

| Polytopes <br> - A polytope is intersection of halfspaces and hyperplanes <br> - Mathematical representation: $\begin{array}{r} C=\left\{x \in \mathbb{R}^{n}: s_{i}^{T} x \leq r_{i} \text { for } i \in\{1, \ldots, m\}\right. \text { and } \\ \left.s_{i}^{T} x=r_{i} \text { for } i \in\{m+1, \ldots, p\}\right\} \end{array}$ <br> - Polytopes convex since intersection of convex sets | Cones <br> - A set $K$ is a cone if for all $x \in K$ and $\alpha \geq 0: \alpha x \in K$ <br> - If $x$ is in cone $K$, so is entire ray from origin passing through $x$ : <br> - Examples: <br> Cone  <br> Cone  <br> Not cone |
| :---: | :---: |
| Convex cones <br> - Cones can be convex or nonconvex: <br> Nonconvex cone  <br> Convex cone <br> - Convex cone examples: <br> - Linear subspaces $\left\{x \in \mathbb{R}^{n}: A x=0\right\}$ (but not affine subspaces) <br> - Halfspaces based on linear (not affine) hyperplanes $\left\{x: s^{T} x \leq 0\right\}$ <br> - Positive semi-definite matrices $\left\{X \in \mathbb{R}^{n \times n}: X\right.$ symmetric and $z^{T} X z \geq 0$ for all $\left.z \in \mathbb{R}^{n}\right\}$ <br> - Nonnegative orthant $\left\{x \in \mathbb{R}^{n}: x \geq 0\right\}$ <br> - Second order cone $\left\{(x, r) \in \mathbb{R}^{n} \times \mathbb{R}:\\|x\\|_{2} \leq r\right\}$ | Sublevel sets <br> - Suppose that $g: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a real-valued function <br> - The (0th) sublevel set of $g$ is defined as $S:=\left\{x \in \mathbb{R}^{n}: g(x) \leq 0\right\}$ <br> - Example: construction giving 1D interval $S=[a, b]$ <br> - $S$ is a convex set if $g$ is a convex function <br> - $S$ is not necessarily nonconvex although $g$ is |
| Sublevel sets - Examples <br> - Levelset of convex quadratic function <br> $\left\{x \in \mathbb{R}^{n}: \frac{1}{2} x^{T} P x+q^{T} x+r \leq 0\right\}$, with $P$ positive definite <br> - Norm balls $\left\{x \in \mathbb{R}^{n}:\\|x\\|-r \leq 0\right\}$ <br> - Second-order cone $\left\{(x, r) \in \mathbb{R}^{n} \times \mathbb{R}:\\|x\\|_{2}-r \leq 0\right\}$ <br> - Halfspaces $\left\{x \in \mathbb{R}^{n}: c^{T} x-r \leq 0\right\}$ | Outline <br> - Definition and convex hull <br> - Examples of convex sets <br> - Convexity preserving operations <br> - Concluding convexity - Examples <br> - Separating and supporting hyperplanes |
| Convexity preserving operations <br> - Intersection (but not union) <br> - Affine image and inverse affine image of a set | Intersection and union <br> - Intersection $C=C_{1} \cap C_{2}$ means $x \in C$ if $x \in C_{1}$ and $x \in C_{2}$ <br> - Union $C=C_{1} \cup C_{2}$ means $x \in C$ if $x \in C_{1}$ or $x \in C_{2}$ <br> Intersection <br> Union <br> - Intersection of any number of, e.g., infinite, convex sets is convex <br> - Union of convex sets need not be convex |

- Let $L(x)=A x+b$ be an affine mapping defined by
- matrix $A \in \mathbb{R}^{m \times n}$
- vector $b \in \mathbb{R}^{m}$
- Let $C$ be a convex set in $\mathbb{R}^{n}$ then the image set of $C$ under $L$

$$
\{A x+b: x \in C\}
$$

is convex

- Let $D$ be a convex set in $\mathbb{R}^{m}$ then the inverse image of $D$ under $L$

$$
\{x: A x+b \in D\}
$$

is convex

- Definition and convex hull
- Examples of convex sets
- Convexity preserving operations
- Concluding convexity - Examples
- Separating and supporting hyperplanes
- Use convexity definition
- Show that set is sublevel set of a convex function
- Show that set constructed by convexity preserving operations


## Example - Nonnegative orthant

- Nonnegative orthant is set $C=\left\{x \in \mathbb{R}^{n}: x \geq 0\right\}$
- Prove convexity from definition:
- Let $x \geq 0$ and $y \geq 0$ be arbitrary points in $C$
- For all $\theta \in[0,1]$ :

$$
\theta x \geq 0 \quad \text { and } \quad(1-\theta) y \geq 0
$$

- All convex combinations therefore also satisfy

$$
\theta x+(1-\theta) y \geq 0
$$

i.e., they belongs to $C$ and the set is convex

Ways to conclude convexity

## Example - Positive semidefinite cone

- The positive semidefinite (PSD) cone is
$\left\{X \in \mathbb{R}^{n \times n}: X\right.$ symmetric $\} \bigcap\left\{X \in \mathbb{R}^{n \times n}: z^{T} X z \geq 0\right.$ for all $\left.z \in \mathbb{R}^{n}\right\}$
- This can be written as the following intersection over all $z \in \mathbb{R}^{n}$

$$
\left\{X \in \mathbb{R}^{n \times n}: X \text { symmetric }\right\} \bigcap_{z \in \mathbb{R}^{n}}\left\{X \in \mathbb{R}^{n \times n}: z^{T} X z \geq 0\right\}
$$

which, by noting that $z^{T} X z=\operatorname{tr}\left(z^{T} X z\right)=\operatorname{tr}\left(z z^{T} X\right)$, is equal to

$$
\left\{X \in \mathbb{R}^{n \times n}: X \text { symmetric }\right\} \bigcap_{z \in \mathbb{R}^{n}}\left\{X \in \mathbb{R}^{n \times n}: \operatorname{tr}\left(z z^{T} X\right) \geq 0\right\}
$$

where $\operatorname{tr}\left(z z^{T} X\right) \geq 0$ is a halfspace in $\mathbb{R}^{n \times n}$ (except when $z=0$ )

- The PSD cone is convex since it is intersection of
- symmetry set, which is a finite set of (convex) linear equalities
- an infinite number of (convex) halfspaces in $\mathbb{R}^{n \times n}$
- Notation: If $X$ belong to the PSD cone, we write $X \succeq 0$


## Example - Linear matrix inequality

- Let us consider a linear matrix inequality (LMI) of the form

$$
\left\{x \in \mathbb{R}^{k}: A+\sum_{i=1}^{k} x_{i} B_{i} \succeq 0\right\}
$$

where $A$ and $B_{i}$ are fixed matrices in $\mathbb{R}^{n \times n}$

- Convex since inverse image of PSD cone under affine mapping


## Outline

- Definition and convex hull
- Examples of convex sets
- Convexity preserving operations
- Concluding convexity - Examples
- Separating and supporting hyperplanes


## Separating hyperplane theorem

- Suppose that $C, D \subseteq \mathbb{R}^{n}$ are two non-intersecting convex sets
- Then there exists hyperplane with $C$ and $D$ in opposite halves

- Mathematical formulation: There exists $s \neq 0$ and $r$ such that

$$
\begin{array}{ll}
s^{T} x \leq r & \text { for all } x \in C \\
s^{T} x \geq r & \text { for all } x \in D
\end{array}
$$

- The hyperplane $\left\{x: s^{T} x=r\right\}$ is called separating hyperplane
- Suppose that $C, D \subseteq \mathbb{R}^{n}$ are non-intersecting closed and convex sets and that one of them is compact (closed and bounded)
- Then there exists hyperplane with strict separation

- Mathematical formulation: There exists $s \neq 0$ and $r$ such that

$$
\begin{array}{ll}
s^{T} x<r & \text { for all } x \in C \\
s^{T} x>r & \text { for all } x \in D
\end{array}
$$

- let $H$ be the intersection of all halfspaces containing $C$
- $\Rightarrow$ : obviously $x \in C \Rightarrow x \in H$
- $\Leftarrow$ : assume $x \notin C$, since $C$ closed and convex and $\{x\}$ compact singleton, there exists a strictly separating hyperplane, i.e., $x \notin H$ :


Supporting hyperplanes

- Supporting hyperplanes touch set and have full set on one side:

- We call the halfspace that contains the set supporting halfspace
- $s$ is called normal vector to $C$ at $x$
- Definition: Hyperplane $\left\{y: s^{T} y=r\right\}$ supports $C$ at $x \in \mathrm{bd} C$ if

$$
s^{T} x=r \quad \text { and } \quad s^{T} y \leq r \text { for all } y \in C
$$

## Normal cone operator

- Normal cone to $C$ at $x \in \operatorname{bd}(C)$ is set of normals at $x$

- Normal cone operator $N_{C}$ to $C$ takes point input and returns set - $x \in \operatorname{bd}(C) \cap C$ : set of normal vectors to supporting halfspaces
- $x \in \operatorname{int}(C)$ : returns zero set $\{0\}$
- $x \notin C$ : returns emptyset $\emptyset$
- Mathematical definition: The normal cone operator to a set $C$ is

$$
N_{C}(x)= \begin{cases}\left\{s: s^{T}(y-x) \leq 0 \text { for all } y \in C\right\} & \text { if } x \in C \\ \emptyset & \text { else }\end{cases}
$$

i.e., vectors that form obtuse angle between $s$ and all $y-x, y \in C$

- For all $x \in C$ : the $N_{C}$ outputs a set that contains 0

Supporting hyperplane theorem

Let $C$ be a nonempty convex set and let $x \in \operatorname{bd}(C)$. Then there exists a supporting hyperplane to $C$ at $x$.

- Does not exist for all point on boundary for nonconvex sets
- Many supporting hyperplanes exist for points of nonsmoothness





Show that $f(x):=\|x\|$ is convex from convexity definition

- Norms satisfy the triangle inequality

$$
\|u+v\| \leq\|u\|+\|v\|
$$

- For arbitrary $x, y$ and $\theta \in[0,1]$ :

$$
\begin{aligned}
f(\theta x+(1-\theta) y) & =\|\theta x+(1-\theta) y\| \\
& \leq\|\theta x\|+\|(1-\theta) y\| \\
& =\theta\|x\|+(1-\theta)\|y\| \\
& =\theta f(x)+(1-\theta) f(y)
\end{aligned}
$$

which is definition of convexity

- Proof uses triangle inequality and $\theta \in[0,1]$


## Example - Matrix fractional function

Show that the matrix fractional function is convex via its epigraph

- The matrix fractional function

$$
f(x, Y)= \begin{cases}x^{T} Y^{-1} x & \text { if } Y \succ 0 \\ \infty & \text { else }\end{cases}
$$

- The epigraph satisfies

$$
\begin{aligned}
\operatorname{epi} f(x, Y, t) & =\{(x, Y, t): f(x, Y) \leq t\} \\
& =\left\{(x, Y, t): x^{T} Y^{-1} x \leq t \text { and } Y \succ 0\right\}
\end{aligned}
$$

- Schur complement condition says for $Y \succ 0$ that

$$
x^{T} Y^{-1} x \leq t \quad \Leftrightarrow \quad\left[\begin{array}{cc}
Y & x \\
x^{T} & t
\end{array}\right] \succeq 0
$$

which is a (convex) linear matrix inequality (LMI) in ( $x, Y, t$ )

- Epigraph is intersection between LMI and positive definite cone


## Example - Composition with matrix

- Let
- $f: \mathbb{R}^{m} \rightarrow \overline{\mathbb{R}}$ be convex
- $L \in \mathbb{R}^{m \times n}$ be a matrix
then composition with a matrix

$$
(f \circ L)(x):=f(L x)
$$

is convex

- Vector composition with convex function and affine mappings


## Example - Image of function under linear mapping

- Let
- $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ be convex
- $L \in \mathbb{R}^{m \times n}$ be a matrix
then image function (sometimes called infimal postcomposition)

$$
(L f)(x):=\inf _{y}\{f(y): L y=x\}
$$

is convex

- Proof: Define

$$
h(x, y)=f(y)+\iota_{\{0\}}(L y-x)
$$

which is convex in $(x, y)$, then

$$
(L f)(x)=\inf _{y} h(x, y)
$$

which is convex since marginal of convex function

## Example - Nested composition

Show that: $f(x):=e^{\|L x-b\|_{2}^{3}}$ is convex where $L$ is matrix $b$ vector:

- Let

$$
g_{1}(u)=\|u\|_{2}, \quad g_{2}(u)=\left\{\begin{array}{ll}
0 & \text { if } u<0 \\
u^{3} & \text { if } u \geq 0
\end{array}, \quad g_{3}(u)=e^{u}\right.
$$

then $f(x)=g_{3}\left(g_{2}\left(g_{1}(L x-b)\right)\right)$

- $g_{1}(L x-b)$ convex: convex $g_{1}$ and $L x-b$ affine
- $g_{2}\left(g_{1}(L x-b)\right)$ convex: cvx nondecreasing $g_{2}$ and $\mathrm{cvx} g_{1}(L x-b)$
- $f(x)$ convex: convex nondecreasing $g_{3}$ and convex $g_{2}\left(g_{1}(L x-b)\right)$


## Example - Conjugate function

Show that the conjugate $f^{*}(s):=\sup _{x \in \mathbb{R}^{n}}\left(s^{T} x-f(x)\right)$ is convex:

- Define (uncountable) index set $J$ and $x_{j}$ such that $\cup_{j \in J} x_{j}=\mathbb{R}^{n}$
- Define $r_{j}:=f\left(x_{j}\right)$ and affine (in $s$ ): $a_{j}(s):=s^{T} x_{j}-r_{j}$
- Therefore $f^{*}(s)=\sup \left(a_{j}(s): j \in J\right)$
- Convex since supremum over family of convex (affine) functions
- Note convexity of $f^{*}$ not dependent on convexity of $f$
- Definition, epigraph, convex envelope
- First- and second-order conditions for convexity
- Convexity preserving operations
- Concluding convexity - Examples
- Strict and strong convexity
- Smoothness


## Strict convexity

- A function is strictly convex if

$$
f(\theta x+(1-\theta) y)<\theta f(x)+(1-\theta) f(y)
$$

for all $x \neq y$ and $\theta \in(0,1)$

- Convexity definition with strict inequality
- No flat (affine) regions
- Example: $f(x)=1 / x$ for $x>0$
$f(x)$



## Strong convexity

- Let $\sigma>0$
- A function $f$ is $\sigma$-strongly convex if $f-\frac{\sigma}{2}\|\cdot\|_{2}^{2}$ is convex
- Alternative equivalent definition of $\sigma$-strong convexity:
$f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) f(y)-\frac{\sigma}{2} \theta(1-\theta)\|x-y\|^{2}$
holds for every $x, y \in \mathbb{R}^{n}$ and $\theta \in[0,1]$
- Strongly convex functions are strictly convex and convex
- Example: $f$ 2-strongly convex since $f-\|\cdot\|_{2}^{2}$ convex:


## Uniqueness of minimizers

- Strictly (strongly) convex functions have unique minimizers
- Strictly convex functions may not have a minimizing point
- Strongly convex functions always have a unique minimizing point




## First-order condition for strict convexity

- Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be differentiable
- $f$ is strictly convex if and only if

$$
f(y)>f(x)+\nabla f(x)^{T}(y-x)
$$

for all $x, y \in \mathbb{R}^{n}$ where $x \neq y$


- Function $f$ has for all $x \in \mathbb{R}^{n}$ an affine minorizer that:
- has slope $s$ defined by $\nabla f$
- coincides with function $f$ only at $x$
- is supporting hyperplane to epigraph of $f$
- defines normal $(\nabla f(x),-1)$ to epigraph of $f$


## First-order condition for strong convexity

- Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be differentiable
- $f$ is $\sigma$-strongly convex with $\sigma>0$ if and only if

$$
f(y) \geq f(x)+\nabla f(x)^{T}(y-x)+\frac{\sigma}{2}\|x-y\|_{2}^{2}
$$

for all $x, y \in \mathbb{R}^{n}$


- Function $f$ has for all $x \in \mathbb{R}^{n}$ a quadratic minorizer that:
- has curvature defined by $\sigma$
- coincides with function $f$ at $x$
- defines normal $(\nabla f(x),-1)$ to epigraph of $f$


## Second-order condition for strict/strong convexity

Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be twice differentiable

- $f$ is strictly convex if

$$
\nabla^{2} f(x) \succ 0
$$

for all $x \in \mathbb{R}^{n}$ (i.e., the Hessian is positive definite)

- $f$ is $\sigma$-strongly convex if and only if

$$
\nabla^{2} f(x) \succeq \sigma I
$$

for all $x \in \mathbb{R}^{n}$

## Examples of strictly/strongly convex functions

Strictly convex

- $f(x)=-\log (x)+\iota_{>0}(x)$
- $f(x)=1 / x+\iota>0(x)$
- $f(x)=e^{-x}$

Strongly convex

- $f(x)=\frac{\lambda}{2}\|x\|_{2}^{2}$
- $f(x)=\frac{1}{2} x^{T} Q x$ where $Q$ positive definite
- $f(x)=f_{1}(x)+f_{2}(x)$ where $f_{1}$ strongly convex and $f_{2}$ convex
- $f(x)=f_{1}(x)+f_{2}(x)$ where $f_{1}, f_{2}$ strongly convex
- $f(x)=\frac{1}{2} x^{T} Q x+\iota_{C}(x)$ where $Q$ positive definite and $C$ convex

Proofs for two examples

Strict convexity of $f(x)=e^{-x}$ :

- $\nabla f(x)=-e^{-x}, \nabla^{2} f(x)=e^{-x}>0$ for all $x \in \mathbb{R}$

Strong convexity of $f(x)=\frac{1}{2} x^{T} Q x$ with $Q$ positive definite

- $\nabla f(x)=Q x, \nabla^{2} f(x)=Q \succeq \lambda_{\min }(Q) I$ where $\lambda_{\min }(Q)>0$


## Outline

- Definition, epigraph, convex envelope
- First- and second-order conditions for convexity
- Convexity preserving operations
- Concluding convexity - Examples
- Strict and strong convexity
- Smoothness


## Smoothness

- A function is called $\beta$-smooth if its gradient is $\beta$-Lipschitz

$$
\|\nabla f(x)-\nabla f(y)\|_{2} \leq \beta\|x-y\|_{2}
$$

for all $x, y \in \mathbb{R}^{n}$ (it is not necessarily convex)

- Alternative equivalent definition of $\beta$-smoothness

$$
f(\theta x+(1-\theta) y) \geq \theta f(x)+(1-\theta) f(y)-\frac{\beta}{2} \theta(1-\theta)\|x-y\|^{2}
$$

$$
f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) f(y)+\frac{\beta}{2} \theta(1-\theta)\|x-y\|^{2}
$$

hold for every $x, y \in \mathbb{R}^{n}$ and $\theta \in[0,1]$

- Smoothness does not imply convexity
- Example:



## First-order condition for smooth convex

- $f$ is $\beta$-smooth with $\beta \geq 0$ and convex if and only if

$$
\begin{aligned}
& f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|x-y\|_{2}^{2} \\
& f(y) \geq f(x)+\nabla f(x)^{T}(y-x)
\end{aligned}
$$

for all $x, y \in \mathbb{R}^{n}$


- Quadratic upper bounds and affine lower bound
- Bounds coincide with function $f$ at $x$
- Quadratic upper bound is called descent lemma

Second-order condition for smoothness

- Quadratic upper/lower bounds with curvatures defined by $\beta$
- Quadratic bounds coincide with function $f$ at $x$
for all $x, y \in \mathbb{R}^{n}$


Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be twice differentiable

- $f$ is $\beta$-smooth if and only if

$$
-\beta I \preceq \nabla^{2} f(x) \preceq \beta I
$$

for all $x \in \mathbb{R}^{n}$

- $f$ is $\beta$-smooth and convex if and only if

$$
0 \preceq \nabla^{2} f(x) \preceq \beta I
$$

for all $x \in \mathbb{R}^{n}$

## Composite optimization form

- We will consider optimization problem on composite form

$$
\underset{x}{\operatorname{minimize}} f(L x)+g(x)
$$

where $f$ and $g$ are convex functions and $L$ is a matrix

- Convex problem due to convexity preserving operations
- Can model constrained problems via indicator function
- This model format is suitable for many algorithms



## Monotonicity of subdifferential

- Subdifferential operator is monotone:

$$
\left(s_{x}-s_{y}\right)^{T}(x-y) \geq 0
$$

for all $s_{x} \in \partial f(x)$ and $s_{y} \in \partial f(y)$

- Proof: Add two copies of subdifferential definition

$$
f(y) \geq f(x)+s_{x}^{T}(y-x)
$$

with $x$ and $y$ swapped

- $\partial f: \mathbb{R} \rightarrow 2^{\mathbb{R}}$ : Minimum slope 0 and maximum slope $\infty$

- Let $A: \mathbb{R}^{n} \rightarrow 2^{\mathbb{R}^{n}}$ be monotone, i.e.:

$$
(u-v)^{T}(x-y) \geq 0
$$

for all $u \in A x$ and $v \in A y$

- If $n=1$, then $A=\partial f$ for some function $f: \mathbb{R} \rightarrow \mathbb{R} \cup\{\infty\}$
- If $n \geq 2$ there exist monotone $A$ that are not subdifferentials


## Minty's theorem

- Let $\partial f: \mathbb{R}^{n} \rightarrow 2^{\mathbb{R}^{n}}$ and $\alpha>0$
- $\partial f$ is maximally monotone if and only if range $(\alpha I+\partial f)=\mathbb{R}^{n}$



- Interpretation: No "holes" in gph $\partial f$


## Outline

- Subdifferential and subgradient - Definition and basic properties
- Monotonicity
- Examples
- Strong monotonicity and cocoercivity
- Fermat's rule
- Subdifferential calculus
- Optimality conditions
- Proximal operators


## Example - Absolute value

- The absolute value:

- Subdifferential
- For $x>0, f$ differentiable and $\nabla f(x)=1$, so $\partial f(x)=\{1\}$
- For $x<0, f$ differentiable and $\nabla f(x)=-1$, so $\partial f(x)=\{-1\}$
- For $x=0, f$ not differentiable, but since $f$ convex:

$$
\partial f(0)=\operatorname{cl}(\operatorname{conv} S(0))=\operatorname{cl}(\operatorname{conv}(\{-1,1\})=[-1,1]
$$

- The subdifferential operator:



## A nonconvex example

- Nonconvex function:

- Subdifferential
- For $x>b, f$ differentiable and $\nabla f(x)=1$, so $\partial f(x)=\{1\}$
- For $x<a, f$ differentiable and $\nabla f(x)=-1$, so $\partial f(x)=\{-1\}$
- For $x \in(a, b)$, no affine minorizer, $\partial f(x)=\emptyset$
- For $x=a, f$ not differentiable, $\partial f(x)=[-1,0]$
- For $x=b, f$ not differentiable, $\partial f(x)=[0,1]$
- The subdifferential operator:



## Example - Separable functions

- Consider the separable function $f(x)=\sum_{i=1}^{n} f_{i}\left(x_{i}\right)$
- Subdifferential

$$
\partial f(x)=\left\{s=\left(s_{1}, \ldots, s_{n}\right): s_{i} \in \partial f_{i}\left(x_{i}\right)\right\}
$$

- The subgradient $s \in \partial f(x)$ if and only if each $s_{i} \in \partial f_{i}\left(x_{i}\right)$
- Proof:
- Assume all $s_{i} \in \partial f\left(x_{i}\right)$ :

$$
f(y)-f(x)=\sum_{i=1}^{n} f_{i}\left(y_{i}\right)-f_{i}\left(x_{i}\right) \geq \sum_{i=1}^{n} s_{i}\left(y_{i}-x_{i}\right)=s^{T}(y-x)
$$

- Assume $s_{j} \notin \partial f\left(x_{j}\right)$ and $x_{i}=y_{i}$ for all $i \neq j$ :

$$
f_{j}\left(y_{j}\right)-f_{j}\left(x_{j}\right)<s_{j}\left(y_{j}-x_{j}\right)
$$

which gives

$$
f(y)-f(x)=f_{j}\left(y_{j}\right)-f_{j}\left(x_{j}\right)<s_{j}\left(y_{j}-x_{j}\right)=s^{T}(y-x)
$$

- Consider the 1-norm $f(x)=\|x\|_{1}=\sum_{i=1}^{n}\left|x_{i}\right|$
- It is a separable function of absolute values
- From previous examples, we conclude that the subdifferential is

$$
\partial f(x)=\left\{\left(s_{1}, \ldots, s_{n}\right):\left\{\begin{array}{ll}
s_{i}=-1 & \text { if } x_{i}<0 \\
s_{i} \in[-1,1] & \text { if } x_{i}=0 \\
s_{i}=1 & \text { if } x_{i}>0
\end{array}\right\}\right.
$$

- Consider the 2-norm $f(x)=\|x\|_{2}=\sqrt{\|x\|_{2}^{2}}$
- The function is differentiable everywhere except for when $x=0$
- Divide into two cases; $x=0$ and $x \neq 0$
- Subdifferential for $x \neq 0: \partial f(x)=\{\nabla f(x)\}$ :
- Let $h(u)=\sqrt{u}$ and $g(x)=\|x\|_{2}^{2}$, then $f(x)=(h \circ g)(x)$
- The gradient for all $x \neq 0$ by chain rule (since $h: \mathbb{R}_{+} \rightarrow \mathbb{R}$ ):

$$
\nabla f(x)=\nabla h(g(x)) \nabla g(x)=\frac{1}{2 \sqrt{\|x\|_{2}}} 2 x=\frac{x}{\|x\|_{2}}
$$

## Example cont'd - 2-norm

Subdifferential of $\|x\|_{2}$ at $x=0$
(i) educated guess of subdifferential from $\partial f(0)=\operatorname{cl}(\operatorname{conv} S(0))$

- recall $S(0)$ is set of all limit points of $\left(\nabla f\left(x_{k}\right)\right)_{k \in \mathbb{N}}$ when $x_{k} \rightarrow 0$
- let $x_{k}=t^{k} d$ with $t \in(0,1)$ and $d \in \mathbb{R}^{n} \backslash 0$, then $\nabla f\left(x_{k}\right)=\frac{d}{\|d\|_{2}}$
- since $d$ arbitrary, $\left(\nabla f\left(x_{k}\right)\right)$ can converge to any unit norm vector
- so $S(0)=\left\{s:\|s\|_{2}=1\right\}$ and $\partial f(0)=\left\{s:\|s\|_{2} \leq 1\right\}$ ?
(ii) verify using subgradient definition $f(y) \geq f(0)+s^{T}(y-0)=s^{T} y$
- Let $\|s\|_{2}>1$, then for, e.g., $y=2 s$

$$
s^{T} y=2\|s\|_{2}^{2}>2\|s\|_{2}=f(y)
$$

so such $s$ are not subgradients

- Let $\|s\|_{2} \leq 1$, then for all $y$ :

$$
s^{T} y \leq\|s\|_{2}\|y\|_{2} \leq\|y\|_{2}=f(y)
$$

so such $s$ are subgradients

## Outline

- Subdifferential and subgradient - Definition and basic properties
- Monotonicity
- Examples
- Strong monotonicity and cocoercivity
- Fermat's rule
- Subdifferential calculus
- Optimality conditions
- Proximal operators


## Strong convexity revisited

- Recall that $f$ is $\sigma$-strongly convex if $f-\frac{\sigma}{2}\|\cdot\|_{2}^{2}$ is convex
- If $f$ is $\sigma$-strongly convex then

$$
f(y) \geq f(x)+s^{T}(y-x)+\frac{\sigma}{2}\|x-y\|_{2}^{2}
$$

holds for all $x \in \operatorname{dom} \partial f, s \in \partial f(x)$, and $y \in \mathbb{R}^{n}$

- The function has convex quadratic minorizers instead of affine

- Multiple lower bounds at $x_{2}$ with subgradients $s_{2,1}$ and $s_{2,2}$


## Strong monotonicity

- If $f \sigma$-strongly convex function, then $\partial f$ is $\sigma$-strongly monotone:

$$
\left(s_{x}-s_{y}\right)^{T}(x-y) \geq \sigma\|x-y\|_{2}^{2}
$$

for all $s_{x} \in \partial f(x)$ and $s_{y} \in \partial f(y)$

- Proof: Add two copies of strong convexity inequality

$$
f(y) \geq f(x)+s_{x}^{T}(y-x)+\frac{\sigma}{2}\|x-y\|_{2}^{2}
$$

with $x$ and $y$ swapped

- $\partial f$ is $\sigma$-strongly monotone if and only if $\partial f-\sigma I$ is monotone
- $\partial f: \mathbb{R} \rightarrow 2^{\mathbb{R}}$ : Minimum slope $\sigma$ and maximum slope $\infty$


Strongly convex functions - An equivalence

The following are equivalent for $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$
(i) $f$ is closed and $\sigma$-strongly convex
(ii) $\partial f$ is maximally monotone and $\sigma$-strongly monotone

## Proof:

(i) $\Rightarrow$ (ii): we know this from before
(ii) $\Rightarrow$ (i): (ii) $\Rightarrow \partial f-\sigma I=\partial\left(f-\frac{\sigma}{2}\|\cdot\|_{2}^{2}\right)$ maximally monotone $\Rightarrow f-\frac{\sigma}{2}\|\cdot\|_{2}^{2}$ closed convex
$\Rightarrow f$ closed and $\sigma$-strongly convex

## Smooth convex functions

- A differentiable function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is convex and $\beta$-smooth if

$$
\begin{aligned}
& f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|x-y\|_{2}^{2} \\
& f(y) \geq f(x)+\nabla f(x)^{T}(y-x)
\end{aligned}
$$

hold for all $x, y \in \mathbb{R}^{n}$

- $f$ has convex quadratic majorizers and affine minorizers

- Quadratic upper bound is called descent lemma


## Cocoercivity of gradient

- Gradient of smooth convex function is monotone and Lipschitz

$$
(\nabla f(x)-\nabla f(y))^{T}(x-y) \geq 0
$$

$$
\|\nabla f(y)-\nabla f(x)\|_{2} \leq \beta\|x-y\|_{2}
$$

- $\nabla f: \mathbb{R} \rightarrow \mathbb{R}$ : Minimum slope 0 and maximum slope $\beta$

- Actually satisfies the stronger $\frac{1}{\beta}$-cocoercivity property

$$
(\nabla f(x)-\nabla f(y))^{T}(x-y) \geq \frac{1}{\beta}\|\nabla f(y)-\nabla f(x)\|_{2}^{2}
$$

due to the Baillon-Haddad theorem

Smooth convex functions - An equivalence

Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be differentiable. The following are equivalent:
(i) $\nabla f$ is $\frac{1}{\beta}$-cocoercive
(ii) $\nabla f$ is maximally monotone and $\beta$-Lipschitz continuous
(iii) $f$ is closed convex and satisfies descent lemma (is $\beta$-smooth)

Will later connect smooth convexity and strong convexity via conjugates

## Smooth strongly convex functions

- Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be differentiable
- $f$ is $\beta$-smooth and $\sigma$-strongly convex with $0<\sigma \leq \beta$ if

$$
\begin{aligned}
& f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|x-y\|_{2}^{2} \\
& f(y) \geq f(x)+\nabla f(x)^{T}(y-x)+\frac{\sigma}{2}\|x-y\|_{2}^{2}
\end{aligned}
$$

hold for all $x, y \in \mathbb{R}^{n}$

- $f$ has quadratic minorizers and quadratic majorizers

- We say that the ratio $\frac{\beta}{\sigma}$ is the condition number for the function


## Gradient of smooth strongly convex function

- Gradient of $\beta$-smooth $\sigma$-strongly convex function $f$ satisfies

$$
\begin{aligned}
\|\nabla f(y)-\nabla f(x)\|_{2} & \leq \beta\|x-y\|_{2} \\
(\nabla f(x)-\nabla f(y))^{T}(x-y) & \geq \sigma\|x-y\|_{2}^{2}
\end{aligned}
$$

so is $\beta$-Lipschitz continuous and $\sigma$-strongly monotone

- $\nabla f: \mathbb{R} \rightarrow \mathbb{R}$ : Minimum slope $\sigma$ and maximum slope $\beta$

- Actually satisfies this stronger property
$(\nabla f(x)-\nabla f(y))^{T}(x-y) \geq \frac{1}{\beta+\sigma}\|\nabla f(y)-\nabla f(x)\|_{2}^{2}+\frac{\sigma \beta}{\beta+\sigma}\|x-y\|_{2}^{2}$ for all $x, y \in \mathbb{R}^{n}$


## Proof of stronger property

- $f$ is $\sigma$-strongly convex if and only if $g:=f-\frac{\sigma}{2}\|\cdot\|_{2}^{2}$ is convex
- Since $f$ is $\beta$-smooth $g$ is $(\beta-\sigma)$-smooth
- Since $g$ convex and $(\beta-\sigma)$-smooth, $\nabla g$ is $\frac{1}{\beta-\sigma}$-cocoercive:

$$
(\nabla g(x)-\nabla g(y))^{T}(x-y) \geq \frac{1}{\beta-\sigma}\|\nabla g(x)-\nabla g(y)\|_{2}^{2}
$$

which by using $\nabla g=\nabla f-\sigma I$ gives
$(\nabla f(x)-\nabla f(y))^{T}(x-y)-\sigma\|x-y\|_{2}^{2} \geq \frac{1}{\beta-\sigma}\|\nabla f(x)-\nabla f(y)-\sigma(x-y)\|_{2}^{2}$
which by expanding the square and rearranging is equivalent to
$(\nabla f(x)-\nabla f(y))^{T}(x-y) \geq \frac{1}{\beta+\sigma}\|\nabla f(x)-\nabla f(y)\|_{2}^{2}+\frac{\sigma \beta}{\beta+\sigma}\|x-y\|_{2}^{2}$

## Outline

- Subdifferential and subgradient - Definition and basic properties
- Monotonicity
- Examples
- Strong monotonicity and cocoercivity
- Fermat's rule
- Subdifferential calculus
- Optimality conditions
- Proximal operators


## Fermat's rule

Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$, then $x$ minimizes $f$ if and only if $0 \in \partial f(x)$

- Proof: $x$ minimizes $f$ if and only if

$$
f(y) \geq f(x)=f(x)+0^{T}(y-x) \quad \text { for all } y \in \mathbb{R}^{n}
$$

which by definition of subdifferential is equivalent to $0 \in \partial f(x)$

- Example: several subgradients at solution, including 0



## Fermat's rule - Nonconvex example

- Fermat's rule holds also for nonconvex functions
- Example:

- $\partial f\left(x_{1}\right)=0$ and $\nabla f\left(x_{1}\right)=0$ (global minimum)
- $\partial f\left(x_{2}\right)=\emptyset$ and $\nabla f\left(x_{2}\right)=0$ (local minimum)
- For nonconvex $f$, we can typically only hope to find local minima




## Prox is (firmly) nonexpansive

- We know 1-cocoercivity implies nonexpansiveness (1-Lipschitz)

$$
\left\|\operatorname{prox}_{\gamma g}(x)-\operatorname{prox}_{\gamma g}(y)\right\|_{2} \leq\|x-y\|_{2}
$$

which was shown using Cauchy-Schwarz inequality

- Actually the stronger firm nonexpansive inequality holds
$\left\|\operatorname{prox}_{\gamma g}(x)-\operatorname{prox}_{\gamma g}(y)\right\|_{2}^{2} \leq\|x-y\|_{2}^{2}$

$$
-\left\|x-\operatorname{prox}_{\gamma g}(x)-\left(y-\operatorname{prox}_{\gamma g}(y)\right)\right\|_{2}^{2}
$$

which implies nonexpansiveness

- Proof:
- take 1-cocoercivity and multiply both sides by 2

$$
2(x-y)^{T}\left(\operatorname{prox}_{\gamma g}(x)-\operatorname{prox}_{\gamma f}(y)\right) \geq 2\left\|\operatorname{prox}_{\gamma g}(x)-\operatorname{prox}_{\gamma f}(y)\right\|_{2}^{2}
$$

- use the following equality with $u=\operatorname{prox}_{\gamma g}(x)$ and $v=\operatorname{prox}_{\gamma g}(y)$ : $(x-y)^{T}(u-v)=\frac{1}{2}\left(\|x-y\|_{2}^{2}+\|u-v\|_{2}^{2}-\|x-y-(u-v)\|_{2}^{2}\right)$


## Proximal operator - Separable functions

- Let $x=\left(x_{1}, \ldots, x_{n}\right)$ and $g(x)=\sum_{i=1}^{n} g_{i}\left(x_{i}\right)$ be separable, then

$$
\operatorname{prox}_{\gamma g}(z)=\left(\operatorname{prox}_{\gamma g_{1}}\left(z_{1}\right), \ldots, \operatorname{prox}_{\gamma g_{n}}\left(z_{n}\right)\right)
$$

decomposes into $n$ individual proxes

- Why? Since also $\|\cdot\|_{2}^{2}$ is separable:

$$
\begin{aligned}
\operatorname{prox}_{\gamma g}(z) & =\underset{x}{\operatorname{argmin}}\left(g(x)+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right) \\
& =\underset{x}{\operatorname{argmin}}\left(\sum_{i=1}^{n}\left(g_{i}\left(x_{i}\right)+\frac{1}{2 \gamma}\left(x_{i}-z_{i}\right)^{2}\right)\right)
\end{aligned}
$$

which gives $n$ independent optimization problems

$$
\underset{x_{i}}{\operatorname{argmin}}\left(g_{i}\left(x_{i}\right)+\frac{1}{2 \gamma}\left(x_{i}-z_{i}\right)^{2}\right)=\operatorname{prox}_{\gamma g_{i}}\left(z_{i}\right)
$$

## Proximal operator - Example 1

- Consider the function $g$ with subdifferential $\partial g$ :

$$
g(x)=\left\{\begin{array}{ll}
-x & \text { if } x \leq 0 \\
0 & \text { if } x \geq 0
\end{array} \quad \partial g(x)= \begin{cases}-1 & \text { if } x<0 \\
{[-1,0]} & \text { if } x=0 \\
0 & \text { if } x>0\end{cases}\right.
$$

- Graphical representations


- Fermat's rule for $x=\operatorname{prox}_{\gamma g}(z)$ :

$$
0 \in \partial g(x)+\gamma^{-1}(x-z)
$$

## Proximal operator - Example 1 cont'd

- Let $x<0$, then Fermat's rule reads

$$
0=-1+\gamma^{-1}(x-z) \quad \Leftrightarrow \quad x=z+\gamma
$$

which is valid $(x<0)$ if $z<-\gamma$

- Let $x=0$, then Fermat's rule reads

$$
0 \in[-1,0]+\gamma^{-1}(0-z)
$$

which is valid $(x=0)$ if $z \in[-\gamma, 0]$

- Let $x>0$, then Fermat's rule reads

$$
0=0+\gamma^{-1}(x-z) \quad \Leftrightarrow \quad x=z
$$

which is valid $(x>0)$ if $z>0$

- The prox satisfies

$$
\operatorname{prox}_{\gamma g}(z)= \begin{cases}z+\gamma & \text { if } z<-\gamma \\ 0 & \text { if } z \in[-\gamma, 0] \\ z & \text { if } z>0\end{cases}
$$

## Proximal operator - Example 2

Let $g(x)=\frac{1}{2} x^{T} P x+q^{T} x$ with $P$ positive semidefinite

- Gradient satisfies $\nabla g(x)=P x+q$
- Fermat's rule for $x=\operatorname{prox}_{\gamma g}(z)$ :

$$
\begin{aligned}
0=\nabla g(x)+\gamma^{-1}(x-z) & \Leftrightarrow \quad 0=P x+q+\gamma^{-1}(x-z) \\
& \Leftrightarrow(I+\gamma P) x=z-\gamma q \\
& \Leftrightarrow x=(I+\gamma P)^{-1}(z-\gamma q)
\end{aligned}
$$

- So $\operatorname{prox}_{\gamma g}(z)=(I+\gamma P)^{-1}(z-\gamma q)$


## Computational cost

- Evaluating prox requires solving optimization problem

$$
\operatorname{prox}_{\gamma g}(z)=\underset{x}{\operatorname{argmin}}\left(g(x)+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right)
$$

- Prox often more expensive to evaluate than gradient
- Example: Quadratic $g(x)=\frac{1}{2} x^{T} P x+q^{T} x$ :

$$
\operatorname{prox}_{\gamma g}(z)=(I+\gamma P)^{-1}(z-\gamma q), \quad \nabla g(z)=P z+q
$$

- But typically cheap to evaluate for separable functions
- Prox often used for nondifferentiable and separable functions

| 两 |  |  |
| :--- | :--- | :--- |


|  | Outline |
| :---: | :---: |
| Conjugate Functions, Optimality Conditions, and Duality <br> Pontus Giselsson | - Conjugate function - Definition and basic properties <br> - Examples <br> - Biconjugate <br> - Fenchel-Young's inequality <br> - Duality correspondence <br> - Moreau decomposition <br> - Duality and optimality conditions <br> - Weak and strong duality |
| 1 2 |  |
| Conjugate Functions | Conjugate function - Definition <br> - The conjugate function of $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ is defined as $f^{*}(s):=\sup _{x}\left(s^{T} x-f(x)\right)$ <br> - Implicit definition via optimization problem |
| 3 | 4 |
| Conjugate function properties <br> - Let $a_{x}(s):=s^{T} x-f(x)$ be affine function parameterized by $x$ : $f^{*}(s)=\sup _{x} a_{x}(s)$ <br> is supremum of family of affine functions <br> - Epigraph of $f^{*}$ is intersection of epigraphs of (below three) $a_{x}$ <br> - $f^{*}$ convex: epigraph intersection of convex halfspaces epi $a_{x}$ <br> - $f^{*}$ closed: epigraph intersection of closed halfspaces epi $a_{x}$ | Conjugate interpretation <br> - Conjugate $f^{*}(s)$ defines affine minorizer to $f$ with slope $s$ : <br> where $-f^{*}(s)$ decides constant offset to get support <br> - Why? $\begin{aligned} f^{*}(s)=\sup _{x}\left(s^{T} x-f(x)\right) \quad & \Leftrightarrow \quad f^{*}(s) \geq s^{T} x-f(x) \text { for all } x \\ & \Leftrightarrow \quad f(x) \geq s^{T} x-f^{*}(s) \text { for all } x \end{aligned}$ <br> - Maximizing argument $x^{*}$ gives support: $f\left(x^{*}\right)=s^{T} x^{*}-f^{*}(s)$ <br> - We have $f\left(x^{*}\right)=s^{T} x^{*}-f^{*}(s)$ if and only if $s \in \partial f\left(x^{*}\right)$ |
| Consequence | Outline |
| - Conjugate of $f$ and $\operatorname{env} f$ are the same, i.e., $f^{*}=(\operatorname{env} f)^{*}$ <br> - Functions have same supporting affine functions <br> - Epigraphs have same supporting hyperplanes | - Conjugate function - Definition and basic properties <br> - Examples <br> - Biconjugate <br> - Fenchel-Young's inequality <br> - Duality correspondence <br> - Moreau decomposition <br> - Duality and optimality conditions <br> - Weak and strong duality |
|  | 8 |




- Consider

Let $g(x)=\frac{1}{2} x^{T} Q x+p^{T} x$ with $Q$ positive definite (invertible)

- Gradient satisfies $\nabla g(x)=Q x+p$
- Fermat's rule for $g^{*}(s)=\sup _{x}\left(s^{T} x-\frac{1}{2} x^{T} Q x-p^{T} x\right)$ :

$$
0=s-Q x-p \quad \Leftrightarrow \quad x=Q^{-1}(s-p)
$$

- So

$$
\begin{aligned}
g^{*}(s) & =s^{T} Q^{-1}(s-p)-\frac{1}{2}(s-p)^{T} Q^{-1} Q Q^{-1}(s-p)+p^{T} Q^{-1}(s-p) \\
& =\frac{1}{2}(s-p)^{T} Q^{-1}(s-p)
\end{aligned}
$$

$$
g(x)= \begin{cases}-x-1 & \text { if } x \leq-1 \\ 0 & \text { if } x \in[-1,1] \\ x-1 & \text { if } x \geq 1\end{cases}
$$



- Subdifferential satisfies

$$
\partial g(x)= \begin{cases}-1 & \text { if } x<-1 \\ {[-1,0]} & \text { if } x=-1 \\ 0 & \text { if } x \in(-1,1) \\ {[0,1]} & \text { if } x=1 \\ 1 & \text { if } x>1\end{cases}
$$



## Example cont'd

- We use $g^{*}(s)=s x-g(x)$ if $s \in \partial g(x)$ :
- $x<-1: s=-1$, hence $g^{*}(-1)=-1 x-(-x-1)=1$
- $x=-1: s \in[-1,0]$ hence $g^{*}(s)=-s-0=-s$
- $x \in(-1,1): s=0$ hence $g^{*}(0)=0 x-0=0$
- $x=1: s \in[0,1]$ hence $g^{*}(s)=s-0=s$
- $x>1$ : $s=1$ hence $g^{*}(1)=x-(x-1)=1$
- That is

$$
g^{*}(s)= \begin{cases}-s & \text { if } s \in[-1,0] \\ s & \text { if } s \in[0,1]\end{cases}
$$

- For $s<-1$ and $s>1, g^{*}(s)=\infty$ :
- $s<-1$ : let $x=t \rightarrow-\infty$ and $g^{*}(s) \geq((s+1) t+1) \rightarrow \infty$
- $s>1$ : let $x=t \rightarrow \infty$ and $g^{*}(s) \geq((s-1) t+1) \rightarrow \infty$


## Example - Separable functions

- Let $f(x)=\sum_{i=1}^{n} f_{i}\left(x_{i}\right)$ be a separable function, then

$$
f^{*}(s)=\sum_{i=1}^{n} f_{i}^{*}\left(s_{i}\right)
$$

is also separable

- Proof:

$$
\begin{aligned}
f^{*}(s) & =\sup _{x}\left(s^{T} x-\sum_{i=1}^{n} f_{i}\left(x_{i}\right)\right) \\
& =\sup _{x}\left(\sum_{i=1}^{n}\left(s_{i} x_{i}-f_{i}\left(x_{i}\right)\right)\right) \\
& =\sum_{i=1}^{n} \sup _{x_{i}}\left(s_{i} x_{i}-f_{i}\left(x_{i}\right)\right) \\
& =\sum_{i=1}^{n} f_{i}^{*}\left(s_{i}\right)
\end{aligned}
$$

## Example - 1-norm

- Let $f(x)=\|x\|_{1}=\sum_{i=1}^{n}\left|x_{i}\right|$ be the 1-norm
- It is a separable sum of absolute values
- Use separable sum formula and that $|\cdot|^{*}=\iota_{[-1,1]}$ :

$$
f^{*}(s)=\sum_{i=1}^{n} f_{i}^{*}\left(s_{i}\right)=\sum_{i=1}^{n} \iota_{[-1,1]}\left(s_{i}\right)= \begin{cases}0 & \text { if } \max _{i}\left(\left|s_{i}\right|\right) \leq 1 \\ \infty & \text { else }\end{cases}
$$

- We have $\max _{i}\left(\left|s_{i}\right|\right)=\|s\|_{\infty}$, let

$$
B_{\infty}(r)=\left\{s:\|s\|_{\infty} \leq r\right\}
$$

be the infinity norm ball of radius $r$, then

$$
f^{*}(s)=\iota_{B_{\infty}(1)}(s)
$$

is the indicator function for the unit infinity norm ball

## Outline

- Conjugate function - Definition and basic properties
- Examples
- Biconjugate
- Fenchel-Young's inequality
- Duality correspondence
- Moreau decomposition
- Duality and optimality conditions
- Weak and strong duality


## Biconjugate

- Biconjuate $f^{* *}:=\left(f^{*}\right)^{*}$ is conjugate of conjugate

$$
f^{* *}(x)=\sup _{s}\left(x^{T} s-f^{*}(s)\right)
$$

- For every $x$, it is largest value of all affine minorizers

- Why?:
- $x^{T} s-f^{*}(s)$ : supporting affine minorizer to $f$ with slope $s$
- $f^{* *}(x)$ picks largest over all these affine minorizers evaluated at $x$

Biconjugate and convex envelope

- Biconjugate is closed convex envelope of $f$

- $f^{* *} \leq f$ and $f^{* *}=f$ if and only if $f$ (closed and) convex


Relation between $\partial f$ and $\partial f^{*}$ - General case
Inverse relation between $\partial f$ and $\partial f^{*}$ - Convex case

Suppose $f$ closed convex, then $s \in \partial f(x) \Longleftrightarrow x \in \partial f^{*}(s)$

- Using implication on previous slide twice and $f^{* *}=f$ :

$$
s \in \partial f(x) \Rightarrow x \in \partial f^{*}(s) \Rightarrow s \in \partial f^{* *}(x) \Rightarrow s \in \partial f(x)
$$

- Another way to write the result is that for closed convex $f$ :

$$
\partial f^{*}=(\partial f)^{-1}
$$

(Definition of inverse of set-valued $A: x \in A^{-1} u \Longleftrightarrow u \in A x$ )

Example 1 - Relation between $\partial f$ and $\partial f^{*}$

- What is $\partial f^{*}$ for below $\partial f$ ?


Example 1 - Relation between $\partial f$ and $\partial f^{*}$

- What is $\partial f^{*}$ for below $\partial f$ ?

- Since $\partial f^{*}=(\partial f)^{-1}$, we flip the figure

Example 2 - Relation between $\partial f$ and $\partial f^{*}$


- region with slope $\sigma$ in $\partial f(x) \Leftrightarrow$ region with slope $\frac{1}{\sigma}$ in $\partial f^{*}(s)$
- Implication: $\partial f \sigma$-strong monotone $\Leftrightarrow \partial f^{*}(s) \sigma$-cocoercive? (Recall: $\sigma$-cocoercivity $\Leftrightarrow \frac{1}{\sigma}$-Lipschitz and monotone)


## Outline

- Conjugate function - Definition and basic properties
- Examples
- Biconjugate
- Fenchel-Young's inequality
- Duality correspondence
- Moreau decomposition
- Duality and optimality conditions
- Weak and strong duality


## Cocoercivity and strong monotonicity

$$
\begin{gathered}
\partial f: \mathbb{R}^{n} \rightarrow 2^{\mathbb{R}^{n}} \text { maximal monotone and } \sigma \text {-strongly monotone } \\
\partial f^{*}=\nabla f^{*}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n} \text { single-valued and } \sigma \text {-cocoercive }
\end{gathered}
$$

- $\sigma$-strong monotonicity: for all $u \in \partial f(x)$ and $v \in \partial f(y)$

$$
\begin{equation*}
(u-v)^{T}(x-y) \geq \sigma\|x-y\|_{2}^{2} \tag{1}
\end{equation*}
$$

or equivalently for all $x \in \partial f^{*}(u)$ and $y \in \partial f^{*}(v)$

- $\partial f^{*}$ is single-valued:
- Assume $x \in \partial f^{*}(u)$ and $y \in \partial f^{*}(u)$, then Ihs of (1) 0 and $x=y$
- $\nabla f^{*}$ is $\sigma$-cocoercive: plug $x=\nabla f^{*}(u)$ and $y=\nabla f^{*}(v)$ into (1)
- That $\partial f^{*}$ has full domain follows from Minty's theorem


## Duality correspondance

Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$. Then the following are equivalent:
(i) $f$ is closed and $\sigma$-strongly convex
(ii) $\partial f$ is maximally monotone and $\sigma$-strongly monotone
(iii) $\nabla f^{*}$ is $\sigma$-cocoercive
(iv) $\nabla f^{*}$ is maximally monotone and $\frac{1}{\sigma}$-Lipschitz continuous
(v) $f^{*}$ is closed convex and satisfies descent lemma (is $\frac{1}{\sigma}$-smooth)
where $\nabla f^{*}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ and $f^{*}: \mathbb{R}^{n} \rightarrow \mathbb{R}$
Comments:

- (i) $\Leftrightarrow$ (ii) and (iii) $\Leftrightarrow$ (iv) $\Leftrightarrow$ (v): Previous lecture
- (ii) $\Leftrightarrow$ (iii): This lecture
- Since $f=f^{* *}$ the result holds with $f$ and $f^{*}$ interchanged
- Full proof available on course webpage
- Prox definition $\operatorname{prox}_{\gamma g}(z)=\operatorname{argmin}_{x}\left(g(x)+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right)$
- Let $r=\gamma g+\frac{1}{2}\|\cdot\|_{2}^{2}$, then

$$
\begin{aligned}
\operatorname{prox}_{\gamma g}(z) & =\underset{x}{\operatorname{argmin}}\left(g(x)+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right) \\
& =\underset{x}{\operatorname{argmax}}\left(-\gamma g(x)-\frac{1}{2}\|x-z\|_{2}^{2}\right) \\
& =\underset{x}{\operatorname{argmax}}\left(z^{T} x-\left(\frac{1}{2}\|x\|_{2}^{2}+\gamma g(x)\right)\right) \\
& =\underset{x}{\operatorname{argmax}}\left(z^{T} x-r(x)\right) \\
& =\nabla r^{*}(z)
\end{aligned}
$$

where last step is subdifferential formula for $r^{*}$ for convex $r$

- Now, $r$ is 1-strongly convex and $\nabla r^{*}=\operatorname{prox}_{\gamma g}$ is 1-cocoercive
- Conjugate function - Definition and basic properties
- Examples
- Biconjugate
- Fenchel-Young's inequality
- Duality correspondence
- Moreau decomposition
- Duality and optimality conditions
- Weak and strong duality

Assume $g$ closed convex, then $\operatorname{prox}_{g}(z)+\operatorname{prox}_{g^{*}}(z)=z$

- When $g$ scaled by $\gamma>0$, Moreau decomposition is
$z=\operatorname{prox}_{\gamma g}(z)+\operatorname{prox}_{(\gamma g)^{*}}(z)=\operatorname{prox}_{\gamma g}(z)+\gamma \operatorname{prox}_{\gamma^{-1} g^{*}}\left(\gamma^{-1} z\right)$
(since $\left.\operatorname{prox}_{(\gamma g)^{*}}=\gamma \operatorname{prox}_{\gamma^{-1} g^{*}} \circ \gamma^{-1} \mathrm{Id}\right)$
- Don't need to know $g^{*}$ to compute $\operatorname{prox}_{\gamma g^{*}}$

Moreau decomposition - Proof

- Let $u=z-x$
- Fermat's rule: $x=\operatorname{prox}_{g}(z)$ if and only if

$$
\begin{array}{rll}
0 \in \partial g(x)+x-z & \Leftrightarrow & z-x \in \partial g(x) \\
& \Leftrightarrow & u \in \partial g(x) \\
& \Leftrightarrow & x \in \partial g^{*}(u) \\
& \Leftrightarrow & z-u \in \partial g^{*}(u) \\
& \Leftrightarrow & 0 \in \partial g^{*}(u)+u-z
\end{array}
$$

if and only if $u=\operatorname{prox}_{g^{*}}(z)$ by Fermat's rule

- Using $z=x+u$, we get

$$
z=x+u=\operatorname{prox}_{g}(z)+\operatorname{prox}_{g^{*}}(z)
$$

## Optimality Conditions and Duality

## Outline

- Conjugate function - Definition and basic properties
- Examples
- Biconjugate
- Fenchel-Young's inequality
- Duality correspondence
- Moreau decomposition
- Duality and optimality conditions
- Weak and strong duality

Composite optimization problem

- Consider primal composite optimization problem

$$
\operatorname{minimize} f(L x)+g(x)
$$

where $f, g$ closed convex and $L$ is a matrix

- We will derive primal-dual optimality conditions and dual problem
- Introduce dual variable $\mu \in \partial f(L x)$, then optimality condition

$$
\text { and assume } C Q \text {, then: }
$$

minimize $f(L x)+g(x)$
is solved by $x^{\star} \in \mathbb{R}^{n}$ if and only if $x^{\star}$ satisfies

$$
0 \in L^{T} \partial f\left(L x^{\star}\right)+\partial g\left(x^{\star}\right)
$$

- Optimality condition implies that vector $s$ exists such that

$$
s \in L^{T} \partial f\left(L x^{\star}\right) \quad \text { and } \quad-s \in \partial g\left(x^{\star}\right)
$$

- So CQ implies a subgradient exists for both functions at solution


## Primal-dual optimality condition 2

- Primal-dual optimality condition

$$
\begin{aligned}
\mu & \in \partial f(L x) \\
-L^{T} \mu & \in \partial g(x)
\end{aligned}
$$

- Using subdifferential inverse:

$$
\mu \in \partial f(L x) \quad \Longleftrightarrow \quad L x \in \partial f^{*}(\mu)
$$

gives equivalent primal dual optimality condition

$$
\begin{aligned}
& L x \in \partial f^{*}(\mu) \\
&-L^{T} \mu \in \partial g(x)
\end{aligned}
$$

## Dual optimality condition

- Using subdifferential inverse on other condition

$$
-L^{T} \mu \in \partial g(x) \quad \Longleftrightarrow \quad x \in \partial g^{*}\left(-L^{T} \mu\right)
$$

gives equivalent primal dual optimality condition

$$
\begin{aligned}
& L x \in \partial f^{*}(\mu) \\
& x
\end{aligned}=\partial g^{*}\left(-L^{T} \mu\right) .
$$

- This is equivalent to that:

$$
0 \in \partial f^{*}(\mu)-L \underbrace{\partial g^{*}\left(-L^{T} \mu\right)}_{x}
$$

which is a dual optimality condition since it involves only $\mu$

## Dual problem

- The dual optimality condition

$$
0 \in \partial f^{*}(\mu)-L \partial g^{*}\left(-L^{T} \mu\right)
$$

is a sufficient condition for solving the dual problem

$$
\operatorname{minimize} f^{*}(\mu)+g^{*}\left(-L^{T} \mu\right)
$$

- Have also necessity under $C Q$ on dual, which is mild


## Why dual problem?

- Sometimes easier to solve than primal
- Only useful if primal solution can be obtained from dual


## Solving primal from dual

- Assume $f, g$ closed convex and CQ holds
- Assume optimal dual $\mu$ known: $0 \in \partial f^{*}(\mu)-L \partial g^{*}\left(-L^{T} \mu\right)$
- Optimal primal $x$ must satisfy any and all primal-dual conditions:

$$
\begin{array}{ll}
\left\{\begin{array}{l}
\mu \in \partial f(L x) \\
-L^{T} \mu \in \partial g(x)
\end{array}\right. & \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
-L^{*} \mu \in \partial g(x)
\end{array}\right. \\
\left\{\begin{array}{l}
\mu \in \partial f(L x) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right. & \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right.
\end{array}
$$

- If one of these uniquely characterizes $x$, then must be solution
- $g^{*}$ is differentiable at $-L^{T} \mu$ for dual solution $\mu$
- $f^{*}$ is differentiable at dual solution $\mu$ and $L$ invertible
- ...


## Optimality conditions - Summary

- Assume $f, g$ closed convex and that CQ holds
- Problem $\min _{x} f(L x)+g(x)$ is solved by $x$ if and only if

$$
0 \in L^{T} \partial f(L x)+\partial g(x)
$$

- Primal dual necessary and sufficient optimality conditions:

$$
\begin{array}{ll}
\left\{\begin{array}{l}
\mu \in \partial f(L x) \\
-L^{T} \mu \in \partial g(x)
\end{array}\right. & \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
-L^{T} \mu \in \partial g(x)
\end{array}\right. \\
\left\{\begin{array}{l}
\mu \in \partial f(L x) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right. & \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right.
\end{array}
$$

- Dual optimality condition

$$
0 \in \partial f^{*}(\mu)-L \partial g^{*}\left(-L^{T} \mu\right)
$$

solves dual problem $\min _{\mu} f^{*}(\mu)+g^{*}\left(-L^{T} \mu\right)$


| 两 |  |  |
| :--- | :--- | :--- |



## Proximal gradient - Example

- Proximal gradient iterations for problem minimize $\frac{1}{2}(x-a)^{2}+|x|$
- $f(x)=\frac{1}{2}(x-a)^{2}$ is smooth term and $g(x)=|x|$ is nonsmooth
- Iteration: $x_{k+1}=\operatorname{prox}_{\gamma g}\left(x_{k}-\gamma \nabla f\left(x_{k}\right)\right)$
- Note: convergence in finite number of iterations (not always)



## Proximal gradient - Example

- Proximal gradient iterations for problem minimize $\frac{1}{2}(x-a)^{2}+|x|$
- $f(x)=\frac{1}{2}(x-a)^{2}$ is smooth term and $g(x)=|x|$ is nonsmooth
- Iteration: $x_{k+1}=\operatorname{prox}_{\gamma g}\left(x_{k}-\gamma \nabla f\left(x_{k}\right)\right)$
- Note: convergence in finite number of iterations (not always)



## Proximal gradient - Example

- Proximal gradient iterations for problem minimize $\frac{1}{2}(x-a)^{2}+|x|$
- $f(x)=\frac{1}{2}(x-a)^{2}$ is smooth term and $g(x)=|x|$ is nonsmooth
- Iteration: $x_{k+1}=\operatorname{prox}_{\gamma g}\left(x_{k}-\gamma \nabla f\left(x_{k}\right)\right)$
- Note: convergence in finite number of iterations (not always)


## Proximal gradient - Special cases

- Proximal gradient method:
- solves minimize $(f(x)+g(x))$
- iteration: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)$
- Proximal gradient method with $g=0$ :
- solves $\underset{x}{\operatorname{minimize}}(f(x))$
- $\operatorname{prox}_{\gamma_{k} g}(z)^{x}=\operatorname{argmin}_{x}\left(0+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right)=z$
- iteration: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)=x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)$
- reduces to gradient method
- Proximal gradient method with $f=0$ :
- solves minimize $(g(x))$
- $\nabla f(x)=0^{x}$
- iteration: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}\right)$
- reduces to proximal point method (which is not very useful)



## Outline

- Introducing proximal gradient method and examples
- Solving composite problem - Fixed-points and convergence
- Application to primal and dual problems


## Proximal gradient method - Fixed-point set

- Proximal gradient step

$$
x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)
$$

- If $x_{k+1}=x_{k}$, they are in proximal gradient fixed-point set

$$
\left\{x: x=\operatorname{prox}_{\gamma g}(x-\gamma \nabla f(x))\right\}
$$

- Under some assumptions, algorithm will satisfy $x_{k+1}-x_{k} \rightarrow 0$
- this means that fixed-point equation will be satisfied in limit
- what does it mean for $x$ to be a fixed-point?


## Proximal gradient - Optimality condition

- Proximal gradient step:

$$
v=\operatorname{prox}_{\gamma g}(x-\gamma \nabla f(x))=\underset{y}{\operatorname{argmin}}(g(y)+\underbrace{\frac{1}{2 \gamma}\|y-(x-\gamma \nabla f(x))\|_{2}^{2}}_{h(y)})
$$

where $v$ is unique due to strong convexity of $h$

- Fermat's rule (since CQ holds) gives $v=\operatorname{prox}_{\gamma g}(x-\gamma \nabla f(x))$ iff:

$$
\begin{aligned}
0 & \in \partial g(v)+\partial h(v) \\
& =\partial g(v)+\gamma^{-1}(v-(x-\gamma \nabla f(x))) \\
& =\partial g(v)+\nabla f(x)+\gamma^{-1}(v-x)
\end{aligned}
$$

since $h$ differentiable

For $\gamma>0$, we have that
$\bar{x}=\operatorname{prox}_{\gamma g}(\bar{x}-\gamma \nabla f(\bar{x})) \quad$ if and only if $\quad 0 \in \partial g(\bar{x})+\nabla f(\bar{x})$

- Proof: the proximal step equivalence
$v=\operatorname{prox}_{\gamma g}(x-\gamma \nabla f(x)) \quad \Leftrightarrow \quad 0 \in \partial g(v)+\nabla f(x)+\gamma^{-1}(v-x)$
evaluated at a fixed-point $x=v=\bar{x}$ reads

$$
\bar{x}=\operatorname{prox}_{\gamma g}(\bar{x}-\gamma \nabla f(\bar{x})) \quad \Leftrightarrow \quad 0 \in \partial g(\bar{x})+\nabla f(\bar{x})
$$

- We call inclusion $0 \in \partial g(\bar{x})+\nabla f(\bar{x})$ fixed-point characterization
- What does fixed-point characterization $0 \in \partial g(\bar{x})+\nabla f(\bar{x})$ mean?
- For convex differentiable $f$, subdifferential $\partial f(x)=\{\nabla f(x)\}$ and

$$
0 \in \partial f(\bar{x})+\partial g(\bar{x})=\partial(f+g)(\bar{x})
$$

(subdifferential sum rule holds), i.e., fixed-points solve problem

- For nonconvex differentiable $f$, we might have $\partial f(\bar{x})=\emptyset$
- Fixed-point are not in general global solutions
- Points $\bar{x}$ that satisfy $0 \in \partial g(\bar{x})+\nabla f(\bar{x})$ are called critical points
- If $g=0$, the condition is $\nabla f(\bar{x})=0$, i.e., a stationary point
- Quality of fixed-points differs between convex and nonconvex $f$


## Conditions on $\gamma_{k}$ for convergence

- We replace in proximal gradient method $f(y)$ by

$$
f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(y-x_{k}\right)+\frac{1}{2 \gamma_{k}}\left\|y-x_{k}\right\|_{2}^{2}
$$

and minimize this plus $g(y)$ over $y$ to get the next iterate

- We know from $\beta$-smoothness of $f$ that for all $x, y$

$$
f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|y-x\|_{2}^{2}
$$

- If $\gamma_{k} \in\left[\epsilon, \frac{1}{\beta}\right]$ with $\epsilon>0$, an upper bound is minimized
- Can use $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta}-\epsilon\right]$ and show convergence of some quantity


## Practical convergence - Example

- Logarithmic $y$ axis of quantity that should go to 0 for convergence
- Linear $x$ axis with iteration number

- Fast convergence to medium accuracy, slow from medium to high
- Many iterations may be required


## Stopping conditions

- For $\beta$-smooth $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$, we can stop algorithm when

$$
\frac{1}{\beta} u_{k}:=\frac{1}{\beta}\left(\gamma_{k}^{-1}\left(x_{k}-x_{k+1}\right)+\nabla f\left(x_{k+1}\right)-\nabla f\left(x_{k}\right)\right)
$$

is small (notation and reason will be motivated in future lecture)

- This is the plotted quantity on the previous slide
- We can use absolute or relative stopping conditions:
- absolute stopping conditions with small $\epsilon_{\text {abs }}>0$

$$
\frac{1}{\beta}\left\|u_{k}\right\|_{2} \leq \epsilon_{\mathrm{abs}} \quad \text { or } \quad \frac{1}{\beta}\left\|u_{k}\right\|_{2} \leq \epsilon_{\mathrm{abs}} \sqrt{n}
$$

- relative stopping condition with small $\epsilon_{\text {rel }}, \epsilon>0$ :

$$
\frac{1}{\beta} \frac{\left\|u_{k}\right\|_{2}}{\left\|x_{k}\right\|_{2}+\beta-1\left\|\nabla f\left(x_{k}\right)\right\|_{2}+\epsilon} \leq \epsilon_{\mathrm{rel}}
$$

- Problem considered solved to optimality if, say, $\frac{1}{\beta}\left\|u_{k}\right\|_{2} \leq 10^{-6}$
- Often lower accuracy of $10^{-3}$ or $10^{-4}$ is enough


## Outline

- Introducing proximal gradient method and examples
- Solving composite problem - Fixed-points and convergence
- Application to primal and dual problems


## Applying proximal gradient to primal problems

## Problem minimize $f(x)+g(x)$ :

## - Assumptions:

- $f$ smooth
- $g$ closed convex and prox friendly ${ }^{1}$
- Algorithm: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)$

Problem $\underset{x}{\operatorname{minimize}} f(L x)+g(x)$ :

- Assumptions:
- $f$ smooth (implies $f \circ L$ smooth)
- $g$ closed convex and prox friendly ${ }^{1}$
- Gradient $\nabla(f \circ L)(x)=L^{T} \nabla f(L x)$
- Algorithm: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} L^{T} \nabla f\left(L x_{k}\right)\right)$

[^0]
## Applying proximal gradient to dual problem

- Let us apply the proximal gradient method to the dual problem

$$
\underset{\mu}{\operatorname{minimize}} f^{*}(\mu)+g^{*}\left(-L^{T} \mu\right)
$$

- Assumptions:
- $f$ : closed convex and prox friendly
- $g$ : $\sigma$-strongly convex
- Why these assumptions?
- $f^{*}$ : closed convex and prox friendly
- $g^{*} \circ-L^{T}: \frac{\|L\|_{2}^{2}}{\sigma}$-smooth and convex
- Algorithm:

$$
\mu_{k+1}=\operatorname{prox}_{\gamma_{k} f^{*}}\left(\mu_{k}-\gamma_{k} \nabla\left(g^{*} \circ-L^{T}\right)\left(\mu_{k}\right)\right)
$$

Dual proximal gradient method - Explicit version 1

- We will make the dual proximal gradient method more explicit

$$
\mu_{k+1}=\operatorname{prox}_{\gamma_{k} f^{*}}\left(\mu_{k}-\gamma_{k} \nabla\left(g^{*} \circ-L^{T}\right)\left(\mu_{k}\right)\right)
$$

- Use $\nabla\left(g^{*} \circ-L^{T}\right)(\mu)=-L \nabla g^{*}\left(-L^{T} \mu\right)$ to get

$$
\begin{aligned}
x_{k} & =\nabla g^{*}\left(-L^{T} \mu_{k}\right) \\
\mu_{k+1} & =\operatorname{prox}_{\gamma_{k} f^{*}}\left(\mu_{k}+\gamma_{k} L x_{k}\right)
\end{aligned}
$$

Dual proximal gradient method - Explicit version 2

- Restating the previous formulation

$$
\begin{aligned}
x_{k} & =\nabla g^{*}\left(-L^{T} \mu_{k}\right) \\
\mu_{k+1} & =\operatorname{prox}_{\gamma_{k} f^{*}}\left(\mu_{k}+\gamma_{k} L x_{k}\right)
\end{aligned}
$$

- Use Moreau decomposition for prox:

$$
\operatorname{prox}_{\gamma f^{*}}(v)=v-\gamma \operatorname{prox}_{\gamma^{-1} f}\left(\gamma^{-1} v\right)
$$

to get

$$
\begin{aligned}
x_{k} & =\nabla g^{*}\left(-L^{T} \mu_{k}\right) \\
v_{k} & =\mu_{k}+\gamma_{k} L x_{k} \\
\mu_{k+1} & =v_{k}-\gamma_{k} \operatorname{prox}_{\gamma_{k}^{-1} f}\left(\gamma_{k}^{-1} v_{k}\right)
\end{aligned}
$$

## Dual proximal gradient method - Explicit version 3

- Restating the previous formulation

$$
\begin{aligned}
x_{k} & =\nabla g^{*}\left(-L^{T} \mu_{k}\right) \\
v_{k} & =\mu_{k}+\gamma_{k} L x_{k} \\
\mu_{k+1} & =v_{k}-\gamma_{k} \operatorname{prox}_{\gamma_{k}^{-1} f}\left(\gamma_{k}^{-1} v_{k}\right)
\end{aligned}
$$

- Use subdifferential formula, since $g^{*}$ differentiable:

$$
\nabla g^{*}(\nu)=\underset{x}{\operatorname{argmax}}\left(\nu^{T} x-g(x)\right)=\underset{x}{\operatorname{argmin}}\left(g(x)-\nu^{T} x\right)
$$

with $\nu=-L^{T} \mu_{k}$ to get

$$
\begin{aligned}
x_{k} & =\underset{x}{\operatorname{argmin}}\left(g(x)+\left(\mu_{k}\right)^{T} L x\right) \\
v_{k} & =\mu_{k}+\gamma_{k} L x_{k} \\
\mu_{k+1} & =v_{k}-\gamma_{k} \operatorname{prox}_{\gamma_{k}^{-1} f}\left(\gamma_{k}^{-1} v_{k}\right)
\end{aligned}
$$

- Can implement method without computing conjugate functions


## Dual proximal gradient method - Primal recovery

- Can we recover a primal solution from dual prox grad method?
- Let us use explicit version 1

$$
\begin{aligned}
x_{k} & =\nabla g^{*}\left(-L^{T} \mu_{k}\right) \\
\mu_{k+1} & =\operatorname{prox}_{\gamma_{k} f^{*}}\left(\mu_{k}+\gamma_{k} L x_{k}\right)
\end{aligned}
$$

and assume we have found fixed-point $(\bar{x}, \bar{\mu})$ : for some $\bar{\gamma}>0$,

$$
\begin{aligned}
& \bar{x}=\nabla g^{*}\left(-L^{T} \bar{\mu}\right) \\
& \bar{\mu}=\operatorname{prox}_{\bar{\gamma} f^{*}}(\bar{\mu}+\bar{\gamma} L \bar{x})
\end{aligned}
$$

- Fermat's rule for proximal step

$$
0 \in \partial f^{*}(\bar{\mu})+\bar{\gamma}^{-1}(\bar{\mu}-(\bar{\mu}+\bar{\gamma} L \bar{x}))=\partial f^{*}(\bar{\mu})-L \bar{x}
$$

is with $\bar{x}=\nabla g^{*}\left(-L^{T} \bar{\mu}\right)$ a primal-dual optimality condition

- So $x_{k}$ will solve primal problem if algorithm converges

Problems that prox-grad cannot solve

- Problem minimize $f(x)+g(x)$
- Assumptions: $f$ and $g$ convex but nondifferentiable
- No term differentiable, another method must be used
- Subgradient method
- Douglas-Rachford splitting
- Primal-dual methods

Problems that prox-grad cannot solve efficiently

- Problem minimize $f(x)+g(L x)$
- Assumptions:
- $f$ smooth
- $g$ nonsmooth convex
- $L$ arbitrary structured matrix
- Can apply proximal gradient method

$$
x_{k+1}=\underset{y}{\operatorname{argmin}}\left(g(L y)+\frac{1}{2 \gamma_{k}}\left\|y-\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)\right\|_{2}^{2}\right)
$$

but proximal operator of $g \circ L$

$$
\operatorname{prox}_{\gamma(g \circ L)}(z)=\underset{x}{\operatorname{argmin}}\left(g(L x)+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right)
$$

often not "prox friendly", i.e., it is expensive to evaluate


| Notation | Outline |
| :---: | :---: |
| - (Primal) Optimization variable notation: <br> - Optimization literature: $x, y, z$ (as in first part of course) <br> - Statistics literature: $\beta$ <br> - Machine learning literature: $\theta, w, b$ <br> - Reason: data, labels in statistics and machine learning are $x, y$ <br> - Will use machine learning notation in these lectures <br> - We collect training data in matrices (one example per row) $X=\left[\begin{array}{c} x_{1}^{T} \\ \vdots \\ x_{N}^{T} \end{array}\right] \quad Y=\left[\begin{array}{c} y_{1}^{T} \\ \vdots \\ y_{N}^{T} \end{array}\right]$ <br> - Columns $X_{j}$ of data matrix $X=\left[X_{1}, \ldots, X_{n}\right]$ are called features | - Supervised learning - Overview <br> - Least squares - Basics <br> - Nonlinear features <br> - Generalization, overfitting, and regularization <br> - Cross validation <br> - Feature selection <br> - Training problem properties |
| 9 | 10 |
| Regression training problem <br> - Objective: Find data model $m$ such that for all $(x, y)$ : $m(x)-y \approx 0$ <br> - Let model output $u=m(x)$; Examples of data misfit losses $\begin{aligned} L(u, y) & =\frac{1}{2}(u-y)^{2} \\ L(u, y) & =\|u-y\| \\ L(u, y) & = \begin{cases}\frac{1}{2}(u-y)^{2} & \text { if }\|u-v\| \leq c \\ c(\|u-y\|-c / 2) & \text { else }\end{cases} \end{aligned}$  <br> Square  <br> 1-norm <br> Huber <br> - Training: find model $m$ that minimizes sum of training set losses $\underset{m}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i}\right), y_{i}\right)$ | Supervised learning - Least squares <br> - Parameterize model $m$ and set a linear (affine) structure $m(x ; \theta)=w^{T} x+b$ <br> where $\theta=(w, b)$ are parameters (also called weights) <br> - Training: find model parameters that minimize training cost $\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)=\frac{1}{2} \sum_{i=1}^{N}\left(w^{T} x_{i}+b-y_{i}\right)^{2}$ <br> (note: optimization over model parameters $\theta$ ) <br> - Once trained, predict response of new input $x$ as $\hat{y}=w^{T} x+b$ |
| Example - Least squares <br> - Find affine function parameters that fit data: | Example - Least squares <br> - Find affine function parameters that fit data: <br> - Data points $(x, y)$ marked with (*), LS model $w x+b(-)$ |
| Example - Least squares <br> - Find affine function parameters that fit data: <br> - Data points $(x, y)$ marked with $(*)$, LS model $w x+b(-)$ <br> - Least squares finds affine function that minimizes squared distance | Solving for constant term <br> - Constant term $b$ also called bias term or intercept <br> - What is optimal $b$ ? $\underset{w, b}{\operatorname{minimize}} \frac{1}{2} \sum_{i=1}^{N}\left(w^{T} x_{i}+b-y_{i}\right)^{2}$ <br> - Optimality condition w.r.t. $b$ (gradient w.r.t. $b$ is 0 ): $0=N b+\sum_{i=1}^{N}\left(w^{T} x_{i}-y_{i}\right) \quad \Leftrightarrow \quad b=\bar{y}-w^{T} \bar{x}$ <br> where $\bar{x}=\frac{1}{N} \sum_{i=1}^{N} x_{i}$ and $\bar{y}=\frac{1}{N} \sum_{i=1}^{N} y_{i}$ are mean values |









| Lasso and correlated features <br> - Assume two equal features exist, e.g., $X_{1}=X_{2}$, lasso problem is $\operatorname{minimize} \frac{1}{2}\left\\|\left(w_{1}+w_{2}\right) X_{1}+\sum_{i=3}^{n} w_{i} X_{i}-Y\right\\|_{2}^{2}+\lambda\left(\left\|w_{1}\right\|+\left\|w_{2}\right\|+\left\\|w_{3: n}\right\\|_{1}\right)$ <br> - Assume $w^{*}$ solves the problem and let $\Delta:=w_{1}^{*}+w_{2}^{*}>0$ (wlog) <br> - Then all $w_{1} \in[0, \Delta]$ with $w_{2}=\Delta-w_{1}$ solves problem: <br> - quadratic cost unchanged since sum $w_{1}+w_{2}$ still $\Delta$ <br> - the remainder of the regularization part reduces to <br> - For almost correlated features: <br> - often only $w_{1}$ or $w_{2}$ nonzero (the one with slightly better fit) <br> - however, features highly correlated, if $X_{1}$ explains data so does $X_{2}$ | Elastic net <br> - Add Tikhonov regularization to the Lasso $\operatorname{minimize} \frac{1}{2}\\|X w-Y\\|^{2}+\lambda_{1}\\|w\\|_{1}+\frac{\lambda_{2}}{2}\\|w\\|_{2}^{2}$ <br> - This problem is called elastic net in statistics <br> - Can perform better with correlated features |
| :---: | :---: |
| Elastic net and correlated features <br> - Assume equal features $X_{1}=X_{2}$ and that $w^{*}$ solves the elastic net <br> - Let $\Delta:=w_{1}^{*}+w_{2}^{*}>0(\mathrm{wlog})$, then $w_{1}^{*}=w_{2}^{*}=\frac{\Delta}{2}$ <br> - Data fit cost still unchanged for $w_{2}=\Delta-w_{1}$ with $w_{1} \in[0, \Delta]$ <br> - Remaining (regularization) part is $\min _{w_{1}} \lambda_{1}\left(\left\|w_{1}\right\|+\left\|\Delta-w_{1}\right\|\right)+\lambda_{2}\left(w_{1}^{2}+\left(\Delta-w_{1}\right)^{2}\right)$ <br> which is minimized in the middle at $w_{1}=w_{2}=\frac{\Delta}{2}$ <br> - For highly correlated features, both (or none) probably selected | Group lasso <br> - Sometimes want groups of variables to be 0 or nonzero <br> - Introduce blocks $w=\left(w_{1}, \ldots, w_{p}\right)$ where $w_{i} \in \mathbb{R}^{n_{i}}$ <br> - The group Lasso problem is $\operatorname{minimize} \frac{1}{2}\\|X w-Y\\|_{2}^{2}+\lambda \sum_{i=1}^{p}\left\\|w_{i}\right\\|_{2}$ <br> (note $\\|\cdot\\|_{2}$-norm without square) <br> - With all $n_{i}=1$, it reduces to the Lasso <br> - Promotes block sparsity, meaning full block $w_{i} \in \mathbb{R}^{n_{i}}$ would be 0 |
| Outline <br> - Supervised learning - Overview <br> - Least squares - Basics <br> - Nonlinear features <br> - Generalization, overfitting, and regularization <br> - Cross validation <br> - Feature selection <br> - Training problem properties | Composite optimization <br> - Least squares problems are convex problems of the form $\underset{\theta}{\operatorname{minimize}} f(X \theta)+g(\theta)$ <br> where <br> - $f=\frac{1}{2}\\|\cdot-Y\\|_{2}^{2}$ is data misfit term <br> - $X$ is training data matrix (potentially extended with features) <br> - $g$ is regularization term (1-norm, squared 2-norm, group lasso) <br> - Function properties <br> - $f$ is 1 -strongly convex and 1 -smooth and $f \circ X$ is $\\|X\\|_{2}^{2}$-smooth <br> - $g$ is convex and possibly nondifferentiable <br> - Gradient $\nabla(f \circ X)(\theta)=X^{T}(X \theta-Y)$ |
|  |  |




## Training problem interpretation

- Every parameter choice $\theta=(w, b)$ gives hyperplane in data space:

$$
H:=\left\{x: w^{T} x+b=0\right\}=\{x: m(x ; \theta)=0\}
$$

- Training problem searches hyperplane to "best" separates classes
- Example - models with different parameters $\theta$ :



## What is "best" separation?

- The "best" separation is the one that minimizes the loss function
- Hyperplane for model $m(\cdot ; \theta)$ with parameter $\theta=\theta_{1}$

- Training loss



## What is "best" separation?

- The "best" separation is the one that minimizes the loss function
- Hyperplane for model $m(\cdot ; \theta)$ with parameter $\theta=\theta_{2}$ :

- Training loss:



## What is "best" separation?

- The "best" separation is the one that minimizes the loss function
- Hyperplane for model $m(\cdot ; \theta)$ with parameter $\theta=\theta_{4}$ :
raing loss:


10

## What is "best" separation?

- The "best" separation is the one that minimizes the loss function
- Hyperplane for model $m(\cdot ; \theta)$ with parameter $\theta=\theta_{3}$

- Training loss:



## What is "best" separation?

- The "best" separation is the one that minimizes the loss function
- Hyperplane for model $m(\cdot ; \theta)$ with parameter $\theta=\theta^{*}$ :

- Training loss:


10

## Fully separable data - Solution

- Let $\bar{\theta}=(\bar{w}, \bar{b})$ give model that separates data:

$$
\stackrel{*}{*}_{*^{*}{ }^{*}{\stackrel{u}{w_{*}^{*}}}_{*_{*}^{*}}^{* *} m(x ; \bar{\theta})=0}
$$

- Let $H_{\bar{\theta}}:=\left\{x: m(x ; \bar{\theta})=\bar{w}^{T} x+\bar{b}=0\right\}$ be hyperplane separates
- Training loss:


Fully separable data - Solution

- Also $2 \bar{\theta}=(2 \bar{w}, 2 \bar{b})$ separates data:

$$
\begin{aligned}
& \text { * } *^{*} * * \quad m(x ; 2 \bar{\theta})=0 \\
& \stackrel{*}{*}{ }^{*}{ }_{2}^{*}{ }^{*}{ }^{*}{ }^{*}{ }^{*}
\end{aligned}
$$

- Hyperplane $H_{2 \bar{\theta}}:=\left\{x: m(x ; 2 \bar{\theta})=2\left(\bar{w}^{T} x+\bar{b}\right)=0\right\}=H_{\bar{\theta}}$ same
- Training loss reduced since input $m(x ; 2 \bar{\theta})=2 m(x ; \bar{\theta})$ further out






| Example - Different polynomial model orders <br> - "Lifting" example with fewer samples and some mislabels <br> - Logistic regression (no regularization) polynomial features of degree: | Example - Different polynomial model orders <br> - "Lifting" example with fewer samples and some mislabels <br> - Logistic regression (no regularization) polynomial features of degree: 2 |
| :---: | :---: |
| Example - Different polynomial model orders <br> - "Lifting" example with fewer samples and some mislabels <br> - Logistic regression (no regularization) polynomial features of degree: 3 | Example - Different polynomial model orders <br> - "Lifting" example with fewer samples and some mislabels <br> - Logistic regression (no regularization) polynomial features of degree: 4 |
| Example - Different polynomial model orders <br> - "Lifting" example with fewer samples and some mislabels <br> - Logistic regression (no regularization) polynomial features of degree: 5 | Example - Different polynomial model orders <br> - "Lifting" example with fewer samples and some mislabels <br> - Logistic regression (no regularization) polynomial features of degree: 6 |
| Outline <br> - Classification <br> - Logistic regression <br> - Nonlinear features <br> - Overfitting and regularization <br> - Multiclass logistic regression <br> - Training problem properties | Overfitting <br> - Models with higher order polynomials overfit <br> - Logistic regression (no regularization) polynomial features of degree 6 <br> - Tikhonov regularization can reduce overfitting |

## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6
- Regularization parameter $\lambda$, training cost $J$, \# mislabels in training

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N}\left(\log \left(1+e^{x_{i}^{T} w+b}\right)-y_{i}\left(x_{i}^{T} w+b\right)\right)+\lambda\|w\|_{2}^{2}
$$

Regularization:

- Regularize only $w$ and not the bias term $b$
- Why? Model looses shift invariance if also $b$ regularized

Problem properties

- Problem is strongly convex in $w \Rightarrow$ optimal $w$ exists and is unique
- Optimal $b$ is bounded if examples from both classes exist



## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6
- Regularization parameter $\lambda$, training cost $J$, \# mislabels in training



## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6
- Regularization parameter $\lambda$, training cost $J$, \# mislabels in training



## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6 - Regularization parameter $\lambda$, training cost $J$, \# mislabels in training



## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6 - Regularization parameter $\lambda$, training cost $J$, \# mislabels in training



## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6 - Regularization parameter $\lambda$, training cost $J$, \# mislabels in training



## Example - Different regularization

- Regularized logistic regression and polynomial features of degree 6
- Regularization parameter $\lambda$, training cost $J$, \# mislabels in training


| Example - Different regularization <br> - Regularized logistic regression and polynomial features of degree 6 <br> - Regularization parameter $\lambda$, training cost $J$, \# mislabels in training | Generalization <br> - Interested in models that generalize well to unseen data <br> - Assess generalization using holdout or $k$-fold cross validation |
| :---: | :---: |
| Example - Validation data <br> - Regularized logistic regression and polynomial features of degree 6 <br> - $J$ and \# mislabels specify training/test values | Example - Validation data <br> - Regularized logistic regression and polynomial features of degree 6 <br> - $J$ and \# mislabels specify training/test values |
| Example - Validation data <br> - Regularized logistic regression and polynomial features of degree 6 <br> - $J$ and \# mislabels specify training/test values | Example - Validation data <br> - Regularized logistic regression and polynomial features of degree 6 <br> - $J$ and \# mislabels specify training/test values |
| Example - Validation data <br> - Regularized logistic regression and polynomial features of degree 6 <br> - $J$ and \# mislabels specify training/test values | Example - Validation data <br> - Regularized logistic regression and polynomial features of degree 6 <br> - $J$ and \# mislabels specify training/test values |



| Multiclass logistic loss function - Example <br> - Multiclass logistic loss for $K=3, u_{1}=1, y=e_{1}$ $L\left(\left(1, u_{2}, u_{3}\right), 1\right)=\log \left(e^{1}+e^{u_{2}}+e^{u_{3}}\right)-1$ <br> - Model outputs $u_{2} \ll 0, u_{3} \ll 0$ give smaller cost for label $y=e_{1}$ | Multiclass logistic loss function - Example <br> - Multiclass logistic loss for $K=3, u_{2}=-1, y=e_{1}$ $L\left(\left(u_{1},-1, u_{3}\right), 1\right)=\log \left(e^{u_{1}}+e^{-1}+e^{u_{3}}\right)-u_{1}$ <br> - Model outputs $u_{1} \gg 0$ and $u_{3} \ll 0$ give smaller cost for $y=e_{1}$ |
| :---: | :---: |
| Multiclass logistic regression - Training problem <br> - Affine data model $m(x ; \theta)=w^{T} x+b$ with $w=\left[w_{1}, \ldots, w_{K}\right] \in \mathbb{R}^{n \times K}, \quad b=\left[b_{1}, \ldots, b_{K}\right]^{T} \in \mathbb{R}^{K}$ <br> - One data model per class $m(x ; \theta)=\left[\begin{array}{c} m_{1}\left(x ; \theta_{1}\right) \\ \vdots \\ m_{K}\left(x ; \theta_{K}\right) \end{array}\right]=\left[\begin{array}{c} w_{1}^{T} x+b_{1} \\ \vdots \\ w_{K}^{T} x+b_{K} \end{array}\right]$ <br> - Training problem: $\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} \log \left(\sum_{j=1}^{K} e^{w_{j}^{T} x_{i}+b_{j}}\right)-y_{i}^{T}\left(w^{T} x_{i}+b\right)$ <br> where $y_{i}$ is "one-hot" encoding of label <br> - Problem is convex since affine model is used <br> - (Alt.: model $\sigma\left(w^{T} x+b\right)$ with $\sigma$ softmax and cross entropy loss) | Multiclass logistic regression - Prediction <br> - Assume model is trained and want to predict label for new data $x$ <br> - Predict class with parameter $\theta$ for $x$ according to: $\underset{j \in\{1, \ldots, K\}}{\operatorname{argmax}} m_{j}(x ; \theta)$ <br> i.e., class with largest model value (since trained to achieve this) |
| Special case - Binary logistic regression <br> - Consider two-class version and let <br> - $\Delta u=u_{1}-u_{2}, \Delta w=w_{1}-w_{2}$, and $\Delta b=b_{1}-b_{2}$ <br> - $\Delta u=m_{\text {bin }}(x ; \theta)=m_{1}\left(x ; \theta_{1}\right)-m_{2}\left(x ; \theta_{2}\right)=\Delta w^{T} x+\Delta b$ <br> - $y_{\text {bin }}=1$ if $y=(1,0)$ and $y_{\text {bin }}=0$ if $y=(0,1)$ <br> - Loss $L$ is equivalent to binary, but with different variables: $\begin{aligned} L(u, y) & =\log \left(e^{u_{1}}+e^{u_{2}}\right)-y_{1} u_{1}-y_{2} u_{2} \\ & =\log \left(1+e^{u_{1}-u_{2}}\right)+\log \left(e^{u_{2}}\right)-y_{1} u_{1}-y_{2} u_{2} \\ & =\log \left(1+e^{\Delta u}\right)-y_{1} u_{1}-\left(y_{2}-1\right) u_{2} \\ & =\log \left(1+e^{\Delta u}\right)-y_{\operatorname{bin}} \Delta u \end{aligned}$ | Example - Linearly separable data <br> - Problem with 7 classes |
| Example - Linearly separable data <br> - Problem with 7 classes and affine multiclass model | Example - Quadratically separable data <br> - Same data, new labels in 6 classes |


$\square$









| Example - Laplacian Kernel <br> - Regularized SVM with Laplacian Kernel with $\sigma=1$ <br> - Regularization parameter: $\lambda=0.12915$ | Example - Laplacian Kernel <br> - Regularized SVM with Laplacian Kernel with $\sigma=1$ <br> - Regularization parameter: $\lambda=0.46416$ |
| :---: | :---: |
| Example - Laplacian Kernel <br> - Regularized SVM with Laplacian Kernel with $\sigma=1$ <br> - Regularization parameter: $\lambda=1.6681$ | Example - Laplacian Kernel <br> - Regularized SVM with Laplacian Kernel with $\sigma=1$ <br> - Regularization parameter: $\lambda=5.9948$ |
| Example - Laplacian Kernel <br> - Regularized SVM with Laplacian Kernel with $\sigma=1$ <br> - Regularization parameter: $\lambda=21.5443$ | Example - Laplacian Kernel <br> - What if there is no structure in data? (Labels are randomly set) |
| Example - Laplacian Kernel <br> - What if there is no structure in data? (Labels are randomly set) <br> - Regularized SVM Laplacian Kernel, regularization parameter: $\lambda=0.01$ <br> - Linearly separable in high dimensional feature space <br> - Can be prone to overfitting $\Rightarrow$ Regularize and use cross validation | Outline <br> - Classification <br> - Support vector machines <br> - Nonlinear features <br> - Overfitting and regularization <br> - Dual problem <br> - Kernel SVM <br> - Training problem properties |




| Prediction | Outline |
| :---: | :---: |
| - Prediction as in least squares and multiclass logistic regression <br> - Assume model $m(x ; \theta)$ trained and "optimal" $\theta^{\star}$ found <br> - Regression: <br> - Predict response for new data $x$ using $\hat{y}=m\left(x ; \theta^{\star}\right)$ <br> - Classification (with no final layer activation): <br> - We have one model $m_{j}\left(x ; \theta^{\star}\right)$ output for each class <br> - Predict class belonging for new data $x$ according to $\underset{j \in\{1, \ldots, K\}}{\operatorname{argmax}} m_{j}\left(x ; \theta^{\star}\right)$ <br> i.e., class with largest model value (since loss designed this way) | - Deep learning <br> - Learning features <br> - Model properties and activation functions <br> - Loss landscape <br> - Residual networks <br> - Overparameterized networks <br> - Generalization and regularization <br> - Generalization - Norm of weights <br> - Generalization - Flatness of minima <br> - Backpropagation <br> - Vanishing and exploding gradients |
| 10 |  |
| Learning features <br> - Convex methods use prespecified feature maps (or kernels) <br> - Deep learning instead learns feature map during training <br> - Define parameter dependent feature vector: $\phi(x ; \theta):=\sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)$ <br> - Model becomes $m(x ; \theta)=W_{n} \phi(x ; \theta)+b_{n}$ <br> - Inserted into training problem: $\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(W_{n} \phi\left(x_{i} ; \theta\right)+b_{n}, y_{i}\right)$ <br> same as before, but with learned (parameter-dependent) features <br> - Learning features at training makes training nonconvex | Learning features - Graphical representation <br> - Fixed features gives convex training problems <br> - Learning features gives nonconvex training problems <br> - Output of last activation function is feature vector |
| Optimizing only final layer <br> - Assume: <br> - that parameters $\bar{\theta}_{f}$ in the layers in the square are fixed <br> - that we optimize only the final layer parameters <br> - that the loss is a (binary) logistic loss <br> - What can you say about the training problem? | Optimizing only final layer <br> - Assume: <br> - that parameters $\bar{\theta}_{f}$ in the layers in the square are fixed <br> - that we optimize only the final layer parameters <br> - that the loss is a (binary) logistic loss <br> - What can you say about the training problem? <br> - It reduces to logistic regression with fixed features $\phi\left(x_{i} ; \bar{\theta}_{f}\right)$ $\operatorname{minimize}_{\theta=\left(W_{n}, b_{n}\right)} \sum_{i=1}^{N} L\left(W_{n} \phi\left(x_{i} ; \bar{\theta}_{f}\right)+b_{n}, y_{i}\right)$ <br> - The training problem is convex |
| Design choices | Outline |
| Many design choices in building model to create good features <br> - Number of layers <br> - Width of layers <br> - Types of layers <br> - Types of activation functions <br> - Different model structures (e.g., residual network) | - Deep learning <br> - Learning features <br> - Model properties and activation functions <br> - Loss landscape <br> - Residual networks <br> - Overparameterized networks <br> - Generalization and regularization <br> - Generalization - Norm of weights <br> - Generalization - Flatness of minima <br> - Backpropagation <br> - Vanishing and exploding gradients |
| 14 | 15 |

- Recall model
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{i}$ and $b_{i}$
- Assume that all activation functions are (Leaky)ReLU
- What can you say about the properties of $m(\cdot ; \theta)$ for fixed $\theta$ ?
- Recall model
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$
where $\theta$ contains all $W_{i}$ and $b_{i}$
- Assume that all activation functions are (Leaky)ReLU
- What can you say about the properties of $m(\cdot ; \theta)$ for fixed $\theta$ ? - It is continuous piece-wise affine


## 1D Regression - Model properties

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyReLU

- Vertical lines show kinks

1D Regression - Model properties

- Fully connected, layers widths: 5,5,5,1,1 (78 params), Tanh

- No kinks for Tanh


## Identity activation

- Do we need nonlinear activation functions?
- What can you say about model if all $\sigma_{j}=\mathrm{Id}$ in $m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{j}$ and $b_{j}$


## Identity activation

- Do we need nonlinear activation functions?
- What can you say about model if all $\sigma_{j}=\mathrm{Id}$ in $m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{j}$ and $b_{j}$
- We then get

$$
\begin{aligned}
m(x ; \theta) & :=W_{n}\left(W_{n-1}\left(\cdots\left(W_{2}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n} \\
& =\underbrace{W_{n} W_{n-1} \cdots W_{2} W_{1}}_{W} x+\underbrace{b_{n}+\sum_{l=2}^{n-1} W_{n} \cdots W_{l} b_{l-1}}_{b} \\
& =W x+b
\end{aligned}
$$

which is linear in $x$ (but training problem nonconvex)

## Network with identity activations - Example

- Fully connected, layers widths: 5,5,5,1,1 (78 params), Identity

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients
- Recall mode
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$
where $\theta$ includes all $W_{j}$ and $b_{j}$ and training problem

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

- If all $\sigma_{j}$ LeakyReLU and $L(u, y)=\frac{1}{2}\|u-y\|_{2}^{2}$, then for fixed $x, y$
- $m(x ; \cdot)$ is continuous piece-wise polynomial (cpp) of degree $n$ in $\theta$
- $L(m(x ; \theta), y)$ is cpp of degree $2 n$ in $\theta$
where both model output and loss can grow fast
- If $\sigma_{j}$ is instead Tanh
- model no longer piece-wise polynomial (but "more" nonlinear)
- model output grows slower since $\sigma_{j}: \mathbb{R} \rightarrow(-1,1)$


## Loss landscape - Leaky ReLU

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- First choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Leaky ReLU

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where - $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Second choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Leaky ReLU

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{*}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Third choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Tanh

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where - $\theta^{*}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- First choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Tanh

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Second choice of $\theta_{1}$ and $\theta_{2}$ :




| Outline | Why overparameterization? |  |
| :---: | :---: | :---: |
| - Deep learning <br> - Learning features <br> - Model properties and activation functions <br> - Loss landscape <br> - Residual networks <br> - Overparameterized networks <br> - Generalization and regularization <br> - Generalization - Norm of weights <br> - Generalization - Flatness of minima <br> - Backpropagation <br> - Vanishing and exploding gradients | - Neural networks are often overparameterized in practice <br> - Why? They often perform better than underparameterized |  |
| 31 l 32 |  |  |
| What is overparameterization? <br> - We mean that many solutions exist that can: <br> - fit all data points ( 0 training loss) in regression <br> - correctly classify all training examples in classification <br> - This requires (many) more parameters than training examples <br> - Need wide and deep enough networks <br> - Can result in overfitting <br> - Questions: <br> - Which of all solutions give best generalization? <br> - (How) can network design affect generalization? | Overparameterization - An example <br> - Assume fully connected network with <br> - input data $x_{i} \in \mathbb{R}^{p}$ <br> - $n$ layers and $N \approx p^{2}$ samples <br> - same width throughout (except last layer, which can be neglected) <br> - What is the relation between number of weights and samples? | 34 |
| Overparameterization - An example | Outline |  |
| - Assume fully connected network with <br> - input data $x_{i} \in \mathbb{R}^{p}$ <br> - $n$ layers and $N \approx p^{2}$ samples <br> - same width throughout (except last layer, which can be neglected) <br> - What is the relation between number of weights and samples? <br> - We have: <br> - Number of parameters approximately: $\left(W_{j}\right)_{l k}: p^{2} n$ and $\left(b_{j}\right)_{l}: p n$ <br> - Then \#weights $\#$ samples $\approx \frac{p^{2} n}{p^{2}}=n$ more weights than samples | - Deep learning <br> - Learning features <br> - Model properties and activation functions <br> - Loss landscape <br> - Residual networks <br> - Overparameterized networks <br> - Generalization and regularization <br> - Generalization - Norm of weights <br> - Generalization - Flatness of minima <br> - Backpropagation <br> - Vanishing and exploding gradients |  |
| 34 ( 35 |  |  |
| Generalization | Regularization |  |
| - Most important for model to generalize well to unseen data <br> - General approach in training <br> - Train a model that is too expressive for the underlying data <br> - Overparameterization in deep learning <br> What regularization techniques in DL are you familiar with? <br> - Use regularization to <br> - find model of appropriate (lower) complexity <br> - favor models with desired properties |  |  |
| 36 |  | 37 |

- Reduce number of parameters
- Sparse weight tensors (e.g., convolutional layers)
- Subsampling (gives fewer parameters deeper in network)
- Explicit regularization term in cost function, e.g., Tikhonov
- Data augmentation - more samples, artificial often OK
- Early stopping - stop algorithm before convergence
- Dropouts
- ...
- Regularization can be explicit or implicit
- Explicit - Introduce something with intent to regularize:
- Add cost function to favor desirable properties
- Design (adapt) network to have regularizing properties
- Implicit - Use something with regularization as byproduct:
- Use algorithm that finds favorable solution among many
- Will look at implicit regularization via SGD

Will here discuss generalization via:

- Norm of parameters - leads to implicit regularization via SGD
- Flatness of minima - leads to implicit regularization via SGD
- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Lipschitz continuity of ReLU networks

- Assume that all activation functions 1-Lipschitz continuous
- The neural network model $m(\cdot ; \theta)$ is Lipschitz continuous in $x$,

$$
\left\|m\left(x_{1} ; \theta\right)-m\left(x_{2} ; \theta\right)\right\|_{2} \leq L\left\|x_{1}-x_{2}\right\|_{2}
$$

for fixed $\theta$, e.g., the $\theta$ obtained after training

- This means output differerences are bounded by input differences
- A Lipschitz constant $L$ is given by

$$
L=\left\|W_{n}\right\|_{2} \cdot\left\|W_{n-1}\right\|_{2} \cdots\left\|W_{1}\right\|_{2}
$$

since activation functions are 1-Lipschitz continuous

- For residual layers each $\left\|W_{j}\right\|_{2}$ replaced by $\left(1+\left\|W_{j}\right\|_{2}\right)$


## Desired Lipschitz constant

- Overparameterization gives many solutions that perfectly fit data
- Would you favor one with high or low Lipschitz constant $L$ ?


## Small norm likely to generalize better

- Smaller Lipschitz constant probably generalizes better if perfect fit
- "Similar inputs give similar outputs", recall

$$
\left\|m\left(x_{1} ; \theta\right)-m\left(x_{2} ; \theta\right)\right\|_{2} \leq L\left\|x_{1}-x_{2}\right\|_{2}
$$

with a Lipschitz constant is given by

$$
L=\left\|W_{n}\right\|_{2} \cdot\left\|W_{n-1}\right\|_{2} \cdots\left\|W_{1}\right\|_{2}
$$

or with $\left\|W_{j}\right\|_{2}$ replaced by $\left(1+\left\|W_{j}\right\|_{2}\right)$ for residual layers

- Smaller weight norms give better generalization if perfect fit
- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ (888 params)
- Norm of weights (with perfect fit): 72




## Generalization from loss landscape

- Training set $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{N}$ and training problem:

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

- Test set $\left\{\left(\hat{x}_{i}, \hat{y}_{i}\right)\right\}_{i=1}^{\hat{N}}, \theta$ generalizes well if test loss small

$$
\sum_{i=1}^{\hat{N}} L\left(m\left(\hat{x}_{i} ; \theta\right), \hat{y}_{i}\right)
$$

- By overparameterization, we can for each $\left(\hat{x}_{i}, \hat{y}_{i}\right)$ find $\hat{\theta}_{i}$ so that

$$
L\left(m\left(\hat{x}_{i} ; \theta\right), \hat{y}_{i}\right)=L\left(m\left(x_{j_{i}} ; \theta+\hat{\theta}_{i}\right), y_{j_{i}}\right)
$$

for all $\theta$ given a (similar) $\left(x_{j_{i}}, y_{j_{i}}\right)$ pair in training set

- Evaluate test loss by training loss at shifted points $\theta+\hat{\theta}_{i}{ }^{1)}$
- Test loss small if original individual loss small at all $\theta+\hat{\theta}_{i}$
- Previous figure used same $\hat{\theta}_{i}=\hat{\theta}$ for all $i$
${ }^{1)}$ Don't compute in practice just thought experiment to connet get genalization to training loss


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?



## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses
- Let us evaluate test loss by shifting individual training losses

- Compute gradient/Jacobian of

$$
L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

w.r.t. $\theta=\left\{\left(W_{j}, b_{j}\right)\right\}_{j=1}^{n}$, where
$m\left(x_{i} ; \theta\right)=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x_{i}+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$

- Rewrite as function with states $z_{j}$

$$
\begin{aligned}
\text { where } & z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right) \text { for } j \in\{1, \ldots, n\} \\
\text { and } & z_{1}=x_{i}
\end{aligned}
$$

where $\sigma_{n}(u) \equiv u$

## Graphical representation

- Per sample loss function

|  | $L\left(z_{n+1}, y_{i}\right)$ |
| ---: | :--- |
| where | $z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)$ for $j \in\{1, \ldots, n\}$ |
| and | $z_{1}=x_{i}$ |

where $\sigma_{n}(u) \equiv u$

- Graphical representation



## Backpropagation - Chain rule

- Jacobian of $L$ w.r.t. $W_{j}$ and $b_{j}$ can be computed as

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}} \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial W_{j}} \\
\frac{\partial L}{\partial b_{j}} & =\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}} \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial b_{j}}
\end{aligned}
$$

where we mean derivative w.r.t. first argument in $L$

- Backpropagation evaluates partial Jacobians as follows

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\left(\left(\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right) \frac{\partial z_{j+1}}{\partial W_{j}} \\
\frac{\partial L}{\partial b_{j}} & =\left(\left(\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right) \frac{\partial z_{j+1}}{\partial b_{j}}
\end{aligned}
$$

## Backpropagation - Forward and backward pass

- Jacobian of $L\left(z_{n+1}, y_{i}\right)$ w.r.t. $z_{n+1}$ (transpose of gradient)
- Computing Jacobian of $L\left(z_{n+1}, y_{i}\right)$ requires $z_{n+1}$
$\Rightarrow$ forward pass: $z_{1}=x_{i}, z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)$
- Backward pass, store $\delta_{j}$

$$
\frac{\partial L}{\partial z_{j+1}}=\underbrace{(\underbrace{(\underbrace{\left.\left.\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right)}_{\delta_{n+1}^{T}}}_{\delta_{n}^{T}})}_{\delta_{j+1}^{T}}
$$

- Compute

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial W_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial W_{j}} \\
\frac{\partial L}{\partial b_{j}} & =\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial b_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial b_{j}}
\end{aligned}
$$

## Dimensions

- Let $z_{j} \in \mathbb{R}^{n_{j}}$, consequently $W_{j} \in \mathbb{R}^{n_{j+1} \times n_{j}}, b_{j} \in \mathbb{R}^{n_{j+1}}$
- Dimensions

- Vector matrix multiplies except for in last step
- Multiplication with tensor $\frac{\partial z_{j+1}}{\partial W_{i}}$ can be simplified
- Backpropagation variables $\delta_{j} \in \mathbb{R}^{n_{j}}$ are vectors (not matrices)


## Partial Jacobian $\frac{\partial z_{j+1}}{\partial z_{j}}$

- Recall relation $z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)$ and let $v_{j}=W_{j} z_{j}+b_{j}$
- Chain rule gives

$$
\begin{aligned}
\frac{\partial z_{j+1}}{\partial z_{j}} & =\frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial z_{j}}=\operatorname{diag}\left(\sigma_{j}^{\prime}\left(v_{j}\right)\right) \frac{\partial v_{j}}{\partial z_{j}} \\
& =\operatorname{diag}\left(\sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) W_{j}
\end{aligned}
$$

where, with abuse of notation (notation overloading)

$$
\sigma_{j}^{\prime}(u)=\left[\begin{array}{c}
\sigma_{j}^{\prime}\left(u_{1}\right) \\
\vdots \\
\sigma_{j}^{\prime}\left(u_{n_{j+1}}\right)
\end{array}\right]
$$

- Reason: $\sigma_{j}(u)=\left[\sigma_{j}\left(u_{1}\right), \ldots, \sigma_{j}\left(u_{n_{j+1}}\right)\right]^{T}$ with $\sigma_{j}: \mathbb{R}^{n_{j+1}} \rightarrow \mathbb{R}^{n_{j+1}}$, gives

$$
\frac{d \sigma_{j}}{d u}=\left[\begin{array}{lll}
\sigma_{j}^{\prime}\left(u_{1}\right) & & \\
& \ddots & \\
& & \sigma_{j}^{\prime}\left(u_{n_{j+1}}\right)
\end{array}\right]=\operatorname{diag}\left(\sigma_{j}^{\prime}(u)\right)
$$

## Partial Jacobian $\delta_{j}^{T}=\frac{\partial L}{\partial z_{j}}$

- For any vector $\delta_{j+1} \in \mathbb{R}^{n_{j+1} \times 1}$, we have

$$
\begin{aligned}
\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial z_{j}} & =\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) W_{j} \\
& =\left(W_{j}^{T}\left(\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right)^{T}\right)^{T} \\
& =\left(W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right)^{T}
\end{aligned}
$$

where $\cdot$ is element-wise (Hadamard) product

- We have defined $\delta_{n+1}^{T}=\frac{\partial L}{\partial z_{n+1}}$, then

$$
\delta_{n}^{T}=\frac{\partial L}{\partial z_{n}}=\delta_{n+1}^{T} \frac{\partial z_{n+1}}{\partial z_{n}}=(\underbrace{W_{n}^{T}\left(\delta_{n+1} \odot \sigma_{n}^{\prime}\left(W_{n} z_{n}+b_{n}\right)\right)}_{\delta_{n}})^{T}
$$

- Consequently, using induction

$$
\delta_{j}^{T}=\frac{\partial L}{\partial z_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial z_{j}}=(\underbrace{W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)}_{\delta_{j}})^{T}
$$

## Information needed to compute $\frac{\partial L}{\partial z_{j}}$

- To compute first Jacobian $\frac{\partial L}{\partial z_{n}}$, we need $z_{n} \Rightarrow$ forward pass
- Computing

$$
\frac{\partial L}{\partial z_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial z_{j}}=\left(W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right)^{T}=\delta_{j}^{T}
$$

is done using a backward pass

$$
\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)
$$

- All $z_{j}$ (or $v_{j}=W_{j} z_{j}+b_{j}$ ) need to be stored for backward pass


| Partial Jacobian $\frac{\partial L}{\partial W_{j}}$ <br> - Computed by $\frac{\partial L}{\partial W_{j}}=\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial W_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial W_{j}}$ <br> where $z_{j+1}=\sigma_{j}\left(v_{j}\right)$ and $v_{j}=W_{j} z_{j}+b_{j}$ <br> - Recall $\frac{\partial z_{j+1}}{\partial W_{l}}$ is 3D tensor, compute Jacobian w.r.t. row $l\left(W_{j}\right)_{l}$ $\begin{aligned} \delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial\left(W_{j}\right)_{l}} & =\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial\left(W_{j}\right)_{l}}=\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(v_{j}\right)\right)\left[\begin{array}{c} 0 \\ \vdots \\ z_{j}^{T} \\ \vdots \\ 0 \end{array}\right] \\ & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)^{T}\left[\begin{array}{c} 0 \\ \vdots \\ z_{j}^{T} \\ \vdots \\ 0 \end{array}\right]=\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) l_{l}^{T} \end{aligned}$ | Partial Jacobian $\frac{\partial L}{\partial W_{j}}$ cont'd <br> - Stack Jacobians w.r.t. rows to get full Jacobian: $\begin{aligned} \frac{\partial L}{\partial W_{j}} & =\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial W_{j}}=\left[\begin{array}{c} \delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial\left(W_{j}\right)_{1}} \\ \vdots \\ \delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial\left(W_{j}\right)_{j+1}} \end{array}\right]=\left[\begin{array}{c} \left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)_{1} z_{j}^{T} \\ \vdots \\ \left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)_{n_{j+1}} z_{j}^{T} \end{array}\right] \\ & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) z_{j}^{T} \end{aligned}$ <br> for all $j \in\{1, \ldots, n-1\}$ <br> - Dimension of result is $n_{j+1} \times n_{j}$, which matches $W_{j}$ <br> - This is used to update $W_{j}$ weights in algorithm |
| :---: | :---: |
| Partial Jacobian $\frac{\partial L}{\partial b_{j}}$ <br> - Recall $z_{j+1}=\sigma_{j}\left(v_{j}\right)$ where $v_{j}=W_{j} z_{j}+b_{j}$ <br> - Computed by $\begin{aligned} \frac{\partial L}{\partial b_{j}} & =\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial b_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial b_{j}}=\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(v_{j}\right)\right) \\ & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)^{T} \end{aligned}$ | Backpropagation summarized <br> 1. Forward pass: Compute and store $z_{j}$ (or $v_{j}=W_{j} z_{j}+b_{j}$ ): $z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)$ <br> where $z_{1}=x_{i}$ and $\sigma_{n}=\mathrm{Id}$ <br> 2. Backward pass: $\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)$ <br> with $\delta_{n+1}=\frac{\partial L}{\partial z_{n+1}}$ <br> 3. Weight update Jacobians (used in SGD) $\begin{aligned} \frac{\partial L}{\partial W_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) z_{j}^{T} \\ \frac{\partial L}{\partial b_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} x_{j}+b_{j}\right)\right)^{T} \end{aligned}$ |
| Backpropagation - Residual networks <br> 1. Forward pass: Compute and store $z_{j}$ (or $v_{j}=W_{j} z_{j}+b_{j}$ ): $z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)+z_{j}$ <br> where $z_{1}=x_{i}$ and $\sigma_{n}=\mathrm{Id}$ <br> 2. Backward pass: $\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)+\delta_{j+1}$ <br> with $\delta_{n+1}=\frac{\partial L}{\partial z_{n+1}}$ <br> 3. Weight update Jacobians (used in SGD) $\begin{aligned} \frac{\partial L}{\partial W_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) z_{j}^{T} \\ \frac{\partial L}{\partial b_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} x_{j}+b_{j}\right)\right)^{T} \end{aligned}$ | Outline <br> - Deep learning <br> - Learning features <br> - Model properties and activation functions <br> - Loss landscape <br> - Residual networks <br> - Overparameterized networks <br> - Generalization and regularization <br> - Generalization - Norm of weights <br> - Generalization - Flatness of minima <br> - Backpropagation <br> - Vanishing and exploding gradients |
| Vanishing and exploding gradient problem <br> - For some activation functions, gradients can vanish <br> - For other activation functions, gradients can explode | Vanishing gradient example: Sigmoid <br> - Assume $\left\\|W_{j}\right\\| \leq 1$ for all $j$ and $\left\\|\delta_{n+1}\right\\| \leq C$ <br> - Maximal derivative of sigmoid $(\sigma)$ is 0.25 <br> - Then $\begin{aligned} \left\\|\frac{\partial L}{\partial z_{j}}\right\\| & =\left\\|\delta_{j}\right\\|=\left\\|W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right\\| \leq 0.25\left\\|\delta_{j+1}\right\\| \\ & \leq 0.25^{n-j+1}\left\\|\delta_{n+1}\right\\| \leq 0.25^{n-j+1} C \end{aligned}$ <br> - Hence, as $n$ grows, gradients can become very small for small $i$ <br> - In general, vanishing gradient if $\sigma^{\prime}<1$ everywhere <br> - Similar reasoning: exploding gradient if $\sigma^{\prime}>1$ everywhere <br> - Hence, need $\sigma^{\prime}=1$ in important regions |
| 71 | 72 |




- Algorithm overview
- Sometimes first-order methods computationally too expensive
- Stochastic gradient methods:
- Use stochastic approximation of gradient
- For finite sum problems, cheaply computed approximation exists
- Coordinate-wise updates:
- Update only one (or block of) coordinates in every iteration
- via direct minimization
- via proximal gradient step
- Can update coordinates in cyclic fashion
- Stronger convergence results if random selection of block
- Efficient if cost of updating one coordinate is $1 / n$ of full update
- Can solve huge scale problems
- Will cover randomized coordinate and stochastic methods


## Outline

- Convergence and convergence rates
- Proving convergence rates


## Types of convergence

## Convergence for stochastic algorithms

- Stochastic algorithms described by stochastic process $\left(x_{k}\right)_{k \in \mathbb{N}}$
- When algorithm is run, we get realization of stochastic process
- Let $x^{\star}$ be solution to composite problem and $p^{\star}=f\left(x^{\star}\right)+g\left(x^{\star}\right)$
- We will see convergence of different quantities in different settings
- For deterministic algorithms that generate $\left(x_{k}\right)_{k \in \mathbb{N}}$, we will see
- Sequence convergence: $x_{k} \rightarrow x^{\star}$
- Function value convergence: $f\left(x_{k}\right)+g\left(x_{k}\right) \rightarrow p^{\star}$
- If $g=0$, gradient norm convergence: $\left\|\nabla f\left(x_{k}\right)\right\|_{2} \rightarrow 0$
- Convergence is stronger as we go up the list
- First two common in convex setting, last in nonconvex
- Expected distance to solution: $\sum_{k=0}^{\infty} \mathbb{E}\left[\left\|x_{k}-x^{\star}\right\|_{2}\right]<\infty$
- Expected function value: $\sum_{k=0}^{\infty} \mathbb{E}\left[f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star}\right]<\infty$
- If $g=0$, expected gradient norm: $\sum_{k=0}^{\infty} \mathbb{E}\left[\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}\right]<\infty$
- Sometimes arrive at weaker conclusion, when $g=0$, that, e.g.,:
- Expected smallest function value: $\mathbb{E}\left[\min _{l \in\{0, \ldots, k\}} f\left(x_{l}\right)-p^{\star}\right] \rightarrow 0$
- Expected smallest gradient norm: $\mathbb{E}\left[\underset{l \in\{0, \ldots, k\}}{l \in\{0, \ldots, k\}} \min _{l}\left\|\nabla f\left(x_{l}\right)\right\|_{2}\right] \rightarrow 0$
- Says what happens with expected value of different quantities


## Algorithm realizations - Summable case

- Will conclude that sequence of expected values containing, e.g.,:
$\mathbb{E}\left[\left\|x_{k}-x^{\star}\right\|_{2}\right]$ or $\mathbb{E}\left[f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star}\right]$ or $\mathbb{E}\left[\left\|\nabla f\left(x_{k}\right)\right\|_{2}\right]$
is summable, where all quantities are nonnegative
- What happens with the actual algorithm realizations?
- We can make conclusions by the following result: If
- $\left(Z_{k}\right)_{k \in \mathbb{N}}$ is a stochastic process with $Z_{k} \geq 0$
- the sequence $\left\{\mathbb{E}\left[Z_{k}\right]\right\}_{k \in \mathbb{N}}$ is summable: $\sum_{k=0}^{\infty} \mathbb{E}\left[Z_{k}\right]<\infty$ then almost sure convergence to 0 :

$$
P\left(\lim _{k \rightarrow \infty} Z_{k}=0\right)=1
$$

i.e., convergence to 0 with probability 1

## Algorithm realizations - Convergent case

- Will conclude that sequence of expected values containing, e.g.,:

$$
\mathbb{E}\left[\min _{l \in\{0, \ldots, k\}} f\left(x_{l}\right)-p^{\star}\right] \quad \text { or } \quad \mathbb{E}\left[\min _{l \in\{0, \ldots, k\}}\left\|\nabla f\left(x_{l}\right)\right\|_{2}\right]
$$

converges to 0 , where all quantities are nonnegative

- What happens with the actual algorithm realizations?
- We can make conclusions by the following result: If
- $\left(Z_{k}\right)_{k \in \mathbb{N}}$ is a stochastic process with $Z_{k} \geq 0$
- the expected value $\mathbb{E}\left[Z_{k}\right] \rightarrow 0$ as $k \rightarrow \infty$
then convergence to 0 in probability; for all $\epsilon>0$

$$
\lim _{k \rightarrow \infty} P\left(Z_{k}>\epsilon\right)=0
$$

which is weaker than almost sure convergence to 0

## Convergence rates

- We have only talked about convergence, not convergence rate
- Rates indicate how fast (in iterations) algorithm reaches solution
- Typically divided into:
- Sublinear rates
- Linear rates (also called geometric rates)
- Quadratic rates (or more generally superlinear rates)
- Sublinear rates slowest, quadratic rates fastest
- Linear rates further divided into Q-linear and R-linear
- Quadratic rates further divided into Q-quadratic and R-quadratic


## Linear rates

- A Q-linear rate with factor $\rho \in[0,1)$ can be:

$$
\begin{aligned}
f\left(x_{k+1}\right)+g\left(x_{k+1}\right)-p^{\star} & \leq \rho\left(f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star}\right) \\
\mathbb{E}\left[\left\|x_{k+1}-x^{\star}\right\|_{2}\right] & \leq \rho \mathbb{E}\left[\left\|x_{k}-x^{\star}\right\|_{2}\right]
\end{aligned}
$$

- An R-linear rate with factor $\rho \in[0,1)$ and some $C>0$ can be:

$$
\left\|x_{k}-x^{\star}\right\|_{2} \leq \rho^{k} C
$$

this is implied by Q-linear rate and has exponential decrease

- Linear rate is superlinear if $\rho=\rho_{k}$ and $\rho_{k} \rightarrow 0$ as $k \rightarrow \infty$
- Examples:
- (Accelerated) proximal gradient with strongly convex cost
- Randomized coordinate descent with strongly convex cost
- BFGS has local superlinear with strongly convex cost
- but SGD with strongly convex cost gives sublinear rate

- If we suspect linear or quadratic convergence for $V_{k} \geq 0$

$$
V_{k+1} \leq \rho V_{k}^{p}
$$

where $\rho \in[0,1)$ and $p=1$ or $p=2$ and $V_{k}$ can, e.g., be
$V_{k}=\left\|x_{k}-x^{\star}\right\|_{2} \quad$ or $\quad V_{k}=f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star} \quad$ or $\quad V_{k}=\left\|\nabla f\left(x_{k}\right)\right\|_{2}$

- Can prove by starting with $V_{k+1}$ (or $V_{k+1}^{2}$ ) and continue using
- function class inequalities
- algorithm equalities
- propeties of norms
- 
- Assume we want to show sublinear convergence of some $R_{k} \geq 0$
- This typically requires finding a Lyapunov inequality:

$$
V_{k+1} \leq V_{k}+W_{k}-R_{k}
$$

where

- $\left(V_{k}\right)_{k \in \mathbb{N}},\left(W_{k}\right)_{k \in \mathbb{N}}$, and $\left(R_{k}\right)_{k \in \mathbb{N}}$ are nonnegative real numbers - $\left(W_{k}\right)_{k \in \mathbb{N}}$ is summable, i.e., $\bar{W}:=\sum_{k=0}^{\infty} W_{k}<\infty$
- Such a Lyapunov inequality can be found by using
- function class inequalities
- algorithm equalities
- propeties of norms
- ...


## Lyapunov inequality consequences

- From the Lyapunov inequality:

$$
V_{k+1} \leq V_{k}+W_{k}-R_{k}
$$

we can conclude that

- $V_{k}$ is nonincreasing if all $W_{k}=0$
- $V_{k}$ converges as $k \rightarrow \infty$ (will not prove)
- Recursively applying the inequality for $l \in\{k, \ldots, 0\}$ gives

$$
V_{k+1} \leq V_{0}+\sum_{l=0}^{k} W_{l}-\sum_{l=0}^{k} R_{l} \leq V_{0}+\bar{W}-\sum_{l=0}^{k} R_{l}
$$

where $\bar{W}$ is infinite sum of $W_{k}$, this implies

$$
\sum_{l=0}^{k} R_{l} \leq V_{0}-V_{k+1}+\sum_{l=0}^{k} W_{l} \leq V_{0}+\sum_{l=0}^{k} W_{l} \leq V_{0}+\bar{W}
$$

from which we can

- conclude that $R_{k} \rightarrow 0$ as $k \rightarrow \infty$ since $R_{k} \geq 0$
- derive sublinear rates of convergence for $R_{k}$ towards 0


## Concluding sublinear convergence

- Lyapunov inequality consequence restated

$$
\sum_{l=0}^{k} R_{l} \leq V_{0}+\sum_{l=0}^{k} W_{l} \leq V_{0}+\bar{W}
$$

- We can derive sublinear convergence for
- Best $R_{k}:(k+1) \min _{l \in\{0, \ldots, k\}} R_{l} \leq \sum_{l=0}^{k} R_{l}$
- Last $R_{k}$ (if $R_{k}$ decreasing): $(k+1) R_{k} \leq \sum_{l=0}^{k} R_{l}$
- Average $R_{k}: \bar{R}_{k}=\frac{1}{k+1} \sum_{l=0}^{k} R_{l}$
- Let $\hat{R}_{k}$ be any of these quantities, and we have

$$
\hat{R}_{k} \leq \frac{\sum_{l=0}^{k} R_{l}}{k+1} \leq \frac{V_{0}+\bar{W}}{k+1}
$$

which shows a $O(1 / k)$ sublinear convergence

Deriving other than $O(1 / k)$ convergence (1/3)

- Other rates can be derived from a modified Lyapunov inequality:

$$
V_{k+1} \leq V_{k}+W_{k}-\lambda_{k} R_{k}
$$

with $\lambda_{k}>0$ when we are interested in convergence of $R_{k}$, then

$$
\sum_{l=0}^{k} \lambda_{l} R_{l} \leq V_{0}+\sum_{l=0}^{k} W_{l} \leq V_{0}+\bar{W}
$$

- We have $R_{k} \rightarrow 0$ as $k \rightarrow \infty$ if, e.g., $\sum_{l=0}^{\infty} \lambda_{l}=\infty$

Deriving other than $O(1 / k)$ convergence (2/3)

- Restating the consequence: $\sum_{l=0}^{k} \lambda_{l} R_{l} \leq V_{0}+\bar{W}$
- We can derive sublinear convergence for
- Best $R_{k}: \min _{l \in\{0, \ldots, k\}} R_{l} \sum_{l=0}^{k} \lambda_{l} \leq \sum_{l=0}^{k} \lambda_{l} R_{l}$
- Last $R_{k}$ (if $R_{k}$ decreasing): $R_{k} \sum_{l=0}^{k} \lambda_{l} \leq \sum_{l=0}^{k} \lambda_{l} R_{l}$
- Weighted average $R_{k}: \bar{R}_{k}=\frac{1}{\sum_{l=0}^{k} \lambda_{l}} \sum_{l=0}^{k} \lambda_{l} R_{l}$
- Let $\hat{R}_{k}$ be any of these quantities, and we have

$$
\hat{R}_{k} \leq \frac{\sum_{l=0}^{k} R_{l}}{\sum_{l=0}^{k} \lambda_{l}} \leq \frac{V_{0}+\bar{W}}{\sum_{l=0}^{k} \lambda_{l}}
$$

Deriving other than $O(1 / k)$ convergence (3/3)

- How to get a rate out of:

$$
\hat{R}_{k} \leq \frac{V_{0}+\bar{W}}{\sum_{l=0}^{k} \lambda_{l}}
$$

- Assume $\psi(k) \leq \sum_{l=0}^{k} \lambda_{l}$, then $\psi(k)$ decides rate:

$$
\hat{R}_{k} \leq \frac{\sum_{l=0}^{k} R_{l}}{\sum_{l=0}^{k} \lambda_{l}} \leq \frac{V_{0}+\bar{W}}{\psi(k)}
$$

which gives a $O\left(\frac{1}{\psi(k)}\right)$ rate

- If $\lambda_{k}=c$ is constant: $\psi(k)=c(k+1)$ and we have $O(1 / k)$ rate
- If $\lambda_{k}$ is decreasing: slower rate than $O(1 / k)$
- If $\lambda_{k}$ is increasing: faster rate than $O(1 / k)$


## Estimating $\psi$ via integrals

- Assume that $\lambda_{k}=\phi(k)$, then $\psi(k) \leq \sum_{l=0}^{k} \phi(l)$ and

$$
\hat{R}_{k} \leq \frac{\sum_{l=0}^{k} R_{l}}{\sum_{l=0}^{k} \phi(l)} \leq \frac{V_{0}+\bar{W}}{\psi(k)}
$$

- To estimate $\psi$, we use the integral inequalities
- for decreasing nonnegative $\phi$ :

$$
\int_{t=0}^{k} \phi(t) d t+\phi(k) \leq \sum_{l=0}^{k} \phi(l) \leq \int_{t=0}^{k} \phi(t) d t+\phi(0)
$$

- for increasing nonnegative $\phi$ :

$$
\int_{t=0}^{k} \phi(t) d t+\phi(0) \leq \sum_{l=0}^{k} \phi(l) \leq \int_{t=0}^{k} \phi(t) d t+\phi(k)
$$

- Remove $\phi(k), \phi(0) \geq 0$ from the lower bounds and use estimate:

$$
\psi(k)=\int_{t=0}^{k} \phi(t) d t \leq \sum_{l=0}^{k} \phi(l)
$$

## Sublinear rate examples

- For Lyapunov inequality $V_{k+1} \leq V_{k}+W_{k}-\lambda_{k} R_{k}$, we get:
$\hat{R}_{k} \leq \frac{V_{0}+\bar{W}}{\psi(k)} \quad$ where $\quad \lambda_{k}=\phi(k)$ and $\psi(k)=\int_{t=0}^{k} \phi(t) d t$
- Let us quantify the rate $\psi$ in a few examples:
- Two examples that are slower than $O(1 / k)$ :
- $\lambda_{k}=\phi(k)=c /(k+1)$ gives slow $O\left(\frac{1}{\log k}\right)$ rate:

$$
\psi(k)=\int_{t=0}^{k} \frac{c}{t+1} d t=c[\log (t+1)]_{t=0}^{k}=c \log (k+1)
$$

- $\lambda_{k}=\phi(k)=c /(k+1)^{\alpha}$ for $\alpha \in(0,1)$, gives faster $O\left(\frac{1}{k^{1-\alpha}}\right)$ rate:

$$
\psi(k)=\int_{t=0}^{k} \frac{c}{(t+1)^{\alpha}} d t=c\left[\frac{(t+1)^{1-\alpha}}{(1-\alpha)}\right]_{t=0}^{k}=\frac{c}{1-\alpha}\left((k+1)^{1-\alpha}-1\right)
$$

- An example that is faster than $O(1 / k)$
- $\lambda_{k}=\phi(k)=c(k+1)$ gives $O\left(\frac{1}{k^{2}}\right)$ rate:

$$
\psi(k)=\int_{t=0}^{k} c(t+1) d t=c\left[\frac{1}{2}(t+1)^{2}\right]_{t=0}^{k}=\frac{c}{2}\left((k+1)^{2}-1\right)
$$

## Stochastic setting and law of total expectation

- In the stochastic setting, we analyze the stochastic process

$$
x_{k+1}=\mathcal{A}_{k}\left(\xi_{k}\right) x_{k}
$$

- We will look for inequalities of the form

$$
\mathbb{E}\left[V_{k+1} \mid x_{k}\right] \leq \mathbb{E}\left[V_{k} \mid x_{k}\right]+\mathbb{E}\left[W_{k} \mid x_{k}\right]-\lambda_{k} \mathbb{E}\left[R_{k} \mid x_{k}\right]
$$

to see what happens in one step given $x_{k}$ (but not given $\xi_{k}$ )

- We use law of total expectation $\mathbb{E}[\mathbb{E}[X \mid Y]]=\mathbb{E}[X]$ to get

$$
\mathbb{E}\left[V_{k+1}\right] \leq \mathbb{E}\left[V_{k}\right]+\mathbb{E}\left[W_{k}\right]-\lambda_{k} \mathbb{E}\left[R_{k}\right]
$$

which is a Lyapunov inequality

- We can draw rate conclusions, as we did before, now for $\mathbb{E}\left[R_{k}\right]$
- For realizations we can say:
- If $\mathbb{E}\left[R_{k}\right]$ is summable, then $R_{k} \rightarrow 0$ almost surely
- If $\mathbb{E}\left[R_{k}\right] \rightarrow 0$, then $R_{k} \rightarrow 0$ in probability


## Rates in stochastic setting

- Lyapunov inequality $\mathbb{E}\left[V_{k+1}\right] \leq \mathbb{E}\left[V_{k}\right]+\mathbb{E}\left[W_{k}\right]-\lambda_{k} \mathbb{E}\left[R_{k}\right]$ implies:

$$
\sum_{l=0}^{k} \lambda_{l} \mathbb{E}\left[R_{l}\right] \leq V_{0}+\sum_{l=0}^{\infty} \mathbb{E}\left[W_{l}\right] \leq V_{0}+\bar{W}
$$

- Same procedure as before gives sublinear rates for
- Best $\mathbb{E}\left[R_{k}\right]: \min _{l \in\{0, \ldots, k\}} \mathbb{E}\left[R_{l}\right] \sum_{l=0}^{k} \lambda_{l} \leq \sum_{l=0}^{k} \lambda_{l} \mathbb{E}\left[R_{l}\right]$
- Last $\mathbb{E}\left[R_{k}\right]$ (if $\mathbb{E}\left[R_{k}\right]$ decreasing): $\mathbb{E}\left[R_{k}\right] \sum_{l=0}^{k} \lambda_{l} \leq \sum_{l=0}^{k} \lambda_{l} \mathbb{E}\left[R_{l}\right]$
- Weighted average: $\mathbb{E}\left[\bar{R}_{k}\right]=\frac{1}{\sum_{l=0}^{k} \lambda_{l}} \sum_{l=0}^{k} \lambda_{l} \mathbb{E}\left[R_{l}\right]$
- Jensen's inequality for concave $\min _{l}$ in best residual reads

$$
\mathbb{E}\left[\min _{l \in\{0, \ldots, k\}} R_{l}\right] \leq \min _{l \in\{0, \ldots, k\}} \mathbb{E}\left[R_{l}\right]
$$

- Let $\hat{R}_{k}$ be any of the above quantities, and we have

$$
\mathbb{E}\left[\hat{R}_{k}\right] \leq \frac{V_{0}+\bar{W}}{\sum_{l=0}^{k} \lambda_{l}}
$$

| 两 |  |
| :--- | :--- |



- The proof continues by using the equality

$$
\begin{aligned}
& \left(x_{k+1}-x_{k}\right)^{T}\left(z-x_{k+1}\right) \\
& \quad=\frac{1}{2}\left(\left\|x_{k}-z\right\|_{2}^{2}-\left\|x_{k+1}-z\right\|_{2}^{2}-\left\|x_{k+1}-x_{k}\right\|_{2}^{2}\right)
\end{aligned}
$$

- Applying to previous inequality gives

$$
\begin{aligned}
f\left(x_{k+1}\right)+ & g\left(x_{k+1}\right) \\
\leq & f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(z-x_{k}\right)+\frac{\beta_{k}}{2}\left\|x_{k+1}-x_{k}\right\|_{2}^{2}+g(z) \\
& +\gamma_{k}^{-1}\left(x_{k+1}-x_{k}\right)^{T}\left(z-x_{k+1}\right) \\
= & f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(z-x_{k}\right)+\frac{\beta_{k}}{2}\left\|x_{k+1}-x_{k}\right\|_{2}^{2}+g(z) \\
& +\frac{1}{2 \gamma_{k}}\left(\left\|x_{k}-z\right\|_{2}^{2}-\left\|x_{k+1}-z\right\|_{2}^{2}-\left\|x_{k}-x_{k+1}\right\|_{2}^{2}\right)
\end{aligned}
$$

which after rearrangement gives the fundamental inequality

- A fundamental inequality
- Nonconvex setting
- Convex setting
- Strongly convex setting
- Backtracking
- Stopping conditions
- Accelerated gradient method
- Scaling
- We will analyze the proximal gradient method

$$
x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)
$$

in a nonconvex setting for solving

$$
\operatorname{minimize} f(x)+g(x)
$$

- Will show sublinear $O(1 / k)$ convergence
- Analysis based on A fundamental inequality


## Nonconvex setting - Assumptions

(i) $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is continuously differentiable (not necessarily convex)
(ii) For every $x_{k}$ and $x_{k+1}$ there exists $\beta_{k} \in\left[\eta, \eta^{-1}\right], \eta \in(0,1]$ :
$f\left(x_{k+1}\right) \leq f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(x_{k+1}-x_{k}\right)+\frac{\beta_{k}}{2}\left\|x_{k}-x_{k+1}\right\|_{2}^{2}$
where $\beta_{k}$ is a sort of local Lipschitz constant
(iii) $g: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ is closed convex
(iv) A minimizer $x^{\star}$ exists and $p^{\star}=f\left(x^{\star}\right)+g\left(x^{\star}\right)$ is optimal value
(v) Algorithm parameters $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta_{k}}-\epsilon\right]$, where $\epsilon>0$

- Differs from assumptions for fundamental inequality only in $(v)$
- Assumption (ii) satisfied with $\beta_{k} \geq \beta$ if $f$ is $\beta$-smooth


## Nonconvex setting - Analysis

- Use fundamental inequality

$$
\begin{aligned}
f\left(x_{k+1}\right)+g\left(x_{k+1}\right) \leq & f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(z-x_{k}\right)-\frac{\gamma_{k}^{-1}-\beta_{k}}{2}\left\|x_{k+1}-x_{k}\right\|_{2}^{2} \\
& +g(z)+\frac{1}{2 \gamma_{k}}\left(\left\|x_{k}-z\right\|_{2}^{2}-\left\|x_{k+1}-z\right\|_{2}^{2}\right)
\end{aligned}
$$

- Set $z=x_{k}$ to get

$$
f\left(x_{k+1}\right)+g\left(x_{k+1}\right) \leq f\left(x_{k}\right)+g\left(x_{k}\right)-\left(\gamma_{k}^{-1}-\frac{\beta_{k}}{2}\right)\left\|x_{k+1}-x_{k}\right\|_{2}^{2}
$$

## Step-size requirements

- Step-sizes $\gamma_{k}$ should be restricted for inequality to be useful:
$f\left(x_{k+1}\right)+g\left(x_{k+1}\right) \leq f\left(x_{k}\right)+g\left(x_{k}\right)-\left(\gamma_{k}^{-1}-\frac{\beta_{k}}{2}\right)\left\|x_{k+1}-x_{k}\right\|_{2}^{2}$
- Requirements $\beta_{k} \in\left[\eta, \eta^{-1}\right]$ and $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta_{k}}-\epsilon\right]$ :
- upper bound $\gamma_{k} \leq \frac{2}{\beta_{k}}-\epsilon$ can be written as

$$
\gamma_{k} \leq \frac{2}{\beta_{k}+2 \delta_{k}} \quad \text { where } \quad \delta_{k}=\frac{\beta_{k} \epsilon}{2\left(\frac{2}{\beta_{k}}-\epsilon\right)} \geq \frac{\beta_{k}^{2} \epsilon}{4} \geq \frac{\eta^{2} \epsilon}{4}>0
$$ since upper bound $\beta_{k} \leq \eta^{-1}$ gives $\frac{2}{\beta_{k}}-\epsilon \geq 2 \eta-\epsilon>0$ and $\epsilon>0$

- Inverting upper step-size bound and letting $\delta:=\frac{\eta^{2} \epsilon}{4} \leq \delta_{k}$ :

$$
\gamma_{k}^{-1} \geq \frac{\beta_{k}+2 \delta_{k}}{2} \geq \frac{\beta_{k}}{2}+\delta \quad \Rightarrow \quad \gamma_{k}^{-1}-\frac{\beta_{k}}{2} \geq \delta>0
$$

- This implies, by subtracting $p^{\star}$ from both sides to have $V_{k} \geq 0$,
$\underbrace{f\left(x_{k+1}\right)+g\left(x_{k+1}\right)-p^{\star}}_{V_{k+1}} \leq \underbrace{f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star}}_{V_{k}}-\underbrace{\delta\left\|x_{k+1}-x_{k}\right\|_{2}^{2}}_{R_{k}}$
where bounds on $\gamma_{k}$ imply that all $R_{k}$ are nonnegative


## Lyapunov inequality consequences

- Restating Lyapunov inequality
$\underbrace{f\left(x_{k+1}\right)+g\left(x_{k+1}\right)-p^{\star}}_{V_{k+1}} \leq \underbrace{f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star}}_{V_{k}}-\underbrace{\delta\left\|x_{k+1}-x_{k}\right\|_{2}^{2}}_{R_{k}}$
- Consequences
- Function value is decreasing sequence (may not converge to $p^{\star}$ )
- Fixed-point residual converges to 0 as $k \rightarrow \infty$

$$
\left\|x_{k+1}-x_{k}\right\|_{2}=\left\|\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)-x_{k}\right\|_{2} \rightarrow 0
$$

- Best fixed-point residual norm square converges as $O(1 / k)$ :

$$
\min _{i \in\{0, \ldots, k\}}\left\|x_{i+1}-x_{i}\right\|_{2}^{2} \leq \frac{f\left(x_{0}\right)+g\left(x_{0}\right)-p^{\star}}{\delta(k+1)}
$$

Lyapunov inequality consequences $-g=0$

- For $g=0$, then $x_{k+1}=x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)$ and

$$
\left\|x_{k+1}-x_{k}\right\|_{2}=\gamma_{k}\left\|\nabla f\left(x_{k}\right)\right\|_{2} \quad \text { and } \quad R_{k}=\delta \gamma_{k}^{2}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}
$$

- Lyapunov inequality consequences in this setting:
- Gradient converges to 0 (since $\gamma_{k} \geq \epsilon$ ): $\left\|\nabla f\left(x_{k}\right)\right\|_{2} \rightarrow 0$
- Smallest gradient norm square converges as:

$$
\min _{i \in\{0, \ldots, k\}}\left\|\nabla f\left(x_{i}\right)\right\|_{2}^{2} \leq \frac{f\left(x_{0}\right)-p^{\star}}{\delta \sum_{i=0}^{k} \gamma_{i}^{2}}
$$

- If, in addition, $f$ is $\beta$-smooth and $\gamma_{k}=\frac{1}{\beta}$.

$$
\min _{i \in\{0, \ldots, k\}}\left\|\nabla f\left(x_{i}\right)\right\|_{2}^{2} \leq \frac{2 \beta\left(f\left(x_{0}\right)-p^{\star}\right)}{k+1}
$$

since then $\beta_{k}=\beta$ and $\gamma_{k}^{-1}-\frac{\beta_{k}}{2}=\frac{\beta}{2}=\delta>0$

- So, will approach local maximum, minimum, or saddle-point





## Scaling invariant stopping condition

- For $\beta$-smooth $f$, use scaled condition $\frac{1}{\beta} u_{k}$

$$
\frac{1}{\beta} u_{k}:=\frac{1}{\beta}\left(\gamma_{k}^{-1}\left(x_{k}-x_{k+1}\right)+\nabla f\left(x_{k+1}\right)-\nabla f\left(x_{k}\right)\right)
$$

that we have seen before

- Let us scale problem by $a$ to get minimize $a f(x)+a g(x)$, then
- smoothness constant $\beta_{a}=a \beta$ scaled by $a \Rightarrow$ use $\gamma_{a, k}=\frac{\gamma_{k}}{a}$
- optimality measure $\frac{1}{\beta_{a}} u_{a, k}=\frac{1}{a \beta} a u_{k}=\frac{1}{\beta} u_{k}$ remains the same so it is scaling invariant
- Problem considered solved to optimality if, say, $\frac{1}{\beta}\left\|u_{k}\right\|_{2} \leq 10^{-6}$
- Often lower accuracy $10^{-3}$ to $10^{-4}$ is enough


## Example - SVM

- Classification problem from SVM lecture, SVM with
- polynomial features of degree 2
- regularization parameter $\lambda=0.00001$



## Example - Optimality measure

- Plots $\beta^{-1}\left\|u_{k}\right\|_{2}=\beta^{-1}\left\|\gamma_{k}^{-1}\left(x_{k}-x_{k+1}\right)+\nabla f\left(x_{k+1}\right)-\nabla f\left(x_{k}\right)\right\|_{2}$
- Shows $\beta^{-1}\left\|u_{k}\right\|_{2}$ up to $20^{\prime} 000$ iterations
- Quite many iterations needed to converge


Example - SVM higher degree polynomial

- Classification problem from SVM lecture, SVM with
- polynomial features of degree 6
- regularization parameter $\lambda=0.00001$



## Example - Optimality measure

- Plots $\beta^{-1}\left\|u_{k}\right\|_{2}=\beta^{-1}\left\|\gamma_{k}^{-1}\left(x_{k}-x_{k+1}\right)+\nabla f\left(x_{k+1}\right)-\nabla f\left(x_{k}\right)\right\|_{2}$
- Shows $\beta^{-1}\left\|u_{k}\right\|_{2}$ up to $200^{\prime} 000$ iterations (10x more than before)
- Many iterations needed for high accuracy


Outline

- A fundamental inequality
- Nonconvex setting
- Convex setting
- Strongly convex setting
- Backtracking
- Stopping conditions
- Accelerated gradient method
- Scaling


## Accelerated proximal gradient method

- Consider convex composite problem

$$
\underset{x}{\operatorname{minimize}} f(x)+g(x)
$$

where

- $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is $\beta$-smooth and convex
- $g: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ is closed and convex
- Proximal gradient descent

$$
x_{k+1}=\operatorname{prox}_{\gamma g}\left(x_{k}-\gamma \nabla f\left(x_{k}\right)\right)
$$

achieves $O(1 / k)$ convergence rate in function value

- Accelerated proximal gradient method

$$
\begin{aligned}
y_{k} & =x_{k}+\theta_{k}\left(x_{k}-x_{k-1}\right) \\
x_{k+1} & =\operatorname{prox}_{\gamma g}\left(y_{k}-\gamma \nabla f\left(y_{k}\right)\right)
\end{aligned}
$$

(with specific $\theta_{k}$ ) achieves faster $O\left(1 / k^{2}\right)$ convergence rate

## Accelerated proximal gradient method - Parameters

- Accelerated proximal gradient method

$$
\begin{aligned}
y_{k} & =x_{k}+\theta_{k}\left(x_{k}-x_{k-1}\right) \\
x_{k+1} & =\operatorname{prox}_{\gamma g}\left(y_{k}-\gamma \nabla f\left(y_{k}\right)\right)
\end{aligned}
$$

- Step-sizes are restricted $\gamma \in\left(0, \frac{1}{\beta}\right]$
- The $\theta_{k}$ parameters can be chosen either as

$$
\theta_{k}=\frac{k-1}{k+2}
$$

or $\theta_{k}=\frac{t_{k-1}-1}{t_{k}}$ where

$$
t_{k}=\frac{1+\sqrt{1+4 t_{k-1}^{2}}}{2}
$$

these choices are very similar

- Algorithm behavior in nonconvex setting not well understood



What is $\|\cdot\|_{H}$ ?

- Requirement: $H \in \mathbb{R}^{n \times n}$ is symmetric positive definite $(H \succ 0)$
- The norm $\|x\|_{H}^{2}:=x^{T} H x$, for $H=I$, we get $\|x\|_{I}^{2}=\|x\|_{2}^{2}$


## Smoothness

- Function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is $\beta$-smooth if for all $x, y \in \mathbb{R}^{n}$ :

$$
\begin{aligned}
& f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|x-y\|_{2}^{2} \\
& f(y) \geq f(x)+\nabla f(x)^{T}(y-x)-\frac{\beta}{2}\|x-y\|_{2}^{2}
\end{aligned}
$$

- We say $f \beta_{H}$-smoothness w.r.t. scaled norm $\|\cdot\|_{H}$ if

$$
f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta_{H}}{2}\|x-y\|_{H}^{2}
$$

$$
f(y) \geq f(x)+\nabla f(x)^{T}(y-x)-\frac{\beta_{H}}{2}\|x-y\|_{H}^{2}
$$

for all $x, y \in \mathbb{R}^{n}$

- If $f$ is smooth (w.r.t. $\|\cdot\|_{2}$ ) it is also smooth w.r.t. $\|\cdot\|_{H}$


## Example - A quadratic

- Let $f(x)=\frac{1}{2} x^{T} H x=\frac{1}{2}\|x\|_{H}^{2}$ with $H \succ 0$
- $f$ is 1 -smooth w.r.t $\|\cdot\|_{H}$ (with equality):

$$
\begin{aligned}
f(x) & +\nabla f(x)^{T}(y-x)+\frac{1}{2}\|x-y\|_{H}^{2} \\
& =\frac{1}{2} x^{T} H x+(H x)^{T}(y-x)+\frac{1}{2}\|x-y\|_{H}^{2} \\
& =\frac{1}{2} x^{T} H x+(H x)^{T}(y-x)+\frac{1}{2}\left(\|x\|_{H}^{2}-2(H x)^{T} y+\|y\|_{H}^{2}\right) \\
& =\frac{1}{2}\|y\|_{H}^{2}=f(y)
\end{aligned}
$$

## Scaled proximal gradient for quadratics

- Let $f(x)=\frac{1}{2} x^{T} H x$ with $H \succ 0$, which is 1 -smooth w.r.t. $\|\cdot\|_{H}$
- Approximation with scaled norm $\|\cdot\|_{H}$ and $\gamma_{k}=1$ satisfies $\forall x_{k}$ :

$$
\hat{f}_{x_{k}}(y)=f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(y-x_{k}\right)+\frac{1}{2}\left\|x_{k}-y\right\|_{H}^{2}=f(y)
$$

since $f$ is 1 -smooth w.r.t. $\|\cdot\|_{H}$ with equality

- An iteration then reduces to solving problem itself:

$$
x_{k+1}=\underset{y}{\operatorname{argmin}}\left(\hat{f}_{x_{k}}(y)+g(y)\right)=\underset{y}{\operatorname{argmin}}(f(y)+g(y))
$$

- Model very accurate, but very expensive iterations

$$
\begin{aligned}
f(y) & =f(x)+\nabla f(x)^{T}(y-x)+\frac{1}{2}\|x-y\|_{H}^{2} \\
& \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\lambda_{\max }(H)}{2}\|x-y\|_{2}^{2}
\end{aligned}
$$

much more conservative estimate of function!

## Scaled proximal gradient method reformulation

- Proximal gradient method with scaled norm $\|\cdot\|_{H}$ :

$$
\begin{aligned}
x_{k+1} & =\underset{y}{\operatorname{argmin}}\left(f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}(y-x)+\frac{1}{2 \gamma_{k}}\left\|y-x_{k}\right\|_{H}^{2}+g(y)\right) \\
& =\underset{y}{\operatorname{argmin}}\left(g(y)+\frac{1}{2 \gamma_{k}} \| y-\left(x_{k}-\gamma_{k} H^{-1} \nabla f\left(x_{k}\right) \|_{H}^{2}\right)\right. \\
& =: \operatorname{prox}_{\gamma_{k} g}^{H}\left(x_{k}-\gamma_{k} H^{-1} \nabla f\left(x_{k}\right)\right)
\end{aligned}
$$

where $H=I$ gives nominal method

- Computational difference per iteration:

1. Need to invert $H^{-1}$ (or solve $H d_{k}=\nabla f\left(x_{k}\right)$ )
2. Need to compute prox with new metric

$$
\operatorname{prox}_{\gamma_{k} g}^{H}(z):=\underset{x}{\operatorname{argmin}}\left(g(x)+\frac{1}{2 \gamma_{k}}\|x-z\|_{H}^{2}\right)
$$

that may be very costly

## Computational cost

- Assume that $H$ is dense or general sparse
- $H^{-1}$ dense: cubic complexity (vs maybe quadratic for gradient)
- $H^{-1}$ sparse: lower than cubic complexity
- $\operatorname{prox}_{\gamma_{k} g}^{H}$ : difficult optimization problem
- Assume that $H$ is diagonal
- $H^{-1}$ : invert diagonal elements - linear complexity
- $\operatorname{prox}_{\gamma_{k} g}^{H}$ : often as cheap as nominal prox (e.g., for separable $g$ )
- this gives individual step-sizes for each coordinate
- Assume that $H$ is block-diagonal with small blocks
- $H^{-1}$ : invert individual blocks - also cheap
- prox ${ }_{\gamma_{k} g}^{H}$ : often quite cheap (e.g., for block-separable $g$ )
- If $H=I$, method is nominal method


## Convergence

- We get similar results as in the nominal $H=I$ case
- We assume $\beta_{H}$ smoothness w.r.t. $\|\cdot\|_{H}$
- We can replace all $\|\cdot\|_{2}$ with $\|\cdot\|_{H}$ and $\nabla f$ with $H^{-1} \nabla f$ :
- Nonconvex setting with $\gamma_{k}=\frac{1}{\beta_{H}}$

$$
\min _{l \in\{0, \ldots, k\}}\left\|\nabla f\left(x_{l}\right)\right\|_{H^{-1}}^{2} \leq \frac{2 \beta_{H}\left(f\left(x_{0}\right)+g\left(x_{0}\right)-p^{\star}\right)}{k+1}
$$

- Convex setting with $\gamma_{k}=\frac{1}{\beta_{H}}$

$$
f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star} \leq \frac{\beta_{H}\left\|x_{0}-x^{\star}\right\|_{H}^{2}}{2(k+1)}
$$

- Strongly convex setting with $f \sigma_{H}$-strongly convex w.r.t. $\|\cdot\|_{H}$

$$
\left\|x_{k+1}-x^{\star}\right\|_{H} \leq \max \left(\beta_{H} \gamma-1,1-\sigma_{H} \gamma\right)\left\|x_{k}-x^{\star}\right\|_{H}
$$

## Example - Logistic regression

- Logistic regression with $\theta=(w, b)$ :

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} \log \left(1+e^{w^{T} \phi\left(x_{i}\right)+b}\right)-y_{i}\left(w^{T} \phi\left(x_{i}\right)+b\right)+\frac{\lambda}{2}\|w\|_{2}^{2}
$$

on the following data set (from logistic regression lecture)

- Polynomial features of degree 6 , Tikhonov regularization $\lambda=0.01$
- Number of decision variables: 28




| Mini-batch stochastic gradient <br> - Let $\mathcal{B}$ be set of $K$-sample mini-batches to choose from: <br> - Example: 2 -sample mini-batches and $N=4$ : $\mathcal{B}=\{\{1,2\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\}\}$ <br> - Number of mini batches $\binom{N}{K}$, each item in $\binom{N-1}{K-1}$ batches <br> - Let $\mathbb{B}$ be $\mathcal{B}$-valued random variable <br> - Let, as before, $\widehat{\nabla} f$ denote stochastic gradient estimator <br> - Realization: let $B$ be drawn from probability distribution of $\mathbb{B}$ $\widetilde{\nabla} f(x)=\frac{1}{K} \sum_{i \in B} \nabla f_{i}(x)$ <br> where we will use uniform probability distribution $p_{B}=p(\mathbb{B}=B)=\frac{1}{\binom{N}{K}}$ <br> - Stochastic gradient is unbiased: $\mathbb{E} \widehat{\nabla} f(x)=\frac{1}{\binom{N}{K}} \sum_{B \in \mathcal{B}} \frac{1}{K} \sum_{i \in B} \nabla f_{i}(x)=\frac{\binom{N-1}{K-1}}{\binom{N}{K} K} \sum_{i=1}^{N} \nabla f_{i}(x)=\frac{1}{N} \sum_{i=1}^{N} \nabla f_{i}(x)=\nabla f(x)$ | Stochastic gradient descent for finite sum problems <br> - The algorithm, choose $x_{0} \in \mathbb{R}^{n}$ and iterate: <br> 1. Sample a mini-batch $B_{k} \in \mathcal{B}$ of $K$ indices uniformly <br> 2. Update $x_{k+1}=x_{k}-\frac{\gamma_{k}}{K} \sum_{j \in B_{k}} \nabla f_{j}\left(x_{k}\right)$ <br> - Can have $\mathcal{B}=\{\{1\}, \ldots,\{N\}\}$ and sample only one function <br> - Gives realization of underlying stochastic process |
| :---: | :---: |
| Outline <br> - Stochastic gradient descent <br> - Convergence and distance to solution <br> - Convergence and solution norms <br> - Overparameterized vs underparameterized setting <br> - Escaping not individually flat minima <br> - SGD step-sizes <br> - SGD convergence | Qualitative convergence behavior <br> - Consider single-function batch setting <br> - Assume that the individual gradients satisfy $\left(\nabla f_{i}(x)\right)^{T}\left(\nabla f_{j}(x)\right) \geq \mu$ <br> for all $i, j$ and for some $\mu \in \mathbb{R}$ (i.e., can be positive or negative) <br> Will larger or smaller $\mu$ likely give better SGD convergence? Why? |
| Qualitative convergence behavior <br> - Consider single-function batch setting <br> - Assume that the individual gradients satisfy $\left(\nabla f_{i}(x)\right)^{T}\left(\nabla f_{j}(x)\right) \geq \mu$ <br> for all $i, j$ and for some $\mu \in \mathbb{R}$ (i.e., can be positive or negative) <br> Will larger or smaller $\mu$ likely give better SGD convergence? Why? <br> - Larger $\mu$ gives more similar to full gradient and faster convergence | Minibatch setting <br> - Larger minibatch gives larger $\mu$ and faster convergence <br> - Comes at the cost of higher per iteration count <br> - Limiting minibatch case is the gradient method <br> - Tradeoff in how large minibatches to use to optimize convergence <br> - Other reasons exist that favor small batches (later) |
| SGD - Example <br> - Let $c_{1}+c_{2}+c_{3}=0$ <br> - Solve minimize $x\left(\frac{1}{2}\left(\left\\|x-c_{1}\right\\|_{2}^{2}+\left\\|x-c_{2}\right\\|_{2}^{2}+\left\\|x-c_{3}\right\\|_{2}^{2}\right)\right)=\frac{3}{2}\\|x\\|_{2}^{2}+c$ <br> - How will trajectory look for SGD with $\gamma_{k}=1 / 3$ ? <br> Levelsets of summands <br> Levelset of sum | SGD - Example <br> - Let $c_{1}+c_{2}+c_{3}=0$ <br> - Solve minimize $x\left(\frac{1}{2}\left(\left\\|x-c_{1}\right\\|_{2}^{2}+\left\\|x-c_{2}\right\\|_{2}^{2}+\left\\|x-c_{3}\right\\|_{2}^{2}\right)\right)=\frac{3}{2}\\|x\\|_{2}^{2}+c$ <br> - How will trajectory look for SGD with $\gamma_{k}=1 / 3$ ? <br> Levelsets of summands <br> Levelset of sum |
| 14 | 14 |

## SGD - Example

- Let $c_{1}+c_{2}+c_{3}=0$
- Solve minimize $x\left(\frac{1}{2}\left(\left\|x-c_{1}\right\|_{2}^{2}+\left\|x-c_{2}\right\|_{2}^{2}+\left\|x-c_{3}\right\|_{2}^{2}\right)\right)=\frac{3}{2}\|x\|_{2}^{2}+c$
- How will trajectory look for SGD with $\gamma_{k}=1 / 3$ ?


Levelsets of summands


Levelset of sum

- Fast convergence outside "triangle" where gradients similar, slow inside
- Constant step SGD converges to noise ball


## SGD - Example zoomed out

- Same example but zoomed out
- Solve minimize $x\left(\frac{1}{2}\left(\left\|x-c_{1}\right\|_{2}^{2}+\left\|x-c_{2}\right\|_{2}^{2}+\left\|x-c_{3}\right\|_{2}^{2}\right)\right)=\frac{3}{2}\|x\|_{2}^{2}+c$
- How will trajectory look with $\gamma_{k}=1 / 3$ from more global view?



## SGD - Example

- Let $c_{1}+c_{2}+c_{3}=0$
- Solve minimize $x\left(\frac{1}{2}\left(\left\|x-c_{1}\right\|_{2}^{2}+\left\|x-c_{2}\right\|_{2}^{2}+\left\|x-c_{3}\right\|_{2}^{2}\right)\right)=\frac{3}{2}\|x\|_{2}^{2}+c$
- How will trajectory look for SGD with $\gamma_{k}=1 / 3$ ?

- Constant step GD converges (in this case straight to) solution (right)
- Difference is noise in stochastic gradient that can be measured by $\mu$ 14


## SGD - Example zoomed out

- Same example but zoomed out
- Solve minimize $\left(\frac{1}{2}\left(\left\|x-c_{1}\right\|_{2}^{2}+\left\|x-c_{2}\right\|_{2}^{2}+\left\|x-c_{3}\right\|_{2}^{2}\right)\right)=\frac{3}{2}\|x\|_{2}^{2}+c$
- How will trajectory look with $\gamma_{k}=1 / 3$ from more global view?

- Far form solution $\nabla f_{i}$ more similar to $\nabla f$, larger $\mu \Rightarrow$ faster convergence


## Drawback of diminishing step-size

- Diminishing step-size typically gives slow convergence
- Often better convergence with constant step (if it works)
- Fixed-step size converges to noise ball in general
- Need diminishing step-size to converge to solution in general
- Stochastic gradient descent
- Convergence and distance to solution
- Convergence and solution norms
- Overparameterized vs underparameterized setting
- Escaping not individually flat minima
- SGD step-sizes
- SGD convergence


## Fixed step-size SGD does not converge to solution

- We can at most hope for finding point $\bar{x}$ such that

$$
\nabla f(\bar{x})=0
$$

- Let $x_{k}=\bar{x}$, and assume $\nabla f_{i}\left(x_{k}\right) \neq 0$, then

$$
x_{k+1}=x_{k}-\gamma_{k} \nabla f_{i}\left(x_{k}\right) \neq x_{k}
$$

i.e., moves away from solution $\bar{x}$

- Only hope with fixed step-size if all $\nabla f_{i}(\bar{x})=0$, since for $x_{k}=\bar{x}$

$$
x_{k+1}=x_{k}-\gamma_{k} \nabla f_{i}\left(x_{k}\right)=x_{k}
$$

independent on $\gamma_{k}$ and algorithm stays at solution

- How does norm of individual gradients affect local convergence?

| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |
| :---: | :---: |
| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |
| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |
| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |


| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |
| :---: | :---: |
| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Large gradients at solution <br> - Individal gradients at solution $0: \nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |
| Example - Large gradients at solution <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.83, \nabla f_{2}(0)=-0.83$ <br> - SGD with $\gamma=0.07$ and cyclic update order: <br> $f\left(x_{10}\right)-f^{\star}=1.07$ <br> - Will not converge to solution with constant step-size | Example - Small gradients at solution <br> - Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem <br> - Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |
| Example - Small gradients at solution <br> - Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem <br> - Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$ <br> - SGD with $\gamma=0.07$ and cyclic update order: | Example - Small gradients at solution <br> - Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem <br> - Individal gradients at solution 0: $\nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$ <br> - SGD with $\gamma=0.07$ and cyclic update order: |

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:



## Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:


Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:



## Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:



## Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:



## Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:



## Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution 0: $\nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:



## Example - Small gradients at solution

- Shift $f_{1}$ and $f_{2}$ "outwards" to get new problem
- Individal gradients at solution $0: \nabla f_{1}(0)=0.02, \nabla f_{2}(0)=-0.02$
- SGD with $\gamma=0.07$ and cyclic update order:








## Kurdyka-Lojasiewicz

- Error bound is instance of the Kurdyka-Lojasiewicz (KL) property
- KL property has exponent $\alpha \in[0,1), \alpha=\frac{1}{2}$ gives error bound
- Examples of KL functions:
- Continuous (on closed domain) semialgebraic functions are KL

$$
\operatorname{graph} f=\cup_{i=1}^{r}\left(\cap_{j=1}^{q}\left\{x: h_{i j}(x)=0\right\} \cap_{l=1}^{p}\left\{x: g_{i l}(x)<0\right\}\right)
$$

graph is union of intersection, where $h_{i j}$ and $g_{i l}$ polynomials

- Continuous piece-wise polynomials (some DL training problems)
- Strongly convex functions
- Often difficult to decide KL-exponent
- Result: descent methods on KL functions converge
- sublinearly if $\alpha \in\left(\frac{1}{2}, 1\right)$
- linearly if $\alpha \in\left(0, \frac{1}{2}\right]$ (the error bound regime)


## Strongly convex functions satisfy error bound

- $s+\sigma x \in \partial f(x)$ with $s \in \partial g(x)$ for convex $g=f-\frac{\sigma}{2}\|\cdot\|_{2}^{2}$
- Therefore

$$
\begin{aligned}
\|s+\sigma x\|_{2}^{2} & =\|s\|_{2}^{2}+2 \sigma s^{T} x+\sigma^{2}\|x\|_{2}^{2} \\
& \geq\|s\|_{2}^{2}+2 \sigma s^{T} x^{\star}+2 \sigma\left(g(x)-g\left(x^{\star}\right)\right)+\sigma^{2}\|x\|_{2}^{2} \\
& =\|s\|_{2}^{2}+2 \sigma s^{T} x^{\star}+\sigma\left\|x^{\star}\right\|_{2}^{2}+2 \sigma\left(f(x)-f\left(x^{\star}\right)\right) \\
& =\left\|s+\sigma x^{\star}\right\|_{2}^{2}+2 \sigma\left(f(x)-f\left(x^{\star}\right)\right) \\
& \geq 2 \sigma\left(f(x)-f\left(x^{\star}\right)\right)
\end{aligned}
$$

where we used

- subgradient definition $g\left(x^{\star}\right) \geq g(x)+s^{T}\left(x^{\star}-x\right)$ in first inequality
- nonnegativity of norms in the second inequality


## Implications of error bound

- Restating error bound for differentiable case

$$
\delta\left(f(x)-f\left(x^{\star}\right)\right) \leq\|\nabla f(x)\|_{2}^{2}
$$

- Assume it holds for all $x$ in some ball $X$ around solution $x^{\star}$
- What can you say about local minima and saddle-points in $X$ ?


## Implications of error bound

- Restating error bound for differentiable case

$$
\delta\left(f(x)-f\left(x^{\star}\right)\right) \leq\|\nabla f(x)\|_{2}^{2}
$$

- Assume it holds for all $x$ in some ball $X$ around solution $x^{\star}$
- What can you say about local minima and saddle-points in $X$ ?
- There are none! Proof by contradiction:
- Assume local minima or saddle-point $\bar{x}$
- Then $\nabla f(\bar{x})=0 \Rightarrow f(\bar{x})=f\left(x^{\star}\right)$ and $\bar{x}$ is global minima


## Convergence analysis - Smoothness and error bound

- Convergence analysis of gradient method
- $\beta$-smoothness and error bound assumptions $\left(f^{\star}=f\left(x^{\star}\right)\right.$ )
$f\left(x_{k+1}\right)-f^{\star} \leq f\left(x_{k}\right)-f^{\star}+\nabla f\left(x_{k}\right)^{T}\left(x_{k+1}-x_{k}\right)+\frac{\beta}{2}\left\|x_{k}-x_{k+1}\right\|_{2}^{2}$ $=f\left(x_{k}\right)-f^{\star}-\gamma_{k}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}+\frac{\beta \gamma_{k}^{2}}{2}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}$
$=f\left(x_{k}\right)-f^{\star}-\gamma_{k}\left(1-\frac{\beta \gamma_{k}}{2}\right)\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}$
$\leq\left(1-\gamma_{k} \delta\left(1-\frac{\beta \gamma_{k}}{2}\right)\right)\left(f\left(x_{k}\right)-f^{\star}\right)$
where
- $\beta$-smoothness of $f$ is used in first inequality
- gradient update $x_{k+1}=x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)$ in first equality
- error bound is used in the final inequality
- Linear convergence in function values if $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta}-\epsilon\right], \epsilon>0$


## Semi-smoothness

- Typical DL training problems are not smooth
- E.g.: overparameterized ReLU networks with smooth loss
- But semi-smooth ${ }^{1}$ in neighborhood around random initialization ${ }^{2}$ :

$$
f(x) \leq f(y)+\nabla f(y)^{T}(x-y)+c\|x-y\|_{2} \sqrt{f(y)}+\frac{\beta}{2}\|x-y\|_{2}^{2}
$$

for some constants $c$ and $\beta$

- Holds locally for large enough $c, \beta$ if cont. piece-wise polynomial
- Constants and neighborhood quantified in $[1]^{2}$
- $c=0$ gives smoothness
- $c$ small gives close to smoothness but allows nondifferentiable


## Convergence - Error bound and semi-smoothness

- Convergence analysis of gradient descent method
- Assumptions: $(c, \beta)$-semi-smooth, $\delta$-error bound, $f^{\star}=0$ (w.l.o.g.)
- Parameters $c \leq \frac{\sqrt{\delta} \gamma \beta}{2}$ and $\gamma \in\left(0, \frac{1}{\beta}\right)$ :
$f\left(x_{k+1}\right)$
$\leq f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(x_{k+1}-x_{k}\right)+c\left\|x_{k+1}-x_{k}\right\| \sqrt{f\left(x_{k}\right)}+\frac{\beta}{2}\left\|x_{k+1}-x_{k}\right\|_{2}^{2}$
$=f\left(x_{k}\right)-\gamma\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}+c \gamma\left\|\nabla f\left(x_{k}\right)\right\| \sqrt{f\left(x_{k}\right)}+\frac{\beta \gamma^{2}}{2}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}$
$\leq f\left(x_{k}\right)-\gamma\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}+\frac{c \gamma}{\sqrt{\delta}}\left\|\nabla f\left(x_{k}\right)\right\|^{2}+\frac{\beta \gamma^{2}}{2}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}$
$\leq f\left(x_{k}\right)-\gamma\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}+\beta \gamma^{2}\left\|\nabla f\left(x_{k}\right)\right\|^{2}$
$\leq f\left(x_{k}\right)-\gamma(1-\beta \gamma)\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}$
$\leq(1-c \gamma(1-\beta \gamma)) f\left(x_{k}\right)$
which shows linear convergence to 0 loss
- Need the nonsmooth part of upper bound $c$ to be small enough
- Can analyze SGD in similar manner


## Convergence in deep learning

- Setting: ReLU network, fully connected, smooth loss
- $c$ is small enough when model overparameterized enough [1] ${ }^{1}$
- Linear convergence (with high prob.) for random initialization [1]
- In practice:
- $\beta$ will be big - relies on small enough ( $\leq \frac{1}{\beta}$ ) constant step-size
- need to find "correct" step-size by diminishing rule
- need to control steps to not depart from linear convergence region
- hopefully achieved by previous step-size rule






## Convergence to projection point

## Graphical interpretation

- The scaled gradient method can be written as

$$
H x_{k+1}=H x_{k}-\gamma_{k} A^{T}\left(A x_{k}-b\right)
$$

if all $\gamma_{k}>\epsilon>0$ are small enough, it converges to a solution $\bar{x}$ :

$$
x_{k} \rightarrow \bar{x} \quad \text { and } \quad A \bar{x}=b
$$

- Letting $\lambda_{k}=-\sum_{l=0}^{k} \gamma_{l}\left(A x_{l}-b\right) \in \mathbb{R}^{m}$ and unfolding iteration:

$$
H x_{k+1}-H x_{0}=-\sum_{l=0}^{k} \gamma_{l} A^{T}\left(A x_{l}-b\right)=A^{T} \lambda_{k} \in \mathcal{R}\left(A^{T}\right)
$$

- In the limit $x_{k} \rightarrow \bar{x}$, we get

$$
H \bar{x}-H x_{0} \in \mathcal{R}\left(A^{T}\right)
$$

which with $A \bar{x}=b$ gives optimality conditions for projection

- If $x_{0}=0$, the algorithm converges to $\underset{x \in X}{\operatorname{argmin}}\left(\|x\|_{H}\right)$

$$
x \in X
$$

- What happens with scaled gradient method?
- Solution set $X$ extends infinitely
- sequence is perpendicular to $X$ in scalar product $(H x)^{T} y$
- algorithm converges to projection point $\operatorname{argmin}_{x \in X}\left(\left\|x-x_{0}\right\|_{H}\right)$

> Gradient method


## SGD - Convergence to projection point

- Least squares problem on finite sum form

$$
\underset{x}{\operatorname{minimize}} \frac{1}{2}\|A x-b\|_{2}^{2}=\frac{1}{2} \sum_{i=1}^{m}\left(a_{i}^{T} x-b_{i}\right)^{2}
$$

where $A=\left[a_{1}, \ldots, a_{m}\right]^{T}$

- Applying single-batch scaled SGD:

$$
x_{k+1}=x_{k}-\gamma_{k} H^{-1} a_{i_{k}}\left(a_{i_{k}}^{T} x_{k}-b_{i_{k}}\right)
$$

- The iteration can be unfolded as

This analysis hints towards that SGD gives smaller norm solutions and better generalization than variable metric Adam. Is this true?
$H x_{k+1}-H x_{0}=-\sum_{l=0}^{k} a_{i_{l}} \gamma_{l}\left(a_{i_{l}}^{T} x_{l}-b_{i_{l}}\right)=A^{T}\left[\begin{array}{c}-\sum_{l=0}^{k} \chi_{i_{l}=1}^{\chi}\left(\gamma_{l}\left(a_{1}^{T} x_{l}-b_{1}\right)\right) \\ \vdots \\ -\sum_{l=0}^{k} \chi_{i_{l}=m}\left(\gamma_{l}\left(a_{m}^{T} x_{l}-b_{m}\right)\right)\end{array}\right]$
where $\underset{i_{l}=j}{\chi}(v)=v$ if $i_{l}=j$, else 0 , so $H x_{k+1}-H x_{0} \in \mathcal{R}\left(A^{T}\right)$

- Assume $x_{k} \rightarrow \bar{x}$ with $A \bar{x}=b \Rightarrow$ convergence to projection point ${ }_{23}$


## SGD vs Adam

## How about deep learning?

- The analysis does not carry over to nonconvex DL settings
- However, often convergence to similar norm as initial point


## Initialization in next example

- Set scaling of weights by $\sigma$
- For the residual layers (all square layers)
- $\left(W_{j}\right)_{i j} \sim \mathcal{N}(0,1)$, normalize $W_{j}$, scale by $\sigma$
- $\left(b_{j}\right)_{i} \sim \mathcal{N}(0,1)$, normalize $b_{j}$, scale by $\sigma$
- For the non-residual layers (non-square layers)
- $\left(W_{j}\right)_{i j} \sim \mathcal{N}(0,1)$, normalize $W_{j}$, scale by $\max (1, \sigma)$
- $\left(b_{j}\right)_{i} \sim \mathcal{N}(0,1)$, normalize $b_{j}$, scale by $\max (1, \sigma)$
- use $\max (1, \sigma)$ for gradient to not vanish in non-residual layers


## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma$ : 0.01 Algorithm: SGD

$$
\left\|\theta_{\text {end }}\right\|_{2}=9.9 \quad \operatorname{loss}\left(\theta_{\text {end }}\right)=0.051
$$

$$
{ }^{*} \delta_{*}^{*}{ }_{*}^{*} \quad \int_{*}^{*}{ }_{*}^{*}
$$

## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma: 0.1$ Algorithm: SGD



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma$ : 1 Algorithm: SGD



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma$ : 5 Algorithm: SGD



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma$ : 10 Algorithm: SGD



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, 15x25,2,1 (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma: 0.01$ Algorithm: Adam



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma$ : 0.1 Algorithm: Adam



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma: 1$ Algorithm: Adam



## Convergence from different initial point

- Classification, hinge loss, ReLU, residual, $15 \times 25,2,1$ (17 layers)
- $L_{m}$ is Lipschitz constant in $x$ of final model $m(x ; \theta)$
- Initialization scaling $\sigma$ : 5 Algorithm: Adam




## Adam vs SGD - Flat or sharp minima

- Data from previous classification example with $\sigma=10$
- Loss landscape around final point $\theta_{\text {end }}$ for SGD and Adam
- SGD and Adam reach 0 loss but Adam minimum much sharper
- Same $\theta_{1}, \theta_{2}$ directions, same axes, $z_{\max }=10^{9}$

SGD


Adam


34


## Supporting hyperplane theorem

## Connection to duality and subgradients

Let $S$ be a nonempty convex set and let $x \in \operatorname{bd}(S)$. Then there exists a supporting hyperplane to $S$ at $x$.

- Does not exist for all point on boundary for nonconvex sets
- Many supporting hyperplanes exist for points of nonsmoothness


Supporting hyperplanes are at the core of convex analysis:

- Subgradients define supporting hyperplanes to epif
- Conjugate functions define supporting hyperplanes to epif
- Duality is based on subgradients, hence supporting hyperplanes:
- Consider minimize $(f(x)+g(x))$ and primal solution $x^{\star}$
- Dual problem minimize $\mu\left(f^{*}(\mu)+g^{*}(-\mu)\right)$ solution $\mu^{\star}$ satisfies

$$
\mu^{\star} \in \partial f\left(x^{\star}\right) \quad-\mu^{\star} \in \partial g\left(x^{\star}\right)
$$

i..e, dual problem finds subgradients at optimal point ${ }^{1}$
${ }^{1}$ When solving $\min _{x}(f(L x)+g(x))$ dual problem finds $\mu$ such that $L^{T} \mu \in \partial(f \circ L)(x)$ and $-L^{T} \mu \in \partial g(x)$. 10

Epigraphs and convexity

- Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$
- Then $f$ is convex if and only epi $f$ is a convex set in $\mathbb{R}^{n} \times \mathbb{R}$

- $f$ is called closed (lower semi-continuous) if epi $f$ is closed set
(in extended valued arithmetics)
- A function $f$ is concave if $-f$ is convex


## First-order condition for convexity

- A differentiable function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is convex if and only if

$$
f(y) \geq f(x)+\nabla f(x)^{T}(y-x)
$$

for all $x, y \in \mathbb{R}^{n}$


- Function $f$ has for all $x \in \mathbb{R}^{n}$ an affine minorizer that:
- has slope $s$ defined by $\nabla f$
- coincides with function $f$ at $x$
- is supporting hyperplane to epigraph of $f$
- defines normal $(\nabla f(x),-1)$ to epigraph of $f$


## Subdifferentials and subgradients

- Subgradients $s$ define affine minorizers to the function that:

- coincide with $f$ at $x$
- define normal vector $(s,-1)$ to epigraph of $f$
- can be one of many affine minorizers at nondifferentiable points $x$
- Subdifferential of $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ at $x$ is set of vectors $s$ satisfying

$$
\begin{equation*}
f(y) \geq f(x)+s^{T}(y-x) \quad \text { for all } y \in \mathbb{R}^{n}, \tag{1}
\end{equation*}
$$

- Notation:
- subdifferential: $\partial f: \mathbb{R}^{n} \rightarrow 2^{\mathbb{R}^{n}}$ (power-set notation $2^{\mathbb{R}^{n}}$ )
- subdifferential at $x: \partial f(x)=\{s:(1)$ holds $\}$
- elements $s \in \partial f(x)$ are called subgradients of $f$ at $x$


## Subgradient existence - Nonconvex example

- Function can be differentiable at $x$ but $\partial f(x)=\emptyset$

- $x_{1}: \partial f\left(x_{1}\right)=\{0\}, \nabla f\left(x_{1}\right)=0$
- $x_{2}: \partial f\left(x_{2}\right)=\emptyset, \nabla f\left(x_{2}\right)=0$
- $x_{3}: \partial f\left(x_{3}\right)=\emptyset, \nabla f\left(x_{3}\right)=0$
- Gradient is a local concept, subdifferential is a global property


## Existence for extended-valued convex functions

- Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ be convex, then:

1. Subgradients exist for all $x$ in relative interior of $\operatorname{dom} f$

Subgradients sometimes exist for $x$ on boundary of $\operatorname{dom} f$
3. No subgradient exists for $x$ outside $\operatorname{dom} f$

- Examples for second case, boundary points of $\operatorname{dom} f$ :


- No subgradient (affine minorizer) exists for left function at $x=1$




## First-order condition for strong convexity

- Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be differentiable
- $f$ is $\sigma$-strongly convex with $\sigma>0$ if and only if

$$
f(y) \geq f(x)+\nabla f(x)^{T}(y-x)+\frac{\sigma}{2}\|x-y\|_{2}^{2}
$$

for all $x, y \in \mathbb{R}^{n}$


- Function $f$ has for all $x \in \mathbb{R}^{n}$ a quadratic minorizer that:
- has curvature defined by $\sigma$
- coincides with function $f$ at $x$
- defines normal $(\nabla f(x),-1)$ to epigraph of $f$


## First-order condition for smoothness

- $f$ is $\beta$-smooth with $\beta \geq 0$ if and only if
$f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|x-y\|_{2}^{2}$
$f(y) \geq f(x)+\nabla f(x)^{T}(y-x)-\frac{\beta}{2}\|x-y\|_{2}^{2}$
for all $x, y \in \mathbb{R}^{n}$

- Quadratic upper/lower bounds with curvatures defined by $\beta$
- Quadratic bounds coincide with function $f$ at $x$


## Smoothness

- A function is called $\beta$-smooth if its gradient is $\beta$-Lipschitz:

$$
\|\nabla f(x)-\nabla f(y)\|_{2} \leq \beta\|x-y\|_{2}
$$

for all $x, y \in \mathbb{R}^{n}$ (it is not necessarily convex)

- Alternative equivalent definition of $\beta$-smoothness

$$
f(\theta x+(1-\theta) y) \geq \theta f(x)+(1-\theta) f(y)-\frac{\beta}{2} \theta(1-\theta)\|x-y\|^{2}
$$

$$
f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) f(y)+\frac{\beta}{2} \theta(1-\theta)\|x-y\|^{2}
$$

hold for every $x, y \in \mathbb{R}^{n}$ and $\theta \in[0,1]$

- Smoothness does not imply convexity
- Example:



## First-order condition for smooth convex

- $f$ is $\beta$-smooth with $\beta \geq 0$ and convex if and only if

$$
\begin{aligned}
& f(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|x-y\|_{2}^{2} \\
& f(y) \geq f(x)+\nabla f(x)^{T}(y-x)
\end{aligned}
$$

for all $x, y \in \mathbb{R}^{n}$


- Quadratic upper bound and affine lower bound
- Bounds coincide with function $f$ at $x$
- Quadratic upper bound is called descent lemma


## Duality correspondance

Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$. Then the following are equivalent
(i) $f$ is closed and $\sigma$-strongly convex
(ii) $\partial f$ is maximally monotone and $\sigma$-strongly monotone
(iii) $\nabla f^{*}$ is $\sigma$-cocoercive
(iv) $\nabla f^{*}$ is maximally monotone and $\frac{1}{\sigma}$-Lipschitz continuous
(v) $f^{*}$ is closed convex and satisfies descent lemma (is $\frac{1}{\sigma}$-smooth)
where $\nabla f^{*}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ and $f^{*}: \mathbb{R}^{n} \rightarrow \mathbb{R}$
Comments:

- Relation (i) $\Leftrightarrow$ (v) most important for us
- Since $f=f^{* *}$ the result holds with $f$ and $f^{*}$ interchanged
- Full proof available on course webpage


## Composite Optimization

## Composite optimization

We consider composite optimization problems of the form

$$
\underset{x}{\operatorname{minimize}} f(L x)+g(x)
$$

## Optimality conditions and dual problem

- Assume $f, g$ closed convex and that CQ holds
- Problem minimize $(f(L x)+g(x))$ is solved by $x$ iff

$$
0 \in L^{T} \underbrace{\partial f(L x)}_{\mu}+\partial g(x)
$$

where dual variable $\mu$ has been defined

- Primal dual necessary and sufficient optimality conditions:

$$
\begin{aligned}
& \left\{\begin{array}{l}
\mu \in \partial f(L x) \\
-L^{T} \mu \in \partial g(x)
\end{array}\right. \\
& \left\{\begin{array}{l}
\mu \in \partial f(L x) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right. \\
& \begin{array}{l}
L x \in \partial f^{*}(\mu) \\
-L^{*} \mu \in \partial g(x)
\end{array} \\
& \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right.
\end{aligned}
$$

- Dual optimality condition

$$
\begin{equation*}
0 \in \partial f^{*}(\mu)+\partial\left(g^{*} \circ-L^{T}\right)(\mu) \tag{1}
\end{equation*}
$$

solves dual problem minimize ${ }_{\mu} f^{*}(\mu)+g^{*}\left(-L^{T} \mu\right)$

- If CQ-D holds, all dual problem solutions satisfy (1)
- Dual searches for $\mu$ such that $L^{T} \mu \in \partial f(x)$ and $-L^{T} \mu \in \partial g(x)$
- Why solve dual? Sometimes easier to solve than primal
- Only interesting if primal solution can be recovered
- Assume $f, g$ closed convex and CQ
- Assume optimal dual $\mu$ known: $0 \in \partial f^{*}(\mu)+\partial\left(g^{*} \circ-L^{T}\right)(\mu)$
- Optimal primal $x$ must satisfy any and all primal-dual conditions

$$
\begin{array}{ll}
\left\{\begin{array}{l}
\mu \in \partial f(L x) \\
-L^{T} \mu \in \partial g(x)
\end{array}\right. & \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
-L^{T} \mu \in \partial g(x)
\end{array}\right. \\
\left\{\begin{array}{l}
\mu \in \partial f(L x) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right. & \left\{\begin{array}{l}
L x \in \partial f^{*}(\mu) \\
x \in \partial g^{*}\left(-L^{T} \mu\right)
\end{array}\right.
\end{array}
$$

- If one of these uniquely characterizes $x$, then must be solution:
- $\partial g^{*}$ is differentiable at $-L^{T} \mu$ for dual solution $\mu$
- $\partial f^{*}$ is differentiable at dual solution $\mu$ and $L$ invertible
- ...


## Algorithms

## Proximal gradient method

- Consider minimize $f(x)+g(x)$ where
- $f$ is $\beta$-smooth $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ (not necessarily convex)
- $g$ is closed convex
- Due to $\beta$-smoothness of $f$, we have

$$
f(y)+g(y) \leq f(x)+\nabla f(x)^{T}(y-x)+\frac{\beta}{2}\|y-x\|_{2}^{2}+g(y)
$$

for all $x, y \in \mathbb{R}^{n}$, i.e., r.h.s. is majorizing function for fixed $x$

- Majorization minimization with majorizer if $\gamma_{k} \in\left[\epsilon, \beta^{-1}\right], \epsilon>0$ :


## Proximal gradient - Fixed-points

- Denote $T_{\mathrm{PG}}^{\gamma}:=\operatorname{prox}_{\gamma g}(I-\gamma \nabla f)$, gives algorithm $x_{k+1}=T_{\mathrm{PG}}^{\gamma} x_{k}$
- Proximal gradient fixed-point set definition

$$
\operatorname{fix}_{\mathrm{PG}}^{\gamma}=\left\{x: x=T_{\mathrm{PG}}^{\gamma} x\right\}=\left\{x: x=\operatorname{prox}_{\gamma g}(x-\gamma \nabla f(x))\right\}
$$

i.e., set of points for which $x_{k+1}=x_{k}$

Let $\gamma>0$. Then $\bar{x} \in \operatorname{fix} T_{\mathrm{PG}}^{\gamma}$ if and only if $0 \in \partial g(\bar{x})+\nabla f(\bar{x})$.

$$
\begin{aligned}
x_{k+1} & =\underset{y}{\operatorname{argmin}}\left(f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}(y-x)+\frac{1}{2 \gamma_{k}}\left\|y-x_{k}\right\|_{2}^{2}+g(y)\right) \\
& =\underset{y}{\operatorname{argmin}}\left(g(y)+\frac{1}{2 \gamma_{k}}\left\|y-\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)\right\|_{2}^{2}\right) \\
& =\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)
\end{aligned}
$$

gives proximal gradient method

- Consequence: fixed-point set same for all $\gamma>0$
- We call inclusion $0 \in \partial g(\bar{x})+\nabla f(\bar{x})$ fixed-point characterization
- For convex problems: global solutions
- For nonconvex problems: critical points


## Applying proximal gradient to primal problems

Problem minimize $f(x)+g(x)$ :

- Assumptions:
- $f \beta$-smooth
- $g$ closed convex and prox friendly ${ }^{1}$
- $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta}-\epsilon\right]$
- Algorithm: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} \nabla f\left(x_{k}\right)\right)$

Problem minimize $f(L x)+g(x)$ :

- Assumptions:
- $f \beta$-smooth (implies $f \circ L \beta\|L\|_{2}^{2}$-smooth)
- $g$ closed convex and prox friendly
- $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta\|L\|_{2}^{2}}-\epsilon\right]$
- Gradient $\nabla(f \circ L)(x)=L^{T} \nabla f(L x)$
- Algorithm: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} L^{T} \nabla f\left(L x_{k}\right)\right)$
${ }^{1}$ Prox friendly: proximal operator cheap to evaluate, e.g., $g$ separable


## Applying proximal gradient to dual problem

Dual problem $\underset{\nu}{\operatorname{minimize}} f^{*}(\nu)+g^{*}\left(-L^{T} \nu\right)$

- Assumptions:
- $f$ closed convex and prox friendly
- $g \sigma$-strongly convex (which implies $g^{*} \circ-L^{T} \frac{\|L\|_{2}^{2}}{\sigma}$-smooth)
- $\gamma_{k} \in\left[\epsilon, \frac{2 \sigma}{\|L\|_{2}^{2}}-\epsilon\right]$
- Gradient: $\nabla\left(g^{*} \circ-L^{T}\right)(\nu)=-L \nabla g^{*}\left(-L^{T} \nu\right)$
- $\operatorname{Prox}($ Moreau $): \operatorname{prox}_{\gamma_{k} f^{*}}(\nu)=\nu-\gamma_{k} \operatorname{prox}_{\gamma_{k}^{-1} f}\left(\gamma_{k}^{-1} \nu\right)$
- Algorithm:

$$
\begin{aligned}
\nu_{k+1} & =\operatorname{prox}_{\gamma_{k} f^{*}}\left(\nu_{k}-\gamma_{k} \nabla\left(g^{*} \circ-L^{T}\right)\left(\nu_{k}\right)\right) \\
& =\left(I-\gamma_{k} \operatorname{prox}_{\gamma_{k}^{-1} f}\left(\gamma_{k}^{-1} \circ I\right)\right)\left(\nu_{k}+\gamma_{k} L \nabla g^{*}\left(-L^{T} \nu_{k}\right)\right)
\end{aligned}
$$

- Problem must be convex to have dual!
- Enough to know prox of $f$

What problems cannot be solved (efficiently)?
Problem minimize $f(x)+g(x)$

- Assumptions: $f$ and $g$ convex and nonsmooth
- No term differentiable, another method must be used
- Subgradient method
- Douglas-Rachford splitting
- Primal-dual methods

Problem minimize $f(x)+g(L x)$

- Assumptions:
- $f$ smooth
- $g$ nonsmooth convex
- $L$ arbitrary structured matrix
- Can apply proximal gradient method, but

$$
\left.\operatorname{prox}_{\gamma_{k}(g \circ L)}(z)=\underset{x}{\operatorname{argmin}} g(L x)+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right)
$$

often not "prox friendly", i.e., it is expensive to evaluate

## Training problems

- Training problem format

where $f$ is data misfit term and $g$ is regularizer
- Regularizers $(\theta=(w, b))$
- Tikhonov $g(\theta)=\|w\|_{2}^{2}$ is prox-friendly
- Sparsity inducing 1-norm $g(\theta)=\|w\|_{1}$ is prox-friendly
- Data misfit terms (with $m(x ; \theta)=\phi(x)^{T} \theta$ for convex problems)
- Least squares $L(u, y)=\|u-y\|_{2}^{2}$ smooth, hence $f$ smooth
- Logistic $L(u, y)=\log \left(1+e^{u}\right)-y u$ smooth, hence $f$ smooth
- SVM $L(u, y)=\max (0,1-y u)$ not smooth, hence $f$ not smooth
- Proximal gradient method
- Least squares: can efficiently solve primal
- Logistic regression: can solve primal
- SVM: add strongly convex regularization and solve dual
- Strongly convex regulariztion to have one conjugate smooth
- If bias term not regularized, only strongly convex in $w$
- SVM with $\|\cdot\|_{1}$-regularization not solvable with prox-grad


## Dual training problem

- Convex training problem

$$
\underset{\theta}{\operatorname{minimize}} \underbrace{\sum_{i=1}^{N} L\left(\phi\left(x_{i}\right)^{T} \theta, y_{i}\right)}_{f(X \theta)}+\underbrace{\sum_{j=1}^{n} g_{j}\left(\theta_{j}\right)}_{g(\theta)}
$$

has dual

$$
\underset{\theta}{\operatorname{minimize}} \underbrace{\sum_{i=1}^{N} L^{*}\left(\mu_{i}\right)}_{f^{*}(\mu)}+\underbrace{\sum_{j=1}^{n} g_{j}^{*}\left(\left(-X^{T} \mu\right)_{j}\right)}_{g^{*}\left(-X^{T} \mu\right)}
$$

where the conjugate of $L$ is w.r.t. first argument

- Dual has same structure as primal, finite-sum plus separable


## Exploiting structure

- Common structure, finite sum plus separable

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} f_{i}\left((X \theta)_{i}\right)+\sum_{j=1}^{n} \psi_{j}\left(\theta_{j}\right)
$$

- Stochastic gradient descent exploits finite-sum structure:
- Computes stochastic gradient of smooth part $f$
- Pick summand $f_{i}$ at random and perform gradient step
- Primal formulations: Pick training example and compute gradient
- Deep learning: evaluted via backpropagation
- Coordinate gradient descent exploits separable structure:
- Coordinate-wise updates if nonsmooth $\phi_{j}$ separable
- Requires efficient coordinate-wise evaluations of $\nabla f$


## Training problem structure

- Primal training problem

$$
\underset{\theta}{\operatorname{minimize}} \underbrace{\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)}_{f(X \theta)}+\underbrace{\sum_{j=1}^{n} g_{j}\left(\theta_{j}\right)}_{g(\theta)}
$$

- Dual training problem

$$
\underset{\theta}{\operatorname{minimize}} \underbrace{\sum_{i=1}^{N} L^{*}\left(\mu_{i}\right)}_{f^{*}(\mu)}+\underbrace{\sum_{j=1}^{n} g_{j}^{*}\left(\left(-X^{T} \mu\right)_{j}\right)}_{g^{*}\left(-X^{T} \mu\right)}
$$

- Common structure, finite sum plus separable:

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} f_{i}\left((X \theta)_{i}\right)+\sum_{j=1}^{n} \psi_{j}\left(\theta_{j}\right)
$$

- Primal: $f_{i}=L\left(m\left(x_{i} ; \cdot\right), y_{i}\right)$ (one summand per training example)
- Dual: $f_{i}=g_{j}^{*}\left(\left(-X^{T} \cdot\right)_{j}\right), \psi_{j}=L^{*}$


[^0]:    ${ }^{1}$ Prox friendly: proximal operator cheap to evaluate, e.g., $g$ separable

