# Proximal Gradient Method 

Pontus Giselsson

## Outline

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


## Composite optimization problems

- We have introduced the composite optimization problem

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

- Need an algorithm that solves it - proximal gradient method
- We will consider the simpler composite optimization problem

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

that gives the former by letting $f \rightarrow f \circ L$

## Problem assumptions

- Proximal gradient method works, e.g., for problems that satisfy
- $f$ is $\beta$-smooth $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ (not necessarily convex)
- $g$ is closed convex
- Recall that if $\beta$-smoothness implies that $f$ satisfies

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

it has convex quadratic upper and concave quadratic lower bounds

- If $f$ in addition is convex, we instead have

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

where the concave quadratic lower bound is replaced by affine

## Minimizing upper bound

- 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 upper bound to l.h.s.

- Minimizing in every iteration the r.h.s. w.r.t. $y$ for given $x$ gives

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

## Proximal gradient method

- Let us replace $\beta$ by $\gamma_{k}^{-1}, x$ by $x_{k}$, and $v$ by $x_{k+1}$ to get:

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

- This is exactly the proximal gradient method
- The method replaces $f$ by quadratic approximation and minimizes
- (Note that we need an initial guess $x_{0}$ to start the iteration)


## 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 - 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 minimize $(f(x))$
- $\operatorname{prox}_{\gamma_{k} g}(z)=\operatorname{argmin}_{x}\left(0+\frac{1}{2 \gamma}\|x-z\|_{2}^{2}\right)=z$
- 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 $(g(x))$
- $\nabla f(x)=0$
- 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

## Proximal gradient - Fixed-point characterization

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

- 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


## Meaning of 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 $\epsilon_{\mathrm{rel}}, \epsilon>0$ :

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

- 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 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: $x_{k+1}=\operatorname{prox}_{\gamma_{k} g}\left(x_{k}-\gamma_{k} L^{T} \nabla f\left(L x_{k}\right)\right)$
${ }^{1}$ Prox friendly: proximal operator cheap to evaluate, e.g., $g$ separable


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

