ECE 555 **Assignment # 5** **(version 2)**

Issued: April 12 Due: ~~April 21~~ May 5, 2011

Reading: See resources page - handouts and bibliography on ODE methods and SA.

Exercises:

13. Compute the conditional distribution of X_1 and X_2 for the Gaussian two-armed bandit problem introduced in class. Recall that the observations are given by,

$$Y(t) = U(t)X_1 + (1 - U(t))X_2 + N(t)$$

where $\{X_i\}$ are independent Gaussian, $N(0, \sigma_i^2)$, $\boldsymbol{N}$ is i.i.d., $N(0, \sigma_N^2)$, and $U(t) \in \{0, 1\}$ is adapted to $\boldsymbol{Y}$.

14. *Avoidance of traps.*

In this exercise you will try out the SA algorithm to minimize a function $f(x)$ over $x \in \mathbb{R}$. You will consider $f(x) = x^2(1 + (x + 10)^2)$. We have access only to noisy measurements of its normalized gradient: If at time t the value x is selected, then $Y(t)$ is observed,

$$Y(t) = \frac{1}{1 + (x + 5)^2} \nabla f\,(x) + N(t)$$

where $\boldsymbol{N}$ is i.i.d., with zero mean, and unit variance.

(i) Apply the stochastic approximation algorithm repeatedly, from various initial conditions, to obtain estimates $\{X(t)\}$ of $x^* = 0$. Obtain an estimate of the probability that $X(\infty) = 0$ when $X(0) = 20$. Repeat, with $X(0) = -20$.

(ii) Compare the sample path behavior of the standard SA algorithm, with the algorithm obtained using Polyak's averaging technique (again for $X(0) = 20$, and $X(0) = -20$).

(iii) Propose a modification the SA algorithm to ensure that your estimates converge to $x^* = 0$ with probability one, from each initial condition. Lai & Robbins' exploration should provide some inspiration.

15. *Variance analysis of stochastic approximation.* Consider the standard SA recursion in d dimensions,

$$X(t+1) = X(t) + \frac{G}{t+1}\Big(f(X(t)) + N(t+1)\Big), \qquad X(0) \in \mathbb{R}^d,$$

where $\boldsymbol{N}$ is i.i.d., with zero mean, and finite variance. The $d \times d$ matrix G is to be optimized as part of this exercise.

Suppose that the recursion $\boldsymbol{X}$ is globally convergent to the origin. To analyze the variance of the algorithm, we might consider a linearization,

$$\mathcal{E}(t+1) = \mathcal{E}(t) + \frac{G}{t+1}\Big(A\mathcal{E}(t) + N(t+1)\Big), \qquad \mathcal{E}(0) = 0.$$

where $A = \nabla f\,(0)$.

(i) Obtain a recursive formula for $\Sigma_G(t) = \mathsf{E}[\mathcal{E}(t)\mathcal{E}(t)^{\mathrm{T}}]$.

(ii) Find conditions on A and G such that the asymptotic variance exists,

$$\Sigma_G^{\infty} = \lim_{t\to\infty} t\Sigma_G(t)$$

(iii) Obtain the matrix G that minimizes $\operatorname{trace}\Sigma_G^{\infty}$.