## Recap

Pontus Giselsson

## Outline

- Convex analysis
- Composite optimization and duality
- Solving composite optimization problems - Algorithms


## Convex Analysis

## Convex sets

- A set $C$ is convex if for every $x, y \in C$ and $\theta \in[0,1]$ :

$$
\theta x+(1-\theta) y \in C
$$

- "Every line segment that connect any two points in $C$ is in $C$ "


Nonconvex


Nonconvex


Convex


Nonconvex

- Will assume that all sets are nonempty and closed


## Separating hyperplane theorem

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


Example


Counter-example $R$ nonconvex

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

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

- The hyperplane $\left\{x: s^{T} x=r\right\}$ is called separating hyperplane


## A strictly separating hyperplane theorem

- Suppose that $R, S \subseteq \mathbb{R}^{n}$ are non-intersecting closed and convex sets and that one of them is compact (closed and bounded)
- Then there exists hyperplane with strict separation


Example


Counter example $R, S$ not compact

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

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

## Consequence $-S$ is intersection of halfspaces

a closed convex set $S$ is the intersection of all halfspaces that contain it
proof:

- let $H$ be the intersection of all halfspaces containing $S$
- $\Rightarrow$ : obviously $x \in S \Rightarrow x \in H$
- $\Leftarrow$ : assume $x \notin S$, since $S$ closed and convex and $x$ compact (a point), 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 $S$ at $x$
- Definition: Hyperplane $\left\{y: s^{T} y=r\right\}$ supports $S$ at $x \in \operatorname{bd} S$ if

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

## Supporting hyperplane theorem

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



## Connection to duality and subgradients

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 ${ }_{x}(f(x)+g(x))$ 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 ${ }^{1}$

[^0]
## Convex functions

- Graph below line connecting any two pairs $(x, f(x))$ and $(y, f(y))$

- Function $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ is convex if for all $x, y \in \mathbb{R}^{n}$ and $\theta \in[0,1]$ :

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

(in extended valued arithmetics)

- A function $f$ is concave if $-f$ is convex


## 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 epif is closed set


## 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 $\operatorname{dom} f$
2. 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 $x=1$


## 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

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

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


## Subdifferential calculus rules

- Subdifferential of sum $\partial\left(f_{1}+f_{2}\right)$
- Subdifferential of composition with matrix $\partial(g \circ L)$


## Subdifferential of sum

If $f_{1}, f_{2}$ closed convex and relint $\operatorname{dom} f_{1} \cap \operatorname{relint} \operatorname{dom} f_{2} \neq \emptyset$ :

$$
\partial\left(f_{1}+f_{2}\right)=\partial f_{1}+\partial f_{2}
$$

- One direction always holds: if $x \in \operatorname{dom} \partial f_{1} \cap \operatorname{dom} \partial f_{2}$ :

$$
\partial\left(f_{1}+f_{2}\right)(x) \supseteq \partial f_{1}(x)+\partial f_{2}(x)
$$

Proof: let $s_{i} \in \partial f_{i}(x)$, add subdifferential definitions:

$$
f_{1}(y)+f_{2}(y) \geq f_{1}(x)+f_{2}(x)+\left(s_{1}+s_{2}\right)^{T}(y-x)
$$

i.e. $s_{1}+s_{2} \in \partial\left(f_{1}+f_{2}\right)(x)$

- If $f_{1}$ and $f_{2}$ differentiable, we have (without convexity of $f$ )

$$
\nabla\left(f_{1}+f_{2}\right)=\nabla f_{1}+\nabla f_{2}
$$

## Subdifferential of composition

$$
\begin{aligned}
& \text { If } f \text { closed convex and relint } \operatorname{dom}(f \circ L) \neq \emptyset: \\
& \qquad \partial(f \circ L)(x)=L^{T} \partial f(L x)
\end{aligned}
$$

- One direction always holds: If $L x \in \operatorname{dom} f$, then

$$
\partial(f \circ L)(x) \supseteq L^{T} \partial f(L x)
$$

Proof: let $s \in \partial f(L x)$, then by definition of subgradient of $f$ :
$(f \circ L)(y) \geq(f \circ L)(x)+s^{T}(L y-L x)=(f \circ L)(x)+\left(L^{T} s\right)^{T}(y-x)$
i.e., $L^{T} s \in \partial(f \circ L)(x)$

- If $f$ differentiable, we have chain rule (without convexity of $f$ )

$$
\nabla(f \circ L)(x)=L^{T} \nabla f(L x)
$$

## A sufficient optimality condition

$$
\begin{align*}
& \text { Let } f: \mathbb{R}^{m} \rightarrow \overline{\mathbb{R}}, g: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}} \text {, and } L \in \mathbb{R}^{m \times n} \text { then: } \\
& \qquad \begin{array}{l}
\text { minimize } f(L x)+g(x) \\
\text { is solved by every } x \in \mathbb{R}^{n} \text { that satisfies } \\
\qquad 0 \in L^{T} \partial f(L x)+\partial g(x)
\end{array} \tag{1}
\end{align*}
$$

- Subdifferential calculus inclusions say:

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

which by Fermat's rule is equivalent to $x$ solution to (1)

- Note: (1) can have solution but no $x$ exists that satisfies (2)


## A necessary and sufficient optimality condition

Let $f: \mathbb{R}^{m} \rightarrow \overline{\mathbb{R}}, g: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}, L \in \mathbb{R}^{m \times n}$ with $f, g$ closed convex and assume relint $\operatorname{dom}(f \circ L) \cap$ relint $\operatorname{dom} g \neq \emptyset$ then:

$$
\begin{equation*}
\operatorname{minimize} f(L x)+g(x) \tag{1}
\end{equation*}
$$

is solved by $x \in \mathbb{R}^{n}$ if and only if $x$ satisfies

$$
\begin{equation*}
0 \in L^{T} \partial f(L x)+\partial g(x) \tag{2}
\end{equation*}
$$

- Subdifferential calculus equality rules say:

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

which by Fermat's rule is equivalent to $x$ solution to (1)

- Algorithms search for $x$ that satisfy $0 \in L^{T} \partial f(L x)+\partial g(x)$


## Evaluating subgradients of convex functions

- Obviously need to evaluate subdifferentials to solve

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

- Explicit evaluation:
- If function is differentiable: $\nabla f$ (unique)
- If function is nondifferentiable: compute element in $\partial f$
- Implicit evaluation:
- Proximal operator (specific element of subdifferential)


## Proximal operator

- Proximal operator of (convex) $g$ defined as:

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

where $\gamma>0$ is a parameter

- Evaluating prox requires solving optimization problem
- Objective is strongly convex $\Rightarrow$ solution exists and is unique


## Prox evaluates the subdifferential

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

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

Hence, $\gamma^{-1}(z-x)$ is element in $\partial g(x)$

- A subgradient in $\partial g(x)$ where $x=\operatorname{prox}_{\gamma g}(z)$ is computed
- Often used in algorithms when $g$ nonsmooth (no gradient exists)


## Conjugate functions

- The conjugate function of $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ is defined as

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

- Implicit definition via optimization problem


## Conjugate interpretation

- Conjugate $f^{*}(s)$ defines affine minorizer to $f$ with slope $s$ :

where $f^{*}(s)$ decides the constant offset to have support at $x^{*}$
- "Affine minorizor generator: Pick slope $s$, get offset for support"
- Why? Consider $f^{*}(s)=\sup _{x}\left(s^{T} x-f(x)\right)$ with maximizer $x^{*}$ :

$$
\begin{aligned}
& f^{*}(s)=s^{T} x^{*}-f\left(x^{*}\right) \quad \Leftrightarrow \\
& \Leftrightarrow \\
& f^{*}(s) \geq s^{T} x-f(x) \text { for all } x \\
& f(x) \geq s^{T} x-f^{*}(s) \text { for all } x
\end{aligned}
$$

- Support at $x^{*}$ since $f\left(x^{*}\right)=s^{T} x^{*}-f^{*}(s)$


## Fenchel Young's equality

- Going back to conjugate interpretation:

- Fenchel's inequality: $f(x) \geq s^{T} x-f^{*}(s)$ for all $x, s$
- Fenchel-Young's equality and equivalence:

$$
f\left(x^{*}\right)=s^{T} x^{*}-f^{*}(s) \text { holds if and only if } s \in \partial f\left(x^{*}\right)
$$

## A subdifferential formula

Assume $f$ closed convex, then $\partial f(x)=\operatorname{Argmax}_{s}\left(s^{T} x-f^{*}(s)\right)$

- Since $f^{* *}=f$, we have $f(x)=\sup _{s}\left(x^{T} s-f^{*}(s)\right)$ and

$$
\begin{aligned}
s^{*} \in \underset{s}{\operatorname{Argmax}}\left(x^{T} s-f^{*}(s)\right) & \Longleftrightarrow f(x)=x^{T} s^{*}-f^{*}\left(s^{*}\right) \\
& \Longleftrightarrow s^{*} \in \partial f(x)
\end{aligned}
$$

- The last equivalence is Fenchel-Young


## Subdifferential of conjugate - Inversion formula

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

- Consequence of Fenchel-Young
- Another way to write the result is that for closed convex $f$ :

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

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

## 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:



## 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 $f$


## 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

$$
\begin{aligned}
& f(\theta x+(1-\theta) y) \geq \theta f(x)+(1-\theta) f(y)-\frac{\beta}{2} \theta(1-\theta)\|x-y\|^{2} \\
& f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) f(y)+\frac{\beta}{2} \theta(1-\theta)\|x-y\|^{2}
\end{aligned}
$$

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

- Smoothness does not imply convexity
- Example:


## First-order condition for smoothness

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

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

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


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


## 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 ${ }_{x}(f(L x)+g(x))$ 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:

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

- 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 $-L^{T} \mu \in \partial g(x)$


## Solving the primal via the dual

- 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 $\underset{x}{\operatorname{minimize}} f(x)+g(x)$ 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 $\gamma_{k} \in\left[\epsilon, \beta^{-1}\right], \epsilon>0$ :

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

gives proximal gradient method

## 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

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

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

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

- 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 ${ }^{1}$
- $\gamma_{k} \in\left[\epsilon, \frac{2}{\beta\|L\|_{2}^{2}}-\epsilon\right]$
- 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

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

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


## 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: $f_{i}=g_{j}^{*}\left(\left(-X^{T} \cdot\right)_{j}\right), \psi_{j}=L^{*}$


## 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 $\nabla f$


[^0]:    ${ }^{1}$ 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 $-L^{T} \mu \in \partial g(x)$.

