# Large Scale Learning and Optimization 

## Aymeric DIEULEVEUT

CMAP, Polytechnique (X)
October 3-12, 2019

Autumn School on Machine Learning - Tbilisi


Slides on my web page: www.cmap.polytechnique.fr/~aymeric.dieuleveut/

## Outline

1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

## Outline

1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient

## Algorithms

3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer
RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

## Large scale learning: multiple contexts and applications

What happened over the last 20 years?

1. Increase in computational power

2 Data everywhere $\rightarrow$ learning from examples.
2 New algorithms, new models

Large scale framework:
Data increase in number $\boldsymbol{n}$ and quality/dimension $\boldsymbol{d}$.

## New Applications: Translation

$\equiv$ Google Translate
$\overline{\text { x }}_{\text {A }}$ Text
Documents

ENGLISH - DETECTED ENGLISH
FRENCH
SPANISI $\vee \stackrel{\text { GEORGIAN }}{ }$
FRENCH
ENGLISH
Hi , thank you for inviting me!

4 (1) 30/5000

NLP tasks:

1. Words representations, sentence representations, etc.
2. Automatic translation
3. Text generation, ...
number $n$ : billions of observations (wikipedia)
features dimension $d$ : high dimensional representations of words

## Advertisement


number $n$ : billions of people
features dimension $d$ : cookies, clicks

## Bio-informatics



Bio-informatics
Input: DNA/RNA sequence,
Output: Drug responsiveness
number n: not always many patients
features dimension $d$ : e.g., number of basis $\rightarrow 10^{6}$.

## Image recognition

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |

Image classification
Input: Images, Videos
Output: Digit , more complex category, action recognition...
number $n$ : millions of images
features dimension $d$ : millions of pixels, potentially thousand of frames in short video.

## Large scale learning : Tons of applications, fewer algorithms \& frameworks



- Sometimes combine supervised + unsupervised
- Many methods for each domain. For example for regression: Nearest neighbours, Linear regression, Kernel Regression, etc.
- Why is optimization about?


## Optimization is a key tool for large scale learning.

What is optimization about?

## $\min _{\theta \in \Theta} f(\theta)$

With $\theta$ a parameter, and $f$ a cost function.

Why?
We formulate our problem as a cost minimization problem.
A few examples:

- Supervised machine learning
- Signal Processing
- Optimal transport
- GANS


## Optimization: some Examples $1 / 4$

## Example 1: Supervised Machine Learning

Consider an input/output pair $(X, Y) \in \mathcal{X} \times \mathcal{Y},(X, Y) \sim \rho$.
Goal: function $\theta: \mathcal{X} \rightarrow \mathbb{R}$, s.t. $\theta(X)$ good prediction for $Y$. Here, as a linear function $\langle\theta, \Phi(X)\rangle$ of features $\boldsymbol{\Phi}(\boldsymbol{X}) \in \mathbb{R}^{\boldsymbol{d}}$.

Consider a loss function $\ell: \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}_{+}$
Define the Generalization risk :

$$
\mathcal{R}(\theta):=\mathbb{E}_{\rho}[\ell(Y,\langle\theta, \Phi(X)\rangle)]
$$

## Empirical Risk minimization (I)

Data: $n$ observations $\left(x_{i}, y_{i}\right) \in \mathcal{X} \times \mathcal{Y}, i=1, \ldots, n$, i.i.d.
Empirical risk (or training error):

$$
\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i},\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)
$$

Empirical risk minimization (ERM) : find $\hat{\boldsymbol{\theta}}$ solution of

$$
\min _{\theta \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(\boldsymbol{y}_{\boldsymbol{i}},\left\langle\boldsymbol{\theta}, \Phi\left(\boldsymbol{x}_{\boldsymbol{i}}\right)\right\rangle\right)+\mu \Omega(\theta)
$$

convex data fitting term + regularizer

## Empirical Risk minimization (II)

For example, least-squares regression:

$$
\min _{\theta \in \mathbb{R}^{d}} \frac{1}{2 \boldsymbol{n}} \sum_{i=1}^{\boldsymbol{n}}\left(\boldsymbol{y}_{\boldsymbol{i}}-\left\langle\boldsymbol{\theta}, \boldsymbol{\Phi}\left(\boldsymbol{x}_{\boldsymbol{i}}\right)\right\rangle\right)^{2}+\mu \Omega(\theta)
$$

and logistic regression:

$$
\min _{\theta \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i}\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)\right)+\mu \Omega(\theta) .
$$

## Optimization: some Examples 2/4

## Example 2: Signal processing

Observe a signal $\boldsymbol{Y} \in \mathbb{R}^{\boldsymbol{n} \times \boldsymbol{q}}$, try to recover the source $B \in \mathbb{R}^{p \times \boldsymbol{q}}$, knowing the "forward matrix" $X \in \mathbb{R}^{\boldsymbol{n} \times p}$. (multi-task regression)

$$
\min _{\beta}\|X \beta-Y\|_{F}^{2}
$$

$\Omega$ sparsity inducing regularization.

How to choose $\lambda$ ?
$\rightarrow$ non smooth optimization, optimization with sparsity inducing norms, etc.

## Optimization: some Examples 3/4

## Example 3: Optimal transport

$$
\min _{\pi \in \Pi} \int c(x, y) \mathrm{d} \pi(x, y)
$$

$\Pi$ set of probability distributions $c(x, y)$ "distance" from $x$ to $\boldsymbol{y}$.

+ regularization

Kantorovic formulation of OT.
$\leftrightarrow$ alternating directions algorithms, ....

## Optimization: some Examples 4/4

GANS

$$
\min _{G} \max _{D}\left\{\mathbb{E}_{x \sim p_{\text {data }}}[\log D(x)]+\mathbb{E}_{z \sim p_{z}}[\log (1-D(G(z))]\}\right.
$$

- Discriminator: tries to discriminate between real and fake images
- G generator: tries to fool the discriminator.
$\leftrightarrow$ minimax optimization, non convex optimization....
- Optimization is at the heart of most Learning methods.
- Is it difficult ?


## Is it a (hard) problem?

for convex optimization, in $99 \%$ of the cases, no.
In the words of Steven Boyd:


## What makes it hard: 1. Convexity

## Why?





Typical non-convex problems:
Empirical risk minimization with 0-1 loss.
$\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} 1_{y_{i} \neq \operatorname{sign}\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle}$.


Neural networks: parametric non-convex functions.

## What makes it hard: 2. Regularity of the function

## a. Smoothness

- A function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is $L$-smooth if and only if it is twice differentiable and

$$
\forall \theta \in \mathbb{R}^{d}, \text { eigenvalues }\left[g^{\prime \prime}(\theta)\right] \leqslant L
$$



For all $\boldsymbol{\theta} \in \mathbb{R}^{\boldsymbol{d}}$ :

$$
f(\theta) \leq f\left(\theta^{\prime}\right)+\left\langle f\left(\theta^{\prime}\right), \theta-\theta^{\prime}\right\rangle+\frac{L}{2}\left\|\theta-\theta^{\prime}\right\|^{2}
$$

## What makes it hard: 2. Regularity of the function

## b. Strong Convexity

- A twice differentiable function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is $\mu$-strongly convex if and only if

$$
\forall \theta \in \mathbb{R}^{d}, \text { eigenvalues }\left[f^{\prime \prime}(\theta)\right] \geqslant \mu
$$



For all $\theta \in \mathbb{R}^{d}$ :

$$
f(\theta) \geq f\left(\theta^{\prime}\right)+\left\langle f\left(\theta^{\prime}\right), \theta-\theta^{\prime}\right\rangle+\frac{\mu}{2}\left\|\theta-\theta^{\prime}\right\|^{2}
$$

## What makes it hard: 2. Regularity of the function

Why?
Rates typically depend on the condition number $\kappa=\frac{L}{\mu}$ :


Large $\kappa$


Small $\kappa$
harder to optimize easier to optimize

## Smoothness and strong convexity in ML

We consider an a.s. convex loss in $\theta$. Thus $\hat{\mathcal{R}}$ and $\mathcal{R}$ are convex. Hessian of $\hat{\mathcal{R}} \approx$ covariance matrix $\frac{1}{n} \sum_{i=1}^{n} \Phi\left(x_{i}\right) \boldsymbol{\Phi}\left(x_{i}\right)^{\top}$ If $\ell$ is smooth, and $\mathbb{E}\left[\|\Phi(X)\|^{2}\right] \leq r^{2}, \mathcal{R}$ is smooth.

If $\ell$ is $\mu$-strongly convex, and data has an invertible covariance matrix (low correlation/dimension), $\mathcal{R}$ is strongly convex.

Importance of regularization: provides strong convexity, and avoids overfitting.

Note: when considering dual formulation of the problem:

- L-smoothness $\leftrightarrow 1 /$-strong convexity.
- $\mu$-strong convexity $\leftrightarrow 1 / \mu$-smoothness


## What makes it hard: 3. Set $\Theta$, complexity of $f$

a. Set $\Theta$ : (if $\boldsymbol{\Theta}$ is a convex set.)

- May be described implicitly (via equations): $\Theta=\left\{\theta \in \mathbb{R}^{d}\right.$ s.t. $\|\theta\|_{2} \leq R$ and $\left.\langle\theta, 1\rangle=r\right\}$. $\rightarrow$ Use dual formulation of the problem.
- Projection might be difficult or impossible.
- Even when $\Theta=\mathbb{R}^{\boldsymbol{d}}, \boldsymbol{d}$ might be very large (typically millions)
$\rightarrow$ use only first order methods
b. Structure of $f$. If $f=\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i},\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)$,
computing a gradient has a cost proportional to $n$.


## Optimization

## Take home

- We express problems as minimizing a function over a set
- Many convex problems are solved
- Difficulties come from non-convexity, lack of regularity, complexity of the set $\Theta$, complexity of computing gradients

Our focus in this course:

- Supervised Machine Learning.
- Stochastic algorithms.

Goals:

- present algorithms (convex, large dimension, high number of observations)
- show how rates depend on smoothness and strong convexity
- show how we can use the structure
- not forgetting the initial problem: Generalization properties


## Roadmap

2. Classical rates for deterministic methods

## Algorithms

3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer
RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

## Goals:

1. Rates
2. Proof techniques

Some slides from this section are adapted from Francis Bach's lecture in Orsay.

## Classical rates for deterministic methods

- Assumption: $f$ convex on $\mathbb{R}^{d}$
- Classical generic algorithms
- Gradient descent and accelerated gradient descent
- Newton method
- Subgradient method (and ellipsoid algorithm)
- Key additional properties of $\boldsymbol{f}$
- Lipschitz continuity, smoothness or strong convexity
- Key references: Nesterov (2004), Bubeck (2015)


## Several criteria for characterizing convergence

- Objective function values

$$
f(\theta)-\inf _{\eta \in \mathbb{R}^{d}} f(\eta)
$$

- Usually weaker condition
- Iterates

$$
\inf _{\eta \in \arg \min f}\|\theta-\eta\|^{2}
$$

- Typically used for strongly-convex problems
- NB 1: relationships between the two types in several situations
- NB 2: similarity with prediction vs. estimation in statistics


## Toolbox

We use a lot a few very useful inequalities.
Convex: the function is above the tangent line:


$$
\begin{equation*}
f(\theta) \geq f\left(\theta^{\prime}\right)+\left\langle f\left(\theta^{\prime}\right), \theta-\theta^{\prime}\right\rangle \tag{1}
\end{equation*}
$$

## Toolbox

We use a lot a few very useful inequalities.
Strongly-convex: function above the tangent line $+\mu^{*}$ quadratic.


$$
\begin{gather*}
f(\theta) \geq f\left(\theta^{\prime}\right)+\left\langle f\left(\theta^{\prime}\right), \theta-\theta^{\prime}\right\rangle+\frac{\mu}{2}\left\|\theta-\theta^{\prime}\right\|^{2}  \tag{2}\\
\left\langle f^{\prime}\left(\theta^{\prime}\right)-f^{\prime}(\theta), \theta^{\prime}-\theta\right\rangle \geq \mu\left\|\theta-\theta^{\prime}\right\|^{2} \tag{3}
\end{gather*}
$$

## Toolbox

We use a lot a few very useful inequalities. Smooth-convex: function below the tangent line $+L^{*}$ quadratic.


$$
\begin{equation*}
f(\theta) \leq f\left(\theta^{\prime}\right)+\left\langle f\left(\theta^{\prime}\right), \theta-\theta^{\prime}\right\rangle+\frac{L}{2}\left\|\theta-\theta^{\prime}\right\|^{2} \tag{4}
\end{equation*}
$$

Co-coercivity:

$$
\begin{equation*}
\left\|f^{\prime}(\theta)-f^{\prime}\left(\theta^{\prime}\right)\right\|^{2} \leq L\left\langle f\left(\theta^{\prime}\right)-f\left(\theta^{\prime}\right), \theta-\theta^{\prime}\right\rangle \tag{5}
\end{equation*}
$$

## Toolbox

3 Starting Points:

1. Expand $\left\|\theta_{t+1}-\theta_{*}\right\|^{2}$ "Lyapunov approach"
2. Expand $f\left(\theta_{t+1}\right)-f\left(\theta_{t}\right)$ (if smooth!)
3. Expand $\theta_{t+1}-\theta_{t}$

## (smooth) Gradient Descent

- Assumptions
- $f$ convex with L-Lipschitz-continuous gradient (e.g., L-smooth)
- Algorithm:

$$
\theta_{t}=\theta_{t-1}-\frac{1}{L} g^{\prime}\left(\theta_{t-1}\right)
$$



## (smooth) Gradient Descent - strong convexity

- Assumptions
- $f$ convex with L-Lipschitz-continuous gradient (e.g., L-smooth)
- f $\mu$-strongly convex
- Algorithm:

$$
\theta_{t}=\theta_{t-1}-\frac{1}{L} f^{\prime}\left(\theta_{t-1}\right)
$$

- Bound:

$$
f\left(\theta_{t}\right)-f\left(\theta_{*}\right) \leqslant(1-\mu / L)^{t}\left[f\left(\theta_{0}\right)-f\left(\theta_{*}\right)\right]
$$

- Three-line proof. Challenge 1 ! (start from $\left(\left\|\theta_{t}-\theta_{*}\right\|^{2}\right)$
- Line search, steepest descent or constant step-size

Proof

## (smooth) Gradient Descent - slow rate <br> - Assumptions

- f convex with L-Lipschitz-continuous gradient (e.g., L-smooth)
- Minimum attained at $\theta_{*}$
- Algorithm:

$$
\theta_{t}=\theta_{t-1}-\frac{1}{L} f^{\prime}\left(\theta_{t-1}\right)
$$

- Bound:

$$
f\left(\theta_{t}\right)-f\left(\theta_{*}\right) \leqslant \frac{2 L\left\|\theta_{0}-\theta_{*}\right\|^{2}}{t+4}
$$

- Five-lines proof
- Adaptivity of gradient descent to problem difficulty



## Gradient descent - Proof for quadratic functions

- Quadratic convex function: $f(\theta)=\frac{1}{2} \theta^{\top} \boldsymbol{H} \theta-c^{\top} \theta$
- $\mu$ and $L$ are smallest largest eigenvalues of $H$
- Global optimum $\theta_{*}=H^{-1} c$ (or $H^{\dagger} c$ ) such that $H \theta_{*}=c$
- Gradient descent with $\gamma=1 / L$ :

$$
\begin{aligned}
\theta_{t} & =\theta_{t-1}-\frac{1}{L}\left(H \theta_{t-1}-c\right)=\theta_{t-1}-\frac{1}{L}\left(H \theta_{t-1}-H \theta_{*}\right) \\
\theta_{t}-\theta_{*} & =\left(I-\frac{1}{L} H\right)\left(\theta_{t-1}-\theta_{*}\right)=\left(I-\frac{1}{L} H\right)^{t}\left(\theta_{0}-\theta_{*}\right)
\end{aligned}
$$

- Strong convexity $\mu>0$ : eigenvalues of $\left(I-\frac{1}{L} H\right)^{t}$ in [0, $\left(1-\frac{\mu}{L}\right)^{t}$ ]
- Convergence of iterates: $\left\|\boldsymbol{\theta}_{\boldsymbol{t}}-\boldsymbol{\theta}_{*}\right\|^{2} \leqslant(1-\mu / L)^{2 t}\left\|\boldsymbol{\theta}_{0}-\boldsymbol{\theta}_{*}\right\|^{2}$
- Function values: $\boldsymbol{f}\left(\boldsymbol{\theta}_{\boldsymbol{t}}\right)-\boldsymbol{f}\left(\boldsymbol{\theta}_{*}\right) \leqslant(1-\mu / L)^{2 t}\left[\boldsymbol{f}\left(\boldsymbol{\theta}_{\mathbf{0}}\right)-\boldsymbol{f}\left(\boldsymbol{\theta}_{*}\right)\right]$


## Gradient descent - Proof for quadratic functions

- Quadratic convex function: $\boldsymbol{f}(\boldsymbol{\theta})=\frac{1}{2} \boldsymbol{\theta}^{\top} \boldsymbol{H} \boldsymbol{\theta}-\boldsymbol{c}^{\top} \boldsymbol{\theta}$
- $\mu$ and $L$ are smallest largest eigenvalues of $H$
- Global optimum $\theta_{*}=H^{-1} c$ (or $H^{\dagger} c$ ) such that $H \theta_{*}=c$
- Gradient descent with $\gamma=1 / L$ :

$$
\begin{aligned}
\theta_{t} & =\theta_{t-1}-\frac{1}{L}\left(H \theta_{t-1}-c\right)=\theta_{t-1}-\frac{1}{L}\left(H \theta_{t-1}-H \theta_{*}\right) \\
\theta_{t}-\theta_{*} & =\left(I-\frac{1}{L} H\right)\left(\theta_{t-1}-\theta_{*}\right)=\left(I-\frac{1}{L} H\right)^{t}\left(\theta_{0}-\theta_{*}\right)
\end{aligned}
$$

- Convexity $\mu=0$ : eigenvalues of $\left(I-\frac{1}{L} H\right)^{t}$ in $[0,1]$
- No convergence of iterates: $\left\|\theta_{\boldsymbol{t}}-\boldsymbol{\theta}_{*}\right\|^{2} \leqslant\left\|\theta_{0}-\boldsymbol{\theta}_{*}\right\|^{2}$
- Function values:

$$
\begin{aligned}
f\left(\theta_{t}\right)-f\left(\theta_{*}\right) & \leqslant \max _{v \in[0, L]} v(1-v / L)^{2 t}\left\|\theta_{0}-\theta_{*}\right\|^{2} \\
& \leqslant \frac{L}{t}\left\|\theta_{0}-\theta_{*}\right\|^{2}
\end{aligned}
$$

## Accelerated gradient methods (Nesterov, 1983)

- Assumptions $\boldsymbol{f}$ convex and smooth $L$
- Algorithm:

$$
\begin{aligned}
\theta_{t} & =\eta_{t-1}-\frac{1}{L} f^{\prime}\left(\eta_{t-1}\right) \\
\eta_{t} & =\theta_{t}+\frac{t-1}{t+2}\left(\theta_{t}-\theta_{t-1}\right)
\end{aligned}
$$



- Bound:

$$
f\left(\theta_{t}\right)-f\left(\theta_{*}\right) \leqslant \frac{2 L\left\|\theta_{0}-\theta_{*}\right\|^{2}}{(t+1)^{2}}
$$

- Ten-line proof (see, e.g., Schmidt et al., 2011)
- Not improvable
- Extension to strongly-convex functions


## Accelerated gradient methods - strong convexity

- Assumptions
- $\boldsymbol{f}$ convex with L-Lipschitz-cont. gradient , min. attained at $\theta_{*}$
- f $\mu$-strongly convex
- Algorithm:

$$
\begin{aligned}
\theta_{t} & =\eta_{t-1}-\frac{1}{L} f^{\prime}\left(\eta_{t-1}\right) \\
\eta_{t} & =\theta_{t}+\frac{1-\sqrt{\mu / L}}{1+\sqrt{\mu / L}}\left(\theta_{t}-\theta_{t-1}\right)
\end{aligned}
$$

- Bound: $\boldsymbol{f}\left(\theta_{t}\right)-\boldsymbol{f}\left(\theta_{*}\right) \leqslant L\left\|\theta_{0}-\theta_{*}\right\|^{2}(1-\sqrt{\mu / L})^{t}$
- Ten-line proof (see, e.g., Schmidt et al., 2011)
- Not improvable
- Relationship with conjugate gradient for quadratic functions

Proof in the quadratic setting: compute the largest eigenvalue of a non-symmetric matrix. Challenge 2! Simple an insightful computation!

## Other methods: Projected gradient descent

- Problems of the form:

```
min
- \(\boldsymbol{\theta}_{\boldsymbol{t}+\boldsymbol{1}}=\arg \min _{\boldsymbol{\theta} \in \mathcal{K}} \boldsymbol{f}\left(\boldsymbol{\theta}_{\boldsymbol{t}}\right)+\left(\boldsymbol{\theta}-\boldsymbol{\theta}_{\boldsymbol{t}}\right)^{\top} \boldsymbol{\nabla} \boldsymbol{f}\left(\boldsymbol{\theta}_{\boldsymbol{t}}\right)+\frac{L}{2}\left\|\theta-\theta_{t}\right\|_{2}^{2}\)
- \(\theta_{t+1}=\arg \min _{\theta \in \mathcal{K}} \frac{1}{2}\left\|\theta-\left(\theta_{t}-\frac{1}{L} \nabla f\left(\theta_{t}\right)\right)\right\|_{2}^{2}\)
- Projected gradient descent
- Similar convergence rates than smooth optimization
- Acceleration methods (Nesterov, 2007; Beck and Teboulle, 2009)

\section*{Other methods: Newton method}
- Given \(\theta_{t-1}\), minimize second-order Taylor expansion
\[
\begin{aligned}
& \tilde{f}(\theta)=f\left(\theta_{t-1}\right)+f^{\prime}\left(\theta_{t-1}\right)^{\top}\left(\theta-\theta_{t-1}\right) \\
& \quad+\frac{1}{2}\left(\theta-\theta_{t-1}\right)^{\top} f^{\prime \prime}\left(\theta_{t-1}\right)^{\top}\left(\theta-\theta_{t-1}\right)
\end{aligned}
\]
- Expensive Iteration: \(\theta_{t}=\theta_{t-1}-f^{\prime \prime}\left(\theta_{t-1}\right)^{-1} f^{\prime}\left(\theta_{t-1}\right)\)
- Running-time complexity: \(O\left(d^{3}\right)\) in general
- Quadratic convergence: If \(\left\|\theta_{\boldsymbol{t}-1}-\boldsymbol{\theta}_{*}\right\|\) small enough, for some constant \(C\), we have
\[
\left(C\left\|\theta_{t}-\theta_{*}\right\|\right)=\left(C\left\|\theta_{t-1}-\theta_{*}\right\|\right)^{2}
\]
- See Boyd and Vandenberghe (2003)

\section*{Summary: minimizing smooth convex functions}
- Assumption: \(f\) convex
- Gradient descent: \(\theta_{t}=\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\)
- \(O(1 / t)\) convergence rate for smooth convex functions
- \(O\left(e^{-t \mu / L}\right)\) convergence rate for strongly smooth convex functions
- Optimal rates \(O\left(1 / t^{2}\right)\) and \(O\left(e^{-t \sqrt{\mu / L}}\right)\) with FOI.
- Newton method: \(\theta_{t}=\theta_{t-1}-f^{\prime \prime}\left(\theta_{t-1}\right)^{-1} f^{\prime}\left(\theta_{t-1}\right)\)
- \(O\left(e^{-\rho 2^{t}}\right)\) convergence rate
- From smooth to non-smooth
- Subgradient method

\section*{Subgradient method/"descent" (Shor et al., 1985)}
- Assumptions
- \(\boldsymbol{f}\) convex and \(B\)-Lipschitz-continuous on \(\left\{\|\boldsymbol{\theta}\|_{2} \leqslant \boldsymbol{D}\right\}\)
- Algorithm: \(\theta_{t}=\Pi_{D}\left(\theta_{t-1}-\frac{2 D}{B \sqrt{t}} f^{\prime}\left(\theta_{t-1}\right)\right)\)
- \(\Pi_{D}\) : orthogonal projection onto \(\left\{\|\theta\|_{2} \leqslant D\right\}\)


\section*{Subgradient method/"descent" (Shor et al., 1985)}
- Assumptions
- \(f\) convex and \(B\)-Lipschitz-continuous on \(\left\{\|\theta\|_{2} \leqslant D\right\}\)
- Algorithm: \(\theta_{t}=\Pi_{D}\left(\theta_{t-1}-\frac{2 D}{B \sqrt{t}} f^{\prime}\left(\theta_{t-1}\right)\right)\)
- \(\Pi_{D}\) : orthogonal projection onto \(\left\{\|\theta\|_{2} \leqslant D\right\}\)
- Bound:
\[
f\left(\frac{1}{t} \sum_{k=0}^{t-1} \theta_{k}\right)-f\left(\theta_{*}\right) \leqslant \frac{2 D B}{\sqrt{t}}
\]
- Three-line proof
- Best possible convergence rate after \(O(d)\) iterations (Bubeck, 2015)

\section*{Need for decaying steps}

Example of \(|x|\)

\section*{Subgradient method/"descent" - proof - I}
- Iteration: \(\theta_{t}=\Pi_{D}\left(\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\right)\) with \(\gamma_{t}=\frac{2 D}{B \sqrt{t}}\)
- Assumption: \(\left\|f^{\prime}(\theta)\right\|_{2} \leqslant B\) and \(\|\theta\|_{2} \leqslant D\)
\(\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2} \leqslant\left\|\theta_{t-1}-\theta_{*}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\right\|_{2}^{2}\) by contractivity of projections
\[
\begin{aligned}
& =\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}+\gamma_{t}^{2}\left\|f^{\prime}\left(\theta_{t-1}\right)\right\|_{2}^{2}-2 \gamma_{t}\left(\theta_{t-1}-\theta_{*}\right)^{\top} g^{\prime}\left(\theta_{t-1}\right) \\
& \leqslant\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}+B^{2} \gamma_{t}^{2}-2 \gamma_{t}\left(\theta_{t-1}-\theta_{*}\right)^{\top} f^{\prime}\left(\theta_{t-1}\right) \\
& \quad \text { because }\left\|f^{\prime}\left(\theta_{t-1}\right)\right\|_{2} \leqslant B \\
& \leqslant\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}+B^{2} \gamma_{t}^{2}-2 \gamma_{t}\left[f\left(\theta_{t-1}\right)-\boldsymbol{f}\left(\theta_{*}\right)\right]
\end{aligned}
\]
(property of subgradients)
- leading to
\[
f\left(\theta_{t-1}\right)-f\left(\theta_{*}\right) \leqslant \frac{B^{2} \gamma_{t}}{2}+\frac{1}{2 \gamma_{t}}\left[\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2}\right]
\]

\section*{Subgradient method/"descent" - proof - II}
- Starting from
\[
f\left(\theta_{t-1}\right)-f\left(\theta_{*}\right) \leqslant \frac{B^{2} \gamma_{t}}{2}+\frac{1}{2 \gamma_{t}}\left[\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2}\right]
\]
- Constant step-size \(\gamma_{t}=\gamma\)
\[
\begin{aligned}
\sum_{u=1}^{t}\left[f\left(\theta_{u-1}\right)-f\left(\theta_{*}\right)\right] & \leqslant \sum_{u=1}^{t} \frac{B^{2} \gamma}{2}+\sum_{u=1}^{t} \frac{1}{2 \gamma}\left[\left\|\theta_{u-1}-\theta_{*}\right\|_{2}^{2}-\left\|\theta_{u}-\theta_{*}\right\|_{2}^{2}\right] \\
& \leqslant t \frac{B^{2} \gamma}{2}+\frac{1}{2 \gamma}\left\|\theta_{0}-\theta_{*}\right\|_{2}^{2} \leqslant t \frac{B^{2} \gamma}{2}+\frac{2}{\gamma} D^{2}
\end{aligned}
\]
- Optimized step-size \(\gamma_{t}=\frac{2 D}{B \sqrt{t}}\) depends on "horizon" t
- Leads to bound of \(2 D B \sqrt{t}\)
- Slightly more complex proof for online setting (decreasing steps)
- Using convexity:
\[
f\left(\frac{1}{t} \sum_{k=0}^{t-1} \theta_{k}\right)-f\left(\theta_{*}\right) \leqslant \frac{1}{t} \sum_{k=0}^{t-1} f\left(\theta_{k}\right)-f\left(\theta_{*}\right) \leqslant \frac{2 D B}{\sqrt{t}}
\]

\section*{Subgradient method/"descent" - proof - III}
- Starting from
\[
f\left(\theta_{t-1}\right)-f\left(\theta_{*}\right) \leqslant \frac{B^{2} \gamma_{t}}{2}+\frac{1}{2 \gamma_{t}}\left[\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2}\right]
\]
- Decreasing step-size
\[
\begin{aligned}
\sum_{u=1}^{t}\left[f\left(\theta_{u-1}\right)-f\left(\theta_{*}\right)\right] & \leqslant \sum_{u=1}^{t} \frac{B^{2} \gamma_{u}}{2}+\sum_{u=1}^{t} \frac{1}{2 \gamma_{u}}\left[\left\|\theta_{u-1}-\theta_{*}\right\|_{2}^{2}-\left\|\theta_{u}-\theta_{*}\right\|_{2}^{2}\right] \\
& =\sum_{u=1}^{t} \frac{B^{2} \gamma_{u}}{2}+\sum_{u=1}^{t-1}\left\|\theta_{u}-\theta_{*}\right\|_{2}^{2}\left(\frac{1}{2 \gamma_{u+1}}-\frac{1}{2 \gamma_{u}}\right)+\frac{\| \theta_{0}-\theta_{*}}{2 \gamma_{1}} \\
& \leqslant \sum_{u=1}^{t} \frac{B^{2} \gamma_{u}}{2}+\sum_{u=1}^{t-1} 4 D^{2}\left(\frac{1}{2 \gamma_{u+1}}-\frac{1}{2 \gamma_{u}}\right)+\frac{4 D^{2}}{2 \gamma_{1}} \\
& =\sum_{u=1}^{t} \frac{B^{2} \gamma_{u}}{2}+\frac{4 D^{2}}{2 \gamma_{t}} \leqslant 3 D B \sqrt{t} \text { with } \gamma_{t}=\frac{2 D}{B \sqrt{t}}
\end{aligned}
\]
- Using convexity: \(f\left(\frac{1}{t} \sum_{k=0}^{t-1} \theta_{k}\right)-f\left(\theta_{*}\right) \leqslant \frac{3 D B}{\sqrt{t}}\)

\section*{Subgradient descent - strong convexity}
- Assumptions
- \(f\) convex and \(B\)-Lipschitz-continuous on \(\left\{\|\theta\|_{2} \leqslant D\right\}\)
- f \(\mu\)-strongly convex
- Algorithm: \(\theta_{t}=\Pi_{D}\left(\theta_{t-1}-\frac{2}{\mu(t+1)} f^{\prime}\left(\theta_{t-1}\right)\right)\)
- Bound:
\[
f\left(\frac{2}{t(t+1)} \sum_{k=1}^{t} k \theta_{k-1}\right)-f\left(\theta_{*}\right) \leqslant \frac{2 B^{2}}{\mu(t+1)}
\]
- Three-line proof
- Best possible convergence rate after \(O(d)\) iterations (Bubeck, 2015)

\section*{Subgradient method - strong convexity - proof - I}
- Iteration: \(\theta_{t}=\Pi_{D}\left(\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\right)\) with \(\gamma_{t}=\frac{2}{\mu(t+1)}\)
- Assumption: \(\left\|f^{\prime}(\theta)\right\|_{2} \leqslant B\) and \(\|\theta\|_{2} \leqslant D\) and \(\mu\)-strong convexity of \(f\)
\[
\left\|\boldsymbol{\theta}_{\boldsymbol{t}}-\boldsymbol{\theta}_{*}\right\|_{2}^{2} \leqslant\left\|\boldsymbol{\theta}_{t-1}-\boldsymbol{\theta}_{*}-\gamma_{t} \boldsymbol{f}^{\prime}\left(\boldsymbol{\theta}_{\boldsymbol{t}-1}\right)\right\|_{2}^{2}
\]
by contractivity of projections
\[
\begin{aligned}
& \leqslant\left\|\boldsymbol{\theta}_{\boldsymbol{t}-1}-\boldsymbol{\theta}_{*}\right\|_{2}^{2}+\boldsymbol{B}^{2} \gamma_{\boldsymbol{t}}^{2}-\mathbf{2} \gamma_{\boldsymbol{t}}\left(\boldsymbol{\theta}_{\boldsymbol{t}-\mathbf{1}}-\boldsymbol{\theta}_{*}\right)^{\top} \boldsymbol{f}^{\prime}\left(\boldsymbol{\theta}_{\boldsymbol{t}-1}\right) \\
& \quad \text { because }\left\|\boldsymbol{f}^{\prime}\left(\theta_{\boldsymbol{t}-1}\right)\right\|_{2} \leqslant \boldsymbol{B} \\
& \leqslant\left\|\boldsymbol{\theta}_{\boldsymbol{t}-1}-\boldsymbol{\theta}_{*}\right\|_{2}^{2}+\boldsymbol{B}^{2} \gamma_{\boldsymbol{t}}^{2}-\mathbf{2} \gamma_{\boldsymbol{t}}\left[\boldsymbol{f}\left(\boldsymbol{\theta}_{\boldsymbol{t}-\mathbf{1}}\right)-\boldsymbol{f}\left(\boldsymbol{\theta}_{*}\right)+\frac{\mu}{2}\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}\right]
\end{aligned}
\]
(property of subgradients and strong convexity)
\(\rightarrow\) leading to
\[
\begin{aligned}
f\left(\theta_{t-1}\right)-f\left(\theta_{*}\right) & \leqslant \frac{B^{2} \gamma_{t}}{2}+\frac{1}{2}\left[\frac{1}{\gamma_{t}}-\mu\right]\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\frac{1}{2 \gamma_{t}}\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2} \\
& \leqslant \frac{B^{2}}{\mu(t+1)}+\frac{\mu}{2}\left[\frac{t-1}{2}\right]\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\frac{\mu(t+1)}{4}\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2}
\end{aligned}
\]

\section*{Subgradient method - strong convexity- proof - II}
\[
\begin{aligned}
& f\left(\theta_{t-1}\right)-f\left(\theta_{*}\right) \leqslant \frac{B^{2} \gamma_{t}}{2}+\frac{1}{2}\left[\frac{1}{\gamma_{t}}-\mu\right]\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\frac{1}{2 \gamma_{t}}\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2} \\
& \leqslant \frac{B^{2}}{\mu(t+1)}+\frac{\mu}{2}\left[\frac{t-1}{2}\right]\left\|\theta_{t-1}-\theta_{*}\right\|_{2}^{2}-\frac{\mu(t+1)}{4}\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2} \\
& \sum_{u=1}^{t} u\left[f\left(\theta_{u-1}\right)-f\left(\theta_{*}\right)\right] \leqslant \sum_{t=1}^{u} \frac{B^{2} u}{\mu(u+1)}+\frac{1}{4} \sum_{u=1}^{t}\left[u(u-1)\left\|\theta_{u-1}-\theta_{*}\right\|_{2}^{2}\right. \\
&\left.-u(u+1)\left\|\theta_{u}-\theta_{*}\right\|_{2}^{2}\right]
\end{aligned} \quad \begin{aligned}
& \quad \leqslant \frac{B^{2} t}{\mu}+\frac{1}{4}\left[0-t(t+1)\left\|\theta_{t}-\theta_{*}\right\|_{2}^{2}\right] \leqslant \frac{B^{2} t}{\mu}
\end{aligned}
\]
- Using convexity: \(\boldsymbol{f}\left(\frac{2}{t(t+1)} \sum_{u=1}^{t} u \theta_{u-1}\right)-\boldsymbol{f}\left(\theta_{*}\right) \leqslant \frac{2 B^{2}}{t+1}\)
- NB: with step-size \(\gamma_{n}=1 /(n \mu)\), extra logarithmic factor

\section*{Summary: minimizing convex functions}

Gradient descent: \(\theta_{t}=\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\)
Convergence rate ( \(=\) speed of convergence)
\(O(1 / \sqrt{t})\) for non-smooth convex functions
\(O(1 / t) \quad\) for smooth convex functions
\(O\left(e^{-t \mu / L}\right)\) for strongly smooth convex functions

\section*{Summary of rates of convergence}
- Problem parameters
- D diameter of the domain
- B Lipschitz-constant
- L smoothness constant
- \(\mu\) strong convexity constant
\begin{tabular}{lll}
\hline & convex & strongly convex \\
\hline nonsmooth & deterministic: \(B D / \sqrt{t}\) & deterministic: \(B^{2} /(t \mu)\) \\
\hline smooth & deterministic: \(L D^{2} / t^{2}\) & deterministic: \(\exp (-t \sqrt{\mu / L})\) \\
& & \\
\hline quadratic & deterministic: \(L D^{2} / t^{2}\) & deterministic: \(\exp (-t \sqrt{\mu / L})\)
\end{tabular}

\section*{Summary of the first session}
1. Optimizing a cost function is at the heart of Large scale learning
2. Difficulty comes from the fact that both the number of examples \(\mathbf{n}\) and the number of dimensions \(d\) are very large.
First method: Gradient descent: \(\theta_{t}=\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\).


\section*{Summary of the first session}

First method: Gradient descent: \(\theta_{t}=\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)\).
Convergence rate ( \(=\) speed of convergence)
\(O(1 / t) \quad\) for smooth convex functions
\(O\left(e^{-t \mu / L}\right)\) for strongly smooth convex functions
Optimal rates \(O\left(1 / t^{2}\right)\) and \(O\left(e^{-t \sqrt{\mu / L}}\right)\) with acceleration (optimal - not seen).

\section*{Spirit - Goals}

Goals:
1. Understand what SGD is.
2. Comparison to GD (cost, convergence speed)
3. Important variants.

Approach:
1. convergence speed helps to choose between algorithms
2. influence of parameters \(\rightarrow\) choice of paramaters (e.g., step size)
3. proofs help to understand assumptions

\section*{Roadmap}
1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer
RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

\section*{Back to Supervised Machine Learning framework}

\section*{Example 1: Supervised Machine Learning}

Consider an input/output pair \((X, Y) \in \mathcal{X} \times \mathcal{Y},(X, Y) \sim \rho\).
Goal: function \(\theta: \mathcal{X} \rightarrow \mathbb{R}\), s.t. \(\theta(X)\) good prediction for \(Y\). Here, as a linear function \(\langle\theta, \Phi(X)\rangle\) of features \(\boldsymbol{\Phi}(\boldsymbol{X}) \in \mathbb{R}^{\boldsymbol{d}}\).

Consider a loss function \(\ell: \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}_{+}\)
Define the Generalization risk :
\[
\mathcal{R}(\theta):=\mathbb{E}_{\rho}[\ell(Y,\langle\theta, \Phi(X)\rangle)]
\]

\section*{Empirical Risk minimization (I)}

Data: \(n\) observations \(\left(x_{i}, y_{i}\right) \in \mathcal{X} \times \mathcal{Y}, i=1, \ldots, n\), i.i.d.
Empirical risk (or training error):
\[
\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i},\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)
\]

Empirical risk minimization (ERM) : find \(\hat{\boldsymbol{\theta}}\) solution of
\[
\min _{\theta \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i},\left\langle\boldsymbol{\theta}, \Phi\left(x_{i}\right)\right\rangle\right)+\mu \Omega(\theta)
\]
convex data fitting term + regularizer

\section*{Empirical Risk minimization (II)}

For example, least-squares regression:
\[
\min _{\theta \in \mathbb{R}^{d}} \frac{1}{2 \boldsymbol{n}} \sum_{i=1}^{n}\left(\boldsymbol{y}_{\boldsymbol{i}}-\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)^{2}+\mu \Omega(\theta)
\]
and logistic regression:
\[
\min _{\theta \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i}\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)\right)+\mu \Omega(\theta)
\]

\section*{Empirical Risk Minimization (ERM) setting.}
\[
\min _{\theta \in \mathbb{R}^{d}}\left\{\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i},\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)\right\}
\]

Two fundamental questions: (a) computing (b) analyzing \(\hat{\boldsymbol{\theta}}\).
"Large scale" framework: number of examples \(n\) and the number of explanatory variables \(d\) are both large.
1. High dimension \(\boldsymbol{d} \Longrightarrow\) First order algorithms Gradient Descent (GD) :
\[
\theta_{t}=\theta_{t-1}-\gamma_{t} \hat{\mathcal{R}}^{\prime}\left(\theta_{t-1}\right)
\]

Problem: computing the gradient costs \(O(d n)\) per iteration.

\section*{Gradient descent for ERM}
- Assumptions ( \(\mathcal{R}\) is the expected risk, \(\hat{\mathcal{R}}\) the empirical risk)
- \(\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i}, \Phi\left(x_{i}\right)^{\top} \theta\right)\)
- \(\ell\) smooth.
- Cost: At each step, compute
\[
\hat{\mathcal{R}}^{\prime}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell^{\prime}\left(y_{i}, \Phi\left(x_{i}\right)^{\top} \theta\right) \Phi\left(x_{i}\right)
\]
cost \(=n d\) each step
- Convergence: after \(t\) iterations of subgradient method
\[
\hat{\mathcal{R}}\left(\theta_{t}\right)-\min _{\eta \in \Theta} \hat{\mathcal{R}}(\eta) \leqslant \frac{L}{t}
\]
- Summary: for \(t=\sqrt{n}\) iterations, convergence \(L / \sqrt{n}\), with total running-time complexity of \(O\left(n^{3 / 2} d\right)\)

\section*{Empirical Risk Minimization (ERM) setting.}
\[
\min _{\theta \in \mathbb{R}^{d}}\left\{\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i},\left\langle\theta, \Phi\left(x_{i}\right)\right\rangle\right)\right\}
\]

Two fundamental questions: (a) computing (b) analyzing \(\hat{\boldsymbol{\theta}}\).
"Large scale" framework: number of examples \(n\) and the number of explanatory variables \(d\) are both large.
1. High dimension \(\boldsymbol{d} \Longrightarrow\) First order algorithms Gradient Descent (GD) :
\[
\theta_{t}=\theta_{t-1}-\gamma_{t} \hat{\mathcal{R}}^{\prime}\left(\theta_{t-1}\right)
\]

Problem: computing the gradient costs \(O(d n)\) per iteration.
2. Large \(n \Longrightarrow\) Stochastic algorithms

Stochastic Gradient Descent (SGD)

\section*{Idea of SGD}

What is our main problem? computing
\[
\hat{\mathcal{R}}^{\prime}(\theta)=\frac{1}{n} \sum_{i=1}^{n} \ell^{\prime}\left(y_{i}, \Phi\left(x_{i}\right)^{\top} \theta\right) \Phi\left(x_{i}\right)=: \frac{1}{n} \sum_{i=1}^{n} f_{i}^{\prime}(\theta)
\]
costs \(n d\) per iteration
Solution?
Use instead for the gradient just one element of the sum!!
\[
f_{i}^{\prime}(\theta) \quad\left(=\ell^{\prime}\left(y_{i}, \Phi\left(x_{i}\right)^{\top} \theta\right) \Phi\left(x_{i}\right)\right)
\]
with \(i \in \mathcal{U}\{1, \ldots, n\}\).
One observation at each step \(\rightarrow\) complexity \(d\) per iteration.

\section*{SGD for ERM: \(f=\hat{\mathcal{R}}\)}

Loss for a single pair of observations, for any \(\boldsymbol{j} \leq \boldsymbol{n}\) :
\[
f_{j}(\theta):=\ell\left(y_{j},\left\langle\theta, \Phi\left(x_{j}\right)\right\rangle\right)
\]

For the empirical risk \(\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{t=1}^{n} \ell\left(y_{t},\left\langle\theta, \Phi\left(x_{t}\right)\right\rangle\right)\).
- At each step \(t \in \mathbb{N}^{*}\), sample \(I_{t} \sim \mathcal{U}\{1, \ldots n\}:\)
\[
\begin{gathered}
\boldsymbol{f}_{l_{t}}^{\prime}\left(\theta_{t-1}\right)=\ell^{\prime}\left(y_{l_{t}},\left\langle\theta_{t-1}, \Phi\left(x_{l_{t}}\right)\right\rangle\right) \\
\mathbb{E}\left[f_{l_{t}}^{\prime}\left(\theta_{t-1}\right)\right]=\frac{1}{n} \sum_{t=1}^{n} f_{i}^{\prime}\left(\theta_{t-1}\right)=\hat{\mathcal{R}}^{\prime}\left(\theta_{t-1}\right)
\end{gathered}
\]

More generally, let's define SGD for a general function \(\boldsymbol{f}\).

\section*{Stochastic Gradient descent}
- Goal:
\[
\min _{\theta \in \mathbb{R}^{d}} f(\theta)
\]
given unbiased gradient estimates \(\boldsymbol{f}_{\boldsymbol{n}}^{\prime}\)
- \(\theta_{*}:=\operatorname{argmin}_{\mathbb{R}^{d}} f(\theta)\).

\section*{Why is randomness not a problem}

Key insights from Bottou and Bousquet (2008)
1. In machine learning, no need to optimize below statistical error
2. In machine learning, cost functions are averages
3. Testing errors are more important than training errors

Take home SGD is :
1. Necessary in the Large Scale setting (complexity)
2. Well suited to Learning problems !

Convergence ?

\section*{Analysis: behaviour of \(\left(\theta_{n}\right)_{n \geq 0}\)}
\[
\theta_{t}=\theta_{t-1}-\gamma_{t} f_{t}^{\prime}\left(\theta_{t-1}\right)
\]

Importance of the learning rate \(\left(\gamma_{t}\right)_{t \geq 0}\).
For smooth and strongly convex problem, \(\theta_{t} \rightarrow \theta_{*}\) a.s. if
\[
\sum_{t=1}^{\infty} \gamma_{t}=\infty \quad \sum_{t=1}^{\infty} \gamma_{t}^{2}<\infty
\]

\section*{Converges as}

- Limit (variance) scales as \(1 / \mu^{2}\)
- Very sensitive to ill-conditioned problems.
- \(\mu\) generally unknown...

\section*{Polyak Ruppert averaging}

Introduced by Polyak and Juditsky (1992) and Ruppert (1988):
\[
\bar{\theta}_{t}=\frac{1}{t+1} \sum_{i=0}^{t} \theta_{i}
\]
- off line averaging reduces the noise effect.
- on line computing: \(\bar{\theta}_{t+1}=\frac{1}{t+1} \theta_{t+1}+\frac{t}{t+1} \bar{\theta}_{t}\).

\section*{Convex stochastic approximation: convergence}

Known global minimax rates for non-smooth problems
- Strongly convex: \(O\left((\mu t)^{-1}\right)\) Attained by averaged stochastic gradient descent with \(\gamma_{t} \propto(\mu t)^{-1}\)
- Non-strongly convex: \(O\left(t^{-1 / 2}\right)\)

Attained by averaged stochastic gradient descent with \(\gamma_{t} \propto t^{-1 / 2}\)

For smooth problems, use larger steps
- Strongly convex: \(O(\mu t)^{-1}\) for \(\gamma_{t} \propto t^{-1 / 2}\) : adapts to strong convexity.

\section*{Convergence rate for \(f\left(\tilde{\theta}_{t}\right)-f\left(\theta_{*}\right)\), smooth \(f\).}
\begin{tabular}{ccc}
\multicolumn{4}{c}{\(\min \hat{\mathcal{R}}\)} \\
\hline & SGD & GD \\
Convex & \(O\left(\frac{1}{\sqrt{t}}\right)\) & \(O\left(\frac{1}{t}\right)\) \\
Stgly-Cvx & \(O\left(\frac{1}{\mu t}\right)\) & \(O\left(e^{-\mu t}\right)\)
\end{tabular}

\section*{Convergence rate for \(f\left(\tilde{\theta}_{t}\right)-f\left(\theta_{*}\right)\), smooth \(f\).}

\author{
\(\min \hat{\mathcal{R}}\) \\ \section*{SGD} \\ ```
                                    GD
``` \\ Convex \(O\left(\frac{1}{\sqrt{t}}\right) \quad O\left(\frac{1}{t}\right)\) \\ Stgly-Cvx \(\quad O\left(\frac{1}{\mu t}\right) \quad O\left(e^{-\mu t}\right)\)
}
\(\ominus\) Gradient descent update costs \(n\) times as much as SGD update.

Which one to choose?
Can we get best of both worlds?

\section*{Stochastic vs. deterministic methods}
- Batch gradient descent:
\[
\theta_{t}=\theta_{t-1}-\gamma_{t} f^{\prime}\left(\theta_{t-1}\right)=\theta_{t-1}-\frac{\gamma_{t}}{n} \sum_{i=1}^{n} f_{i}^{\prime}\left(\theta_{t-1}\right)
\]

- Stochastic gradient descent: \(\theta_{t}=\theta_{t-1}-\gamma_{t} f_{i(t)}^{\prime}\left(\theta_{t-1}\right)\)


\section*{Comparison of convergence : SGD vs GD}

\section*{Which one to choose?}
1. Depends on the precision we want.


Example: non strongly convex case.
2. If our goal is to get a convergence of \(1 / \sqrt{n}\), then
- Complexity of GD: \(n^{3 / 2} d\)
- Complexity of SGD: nd.
3. If our goal is to get a convergence of \(1 / n^{2}\), then
- Complexity of GD: \(n^{3} d\) ( \(n^{2}\) iterations)
- Complexity of SGD: \(n^{4} d\) ( \(n^{4}\) iterations).

Why one is the most likely in Learning ? (Details later...)

\section*{Take home}
1. SGD is a great algorithm
2. Exactly suited for Large Scale Learning
2.1 Low complexity per iteration
\(2.2 \rightarrow\) rapid convergence to a correct precision
Question 2: Can we get best of both worlds?
1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer
RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

\section*{Methods for finite sum minimization}
- GD: at step \(t\), use \(\frac{1}{n} \sum_{i=0}^{n} f_{i}^{\prime}\left(\theta_{t}\right)\)
- SGD: at step \(t\), sample \(i_{t} \sim \mathcal{U}[1 ; n]\), use \(\boldsymbol{f}_{i_{t}}^{\prime}\left(\theta_{t}\right)\)
- SAG: at step \(t\),
- keep a "full gradient" \(\frac{1}{n} \sum_{i=0}^{n} f_{i}^{\prime}\left(\theta_{t_{i}}\right)\), with \(\theta_{t_{i}} \in\left\{\theta_{1}, \ldots \theta_{t}\right\}\)
- sample \(i_{t} \sim \mathcal{U}[1 ; n]\), use
\[
\frac{1}{n}\left(\sum_{i=0}^{n} \boldsymbol{f}_{i}^{\prime}\left(\theta_{t_{i}}\right)-\boldsymbol{f}_{i_{t}}^{\prime}\left(\theta_{t_{i_{t}}}\right)+\boldsymbol{f}_{i_{t}}^{\prime}\left(\theta_{t}\right)\right)
\]

In other words:
- Keep in memory past gradients of all functions \(f_{i}, i=1, \ldots, n\)
- Random selection \(i(t) \in\{1, \ldots, n\}\) with replacement
- Iteration: \(\theta_{t}=\theta_{t-1}-\frac{\gamma_{t}}{n} \sum_{i=1}^{n} y_{i}^{t}\) with
\[
y_{i}^{t}= \begin{cases}f_{i}^{\prime}\left(\theta_{t-1}\right) & \text { if } i=i(t) \\ y_{i}^{t-1} & \text { otherwise }\end{cases}
\]

\section*{SAG}
- Keep in memory past gradients of all functions \(f_{i}, i=1, \ldots, n\)
- Random selection \(i(t) \in\{1, \ldots, n\}\) with replacement
- Iteration: \(\theta_{t}=\theta_{t-1}-\frac{\gamma_{t}}{n} \sum_{i=1}^{n} y_{i}^{t}\) with
\(y_{i}^{t}= \begin{cases}f_{i}^{\prime}\left(\theta_{t-1}\right) & \text { if } i=i(t) \\ y_{i}^{t-1} & \text { otherwise }\end{cases}\)


\section*{SAG}
- Keep in memory past gradients of all functions \(f_{i}, i=1, \ldots, n\)
- Random selection \(i(t) \in\{1, \ldots, n\}\) with replacement
- Iteration: \(\theta_{t}=\theta_{t-1}-\frac{\gamma_{t}}{n} \sum_{i=1}^{n} y_{i}^{t}\) with
\(y_{i}^{t}= \begin{cases}f_{i}^{\prime}\left(\theta_{t-1}\right) & \text { if } i=i(t) \\ y_{i}^{t-1} & \text { otherwise }\end{cases}\)


\section*{SAG}
- Keep in memory past gradients of all functions \(f_{i}, i=1, \ldots, n\)
- Random selection \(i(t) \in\{1, \ldots, n\}\) with replacement
- Iteration: \(\theta_{t}=\theta_{t-1}-\frac{\gamma_{t}}{n} \sum_{i=1}^{n} y_{i}^{t}\) with
\(y_{i}^{t}= \begin{cases}f_{i}^{\prime}\left(\theta_{t-1}\right) & \text { if } i=i(t) \\ y_{i}^{t-1} & \text { otherwise }\end{cases}\)

\(\leftrightarrow \oplus\) update costs the same as SGD
\(\rightarrow \ominus\) needs to store all gradients \(\boldsymbol{f}_{\boldsymbol{i}}^{\prime}\left(\boldsymbol{\theta}_{\boldsymbol{t} \boldsymbol{i}}\right)\) at "points in the past"

\section*{Variance reduced methods}

Some references:
- SAG Schmidt et al. (2013), SAGA Defazio et al. (2014a)
- SVRG Johnson and Zhang (2013) (reduces memory cost but 2 epochs...)
- FINITO Defazio et al. (2014b)
- S2GD Konečnỳ and Richtárik (2013)...

And many others... See for example Niao He's lecture notes for a nice overview.

\section*{Convergence rate for \(f\left(\tilde{\theta}_{t}\right)-f\left(\theta_{*}\right)\), smooth objective \(f\).}

\author{
\(\min \hat{\mathcal{R}}\) \\ \section*{SGD \\ \\ GD \\ \\ SAG} \\ Convex \(O\left(\frac{1}{\sqrt{t}}\right) \quad O\left(\frac{1}{t}\right)\) \\ Stgly-Cvx \\ \(O\left(\frac{1}{\mu t}\right)\) \\ \(O\left(e^{-\mu t}\right)\) \\ \(O\left(1-\left(\mu \wedge \frac{1}{n}\right)\right)^{t}\) \\ 
}

GD, SGD, SAG (Fig. from Schmidt et al. (2013))

\section*{Summary}

\section*{Take home}
1. Variance reduced algorithms can have both:
- low iteration cost
- fast asymptotic convergence

How precisely do I need to converge?
1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer
RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

\section*{Generalization gap: the overfitting problem ?}

My true goal is to control \(\mathcal{R}\) :
\[
\mathcal{R}(\theta):=\mathbb{E}_{\rho}[\ell(Y,\langle\theta, \Phi(X)\rangle)]
\]

Optimization: after \(t\) iterations of one method
\[
\hat{\mathcal{R}}(\hat{\theta})-\hat{\mathcal{R}}\left(\theta_{*}\right) \leqslant \frac{C}{t^{?}}
\]

Statistics: with probability greater than \(1-\delta\)
\[
\sup _{\theta \in \Theta}|\hat{\mathcal{R}}(\theta)-\mathcal{R}(\theta)| \leqslant \frac{G R D}{\sqrt{n}}\left[2+\sqrt{2 \log \frac{2}{\delta}}\right]
\]

SGD for the generalization risk: \(f=\mathcal{R}\)
SGD: key assumption \(\mathbb{E}\left[f_{n}^{\prime}\left(\theta_{n-1}\right) \mid \mathcal{F}_{n-1}\right]=f^{\prime}\left(\theta_{n-1}\right)\).
For the risk
\[
\mathcal{R}(\theta)=\mathbb{E}_{\rho}[\ell(Y,\langle\theta, \Phi(X)\rangle)]
\]
- At step \(0<\boldsymbol{k} \leq n\), use a new point independent of \(\theta_{k-1}\) :
\[
\boldsymbol{f}_{k}^{\prime}\left(\theta_{k-1}\right)=\ell^{\prime}\left(y_{k},\left\langle\theta_{k-1}, \Phi\left(x_{k}\right)\right\rangle\right)
\]
- For \(0 \leq k \leq n, \mathcal{F}_{k}=\sigma\left(\left(x_{i}, y_{i}\right)_{1 \leq i \leq k}\right)\).
\[
\begin{aligned}
\mathbb{E}\left[f_{k}^{\prime}\left(\theta_{k-1}\right) \mid \mathcal{F}_{k-1}\right] & =\mathbb{E}_{\rho}\left[\ell^{\prime}\left(y_{k},\left\langle\theta_{k-1}, \Phi\left(x_{k}\right)\right\rangle\right) \mid \mathcal{F}_{k-1}\right] \\
& =\mathbb{E}_{\rho}\left[\ell^{\prime}\left(Y,\left\langle\theta_{k-1}, \Phi(X)\right\rangle\right)\right]=\mathcal{R}^{\prime}\left(\theta_{k-1}\right)
\end{aligned}
\]
- Single pass through the data, Running-time \(=O(n d)\),
- "Automatic" regularization.

\section*{SGD for the generalization risk: \(f=\mathcal{R}\)}


\section*{Convergence rate for \(f\left(\tilde{\theta}_{\boldsymbol{k}}\right)-\boldsymbol{f}\left(\boldsymbol{\theta}_{*}\right)\), smooth objective \(f\).}
\begin{tabular}{ccccc} 
& \multicolumn{3}{c}{\(\min \hat{\mathcal{R}}\)} & \(\sin \mathcal{R}\) \\
& SGD & GD & SAG & SGD \\
Convex & \(O\left(\frac{1}{\sqrt{k}}\right)\) & \(O\left(\frac{1}{k}\right)\) & & \(O\left(\frac{1}{\sqrt{k}}\right)\) \\
Stgly-Cvx & \(O\left(\frac{1}{\mu k}\right)\) & \(O\left(e^{-\mu k}\right)\) & \(O\left(1-\left(\mu \wedge \frac{1}{n}\right)\right)^{k}\) & \(O\left(\frac{1}{\mu k}\right)\)
\end{tabular}

\section*{Convergence rate for \(f\left(\tilde{\theta}_{\boldsymbol{k}}\right)-\boldsymbol{f}\left(\boldsymbol{\theta}_{*}\right)\), smooth objective \(f\).}


\section*{Gradient is unknown}

\section*{Take home}
- In the context of large scale learning, we have to use SGD
- It is a stochastic algorithm
- Typically, steps sizes have to decay to 0
- For smooth problems, larger steps are allowed and adapts to strong convexity.

Moreover: "one epoch = one pass over my observations"

\section*{Take home}
- It is possible to use variance reduced algorithms to have a faster convergence rate after many epochs.
- During the first epoch, we optimize the (unknown!) generalization error!!
- powerful remark
- e.g., streaming setting.

\section*{Next Goals}
1. Even larger steps ?
2. Mini-batch algorithms.
3. Adaptive algorithms.

\section*{Summary of the first two days}
1. Large Scale Learning framework
2. Optimization
- First order methods: speed of convergence of GD
- SGD vs GD: SGD is fast \& low precision

- Variance reduced SGD
- Generalization with SGD: we can optimize an unknown function!


\section*{Convergence rate \(f\left(\tilde{\theta}_{k}\right)-f\left(\theta_{*}\right)\), smooth objective \(f\).}
\begin{tabular}{|c|c|c|c|c|}
\hline & \multicolumn{3}{|c|}{\(\min \hat{\mathcal{R}}\)} & \(\min \mathcal{R}\) \\
\hline & SGD & GD & SAG & SGD \\
\hline Convex & \(o\left(\frac{1}{\sqrt{k}}\right)\) & \(O\left(\frac{1}{k}\right)\) & & \(o\left(\frac{1}{\sqrt{n}}\right)\) \\
\hline Stgly-Cux & \(o\left(\frac{1}{\mu k}\right)\) & \(O\left(e^{-\mu k}\right)\) & \(O\left(1-\left(\mu \wedge \frac{1}{n}\right)\right)^{k}\) & O ( \(\frac{1}{\mu n}\) ) \\
\hline
\end{tabular}

\section*{Today}
1. Mini-batch algorithms
2. Adaptive algorithms
3. (Markov chain point of view)

\section*{Outline}
1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient

Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

\section*{ADAGrad Optimizer \\ AdaDelta Optimizer \\ RMSprop optimizer}
5. Wednesday: python practical
6. Larger steps

See the very good post: http://ruder.io/optimizing-gradient-descent/

\section*{Minibatch SGD for ERM: \(f=\hat{\mathcal{R}}\)}

Loss for a single pair of observations, for any \(\boldsymbol{j} \leq \boldsymbol{n}\) :
\[
f_{j}(\theta):=\ell\left(y_{j},\left\langle\theta, \Phi\left(x_{j}\right)\right\rangle\right)
\]

Empirical risk \(\hat{\mathcal{R}}(\theta)=\frac{1}{n} \sum_{t=1}^{n} \ell\left(y_{t},\left\langle\theta, \Phi\left(x_{t}\right)\right\rangle\right)\). SGD:
- At each step \(t \in \mathbb{N}^{*}\), sample \(I_{t} \sim \mathcal{U}\{1, \ldots n\}:\)
\[
\theta_{t}=\theta_{t-1}-\gamma_{t} f_{t_{t}}^{\prime}\left(\theta_{t-1}\right)
\]

Mini-batch SGD: choose \(\boldsymbol{m} \leq \boldsymbol{n}\)
- At each step \(t \in \mathbb{N}^{*}\), sample \(\left(I_{1, t}, \ldots, I_{m, t}\right) \sim \mathcal{U}\{1, \ldots n\}^{\otimes m}\) :
\[
\theta_{t}=\theta_{t-1}-\gamma_{t} \frac{1}{m} \sum_{i=1}^{m} f_{i, t}^{\prime}\left(\theta_{t-1}\right)
\]

\section*{Minibatch SGD : behavior}
1. Gradient is still stochastic (if \(m<n\) )
2. Level of noise in the gradient is reduced: formally
\[
\operatorname{var}\left(\frac{1}{m} \sum_{i=1}^{m} f_{i_{i, t}}^{\prime}\left(\theta_{t-1}\right)=\frac{1}{m} \operatorname{var}\left(f_{t_{t}}^{\prime}\left(\theta_{t-1}\right)\right)\right.
\]
3. Cost/time per iteration?
- cost/complexity: \(\mathbf{O}(\boldsymbol{m d})\) per iteration
- In practice, distribution of the computation over many cores can reduce the time par iteration to less than \(O(m d)\).
4. Convergence ?

We denote \(\sigma^{2}=\operatorname{var}\left(f_{t_{t}}^{\prime}\left(\theta_{t-1}\right)\right)\).

\section*{Convergence of SGD for smooth smooth \(f\)}

\section*{SGD:}
1. What matters? For smooth functions - the Variance of stochastic gradient. Bound \(\simeq\) :
\[
f\left(\bar{\theta}_{t}\right)-f\left(\theta_{*}\right) \leq \frac{\left\|\theta_{0}-\theta_{*}\right\|^{2}}{\gamma_{t} t}+\gamma_{t} \operatorname{var}\left(f_{t}^{\prime}\left(\theta_{t-1}\right)\right)
\]
2. "Optimal" step size: \(\gamma_{t}=\sqrt{\frac{\left\|\theta_{0}-\theta_{*}\right\|^{2}}{\sigma^{2} t}}\) : gives a rate
\[
f\left(\bar{\theta}_{t}\right)-f\left(\theta_{*}\right) \leq 2 \sqrt{\frac{\sigma^{2}\left\|\theta_{0}-\theta_{*}\right\|^{2}}{t}}
\]

Step size has always to be \(\leq \frac{2}{L}\) otherwise SGD diverges.

\section*{Convergence of mini-batch SGD for smooth \(\boldsymbol{f}\)}

\section*{Mini-batch SGD:}
- to keep same total complexity: \(\boldsymbol{t} \leftarrow \boldsymbol{t} / \mathrm{m}\)
- Reduced variance : \(\sigma^{2} \leftarrow \sigma^{2} / m\)
1. For smooth functions - the Variance of stochastic gradient. Bound \(\simeq\) :
\[
\boldsymbol{f}\left(\bar{\theta}_{t / m}\right)-\boldsymbol{f}\left(\theta_{*}\right) \leq \frac{\left\|\theta_{0}-\theta_{*}\right\|^{2}}{\gamma_{t} \boldsymbol{t} / m}+\frac{\gamma_{t} \sigma^{2}}{m} .
\]
2. "Optimal" step size: \(\gamma_{\boldsymbol{t}}=\sqrt{\frac{\left\|\boldsymbol{\theta}_{0}-\boldsymbol{\theta}_{*}\right\|^{2}}{\sigma^{2} / m} \boldsymbol{t} / m}=\boldsymbol{m} \sqrt{\frac{\left\|\boldsymbol{\theta}_{0}-\boldsymbol{\theta}_{*}\right\|^{2}}{\sigma^{2} \boldsymbol{t}}}\) : gives a rate
\[
f\left(\bar{\theta}_{t}\right)-f\left(\theta_{*}\right) \leq 2 \sqrt{\frac{\sigma^{2}\left\|\theta_{0}-\theta_{*}\right\|^{2}}{t}}
\]

Step size has always to be \(\leq \frac{2}{L}\) otherwise SGD diverges.

\section*{Convergence of mini-batch SGD for smooth \(\boldsymbol{f}\)}
\begin{tabular}{lccl} 
& SGD & \(m\)-Mini-batch SGD & \\
\hline Steps \(\mathbb{C}=\boldsymbol{O}(\boldsymbol{t d})\) & \(\boldsymbol{t}\) & \(\frac{t}{m}\) & \\
Gradient Variance & \(\sigma^{2}\) & \(\frac{\sigma^{2}}{m}\) & \\
Optimal step & \(\frac{c_{\theta_{0}, \sigma^{2}}}{\sqrt{t}}\) & \(m \frac{c_{\theta_{0}, \sigma^{2}}^{\sqrt{t}}}{t}\) & \(\wedge 2 L^{-1}!\) \\
Global rate & & \(\sqrt{\frac{\sigma^{2}\left\|\theta_{0}-\theta_{*}\right\|^{2}}{t}}\) &
\end{tabular}
1. Same Global convergence rate
2. If mini-batch size starts being too large, saturation because of the upper bound on the step size
3. Reasonable ( \(n\)-minibatch \(=G D!\) )
4. In practice, used a lot because time < complexity.

\section*{Convergence of SGD for smooth non-smooth \(f\)}

\section*{SGD:}
1. What matters? For non-smooth functions - the upper bound \(B^{2}\) on stochastic gradient. Bound \(\simeq\) :
\[
f\left(\bar{\theta}_{t}\right)-\boldsymbol{f}\left(\theta_{*}\right) \leq \frac{\left\|\theta_{0}-\theta_{*}\right\|^{2}}{\gamma_{t} t}+\gamma_{t} \sup \mathbb{E}\left\|\boldsymbol{f}_{t}^{\prime}\left(\theta_{t-1}\right)\right\|^{2}
\]

Mini-batch SGD:
\[
\sup \mathbb{E}\left\|\frac{1}{m} \sum_{i=1}^{m} f_{l_{i, t}}^{\prime}\left(\theta_{t-1}\right)\right\|^{2} \lesssim \sup \mathbb{E}\left\|f_{l_{t}}^{\prime}\left(\theta_{t-1}\right)\right\|^{2}
\]
1. Same bound for same number of iterations
2. Higher cost par iteration

Using mini-batch is a bad idea.

\section*{Convergence of minibatch SAG}

In variance reduced method:
1. The variance is already reduced by the method itself
2. No need to use mini-batch

\section*{Take home}

Mini-batch gradient descent:
1. Simple algorithm derived for SGD using a small "batch" of examples
2. Reduces the variance of the random gradients
3. Helps when
- 1. Function is smooth, \&
- 2. \(m\) not too large (Saturation) \&
- 3. Time < Complexity
4. Does not help much for non smooth function or Variance reduced methods.

Remark: all these insights come from theory and proofs.

Take home
Read papers or ask people with theoretical knowledge :)

\section*{Outline}
1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient

Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
4. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

\section*{ADAGrad Optimizer \\ AdaDelta Optimizer \\ RMSprop optimizer}
5. Wednesday: python practical
6. Larger steps

See the very good post: http://ruder.io/optimizing-gradient-descent/

\section*{Challenge number 1: Acceleration}
1. Earlier we saw that we could accelerate GD getting a better rate
2. Similar process for SGD.
- Might cause instability or divergence
- Not fully understood theoretically
- Used a lot in practice

\section*{Momentum algorithmI I}

Aim: related to Nesterov Acceleration but older (1964) Particularly useful for stochastic gradient descent.
https://distill.pub/2017/momentum/


\section*{Momentum algorithm II}

Polyak's momentum algorithm - Heavy ball method
1. starting point \(\boldsymbol{\theta}^{(0)}\),
2. learning rate \(\gamma_{t}>0\),
3. momentum \(\beta \in[0,1]\) (default \(\beta=0.9\) ).

Iterate
\[
\boldsymbol{\theta}_{\boldsymbol{t}+\boldsymbol{1}}=\boldsymbol{\theta}_{\boldsymbol{t}}-\gamma_{\boldsymbol{t}} \boldsymbol{\nabla} \boldsymbol{f}\left(\boldsymbol{\theta}_{\boldsymbol{t}}\right)+\beta\left(\theta_{t}-\theta_{t-1}\right)
\]

Return last \(\theta^{(t+1)}\).


\section*{Challenge number 2: Adaptation}
1. Same learning rate for all coordinates. Could we use a different learning rate for all coordinates ?
i.e., for \(1 \leq \boldsymbol{j} \leq \boldsymbol{d}\) :
\[
\left(\theta_{t}\right)_{j}=\left(\theta_{t-1}\right)_{j}-\gamma_{t, j}\left(f_{t_{t}}^{\prime}\left(\theta_{t-1}\right)\right)_{j}
\]

\section*{Intuition: Gradient descent}
- Quadratic convex function: \(f(\theta)=\frac{1}{2} \theta^{\top} H \theta-c^{\top} \theta\)
- \(\mu\) and \(L\) are smallest largest eigenvalues of \(H\)
- Global optimum \(\theta_{*}=H^{-1} c\left(\right.\) or \(H^{\dagger} c\) ) such that \(H \theta_{*}=c\)
- Gradient descent with learning rate \(\gamma\) :
\[
\theta_{t}-\theta_{*}=(I-\gamma H)\left(\theta_{t-1}-\theta_{*}\right)=(I-\gamma H)^{t}\left(\theta_{0}-\theta_{*}\right)
\]
- If \(H=\operatorname{Diag}\left(\alpha_{1}, \ldots, \alpha_{d}\right), \alpha_{1}=L, \alpha_{d}=\mu\)
- For coordinate \(\boldsymbol{j}\), we have:
\[
\left(\theta_{t}\right)_{j}=\left(1-\gamma \alpha_{j}\right)^{t}\left(\theta_{0}-\theta_{*}\right)_{j}
\]
- \(\rightarrow\) step size cannot be larger than \(2 / \alpha_{1}=2 / L\) otherwise first coefficient \(\left|\left(1-\gamma \alpha_{1}\right)\right|>1\) and this coordinate diverges.
- \(\rightarrow\) Rate is dictated by the smallest coordinate: rate
\[
\left(1-\alpha_{d} / \alpha_{1}\right)^{t}=(1-\mu / L)^{t}
\]

\section*{Notations}
\[
\left(\theta_{t}\right)_{j}=\left(\theta_{t-1}\right)_{j}-\gamma_{t, k}\left(f_{t_{t}}^{\prime}\left(\theta_{t-1}\right)\right)_{j}
\]
1. \(g_{t}=f_{t}^{\prime}\left(\theta_{t-1}\right)\) stochastic gradient at time \(t\)
\[
\left(\theta_{t}\right)_{j}=\left(\theta_{t-1}\right)_{j}-\gamma_{t, j}\left(g_{t}\right)_{j}
\]
2. Avoiding double subscript:
\[
\begin{gathered}
\left(\theta^{t}\right)_{j}=\left(\theta^{t-1}\right)_{j}-\gamma_{j}^{t}\left(g^{t}\right)_{j} \\
\theta_{j}^{t}=\theta_{j}^{t-1}-\gamma_{j}^{t} g_{j}^{t}
\end{gathered}
\]

\section*{ADAGRAD}
\[
\theta_{j}^{t}=\theta_{j}^{t-1}-\gamma_{j}^{t} g_{j}^{t}
\]

Special choice for step-sizes:
\[
\boldsymbol{\theta}_{j}^{\boldsymbol{t}}=\boldsymbol{\theta}_{j}^{\boldsymbol{t - 1}}-\frac{\gamma}{\sqrt{C_{t, j}+\varepsilon}} \boldsymbol{g}_{j}^{\boldsymbol{t}}
\]

\section*{ADAptive GRADient algorithm}
1. starting point \(\theta^{0}\),
2. learning rate \(\gamma>\mathbf{0}\), (default value of 0.01 )
3. momentum \(\beta\), constant \(\varepsilon\).

For \(t=1,2, \ldots\) until convergence do for \(1 \leq j \leq d\)
\[
\theta_{j}^{t+1} \leftarrow \theta_{j}^{t}-\frac{\gamma}{\sqrt{\sum_{\tau=1}^{t}\left(g_{j}^{\tau}\right)^{2}+\varepsilon}} g_{j}^{t}
\]

Return last \(\boldsymbol{\theta}^{\boldsymbol{t}}\)

\section*{ADAGRAD}

Update equation for ADAGRAD \(\theta_{j}^{t+1} \leftarrow \theta_{j}^{t}-\frac{\gamma}{\sqrt{\sum_{\tau=1}^{t}\left(g_{j}^{\tau}\right)^{2}+\varepsilon}} \boldsymbol{g}_{j}^{t}\) Pros:
- Different dynamic rates on each coordinate
- Dynamic rates grow as the inverse of the gradient magnitude:
1. Large/small gradients have small/large learning rates
2. The dynamic over each dimension tends to be of the same order
3. Interesting for NN in which gradient at different layers can be of different order of magnitude.
- Accumulation of gradients in the denominator act as a decreasing learning rate.
Cons:
- Very sensitive to initial condition: large initial gradients lead to small learning rates.
- Can be fought by increasing the learning rate thus making the algorithm sensitive to the choice of the learning rate.

\section*{Improving upon AdaGrad: AdaDelta}

Idea : restricts the window of accumulated past gradients to some fixed size.
1. starting point \(\boldsymbol{\theta}^{0}\), constant \(\varepsilon\),
2. new params: decay rate \(\rho>0\)

Update:
\[
\theta_{j}^{t+1}=\theta_{j}^{t}-\frac{\gamma_{j}^{t}}{\sqrt{C_{j, t}+\varepsilon}} g_{j}^{t}
\]

Before: \(C_{j, t}=\sum_{\tau=1}^{t}\left(g_{j}^{\tau}\right)^{2}\)
Now: \(\quad C_{j, t}=\rho C_{j}^{t-1}+(1-\rho)\left(g_{j}^{t}\right)^{2}\)

\section*{Adadelta}

Interpretation:
- Less sensitivity to initial parameters than Adagrad.
- \(\gamma_{j}^{t}\) is chosen to by size the previous step in memory and enforce larger steps along directions in which large steps were made.
- The denominator keeps the size of the previous gradients in memory and acts as a decreasing learning rate. Weights are lower than in Adagrad due to the decay rate \(\rho\).

\section*{RMSprop}

Unpublished methode, from the online course of Geoff Hinton
\[
\begin{gathered}
\text { http://www.cs.toronto.edu/~tijmen/csc321/slides/ } \\
\text { lecture_slides_lec6.pdf }
\end{gathered}
\]
1. starting point \(\theta^{0}\), constant \(\varepsilon\),
2. decay rate \(\rho>0\)
3. "new" step size \(\gamma(\) default \(=\mathbf{0 . 0 0 1})\)

Update:
\[
\theta_{j}^{t+1}=\theta_{j}^{t}-\frac{\gamma}{\sqrt{C_{j, t}+\varepsilon}} \boldsymbol{g}_{j}^{t}
\]

\section*{Animation of Stochastic Gradient algorithms}


\section*{Wednesday}

Goal: Code:
1. gradient descent (GD)
2. accelerated gradient descent (AGD)
3. coordinate gradient descent (CD)
4. stochastic gradient descent (SGD)
5. stochastic variance reduced gradient descent (SAG)
6. Adagrad
for the linear regression and logistic regression models, with the ridge penalization.


\section*{Wednesday}
1. Who knows python ?
2. Who's using anaconda?
1. Motivation: Large scale learning and Optimization
2. Classical rates for deterministic methods
3. Supervised learning setting - Stochastic Gradient

Algorithms
3.1 SGD vs GD
3.2 Variance reduced SGD
3.3 SGD to avoid overfitting (Generalization Risk)
1. Mini-batch, Adaptive algorithms
4.1 Mini-batch Algotirhms
4.2 Adaptive algorithms

ADAGrad Optimizer
AdaDelta Optimizer
RMSprop optimizer
5. Wednesday: python practical
6. Larger steps

\section*{Least Mean Squares: rate independent of \(\mu\)}

Least-squares: \(\mathcal{R}(\theta)=\frac{1}{2} \mathbb{E}\left[(Y-\langle\Phi(X), \theta\rangle)^{2}\right]\)
Analysis for averaging and constant step-size \(\gamma=1 /\left(4 R^{2}\right)\)
(?)
- Assume \(\left\|\Phi\left(x_{n}\right)\right\| \leqslant r\) and \(\left|y_{n}-\left\langle\Phi\left(x_{n}\right), \theta_{*}\right\rangle\right| \leqslant \sigma\)
- No assumption regarding lowest eigenvalues of the Hessian
\[
\mathbb{E} \mathcal{R}\left(\overline{\boldsymbol{\theta}}_{\boldsymbol{n}}\right)-\mathcal{R}\left(\boldsymbol{\theta}_{*}\right) \leqslant \frac{4 \sigma^{2} d}{n}+\frac{\left\|\boldsymbol{\theta}_{0}-\boldsymbol{\theta}_{*}\right\|^{2}}{\gamma \boldsymbol{n}}
\]
- Matches statistical lower bound (Tsybakov, 2003).
- Optimal rate with "large" step sizes

Take home
- SGD can be used to minimize the true risk directly
- Stochastic algorithm to minimize unknown function
- No regularization needed, only one pass
- For Least Squares, with constant step, optimal rate .

\section*{Beyond least squares. Logistic regression}
\[
\min _{\theta \in \mathbb{D}^{d}} \mathbb{E} \log (1+\exp (-Y\langle\theta, \Phi(X)\rangle)) .
\]


Logistic regression. Final iterate (dashed), and averaged recursion (plain).

\section*{Motivation 2/ 2. Difference between quadratic and logistic loss}


Logistic Regression
\(\mathbb{E} \mathcal{R}\left(\bar{\theta}_{n}\right)-\mathcal{R}\left(\theta_{*}\right)=O\left(\gamma^{2}\right)\) with \(\gamma=1 /\left(4 R^{2}\right)\)


Least-Squares Regression \(\mathbb{E} \mathcal{R}\left(\bar{\theta}_{n}\right)-\mathcal{R}\left(\theta_{*}\right)=O\left(\frac{1}{n}\right)\) with \(\gamma=1 /\left(4 R^{2}\right)\)

\section*{SGD: an homogeneous Markov chain}

Consider a \(L\)-smooth and \(\mu\)-strongly convex function \(\mathcal{R}\).
SGD with a step-size \(\gamma>\mathbf{0}\) is an homogeneous Markov chain:
\[
\boldsymbol{\theta}_{k+1}^{\gamma}=\boldsymbol{\theta}_{k}^{\gamma}-\gamma\left[\mathcal{R}^{\prime}\left(\boldsymbol{\theta}_{k}^{\gamma}\right)+\varepsilon_{k+1}\left(\boldsymbol{\theta}_{k}^{\gamma}\right)\right],
\]
- satisfies Markov property
- is homogeneous, for \(\gamma\) constant, \(\left(\varepsilon_{k}\right)_{k \in \mathbb{N}}\) i.i.d.

Also assume:
- \(\mathcal{R}_{k}^{\prime}=\mathcal{R}^{\prime}+\varepsilon_{k+1}\) is almost surely L-co-coercive.
- Bounded moments
\[
\mathbb{E}\left[\left\|\varepsilon_{k}\left(\theta_{*}\right)\right\|^{4}\right]<\infty
\]

\section*{Stochastic gradient descent as a Markov Chain: Analysis framework \({ }^{\dagger}\)}
- Existence of a limit distribution \(\pi_{\gamma}\), and linear convergence to this distribution:
\[
\theta_{k}^{\gamma} \xrightarrow{d} \pi_{\gamma} .
\]
- Convergence of second order moments of the chain,
\[
\bar{\theta}_{k}^{\gamma} \xrightarrow[k \rightarrow \infty]{L^{2}} \bar{\theta}_{\gamma}:=\mathbb{E}_{\pi_{\gamma}}[\theta]
\]
- Behavior under the limit distribution \((\gamma \rightarrow \mathbf{0}): \bar{\theta}_{\gamma}=\theta_{*}+\) ?.
\(\rightarrow\) Provable convergence improvement with extrapolation tricks.

\section*{Existence of a limit distribution \(\gamma \rightarrow \mathbf{0}\)}

Goal:
\[
\left(\theta_{k}^{\gamma}\right)_{k \geq 0} \xrightarrow{d} \pi_{\gamma} .
\]

\section*{Theorem}

For any \(\gamma<\boldsymbol{L}^{-1}\), the chain \(\left(\boldsymbol{\theta}_{\boldsymbol{k}}^{\gamma}\right)_{k \geq 0}\) admits a unique stationary distribution \(\pi_{\gamma}\). In addition for all \(\boldsymbol{\theta}_{0} \in \mathbb{R}^{\boldsymbol{d}}, \boldsymbol{k} \in \mathbb{N}\) :
\[
W_{2}^{2}\left(\theta_{k}^{\gamma}, \pi_{\gamma}\right) \leq(1-2 \mu \gamma(1-\gamma L))^{k} \int_{\mathbb{R}^{d}}\left\|\theta_{0}-\vartheta\right\|^{2} \mathrm{~d} \pi_{\gamma}(\vartheta)
\]

Wasserstein metric: distance between probability measures.

\section*{Behavior under limit distribution.} Ergodic theorem: \(\bar{\theta}_{k} \rightarrow \mathbb{E}_{\pi_{\gamma}}[\theta]=: \bar{\theta}_{\gamma}\). Where is \(\overline{\theta_{\gamma}}\) ?

If \(\theta_{0} \sim \pi_{\gamma}\), then \(\theta_{1} \sim \pi_{\gamma}\).
\[
\theta_{1}^{\gamma}=\theta_{0}^{\gamma}-\gamma\left[\mathcal{R}^{\prime}\left(\theta_{0}^{\gamma}\right)+\varepsilon_{1}\left(\theta_{0}^{\gamma}\right)\right] .
\]
\[
\mathbb{E}_{\pi_{\gamma}}\left[\mathcal{R}^{\prime}(\theta)\right]=0
\]

In the quadratic case (linear gradients) \(\boldsymbol{\Sigma} \mathbb{E}_{\boldsymbol{\pi}_{\gamma}}\left[\boldsymbol{\theta}-\boldsymbol{\theta}_{*}\right]=\mathbf{0}: \bar{\theta}_{\gamma}=\theta_{*}\) !

\section*{Constant learning rate SGD: convergence in the quadratic case}


\section*{Constant learning rate SGD: convergence in the quadratic case}


\section*{Constant learning rate SGD: convergence in the quadratic case}


\section*{Constant learning rate SGD: convergence in the quadratic case}


\section*{Behavior under limit distribution.}

Ergodic theorem: \(\overline{\boldsymbol{\theta}}_{\boldsymbol{n}} \rightarrow \mathbb{E}_{\boldsymbol{\pi}_{\gamma}}[\theta]=: \overline{\boldsymbol{\theta}_{\gamma}}\). Where is \(\overline{\theta_{\gamma}}\) ?
If \(\theta_{0} \sim \pi_{\gamma}\), then \(\theta_{1} \sim \pi_{\gamma}\).
\[
\begin{gathered}
\theta_{1}^{\gamma}=\theta_{0}^{\gamma}-\gamma\left[\mathcal{R}^{\prime}\left(\theta_{0}^{\gamma}\right)+\varepsilon_{1}\left(\theta_{0}^{\gamma}\right)\right] . \\
\mathbb{E}_{\boldsymbol{\pi}_{\gamma}}\left[\mathcal{R}^{\prime}(\theta)\right]=\mathbf{0}
\end{gathered}
\]

In the quadratic case (linear gradients) \(\boldsymbol{\Sigma} \mathbb{E}_{\boldsymbol{\pi}_{\gamma}}\left[\boldsymbol{\theta}-\boldsymbol{\theta}_{*}\right]=\mathbf{0}\) : \(\bar{\theta}_{\gamma}=\theta_{*}\) !
In the general case, Taylor expansion of \(\mathcal{R}\), and same reasoning on higher moments of the chain leads to
\(\bar{\theta}_{\gamma}-\theta_{*} \simeq \gamma \mathcal{R}^{\prime \prime}\left(\theta_{*}\right)^{-1} \mathcal{R}^{\prime \prime \prime}\left(\theta_{*}\right)\left(\left[\mathcal{R}^{\prime \prime}\left(\theta_{*}\right) \otimes I+\prime \otimes \mathcal{R}^{\prime \prime}\left(\theta_{*}\right)\right]^{-1} \mathbb{E}_{\varepsilon}\left[\varepsilon\left(\theta_{*}\right)^{\otimes 2}\right]\right)\)
Overall, \(\bar{\theta}_{\gamma}-\theta_{*}=\gamma \Delta+O\left(\gamma^{2}\right)\).

\section*{Constant learning rate SGD: convergence in the non-quadratic case}


\section*{Constant learning rate SGD: convergence in the non-quadratic case}


\section*{Constant learning rate SGD: convergence in the non-quadratic case}


\section*{Constant learning rate SGD: convergence in the non-quadratic case}


\section*{Richardson extrapolation}


\section*{Richardson extrapolation}
\[
\begin{aligned}
& . \theta_{*} \\
& \bar{\theta}_{\gamma} \bigodot \theta_{*}+\gamma \Delta
\end{aligned}
\]

\section*{Richardson extrapolation}


\section*{Richardson extrapolation}


\section*{Richardson extrapolation}


\section*{Richardson extrapolation}


Recovering convergence closer to \(\theta_{*}\) by Richardson extrapolation \(2 \bar{\theta}_{n}^{\gamma}-\bar{\theta}_{n}^{2 \gamma}\)

\section*{Experiments: smaller dimension}


Synthetic data, logistic regression, \(n=8.10^{6}\)

\section*{Experiments: Double Richardson}


Synthetic data, logistic regression, \(n=8.10^{6}\) "Richardson \(3 \gamma\) ": estimator built using Richardson on 3 different sequences: \(\tilde{\theta}_{n}^{3}=\frac{8}{3} \bar{\theta} \bar{n}_{n}^{\gamma}-2 \bar{\theta}_{n}^{2 \gamma}+\frac{1}{3} \bar{\theta}_{n}^{4 \gamma}\)

\section*{Conclusion MC}

\section*{Take home}
- Asymptotic sometimes matter less than first iterations: consider large step size.
- Constant step size SGD is a homogeneous Markov chain.
- Difference between LS and general smooth loss is intuitive.

For smooth strongly convex loss:
- Convergence in terms of Wasserstein distance.
- Decomposition as three sources of error: variance, initial conditions, and "drift"
- Detailed analysis of the position of the limit point: the direction does not depend on \(\gamma\) at first order \(\Longrightarrow\) Extrapolation tricks can help.

\section*{Further references}

Many stochastic algorithms not covered in this talk (coordinate descent, online Newton, composite optimization, non convex learning) ...
- Good introduction: Francis's lecture notes at Orsay
- Book:

Convex Optimization: Algorithms and Complexity, Sébastien Bubeck

Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences, 2(1):183-202.

Bottou, L. and Bousquet, O. (2008). The tradeoffs of large scale learning. In Adv. NIPS.
Boyd, S. and Vandenberghe, L. (2003). Convex Optimization. Cambridge University Press.
Bubeck, S. (2015). Convex optimization: Algorithms and complexity. Foundations and Trends \(®\) in Machine Learning, 8(3-4):231-357.
Defazio, A., Bach, F., and Lacoste-Julien, S. (2014a). Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In Advances in Neural Information Processing Systems, pages 1646-1654.
Defazio, A., Domke, J., and Caetano, T. (2014b). Finito: A faster, permutable incremental gradient method for big data problems. In Proceedings of the 31st international conference on machine learning (ICML-14), pages 1125-1133.
Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In Advances in neural information processing systems, pages 315-323.
Konečnỳ, J. and Richtárik, P. (2013). Semi-stochastic gradient descent methods. arXiv preprint arXiv:1312.1666.

Nesterov, Y. (1983). A method of solving a convex programming problem with convergence rate o (1/k2). In Soviet Mathematics Doklady, volume 27, pages 372-376.

Nesterov, Y. (2004). Introductory lectures on convex optimization: A basic course. Springer.

Nesterov, Y. (2007). Gradient methods for minimizing composite objective function. Center for Operations Research and Econometrics (CORE), Catholic University of Louvain, Tech. Rep, 76.
Polyak, B. T. and Juditsky, A. B. (1992). Acceleration of stochastic approximation by averaging. SIAM J. Control Optim., 30(4):838-855.
Robbins, H. and Monro, S. (1951). A stochastic approxiation method. The Annals of mathematical Statistics, 22(3):400-407.
Ruppert, D. (1988). Efficient estimations from a slowly convergent Robbins-Monro process. Technical report, Cornell University Operations Research and Industrial Engineering.
Schmidt, M., Le Roux, N., and Bach, F. (2011). Convergence rates for inexact proximal-gradient method. In Adv. NIPS.
Schmidt, M., Le Roux, N., and Bach, F. (2013). Minimizing finite sums with the stochastic average gradient. Mathematical Programming, 162(1-2):83-112.
Shor, N. Z., Kiwiel, K. C., and Ruszcay?ski, A. (1985). Minimization methods for non-differentiable functions. Springer-Verlag New York, Inc.
Tsybakov, A. B. (2003). Optimal rates of aggregation. In Proceedings of the Annual Conference on Computational Learning Theory.```

