CS 277 (W24): Control and Reinforcement Learning
Quiz 1: Mathematical Background

Due date: Wednesday, January 17, 2024 (Pacific Time)

Roy Fox
`https://royf.org/crs/CS277/W24`

Instructions: please solve the quiz in the marked spaces and submit this PDF to Gradescope.

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 lowest that always holds)

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

Briefly justify:

Question 3 Consider the following *Bayesian network*:

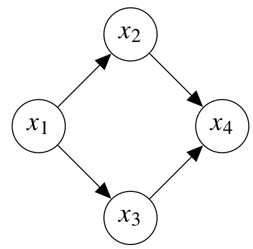

Here $p(x_1, x_2, x_3, x_4) = p(x_1)p(x_2|x_1)p(x_3|x_1)p(x_4|x_2, x_3)$. Check all that hold:

- □ 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

Question 4 Check all that hold:

- □ A geometric random variable t, i.e. having distribution $p(t) = (1 - \gamma)\gamma^t$ for $t \geq 0$, is time-invariant, i.e. $p(t|t \geq t_0) = p(t - t_0)$ for all $t \geq t_0$.
- □ For random variables x and y (not necessarily independent), $x + y$ and $x + \mathbb{E}[y]$ always have the same expectation but the latter always has lower variance.
- □ If a function $g_\theta(x)$ approximates another function $f_\theta(x)$, i.e. there exists some ϵ such that $|g_\theta(x) - f_\theta(x)| \leq \epsilon$ for all x, then the gradient of g w.r.t. θ approximates the gradient of f.
- □ If a distribution $p_\theta(x)$ and a real function $f_\theta(x)$ both depend on the same parameter θ, then the gradient of the expectation w.r.t. θ equals the expectation of the gradient, i.e. $\nabla_\theta \mathbb{E}_{x\sim p_\theta}[f_\theta(x)] = \mathbb{E}_{x\sim p_\theta}[\nabla_\theta f_\theta(x)]$.