# Introduction to Machine Learning (67577) Lecture 6 

## Shai Shalev-Shwartz

School of CS and Engineering, The Hebrew University of Jerusalem

Convexity,Optimization,Surrogates,SGD

## Outline

(1) Convexity
(2) Convex Optimization

- Ellipsoid
- Gradient Descent
(3) Convex Learning Problems

4 Surrogate Loss Functions
(5) Learning Using Stochastic Gradient Descent

## Definition (Convex Set)

A set $C$ in a vector space is convex if for any two vectors $\mathbf{u}, \mathbf{v}$ in $C$, the line segment between $\mathbf{u}$ and $\mathbf{v}$ is contained in $C$. That is, for any $\alpha \in[0,1]$ we have that the convex combination $\alpha \mathbf{u}+(1-\alpha) \mathbf{v}$ is in $C$.


## Definition (Convex function)

Let $C$ be a convex set. A function $f: C \rightarrow \mathbb{R}$ is convex if for every $\mathbf{u}, \mathbf{v} \in C$ and $\alpha \in[0,1]$,

$$
f(\alpha \mathbf{u}+(1-\alpha) \mathbf{v}) \leq \alpha f(\mathbf{u})+(1-\alpha) f(\mathbf{v})
$$



## Epigraph

A function $f$ is convex if and only if its epigraph is a convex set:

$$
\operatorname{epigraph}(\mathrm{f})=\{(\mathbf{x}, \beta): f(\mathbf{x}) \leq \beta\}
$$



## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

- let $B(\mathbf{u}, r)=\{\mathbf{v}:\|\mathbf{v}-\mathbf{u}\| \leq r\}$


## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

- let $B(\mathbf{u}, r)=\{\mathbf{v}:\|\mathbf{v}-\mathbf{u}\| \leq r\}$
- $f(\mathbf{u})$ is a local minimum of $f$ at $\mathbf{u}$ if $\exists r>0$ s.t. $\forall \mathbf{v} \in B(\mathbf{u}, r)$ we have $f(\mathbf{v}) \geq f(\mathbf{u})$


## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

- let $B(\mathbf{u}, r)=\{\mathbf{v}:\|\mathbf{v}-\mathbf{u}\| \leq r\}$
- $f(\mathbf{u})$ is a local minimum of $f$ at $\mathbf{u}$ if $\exists r>0$ s.t. $\forall \mathbf{v} \in B(\mathbf{u}, r)$ we have $f(\mathbf{v}) \geq f(\mathbf{u})$
- It follows that for any $\mathbf{v}$ ( not necessarily in $B$ ), there is a small enough $\alpha>0$ such that $\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}) \in B(\mathbf{u}, r)$ and therefore

$$
f(\mathbf{u}) \leq f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))
$$

## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

- let $B(\mathbf{u}, r)=\{\mathbf{v}:\|\mathbf{v}-\mathbf{u}\| \leq r\}$
- $f(\mathbf{u})$ is a local minimum of $f$ at $\mathbf{u}$ if $\exists r>0$ s.t. $\forall \mathbf{v} \in B(\mathbf{u}, r)$ we have $f(\mathbf{v}) \geq f(\mathbf{u})$
- It follows that for any $\mathbf{v}$ ( not necessarily in $B$ ), there is a small enough $\alpha>0$ such that $\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}) \in B(\mathbf{u}, r)$ and therefore

$$
f(\mathbf{u}) \leq f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))
$$

- If $f$ is convex, we also have that

$$
f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))=f(\alpha \mathbf{v}+(1-\alpha) \mathbf{u}) \leq(1-\alpha) f(\mathbf{u})+\alpha f(\mathbf{v}) .
$$

## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

- let $B(\mathbf{u}, r)=\{\mathbf{v}:\|\mathbf{v}-\mathbf{u}\| \leq r\}$
- $f(\mathbf{u})$ is a local minimum of $f$ at $\mathbf{u}$ if $\exists r>0$ s.t. $\forall \mathbf{v} \in B(\mathbf{u}, r)$ we have $f(\mathbf{v}) \geq f(\mathbf{u})$
- It follows that for any $\mathbf{v}$ ( not necessarily in $B$ ), there is a small enough $\alpha>0$ such that $\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}) \in B(\mathbf{u}, r)$ and therefore

$$
f(\mathbf{u}) \leq f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))
$$

- If $f$ is convex, we also have that

$$
f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))=f(\alpha \mathbf{v}+(1-\alpha) \mathbf{u}) \leq(1-\alpha) f(\mathbf{u})+\alpha f(\mathbf{v}) .
$$

- Combining, we obtain that $f(\mathbf{u}) \leq f(\mathbf{v})$.


## Property I: local minima are global

If $f$ is convex then every local minimum of $f$ is also a global minimum.

- let $B(\mathbf{u}, r)=\{\mathbf{v}:\|\mathbf{v}-\mathbf{u}\| \leq r\}$
- $f(\mathbf{u})$ is a local minimum of $f$ at $\mathbf{u}$ if $\exists r>0$ s.t. $\forall \mathbf{v} \in B(\mathbf{u}, r)$ we have $f(\mathbf{v}) \geq f(\mathbf{u})$
- It follows that for any $\mathbf{v}$ ( not necessarily in $B$ ), there is a small enough $\alpha>0$ such that $\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}) \in B(\mathbf{u}, r)$ and therefore

$$
f(\mathbf{u}) \leq f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))
$$

- If $f$ is convex, we also have that

$$
f(\mathbf{u}+\alpha(\mathbf{v}-\mathbf{u}))=f(\alpha \mathbf{v}+(1-\alpha) \mathbf{u}) \leq(1-\alpha) f(\mathbf{u})+\alpha f(\mathbf{v})
$$

- Combining, we obtain that $f(\mathbf{u}) \leq f(\mathbf{v})$.
- This holds for every $\mathbf{v}$, hence $f(\mathbf{u})$ is also a global minimum of $f$.


## Property II: tangents lie below $f$

If $f$ is convex and differentiable, then

$$
\forall \mathbf{u}, \quad f(\mathbf{u}) \geq f(\mathbf{w})+\langle\nabla f(\mathbf{w}), \mathbf{u}-\mathbf{w}\rangle
$$

(recall, $\nabla f(\mathbf{w})=\left(\frac{\partial f(\mathbf{w})}{\partial w_{1}}, \ldots, \frac{\partial f(\mathbf{w})}{\partial w_{d}}\right)$ is the gradient of $f$ at $\mathbf{w}$ )


## Sub-gradients

- $\mathbf{v}$ is sub-gradient of $f$ at $\mathbf{w}$ if $\forall \mathbf{u}, \quad f(\mathbf{u}) \geq f(\mathbf{w})+\langle\mathbf{v}, \mathbf{u}-\mathbf{w}\rangle$
- The differential set, $\partial f(\mathbf{w})$, is the set of sub-gradients of $f$ at $\mathbf{w}$
- Lemma: $f$ is convex iff for every $\mathbf{w}, \partial f(\mathbf{w}) \neq \emptyset$




## Property II: tangents lie below $f$

$f$ is "locally flat" around $\mathbf{w}$ (i.e. $\mathbf{0}$ is a sub-gradient) iff $\mathbf{w}$ is a global minimizer


## Lipschitzness

## Definition (Lipschitzness)

A function $f: C \rightarrow \mathbb{R}$ is $\rho$-Lipschitz if for every $\mathbf{w}_{1}, \mathbf{w}_{2} \in C$ we have that $\left|f\left(\mathbf{w}_{1}\right)-f\left(\mathbf{w}_{2}\right)\right| \leq \rho\left\|\mathbf{w}_{1}-\mathbf{w}_{2}\right\|$.

## Lipschitzness

## Definition (Lipschitzness)

A function $f: C \rightarrow \mathbb{R}$ is $\rho$-Lipschitz if for every $\mathbf{w}_{1}, \mathbf{w}_{2} \in C$ we have that $\left|f\left(\mathbf{w}_{1}\right)-f\left(\mathbf{w}_{2}\right)\right| \leq \rho\left\|\mathbf{w}_{1}-\mathbf{w}_{2}\right\|$.

## Lemma

If $f$ is convex then $f$ is $\rho$-Lipschitz iff the norm of all sub-gradients of $f$ is at most $\rho$

## Outline

(1) Convexity
(2) Convex Optimization

- Ellipsoid
- Gradient Descent


## (3) Convex Learning Problems

4 Surrogate Loss Functions
(5) Learning Using Stochastic Gradient Descent

## Convex optimization

Approximately solve:

$$
\underset{\mathbf{w} \in C}{\operatorname{argmin}} f(\mathbf{w})
$$

where $C$ is a convex set and $f$ is a convex function.

## Convex optimization

Approximately solve:

$$
\underset{\mathbf{w} \in C}{\operatorname{argmin}} f(\mathbf{w})
$$

where $C$ is a convex set and $f$ is a convex function.

Special cases:

- Feasibility problem: $f$ is a constant function


## Convex optimization

Approximately solve:

$$
\underset{\mathbf{w} \in C}{\operatorname{argmin}} f(\mathbf{w})
$$

where $C$ is a convex set and $f$ is a convex function.

Special cases:

- Feasibility problem: $f$ is a constant function
- Unconstrained minimization: $C=\mathbb{R}^{d}$


## Convex optimization

Approximately solve:

$$
\underset{\mathbf{w} \in C}{\operatorname{argmin}} f(\mathbf{w})
$$

where $C$ is a convex set and $f$ is a convex function.

Special cases:

- Feasibility problem: $f$ is a constant function
- Unconstrained minimization: $C=\mathbb{R}^{d}$
- Can reduce one to another:


## Convex optimization

Approximately solve:

$$
\underset{\mathbf{w} \in C}{\operatorname{argmin}} f(\mathbf{w})
$$

where $C$ is a convex set and $f$ is a convex function.

Special cases:

- Feasibility problem: $f$ is a constant function
- Unconstrained minimization: $C=\mathbb{R}^{d}$
- Can reduce one to another:
- Adding the function $I_{C}(\mathbf{w})$ to the objective eliminates the constraint


## Convex optimization

Approximately solve:

$$
\underset{\mathbf{w} \in C}{\operatorname{argmin}} f(\mathbf{w})
$$

where $C$ is a convex set and $f$ is a convex function.

Special cases:

- Feasibility problem: $f$ is a constant function
- Unconstrained minimization: $C=\mathbb{R}^{d}$
- Can reduce one to another:
- Adding the function $I_{C}(\mathbf{w})$ to the objective eliminates the constraint
- Adding the constraint $f(\mathbf{w}) \leq f^{*}+\epsilon$ eliminates the objective


## Outline

(1) Convexity
(2) Convex Optimization

- Ellipsoid
- Gradient Descent
(3) Convex Learning Problems

4 Surrogate Loss Functions
(5) Learning Using Stochastic Gradient Descent

## The Ellipsoid Algorithm

- Consider a feasibility problem: find $\mathbf{w} \in C$
- Assumptions:
- $B\left(\mathbf{w}^{*}, r\right) \subseteq C \subset B(0, R)$
- Separation oracle: Given $\mathbf{w}$, the oracle tells if it's in $C$ or not. If $\mathbf{w} \notin C$ then the oracle finds $\mathbf{v}$ s.t. for every $\mathbf{w}^{\prime} \in C$ we have $\langle\mathbf{w}, \mathbf{v}\rangle<\left\langle\mathbf{w}^{\prime}, \mathbf{v}\right\rangle$



## The Ellipsoid Algorithm

- We implicitly maintain an ellipsoid: $\mathcal{E}_{t}=\mathcal{E}\left(A_{t}^{1 / 2}, \mathbf{w}_{t}\right)$
- Start with $\mathbf{w}_{1}=\mathbf{0}, A_{1}=I$
- For $t=1,2, \ldots$
- Call oracle with $\mathbf{w}_{t}$
- If $\mathbf{w}_{t} \in C$, break and return $\mathbf{w}_{t}$
- Otherwise, let $\mathbf{v}_{t}$ be the vector defining a separating hyperplane - Update:

$$
\begin{aligned}
\mathbf{w}_{t+1} & =\mathbf{w}_{t}+\frac{1}{d+1} \frac{A_{t} \mathbf{v}_{t}}{\sqrt{\mathbf{v}_{t}^{\top} A_{t} \mathbf{v}_{t}}} \\
A_{t+1} & =\frac{d^{2}}{d^{2}-1}\left(A_{t}-\frac{2}{d+1} \frac{A_{t} \mathbf{v}_{t} \mathbf{v}_{t}^{\top} A_{t}}{\mathbf{v}_{t}^{\top} A_{t} \mathbf{v}_{t}}\right)
\end{aligned}
$$

## The Ellipsoid Algorithm

- We implicitly maintain an ellipsoid: $\mathcal{E}_{t}=\mathcal{E}\left(A_{t}^{1 / 2}, \mathbf{w}_{t}\right)$
- Start with $\mathbf{w}_{1}=\mathbf{0}, A_{1}=I$
- For $t=1,2, \ldots$
- Call oracle with $\mathbf{w}_{t}$
- If $\mathbf{w}_{t} \in C$, break and return $\mathbf{w}_{t}$
- Otherwise, let $\mathbf{v}_{t}$ be the vector defining a separating hyperplane - Update:

$$
\begin{aligned}
\mathbf{w}_{t+1} & =\mathbf{w}_{t}+\frac{1}{d+1} \frac{A_{t} \mathbf{v}_{t}}{\sqrt{\mathbf{v}_{t}^{\top} A_{t} \mathbf{v}_{t}}} \\
A_{t+1} & =\frac{d^{2}}{d^{2}-1}\left(A_{t}-\frac{2}{d+1} \frac{A_{t} \mathbf{v}_{t} \mathbf{v}_{t}^{\top} A_{t}}{\mathbf{v}_{t}^{\top} A_{t} \mathbf{v}_{t}}\right)
\end{aligned}
$$

## Theorem

The Ellipsoid converges after at most $2 d(2 d+2) \log (R / r)$ iterations.

## Implementing the separation oracle using sub-gradients

- Suppose $C=\cap_{i=1}^{n}\left\{\mathbf{w}: f_{i}(\mathbf{w}) \leq 0\right\}$ where each $f_{i}$ is a convex function.
- Given w, we can check if $f_{i}(\mathbf{w}) \leq 0$ for every $i$
- If $f_{i}(\mathbf{w})>0$ for some $i$, consider $\mathbf{v} \in \partial f_{i}(\mathbf{w})$, then, for every $\mathbf{w}^{\prime} \in C$

$$
0 \geq f_{i}\left(\mathbf{w}^{\prime}\right) \geq f_{i}(\mathbf{w})+\left\langle\mathbf{w}^{\prime}-\mathbf{w}, \mathbf{v}\right\rangle>\left\langle\mathbf{w}^{\prime}-\mathbf{w}, \mathbf{v}\right\rangle
$$

- So, the oracle can return - $\mathbf{v}$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$
- Analysis:


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$
- Analysis:
- Let $r$ be s.t. $B\left(\mathbf{w}^{*}, r\right) \subseteq C$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$
- Analysis:
- Let $r$ be s.t. $B\left(\mathbf{w}^{*}, r\right) \subseteq C$
- For example, if $f$ is $\rho$-Lipschitz we can take $r=\epsilon / \rho$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$
- Analysis:
- Let $r$ be s.t. $B\left(\mathbf{w}^{*}, r\right) \subseteq C$
- For example, if $f$ is $\rho$-Lipschitz we can take $r=\epsilon / \rho$
- Let $R=\left\|\mathbf{w}^{*}\right\|+r$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$
- Analysis:
- Let $r$ be s.t. $B\left(\mathbf{w}^{*}, r\right) \subseteq C$
- For example, if $f$ is $\rho$-Lipschitz we can take $r=\epsilon / \rho$
- Let $R=\left\|\mathbf{w}^{*}\right\|+r$
- Then, after $2 d(2 d+2) \log (R / r)$ iterations, $\mathbf{w}_{t}$ must be in $C$


## The Ellipsoid Algorithm for unconstrained minimization

- Consider $\min _{\mathbf{w}} f(\mathbf{w})$ and let $\mathbf{w}^{*}$ be a minimizer
- Let $C=\left\{\mathbf{w}: f(\mathbf{w})-f\left(\mathbf{w}^{*}\right)-\epsilon \leq 0\right\}$
- We can apply the Ellipsoid algorithm while letting $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}_{t}\right)$
- Analysis:
- Let $r$ be s.t. $B\left(\mathbf{w}^{*}, r\right) \subseteq C$
- For example, if $f$ is $\rho$-Lipschitz we can take $r=\epsilon / \rho$
- Let $R=\left\|\mathbf{w}^{*}\right\|+r$
- Then, after $2 d(2 d+2) \log (R / r)$ iterations, $\mathbf{w}_{t}$ must be in $C$
- For $f$ being $\rho$-Lipschitz, we obtain the iteration bound

$$
2 d(2 d+2) \log \left(\frac{\rho\left\|\mathbf{w}^{*}\right\|}{\epsilon}+1\right)
$$

## Outline

(1) Convexity
(2) Convex Optimization

- Ellipsoid
- Gradient Descent
(3) Convex Learning Problems

4 Surrogate Loss Functions
(5) Learning Using Stochastic Gradient Descent

## Gradient Descent

- Start with initial $\mathbf{w}^{(1)}$ (usually, the zero vector)


## Gradient Descent

- Start with initial $\mathbf{w}^{(1)}$ (usually, the zero vector)
- At iteration $t$, update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla f\left(\mathbf{w}^{(t)}\right),
$$

where $\eta>0$ is a parameter

## Gradient Descent

- Start with initial $\mathbf{w}^{(1)}$ (usually, the zero vector)
- At iteration $t$, update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla f\left(\mathbf{w}^{(t)}\right)
$$

where $\eta>0$ is a parameter

- Intuition:


## Gradient Descent

- Start with initial $\mathbf{w}^{(1)}$ (usually, the zero vector)
- At iteration $t$, update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla f\left(\mathbf{w}^{(t)}\right)
$$

where $\eta>0$ is a parameter

- Intuition:
- By Taylor's approximation, if $\mathbf{w}$ close to $\mathbf{w}^{(t)}$ then

$$
f(\mathbf{w}) \approx f\left(\mathbf{w}^{(t)}\right)+\left\langle\mathbf{w}-\mathbf{w}^{(t)}, \nabla f\left(\mathbf{w}^{(t)}\right)\right\rangle
$$

## Gradient Descent

- Start with initial $\mathbf{w}^{(1)}$ (usually, the zero vector)
- At iteration $t$, update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla f\left(\mathbf{w}^{(t)}\right),
$$

where $\eta>0$ is a parameter

- Intuition:
- By Taylor's approximation, if $\mathbf{w}$ close to $\mathbf{w}^{(t)}$ then

$$
f(\mathbf{w}) \approx f\left(\mathbf{w}^{(t)}\right)+\left\langle\mathbf{w}-\mathbf{w}^{(t)}, \nabla f\left(\mathbf{w}^{(t)}\right)\right\rangle
$$

- Hence, we want to minimize the approximation while staying close to $\mathbf{w}^{(t)}$ :

$$
\mathbf{w}^{(t+1)}=\underset{\mathbf{w}}{\operatorname{argmin}} \frac{1}{2}\left\|\mathbf{w}-\mathbf{w}^{(t)}\right\|^{2}+\eta\left(f\left(\mathbf{w}^{(t)}\right)+\left\langle\mathbf{w}-\mathbf{w}^{(t)}, \nabla f\left(\mathbf{w}^{(t)}\right)\right\rangle\right) .
$$

## Gradient Descent

- Initialize $\mathbf{w}^{(1)}=\mathbf{0}$
- Update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla f\left(\mathbf{w}^{(t)}\right)
$$

- Output $\overline{\mathbf{w}}=\frac{1}{T} \sum_{t=1}^{T} \mathbf{w}^{(t)}$


## Sub-Gradient Descent

Replace gradients with sub-gradients:

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \mathbf{v}_{t}
$$

where $\mathbf{v}_{t} \in \partial f\left(\mathbf{w}^{(t)}\right)$

## Analyzing sub-gradient descent

## Lemma

$$
\begin{aligned}
& \sum_{t=1}^{T}\left(f\left(\mathbf{w}^{(t)}\right)-f\left(\mathbf{w}^{\star}\right)\right) \leq \sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \\
& \quad=\frac{\left\|\mathbf{w}^{(1)}-\mathbf{w}^{\star}\right\|^{2}-\left\|\mathbf{w}^{(T+1)}-\mathbf{w}^{\star}\right\|^{2}}{2 \eta}+\frac{\eta}{2} \sum_{t=1}^{T}\left\|\mathbf{v}_{t}\right\|^{2}
\end{aligned}
$$

## Proof:

- The inequality is by the definition of sub-gradients
- The equality follows from the definition of the update using algebraic manipulations


## Analyzing sub-gradient descent for Lipschitz functions

- Since $f$ is convex and $\rho$-Lipschitz, $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for every $t$


## Analyzing sub-gradient descent for Lipschitz functions

- Since $f$ is convex and $\rho$-Lipschitz, $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for every $t$
- Therefore,

$$
\frac{1}{T} \sum_{t=1}^{T}\left(f\left(\mathbf{w}_{t}\right)-f\left(\mathbf{w}^{\star}\right)\right) \leq \frac{\left\|\mathbf{w}^{\star}\right\|^{2}}{2 \eta T}+\frac{\eta \rho^{2}}{2}
$$

## Analyzing sub-gradient descent for Lipschitz functions

- Since $f$ is convex and $\rho$-Lipschitz, $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for every $t$
- Therefore,

$$
\frac{1}{T} \sum_{t=1}^{T}\left(f\left(\mathbf{w}_{t}\right)-f\left(\mathbf{w}^{\star}\right)\right) \leq \frac{\left\|\mathbf{w}^{\star}\right\|^{2}}{2 \eta T}+\frac{\eta \rho^{2}}{2}
$$

- For every $\mathbf{w}^{\star}$, if $T \geq \frac{\left\|\mathbf{w}^{*}\right\|^{2} \rho^{2}}{\epsilon^{2}}$, and $\eta=\sqrt{\frac{\left\|\mathbf{w}^{*}\right\|^{2}}{\rho^{2} T}}$, then the right-hand side is at most $\epsilon$


## Analyzing sub-gradient descent for Lipschitz functions

- Since $f$ is convex and $\rho$-Lipschitz, $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for every $t$
- Therefore,

$$
\frac{1}{T} \sum_{t=1}^{T}\left(f\left(\mathbf{w}_{t}\right)-f\left(\mathbf{w}^{\star}\right)\right) \leq \frac{\left\|\mathbf{w}^{\star}\right\|^{2}}{2 \eta T}+\frac{\eta \rho^{2}}{2}
$$

- For every $\mathbf{w}^{\star}$, if $T \geq \frac{\left\|\mathbf{w}^{*}\right\|^{2} \rho^{2}}{\epsilon^{2}}$, and $\eta=\sqrt{\frac{\left\|\mathbf{w}^{*}\right\|^{2}}{\rho^{2} T}}$, then the right-hand side is at most $\epsilon$
- By convexity, $f(\overline{\mathbf{w}}) \leq \frac{1}{T} \sum_{t=1}^{T} f\left(\mathbf{w}_{t}\right)$, hence $f(\overline{\mathbf{w}})-f\left(\mathbf{w}^{\star}\right) \leq \epsilon$


## Analyzing sub-gradient descent for Lipschitz functions

- Since $f$ is convex and $\rho$-Lipschitz, $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for every $t$
- Therefore,

$$
\frac{1}{T} \sum_{t=1}^{T}\left(f\left(\mathbf{w}_{t}\right)-f\left(\mathbf{w}^{\star}\right)\right) \leq \frac{\left\|\mathbf{w}^{\star}\right\|^{2}}{2 \eta T}+\frac{\eta \rho^{2}}{2}
$$

- For every $\mathbf{w}^{\star}$, if $T \geq \frac{\left\|\mathbf{w}^{*}\right\|^{2} \rho^{2}}{\epsilon^{2}}$, and $\eta=\sqrt{\frac{\left\|\mathbf{w}^{*}\right\|^{2}}{\rho^{2} T}}$, then the right-hand side is at most $\epsilon$
- By convexity, $f(\overline{\mathbf{w}}) \leq \frac{1}{T} \sum_{t=1}^{T} f\left(\mathbf{w}_{t}\right)$, hence $f(\overline{\mathbf{w}})-f\left(\mathbf{w}^{\star}\right) \leq \epsilon$
- Corollary: Sub-gradient descent needs $\frac{\left\|\mathbf{w}^{*}\right\|^{2} \rho^{2}}{\epsilon^{2}}$ iterations to converge


## Example: Finding a Separating Hyperplane

Let $\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$ we would like to find a separating $\mathbf{w}$ :

$$
\forall i, \quad y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle>0
$$

Notation:

- Denote by $\mathbf{w}^{*}$ a separating hyperplane of unit norm and let $\gamma=\min _{i} y_{i}\left\langle\mathbf{w}^{*}, \mathbf{x}_{i}\right\rangle$
- W.l.o.g. assume $\left\|\mathbf{x}_{i}\right\|=1$ for every $i$.



## Separating Hyperplane using the Ellipsoid

- We can take the initial ball to be the unit ball
- The separation oracle looks for $i$ s.t. $y_{i}\left\langle\mathbf{w}^{(t)}, \mathbf{x}_{i}\right\rangle \leq 0$
- If there's no such $i$, we're done. Otherwise, the oracle returns $y_{i} \mathbf{x}_{i}$
- The algorithm stops after at most $2 d(2 d+2) \log (1 / \gamma)$ iterations


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex
- A sub-gradient of $f$ at $\mathbf{w}$ is $-y_{i} \mathbf{x}_{i}$ for some $i \in \operatorname{argmax}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex
- A sub-gradient of $f$ at $\mathbf{w}$ is $-y_{i} \mathbf{x}_{i}$ for some $i \in \operatorname{argmax}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$
- $f$ is 1-Lipschitz


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex
- A sub-gradient of $f$ at $\mathbf{w}$ is $-y_{i} \mathbf{x}_{i}$ for some $i \in \operatorname{argmax}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$
- $f$ is 1 -Lipschitz
- $f\left(\mathbf{w}^{*}\right)=-\gamma$


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex
- A sub-gradient of $f$ at $\mathbf{w}$ is $-y_{i} \mathbf{x}_{i}$ for some $i \in \operatorname{argmax}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$
- $f$ is 1 -Lipschitz
- $f\left(\mathbf{w}^{*}\right)=-\gamma$
- Therefore, after $t>\frac{1}{\gamma^{2}}$ iterations, we have $f\left(\mathbf{w}^{(t)}\right)<f\left(\mathbf{w}^{*}\right)+\gamma=0$


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex
- A sub-gradient of $f$ at $\mathbf{w}$ is $-y_{i} \mathbf{x}_{i}$ for some $i \in \operatorname{argmax}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$
- $f$ is 1 -Lipschitz
- $f\left(\mathbf{w}^{*}\right)=-\gamma$
- Therefore, after $t>\frac{1}{\gamma^{2}}$ iterations, we have $f\left(\mathbf{w}^{(t)}\right)<f\left(\mathbf{w}^{*}\right)+\gamma=0$
- So, $\mathbf{w}^{(t)}$ is a separating hyperplane


## Separating Hyperplane using Sub-gradient Descent

Consider the problem:

$$
\min _{\mathbf{w}} f(\mathbf{w}) \quad \text { where } \quad f(\mathbf{w})=\max _{i}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle
$$

Observe:

- $f$ is convex
- A sub-gradient of $f$ at $\mathbf{w}$ is $-y_{i} \mathbf{x}_{i}$ for some $i \in \operatorname{argmax}-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$
- $f$ is 1 -Lipschitz
- $f\left(\mathbf{w}^{*}\right)=-\gamma$
- Therefore, after $t>\frac{1}{\gamma^{2}}$ iterations, we have $f\left(\mathbf{w}^{(t)}\right)<f\left(\mathbf{w}^{*}\right)+\gamma=0$
- So, $\mathbf{w}^{(t)}$ is a separating hyperplane
- The resulting algorithm is closely related to the Batch Perceptron


## The Batch Perceptron

- Initialize, $\mathbf{w}^{(1)}=\mathbf{0}$
- While exists $i$ s.t. $y_{i}\left\langle\mathbf{w}^{(t)}, \mathbf{x}_{i}\right\rangle \leq 0$ update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}+y_{i} \mathbf{x}_{i}
$$

## The Batch Perceptron

- Initialize, $\mathbf{w}^{(1)}=\mathbf{0}$
- While exists $i$ s.t. $y_{i}\left\langle\mathbf{w}^{(t)}, \mathbf{x}_{i}\right\rangle \leq 0$ update

$$
\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}+y_{i} \mathbf{x}_{i}
$$

Exercise: why did we eliminate $\eta$ ?

## Ellipsoid vs. Sub-gradient

For $f$ convex and $\rho$-Lipschitz:

|  | iterations | cost of iteration |
| :--- | :---: | :---: |
| Ellipsoid | $d^{2} \log \left(\frac{\rho\left\\|\mathbf{w}^{*}\right\\|}{\epsilon}\right)$ | $d^{2}+$ "gradient oracle" |
| Sub-gradient descent | $\frac{\left\\|\mathbf{w}^{*}\right\\|^{2} \rho^{2}}{\epsilon^{2}}$ | $d+$ "gradient oracle" |

## Ellipsoid vs. Sub-gradient

For $f$ convex and $\rho$-Lipschitz:

|  | iterations | cost of iteration |
| :--- | :---: | :---: |
| Ellipsoid | $d^{2} \log \left(\frac{\rho\left\\|\mathbf{w}^{*}\right\\|}{\epsilon}\right)$ | $d^{2}+$ "gradient oracle" |
| Sub-gradient descent | $\frac{\left\\|\mathbf{w}^{*}\right\\|^{2} \rho^{2}}{\epsilon^{2}}$ | $d+"$ gradient oracle" |

For separating hyperplane:

|  | iterations | cost of iteration |
| :--- | :---: | :---: |
| Ellipsoid | $d^{2} \log \left(\frac{1}{\gamma}\right)$ | $d^{2}+d m$ |
| Sub-gradient descent | $\frac{1}{\gamma^{2}}$ | $d m$ |

## Outline

## (1) Convexity

(2) Convex Optimization

- Ellipsoid
- Gradient Descent
(3) Convex Learning Problems

4 Surrogate Loss Functions
(5) Learning Using Stochastic Gradient Descent

## Convex Learning Problems

## Definition (Convex Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called convex if the hypothesis class $\mathcal{H}$ is a convex set and for all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex function (where, for any $z, \ell(\cdot, z)$ denotes the function $f: \mathcal{H} \rightarrow \mathbb{R}$ defined by $f(\mathbf{w})=\ell(\mathbf{w}, z))$.

## Convex Learning Problems

## Definition (Convex Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called convex if the hypothesis class $\mathcal{H}$ is a convex set and for all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex function (where, for any $z, \ell(\cdot, z)$ denotes the function $f: \mathcal{H} \rightarrow \mathbb{R}$ defined by $f(\mathbf{w})=\ell(\mathbf{w}, z))$.

- The $E R M_{\mathcal{H}}$ problem w.r.t. a convex learning problem is a convex optimization problem: $\min _{\mathbf{w} \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^{m} \ell\left(\mathbf{w}, z_{i}\right)$


## Convex Learning Problems

## Definition (Convex Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called convex if the hypothesis class $\mathcal{H}$ is a convex set and for all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex function (where, for any $z, \ell(\cdot, z)$ denotes the function $f: \mathcal{H} \rightarrow \mathbb{R}$ defined by $f(\mathbf{w})=\ell(\mathbf{w}, z))$.

- The $E R M_{\mathcal{H}}$ problem w.r.t. a convex learning problem is a convex optimization problem: $\min _{\mathbf{w} \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^{m} \ell\left(\mathbf{w}, z_{i}\right)$
- Example - least squares: $\mathcal{H}=\mathbb{R}^{d}, Z=\mathbb{R}^{d} \times \mathbb{R}$, $\ell(\mathbf{w},(\mathbf{x}, y))=(\langle\mathbf{w}, \mathbf{x}\rangle-y)^{2}$


## Learnability of convex learning problems

- Claim: Not all convex learning problems over $\mathbb{R}^{d}$ are learnable
- The intuitive reason is numerical stability
- But, with two additional mild conditions, we obtain learnability
- $\mathcal{H}$ is bounded
- The loss function (or its gradient) is Lipschitz


## Convex-Lipschitz-bounded learning problem

## Definition (Convex-Lipschitz-Bounded Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called Convex-Lipschitz-Bounded, with parameters $\rho, B$ if the following holds:

- The hypothesis class $\mathcal{H}$ is a convex set and for all $\mathbf{w} \in \mathcal{H}$ we have $\|\mathbf{w}\| \leq B$.
- For all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex and $\rho$-Lipschitz function.


## Convex-Lipschitz-bounded learning problem

## Definition (Convex-Lipschitz-Bounded Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called Convex-Lipschitz-Bounded, with parameters $\rho, B$ if the following holds:

- The hypothesis class $\mathcal{H}$ is a convex set and for all $\mathbf{w} \in \mathcal{H}$ we have $\|\mathbf{w}\| \leq B$.
- For all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex and $\rho$-Lipschitz function.

Example:

- $\mathcal{H}=\left\{\mathbf{w} \in \mathbb{R}^{d}:\|\mathbf{w}\| \leq B\right\}$
- $\mathcal{X}=\left\{\mathbf{x} \in \mathbb{R}^{d}:\|\mathbf{x}\| \leq \rho\right\}, \mathcal{Y}=\mathbb{R}$,
- $\ell(\mathbf{w},(\mathbf{x}, y))=|\langle\mathbf{w}, \mathbf{x}\rangle-y|$


## Convex-Smooth-bounded learning problem

A function $f$ is $\beta$-smooth if it is differentiable and its gradient is $\beta$-Lipschitz.

## Definition (Convex-Smooth-Bounded Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called Convex-Smooth-Bounded, with parameters $\beta, B$ if the following holds:

- The hypothesis class $\mathcal{H}$ is a convex set and for all $\mathbf{w} \in \mathcal{H}$ we have $\|\mathbf{w}\| \leq B$.
- For all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex, non-negative, and $\beta$-smooth function.


## Convex-Smooth-bounded learning problem

A function $f$ is $\beta$-smooth if it is differentiable and its gradient is $\beta$-Lipschitz.

## Definition (Convex-Smooth-Bounded Learning Problem)

A learning problem, $(\mathcal{H}, Z, \ell)$, is called Convex-Smooth-Bounded, with parameters $\beta, B$ if the following holds:

- The hypothesis class $\mathcal{H}$ is a convex set and for all $\mathbf{w} \in \mathcal{H}$ we have $\|\mathbf{w}\| \leq B$.
- For all $z \in Z$, the loss function, $\ell(\cdot, z)$, is a convex, non-negative, and $\beta$-smooth function.

Example:

- $\mathcal{H}=\left\{\mathbf{w} \in \mathbb{R}^{d}:\|\mathbf{w}\| \leq B\right\}$
- $\mathcal{X}=\left\{\mathbf{x} \in \mathbb{R}^{d}:\|\mathbf{x}\| \leq \beta / 2\right\}, \mathcal{Y}=\mathbb{R}$,
- $\ell(\mathbf{w},(\mathbf{x}, y))=(\langle\mathbf{w}, \mathbf{x}\rangle-y)^{2}$


## Learnability

We will later show that all Convex-Lipschitz-Bounded and Convex-Smooth-Bounded learning problems are learnable, with sample complexity that depends only on $\epsilon, \delta, B$, and $\rho$ or $\beta$.

## Outline

(1) Convexity
(2) Convex Optimization

- Ellipsoid
- Gradient Descent
(3) Convex Learning Problems

4 Surrogate Loss Functions

## (5) Learning Using Stochastic Gradient Descent

## Surrogate Loss Functions

- In many natural cases, the loss function is not convex
- For example, the $0-1$ loss for halfspaces

$$
\ell^{0-1}(\mathbf{w},(\mathbf{x}, y))=\mathbb{1}_{[y \neq \operatorname{sign}(\langle\mathbf{w}, \mathbf{x}\rangle)]}=\mathbb{1}_{[y\langle\mathbf{w}, \mathbf{x}\rangle \leq 0]}
$$

- Non-convex loss function usually yields intractable learning problems
- Popular approach: circumvent hardness by upper bounding the non-convex loss function using a convex surrogate loss function


## Hinge-loss

$$
\ell^{\text {hinge }}(\mathbf{w},(\mathbf{x}, y)) \stackrel{\text { def }}{=} \max \{0,1-y\langle\mathbf{w}, \mathbf{x}\rangle\}
$$



## Error Decomposition Revisited

- Suppose we have a learner for the hinge-loss that guarantees:

$$
L_{\mathcal{D}}^{\text {hinge }}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

## Error Decomposition Revisited

- Suppose we have a learner for the hinge-loss that guarantees:

$$
L_{\mathcal{D}}^{\text {hinge }}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

- Using the surrogate property,

$$
L_{\mathcal{D}}^{0-1}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

## Error Decomposition Revisited

- Suppose we have a learner for the hinge-loss that guarantees:

$$
L_{\mathcal{D}}^{\text {hinge }}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

- Using the surrogate property,

$$
L_{\mathcal{D}}^{0-1}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

- We can further rewrite the upper bound as:

$$
\begin{aligned}
L_{\mathcal{D}}^{0-1}(A(S)) & \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{0-1}(\mathbf{w})+\left(\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})-\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{0-1}(\mathbf{w})\right)+\epsilon \\
& =\epsilon_{\text {approximation }}+\epsilon_{\text {optimization }}+\epsilon_{\text {estimation }}
\end{aligned}
$$

## Error Decomposition Revisited

- Suppose we have a learner for the hinge-loss that guarantees:

$$
L_{\mathcal{D}}^{\text {hinge }}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

- Using the surrogate property,

$$
L_{\mathcal{D}}^{0-1}(A(S)) \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})+\epsilon
$$

- We can further rewrite the upper bound as:

$$
\begin{aligned}
L_{\mathcal{D}}^{0-1}(A(S)) & \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{0-1}(\mathbf{w})+\left(\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{\text {hinge }}(\mathbf{w})-\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}^{0-1}(\mathbf{w})\right)+\epsilon \\
& =\epsilon_{\text {approximation }}+\epsilon_{\text {optimization }}+\epsilon_{\text {estimation }}
\end{aligned}
$$

- The optimization error is a result of our inability to minimize the training loss with respect to the original loss.


## Outline

## (1) Convexity

(2) Convex Optimization

- Ellipsoid
- Gradient Descent
(3) Convex Learning Problems

4 Surrogate Loss Functions
(5) Learning Using Stochastic Gradient Descent

## Learning Using Stochastic Gradient Descent

- Consider a convex-Lipschitz-bounded learning problem.


## Learning Using Stochastic Gradient Descent

- Consider a convex-Lipschitz-bounded learning problem.
- Recall: our goal is to (probably approximately) solve:

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

## Learning Using Stochastic Gradient Descent

- Consider a convex-Lipschitz-bounded learning problem.
- Recall: our goal is to (probably approximately) solve:

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- So far, learning was based on the empirical risk, $L_{S}(\mathbf{w})$


## Learning Using Stochastic Gradient Descent

- Consider a convex-Lipschitz-bounded learning problem.
- Recall: our goal is to (probably approximately) solve:

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- So far, learning was based on the empirical risk, $L_{S}(\mathbf{w})$
- We now consider directly minimizing $L_{\mathcal{D}}(\mathbf{w})$


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$
- Observe: $\nabla L_{\mathcal{D}}(\mathbf{w})=\mathbb{E}_{z \sim \mathcal{D}}[\nabla \ell(\mathbf{w}, z)]$


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$
- Observe: $\nabla L_{\mathcal{D}}(\mathbf{w})=\mathbb{E}_{z \sim \mathcal{D}}[\nabla \ell(\mathbf{w}, z)]$
- We can't calculate $\nabla L_{\mathcal{D}}(\mathbf{w})$ because we don't know $\mathcal{D}$


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$
- Observe: $\nabla L_{\mathcal{D}}(\mathbf{w})=\mathbb{E}_{z \sim \mathcal{D}}[\nabla \ell(\mathbf{w}, z)]$
- We can't calculate $\nabla L_{\mathcal{D}}(\mathbf{w})$ because we don't know $\mathcal{D}$
- But we can estimate it by $\nabla \ell(\mathbf{w}, z)$ for $z \sim \mathcal{D}$


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$
- Observe: $\nabla L_{\mathcal{D}}(\mathbf{w})=\mathbb{E}_{z \sim \mathcal{D}}[\nabla \ell(\mathbf{w}, z)]$
- We can't calculate $\nabla L_{\mathcal{D}}(\mathbf{w})$ because we don't know $\mathcal{D}$
- But we can estimate it by $\nabla \ell(\mathbf{w}, z)$ for $z \sim \mathcal{D}$
- If we take a step in the direction $\mathbf{v}=\nabla \ell(\mathbf{w}, z)$ then in expectation we're moving in the right direction


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$
- Observe: $\nabla L_{\mathcal{D}}(\mathbf{w})=\mathbb{E}_{z \sim \mathcal{D}}[\nabla \ell(\mathbf{w}, z)]$
- We can't calculate $\nabla L_{\mathcal{D}}(\mathbf{w})$ because we don't know $\mathcal{D}$
- But we can estimate it by $\nabla \ell(\mathbf{w}, z)$ for $z \sim \mathcal{D}$
- If we take a step in the direction $\mathbf{v}=\nabla \ell(\mathbf{w}, z)$ then in expectation we're moving in the right direction
- In other words, $\mathbf{v}$ is an unbiased estimate of the gradient


## Stochastic Gradient Descent

$$
\min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w}) \quad \text { where } \quad L_{\mathcal{D}}(\mathbf{w})=\underset{z \sim \mathcal{D}}{\mathbb{E}}[\ell(\mathbf{w}, z)]
$$

- Recall the gradient descent method in which we initialize $\mathbf{w}^{(1)}=\mathbf{0}$ and update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \nabla L_{\mathcal{D}}(\mathbf{w})$
- Observe: $\nabla L_{\mathcal{D}}(\mathbf{w})=\mathbb{E}_{z \sim \mathcal{D}}[\nabla \ell(\mathbf{w}, z)]$
- We can't calculate $\nabla L_{\mathcal{D}}(\mathbf{w})$ because we don't know $\mathcal{D}$
- But we can estimate it by $\nabla \ell(\mathbf{w}, z)$ for $z \sim \mathcal{D}$
- If we take a step in the direction $\mathbf{v}=\nabla \ell(\mathbf{w}, z)$ then in expectation we're moving in the right direction
- In other words, $\mathbf{v}$ is an unbiased estimate of the gradient
- We'll show that this is good enough


## Stochastic Gradient Descent

- initialize: $\mathbf{w}^{(1)}=\mathbf{0}$
- for $t=1,2, \ldots, T$
- choose $z_{t} \sim \mathcal{D}$
- let $\mathbf{v}_{t} \in \partial \ell\left(\mathbf{w}^{(t)}, z_{t}\right)$ update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \mathbf{v}_{t}$
- output $\overline{\mathbf{w}}=\frac{1}{T} \sum_{t=1}^{T} \mathbf{w}^{(t)}$


## Stochastic Gradient Descent

- initialize: $\mathbf{w}^{(1)}=\mathbf{0}$
- for $t=1,2, \ldots, T$
- choose $z_{t} \sim \mathcal{D}$
- let $\mathbf{v}_{t} \in \partial \ell\left(\mathbf{w}^{(t)}, z_{t}\right)$ update $\mathbf{w}^{(t+1)}=\mathbf{w}^{(t)}-\eta \mathbf{v}_{t}$
- output $\overline{\mathbf{w}}=\frac{1}{T} \sum_{t=1}^{T} \mathbf{w}^{(t)}$



## Analyzing SGD for convex-Lipschitz-bounded

By algebraic manipulations, for any sequence of $\mathbf{v}_{1}, \ldots, \mathbf{v}_{T}$, and any $\mathbf{w}^{\star}$,

$$
\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle=\frac{\left\|\mathbf{w}^{(1)}-\mathbf{w}^{\star}\right\|^{2}-\left\|\mathbf{w}^{(T+1)}-\mathbf{w}^{\star}\right\|^{2}}{2 \eta}+\frac{\eta}{2} \sum_{t=1}^{T}\left\|\mathbf{v}_{t}\right\|^{2}
$$

## Analyzing SGD for convex-Lipschitz-bounded

By algebraic manipulations, for any sequence of $\mathbf{v}_{1}, \ldots, \mathbf{v}_{T}$, and any $\mathbf{w}^{\star}$,

$$
\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle=\frac{\left\|\mathbf{w}^{(1)}-\mathbf{w}^{\star}\right\|^{2}-\left\|\mathbf{w}^{(T+1)}-\mathbf{w}^{\star}\right\|^{2}}{2 \eta}+\frac{\eta}{2} \sum_{t=1}^{T}\left\|\mathbf{v}_{t}\right\|^{2}
$$

Assume that $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for all $t$ and that $\left\|\mathbf{w}^{\star}\right\| \leq B$ we obtain

$$
\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \leq \frac{B^{2}}{2 \eta}+\frac{\eta \rho^{2} T}{2}
$$

## Analyzing SGD for convex-Lipschitz-bounded

By algebraic manipulations, for any sequence of $\mathbf{v}_{1}, \ldots, \mathbf{v}_{T}$, and any $\mathbf{w}^{\star}$,

$$
\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle=\frac{\left\|\mathbf{w}^{(1)}-\mathbf{w}^{\star}\right\|^{2}-\left\|\mathbf{w}^{(T+1)}-\mathbf{w}^{\star}\right\|^{2}}{2 \eta}+\frac{\eta}{2} \sum_{t=1}^{T}\left\|\mathbf{v}_{t}\right\|^{2}
$$

Assume that $\left\|\mathbf{v}_{t}\right\| \leq \rho$ for all $t$ and that $\left\|\mathbf{w}^{\star}\right\| \leq B$ we obtain

$$
\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \leq \frac{B^{2}}{2 \eta}+\frac{\eta \rho^{2} T}{2}
$$

In particular, for $\eta=\sqrt{\frac{B^{2}}{\rho^{2} T}}$ we get

$$
\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \leq B \rho \sqrt{T}
$$

## Analyzing SGD for convex-Lipschitz-bounded

Taking expectation of both sides w.r.t. the randomness of choosing $z_{1}, \ldots, z_{T}$ we obtain:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle\right] \leq B \rho \sqrt{T} .
$$

## Analyzing SGD for convex-Lipschitz-bounded

Taking expectation of both sides w.r.t. the randomness of choosing $z_{1}, \ldots, z_{T}$ we obtain:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle\right] \leq B \rho \sqrt{T}
$$

The law of total expectation: for every two random variables $\alpha, \beta$, and a function $g, \mathbb{E}_{\alpha}[g(\alpha)]=\mathbb{E}_{\beta} \mathbb{E}_{\alpha}[g(\alpha) \mid \beta]$.

## Analyzing SGD for convex-Lipschitz-bounded

Taking expectation of both sides w.r.t. the randomness of choosing $z_{1}, \ldots, z_{T}$ we obtain:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle\right] \leq B \rho \sqrt{T} .
$$

The law of total expectation: for every two random variables $\alpha, \beta$, and a function $g, \mathbb{E}_{\alpha}[g(\alpha)]=\mathbb{E}_{\beta} \mathbb{E}_{\alpha}[g(\alpha) \mid \beta]$. Therefore

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle\right]=\underset{z_{1}, \ldots, z_{t-1}}{\mathbb{E}} \underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \mid z_{1}, \ldots, z_{t-1}\right]
$$

## Analyzing SGD for convex-Lipschitz-bounded

Taking expectation of both sides w.r.t. the randomness of choosing $z_{1}, \ldots, z_{T}$ we obtain:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle\right] \leq B \rho \sqrt{T} .
$$

The law of total expectation: for every two random variables $\alpha, \beta$, and a function $g, \mathbb{E}_{\alpha}[g(\alpha)]=\mathbb{E}_{\beta} \mathbb{E}_{\alpha}[g(\alpha) \mid \beta]$. Therefore

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle\right]=\underset{z_{1}, \ldots, z_{t-1}}{\mathbb{E}} \underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \mid z_{1}, \ldots, z_{t-1}\right] .
$$

Once we know $\beta$ the value of $\mathbf{w}^{(t)}$ is not random, hence,

$$
\begin{aligned}
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \mathbf{v}_{t}\right\rangle \mid z_{1}, \ldots, z_{t-1}\right] & =\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \underset{z_{t}}{\mathbb{E}}\left[\nabla \ell\left(\mathbf{w}^{(t)}, z_{t}\right)\right]\right\rangle \\
& =\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \nabla L_{\mathcal{D}}\left(\mathbf{w}^{(t)}\right)\right\rangle
\end{aligned}
$$

## Analyzing SGD for convex-Lipschitz-bounded

We got:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \nabla L_{\mathcal{D}}\left(\mathbf{w}^{(t)}\right)\right\rangle\right] \leq B \rho \sqrt{T}
$$

## Analyzing SGD for convex-Lipschitz-bounded

We got:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \nabla L_{\mathcal{D}}\left(\mathbf{w}^{(t)}\right)\right\rangle\right] \leq B \rho \sqrt{T}
$$

By convexity, this means

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left(L_{\mathcal{D}}\left(\mathbf{w}^{(t)}\right)-L_{\mathcal{D}}\left(\mathbf{w}^{\star}\right)\right)\right] \leq B \rho \sqrt{T}
$$

## Analyzing SGD for convex-Lipschitz-bounded

We got:

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left\langle\mathbf{w}^{(t)}-\mathbf{w}^{\star}, \nabla L_{\mathcal{D}}\left(\mathbf{w}^{(t)}\right)\right\rangle\right] \leq B \rho \sqrt{T}
$$

By convexity, this means

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[\sum_{t=1}^{T}\left(L_{\mathcal{D}}\left(\mathbf{w}^{(t)}\right)-L_{\mathcal{D}}\left(\mathbf{w}^{\star}\right)\right)\right] \leq B \rho \sqrt{T}
$$

Dividing by $T$ and using convexity again,

$$
\underset{z_{1}, \ldots, z_{T}}{\mathbb{E}}\left[L_{\mathcal{D}}\left(\frac{1}{T} \sum_{t=1}^{T} \mathbf{w}^{(t)}\right)\right] \leq L_{\mathcal{D}}\left(\mathbf{w}^{\star}\right)+\frac{B \rho}{\sqrt{T}}
$$

## Learning convex-Lipschitz-bounded problems using SGD

## Corollary

Consider a convex-Lipschitz-bounded learning problem with parameters $\rho, B$. Then, for every $\epsilon>0$, if we run the $S G D$ method for minimizing $L_{\mathcal{D}}(\mathbf{w})$ with a number of iterations (i.e., number of examples)

$$
T \geq \frac{B^{2} \rho^{2}}{\epsilon^{2}}
$$

and with $\eta=\sqrt{\frac{B^{2}}{\rho^{2} T}}$, then the output of SGD satisfies:

$$
\mathbb{E}\left[L_{\mathcal{D}}(\overline{\mathbf{w}})\right] \leq \min _{\mathbf{w} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{w})+\epsilon .
$$

## Summary

- Convex optimization
- Convex learning problems
- Learning using SGD

