CS 277 (W22): Control and Reinforcement Learning
Quiz 1: Mathematical Background

Due date: Friday, January 7, 2022 (Pacific Time)

Roy Fox
https://royf.org/crs/W22/CS277

Question 1 The *hybrid argument* is a proof technique that will occasionally be useful in this course. Let $x_0, \dots, x_T$ be a sequence of $T + 1$ real numbers. For some $\epsilon > 0$, suppose that $|x_{t+1} - x_t| \leq \epsilon$ for all $t = 0, \dots, T - 1$. Then of the following bounds on $|x_T - x_0|$, the tightest that always holds is:

- ☐ $|x_T - x_0| \leq \epsilon$
- ☐ $|x_T - x_0| \leq \epsilon T$
- ☐ $|x_T - x_0| \leq \epsilon(T + 1)$
- ☐ $|x_T - x_0| \leq 2\epsilon T$
- ☐ None of the above always holds.

Briefly justify:

Question 2 Let A be an $n \times n$ matrix and $p_A(\lambda) = |\lambda I - A|$ its characteristic polynomial of degree n. The *Cayley–Hamilton theorem* states that $p_A(A) = 0$. This implies that the columns of A^t are always spanned by the columns of $\begin{bmatrix} I & A & A^2 & \cdots & A^{t-1} \end{bmatrix}$ when t is: (check the highest that always holds)

- ☐ $t \geq n - 1$
- ☐ $t \geq n$
- ☐ $t \geq n + 1$
- ☐ None of the above always hold.

Briefly justify:

Question 3 Let $t \geq 0$ be a random variable with a *geometric distribution* with parameter ("success probability") $1 - \gamma$. Check all that hold:

- □ $p(t) = \gamma(1-\gamma)^t$
- □ $p(t|t \geq t_0) = p(t - t_0)$
- □ $\mathbb{E}[t] = \sum_{t_0=0}^{\infty} p(t \geq t_0)$
- □ $\lim_{\gamma \to 1} p\left(t \geq \frac{1}{1-\gamma}\right) = e^{-1}$

Question 4 Consider the optimization problem

$$\begin{aligned} \min_{\theta \in \Theta} \quad & f(\theta) \\ \text{s.t.} \quad & g(\theta) \leq 0, \end{aligned}$$

with $g(\theta)$ an n-dimensional real vector, constrained to be non-positive element-wise. To assure existence of the optimum, assume that Θ is compact, that f is defined and bounded over Θ, and that there exists $\theta \in \Theta$ with $g(\theta) \leq 0$. The *Lagrange method* introduces an n-dimensional vector λ of *Lagrange multipliers*, and optimizes

$$\min_{\theta \in \Theta} \max_{\lambda \geq 0} f(\theta) + \lambda^\intercal g(\theta).$$

Check all that hold for optimum-achieving θ and λ:

- □ $g(\theta) \leq 0$
- □ $\lambda_i g_i(\theta) = 0$ for all i
- □ There is no i for which $\lambda_i = g_i(\theta) = 0$
- □ $\lambda^\intercal g(\theta) = 0$ and thus $\theta = \operatorname{argmin}_{\theta \in \Theta} f(\theta)$

Question 5 A discrete-time *stochastic process* is a sequence $x_0, x_1, \ldots$ of random variables with joint distribution

$$p(x_0, x_1, \ldots) = p(x_0) \prod_{t \geq 0} p(x_{t+1}|x_{\leq t}),$$

where $x_{\leq t} = x_0, \ldots, x_t$. The process has the *Markov property* (and called a Markov chain) if, for all $t \geq 0$, $p(x_{t+1}|x_{\leq t}) = p(x_{t+1}|x_t)$. The process is called *time-invariant* if $p(x_{t+1}|x_t)$ is independent of t (it may still be a function of the realizations of x_t and x_{t+1}). Check all that hold in this case:

- □ $p(x_{t_1}, x_{t_3}|x_{t_2}) = p(x_{t_1}|x_{t_2})p(x_{t_3}|x_{t_2})$ for all $t_1 < t_2 < t_3$; that is to say, x_{t_1} and x_{t_3} are independent given x_{t_2}
- □ $p(x_{t+1}) = \mathbb{E}_{x_t}[p(x_{t+1}|x_t)]$
- □ There always exists a *stationary distribution* $\bar{p}$ such that if $p(x_t) = \bar{p}$ then also $p(x_{t+1}) = \bar{p}$
- □ The limit $\lim_{t \to \infty} p(x_t)$ is always a stationary distribution

Question 6 Consider the following *Bayesian network*:

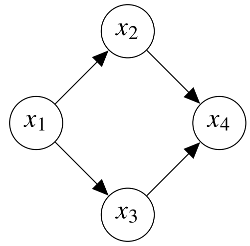

Here $p(x_1, x_2, x_3, x_4) = p(x_1)p(x_2|x_1)p(x_3|x_1)p(x_4|x_2, x_3)$. Check all that hold:

- ☐ x_1 and x_4 are independent
- ☐ x_1 and x_4 are independent given x_2
- ☐ x_1 and x_4 are independent given x_2 and x_3
- ☐ x_2 and x_3 are independent
- ☐ x_2 and x_3 are independent given x_1
- ☐ x_2 and x_3 are independent given x_4
- ☐ x_2 and x_3 are independent given x_1 and x_4