# Algorithms and Convergence 

Pontus Giselsson

## Outline

- Algorithm overview
- Convergence and convergence rates
- Proving convergence rates


## What is an algorithm?

- We are interested in algorithms that solve composite problems

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

- An algorithm:
- generates a sequence $\left(x_{k}\right)_{k \in \mathbb{N}}$ that hopefully converges to solution
- often creates next point in sequence according to

$$
x_{k+1}=\mathcal{A}_{k} x_{k}
$$

where

- $\mathcal{A}_{k}$ is a mapping that gives the next point from the current
- $\mathcal{A}_{k}=\operatorname{prox}_{\gamma_{k} g} \circ\left(I-\gamma_{k} \nabla f\right)$ for proximal gradient method


## Deterministic and stochastic algorithms

- We have deterministic algorithms

$$
x_{k+1}=\mathcal{A}_{k} x_{k}
$$

that given initial $x_{0}$ will give the same sequence $\left(x_{k}\right)_{k \in \mathbb{N}}$

- We will also see stochastic algorithms that iterate

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

where $\xi_{k}$ is a random variable that also decides the mapping

- $\left(x_{k}\right)_{k \in \mathbb{N}}$ is a stochastic process, i.e., collection of random variables
- when running the algorithm, we evaluate $\xi_{k}$ and get a realization
- different realization $\left(x_{k}\right)_{k \in \mathbb{N}}$ every time even if started at same $x_{0}$
- Stochastic algorithms useful although problem is deterministic


## Optimization algorithm overview

- Algorithms can roughly be divided into the following classes:
- Second-order methods
- Quasi second-order methods
- First-order methods
- Stochastic and coordinate-wise first-order methods
- The first three are typically deterministic and the last stochastic
- Cost of computing one iteration decreases down the list


## Second-order methods

- Solves problems using second-order (Hessian) information
- Requires smooth (twice continuously differentiable) functions
- Example: Newton's method to minimize smooth function $f$ :

$$
x_{k+1}=x_{k}-\gamma_{k}\left(\nabla^{2} f\left(x_{k}\right)\right)^{-1} \nabla f\left(x_{k}\right)
$$

- Constraints can be incorporated via barrier functions:
- Use sequence of smooth constraint barrier functions
- Make barriers increasingly well approximate constraint set
- For each barrier, solve smooth problem using Newton's method
- Resulting scheme called interior point method
- (Can be applied to directly solve primal-dual optimality condition)
- Computational backbone: solving linear systems $O\left(n^{3}\right)$
- Often restricted to small to medium scale problems
- We will cover Newton's method


## Quasi second-order methods

- Estimates second-order information from first-order
- Solves problems using estimated second-order information
- Requires smooth (twice continuously differentiable) functions
- Quasi-Newton method for smooth $f$

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

where $B_{k}$ is:

- estimate of Hessian inverse (not Hessian to avoid inverse)
- cheaply computed from gradient information
- Computational backbone: forming $B_{k}$ and matrix multiplication
- Limited memory versions exist with cheaper iterations
- Can solve large-scale smooth problems
- Will briefly look into most common method (BFGS)


## First-order methods

- Solves problems using first-order (sub-gradient) information
- Computational primitives: (sub)gradients and proximal operators
- Use gradient if function differentiable, prox if nondifferentiable
- Examples for solving $\underset{x}{\operatorname{minimize}} f(x)+g(x)$
- Proximal gradient method (requires smooth $f$ since gradient used)

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

- Douglas-Rachford splitting (no smoothness requirement)

$$
z_{k+1}=\frac{1}{2} z_{k}+\frac{1}{2}\left(2 \operatorname{prox}_{\gamma g}-I\right)\left(2 \operatorname{prox}_{\gamma f}-I\right) z_{k}
$$

and $x_{k}=\operatorname{prox}_{\gamma f}\left(z_{k}\right)$ converges to solution

- Iteration often cheaper than second-order if function split wisely
- Can solve large-scale problems
- Will look at proximal gradient method and accelerated version


## Stochastic and coordinate-wise first-order methods

- 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

- Algorithm overview
- Convergence and convergence rates
- Proving convergence rates


## Types of convergence

- 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


## 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
- We analyze stochastic process and will see summability, e.g., of:
- 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: $\mathbb{E}\left[\min _{l \in\{0, \ldots, k\}}\left\|\nabla f\left(x_{l}\right)\right\|_{2}\right] \rightarrow 0$
- Says what happens with expected value of different quantities


## Algorithm realizations - Summable case

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

$$
\mathbb{E}\left[\left\|x_{k}-x^{\star}\right\|_{2}\right] \quad \text { or } \quad \mathbb{E}\left[f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star}\right] \quad \text { or } \mathbb{E}\left[\left\|\nabla f\left(x_{k}\right)\right\|_{2}\right]
$$

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


## Linear rates - Comparison

- Different rates in log-lin plot

- Called linear rate since linear in log-lin plot


## Quadratic rates

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

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

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

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

- Quadratic $\left(\rho^{2^{k}}\right)$ vs linear $\left(\rho^{k}\right)$ rate with factor $\rho=0.9$ :

|  |
| :---: |
|  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |


| Linear |
| :---: |
| 1.000000000000 |
| 0.900000000000 |
| 0.810000000000 |
| 0.659000000000 |
| 0.59049945000 |
| 0.531440050000 |
| 0.478296936000 |
| 0.430467270000 |

- Example: Locally for Newton's method with strongly convex cost


## Quadratic rates - Comparison

- Different rates in log-lin scale

- Quadratic convergence is superlinear


## Sublinear rates

- A rate is sublinear if it is slower than linear
- A sublinear rate can, for instance, be of the form

$$
\begin{aligned}
f\left(x_{k}\right)+g\left(x_{k}\right)-p^{\star} & \leq \frac{C}{\psi(k)} \\
\left\|x_{k+1}-x_{k}\right\|_{2}^{2} & \leq \frac{C}{\psi(k)} \\
\min _{l=0, \ldots, k} \mathbb{E}\left[\left\|\nabla f\left(x_{l}\right)\right\|_{2}^{2}\right] & \leq \frac{C}{\psi(k)}
\end{aligned}
$$

where $C>0$ and $\psi$ decides how fast it decreases, e.g.,

- $\psi(k)=\log k$ : Stochastic gradient descent $\gamma_{k}=c / k$
- $\psi(k)=\sqrt{k}$ : Stochastic gradient descent: optimal $\gamma_{k}$
- $\psi(k)=k$ : Proximal gradient, coordinate proximal gradient
- $\psi(k)=k^{2}$ : Accelerated proximal gradient method with improved rate further down the list
- We say that the rate is $O\left(\frac{1}{\psi(k)}\right)$ for the different $\psi$
- To be sublinear $\psi$ has slower than exponential growth


## Sublinear rates - Comparison

- Different rates on log-lin scale

- Many iterations may be needed for high accuracy


## Rate vs iteration cost

- Consider these classes of algorithms
- Second-order methods
- Quasi second-order methods
- First-order methods
- Stochastic and coordinate-wise first-order methods
- Rate deteriorates and iterations increase as we go down the list $\Downarrow$
- Iteration cost increases as we go up the list $\Uparrow$
- Performance is roughly (\# iterations) $\times$ (iteration cost)
- This gives a tradeoff when selecting algorithm
- Rough advise for problem size: small $(\Uparrow)$ medium ( $\Uparrow \Downarrow)$ large $(\Downarrow)$


## Outline

- Algorithm overview
- Convergence and convergence rates
- Proving convergence rates


## Proving convergence rates

- To prove a convergence rate typically requires
- Using inequalities that describe problem class
- Using algorithm definition equalities (or inclusions)
- Combine these to a form so that convergence can be concluded
- Linear and quadratic rates proofs conceptually straightforward
- Sublinear rates implicit via a Lyapunov inequality


## Proving linear or quadratic rates

- 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
- ...


## Sublinear convergence - Lyapunov inequality

- 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., $W:=\sum_{k=0}^{\infty} W_{k}<\infty$
- 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., $\inf _{k \in \mathbb{N}} \lambda_{k}>0$


## 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

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

## 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:

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

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

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

- 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:

$$
\hat{R}_{k} \leq \frac{V_{0}+\bar{W}}{\psi(k)} \quad \text { where } \quad \lambda_{k}=\phi(k) \text { and } \psi(k)=\int_{t=0}^{k} \phi(t) d t
$$

- 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:

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

- 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

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

