Convergence of Stochastic Approximation Algorithms via Martingale and Converse Lyapunov Methods

M. Vidyasagar FRS

SERB National Science Chair, IIT Hyderabad

RLK Conference, 29 March 2022

Outline

IIT Hyderabad

Problem Formulation – 1

Original problem formulation by Robbins and Monro (1951).

Suppose $\mathbf{f} : \mathbb{R}^d \to \mathbb{R}^d$. The aim is to find a solution to $\mathbf{f}(\boldsymbol{\theta}) = \mathbf{0}$, when *only noisy measurements* of $\mathbf{f}(\cdot)$ are available.

Start with an initial guess $\boldsymbol{\theta}_0 \in \mathbb{R}^d$. At step $t \geq 0$, let

$$\mathbf{y}_{t+1} = \mathbf{f}(\boldsymbol{\theta}_t) + \boldsymbol{\xi}_{t+1},$$

where $\boldsymbol{\xi}_{t+1}$ is the measurement error. Update via

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t \mathbf{y}_{t+1} = \boldsymbol{\theta}_t + \alpha_t (\mathbf{f}(\boldsymbol{\theta}_t) + \boldsymbol{\xi}_{t+1}),$$

where $\{\alpha_t\}_{t \geq 1}$ is a predetermined sequence of step sizes.

Question: When does $\boldsymbol{\theta}_t \to \boldsymbol{\theta}^*$, where $\mathbf{f}(\boldsymbol{\theta}^*) = \mathbf{0}$?

Problem Formulation – 2

Modified problem formulation by Keifer-Wolfowitz (1952) and Blum (1954).

Suppose $J : \mathbb{R}^d \to \mathbb{R}$ is $\mathcal{C}^1$ (or $\mathcal{C}^2$). We wish to find a stationary point of J, i.e., a solution of $\nabla J(\boldsymbol{\theta}) = \mathbf{0}$.

If measurements are

$$\mathbf{y}_{t+1} = -\nabla J(\boldsymbol{\theta}_t) + \boldsymbol{\xi}_{t+1},$$

then this is the same problem as earlier.

Instead, define

$$y_{t+1,i} = \frac{[J(\boldsymbol{\theta}_t - c_t \mathbf{e}_i) + \xi^-_{t+1,i}] - [J(\boldsymbol{\theta}_t + c_t \mathbf{e}_i) + \xi^+_{t+1,i}]}{2c_t}, i = 1, \cdots, d.$$

Problem Formulation – 2 (Cont'd)

Elaboration: At iteration t, perturb *each coordinate* of $\boldsymbol{\theta}_t$ by a predetermined increment c_t, and measure $J(\cdot)$ at each of the $2d$ vectors (subject to measurement error).

Using these $2d$ measurements, *approximate* the gradient $\nabla J(\boldsymbol{\theta}_t)$ using a first-order approximation. Use this to update $\boldsymbol{\theta}_t$ to $\boldsymbol{\theta}_{t+1}$.

Interpretation: Rewrite as

$$y_{t+1,i} = \frac{J(\boldsymbol{\theta}_t - c_t \mathbf{e}_i) - J(\boldsymbol{\theta}_t + c_t \mathbf{e}_i)}{2c_t} + \frac{\xi^-_{t+1,i} - \xi^+_{t+1,i}}{2c_t},$$

Then as $c_t \to 0$ as $t \to \infty$, the first term approaches $-[\nabla J(\boldsymbol{\theta}_t)]_i$. But the variance of the second (noise) term approaches infinity!

Standard Sufficient Conditions

Robbins-Monro Conditions:

$$\sum_{t=0}^{\infty} \alpha_t^2 < \infty, \sum_{t=0}^{\infty} \alpha_t = \infty.$$

Blum Conditions:

$$c_t \to 0 \text{ as } t \to \infty, \sum_{t=0}^{\infty} \left(\frac{\alpha_t}{c_t} \right)^2 < \infty, \sum_{t=0}^{\infty} \alpha_t c_t < \infty, \sum_{t=0}^{\infty} \alpha_t = \infty.$$

My Observation: The condition $c_t \to 0$ as $t \to \infty$ is (more or less) implied by the rest.

If $\alpha_t = (t+1)^{-p}$ and $c_t = (t+1)^{-r}$. then Robbins-Monro conditions are $0.5 < p \leq 1$. and Blum conditions are $p + r < 1$ and $p - r > 1/2$.

Outline

IIT Hyderabad

Some Standard Assumptions

(F). $\boldsymbol{\theta}^*$ is the unique solution of $\mathbf{f}(\boldsymbol{\theta}) = \mathbf{0}$.

(N). Define $\boldsymbol{\theta}_0^t = \{\boldsymbol{\theta}_0, \cdots, \boldsymbol{\theta}_t\}$, and let $\mathcal{F}_t = \sigma(\boldsymbol{\theta}_0^t, \boldsymbol{\xi}_1^t)$. Then (i) the measurements are unbiased, i.e.,

$$E(\boldsymbol{\xi}_{t+1} | \mathcal{F}_t) = \mathbf{0} \text{ a.s.},$$

and (ii) the conditional variance grows quadratically, i.e., $\exists d < \infty$ such that

$$E(\|\boldsymbol{\xi}_{t+1}\|_2^2 | \mathcal{F}_t) \leq d(1 + \|\boldsymbol{\theta}_t\|_2^2).$$

(S). Robbins-Monro (RM) conditions:

$$\sum_{t=0}^{\infty} \alpha_t = \infty, \sum_{t=0}^{\infty} \alpha_t^2 < \infty.$$

A Typical Theorem

Theorem

Suppose (F), (N), and (S) hold. If $\mathbf{f}(\cdot)$ satisfies some more conditions, and if the iterates $\{\boldsymbol{\theta}_t\}$ are bounded almost surely, then $\boldsymbol{\theta}_t \to \boldsymbol{\theta}^*$, a.s. as $t \to \infty$.

Almost sure boundedness of the iterates ("stability") is a part of the hypothesis, not a conclusion.

Thus stability plus other conditions imply convergence.

Question: Can the stability of the iterates be made a *conclusion, instead of being a part of the hypotheses?*

Borkar-Meyn Theorem (2000)

Assumptions:

- All the standard assumptions (F), (N), (S).
- $\mathbf{f}(\cdot)$ is globally Lipschitz continuous, i.e., $\exists L < \infty$ such that

$$\|\mathbf{f}(\boldsymbol{\theta}) - \mathbf{f}(\boldsymbol{\phi})\|_2 \leq L\|\boldsymbol{\theta} - \boldsymbol{\phi}\|_2, \ \forall \boldsymbol{\theta}, \boldsymbol{\phi} \in \mathbb{R}^d.$$

- There is a "limit function" $\mathbf{f}_\infty$ such that

$$\frac{\mathbf{f}(r\boldsymbol{\theta})}{r} \to \mathbf{f}_\infty(\boldsymbol{\theta}) \text{ as } r \to \infty,$$

  uniformly over compact subsets of $\mathbb{R}^d$.
- $\mathbf{0}$ is a globally exponentially stable equilibrium of

$$\dot{\boldsymbol{\theta}} = \mathbf{f}_\infty(\boldsymbol{\theta}).$$

Borkar-Meyn Theorem (2000) – Cont'd

Elaboration:

- Global Lipschitz continuity of $\mathbf{f}(\cdot)$ implies that there is a function $\mathbf{s} : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^d$ such that $\mathbf{s}(\cdot, \boldsymbol{\theta})$ is the unique solution of

$$\frac{d\mathbf{s}(t, \boldsymbol{\theta})}{dt} = \mathbf{f}(\mathbf{s}(t, \boldsymbol{\theta})), \mathbf{s}(0, \boldsymbol{\theta}) = \boldsymbol{\theta}.$$

- The equilibrium $\boldsymbol{\theta}^*$ is **globally exponentially stable (GES)** if there exist constants $\mu \geq 1, \gamma > 0$ such that

$$\|\mathbf{s}(t, \boldsymbol{\theta}) - \boldsymbol{\theta}^*\|_2 \leq \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \exp(-\gamma t), \ \forall t \geq 0, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

Borkar-Meyn Theorem (2000) – Cont'd

Theorem

Under the stated assumptions,

1. $\{\boldsymbol{\theta}_t\}$ *is bounded almost surely.*
2. $\boldsymbol{\theta}_t \to \boldsymbol{\theta}^*$ *as* $t \to \infty$.

The a.s. boundedness of $\{\boldsymbol{\theta}_t\}$ is a *conclusion*, not a hypothesis.

Proof is based on the ODE method, which states that the sample paths of the iterates "converge" to the *deterministic* solution trajectories of the ODE $\dot{\boldsymbol{\theta}} = \mathbf{f}_\infty(\boldsymbol{\theta})$.

Method pioneered by Ljung (1974), Deveritskii and Fradkov (1974), Kushner-Clark (1978); see also Métivier-Priouret (1984).

Rather technical – worthwhile to find an easier proof.

Gladyshev's Theorem (1965)

Theorem

Assumptions (F), (N), but not (S). In addition

$$\inf_{\epsilon < \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 < 1/\epsilon} \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \mathbf{f}(\boldsymbol{\theta}) \rangle < 0, 0 < \epsilon < 1,$$

$$\|\mathbf{f}(\boldsymbol{\theta})\|_2 \leq K \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2, K < \infty.$$

Then

1. *If $\sum_{t=0}^{\infty} \alpha_t^2 < \infty$, then $\{\boldsymbol{\theta}_t\}$ is bounded almost surely.*
2. *If in addition $\sum_{t=0}^{\infty} \alpha_t = \infty$, then $\boldsymbol{\theta}_t \to \boldsymbol{\theta}^*$ almost surely as $t \to \infty$.*

If $\mathbf{f}(\cdot)$ is continuous, the above is equivalent to

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \mathbf{f}(\boldsymbol{\theta}) \rangle < 0, \ \forall \boldsymbol{\theta} \neq \boldsymbol{\theta}^*.$$

Aspects of Gladyshev's Theorem

- Approach: Define $Z_t = a_t \|\boldsymbol{\theta}_t\|_2^2 + b_t$, and choose a_t, b_t such that $\{Z_t\}$ is a nonnegative supermartingale.
- Far less "technical" than the ODE method.
- Clear "division of labor": Square-summability of step sizes gives stability, and divergence of step sizes gives convergence.

Restriction: Applies *only to* "passive" functions.

My Motivation: Can this approach be extended *more general* functions?

Yes, by using "converse" Lyapunov theory (topic of this lecture).

But first, a general "framework" for proving the convergence of SA.

Outline

IIT Hyderabad

Lemma on Almost Sure Boundedness

Lemma

Suppose (Ω, Σ, P) *is a probability space,* $\{\mathcal{F}_t\}$ *is a filtration, and* $\{X_t\}$ *is a nonnegative-valued stochastic process that is adapted to* $\{\mathcal{F}_t\}$*. Suppose* $\{w_t\}$ *is a summable sequence of positive numbers, and that*

$$E(X_{t+1}|\mathcal{F}_t) \leq (1 + w_t)X_t + w_t \text{ a.s., } \forall t \geq 0.$$

Then

1. $\{X_t\}$ *is bounded almost surely, and*
2. *There is a random variable* ζ *such that* $X_t \to \zeta$ *almost surely as* $t \to \infty$*.*

Lemma on Almost Sure Convergence

Lemma

Let (Ω, Σ, p) *and* $\{\mathcal{F}_t\}$ *be as in Lemma 4. Suppose* $\{w_t\}, \{u_t\}$ *are sequences of positive numbers such that*

$$\sum_{t=0}^{\infty} w_t < \infty, \sum_{t=0}^{\infty} u_t = \infty.$$

Suppose $\{X_t\}$ *is a nonnegative stochastic process that is adapted to* $\{\mathcal{F}_t\}$*, and that satisfies*

$$E(X_{t+1}|\mathcal{F}_t) \leq (1 + w_t - u_t)X_t + w_t, \text{ a.s., } \forall t \geq 0.$$

Then $\{X_t\}$ *is bounded almost surely, and converges almost surely to* 0 *as* $t \to \infty$*.*

History of the "Framework" Lemmas

- Robbins and Siegmund (1971) proved the first lemma and a version of the second lemma. (They had sequences and not random variables.)
- They were unaware of Gladyshev's 1965 paper.
- I was motivated to extract the "essence" of Gladyshev's paper.
- I proved the two lemmas independently of Robbins-Siegmund.

IIT Hyderabad

Outline

IIT Hyderabad

Rate of Change Function

Consider the ODE $\dot{\boldsymbol{\theta}} = \mathbf{f}(\boldsymbol{\theta})$ and suppose $V : \mathbb{R}^d \to \mathbb{R}$ is $\mathcal{C}^1$. Then the function $\dot{V} : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$\dot{V}(\boldsymbol{\theta}) = \langle \nabla V(\boldsymbol{\theta}), \mathbf{f}(\boldsymbol{\theta}) \rangle.$$

Suppose $\mathbf{s}(t, \boldsymbol{\theta})$ satisfies

$$\frac{d\mathbf{s}(t, \boldsymbol{\theta})}{dt} = \mathbf{f}(\mathbf{s}(t, \boldsymbol{\theta})), \mathbf{s}(0, \boldsymbol{\theta}) = \boldsymbol{\theta}.$$

Then

$$\frac{dV(\mathbf{s}(t, \boldsymbol{\theta}))}{dt} = \dot{V}(\mathbf{s}(t, \boldsymbol{\theta}))$$

is the *rate of change* of V along the solution trajectories.

Forward vs. Converse Lyapunov Theory

Given the the ODE $\dot{\boldsymbol{\theta}} = \mathbf{f}(\boldsymbol{\theta})$:

"Forward" Lyapunov theory: If there exists a function V with certain properties, then $\boldsymbol{\theta}^*$ has certain stability properties.

"Converse" Lyapunov theory: If the equilibrium $\boldsymbol{\theta}^*$ has certain stability properties, then there exists a function V with certain properties.

Global Exponential Stability: Reprise

Suppose $\mathbf{f}$ is globally Lipschitz continuous, and define $\mathbf{s} : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^d$ via: $\mathbf{s}(t, \boldsymbol{\theta})$ is the unique solution of

$$\frac{d\mathbf{s}(t, \boldsymbol{\theta})}{dt} = \mathbf{f}(\mathbf{s}(t, \boldsymbol{\theta})), \mathbf{s}(0, \boldsymbol{\theta}) = \boldsymbol{\theta}.$$

Suppose $\mathbf{f}(\boldsymbol{\theta}^*) = \mathbf{0}$. The equilibrium $\boldsymbol{\theta}^*$ is **globally exponentially stable (GES)** if there exist $\mu < \infty, \gamma > 0$ such that

$$\|\mathbf{s}(t, \boldsymbol{\theta}) - \boldsymbol{\theta}^*\|_2 \leq \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \exp(-\gamma t), \ \forall t \geq 0, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

Standard Forward Lyapunov Theorem for GES

Theorem

Suppose $\mathbf{f}$ *is globally Lipschitz continuous, and that* $\mathbf{f}(\boldsymbol{\theta}^*) = \mathbf{0}$. *Suppose there exists a* $\mathcal{C}^1$ *function* $V : \mathbb{R}^d \to \mathbb{R}_+$ *that satisfies the following: There exist* $c_1, c_2, c_3 > 0$ *such that*

$$c_1 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \leq V(\boldsymbol{\theta}) \leq c_2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2,$$

$$\dot{V}(\boldsymbol{\theta}) \leq -c_3 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

Then $\boldsymbol{\theta}^*$ *is a GES equilibrium.*

Standard Converse Lyapunov Theorem for GES

Theorem

Suppose $\mathbf{f}$ *is globally Lipschitz continuous, that* $\boldsymbol{\theta}^*$ *is a GES equilibrium. There exists a* $\mathcal{C}^1$ *function* $V : \mathbb{R}^d \to \mathbb{R}_+$ *that satisfies the following: There exist* $c_1, c_2, c_3 > 0$ *such that*

$$c_1 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \leq V(\boldsymbol{\theta}) \leq c_2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2,$$

$$\dot{V}(\boldsymbol{\theta}) \leq -c_3 \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

Need for a New Converse Theorem

Common choice:

$$V(\boldsymbol{\theta}) := \int_0^\infty \|\mathbf{s}(t, \boldsymbol{\theta})\|_2^2 dt.$$

This is *not good enough* for current application. We require a V function with *globally bounded Hessian.*

Such a theorem has been proved, building on earlier work of Corless and Glielmo (1998).

New Converse Lyapunov Theorem for GES

Theorem

Suppose in addition that $\mathbf{f} \in \mathcal{C}^2$, *and that*[a]

$$\sup_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\nabla^2 f_i(\boldsymbol{\theta})\|_S \cdot \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 < \infty, \ \forall i \in [d].$$

Choose

$$0 < \kappa < \gamma, \frac{\ln \mu}{\gamma - \kappa} \leq T < \infty, V(\boldsymbol{\theta}) := \int_0^T e^{\kappa \tau} \|\mathbf{s}(\tau, \boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \, d\tau$$

Then V *is* $\mathcal{C}^2$, *and also satisfies*

$$\|\nabla^2 V(\boldsymbol{\theta})\|_S \leq 2M, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

[a] Here $\|\cdot\|_S$ denotes the spectral norm, and $[d] = \{1, \ldots, d\}$.

Outline

IIT Hyderabad

Convergence Theorem for Robbins-Monro Formulation

Theorem

Suppose (i) $\boldsymbol{\theta}^*$ is the only zero of $\mathbf{f}(\cdot)$, (ii) $\boldsymbol{\theta}^*$ is a GES equilibrium of $\dot{\boldsymbol{\theta}} = \mathbf{f}(\boldsymbol{\theta})$, (iii) $\mathbf{f}(\cdot)$ is globally Lipschitz continuous, and (iv)

$$\sup_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\nabla^2 f_i(\boldsymbol{\theta})\|_S \cdot \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 < \infty, \ \forall i \in [d].$$

Suppose further that $\{\boldsymbol{\xi}_t\}$ satisfies (N). Then

1. *If $\sum_{t=0}^{\infty} \alpha_t^2 < \infty$, then $\{\boldsymbol{\theta}_t\}$ is bounded almost surely.*
2. *If in addition $\sum_{t=0}^{\infty} \alpha_t = \infty$, then $\boldsymbol{\theta}_t \to \boldsymbol{\theta}^*$ almost surely as $t \to \infty$.*

IIT Hyderabad

Comparison with Borkar and Meyn (2000)

Nice "division of labor" as in Gladyshev (1965). Also we don't need the existence of

$$\mathbf{f}_\infty := \lim_{r \to \infty} \mathbf{f}(r\boldsymbol{\theta})/r.$$

But Borkar-Meyn (2000) don't need

$$\sup_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\nabla^2 f_i(\boldsymbol{\theta})\|_S \cdot \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 < \infty, \ \forall i \in [d].$$

However, $\mathbf{f}_\infty$ exists *and has a globally bounded Hessian*, then above condition holds.

So our theorem (more or less) contains Borkar-Meyn (2000) as a special case. (Hat tip: Sean Meyn).

Kiefer-Wolfowitz-Blum Formulation: Reprise

Objective: Find a stationary point of a $\mathcal{C}^2$ function $J : \mathbb{R}^d \to \mathbb{R}$.

Measurements:

$$y_{t+1,i} = \frac{[J(\boldsymbol{\theta}_t - c_t \mathbf{e}_i) + \xi^-_{t+1,i}] - [J(\boldsymbol{\theta}_t + c_t \mathbf{e}_i) + \xi^+_{t+1,i}]}{2c_t}, i = 1, \cdots, d.$$

Updates:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t \mathbf{y}_{t+1}.$$

IIT Hyderabad

KW-B Formulation: Assumptions on Objectve Function

- $J(\cdot)$ is in $\mathcal{C}^2$, and has a unique global minimizer $\boldsymbol{\theta}^*$, which can be taken as $\mathbf{0}$ by translating coordinates.
- There is a constant $h > 0$ such that

$$\langle \boldsymbol{\theta}, \nabla J(\boldsymbol{\theta}) \rangle \geq h \|\boldsymbol{\theta}\|_2^2, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

- $J(\cdot)$ is Lipschitz continuous at $\mathbf{0}$ with constant L. Hence

$$\|\nabla J(\boldsymbol{\theta})\|_2 \leq L \|\boldsymbol{\theta}\|_2, \ \forall \boldsymbol{\theta} \in \mathbb{R}^d.$$

- The Hessian $\nabla^2 J(\cdot)$ is globally bounded.

KW-B Formulation: Assumptions on Noise

Let $\mathcal{F}_t$ denote the σ-algebra generated by $\boldsymbol{\theta}_0^t$, $\boldsymbol{\xi}_{+,1}^{+,t}$ and $\boldsymbol{\xi}_{-,1}^{-,t}$. Then

- We have that

$$E(\xi_{t+1,i}^{+}|\mathcal{F}_t) = 0, E(\xi_{t+1,i}^{-}|\mathcal{F}_t) = 0, \text{ a.s.}, \ \forall i \in [d], \ \forall t \geq 0.$$

- Suppose that $\boldsymbol{\xi}_t^+$ and $\boldsymbol{\xi}_t^-$ are independent for each t, and define $\boldsymbol{\xi}_t = \boldsymbol{\xi}_t^+ + \boldsymbol{\xi}_t^-$. There there is a finite constant d such that

$$E(\|\boldsymbol{\xi}_{t+1}\|_2^2|\mathcal{F}_t) \leq d(1 + \|\boldsymbol{\theta}_t\|_2^2), \ \forall t.$$

Convergence Theorem for KW-B Formulation

Theorem

Suppose $J : \mathbb{R}^d \to \mathbb{R}$ and the noise sequences satisfy the above assumptions. Under these conditions:

1. *If the sequences $\{\alpha_t\}, \{c_t\}$ together satisfy*

$$\sum_{t=0}^{\infty} \left(\frac{\alpha_t}{c_t} \right)^2 < \infty, \sum_{t=0}^{\infty} \alpha_t c_t < \infty,$$

   then $\{\boldsymbol{\theta}_t\}$ is bounded almost surely.

2. *If, in addition, $\sum_{t=0}^{\infty} \alpha_t = \infty$, then $\boldsymbol{\theta}_t \to \mathbf{0}$ almost surely as $t \to \infty$.*

Future Direction – 1

We still haven't replicated the original Gladyshev theorem!

Consider for example a function $f : \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$f(\theta) = \begin{cases} -\theta, & -1 \leq \theta \leq 0, \\ -1, & \theta \leq -1, \end{cases}$$

with $f(\theta) = -f(\theta)$ when $\theta \geq 0$. Then $f(\cdot)$ satisfies $\theta f(\theta) < 0$ and $|f(\theta)| \leq |\theta|$ for all $\theta \neq 0$, and convergence of SA follows from Gladyshev's theorem.

However, the equilibrium $\theta = 0$ is only globally *asymptotically* stable but not globally *exponentially* stable. So current theory doesn't apply.

Work is under way to extend the Lyapunov approach to cover this situation.

Future Direction – 2

In many applications in control theory, the assumption

$$E(\boldsymbol{\xi}_{t+1} | \mathcal{F}_t) = 0$$

does not hold. Thus the measurement

$$\mathbf{y}_{t+1} = \mathbf{f}(\boldsymbol{\theta}_t) + \boldsymbol{\xi}_{t+1}$$

could be *biased!*

Ljung (1978) has addressed this when $\mathbf{f}(\boldsymbol{\theta}) = -\nabla J(\boldsymbol{\theta})$, i.e., a gradient vector field. But no general approach as yet.

Joint Work with RLK on Reinforcement Learning

- BASA (Batch Asynchronous Stochastic Approximation)
  - At eath time t update *some but not all* components of $\boldsymbol{\theta}_t$.
  - The choice of components to be updated can be driven by another process.
  - Includes well-known "temporal difference" algorithm.
- Stochastic approximation with random step sizes.
  - Easy to analyze if step size α_t is independent of the measurement noise.
  - This is *not the case* in Reinforcement Learning.
- Paper is ready.

Hearty Congratulations!

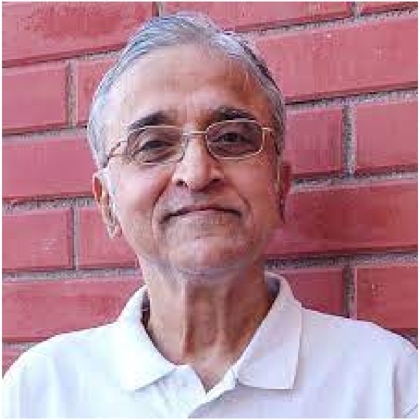

Remember: Life Begins at 65

IIT Hyderabad

Thank You!

IIT Hyderabad