# ORIE 4741: Learning with Big Messy Data Proximal Gradient Method 

Professor Udell<br>Operations Research and Information Engineering Cornell

October 16, 2021

## Announcements

- Homework 5 due next Thursday 11/14
- Pick up midterm exams from Prof. Udell's office hours
- Bug fix!
- update package repository:
using Pkg
Pkg.update()
- download new version of proxgrad.jl from demos repository


## Demo: proximal gradient

https://github.com/ORIE4741/demos/blob/master/
proxgrad-starter-code.ipynb

## Regularized empirical risk minimization

choose model by solving

$$
\operatorname{minimize} \sum_{i=1}^{n} \ell\left(x_{i}, y_{i} ; w\right)+r(w)
$$

with variable $w \in \mathbf{R}^{d}$

- parameter vector $w \in \mathbf{R}^{d}$
- loss function $\ell: \mathcal{X} \times \mathcal{Y} \times \mathbf{R}^{d} \rightarrow \mathbf{R}$
- regularizer $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$


## Regularized empirical risk minimization

choose model by solving

$$
\operatorname{minimize} \sum_{i=1}^{n} \ell\left(x_{i}, y_{i} ; w\right)+r(w)
$$

with variable $w \in \mathbf{R}^{d}$

- parameter vector $w \in \mathbf{R}^{d}$
- loss function $\ell: \mathcal{X} \times \mathcal{Y} \times \mathbf{R}^{d} \rightarrow \mathbf{R}$
$\rightarrow$ regularizer $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$
why?
- want to minimize the risk $\mathbb{E}_{(x, y) \sim p} \ell(x, y ; w)$
- approximate it by the empirical risk $\sum_{i=1}^{n} \ell(x, y ; w)$
- add regularizer to help model generalize


## Solving regularized risk minimization

how should we fit these models?

- with a different software package for each model?
- with a different algorithm for each model?
- with a general purpose optimization solver?
desiderata
- fast
- flexible


## What's wrong with gradient descent?

Q: Why can't we use gradient descent to solve all our problems?

## What's wrong with gradient descent?

Q: Why can't we use gradient descent to solve all our problems? A: Because some regularizers and loss functions aren't differentiable!

## Subgradient

## Definition

The vector $g \in \mathbf{R}^{d}$ is a subgradient of $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$ at $x$ if

$$
f(y) \geq f(x)+g^{\top}(y-x), \quad \forall y \in \mathbf{R}^{d}
$$

The subdifferential of $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$ at $x$ is the set of all subgradients of $f$ at $x$ :

$$
\partial f(x)=\left\{g: f(y) \geq f(x)+g^{\top}(y-x) \forall y \in \mathbf{R}^{d}\right\}
$$

## Subgradient

## Definition

The vector $g \in \mathbf{R}^{d}$ is a subgradient of $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$ at $x$ if

$$
f(y) \geq f(x)+g^{\top}(y-x), \quad \forall y \in \mathbf{R}^{d}
$$

The subdifferential of $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$ at $x$ is the set of all subgradients of $f$ at $x$ :

$$
\partial f(x)=\left\{g: f(y) \geq f(x)+g^{\top}(y-x) \forall y \in \mathbf{R}^{d}\right\}
$$

- one subgradient for each supporting hyperplane of $f$ at $x$


## Subgradient

## Definition

The vector $g \in \mathbf{R}^{d}$ is a subgradient of $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$ at $x$ if

$$
f(y) \geq f(x)+g^{\top}(y-x), \quad \forall y \in \mathbf{R}^{d}
$$

The subdifferential of $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$ at $x$ is the set of all subgradients of $f$ at $x$ :

$$
\partial f(x)=\left\{g: f(y) \geq f(x)+g^{\top}(y-x) \forall y \in \mathbf{R}^{d}\right\}
$$

- one subgradient for each supporting hyperplane of $f$ at $x$
- the subdifferential $\partial f$ maps points to sets


## Subgradient

for $f: \mathbf{R} \rightarrow \mathbf{R}$ and convex, here's a simpler equivalent condition:

- if $f$ is differentiable at $x, \partial f(x)=\{\nabla f(x)\}$
- if $f$ is not differentiable at $x$, it will still be differentiable just to the left and the right of $x^{1}$, so
- let $g^{+}=\lim _{\epsilon \rightarrow 0} \nabla f(x+\epsilon)$
- let $g^{-}=\lim _{\epsilon \rightarrow 0} \nabla f(x-\epsilon)$
- $\partial f(x)$ is any convex combination (i.e., any weighted average) of those gradients:

$$
\partial f(x)=\left\{\alpha g^{+}+(1-\alpha) g^{-}: \alpha \in[0,1]\right\}
$$

## Subgradient: examples

compute subgradient wrt prediction vector $z \in \mathbf{R}$ :

- quadratic loss: $\ell(y, z)=(y-z)^{2}$
- $\ell_{1}$ loss: $\ell(y, z)=|y-z|$
- hinge loss: $\ell(y, z)=(1-y z)_{+}$
- logistic loss: $\ell(y, z)=\log (1+\exp (-y z))$


## Important properties of subdifferential

- Linearity.

$$
\partial_{w} \sum_{(x, y) \in \mathcal{D}} \ell\left(y, w^{\top} x\right)=\sum_{(x, y) \in \mathcal{D}} \partial_{w} \ell\left(y, w^{\top} x\right)
$$

## Important properties of subdifferential

- Linearity.

$$
\partial_{w} \sum_{(x, y) \in \mathcal{D}} \ell\left(y, w^{\top} x\right)=\sum_{(x, y) \in \mathcal{D}} \partial_{w} \ell\left(y, w^{\top} x\right)
$$

- Chain rule. If $f=h \circ g, h: \mathbf{R} \rightarrow \mathbf{R}$, and $g: \mathbf{R}^{d} \rightarrow \mathbf{R}$ is differentiable,

$$
\partial f(x)=\partial h(g(x)) \nabla g(x)
$$

## Important properties of subdifferential

- Linearity.

$$
\partial_{w} \sum_{(x, y) \in \mathcal{D}} \ell\left(y, w^{\top} x\right)=\sum_{(x, y) \in \mathcal{D}} \partial_{w} \ell\left(y, w^{\top} x\right)
$$

- Chain rule. If $f=h \circ g, h: \mathbf{R} \rightarrow \mathbf{R}$, and $g: \mathbf{R}^{d} \rightarrow \mathbf{R}$ is differentiable,

$$
\partial f(x)=\partial h(g(x)) \nabla g(x)
$$

Example. For $z=w^{\top} x$,

$$
\partial_{w} \ell\left(y, w^{\top} x\right)=\partial_{z} \ell(y, z) \nabla_{w}\left(w^{\top} x\right)=x \partial_{z} \ell(y, z)
$$

## Subgradient method

$$
\text { minimize } \quad \ell(w)
$$

Algorithm Subgradient method
Given: function $\ell: \mathbf{R}^{d} \rightarrow \mathbf{R}$, stepsize sequence $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters

- compute subgradient $g \in \partial \ell(w)$
- update $w$ :

$$
w \leftarrow w-\alpha^{t} g
$$

## Stochastic subgradient method

stochastic subgradient obeys

$$
\mathbb{E} \tilde{\partial} \ell(w) \in \partial \ell(w)
$$

examples: for $\ell(w)=\sum_{i=1}^{n} \ell\left(y_{i}, w^{\top} x_{i}\right)$,

## Stochastic subgradient method

stochastic subgradient obeys

$$
\mathbb{E} \tilde{\partial} \ell(w) \in \partial \ell(w)
$$

examples: for $\ell(w)=\sum_{i=1}^{n} \ell\left(y_{i}, w^{\top} x_{i}\right)$,

- single stochastic gradient. pick a random example $i$. set $z_{i}=w^{\top} x_{i}$ and compute

$$
\tilde{\partial} \ell(w)=n x_{i} \partial_{z} \ell\left(y_{i}, z_{i}\right)
$$

## Stochastic subgradient method

stochastic subgradient obeys

$$
\mathbb{E} \tilde{\partial} \ell(w) \in \partial \ell(w)
$$

examples: for $\ell(w)=\sum_{i=1}^{n} \ell\left(y_{i}, w^{\top} x_{i}\right)$,

- single stochastic gradient. pick a random example $i$. set $z_{i}=w^{\top} x_{i}$ and compute

$$
\tilde{\partial} \ell(w)=n x_{i} \partial_{z} \ell\left(y_{i}, z_{i}\right)
$$

- minibatch stochastic gradient. pick a random set of examples $S$. set $z_{i}=w^{\top} x_{i}$ for $i \in S$ and compute

$$
\begin{aligned}
\tilde{\partial} \ell(w) & =\frac{n}{|S|} \partial\left(\sum_{i \in S} \ell\left(y_{i}, w^{\top} x_{i}\right)\right) \\
& =\frac{n}{|S|} \sum_{i \in S} x_{i} \partial_{z} \ell\left(y_{i}, z_{i}\right)
\end{aligned}
$$

## Convergence for stochastic subgradient method

suppose $\ell$ is convex, subdifferentiable, Lipschitz continuous. convergence results:

- stochastic (sub)gradient, fixed step size $\alpha^{t}=\alpha$ :
- iterates converge quickly, then wander within a small ball
- stochastic (sub)gradient, decreasing step size $\alpha^{t}=1 / t$ :
- iterates converge slowly to solution
proofs: [Bertsekas, 2010] https://arxiv.org/pdf/1507.01030v1.pdf


## What's wrong with the subgradient method?

Q: Why can't we use the subgradient method to solve all our problems?

## What's wrong with the subgradient method?

Q: Why can't we use the subgradient method to solve all our problems?
A:

1) because some of our regularizers don't have subgradients everywhere (e.g., $\mathbf{1}_{+}$).
2) proximal gradient is way faster.

## Proximal operator

define the proximal operator of the function $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$

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

## Proximal operator

define the proximal operator of the function $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$

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

$-\operatorname{prox}_{r}: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$

## Proximal operator

define the proximal operator of the function $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$

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

$\rightarrow \operatorname{prox}_{r}: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$

- generalized projection: if $\mathbf{1}_{C}$ is the indicator of set $C$,

$$
\operatorname{prox}_{1_{C}}(w)=\Pi_{C}(w)
$$

## Proximal operator

define the proximal operator of the function $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$

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

$\rightarrow \operatorname{prox}_{r}: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$

- generalized projection: if $\mathbf{1}_{C}$ is the indicator of set $C$,

$$
\operatorname{prox}_{1_{C}}(w)=\Pi_{C}(w)
$$

- implicit gradient step: if $w=\operatorname{prox}_{r}(z)$ and $r$ is smooth,

$$
\begin{aligned}
\nabla r(w)+w-z & =0 \\
w & =z-\nabla r(w)
\end{aligned}
$$

## Proximal operator

define the proximal operator of the function $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$

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

$-\operatorname{prox}_{r}: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$

- generalized projection: if $\mathbf{1}_{C}$ is the indicator of set $C$,

$$
\operatorname{prox}_{1_{C}}(w)=\Pi_{C}(w)
$$

- implicit gradient step: if $w=\operatorname{prox}_{r}(z)$ and $r$ is smooth,

$$
\begin{aligned}
\nabla r(w)+w-z & =0 \\
w & =z-\nabla r(w)
\end{aligned}
$$

- simple to evaluate: closed form solutions for many functions


## Maps from functions to functions

no consistent notation for map from functions to functions.
for a function $f: \mathbf{R}^{d} \rightarrow \mathbf{R}$,

- prox maps $f$ to a new function prox $_{f}: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$
- $\operatorname{prox}_{f}(x)$ evaluates this function at the point $x$
$-\nabla$ maps $f$ to a new function $\nabla f: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$
- $\nabla f(x)$ evaluates this function at the point $x$
$-\frac{\partial}{\partial x}$ maps $f$ to a new function $\frac{\partial f}{\partial x}: \mathbf{R}^{d} \rightarrow \mathbf{R}^{d}$
- $\left.\frac{\partial f}{\partial x}(x)\right|_{x=\bar{x}}$ evaluates this function at the point $\bar{x}$
- this one has the most confusing notation of all...


## Let's evaluate some proximal operators!

define the proximal operator of the function $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$

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

- $r(w)=0$ (identity)
- $r(w)=\sum_{i=1}^{d} r_{i}\left(w_{i}\right)$ (separable)
- $r(w)=\|w\|_{2}^{2}$ (shrinkage)
- $r(w)=\|w\|_{1}$ (soft-thresholding)
- $r(w)=\mathbf{1}(w \geq 0)$ (projection)
- $r(w)=\sum_{i=1}^{d-1}\left(w_{i+1}-w_{i}\right)^{2}$ (smoothing)


## Proximal (sub)gradient method

want to solve

$$
\operatorname{minimize} \quad \ell(w)+r(w)
$$

$-\ell: \mathbf{R}^{d} \rightarrow \mathbf{R}$ subdifferentiable
$-r: \mathbf{R}^{d} \rightarrow \mathbf{R}$ with a fast prox operator
Algorithm Proximal (sub)gradient method
Given: loss $\ell: \mathbf{R}^{d} \rightarrow \mathbf{R}$, regularizer $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$, stepsizes $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters
Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters

- compute subgradient $g \in \partial \ell(w)$
(O(nd) flops)
- update w:
( $\mathcal{O}(d)$ flops)

$$
w \leftarrow \operatorname{prox}_{\alpha^{t} r}\left(w-\alpha^{t} g\right)
$$

## Example: NNLS

$$
\operatorname{minimize} \frac{1}{2}\|y-X w\|^{2}+\mathbf{1}(w \geq 0)
$$

recall

- $\nabla\left(\frac{1}{2}\|y-X w\|^{2}\right)=-X^{T}(y-X w)$
$-\operatorname{prox}_{1(\cdot \geq 0)}(w)=\max (0, w)$
Algorithm Proximal gradient method for NNLS
Given: $X \in \mathbf{R}^{n \times d}, y \in \mathbf{R}^{n}$, stepsize sequence $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters
- compute gradient $g=X^{T}(X w-y)$
( $\mathcal{O}(n d)$ flops)
- update $w$ :
( $\mathcal{O}(d)$ flops)

$$
w \leftarrow \max \left(0, w-\alpha^{t} g\right)
$$

## Example: NNLS

$$
\operatorname{minimize} \frac{1}{2}\|y-X w\|^{2}+\mathbf{1}(w \geq 0)
$$

recall

- $\nabla\left(\frac{1}{2}\|y-X w\|^{2}\right)=-X^{T}(y-X w)$
$-\operatorname{prox}_{1(\cdot \geq 0)}(w)=\max (0, w)$
Algorithm Proximal gradient method for NNLS
Given: $X \in \mathbf{R}^{n \times d}, y \in \mathbf{R}^{n}$, stepsize sequence $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters
- compute gradient $g=X^{T}(X w-y)$
( $\mathcal{O}(n d)$ flops)
- update $w$ :
( $\mathcal{O}(d)$ flops)

$$
w \leftarrow \max \left(0, w-\alpha^{t} g\right)
$$

## Example: NNLS

option: do work up front to reduce per-iteration complexity
Algorithm Proximal gradient method for NNLS
Given: $X \in \mathbf{R}^{n \times d}, y \in \mathbf{R}^{n}$, stepsize sequence $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
compute:

- $b=X^{\top} y$
( $\mathcal{O}(n d)$ flops)
- $G=X^{\top} X$
$\left(\mathcal{O}\left(n d^{2}\right)\right.$ flops $)$

For: $t=1, \ldots$, maxiters

- compute gradient $g=G w-b$
( $\mathcal{O}\left(d^{2}\right)$ flops $)$
- update $w$ :
( $\mathcal{O}(d)$ flops)

$$
w \leftarrow \max \left(0, w-\alpha^{t} g\right)
$$

$\mathcal{O}\left(n d^{2}\right)$ flops to begin, $\mathcal{O}\left(d^{2}\right)$ flops per iteration

## Example: Lasso

$$
\operatorname{minimize} \quad \frac{1}{2}\|y-X w\|^{2}+\lambda\|w\|_{1}
$$

recall

- $\nabla\left(\frac{1}{2}\|y-X w\|^{2}\right)=-X^{T}(y-X w)$
$-\operatorname{prox}_{\mu\|\cdot\|_{1}}(w)=s_{\mu}(w)$ where

$$
\left(s_{\mu}(w)\right)_{i}= \begin{cases}w_{i}-\mu & w_{i} \geq \mu \\ 0 & \left|w_{i}\right| \leq \mu \\ w_{i}+\mu & w_{i} \leq-\mu\end{cases}
$$

## Example: Lasso

$$
\operatorname{minimize} \quad \frac{1}{2}\|y-X w\|^{2}+\lambda\|w\|_{1}
$$

Algorithm Proximal gradient method for Lasso
Given: $X \in \mathbf{R}^{n \times d}, y \in \mathbf{R}^{n}$, stepsize sequence $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters

- compute gradient $g=X^{T}(X w-y)$
( $\mathcal{O}(n d)$ flops)
- update $w$ :
( $\mathcal{O}(d)$ flops)

$$
w \leftarrow s_{\alpha^{t} \lambda}\left(w-\alpha^{t} g\right)
$$

## Example: Lasso

$$
\operatorname{minimize} \quad \frac{1}{2}\|y-X w\|^{2}+\lambda\|w\|_{1}
$$

Algorithm Proximal gradient method for Lasso
Given: $X \in \mathbf{R}^{n \times d}, y \in \mathbf{R}^{n}$, stepsize sequence $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters

- compute gradient $g=X^{T}(X w-y)$
( $\mathcal{O}(n d)$ flops)
- update $w$ :

$$
w \leftarrow s_{\alpha^{t} \lambda}\left(w-\alpha^{t} g\right)
$$

notice: the hard part is computing the gradient! (can speed this up by precomputation, as for NNLS)

## Convergence

two questions to ask:

- will the iteration ever stop?
- what kind of point will it stop at?
if the iteration stops, we say it has converged


## Convergence: what kind of point will it stop at?

- let's suppose $r$ is differentiable ${ }^{2}$
- if we find $w$ so that

$$
w=\operatorname{prox}_{\alpha^{t} r}\left(w-\alpha^{t} \nabla \ell(w)\right)
$$

then

$$
\begin{aligned}
w & =\underset{w^{\prime}}{\operatorname{argmin}}\left(\alpha^{t} r\left(w^{\prime}\right)+\frac{1}{2}\left\|w^{\prime}-\left(w-\alpha^{t} \nabla \ell(w)\right)\right\|_{2}^{2}\right) \\
0 & =\nabla \alpha^{t} r(w)+w-w+\alpha^{t} \nabla \ell(w) \\
& =\nabla(r(w)+\ell(w))
\end{aligned}
$$

- so the gradient of the objective is 0
- if $\ell$ and $r$ are convex, that means $w$ minimizes $\ell+r$

[^0]
## Convergence: will it stop?

definitions:

- $p^{\star}=\inf _{w} \ell(w)+r(w)$
assumptions:
- loss function is continuously differentiable and $L$-smooth:

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

- for simplicity, consider constant step size $\alpha^{t}=\alpha$


## Proximal point method converges

prove it for $\ell=0$ first (aka the proximal point method) for any $t=0,1, \ldots$,

$$
w^{t+1}=\underset{w}{\operatorname{argmin}} \alpha r(w)+\frac{1}{2}\left\|w-w^{t}\right\|^{2}
$$

so in particular,

$$
\begin{aligned}
\alpha r\left(w^{t+1}\right)+\frac{1}{2}\left\|w^{t+1}-w^{t}\right\|^{2} & \leq \alpha r\left(w^{t}\right)+\frac{1}{2}\left\|w^{t}-w^{t}\right\|^{2} \\
\frac{1}{2 \alpha}\left\|w^{t+1}-w^{t}\right\|^{2} & \leq r\left(w^{t}\right)-r\left(w^{t+1}\right)
\end{aligned}
$$

now add up these inequalities for $t=0,1, \ldots, T$ :

$$
\begin{aligned}
\frac{1}{2 \alpha} \sum_{t=0}^{T}\left\|w^{t+1}-w^{t}\right\|^{2} & \leq \sum_{t=0}^{T}\left(r\left(w^{t}\right)-r\left(w^{t+1}\right)\right) \\
& \leq r\left(w^{0}\right)-p^{\star}
\end{aligned}
$$

it converges!

## Proximal gradient method converges (I)

now prove it for $\ell \neq 0$. for any $t=0,1, \ldots$,

$$
w^{t+1}=\underset{w}{\operatorname{argmin}} \alpha r(w)+\frac{1}{2}\left\|w-\left(w^{t}-\alpha \nabla \ell\left(w^{t}\right)\right)\right\|^{2}
$$

so in particular,

$$
\begin{aligned}
r\left(w^{t+1}\right) & +\frac{1}{2 \alpha}\left\|w^{t+1}-\left(w^{t}-\alpha \nabla \ell\left(w^{t}\right)\right)\right\|^{2} \\
& \leq r\left(w^{t}\right)+\frac{1}{2 \alpha}\left\|w^{t}-\left(w^{t}-\alpha \nabla \ell\left(w^{t}\right)\right)\right\|^{2} \\
r\left(w^{t+1}\right) & \left.+\frac{1}{2 \alpha}\left\|w^{t+1}-w^{t}\right\|^{2}+\frac{\alpha}{2} \| \nabla \ell\left(w^{t}\right)\right) \|^{2}+\left\langle\nabla \ell\left(w^{t}\right), w^{t+1}-w^{t}\right\rangle \\
& \leq r\left(w^{t}\right)+\frac{\alpha}{2}\left\|\nabla \ell\left(w^{t}\right)\right\|^{2} \\
r\left(w^{t+1}\right) & +\frac{1}{2 \alpha}\left\|w^{t+1}-w^{t}\right\|^{2}+\left\langle\nabla \ell\left(w^{t}\right), w^{t+1}-w^{t}\right\rangle \\
& \leq r\left(w^{t}\right)
\end{aligned}
$$

## Proximal gradient method converges (II)

now use $\ell\left(w^{\prime}\right) \leq \ell(w)+\left\langle\nabla \ell(w), w^{\prime}-w\right\rangle+\frac{L}{2}\left\|w^{\prime}-w\right\|^{2}$
with $w^{\prime}=w^{t+1}, w=w^{t}$

$$
\begin{aligned}
\ell\left(w^{t+1}\right) & +r\left(w^{t+1}\right)+\frac{1}{2 \alpha}\left\|w^{t+1}-w^{t}\right\|^{2}+\left\langle\nabla \ell\left(w^{t}\right), w^{t+1}-w^{t}\right\rangle \\
& \leq \ell\left(w^{t}\right)+r\left(w^{t}\right)+\left\langle\nabla \ell\left(w^{t}\right), w^{t+1}-w^{t}\right\rangle+\frac{L}{2}\left\|w^{t+1}-w^{t}\right\|^{2} \\
\ell\left(w^{t+1}\right) & +r\left(w^{t+1}\right)+\frac{1}{2 \alpha}\left\|w^{t+1}-w^{t}\right\|^{2} \\
& \leq \ell\left(w^{t}\right)+r\left(w^{t}\right)+\frac{L}{2}\left\|w^{t+1}-w^{t}\right\|^{2} \\
\frac{1}{2 \alpha} \| w^{t+1} & -w^{t}\left\|^{2}-\frac{L}{2}\right\| w^{t+1}-w^{t} \|^{2} \\
& \leq \ell\left(w^{t}\right)+r\left(w^{t}\right)-\left(\ell\left(w^{t+1}\right)+r\left(w^{t+1}\right)\right)
\end{aligned}
$$

## Proximal gradient method converges (III)

now add up these inequalities for $t=0,1, \ldots, T$ :

$$
\begin{aligned}
\frac{1}{2}\left(\frac{1}{\alpha}-L\right) & \sum_{t=0}^{T}\left\|w^{t+1}-w^{t}\right\|^{2} \\
& \left.\leq \sum_{t=0}^{T} \ell\left(w^{t}\right)+r\left(w^{t}\right)-\left(\ell\left(w^{t+1}\right)+r\left(w^{t+1}\right)\right)\right) \\
& \leq \ell\left(w^{0}\right)+r\left(w^{0}\right)-p^{\star}
\end{aligned}
$$

if

$$
\begin{aligned}
& \frac{1}{\alpha}-L \geq 0 \\
& \Longrightarrow \alpha \leq \frac{1}{L}
\end{aligned}
$$

it converges!

## Stochastic proximal subgradient

Algorithm Stochastic proximal subgradient method
Given: loss $\ell: \mathbf{R}^{d} \rightarrow \mathbf{R}$, regularizer $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$, stepsizes $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters
Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters

- pick $S \subseteq\{1, \ldots, n\}$ uniformly at random
- pick $g \in \tilde{\partial} \ell(w)$
( $\mathcal{O}(|S| d)$ flops $)$
- update $w$ : ( $\mathcal{O}(d)$ flops)

$$
w \leftarrow \operatorname{prox}_{\alpha^{t} r}\left(w-\alpha^{t} g\right)
$$

## Stochastic proximal subgradient

Algorithm Stochastic proximal subgradient method
Given: loss $\ell: \mathbf{R}^{d} \rightarrow \mathbf{R}$, regularizer $r: \mathbf{R}^{d} \rightarrow \mathbf{R}$, stepsizes $\left\{\alpha^{t}\right\}_{t=1}^{\infty}$, maxiters
Initialize: $w \in \mathbf{R}^{d}$ (often, $w=0$ )
For: $t=1, \ldots$, maxiters

- pick $S \subseteq\{1, \ldots, n\}$ uniformly at random
- pick $g \in \tilde{\partial} \ell(w)$
( $\mathcal{O}(|S| d)$ flops $)$
- update $w$ : ( $\mathcal{O}(d)$ flops)

$$
w \leftarrow \operatorname{prox}_{\alpha^{t} r}\left(w-\alpha^{t} g\right)
$$

per iteration complexity: $\mathcal{O}(|S| d)$

## Convergence for stochastic proximal (sub)gradient

pick your poison:

- stochastic (sub)gradient, fixed step size $\alpha^{t}=\alpha$ :
- iterates converge quickly, then wander within a small ball
- stochastic (sub)gradient, decreasing step size $\alpha^{t}=1 / t$ :
- iterates converge slowly to solution
- minibatch stochastic (sub)gradient with increasing minibatch size, fixed step size $\alpha^{t}=\alpha$ :
- iterates converge quickly to solution
- later iterations take (much) longer
proofs: [Bertsekas, 2010] https://arxiv.org/pdf/1507.01030v1.pdf conditions:
- $\ell$ is convex, subdifferentiable, Lipschitz continuous, and
- $r$ is convex and Lipschitz continuous where it is $<\infty$
or
- all iterates are bounded


## References

- Beck: book chapter on proximal operators (lots of examples!) https:
//archive.siam.org/books/mo25/mo25_ch6.pdf
- Vandenberghe: lecture on proximal gradient method.
http://www.seas.ucla.edu/~vandenbe/236C/ lectures/proxgrad.pdf
- Yin: lecture on proximal method. http://www.math.ucla.edu/~wotaoyin/summer2013/ slides/Lec05_ProximalOperatorDual.pdf
- Parikh and Boyd: paper on proximal algorithms. https: //stanford.edu/~boyd/papers/pdf/prox_algs.pdf
- Bertsekas: convergence proofs for every proximal gradient style method you can dream of.
https://arxiv.org/pdf/1507.01030v1.pdf


[^0]:    ${ }^{2}$ take Convex Optimization for the proof for non-differentiable $r$

