Midterm Exam **CS226**

Stanford CS226 Statistical Algorithms in Robotics, Winter 2009/10

Full Name: ______________________________

Email: ______________________________

Question	Points
1 (20 max)	
2 (20 max)	
3 (20 max)	
4 (20 max)	
5 (20 max)	
total	

Welcome to the CS223B Midterm Exam!

- The exam is 5 pages long. Make sure your exam is not missing any sheets. The exam has a maximum score of 100 points. You have 60 minutes.
- The exam is open book, open notes, but closed cell phones, etc.
- Write your answers in the space provided. If you need extra space, use the back of the preceding sheet.
- Write clearly and be concise.
- All points will be manually counted before certification.

1 Occupancy Grid Maps and Binary Bayes Filters 20pts

You learned about the log-odds update rule of binary Bayes filters.

1. What happens if we force the log-odds to lie inside an interval $[-b; b]$ for some positive, small b? In what type of situation would such a bound look like a plausible choice?

2. In the log-odds calculation, we have to constantly offset for the prior. Suppose we omit this term. Under what condition will this be mathematically justified? And, conversely, what is the worst thing that can happen when we drop this term?

3. In occupancy grid maps, any cell outside the field of view of a sensor is never updated. While this makes intuitive sense, we want you to prove that this follows from the basic log-odds update rule.

2 Particle filters 20pts

1. Consider a robot whose motion is deterministic (=no motion noise). Will a particle filter (if implemented exactly, that is, without assumed motion noise) approximate the correct posterior?

2. Say we are resampling over and over again, without any motion update. What will happen in the limit?

3. What happens to the particle filter if we use exactly one particle, and why?

3 Gaussians and EKF 20pts

1. In class, we learned that the product of two Gaussians is once again Gaussian. Derive the variance ϕ^2 of the product of two 1-D Gaussians (don't derive the new mean).

$$\frac{1}{\sqrt{2\pi\ \sigma^2}} \exp\left\{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right\} \qquad \frac{1}{\sqrt{2\pi\ \rho^2}} \exp\left\{-\frac{1}{2}\frac{(x-\nu)^2}{\rho^2}\right\}$$

2. Prove that the resulting variance ϕ^2 you just derived is smaller than σ^2 and smaller than ρ^2 (feel free to ignore the degenerate case of infinite variance).

3. In problems like localization, EKFs work poorly if the initial state is unknown (global localization), whereas they often work well if the initial state is known (tracking). Explain how this relates to the "E" in "EKF."

4 FastSLAM 20pts

The FastSLAM algorithm as discussed in class uses Gaussians to represent the 2-D landmark estimates (for notation: N denotes the number of path particles, and M the number of landmarks). Say we wish to replace those Gaussians by particle filters.

1. What would be the number of state variables in each of those new particle filters?

2. How many such new particle filters would there be?

3. When is this a good idea?

5 GraphSLAM 20pts

Say we are using GraphSLAM to build a 2-D map with K robots. For simplicity, we will assume that all measurement and motion equations are linear (so don't worry about linearization). Also, assume that all initial poses are known. Let us consider the case in which all robots communicate and maintain a single, joint belief.

1. How will the joint information matrix be structured? Draw a diagram of this matrix, and indicate which parts of this matrix will be sparse.

2. Will it be symmetric? Will it be positive semi-definite?

3. Say robots can sense each others relative distance; and again everything is linear. Does this affect the information matrix, and if so, how?

4. Say the robots cannot communicate for a while, and then communication is re-established. Will this affect the posterior estimate (at the time the communication is available again)? Argue your answer.