Final Exam

CSE 410/510, Robot Algorithms
Available: 10am, 5/8/2012
Due: noon, 5/9/2012

Please turn in your answers to the TA's mail box in the CSE main office in Davis Hall or email them to the TA. Please email me with your questions. I will answer questions as promptly as possible, but I will be unavailable between 5pm and 10pm Tuesday evening. All answers to emailed questions will be posted on the exam webpage at:

http://people.csail.mit.edu/rplatt/exam.html

You may use any materials you want. However, you may not discuss these questions with any other person besides me or the TA.

1. Consider a system with state, $x_t = \begin{pmatrix} y_t \\ z_t \end{pmatrix}$.

a) (5 pts) Suppose that on every time step, these variables experience a process update of

$$y_{t+1} = y_t + u_t$$

and

$$z_{t+1} = z_t + 2y_t$$

Write this update in state space form,

$$x_{t+1} = Ax_t + Bu_t.$$

(That is, give me the matrices A and B that express the dynamics in the above equation.)

b) (6 pts) Now, suppose that on every time step, we have the following process update:

$$y_{t+1} = y_t + u_t$$

and

$$z_{t+1} = z_t - \frac{z_t^2}{z_t + 0.5}.$$

Write this update in state space form,

$$x_{t+1} = A(m_t)(x_t - m_t) + Bu_t + c(m_t),$$

where the system is linearized about m_t.

2. Consider linear quadratic regulation (LQR) for a stochastic linear system with Gaussian process noise in comparison to LQR for a deterministic linear system.

(a) (3 pts) How does the LQR value function for the stochastic system differ from the LQR value function for the deterministic system?

(b) (3 pts) What effect does this difference have on the optimal policy?

(c) (4 pts) Write the Hamilton-Jacobi-Bellman equation for this system.

3. (11 pts) Consider Figure below. At each time step, the system may be in one of four states corre-

cliff	cliff	cliff	cliff	cliff
cliff				$\times$
cliff	cliff	cliff	cliff	cliff

sponding to the four squares not labelled "cliff". One each time step, the system moves randomly (with equal probability) in one of the four directions: up, down, left, or right. If the system moves into a cliff state, the system gets zero reward and execution ends. If the system moves into the state labeled with an "X", then it gets a reward of $+1$ and execution ends. If the system moves into an unlabeled state, then the system gets zero reward. Calculate the expected reward in the three unlabeled states.

4. (11 pts) Consider the function, $f(x) = -\log(x)$, in the domain $x > 0$. Is this function convex? Prove your answer is correct using the second derivative.

5. Suppose we are given a deterministic linear system, $x_{t+1} = Ax_t + Bu_t$.

(a) (4 pts) Write x_4 (the state at time $t = 4$) in terms of x_1 (the state at time $t = 1$), u_1, u_2, and u_3.

(b) (4 pts) Given an initial state, x_1, and a final state, x_4, write an equation for the least-squares solution, $u_{1:3} = \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}$ (the solution that minimizes the L2 norm of the vector $u_{1:3}$).

(c) (3 pts) Under what conditions may this solution not exist?

6. An object is launched vertically upward into the air with an initial velocity, $v_1 \geq 0$. The object experiences a constant downward gravitational force of g and zero wind resistance. The discretized equation of motion of the object is: $v_{t+1} = v_t - g$. The starting position at time, $t = 1$, is x_1. The ending position at time τ is constrained to be $x_T \neq x_1$.

a) (6 pts) Write an optimization problem that calculates the following: the initial velocity that minimizes the sum of the squared flight time and the squared initial velocity, *i.e.* minimize $v_1^2 + \tau^2$.

b) (5 pts) What kind of optimizaton problem is this? What kind of optimization problem does it become if τ is held fixed?

7. Consider the probability density function (the pdf) shown below:

The pdf is zero everywhere except between zero and one where the pdf is as shown in the figure.

a) (4 pts) Calculate the expectation of this distribution using calculus.

b) (3 pts) If you were to use importance sampling to approximate the expectation, what proposal distribution would you use?

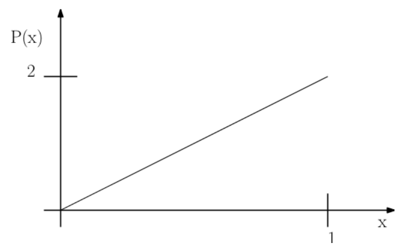

c) (4 pts) Give pseudo-code for using importance sampling to calculate the expectation. Be as specific as possible regarding the weighting factor.

8. (11 pts) Consider the four-state world illustrated below where each square is one state. At a

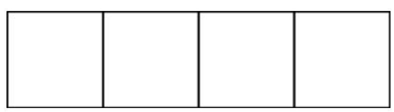

given point in time, an agent may occupy one state. Suppose that initially, it is equally likely that the agent is in any one of the four states. Suppose that the agent executes a sequence of two moves. On each move, the state of the agent moves to the right with probability 0.5, moves to the left with probability 0.25, and stays where it is with probability 0.25. If the move would take the agent past the last state on the left or the right, then the state of the agent does not change. Without using a computer (but a calculator is okay), write down the probability distribution prior to the moves, and after each move (total of three distributions).

9. A robot is equipped with a single-beam laser scanner pointing to the right as illustrated below. The horizontal position of the robot, measured from the left wall, is denoted x. The distance between the walls is d. The laser scanner measures the distance between the robot and the right wall. Distance measurements, z, are corrupted by Gaussian noise with variance, σ^2.

a) (6 pts) Write an expression for the probability density function, $P(z|x)$, to within a constant factor.

b) (6 pts) Suppose that distance measurements larger than a threshold, $\theta < d$, are reported to

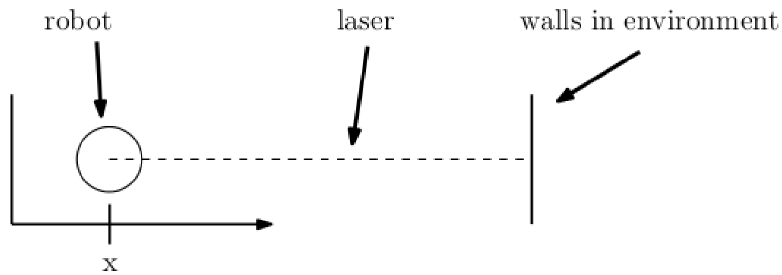

be equal to θ: $z = \theta$. Write a new expression for the probability density function, $P(z|x)$, to within a constant factor.