=\left\|U^{\top} A^{\top} A \omega\right\|_{2}^{2} \leq \epsilon^{\prime 2}\|U\|_{F}^{2}\|\omega\|_{2}^{2}
\end{gathered}
$$

where the last inequality uses the matrix products approximation shown in next slide. Since $\|U\|_{F}^{2} \leq d$, setting $\epsilon^{\prime}=\sqrt{\frac{\epsilon}{d}}$ suffices.

## Approximate Matrix Products

Given $X \in \mathbb{R}^{n \times d}$ and $Y \in \mathbb{R}^{d \times p}$, let $A \in \mathbb{R}^{m \times d}$ one of the following matrices

- a JL transform matrix with $m=\Theta\left(\epsilon^{-2} \log ((n+p) / \delta)\right)$
- the sparse subspace embedding with $m=\Theta\left(\epsilon^{-2}\right)$
- leverage-score sampling matrix based on $p_{i} \geq \frac{\left\|X_{i *}\right\|_{2}^{2}}{2\|X\|_{F}^{2}}$ and $m=\Theta\left(\epsilon^{-2}\right)$ Then w.h.p $1-\delta$

$$
\left\|X A^{\top} A Y-X Y\right\|_{F} \leq \epsilon\|X\|_{F}\|Y\|_{F}
$$

## Analysis for Randomized Least-square regression

$A \in \mathbb{R}^{m \times n}$

1. Subspace embedding: $A U$ full column rank
2. Matrix product approximation: $\sqrt{\epsilon / d}$

Order of $m$

- JL transforms: 1. $O(d \log (d))$, 2. $O\left(d \log (d) \epsilon^{-1}\right) \Rightarrow O\left(d \log (d) \epsilon^{-1}\right)$
- Sparse subspace embedding: 1. $O\left(d^{2}\right)$, 2. $O\left(d \epsilon^{-1}\right) \Rightarrow O\left(d^{2} \epsilon^{-1}\right)$ If we use $\operatorname{SSE}\left(A_{1} \in \mathbb{R}^{m_{1} \times n}\right)$ and JL transform $A_{2} \in \mathbb{R}^{m_{2} \times m_{1}}$

$$
\begin{aligned}
\left\|A_{2} A_{1}\left(X \mathbf{w}_{*}^{2}-b\right)\right\|_{2} & \leq(1+\epsilon)\left\|A_{1}\left(X \mathbf{w}_{*}^{1}-b\right)\right\|_{2} \\
& \leq(1+\epsilon)\left\|A_{1}\left(X \mathbf{w}_{*}-b\right)\right\|_{2} \leq(1+\epsilon)^{2}\left\|X \mathbf{w}_{*}-b\right\|
\end{aligned}
$$

with $m_{1}=O\left(d^{2} \epsilon^{-2}\right)$ and $m_{2}=d \log (d) \epsilon^{-1}, \mathbf{w}_{*}^{2}$ is the optimal solution using $A_{2} A_{1}$ and $\mathbf{w}_{*}^{1}$ is the optimal using $A_{1}$ and $\mathbf{w}_{*}$ is the original optimal solution.

## Randomized Least-squares regression

Theoretical Guarantees (Sarlós, 2006; Drineas et al., 2011; Nelson \& Nguyen, 2012):

$$
\left\|X \widehat{\mathbf{w}}_{*}-b\right\|_{2} \leq(1+\epsilon)\left\|X \mathbf{w}_{*}-b\right\|_{2}
$$

- If $A$ is a fast JL transform with $m=\Theta\left(\epsilon^{-1} d \log (d)\right)$ : Total Time $O\left(n d \log (m)+d^{3} \log (d) \epsilon^{-1}\right)$
- If $A$ is a Sparse Subspace Embedding with $m=\Theta\left(d^{2} \epsilon^{-1}\right)$ : Total Time $O\left(n n z(X)+d^{4} \epsilon^{-1}\right)$
- If $A=A_{1} A_{2}$ combine fast $\mathrm{JL}\left(m_{1}=\Theta\left(\epsilon^{-1} d \log (d)\right)\right)$ and SSE $\left(m_{2}=\Theta\left(d^{2} \epsilon^{-2}\right)\right)$ : Total Time $O\left(n n z(X)+d^{3} \log (d / \epsilon) \epsilon^{-2}\right)$


## Matrix Chernoff bound

## Lemma (Matrix Chernoff (Tropp, 2012))

Let $\mathcal{X}$ be a finite set of PSD matrices with dimension $k$, and suppose that $\max _{X \in \mathcal{X}} \lambda_{\max }(X) \leq B$. Sample $\left\{X_{1}, \ldots, X_{\ell}\right\}$ independently from $\mathcal{X}$.
Compute

$$
\mu_{\max }=\ell \lambda_{\max }\left(\mathrm{E}\left[X_{1}\right]\right), \quad \mu_{\min }=\ell \lambda_{\min }\left(\mathrm{E}\left[X_{1}\right]\right)
$$

Then

$$
\begin{aligned}
& \operatorname{Pr}\left\{\lambda_{\max }\left(\sum_{i=1}^{\ell} X_{i}\right) \geq(1+\delta) \mu_{\max }\right\} \leq k\left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right]^{\frac{\mu_{\max }}{B}} \\
& \operatorname{Pr}\left\{\lambda_{\min }\left(\sum_{i=1}^{\ell} X_{i}\right) \leq(1-\delta) \mu_{\min }\right\} \leq k\left[\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right]^{\frac{\mu_{\min }}{B}}
\end{aligned}
$$

To simplify the usage of Matrix Chernoff bound, we note that

$$
\begin{aligned}
& {\left[\frac{e^{-\delta}}{[1-\delta]^{1-\delta}}\right]^{\mu} \leq \exp \left(-\frac{\delta^{2}}{2}\right)} \\
& {\left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right]^{\mu} \leq \exp \left(-\mu \delta^{2} / 3\right), \delta \leq 1} \\
& {\left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right]^{\mu} \leq \exp (-\mu \delta \log (\delta) / 2), \delta>1}
\end{aligned}
$$

## Noncommutative Bernstein Inequality

## Lemma (Noncommutative Bernstein Inequality (Recht, 2011))

Let $Z_{1}, \ldots, Z_{L}$ be independent zero-mean random matrices of dimension $d_{1} \times d_{2}$. Suppose $\tau_{j}^{2}=\max \left\{\left\|\mathrm{E}\left[Z_{j} Z_{j}^{\top}\right]\right\|_{2}, \| \mathrm{E}\left[Z_{j}^{\top} Z_{j} \|_{2}\right\}\right.$ and $\left\|Z_{j}\right\|_{2} \leq M$ almost surely for all $k$. Then, for any $\epsilon>0$,

$$
\operatorname{Pr}\left[\left\|\sum_{j=1}^{L} Z_{j}\right\|_{2}>\epsilon\right] \leq\left(d_{1}+d_{2}\right) \exp \left[\frac{-\epsilon^{2} / 2}{\sum_{j=1}^{L} \tau_{j}^{2}+M \epsilon / 3}\right]
$$

## Randomized Algorithms for K-means Clustering

K-means:

$$
\sum_{j=1}^{k} \sum_{\mathbf{x}_{i} \in C_{j}}\left\|\mathbf{x}_{i}-\mu_{j}\right\|_{2}^{2}=\left\|X-C C^{\top} X\right\|_{F}^{2}
$$

where $C \in \mathbb{R}^{n \times k}$ is the scaled cluster indicator matrix such that $C^{\top} C=I$.

Constrained Low-rank Approximation (Cohen et al., 2015)

where $\mathcal{S}=Q Q^{\top}$ is any set of rank $k$ orthogonal projection matrix with orthonormal $Q \in \mathbb{R}^{n \times k}$ Low-rank Approximation: $\mathcal{S}$ is the set of all rank $k$ orthogonal projection

## Randomized Algorithms for K-means Clustering

K-means:

$$
\sum_{j=1}^{k} \sum_{\mathbf{x}_{i} \in C_{j}}\left\|\mathbf{x}_{i}-\mu_{j}\right\|_{2}^{2}=\left\|X-C C^{\top} X\right\|_{F}^{2}
$$

where $C \in \mathbb{R}^{n \times k}$ is the scaled cluster indicator matrix such that $C^{\top} C=I$.

Constrained Low-rank Approximation (Cohen et al., 2015)

$$
\min _{P \in \mathcal{S}}\|X-P X\|_{F}^{2}
$$

where $\mathcal{S}=Q Q^{\top}$ is any set of rank $k$ orthogonal projection matrix with orthonormal $Q \in \mathbb{R}^{n \times k}$

## Randomized Algorithms for K-means Clustering

K-means:

$$
\sum_{j=1}^{k} \sum_{\mathbf{x}_{i} \in C_{j}}\left\|\mathbf{x}_{i}-\mu_{j}\right\|_{2}^{2}=\left\|X-C C^{\top} X\right\|_{F}^{2}
$$

where $C \in \mathbb{R}^{n \times k}$ is the scaled cluster indicator matrix such that $C^{\top} C=I$.

Constrained Low-rank Approximation (Cohen et al., 2015)

$$
\min _{P \in \mathcal{S}}\|X-P X\|_{F}^{2}
$$

where $\mathcal{S}=Q Q^{\top}$ is any set of rank $k$ orthogonal projection matrix with orthonormal $Q \in \mathbb{R}^{n \times k}$
Low-rank Approximation: $\mathcal{S}$ is the set of all rank $k$ orthogonal projection matrix. $P^{*}=U_{k} U_{k}^{\top}$

## Randomized Algorithms for K-means Clustering

Define

$$
\begin{aligned}
\widehat{P}^{*} & =\min _{P \in \mathcal{S}}\|\widehat{X}-P \widehat{X}\|_{F}^{2} \\
P^{*} & =\min _{P \in \mathcal{S}}\|X-P X\|_{F}^{2}
\end{aligned}
$$

Guarantees on Approximation

$$
\left\|X-\widehat{P}^{*} X\right\|_{F}^{2} \leq \frac{1+\epsilon}{1-\epsilon}\left\|X-P^{*} X\right\|_{F}^{2}
$$

## Properties of Leverage-score sampling

We prove the properties using Matrix Chernoff bound. Let $\Omega=A U$.

$$
\Omega^{\top} \Omega=(A U)^{\top}(A U)=\sum_{j=1}^{m} \frac{1}{m p_{i_{j}}} \mathbf{u}_{i_{j}} \mathbf{u}_{i_{j}}^{\top}
$$

Let $X_{i}=\frac{1}{m p_{i}} \mathbf{u}_{i} \mathbf{u}_{i}^{\top} . \mathrm{E}\left[X_{i}\right]=\frac{1}{m} I_{k}$. Therefore $\lambda_{\max }\left(X_{i}\right)=\lambda_{\text {min }}\left(X_{i}\right)=\frac{1}{m}$.
And $\lambda_{\max }\left(X_{i}\right) \leq \max _{i} \frac{\left\|\mathbf{u}_{i}\right\|_{2}^{2}}{m p_{i}}=\frac{k}{m}$. Applying the Matrix Chernoff bound for the minimum and maximum eigen-value, we have

$$
\begin{aligned}
& \operatorname{Pr}\left(\lambda_{\min }\left(\Omega^{\top} \Omega\right) \leq(1-\epsilon)\right) \leq k \exp \left(-\frac{m \epsilon^{2}}{2 k}\right) \leq k \exp \left(-\frac{m \epsilon^{2}}{3 k}\right) \\
& \operatorname{Pr}\left(\lambda_{\max }\left(\Omega^{\top} \Omega\right) \geq(1+\epsilon)\right) \leq k \exp \left(-\frac{m \epsilon^{2}}{3 k}\right)
\end{aligned}
$$

## When uniform sampling makes sense?

Coherence measure

$$
\mu_{k}=\frac{d}{k} \max _{1 \leq i \leq d}\left\|U_{i *}\right\|_{2}^{2}
$$

When $\mu_{k} \leq \tau$ and $m=\Theta\left(\frac{k \tau}{\epsilon^{2}} \log \left[\frac{2 k}{\delta}\right]\right)$ w.h.p $1-\delta$,

- $A$ formed by uniform sampling (and scaling)
- $A U \in \mathbb{R}^{m \times k}$ is full column rank
- $\sigma_{i}^{2}(A U) \geq(1-\epsilon) \geq(1-\epsilon)^{2}$
- $\sigma_{i}^{2}(A U) \leq(1+\epsilon) \leq(1+\epsilon)^{2}$
- Valid when the coherence measure is small (some real data mining datasets have small coherence measures)
- The Nyström method usually uses uniform sampling (Gittens, 2011)

