MASSACHUSETTS INSTITUTE OF TECHNOLOGY
Department of Electrical Engineering and Computer Science

6.438 ALGORITHMS FOR INFERENCE
Fall 2013

Final Quiz
Tuesday, December 10, 2013
7:00pm–10:00pm

- This is a closed book exam, but two $8\frac{1}{2}'' \times 11''$ sheets of notes (4 sides total) are allowed.
- Calculators are **not** allowed.
- There are **3** problems.
- The problems are not necessarily in order of difficulty. We recommend that you read through all the problems first, then do the problems in whatever order suits you best.
- Record all your solutions in the answer booklet provided. **NOTE: Only the answer booklet is to be handed in—no additional pages will be considered in the grading**. You may want to first work things through on the scratch paper provided and then neatly transfer to the answer sheet the work you would like us to look at. Let us know if you need additional scratch paper.
- A correct answer does not guarantee full credit, and a wrong answer does not guarantee loss of credit. You should clearly but concisely indicate your reasoning and **show all relevant work**. Your grade on each problem will be based on our best assessment of your level of understanding as reflected by what you have written in the answer booklet.
- Please be neat—we can't grade what we can't decipher!

Question 1

Consider three random variables X_1, X_2 and X_3. Let each of them be binary valued, i.e. $X_i \in \{0, 1\}$ for $1 \leq i \leq 3$. Let their joint distribution be given by

$$\mathbb{P}(x_1, x_2, x_3) \propto \phi_1(x_1)\phi_2(x_2)\phi_3(x_3)\psi_{12}(x_1, x_2)\psi_{13}(x_1, x_3)\psi_{23}(x_2, x_3), \tag{1}$$

for any $(x_1, x_2, x_3) \in \{0, 1\}^3$. Let $\phi_i,\ 1 \leq i \leq 3$, and $\psi_{ij},\ 1 \leq i < j \leq 3$ be strictly positive valued potentials. Therefore, each of $\{0, 1\}^3$ has strictly positive probability.

(a) Draw the corresponding graphical model.

(b) Write down sum-product equations.

(c) Will sum-product always converge for this graphical model ? If no, provide a counter-example. If yes, provide a detailed proof.

Question 2

We learnt about Kalman filtering (inference in Gaussian Hidden Markov Model) in class. The goal of this question is to understand how it can be used in a practical scenario through a stylized, but meaningful toy example.

We wish to control (i.e. navigate) a vehicle in an one dimensional space. Let $x_t \in \mathbb{R}$ be the position of vehicle at time $t = 0, 1, 2, \ldots$. We observe noisy position, denoted by y_t, which is given by

$$y_t = x_t + z_t, \tag{2}$$

where z_t is distributed as per Gaussian distribution with mean 0 and variance 1; and independent of everything else, for all $t \geq 0$. At each time t, we apply "control" $u_t \in \mathbb{R}$ to the vehicle resulting in the change of its location as

$$x_{t+1} = x_t + u_t. \tag{3}$$

We shall assume that the control u_t, at time t, is decided based on observations $y_0^t = (y_0, \ldots, y_t)$ and past control decisions $u_0^{t-1} = (u_0, \ldots, u_{t-1})$; that is, entire history $F^t = (y_0^t, u_0^{t-1})$.

The goal of the controller is to move the vehicle to as close to 0 as possible. However, moving (controlling) the vehicle requires controller to spend energy. And, ideally controller wants to achieve the goal of moving vehicle to 0 at minimal cost. Given this, a reasonable objective, over time $0 \leq t \leq T$ is given by

$$\min_{u_0^{T-1}} \mathbb{E}[x_T^2] + \sum_{t=0}^{T-1} u_t^2, \tag{4}$$

Naturally, solving this problem requires two key components: (i) estimating the state of the system at each time instance, and (ii) use it to decide the amount of control that we wish to exert at a given time. This is precisely what we shall resolve, in that order.

Estimation. This part should make you realize the value of Kalman filtering.

(a) Draw the graphical model of $\{x_0, \ldots, x_T, y_0, \ldots, y_T\}$.
Note: Since control is designed by you, $u_0, \ldots$ are *observed*.

(b) Briefly state how you would produce maximum likelihood estimation of state x_t at time t, given all the history F^t. We shall denote it by $\hat{x}_t$.

(c) Now, we would like to understand the "error" in our estimation.
Assume that the initial location x_0 follows a Gaussian distribution with mean 0 and variance V_0. Let

$$\Delta_t = \hat{x}_t - x_t,$$

denote the estimation error in the maximum likelihood estimation given history up to time t. We shall derive the error distribution inductively. To that end, suppose you are told that Δ_{t-1} has Gaussian distribution with mean 0 and variance V_{t-1}, conditioned on F^{t-1}. Given this, obtain the distribution of Δ_t conditioned on F^t (and of course, you know $\hat{x}_{t-1}$). In particular, show that the variance V_t, of Δ_t, satisfies

$$V_t^{-1} = V_0^{-1} + t \tag{5}$$

Note: It may seem surprising that the variance in error is independent of control sequence.

Control. In this part, you will be guided to derive the optimal control.

(d) Consider scenario when T is very large. Using (c), argue that effectively one can implement control with 0 objective cost as $T \to \infty$.

(e) The more interesting scenario is that of finite T. To that end, we need to carefully evaluate the objective each time and make decision consequently. A "dynamic programming" approach to this is given below: recursively, define

$$W_T = \lambda x_T^2,$$

and for $0 \leq t < T$,

$$W_t = \min_{u_t} \left\{ u_t^2 + \lambda \, \mathbb{E}[W_{t+1} | F^t] \right\}$$

We shall inductively argue the following: for $0 < t < T$, assume that

$$W_t = d_t \hat{x}_t^2 + c_t,$$

where c_t, d_t are constants that may vary with t. Then, using this form and by solving optimization problem

$$W_{t-1} = \min_{u_{t-1}} \left(u_{t-1}^2 + \mathbb{E}[W_{t+1} | F^t] \right),$$

argue that W_{t-1} has a form $d_{t-1} \hat{x}_{t-1}^2 + c_{t-1}$. Determine d_t using these recursive equations. Notice that, the solution of the above optimization problem provides the solution for optimal control u_t, as function of your state estimate, $\hat{x}_t$.

Question 3

We are re-visiting the crowd-sourcing problem from earlier quiz with some additions. Let us quickly remind ourselves of the crowd-sourcing setting. We have M workers and N tasks.

Tasks. Each task i, $1 \leq i \leq N$, has a true answer $\mathsf{z}_i \in \{+1, -1\}$ – e.g. task "Is Kampala the capital of Uganda?" has answer *Yes* or $+1$. We shall assume that each of the N task has true answer $+1$ with probability $\theta \in (0, 1)$, independent of everything else. That is, i.i.d. Bernoulli prior on tasks with parameter θ.

Workers. Let $\mathsf{L}_{ij} \in \{-1, 0, +1\}$ denote answer of worker j to task i – it will be ± 1 if worker j is assigned task i, and 0 otherwise. Let $\mathsf{q}_j \in (0, 1)$ be the *honesty* of a worker: i.e. if worker j is assigned task i, then

$$\mathsf{q}_j = \mathbb{P}[\mathsf{L}_{ij} = \mathsf{z}_i]$$

A worker with $q_j \approx 1$ is an *expert*, while with $q_j \approx 0$ an *adversary*[1]. We assume that worker honesties q_j, $1 \leq j \leq M$, have i.i.d. Beta distribution with parameters α, β: for $q \in (0, 1)$, the density is given by

$$\mathbb{P}_{\mathsf{q}_j}(q; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} q^{\alpha - 1}(1 - q)^{\beta - 1}, \tag{6}$$

where $\alpha, \beta > 0$ and $B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$. Due to known property of Beta distribution, we have

$$\mathbb{E}[\mathsf{q}_j] = \mathbb{E}[\mathsf{q}] = \frac{\alpha}{\alpha + \beta}.$$

[1] I.e., he/she intentionally gives wrong answers to serve his/her own personal interests.

Part 1. For this part, let us suppose that tasks arrive to the system one-by-one. Upon arrival of task i, we make a decision which set of workers get assigned to it. Let $\mathsf{s}_{i,j} \in \{0,1\}$ denote whether worker j is assigned task i (i.e. $\mathsf{s}_{i,j} = 1$) or not. The task assignments are done as follow: (i) if worker j is assigned task $i-1$, then s/he will not be assigned task i (giving rest to the worker); (ii) else, worker j is assigned to task i with probability $\rho_j \in (0,1)$ (activity parameter of worker j). Formally, $\mathbb{P}[\mathsf{s}_{i,j} = 1 | \mathsf{s}_{i-1,j} = 1] = 0$, while $\mathbb{P}[\mathsf{s}_{i,j} = 1 | \mathsf{s}_{i-1,j} = 0] = \rho_j$.

(a) Draw an appropriate graphical model for this problem.

(b) Compute $\mathbb{P}[\mathsf{s}_{i,j} = 1]$ for each i. Observe that it simplifies for very large i (i.e., $i \to \infty$).

(c) Compute the average response of workers (again, assume i very large), i.e., $\lim_{i\to\infty} \mathbb{E}\left[\frac{1}{M}\sum_{j=1}^{M} L_{ij}\right]$.

(d) When does the sign of the number computed in (c) agree with the true answer of the task for very large i (i.e., $i \to \infty$)?

Part 2. In this part, we assume that workers and tasks are pre-assigned. Let $\mathcal{N}_j \subset \{1, \dots, N\}$ denote tasks worker j is assigned to, and $\mathcal{M}_i \subset \{1, \dots, M\}$ denote workers task i is assigned. Let $\mathbf{L}$ be the answers provided by workers to tasks that we observe. Given this, and model parameters α, β and θ, the goal is to estimate $\mathbf{z} = (z_1, \dots, z_N)$ and $\mathbf{q} = (q_1, \dots, q_M)$.

(e) Utilize mean-field variational approximation to estimate marginals for each of **z**, and **q**, assuming a fully factorized distribution

$$b(\mathbf{z}, \mathbf{q}) = \prod_{i=1}^{N} \mu_i(z_i) \prod_{j=1}^{M} \nu_j(q_j).$$

The approximating distribution is $p_{\mathbf{z},\mathbf{q}|\mathbf{L}}(\mathbf{z}, \mathbf{q}|\mathbf{L})$. Provide the mean-field update equations for $\mu_i(z_i), \forall i$.

(f) Instead of mean-field, now state sum-product update equations for estimating marginal distributions for z_i, $p_{z_i|\mathbf{L}}(z_i|\mathbf{L}), \forall i$. For simplicity, assume $\theta = 1/2$ from now on.

Hint: By doing a clever manipulation, you might be able to express the full distribution $p_{\mathbf{z}|\mathbf{L}}(\mathbf{z}|\mathbf{L})$ as a product of M factors.