FINITE-SAMPLE GUARANTEES OF CONTRACTIVE STOCHASTIC APPROXIMATION WITH APPLICATIONS IN REINFORCEMENT LEARNING

Siva Theja Maguluri

Industrial and Systems Engineering
Georgia Institute of Technology

Zaiwei Chen

Computing and Mathematical Sciences
California Institute of Technology

PART I
STOCHASTIC APPROXIMATION

Siva Theja Maguluri

Industrial and Systems Engineering
Georgia Institute of Technology

PART II
REINFORCEMENT LEARNING

Zaiwei Chen

Computing and Mathematical Sciences
California Institute of Technology

JOINT WORK WITH

Martin Zubeldia, Univ of Minnesota

Karthikeyan Shanmugam Google Research

Sanjay Shakkottai UT Austin

PART I

STOCHASTIC APPROXIMATION

BANACH FIXED POINT THEOREM

Want to find $\mathbf{x}^*$ that solves

$$\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$$

A simple iteration

$$\mathbf{x}_{k+1} = \bar{\mathbf{F}}(\mathbf{x}_k)$$

Banach Fixed Point Theorem

$\mathbf{x}_k$ converges to $\mathbf{x}^*$ geometrically fast (linearly) if $\bar{\mathbf{F}}\,(.)$ is a contraction

Contraction: For all $\mathbf{x}$ and $\mathbf{y}$, $\quad \left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \leq \gamma \,\|\mathbf{x} - \mathbf{y}\|$

Works for any norm

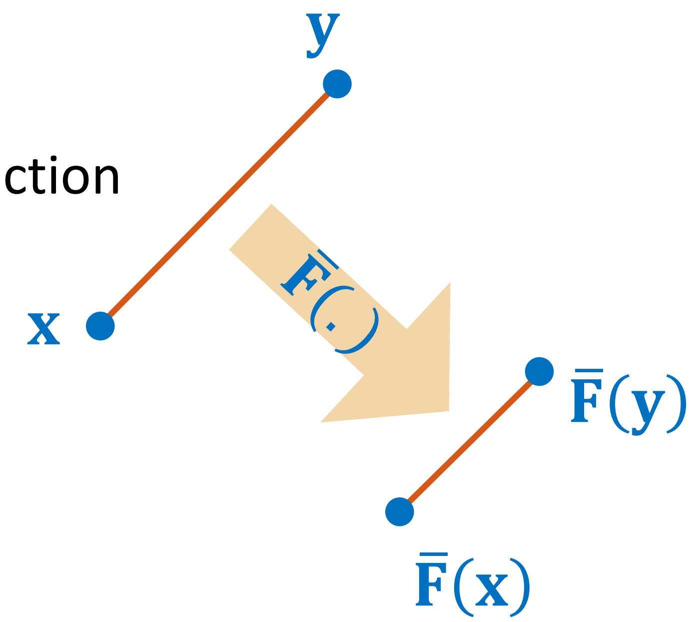

BANACH FIXED POINT THEOREM

Want to find $\mathbf{x}^*$ that solves

$$\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$$

A simple iteration

Noisy Oracle

$$\mathbf{x}_{k+1} = \bar{\mathbf{F}}(\mathbf{x}_k) + \mathbf{w}_k$$

Banach Fixed Point Theorem

$\mathbf{x}_k$ converges to $\mathbf{x}^*$ geometrically fast (linearly) if $\bar{\mathbf{F}}(\cdot)$ is a pseudo-contraction

Pseudo-Contraction: For all $\mathbf{x}$, $\|\bar{\mathbf{F}}(\mathbf{x}) - \mathbf{x}^*\| \leq \gamma \, \|\mathbf{x} - \mathbf{x}^*\|$

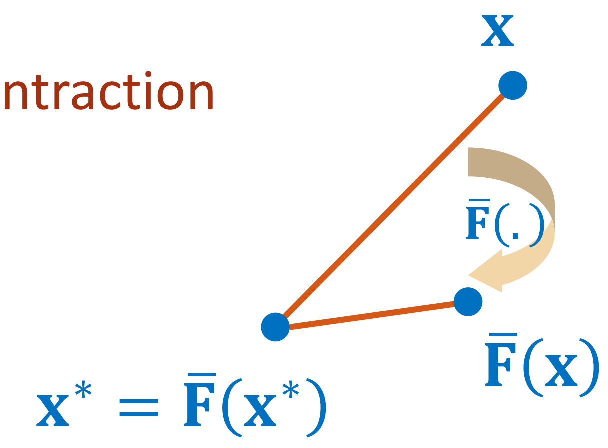

STOCHASTIC APPROXIMATION

Want to find $\mathbf{x}^*$ that solves

$$\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$$

A simple iteration

$$\mathbf{x}_{k+1} = \bar{\mathbf{F}}(\mathbf{x}_k) + \mathbf{w}_k$$

Noisy Oracle

Stochastic Approximation[Robbins, Monro '51]

$$\begin{aligned}\mathbf{x}_{k+1} &= (1 - \alpha_k)\mathbf{x}_k + \alpha_k(\bar{\mathbf{F}}(\mathbf{x}_k) + \mathbf{w}_k) \\ &= \mathbf{x}_k + \alpha_k(\bar{\mathbf{F}}(\mathbf{x}_k) + \mathbf{w}_k - \mathbf{x}_k)\end{aligned}$$

Question: How well does this work?

OUTLINE

- Stochastic Approximation Introduction
- Finite Sample bounds on the mean-square error $\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}^*\|^2]$
  - Proof Sketch - A Lyapunov function
- High Probability bounds on $\|\mathbf{x}_k - \mathbf{x}^*\|$ (Exponentially decaying)
  - Proof Sketch – Exponential Supermartingale and Bootstrapping

STOCHASTIC APPROXIMATION

FIXED POINT PROBLEMS

Stochastic Approximation to solve $\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_k(\bar{\mathbf{F}}(\mathbf{x}_{\mathrm{k}}) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

Optimization:

$$\min f(\mathbf{x})$$

$$-\eta\boldsymbol{\nabla} f(\mathbf{x}) + \mathbf{x} = \mathbf{x}$$

When f is smooth strongly convex, $\bar{\mathbf{F}}(\mathbf{x}) = -\eta\boldsymbol{\nabla} f(\mathbf{x}) + \mathbf{x}$ is contraction wrt ℓ_2-norm

SGD: $\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} - \alpha_{\mathrm{k}}(\boldsymbol{\nabla} f(\mathbf{x}_{\mathrm{k}}) + \mathbf{w}_{\mathrm{k}})$

FIXED POINT PROBLEMS

Stochastic Approximation to solve $\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_k(\bar{\mathbf{F}}(\mathbf{x}_{\mathrm{k}}) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

Markov Decision Processes and RL:

$\bar{\mathbf{F}}\,(\cdot)$ is related to the Bellman operator.

TD learning, Q learning and their variants can be modeled as SA

The underlying norm is weighted ℓ_p (for TD) and ℓ_∞ (for Q learning)

More details in **Part II**

FIXED POINT PROBLEMS

Stochastic Approximation to solve $\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_k(\bar{\mathbf{F}}(\mathbf{x}_{\mathrm{k}}) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

Linear Equations:

$$\mathbf{Ax} = \mathbf{b}$$

$$(\mathbf{I} + \eta\mathbf{A})\mathbf{x} - \boldsymbol{\eta}\mathbf{b} = \mathbf{x}$$

When $\mathbf{A}$ is Hurwitz ($\mathrm{Re}(\lambda_{\mathrm{i}}) < 0$), $\bar{\mathbf{F}}(\mathbf{x}) = (\mathbf{I} + \eta\mathbf{A})\mathbf{x} - \boldsymbol{\eta}\mathbf{b}$ is contraction wrt weighted ℓ_2-norm

Linear SA: $\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_{\mathrm{k}}(\mathbf{Ax}_{\mathrm{k}} - \mathbf{b}_{\mathrm{k}})$

MARKOVIAN STOCHASTIC APPROXIMATION

Want to find $\mathbf{x}^*$ that solves

$$\bar{\mathbf{F}}(\mathbf{x}) = \mathbb{E}_{\mathbf{Y}\sim\boldsymbol{\mu}}\left[\mathbf{F}(\mathbf{x},\mathbf{Y})\right] = \mathbf{x}$$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k(\mathbf{A}_k\mathbf{x}_k - \mathbf{b})$$

Markovian Stochastic Approximation

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k(\mathbf{F}(\mathbf{x}_k,\mathbf{Y_k}) + \mathbf{w}_k - \mathbf{x}_k)$$

Multiplicative Noise   Additive Noise

(Main) Assumptions

- $\mathbf{Y_k}$ is a finite state Ergodic Markov chain with stationary distribution $\boldsymbol{\mu}$
  - $\mathbf{Y_k}$ is geometrically mixing
- Noise $\mathbf{w}_k$ - iid or martingale difference, mean zero, $\|\mathbf{w}_k\| \leq B(\|\mathbf{x}_k\| + 1)$
- $\bar{\mathbf{F}}(.)$ is a contraction w.r.t arbitrary norm $\left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \leq \gamma\,\|\mathbf{x}-\mathbf{y}\|$

MEAN SQUARE BOUNDS

FIXED STEP SIZE

Markovian Stochastic Approximation $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha \left(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k \right)$

$$\left\| \bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y}) \right\| \leq \gamma \left\| \mathbf{x} - \mathbf{y} \right\|$$

ℓ_∞-norm contraction

$\log d$

Theorem[Chen, M, Shakkottai, Shanmugam '21]: If the step-size α is small enough,

$$\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}^*\|^2] \leq c_1 (1 - c_2 \alpha)^{k - \log \alpha^{-1}} + c_3 \alpha \log \alpha^{-1}$$

$\|\mathbf{x}_0 - \mathbf{x}^*\|^2$

$\frac{1-\gamma}{2}$

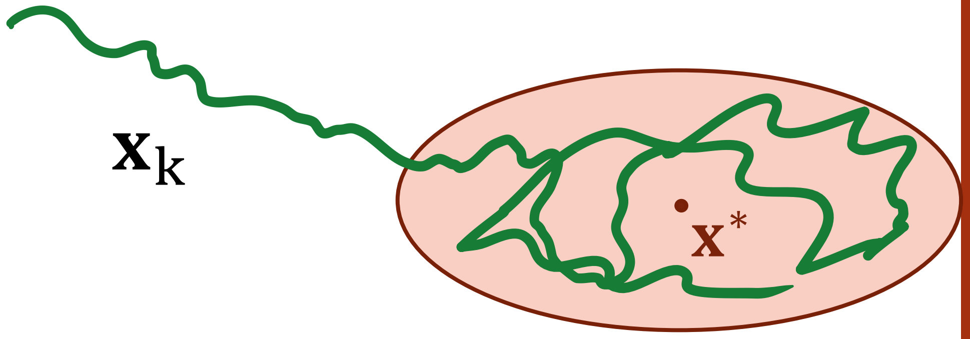

FIXED STEP SIZE

Markovian Stochastic Approximation $\quad \mathbf{x}_{k+1} = \mathbf{x}_k + \alpha\left(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k\right)$

$$\left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \leq \gamma \left\|\mathbf{x} - \mathbf{y}\right\|$$

ℓ_∞-norm contraction

$\log d$

Theorem[Chen, M, Shakkottai, Shanmugam '21]: If the step-size α is small enough,

$$\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}^*\|^2] \leq c_1(1 - c_2\alpha)^{k - \log \alpha^{-1}} + c_3 \alpha \log \alpha^{-1}$$

- Given a target error ϵ, one can pick small enough step size so that eventually the error is ϵ.
  - Sample complexity of $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$

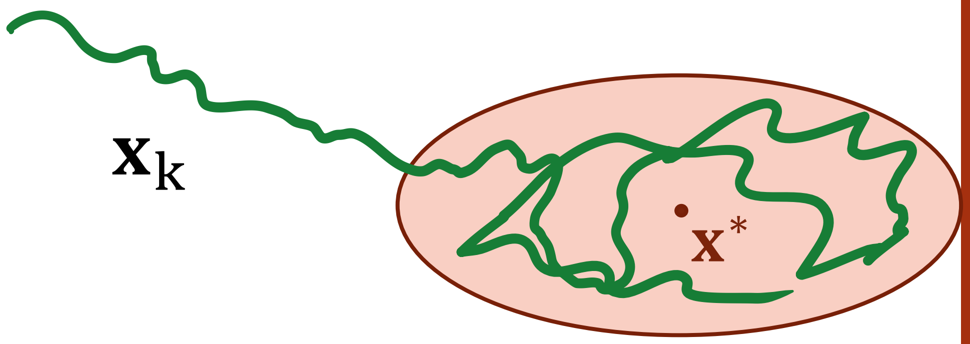

DIMINISHING STEP SIZES

Markovian Stochastic Approximation $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k)$

$$\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\| \leq \gamma \|\mathbf{x} - \mathbf{y}\|$$

$\alpha_k \sim \alpha / k^{\xi}$

Theorem[Chen, M, Shakkottai, Shanmugam '21]:

$$\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}^*\|^2] \leq \begin{cases} c_4 \dfrac{\ln k}{k^{\xi}} & \xi \in (0,1) \\ c_5 \dfrac{(\ln k)^2}{k^{\alpha c_2}} & \xi = 1, \alpha c_2 \leq 1 \\ \hat{c}_6 \left(\dfrac{\log \mathrm{d}}{(1-\gamma)^3}\right) \dfrac{\ln k}{\mathrm{k}} & \xi = 1, \alpha c_2 > 1 \end{cases}$$

$\|\mathbf{x}_0 - \mathbf{x}^*\|^2$

$\dfrac{1-\gamma}{2}$

DIMINISHING STEP SIZES

Markovian Stochastic Approximation $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k)$

$$\left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \leq \gamma \left\|\mathbf{x} - \mathbf{y}\right\|$$

$\alpha_k \sim \alpha / k^{\xi}$

Theorem [Chen, M, Shakkottai, Shanmugam '21]:

$$\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}^*\|^2] \leq \begin{cases} c_4 \dfrac{\ln k}{k^{\xi}} & \xi \in (0,1) \\ c_5 \dfrac{(\ln k)^2}{k^{\alpha c_2}} & \xi = 1, \alpha c_2 \leq 1 \\ \hat{c}_6 \left(\dfrac{\log \mathrm{d}}{(1-\gamma)^3}\right) \dfrac{\ln k}{\mathrm{k}} & \xi = 1, \alpha c_2 > 1 \end{cases}$$

$\frac{1-\gamma}{2}$

- This leads to a sample complexity of $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$
  - With continual improvement beyond this.
  - Algorithm does not depend on ϵ

RELATED WORK

SA mode	Operator	Context	Literature
No Mult noise	$\|\|.\|\|_2$-contraction	SGD	[Bottou et al 18]
Mult noise with boundedness	$\|\|.\|\|_\infty$-contraction	Q-learning	[Beck, Srikant 12,13] (poly d) (Need iterates to be bounded)
Linear	Hurwitz	TD-learning	[Srikant, Ying 19] (Markov Noise), [Lakshminarayanan and Szepesvari 18] (iid noise)
Markovian and Mult noise	Any norm contraction	SGD Q-learning TD-learning Off-policy TD	Our work Also recovers all prior results

PROOF SKETCH

STOCHASTIC APPROXIMATION: INTUITION

Stochastic Approximation $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k)$

Stochastic Approximation

$$\frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\alpha_k} = (\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k)$$

ODE

$$\dot{\mathbf{x}} = (\bar{\mathbf{F}}(\mathbf{x}) - \mathbf{x})$$

- ODE Method [Borkar '09]:
  - Stochastic Approximation converges asymptotically if the ODE is globally asymptotically stable (gas)
  - Show gas using a Lyapunov function, $\mathrm{M}(\mathbf{x}) = \|\mathbf{x}\|_\infty^2$: $\frac{d\mathrm{M}(\mathbf{x}-\mathbf{x}^*)}{dt} \le -\gamma \mathrm{M}(\mathbf{x} - \mathbf{x}^*)$
- Want: Error bounds on original SA. We do not use the ODE method.
- Challenge: We need to handle error terms

Control the Errors

$$\underbrace{\mathbf{x}_{k+1} - \mathbf{x}_k}_{\text{Discretization Error}} = \alpha_k \left(\underbrace{\bar{\mathbf{F}}(\mathbf{x}_k) - \mathbf{x}_k}_{\text{ODE Term}} + \underbrace{\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) - \bar{\mathbf{F}}(\mathbf{x}_k)}_{\text{Markovian Error}} + \underbrace{\mathbf{w}_k}_{\text{Additive Noise Error}} \right)$$

ODE VS STOCHASTIC APPROXIMATION

Stochastic Approximation

$$\mathbf{x}_{k+1} - \mathbf{x}_k = \alpha_k(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k)$$

ODE

$$\dot{\mathbf{x}} = \left(\bar{\mathbf{F}}(\mathbf{x}) - \mathbf{x}\right)$$

BAD NEWS

Lyapunov function $M(\mathbf{x}) = \|\mathbf{x}\|_\infty^2$ is not smooth

Smoothness: $M(\mathbf{y}) \leq M(\mathbf{x}) + \langle \nabla M(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|_\infty^2$

Approximation: $M(\mathbf{x}) \leq \|\mathbf{x}\|_\infty^2 \leq cM(\mathbf{x})$

$$\frac{dM(\mathbf{x} - \mathbf{x}^*)}{dt} \leq -\gamma M(\mathbf{x} - \mathbf{x}^*)$$

$$M(\mathbf{x}_{k+1} - \mathbf{x}^*) - M(\mathbf{x}_k - \mathbf{x}^*) \leq -\gamma\alpha_k M(\mathbf{x}_k - \mathbf{x}^*) + o(\alpha_k)$$

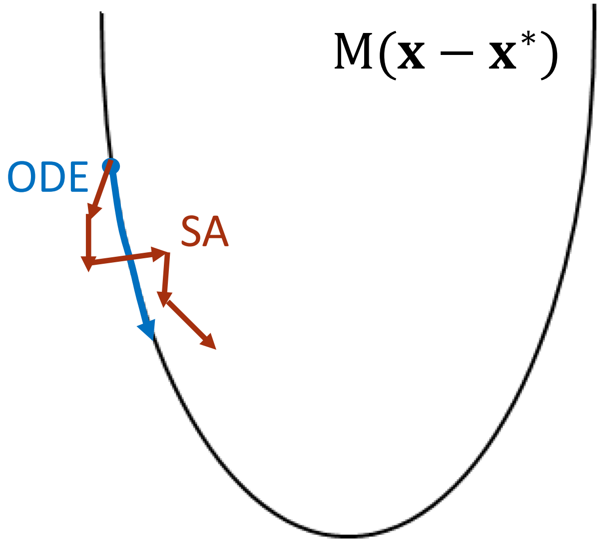

THE LYAPUNOV FUNCTION

WISHLIST

Smoothness: $M(\mathbf{y}) \leq M(\mathbf{x}) + \langle \nabla M(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|_\infty^{\mathbf{2}}$

Approximation: $M(\mathbf{x}) \leq \|\mathbf{x}\|_\infty^2 \leq cM(\mathbf{x})$

$$M(\mathbf{x}) = \|\mathbf{x}\|_\infty^2 \square \frac{1}{\mu} g(\mathbf{x}) = \min_{\mathbf{u}} \left\{ \|\mathbf{u}\|_\infty^2 + \frac{1}{\mu} g(\mathbf{x} - \mathbf{u}) \right\}$$

Moreau Envelope

$$\|\mathbf{x}\|_\infty^2 \square \frac{1}{2\mu} \|\mathbf{x}\|_{\mathbf{2}}^2$$

HANDLING THE ERRORS

Smoothness

$\|\mathbf{w}_k\| \leq A(\|\mathbf{x}_k\| + 1)$

$$\underbrace{\mathbf{x}_{k+1} - \mathbf{x}_k}_{\text{Discretization Error}} = \alpha_k \Big(\underbrace{\bar{\mathbf{F}}(\mathbf{x}_k) - \mathbf{x}_k}_{\text{ODE Term}} + \underbrace{\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) - \bar{\mathbf{F}}(\mathbf{x}_k)}_{\text{Markovian Error}} + \underbrace{\mathbf{w}_k}_{\text{Additive Noise Error}} \Big)$$

Due to smoothness, we are good, if we have a handle on error terms

Markovian Error:

- $\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k)$ is not same as its steady-state $\bar{\mathbf{F}}(\mathbf{x}_k)$
- The key term turns out to be a cross term

Fast mixing

$$\mathbb{E}\big[\langle \mathbf{x}_k, \mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) - \bar{\mathbf{F}}(\mathbf{x}_k)\rangle\big] = \mathbb{E}\big[\mathbb{E}\big[\langle \mathbf{x}_k, \mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) - \bar{\mathbf{F}}(\mathbf{x}_k)\rangle \,\big|\, \mathbf{x}_{k-\tau}, \mathbf{Y}_{k-\tau}\big]\big]$$

Mixing time

- For linear SA this was used in [Srikant, Ying '19] [Bertsikas, Tsitsiklis '96]

TAIL BOUNDS

TAIL BOUNDS

Stochastic Approximation to solve $\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha \quad (\mathbf{F}(\mathbf{x}_{\mathrm{k}}, \mathbf{Y}_k) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

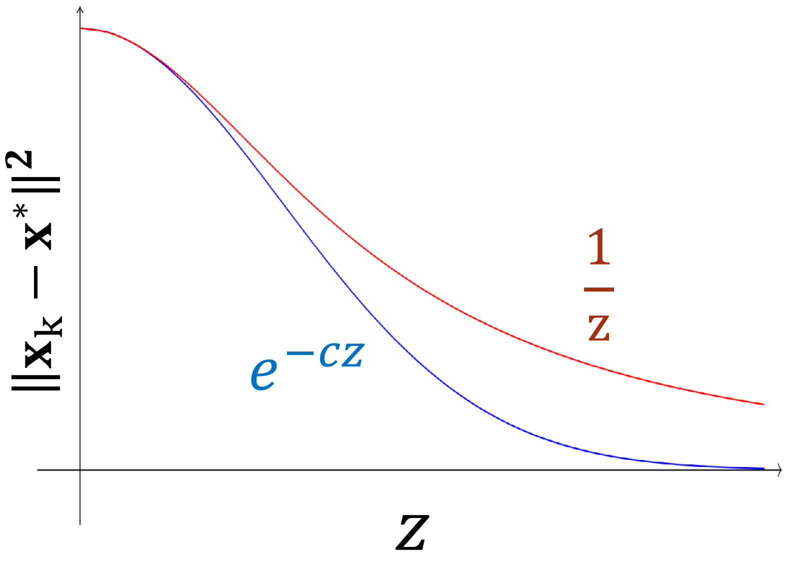

Mean Square Bounds:

$$\mathbb{E}\left[\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\|^2\right] \le O\left(\frac{1}{k}\right)$$

Using Markov Inequality, we get $\mathbb{P}\left(\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\|^2 \ge O\left(\frac{1}{k}\right) z\right) \le \frac{1}{z}$

Question: Can we get exponential tail bounds of the form

$$\mathbb{P}\left(\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\|^2 \ge O\left(\frac{1}{k}\right) \log\left(\frac{1}{\delta}\right)\right) \le \delta?$$

Yes

This implies sample complexity of $O\left(\frac{1}{\epsilon^2}\right) \log\left(\frac{1}{\delta}\right)$ to ensure $\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\| \le \epsilon$ w.p. $(1-\delta)$

LIMITATION OF CONSTANT STEP SIZES

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha \quad (\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k)$$

- Stationary distribution is heavy-tailed (Higher moments don't exist after a point) [Srikant, Ying '20].
  - Large enough moments keep increasing over time and become infinite in the limit.
  - While the mean square error converges to a constant, the tail is getting worse
- Several recent works obtain sample complexity of $O\left(\frac{1}{\epsilon^2}\right)\log\left(\frac{1}{\delta}\right)$ by picking constant step size as a function of ϵ and δ
  - [Telgarsky '22], [Mou et al '22], [Li et al '21], …
  - ϵ and δ have to be picked ahead of time and the algorithm (step size) is tuned for these (So cannot change mind later)
  - No improvement if it is run longer
  - The tail (beyond δ) can get worse the longer it is run
  - Bound only on specific point of the tail or a window and not the entire tail

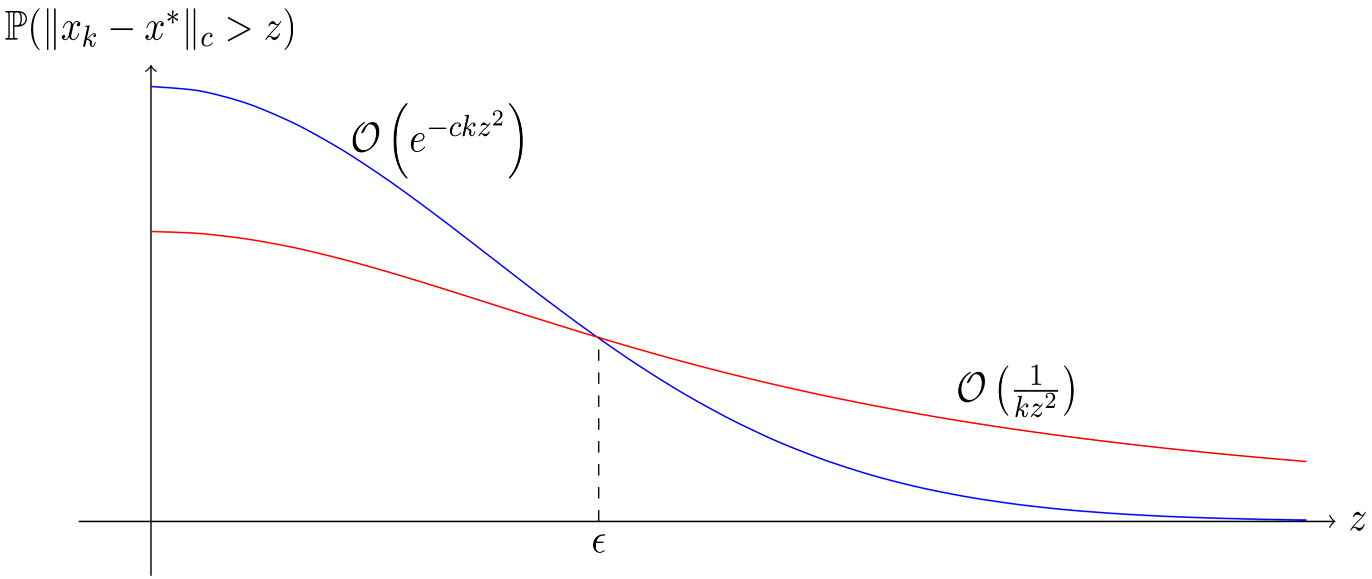

THE CHALLENGE

- Linear SA to solve $\mathbf{Ax} = \mathbf{b}$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_{\mathrm{k}}(\mathbf{A}_{\mathrm{k}}\mathbf{x}_{\mathrm{k}} - \mathbf{b}_{\mathrm{k}})$$

- Focus on multiplicative noise. Set $\mathrm{b}_k = 0$, we get product of matrices

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}}(\mathrm{I} + \alpha_{\mathrm{k}}\mathbf{A}_{\mathrm{k}})$$

$\mathbb{E}[\mathbf{A}_{\mathrm{k}}]$ is Hurwitz and $\mathbb{E}[(\mathrm{I} + \alpha_{\mathrm{k}}\mathbf{A}_{\mathrm{k}})]$ is contraction

The matrix $(\mathrm{I} + \alpha_{\mathrm{k}}\mathbf{A}_{\mathrm{k}})$ is not a contraction. It is a contraction only in **expectation**.

- Mean Square bounds under constant step sizes: [Lakshminarayanan, Szepeswari ‘18] [Srikant, Ying ‘19]
  - Tails are heavy
- Tail Bounds under constant step sizes [Durmus et al ‘21]
  - Exponential tails if $\mathbf{A}_{\mathrm{k}}$ is Hurwitz for all k. (i.e., if it is contractive at all times)
  - Polynomial tails otherwise

We get exponential tails with diminishing step sizes and do it for general contractive SA

STOCHASTIC APPROXIMATION

Want to find $\mathbf{x}^*$ that solves

$$\bar{\mathbf{F}}(\mathbf{x}) = \mathbb{E}_{\mathbf{Y}\sim\boldsymbol{\mu}}[\mathbf{F}(\mathbf{x},\mathbf{Y})] = \mathbf{x}$$

$$\alpha_k = \frac{\alpha}{k+h}$$

Stochastic Approximation

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_{\mathrm{k}}(\mathbf{F}(\mathbf{x}_{\mathrm{k}},\mathbf{Y}_{\mathbf{k}}) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

(Main) Assumptions

- $\mathbf{Y}_{\mathbf{k}}$ is an iid process with stationary distribution $\boldsymbol{\mu}$
- With bounded support

- $\mathbf{Y}_{\mathbf{k}}$ is such that $\mathbb{E}[\mathbf{F}(\mathbf{x},\mathbf{Y}_{\mathrm{k+1}})|\mathcal{F}_k] = \bar{\mathbf{F}}(\mathbf{x})$
- $\|\mathbf{F}(\mathbf{x},\mathbf{Y}_{\mathrm{k}}) - \bar{\mathbf{F}}(\mathbf{y})\| \leq B_1(\|\mathbf{x}_{\mathrm{k}}\|+1)$

- Noise $\mathbf{w}_{\mathrm{k}}$ - iid or martingale difference, mean zero, $\|\mathbf{w}_{\mathrm{k}}\| \leq B(\|\mathbf{x}_{\mathrm{k}}\| + 1)$
- $\bar{\mathbf{F}}(.)$ is a contraction w.r.t arbitrary norm $\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\| \leq \gamma\,\|\mathbf{x} - \mathbf{y}\|$

STOCHASTIC APPROXIMATION

Want to find $\mathbf{x}^*$ that solves

$$\bar{\mathbf{F}}(\mathbf{x}) = \mathbb{E}_{\mathbf{Y}\sim\boldsymbol{\mu}}\left[\mathbf{F}(\mathbf{x},\mathbf{Y})\right] = \mathbf{x}$$

$$\alpha_k = \frac{\alpha}{k+h}$$

Stochastic Approximation

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_{\mathrm{k}}(\mathbf{F}(\mathbf{x}_{\mathrm{k}},\mathbf{Y}_{\mathbf{k}}) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

(Main) Assumptions

- $\mathbf{Y}_{\mathbf{k}}$ is an iid process with stationary distribution $\boldsymbol{\mu}$
- With bounded support
- Noise $\mathbf{w}_{\mathrm{k}}$ - iid or martingale difference, mean zero, $\|\mathbf{w}_{\mathrm{k}}\| \leq B(\|\mathbf{x}_{\mathrm{k}}\| + 1)$
- $\bar{\mathbf{F}}(.)$ is a contraction w.r.t arbitrary norm $\left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \leq \gamma\,\|\mathbf{x}-\mathbf{y}\|$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha_{\mathrm{k}}(\mathbf{A}_{\mathrm{k}}\mathbf{x}_{\mathrm{k}} - \mathbf{b})$$

If $\mathbf{A}_{\mathrm{k}}$ is Gaussian, then, the MGF does not exist for $k \geq 3$

EXPONENTIAL TAILS

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \frac{\alpha}{k+h}\left(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k\right)$$

$\tilde{O}\left(\frac{1}{\epsilon^2}\right)\log\left(\frac{1}{\delta}\right)$ sample complexity
Don't need to fix ϵ and δ ahead

General Norm Contraction: $\left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \leq \gamma \left\|\mathbf{x} - \mathbf{y}\right\|$

Theorem[Zubeldia, Chen, Maguluri '22]: If α is large enough, for a given k, w.p. $(1-\delta)$,

$$\|\mathbf{x}_k - \mathbf{x}^*\|^2 \leq \begin{cases} \frac{c}{k}\left(1 + \log\left(\frac{1}{\delta}\right)\right) & if\ k \geq O\left(\log\left(\frac{1}{\delta}\right)\right) \\ k^\beta & \text{otherwise} \end{cases}$$

Why does the bound go up in the beginning?

WHY DOES THE ERROR GO UP?

Need enough samples for averaging to kick in to make sure the product of matrices becomes contractive.

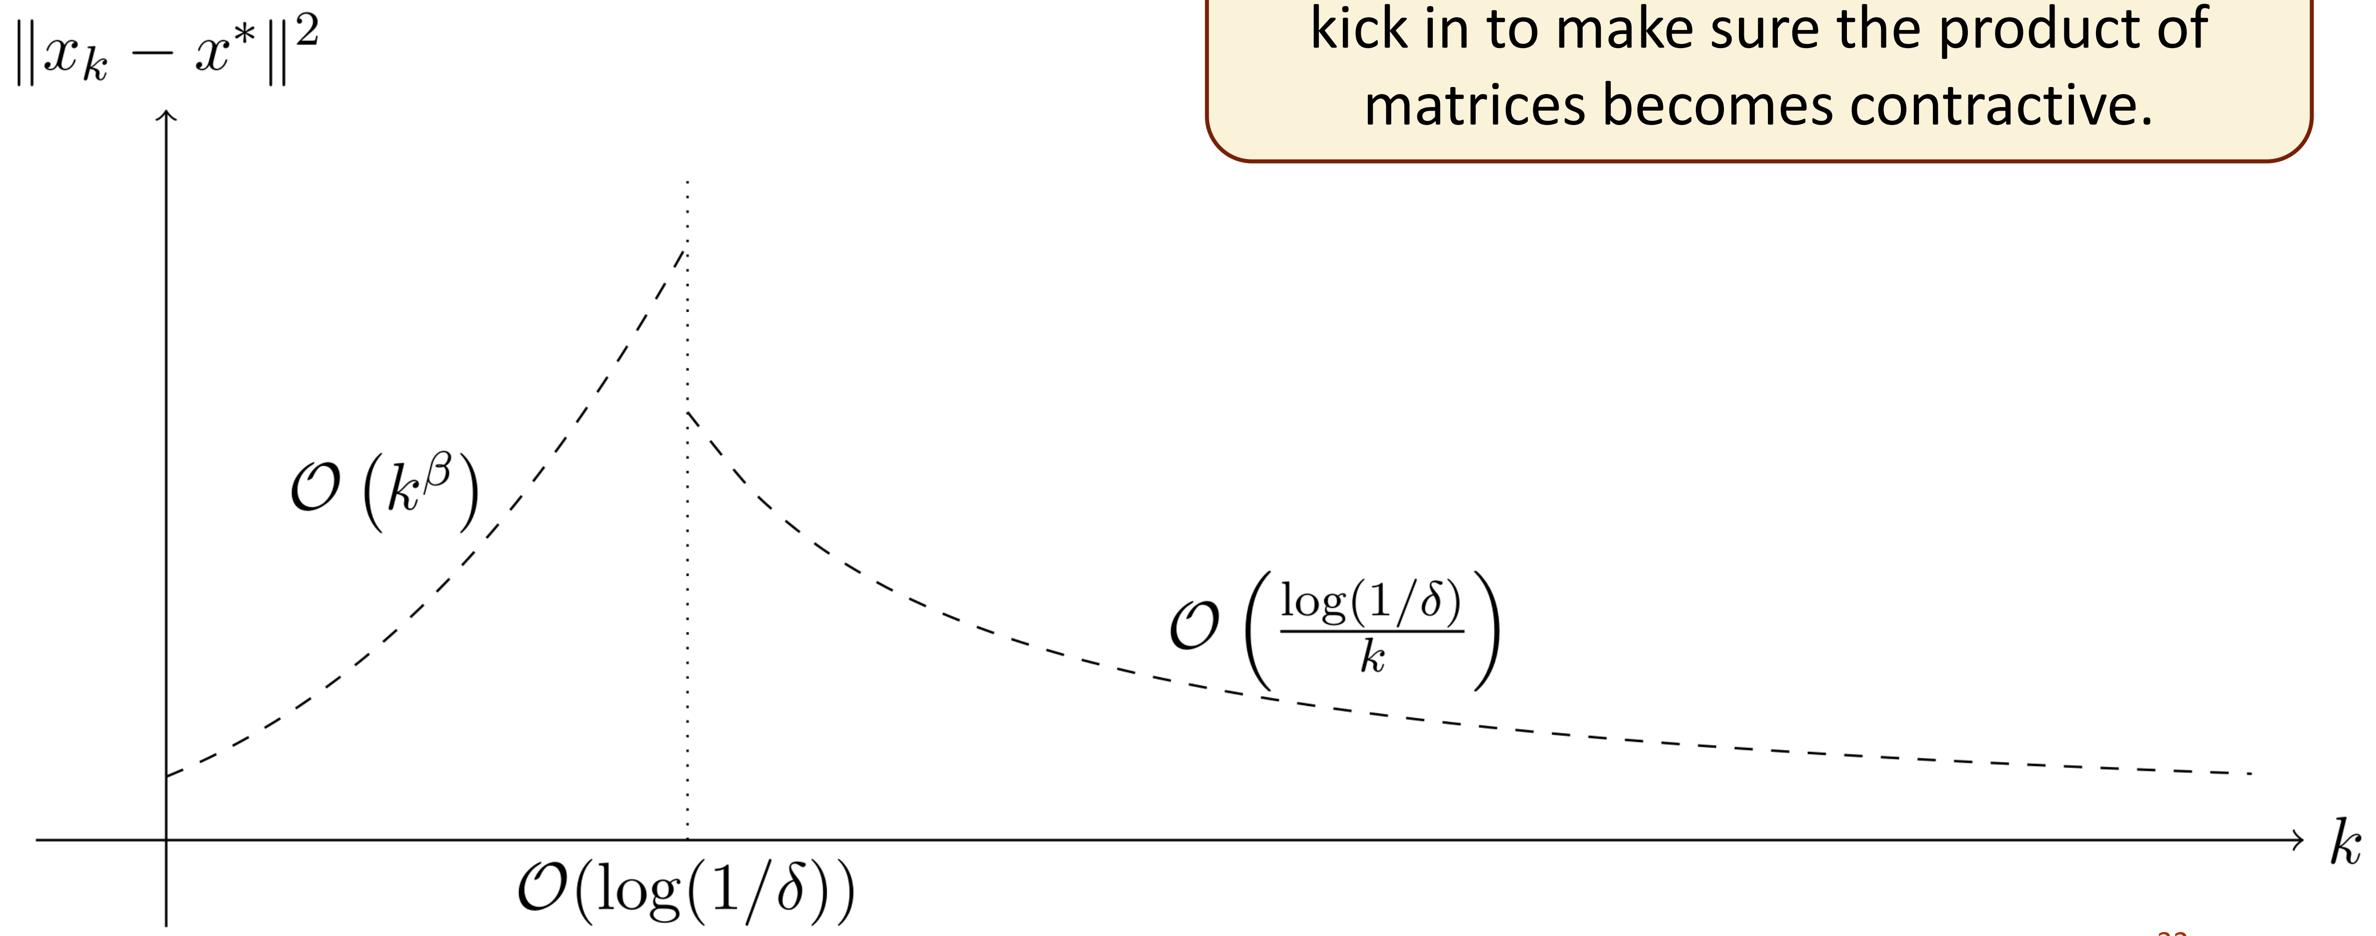

ERROR GOES UP INDEED

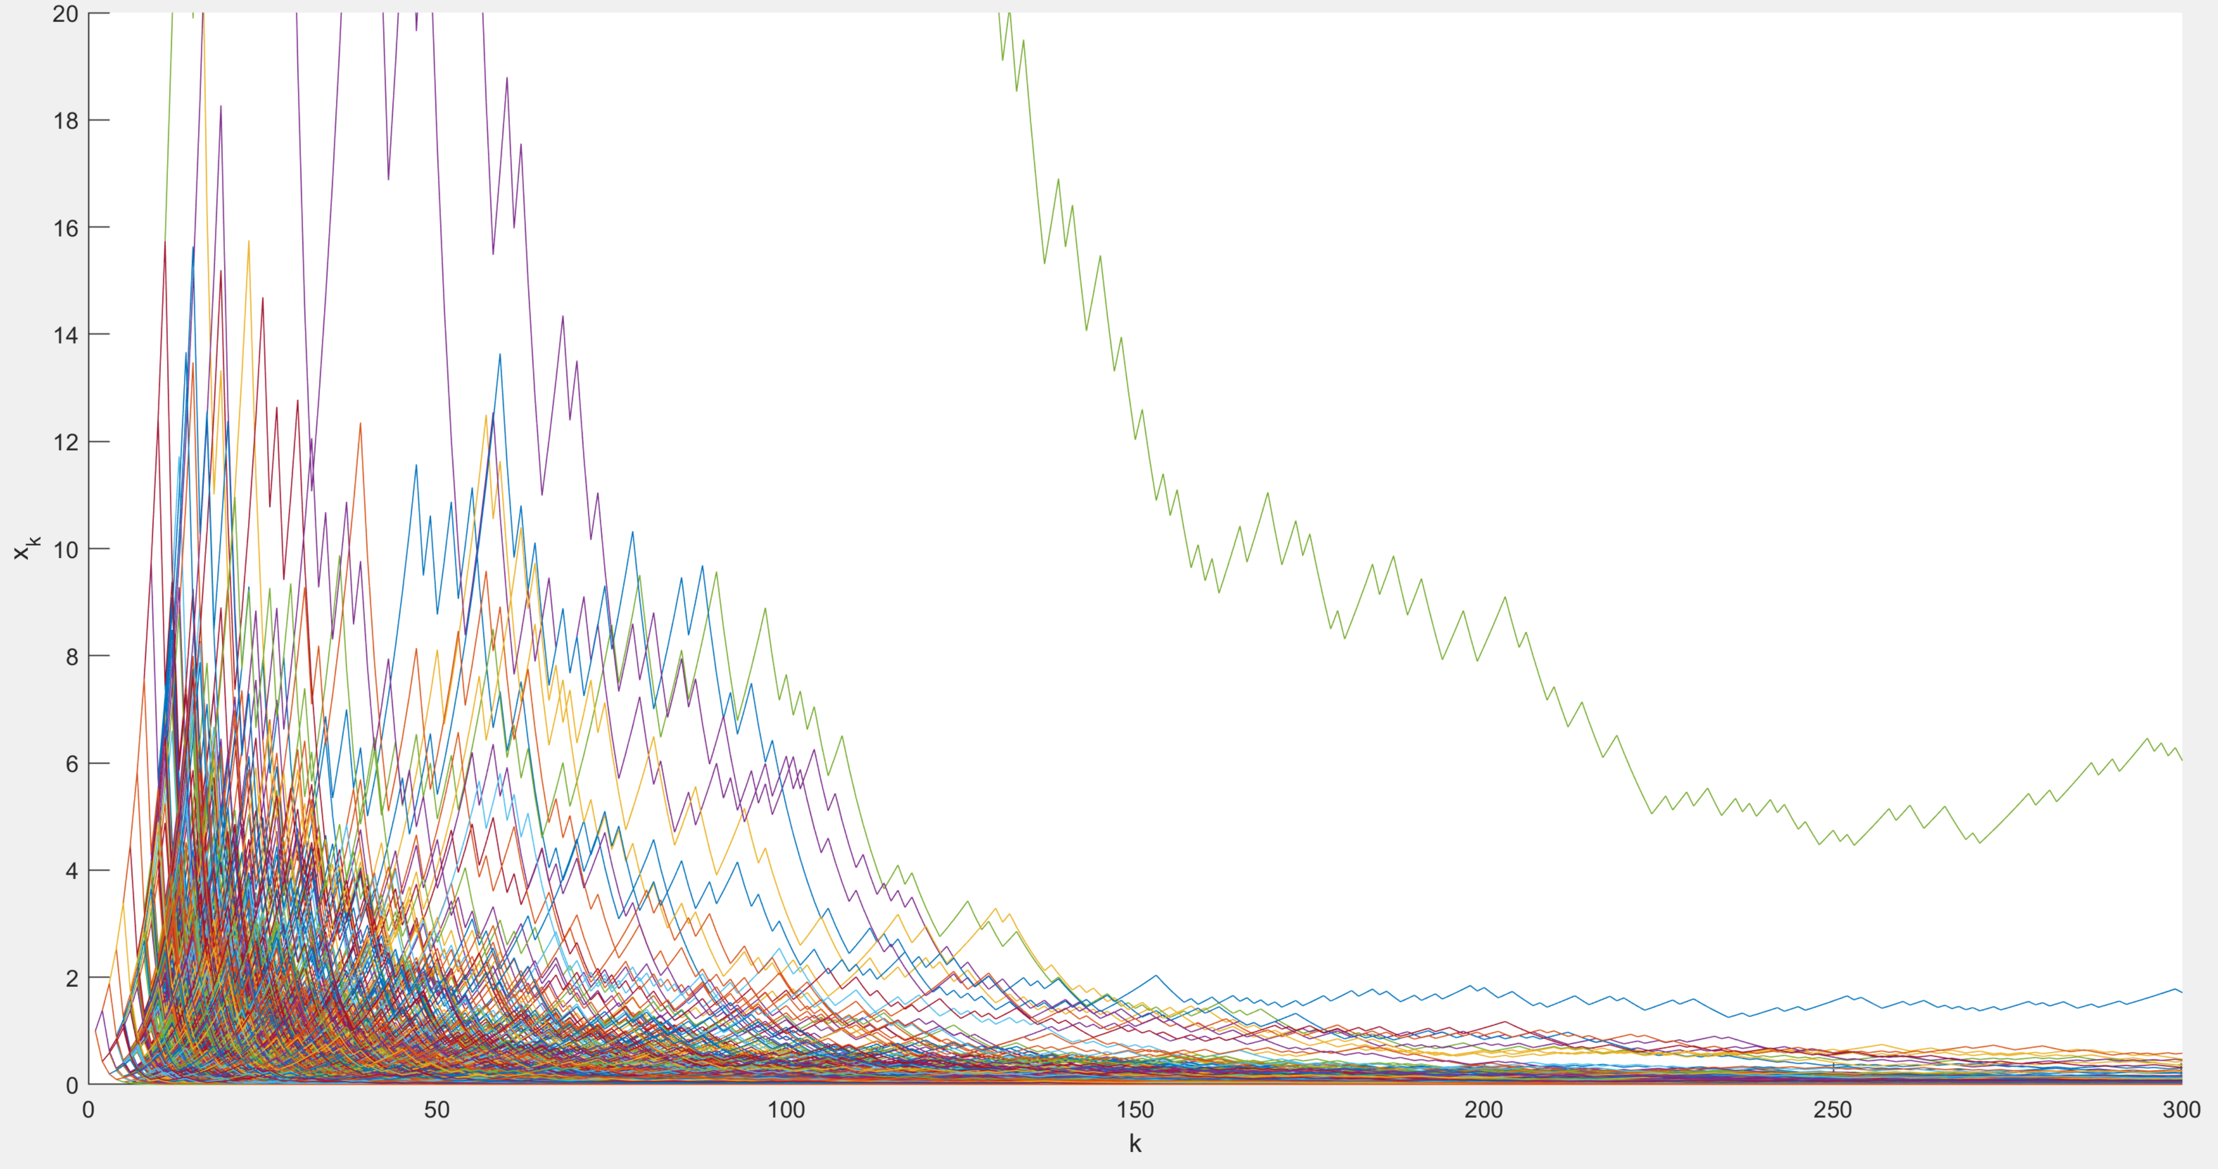

ANY TIME CONCENTRATION

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \frac{\alpha}{k+h}\left(\mathbf{F}(\mathbf{x}_k, \mathbf{Y}_k) + \mathbf{w}_k - \mathbf{x}_k\right)$$

General Norm Contraction: $\left\|\bar{\mathbf{F}}(\mathbf{x}) - \bar{\mathbf{F}}(\mathbf{y})\right\| \quad \leq \gamma \, \|\mathbf{x} - \mathbf{y}\|$

Theorem[Zubeldia, Chen, Maguluri '22]: If α is large enough, for a given K,

$$\mathbb{P}\left(\|\mathbf{x}_k - \mathbf{x}^*\|^2 \leq \begin{cases} \frac{c}{k}\left(1 + \log\left(\frac{1}{\delta}\right) + \log\left(\frac{k+h}{K+h}\right)\right) & if\ k \geq O\left(\log\left(\frac{1}{\delta}\right)\right) \\ k^{\beta} & \text{otherwise} \end{cases} \text{for all } k \geq K \right) \geq (1-\delta)$$

ANY TIME CONCENTRATION

With a small blowup factor of $\log\left(\frac{k+h}{K+h}\right)$, we have bound that is uniform over time

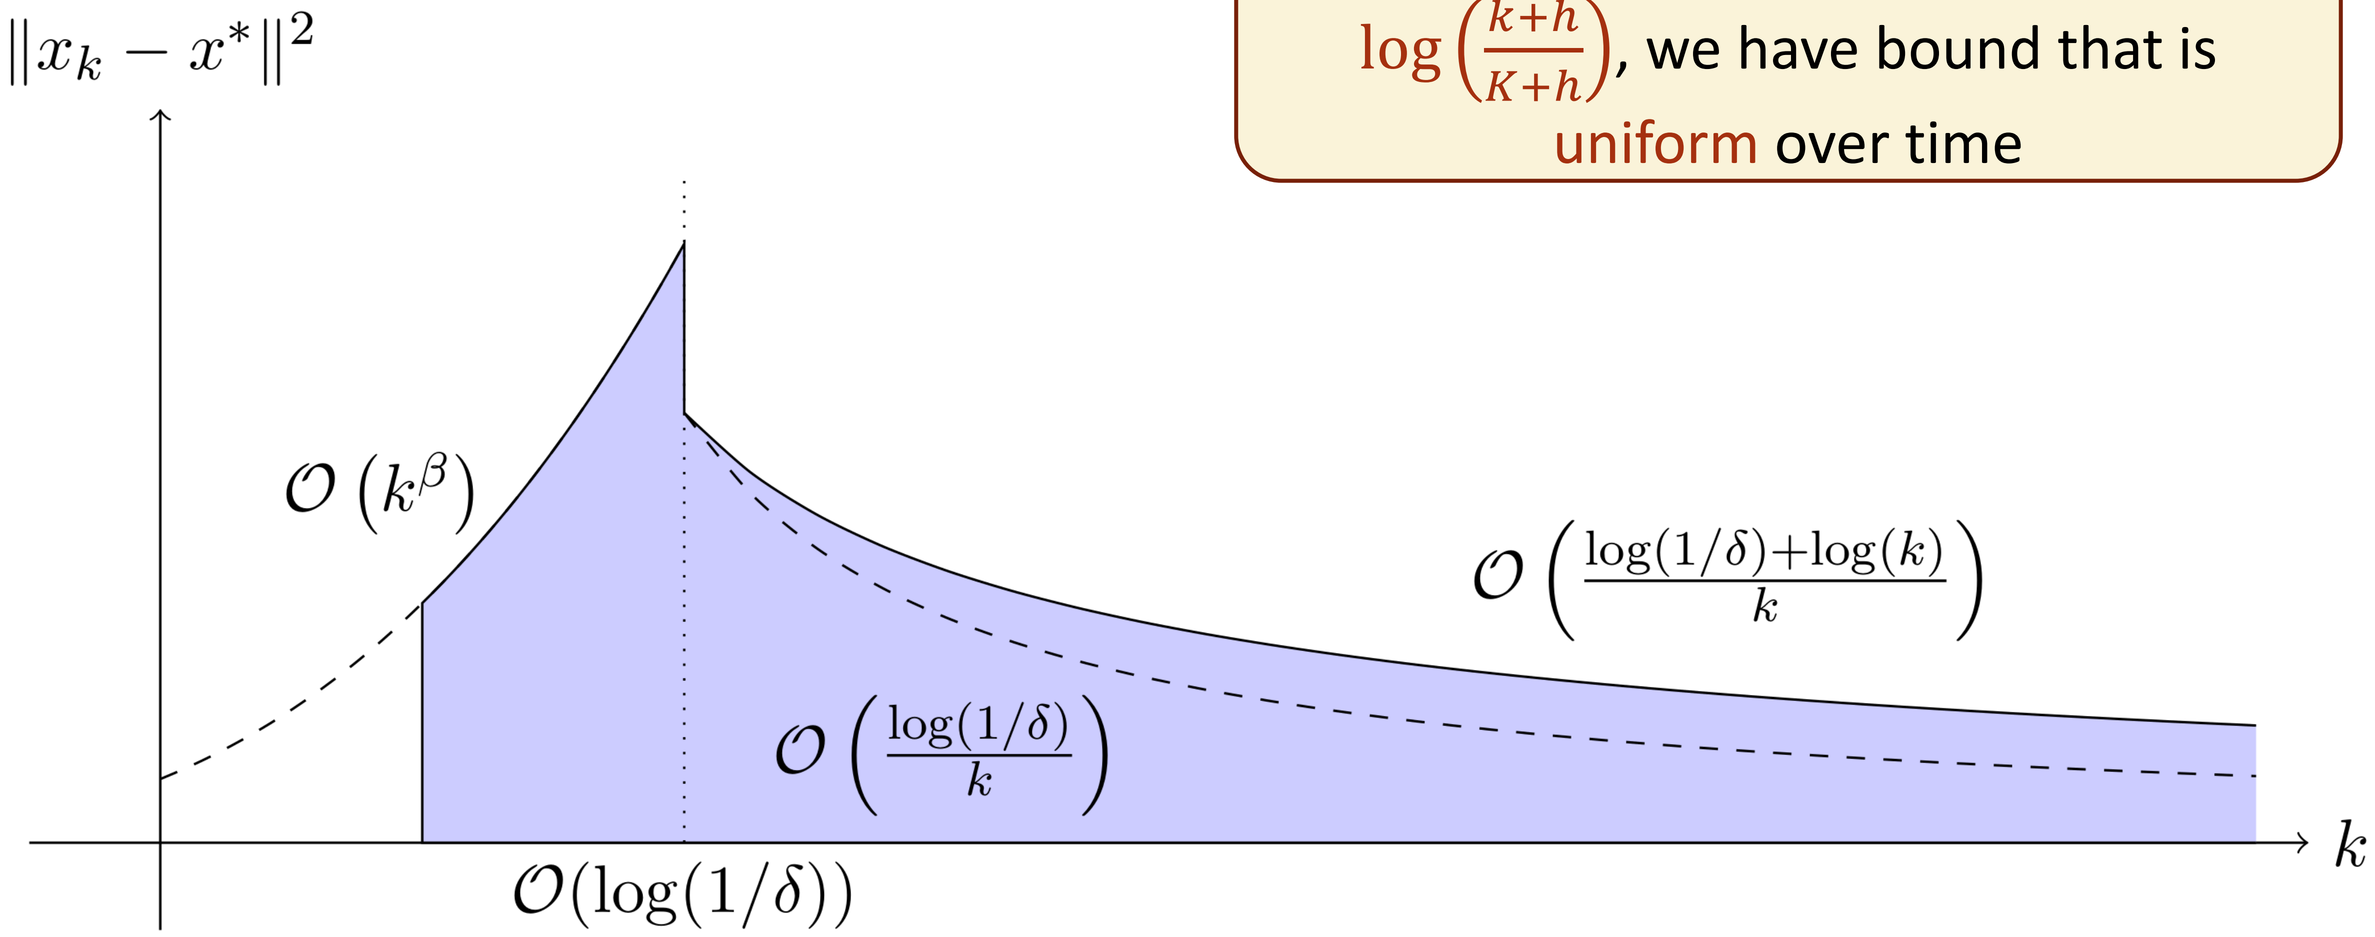

ANY TIME CONCENTRATION

With a small blowup factor of $\log\left(\frac{k+h}{K+h}\right)$, we have bound that is uniform over time

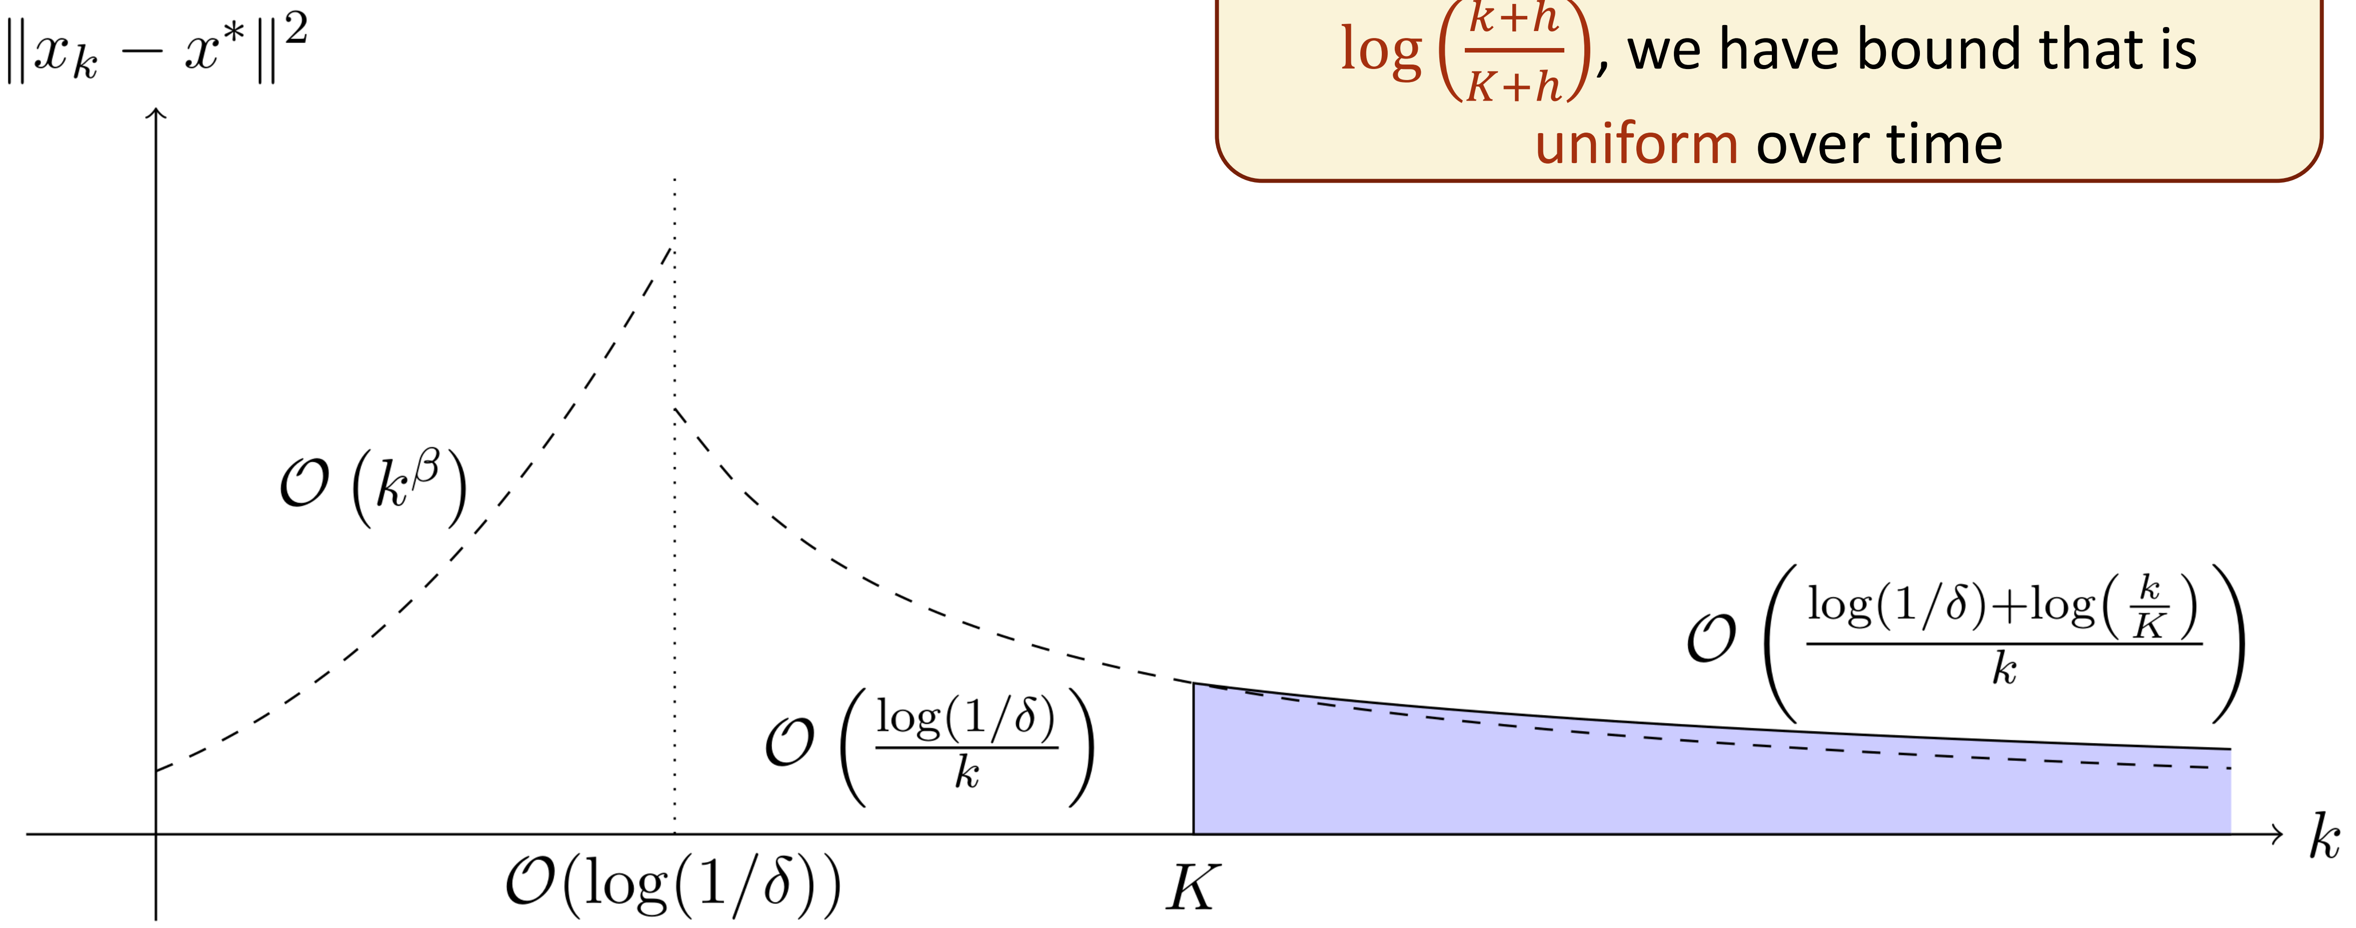

RELATED WORK

- Under boundedness
  - Either due to iterates being in compact set such as constrained optimization [Duchi et al ‘12], [Lan ‘20]
  - Or iterates are bounded due to other structural properties such as in Q Learning, [Evan-Dar et al ‘17], [Li et al ‘21], [Qu et al ‘20] or other related settings [Prashanth et al ‘21] [Thoppe et al ‘19], [Chandak ‘22]

- Constant Step Size that is picked as a function of ϵ and δ by obtaining a bound on just one point (or a window) of the tail
  - [Telgarsky ’22], [Mou et al ‘22], [Li et al ‘21]

- Result needs a bound on the iterates at some time n_0
  - [Thuppe et al ‘19], [Dalal ‘18]

- Our results in contrast, hold for potentially unbounded iterates, with diminishing step sizes and we bound the entire tail, without assuming any future bound.
  - Moreover, we allow for general norm contractions and we get anytime concentration.

PROOF SKETCH

PROOF SKETCH

- **Step 1 - Bounded Case**
  - Develop a proof framework based on Moreau envelope Lyapunov function to get exponential tails at a given time k (assuming the iterates are bounded).
- **Step 2 - Anytime concentration**
  - Generalize the result from Step 1 to get anytime concentration using Supermartingales and Ville's (Doob's) maximal inequality.
- **Step 3 - Bootstrapping**
  - Finally consider the real case of unbounded iterates, and use the previous two steps to inductively bootstrap from the worst case upper bound.

RECALL

Stochastic Approximation to solve $\bar{\mathbf{F}}(\mathbf{x}) = \mathbf{x}$

$$\mathbf{x}_{\mathrm{k+1}} = \mathbf{x}_{\mathrm{k}} + \alpha \quad (\mathbf{F}(\mathbf{x}_{\mathrm{k}}, \mathbf{Y}_k) + \mathbf{w}_{\mathrm{k}} - \mathbf{x}_{\mathrm{k}})$$

Mean Square Bounds:

$$\mathbb{E}[\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\|^2] \leq O\left(\frac{1}{k}\right)$$

Obtained using $\mathrm{M}(\mathbf{x})$ as yapunov function

Using Markov Inequality, we get $\mathbb{P}\left(\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\|^2 \geq O\left(\frac{1}{k}\right) z\right) \leq \frac{1}{z}$

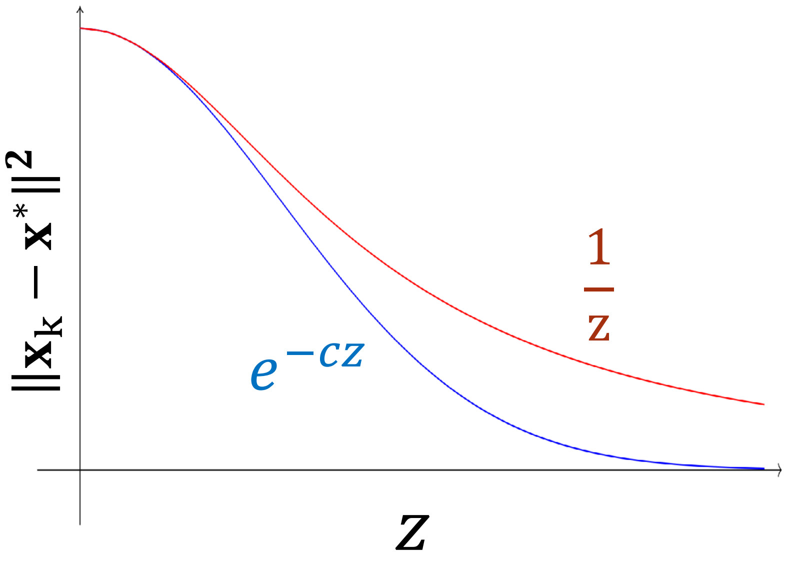

Question: Can we get exponential tail bounds of the form

$$\mathbb{P}\left(\|\mathbf{x}_{\mathrm{k}} - \mathbf{x}^*\|^2 \geq O\left(\frac{1}{k}\right) z\right) \leq e^{-cz}?$$

STEP 1: EXPONENTIAL TAIL BOUNDS

- Use $e^{\mathrm{M}(\mathbf{x})}$ as Lyapunov function to bound $\mathbb{E}\left[e^{\mathrm{M}(\mathbf{x})}\right]$ and obtain tail bounds.
  - Doesn't work – we don't get a recursion

Goal: $\mathbb{P}\left(k\|\mathbf{x}_\mathrm{k} - \mathbf{x}^*\|^2 \geq z\right) \leq e^{-cz}$

- Use $e^{\frac{k\mathrm{M}(\mathbf{x})}{\mathcal{B}}}$ as Lyapunov function to bound $\mathbb{E}\left[e^{\frac{k\mathrm{M}(\mathbf{x}_k)}{\mathcal{B}}}\right]$
  - $\mathcal{B}$ is the bound we assume on the iterates
  - Common Trick: Incorporate the rate into the Lyapunov function
  - It works – We get a recursion (In the bounded case). Solving it, we get

$$\mathbb{E}\left[e^{k\mathrm{M}(\mathbf{x}_k)}\right] \leq ce^{o\left(\frac{1}{k}\right)\mathrm{M}(\mathbf{x}_0)}$$

- Applying Markov inequality, we get the exponential tail bounds.

STEP 2: ANY TIME CONCENTRATION

- Supermartingale - $\mathbb{E}[Z_{k+1}|\mathcal{F}_k] \leq Z_k$

$$\mathbb{P}\left(\sup_{k\geq K} Z_k > z\right) \leq \frac{\mathbb{E}[Z_K]}{z}$$

- Ville's (or Doob's) maximal inequality
- Lyapunov function, $e^{\frac{k\mathrm{M}(\mathbf{x}_k)}{\mathcal{B}}}$ is (almost) decreasing in expectation
  - because we incorporated the rate in it
  - Not quite – need to add a compensator term

$e^{\frac{k\mathrm{M}(\mathbf{x}_k)}{\mathcal{B}} - c\log(k)}$ is a supermartingale

- We get Anytime concentration (still assuming bounded iterates) using the maximal inequality
  - The compensator $\log\left(\frac{k}{K}\right)$ term gives the blowup factor of log in the result

STEP 3: BOOTSTRAPPING

$\mathbf{x}_k \leq \mathcal{B}$ for all k → **Step 1 and Step 2** → $\mathbf{x}_k \leq \tilde{O}\left(\frac{\mathcal{B}}{k}\right)$ for all k whp

When iterates $\mathbf{x}_k$ are not bounded, start with a worst case upper bound $\mathbf{x}_k \leq O\left(k^{\beta}\right)$ for all k

$\mathbf{x}_k \leq \tilde{O}\left(k^{\beta}\right)$ for all k whp → **Step 1 and Step 2** → $\mathbf{x}_k \leq \tilde{O}\left(k^{\beta-1}\right)$for all k whp

⋮

Bootstrap Inductively

Need Anytime Concentration

$\mathbf{x}_k \leq \tilde{O}(1)$ for all k whp → **Step 1 and Step 2** → $\mathbf{x}_k \leq \tilde{O}\left(\frac{1}{k}\right)$ for all k whp.

STEP 3: BOOTSTRAPPING

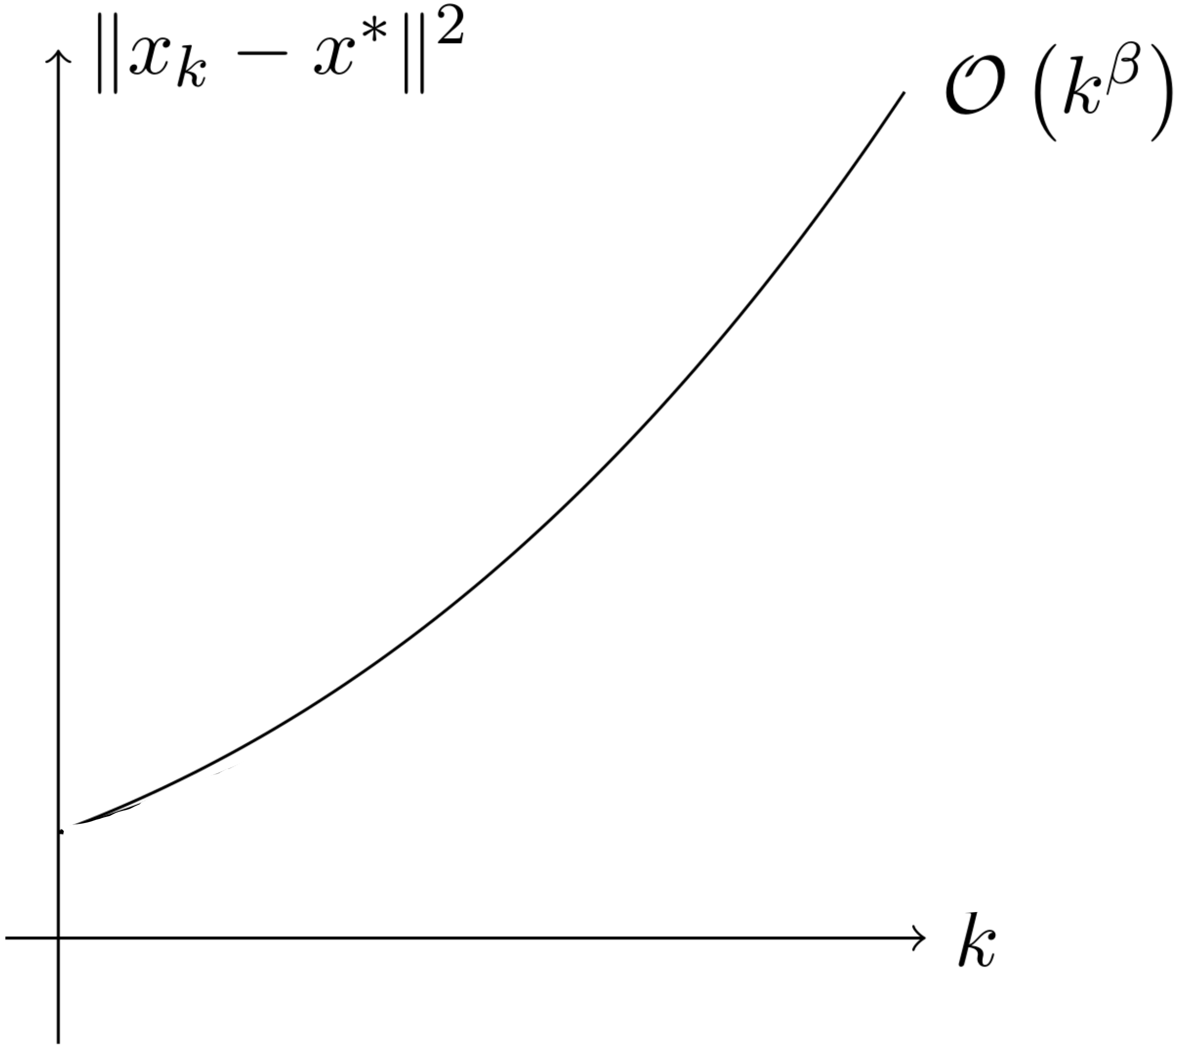

STEP 3: BOOTSTRAPPING

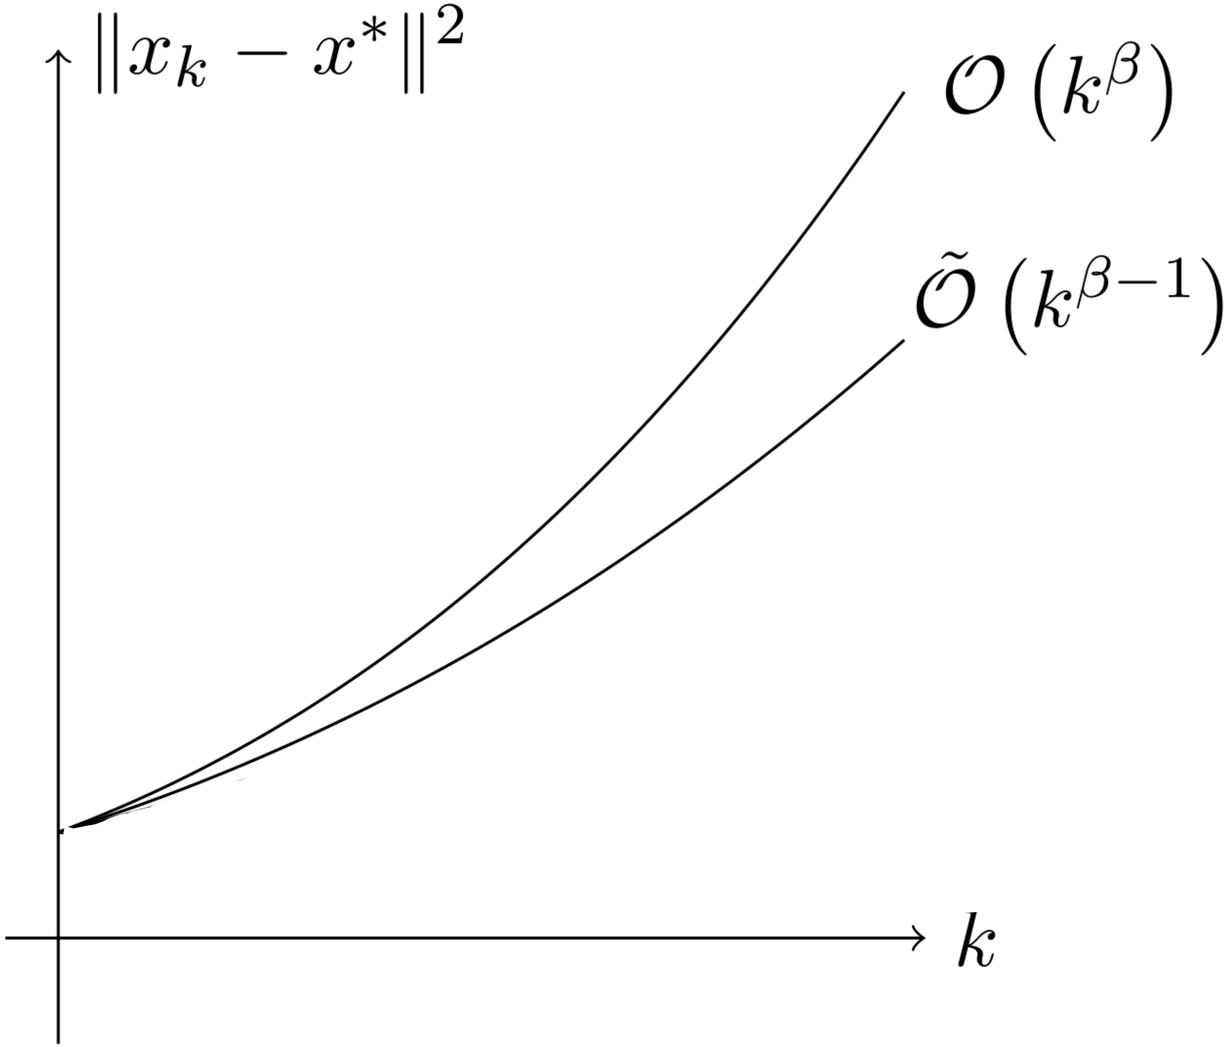

STEP 3: BOOTSTRAPPING

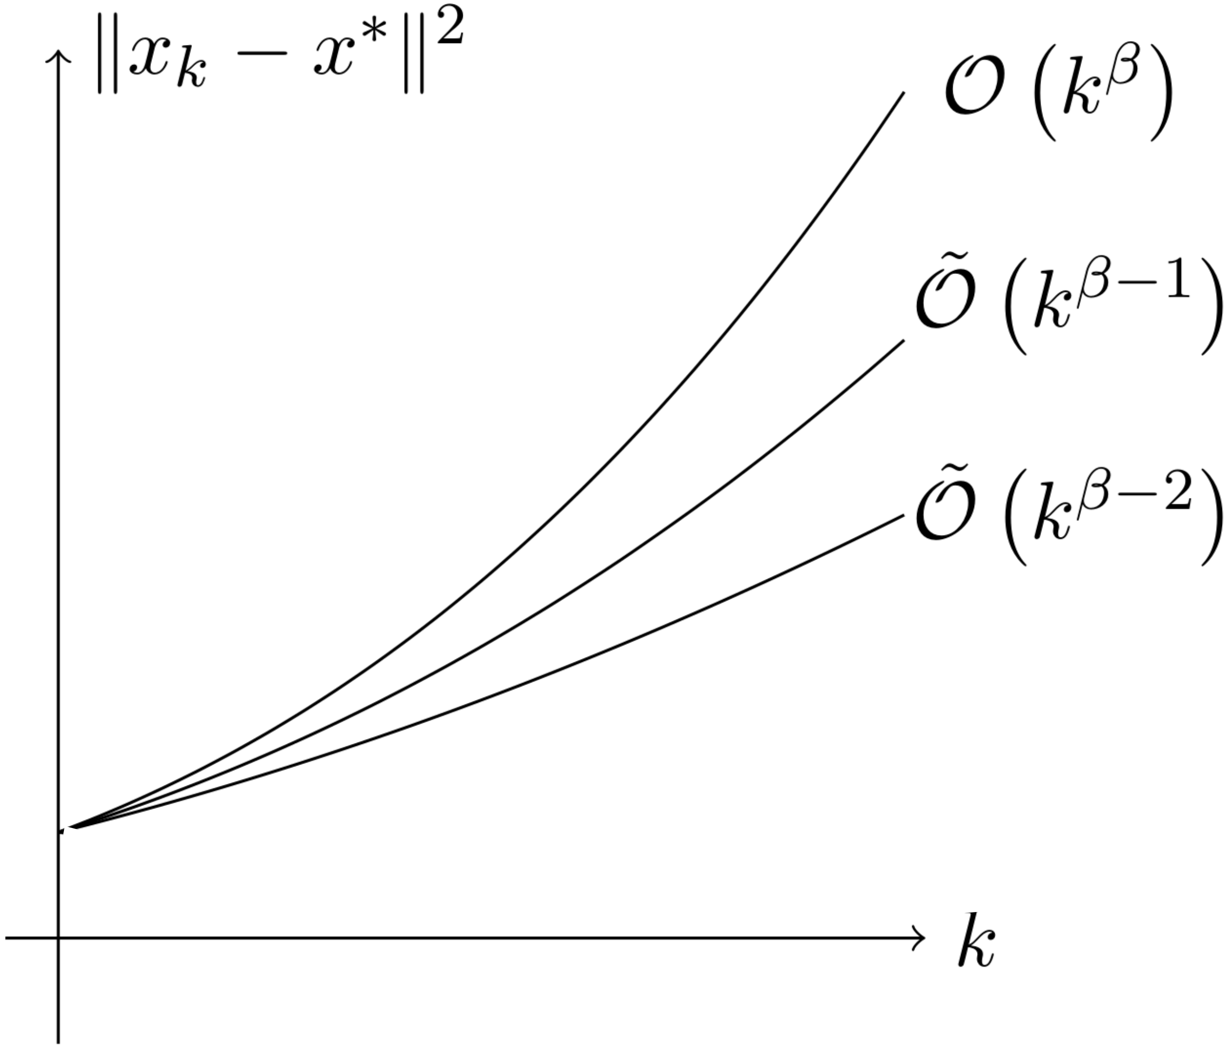

STEP 3: BOOTSTRAPPING

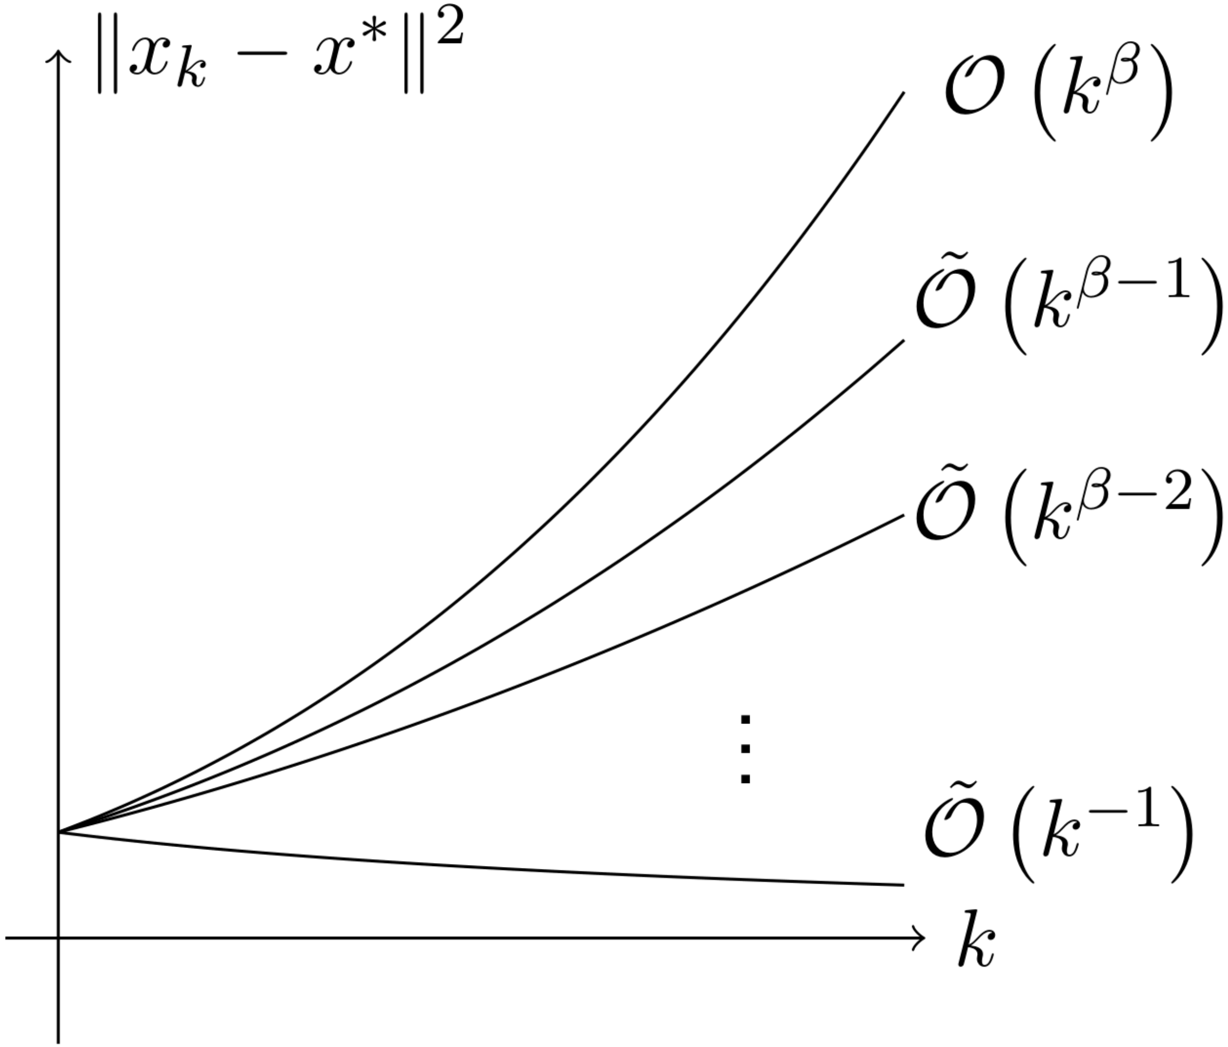

CONCLUSION

- Stochastic Approximation of a contractive operator under general norm
  - Both Additive and Multiplicative Noise

- Mean Square Convergence under Markovian Noise
  - $\tilde{O}\left(\frac{1}{k}\right)$ rate of convergence and $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$ mean square sample complexity
  - Moreau Envelope of the norm square as the Lyapunov function

- Anytime Exponential Concentration under iid Noise
  - $O\left(\frac{1}{k}\right)$ rate Exponential tails and $O\left(\frac{1}{\epsilon^2}\right)\log\left(\frac{1}{\delta}\right)$ sample complexity
  - Proof based on Exponential supermartingales and Bootstrapping

THANK YOU

Questions?