Lecture 6 - *Convergence of Stochastic Approximation*

OBJECTIVE: In this lecture we introduce the powerful theoretical concepts of **Lyapunov (potential) functions** and **martingales**. These are used to prove convergence of the basic stochastic approximation algorithm introduced in the previous lecture.

◇ FIXED POINT EQUATION WITH OPERATORS

For a generic operator F, which could be Bellman's or simply the conditional expectation operator, the fixed point equation at the optimum $\boldsymbol{\theta}^\star$ takes the form:

$$F\boldsymbol{\theta}^\star = \boldsymbol{\theta}^\star$$

As before, $\boldsymbol{\theta}$ can be estimated using the recursive SA algorithm:

$$\begin{aligned}
\boldsymbol{\theta}^{(t+1)} &= \boldsymbol{\theta}^{(t)} + \alpha^{(t)} \left[F\boldsymbol{\theta}^{(t)} - \boldsymbol{\theta}^{(t)}\right] \\
&= \boldsymbol{\theta}^{(t)} + \alpha^{(t)} \left[(F\boldsymbol{\theta}^{(t)} - \boldsymbol{\theta}^\star) + (\boldsymbol{\theta}^\star - \boldsymbol{\theta}^{(t)})\right] \\
&= \boldsymbol{\theta}^{(t)} + \alpha^{(t)} \left[w^{(t)} + (\boldsymbol{\theta}^\star - \boldsymbol{\theta}^{(t)})\right] \\
&= \boldsymbol{\theta}^{(t)} + \alpha^{(t)} s(w^{(t)}, \boldsymbol{\theta}^{(t)})
\end{aligned}$$

where $w^{(t)}$ is the Monte Carlo error resulting from using only one sample to approximate the expectation. The random variable $s(w^{(t)}, \boldsymbol{\theta}^{(t)})$ is the direction of update (descent).

For later developments, we need to introduce a variable denoting the history of the algorithm:

$$\mathcal{H}_t \triangleq \left\{\boldsymbol{\theta}^{(0:t)}, s^{(0:t-1)}, \alpha^{(0:t)}\right\}$$

◇ LYAPUNOV FUNCTIONS

A Lyapunov function is a potential function that is zero at the optimum and unbounded away from it. If this function has a single optimum and we can show that our algorithm descends on it, then we have shown that our algorithm is converging.

More formally, a Lyapunov function $V : \mathbb{R}^n \mapsto \mathbb{R}$ must satisfy:

1. $V(\boldsymbol{\theta}) \geq 0$.
2. $\nabla V(\boldsymbol{\theta}^\star) = 0$.

In particular consider the following Lyapunov function:

$$V(\boldsymbol{\theta}^{(t)}) = \frac{1}{2}\|\boldsymbol{\theta}^{(t)} - \boldsymbol{\theta}^\star\|_2^2$$

Then

$$\nabla V(\boldsymbol{\theta}^{(t)}) = \boldsymbol{\theta}^{(t)} - \boldsymbol{\theta}^\star$$

and our SA algorithm takes the standard stochastic descent

form:

$$\begin{aligned}\boldsymbol{\theta}^{(t+1)} &= \boldsymbol{\theta}^{(t)} + \alpha^{(t)} s(w^{(t)}, \boldsymbol{\theta}^{(t)}) \\ &= \boldsymbol{\theta}^{(t)} + \alpha^{(t)} \left[w^{(t)} - \nabla V(\boldsymbol{\theta}^{(t)}) \right]\end{aligned}$$

We now prove an important lemma about potential functions, which basically states that the angle between the descent direction and the outward gradient is more than ninety degrees and that the steps are not "too large".

Lemma 4 *Assume* $\mathbb{E}[w^{(t)}|\mathcal{H}_t] = 0$ *and* $\mathbb{E}[\|w^{(t)}\|^2|\mathcal{H}_t] \leq A + B\|\nabla V(\boldsymbol{\theta})\|^2$. *Then there exist constants* c_1 *and* c_2 *such that*

1. $\nabla V(\boldsymbol{\theta}^{(t)})^T \mathbb{E}[s(w^{(t)}, \boldsymbol{\theta}^{(t)})|\mathcal{H}_t] \leq -c_1\|\nabla V(\boldsymbol{\theta})\|^2$.
2. $\mathbb{E}[\|s(w^{(t)}, \boldsymbol{\theta}^{(t)})\|^2|\mathcal{H}_t] \leq c_2(1 + \|\nabla V(\boldsymbol{\theta})\|^2)$.

$\star$

$\star$ Proof:

⋆ Proof:

The lemma can also be proved if we instead assume that F is a contraction:

⋆ Proof:

◇ MARTINGALE THEOREMS

If you ever decide to gamble, you should know about martingales. Here we simply state a powerful result about martingales:

Theorem 6 super-martingale theorem

1. *Consider a non-negative random variable* $\mathbf{x}_t \geq 0$*, such that* $\mathbb{E}[\mathbf{x}_{t+1}|\mathcal{H}_t] \leq \mathbf{x}_t$*, then* $\mathbf{x}_t \longrightarrow \overline{\mathbf{x}} \geq 0 \quad w.p.\ 1.$
2. *Consider* $\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t \geq 0$*, such that* $\sum \mathbf{y}_t < \infty$ *and* $\mathbb{E}[\mathbf{x}_{t+1}|\mathcal{H}_t] \leq \mathbf{x}_t + \mathbf{y}_t - \mathbf{z}_t$*, then* $\mathbf{x}_t \longrightarrow \overline{\mathbf{x}} \geq 0$ *and* $\mathbf{z}_t \longrightarrow 0 \quad w.p.\ 1.$

◇ STOCHASTIC CONVERGENCE

We can now state the main result of this lecture:

Theorem 7 *If the SA algorithm satisfies the following assumptions under an appropriate choice of Lyapunov function:*

1. $\nabla V(\boldsymbol{\theta}^{(t)})^T \mathbb{E}[s(w^{(t)}, \boldsymbol{\theta}^{(t)})|\mathcal{H}_t] \leq -c_1 \|\nabla V(\boldsymbol{\theta})\|^2.$
2. $\mathbb{E}[\|s(w^{(t)}, \boldsymbol{\theta}^{(t)})\|^2|\mathcal{H}_t] \leq c_2(1 + \|\nabla V(\boldsymbol{\theta})\|^2).$
3. $\alpha^{(t)}$ *is a diminishing positive sequence with* $\sum_t \alpha^{(t)} = \infty$ *and* $\sum_t \alpha^{(t)2} < \infty$.

Then, $\boldsymbol{\theta}^{(t)} \longrightarrow \boldsymbol{\theta}^\star \quad w.p.\ 1.$

⋆ Proof:

⋆ Proof:

◇ LIMIT DIFFERENTIAL EQUATION

⋆