Problems to be handed in

For the first two problems, you will need the definitions of positive semidefinite and positive definite matrices: A square matrix $A \in \mathbb{R}^{n \times n}$ is positive semidefinite (respectively, positive definite) if $x^\mathsf{T} Ax \ge 0$ for all nonzero $x \in \mathbb{R}^n$ (respectively, $x^\mathsf{T} Ax > 0$ for all nonzero $x \in \mathbb{R}^n$). We use the notation $A \succeq 0$ and $A \succ 0$ to indicate that A is positive semidefinite (respectively, positivie definite).

1 (convex and strictly convex functions) A function $f : \mathbb{R}^n \to \mathbb{R}$ is *convex* if, for any two points $x, y \in \mathbb{R}^n$ and any $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)y) \le \lambda f(x) + (1 - \lambda)f(y), \tag{1}$$

and *strictly convex* if the inequality in (1) is strict unless $x = y$ or $\lambda \in \{0, 1\}$.

(i) Prove that f is convex (respectively, strictly convex) if and only if the function $\lambda \mapsto f(\lambda x + (1 - \lambda)y)$ of $\lambda \in [0, 1]$ is convex (respectively, strictly convex) for any fixed pair of points $x, y \in \mathbb{R}^n$.

(ii) Suppose that f is twice differentiable. Use the result of part (i) to show that f is convex (respectively, strictly convex) if its Hessian $\nabla^2 f(x)$ is positive semidefinite (respectively, positive definite) at every point $x \in \mathbb{R}^n$.

(iii) Let f be a strictly convex function which is bounded from below, i.e., there exists some $c \in \mathbb{R}$ such that $f(x) \ge c$ for all x. Prove that any minimizer of f, if it exists, is unique. Give an example of a strictly convex function bounded from below that does not have any minimizers.

(iv) Consider the function

$$f : \mathbb{R}^n \to \mathbb{R}, \qquad f(x) := x^\mathsf{T} Ax + x^\mathsf{T} v + \alpha \tag{2}$$

for some $A \in \mathbb{R}^{n \times n}$, $v \in \mathbb{R}^n$, $\alpha \in \mathbb{R}$. Prove that f is convex (respectively, strictly convex and bounded from below) if and only if $A \succeq 0$ (respectively, $A \succ 0$).

(v) Use the result of part (iv) to give an alernative proof of the 'completion-of-squares' lemma.

2 (Schur complements and the LQR) In this problem, we will give the full rigorous derivation of the optimal controller and the value functions for the linear quadratic regulator.

(i) Let X be a block matrix of the form

$$X = \begin{pmatrix} A & B \\ B^\mathsf{T} & C \end{pmatrix}, \tag{3}$$

where A, B, C are matrices of appropriate shapes, and A and C are both symmetric (note that this implies that X itself is symmetric, $X = X^\mathsf{T}$). Prove that, if $C \succ 0$, then $X \succeq 0$ if and only if $S := A - BC^{-1}B^\mathsf{T} \succeq 0$ (the matrix S is called the *Schur complement* of C in X).

(ii) Let the matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $C \in \mathbb{R}^{m \times m}$ be given, such that $C \succ 0$, and consider the function

$$f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}, \qquad f(x,u) := x^\top A x + u^\top C u + 2x^\top B u. \tag{4}$$

Fix $x \in \mathbb{R}^n$ and consider the problem of minimizing $f(x,u)$ over $u \in \mathbb{R}^m$. Prove that $u^* = -C^{-1}B^\top x$ is the unique minimizer and that

$$f^*(x) := \min_{u \in \mathbb{R}^m} f(x,u) = x^\top S x, \tag{5}$$

where S is the Schur complement of C in the matrix X defined in (3). Prove that if $X \succeq 0$, then $f^*(x) = x^\top S x$ is a convex function of x.

(iii) Use the above results to give a complete solution of the dynamic program for the LQR problem

$$X_{t+1} = AX_t + BU_t + W_t, \tag{6}$$

where the disturbances $W_0, W_1, \ldots$ are i.i.d. zero mean random vectors with covariance matrix $\Sigma = \mathbf{E}[W_t W_t^\top]$ and the costs are given by

$$c_0(x,u) = \ldots = c_{T-1}(x,u) = x^\top Q x + u^\top R u, \qquad c_T(x) = x^\top Q_T x \tag{7}$$

with $Q, Q_T \succeq 0$ and $R \succ 0$. In particular, show that the optimal controller is linear, i.e., $g_t^*(x) = G_t x$ for some matrices $G_t \in \mathbb{R}^{m \times n}$, and the value functions are convex quadratic, i.e., $V_t(x) = x^\top K_t x + \alpha_t$ with $K_t \succeq 0$ and $\alpha_t \geq 0$.

Hint: Use the Schur complement condition to establish convexity of V_t. In particular, first show that $V_t \geq 0$ and then use this to show that K_t must be positive semidefinite.

3 (hidden Markov models) Consider a discrete-time stochastic process $\{(X_t, Y_t)\}_{t=0}^\infty$, where X_t and Y_t take values in finite sets $\mathcal{X}$ and $\mathcal{Y}$, respectively. We say that this process is a *hidden Markov model* if, for each t and for all tuples $x_0^t \in \mathcal{X}_0^t$ and $y_0^t \in \mathcal{Y}_0^t$,

$$\begin{aligned}\mathbf{P}[(X_t,Y_t) = (x_t,y_t) | (X_0,Y_0) = (x_0,y_0), \ldots, (X_{t-1},Y_{t-1}) = (x_{t-1},y_{t-1})] \\ = \mathbf{P}[(X_t,Y_t) = (x_t,y_t) | (X_{t-1},Y_{t-1}) = (x_{t-1},y_{t-1})].\end{aligned} \tag{8}$$

and

$$\mathbf{P}[Y_0 = y_0, \ldots, Y_t = y_t | X_0 = x_0, \ldots, X_t = x_t] = \prod_{s=0}^{t} \mathbf{P}[Y_t = y_t | X_t = x_t]. \tag{9}$$

The process $\{X_t\}_{t \geq 0}$ is called the *hidden state process* or the *signal process*, while $\{Y_t\}_{t \geq 0}$ is the *obseration process*. In other words, (8) states that $\{(X_t, Y_t)\}_{t \geq 0}$ is a Markov chain with state space $\mathcal{X} \times \mathcal{Y}$, while (9) states that the observations Y_0^t are conditionally independent given the corresponding signals X_0^t.

(i) Prove that a hidden Markov model is completely specified by the probability law μ of X_0 and by two sequences of nonnegative matrices $\{P^{(t)}\}_{t\geq 0}$ and $\{M^{(t)}\}_{t\geq 0}$, where the rows and columns of $P^{(t)}$ are indexed by the elements of $\mathcal{X}$ and the entries are given by

$$P^{(t)}(x,x') = \mathbf{P}[X_{t+1} = x' | X_t = x], \tag{10}$$

while the rows (respectively, columns) of $M^{(t)}$ are indexed by $\mathcal{X}$ (respectively, by $\mathcal{Y}$) and the entries are given by

$$M^{(t)}(x,y) = \mathbf{P}[Y_t = y | X_t = x]. \tag{11}$$

(ii) Let $\{(X_t, Y_t)\}_{t\geq 0}$ be a discrete-time stochastic process, where each (X_t, Y_t) takes values in the Cartesian product $\mathcal{X} \times \mathcal{Y}$ of two finite sets, $X_0 \sim \mu$, and there exist sequences of functions $\{f_t\}_{t\geq 0}$ and $\{h_t\}_{t\geq 0}$ and two mutually independent sequences $\{W_t\}_{t\geq 0}$ and $\{V_t\}_{t\geq 0}$ of i.i.d. random variables[1] that are also independent of X_0, such that

$$X_{t+1} = f_t(X_t, W_t) \qquad \text{and} \qquad Y_t = h_t(X_t, V_t) \qquad \text{for all } t \geq 0. \tag{12}$$

Prove that $\{(X_t, Y_t)\}_{t\geq 0}$ is a hidden Markov model.

(iii) Consider a hidden Markov model specified by μ, $\{P^{(t)}\}_{t\geq 0}$, and $\{M^{(t)}\}_{t\geq 0}$. Prove that it admits a realization of the form described in part (ii), i.e., one can always find two mutually independent sequences $\{W_t\}_{t\geq 0}$ and $\{V_t\}_{t\geq 0}$ of i.i.d. random variables that are also independent of X_0, as well as sequences of functions $\{f_t\}_{t\geq 0}$ and $\{h_t\}_{t\geq 0}$, such that $X_0 \sim \mu$ and

$$X_{t+1} = f_t(X_t, W_t) \qquad \text{and} \qquad Y_t = h_t(X_t, V_t) \qquad \text{for all } t \geq 0. \tag{13}$$

(iv) Let $\{(X_t, Y_t)\}_{t\geq 0}$ be a hidden Markov model. Prove that the signal process $\{X_t\}_{t\geq 0}$ is a Markov chain. (On the other hand, $\{Y_t\}_{t\geq 0}$ may not be a Markov chain.)

4 (unnormalized nonlinear filter) Let $\{(X_t, Y_t)\}_{t\geq 0}$ be a hidden Markov model specified by μ, $\{P^{(t)}\}_{t\geq 0}$, and $\{M^{(t)}\}_{t\geq 0}$. Let $y_0, y_1, \ldots$ be a fixed sequence of observations. For each $t \geq 0$, define the *unnormalized filtering distributions*

$$\sigma_t(x_t) := \sum_{x_0^{t-1} \in \mathcal{X}_0^{t-1}} \mu(x_0) \prod_{s=0}^{t-1} P^{(s)}(x_s, x_{s+1}) \prod_{s=0}^{t} M^{(s)}(x_s, y_s). \tag{14}$$

(i) Prove that the unnormalized filtering distributions can be computed recursively according to

$$\sigma_{t+1}(x_{t+1}) = \sigma_t P^{(t)}(x_{t+1}) M^{(t+1)}(x_{t+1}, y_{t+1}) \tag{15}$$

with the initial condition

$$\sigma_0(x_0) = \mu(x_0) M^{(0)}(x_0, y_0). \tag{16}$$

[1]The distributions of W_0 and V_0 need not be the same; in fact, W_0 and V_0 need not take values in the same set.

(ii) Prove that the (normalized) filtering distributions π_t can be computed in terms of the unnormalized filtering distributions as

$$\pi_t(x_t) = \frac{\sigma_t(x_t)}{\sum_{x \in \mathcal{X}} \sigma_t(x)}. \tag{17}$$