

Homework 5

CMU 10-708: Probabilistic Graphical Models (Fall 2020)

November 20, 2020

Instructions:

- **Collaboration policy:** Collaboration on solving the homework is allowed, after you have thought about the problems on your own. It is also OK to get clarification (but not solutions) from books or online resources, again after you have thought about the problems on your own. There are two requirements: first, cite your collaborators fully and completely (e.g., “Bob explained to me what is asked in Question 4.3”). Second, write your solution *independently*: close the book and all of your notes, and send collaborators out of the room, so that the solution comes from you only.
- **Submitting your work:** Assignments should be submitted as PDFs using Gradescope unless explicitly stated otherwise. Each derivation/proof should be completed on a separate page. Submissions can be handwritten, but should be labeled and clearly legible. Else, submission can be written in LaTeX. Upon submission, label each question using the template provided by Gradescope.

1 Exponential Family (Chih-Kuan) [20 points]

Recall that exponential family has density of the form $p_\theta(x) = \exp\{\langle \theta, \phi(x) \rangle - A(\theta)\}$ with respect to some base measure $v(x)$, and where its log-normalization constant $A(\theta) = \log\{\int_x \exp\{\langle \theta, \phi(x) \rangle\} v(x) dx\}$.

1. [3 points] Show that $\frac{\partial A(\theta)}{\partial \theta_i} = E[\phi_i(x)]$.
2. [3 points] Show that $\frac{\partial^2 A(\theta)}{\partial \theta_i \partial \theta_j} = \text{Cov}(\phi_i(x), \phi_j(x))$.
3. [2 points] Show that $A(\theta)$ is convex.
4. [4 points] Show that $A^*(\mu) = \sup_{\theta \in \Omega} (\langle \theta, \mu \rangle - A(\theta))$ is convex.
5. [2 points] Suppose we are given n samples $\{x_1, x_2, \dots, x_n\}$. Then the likelihood of the data with respect to the exponential family above can be written as

$$\prod_{i=1}^n p_\theta(x_i).$$

Give a compact form of the likelihood of the data based on $\theta, \phi(x_i), A(\theta)$.

6. [2 points] Show that the log-likelihood is concave.
7. [2 points] Calculate the derivative of the likelihood of the data with respect to θ .
8. [2 points] Since the likelihood is concave, you can obtain the maximum likelihood distribution θ by setting the derivative to 0, and solving for θ . Explain the relationship between the mean parameter of the sufficient statistics (under the maximum likelihood distribution θ) and the empirical average of the sufficient statistics over the given samples.

2 Variational Inference [40 points] (Chang)

In this problem, we are going to work with approximate posterior inference via variational inference for a given topic model.

The standard Latent Dirichlet Allocation model only models the word co-occurrences, without considering temporal information, i.e. the time when a document is generated. However, a large number of subjects in documents change dramatically over time. It is important to interpret the topics in the context of the timestamps of the documents. To address how topics occur and shift over time, Topics on Time (TOT) model was proposed, by explicitly modeling of time jointly with word co-occurrence patterns [Wang and McCallum, 2006]. The model is shown in Figure 1.




Figure 1: TOT Model

In the model, there are D documents. Each document d contains N_d words $w_{d1}, w_{d2}, \dots, w_{dN_d}$. Each word w_{di} has a timestamp $t_{di} \in (0, 1)$, indicating when the document is generated in a relative time scale $(0, 1)$. All words in the same document have the same timestamp. There are K topics (also $T = K$ topics for the notation in the paper and Figure 1) in the document corpora. Each topic follows a multinomial distribution ϕ over the V words in the vocabulary. Each document follows a multinomial distribution θ over the K topics. The prior distribution for ϕ and θ are Dirichlet distributions with parameters β and α respectively. For each topic k , the temporal occurrence follows a Beta distribution $Beta(\psi_{k1}, \psi_{k2})$, where $\psi_k = (\psi_{k1}, \psi_{k2})$ and we use $\psi \in \mathbb{R}_+^{K \times 2}$ to denote ψ_k for all topics. Each word w_{di} and its timestamp t_{di} are assumed to be generated from a topic, with a topic label $z_{di} \in \{1, \dots, K\}$.

The generative process of this model is described as follows.

1. Draw K multinomials ϕ_k from a Dirichlet prior β , one for each topic k .
2. For each document d ,
 - Draw a multinomial θ_d from a Dirichlet prior α ;
 - For each word w_{di} in document d ,
 - (a) Sample a topic z_{di} from multinomial θ_d ;
 - (b) Sample a word w_{di} from multinomial $\phi_{z_{di}}$;
 - (c) Sample a timestamp t_{di} from Beta $\psi_{z_{di}}$.

We use variational EM to approximate the posterior of latent variables and learn model parameters. To do this, a mean field variational distribution needs to be defined, which is parameterized by some parameters called variational parameters. The variational EM algorithm iteratively performs two steps: 1) in the E step, variational parameters are updated; 2) in the M step, model parameters are optimized. Same as in the paper, we consider α and β are predefined fixed hyperparameters with no need to update. Therefore, in M step, only the other model parameters are optimized. The pseudo-code for the proposed algorithm is shown in Algorithm 1.

Algorithm 1 Pseudo-code of variational EM algorithm for TOT model

- 1: **Input:** Observations, Topic number K , MaxIter, and other optional parameters
 - 2: **Output:** Posterior distributions for latent variables, optimized model parameters
 - 3: Initialize parameters;
 - 4: Compute and record ELBO with initial parameters
 - 5: **for** $k \leftarrow 1$ to $MaxIter$ **do**
 - 6: Update variational variables ▷ Stage 1: E-Step
 - 7: Update ψ with projected Newton method ▷ Stage 2: M-Step
 - 8: Compute and record ELBO
 - 9: **end for**
-

In the TOT model, θ, \mathbf{z}, ϕ are latent variables and ψ is the model parameter to be learned. As a start, we use mean-field variational inference and the variational distribution has the form:

$$q(\theta, \phi, \mathbf{z} | \gamma, \lambda, \pi) = \prod_{k=1}^K q(\phi_k | \lambda_k) \prod_{d=1}^D [q(\theta_d | \gamma_d) \prod_{n=1}^{N_d} q(z_{dn} | \pi_{dn})], \quad (1)$$

where γ, λ, π are variational parameters that need to be updated in E-step.

Next, we write out the joint distribution of latent and observed variables:

$$p(\mathbf{x}, \mathbf{t}, \phi, \theta, \mathbf{z} | \alpha, \beta, \psi) = \prod_{k=1}^K p(\phi_k | \beta) \prod_{d=1}^D [p(\theta_d | \alpha) \prod_{n=1}^{N_d} p(z_{dn} | \theta_d) p(x_{dn} | z_{dn}, \phi) p(t_{dn} | z_{dn}, \psi)] \quad (2)$$

Given Eq. 1 and 2, we can write out ELBO with:

$$\text{ELBO} = \mathbb{E}_{q(\theta, \phi, \mathbf{z})} [\log p(\mathbf{x}, \mathbf{t}, \phi, \theta, \mathbf{z}) - \log q(\theta, \phi, \mathbf{z})]. \quad (3)$$

Variational EM basically maximizes ELBO w.r.t. variational parameters and model parameters in E- and M-step respectively.

Questions:

1. [10 points] Update variational parameters
Derive the update equations of variational parameters, and also specify their distributions. Here you can directly use the conclusion below for the derivation.

$$q_j^* \propto \exp\{\mathbb{E}_{q_{-j}}[\log p(\mathbf{x}, \mathbf{t}, \phi, \theta, \mathbf{z})]\},$$

where $\mathbb{E}_{q_{-j}}$ denotes expectation over all latent variables excluding variable j .

2. [10 points] Update model parameters
Derive the update equations of model parameters, as mentioned before, there is no need to update α and β . For the updating rule of ψ , please be careful that ψ should be constrained as positive. Hint: For a problem with positive solution ($x > 0$), a projected Newton method could be applied:

$$y = (x - H^{-1}g)_+$$

$$x^+ = x + \lambda(y - x)$$

where x is the current variable, y is the projected update, g and H are gradient and Hessian matrix respectively, λ is the step size and x^+ is the updated variable. $(\cdot)_+$ is defined as $s_+ := \max(0, s)$.

3. [20 points] Detailed variational lower bound

Based on the variational distributions, expand Eq. 3 to obtain detailed variational lower bound. The result should be as specific as possible, that is, it can be directly used in the implementation.

Hint: the problem is designed based on the paper [Wang and McCallum, 2006]. In the paper, Gibbs sampling was used for posterior inference, and here we are working with variational inference. You may gain better understanding of the model and get some ideas of how to solve the problem by reading the paper.

3 Markov Chain Monte Carlo [20 pts] (Chirag)

Nowadays, statistical modelling of sport data has become an important part of sports analytics and is often a critical reference for the managers in their decision-making process. In this part, we will work on a real world example in professional sports. Specifically, we are going to use the data from the 2013-2014 Premier League, the top-flight English professional league for men’s football clubs, and build a predictive model on the number of goals scored in a single game by the two opponents. A Bayesian hierarchical model is a good candidate for this kind of modeling task. We model each team’s strength (both attacking and defending) as latent variables. Then in each game, the goals scored by the home team is a random variable conditioned on the attacking strength of the home team and the defending strength of the away team. Similarly, the goals scored by the away team is a random variable conditioned on the attack strength of the away team and the defense strength of the home team. Therefore, the distribution of the scoreline of a specific game is dependent on the relative strength between the home team A and the away team B, which also depends on the relative strength between those teams with their other opponents.

Table 1: 2013-2014 Premier League Teams

Index	0	1	2	3	4
Team	Arsenal	Aston Villa	Cardiff City	Chelsea	Crystal Palace
Index	5	6	7	8	9
Team	Everton	Fulham	Hull City	Liverpool	Manchester City
Index	10	11	12	13	14
Team	Manchester United	Newcastle United	Norwich City	Southampton	Stoke City
Index	15	16	17	18	19
Team	Sunderland	Swansea City	Tottenham Hotspurs	West Bromwich Albion	West Ham United

Here we consider using the same model as described by Baio and Blangiardo [2010]. The Premier League has 20 teams, and we index them as in Table 1. Each team would play 38 matches every season (playing each of the other 19 teams home and away), which totals 380 games in the entire season. For the g -th game, assume that the index of home team is $h(g)$ and the index of the away team is $a(g)$. The observed number of goals (y_{g0}, y_{g1}) of home and away team is modeled as independent Poisson random variables:

$$y_{gj} | \theta_{gj} \sim \text{Poisson}(\theta_{gj}), \quad j = 0, 1 \quad (4)$$

where $\theta = (\theta_{g0}, \theta_{g1})$ represents the scoring intensity in the g -th game for the team playing at home ($j = 0$) and away ($j = 1$), respectively. We use a log-linear model for the θ s:

$$\log \theta_{g0} = \text{home} + \text{att}_{h(g)} - \text{def}_{a(g)} \quad (5)$$

$$\log \theta_{g1} = \text{att}_{a(g)} - \text{def}_{h(g)} \quad (6)$$

Note that team strength is broken into attacking and defending strength. And home represents home-team advantage, and in this model is assumed to be constant across teams. The prior on the home is a normal distribution:

$$\text{home} \sim \mathcal{N}(0, \tau_0^{-1}) \quad (7)$$

where we set the precision $\tau_0 = 0.0001$.

The team-specific attacking and defending effects are modeled as:

$$\text{att}_t \sim \mathcal{N}(\mu_{\text{att}}, \tau_{\text{att}}^{-1}) \quad (8)$$

$$\text{def}_t \sim \mathcal{N}(\mu_{\text{def}}, \tau_{\text{def}}^{-1}) \quad (9)$$

We use conjugate priors as the hyper-priors on the attack and defense means and precisions:

$$\mu_{att} \sim \mathcal{N}(0, \tau_1^{-1}) \quad (10)$$

$$\mu_{def} \sim \mathcal{N}(0, \tau_1^{-1}) \quad (11)$$

$$\tau_{att} \sim \text{Gamma}(\alpha, \beta) \quad (12)$$

$$\tau_{def} \sim \text{Gamma}(\alpha, \beta) \quad (13)$$

where the precision $\tau_1 = 0.0001$, and we set parameters $\alpha = \beta = 0.1$.

This hierarchical Bayesian model can be represented using a directed acyclic graph as shown in Figure 2.




Figure 2: The DAG representation of the hierarchical Bayesian model. Figure adapted from ?.

The goals of each game are $\mathbf{y} = \{y_{gj} | g = 0, 1, \dots, 379, j = 0, 1\}$ are the observed variables, and parameters $\boldsymbol{\theta} = \{home, att_0, def_0, \dots, att_{19}, def_{19}\}$ and hyper-parameters $\boldsymbol{\eta} = (\mu_{att}, \mu_{def}, \tau_{att}, \tau_{def})$ are unobserved variables that we need to make inference on. To ensure identifiability, we enforce a corner constraint on the parameters (pinning one team's parameters to 0,0). Here we use the first team as reference and assign its attacking and defending strength to be 0:

$$att_0 = def_0 = 0 \quad (14)$$

In this question, we want to estimate the posterior mean of the attacking and defending strength for each team, i.e. $\mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{\eta} | \mathbf{y})}[att_i]$, $\mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{\eta} | \mathbf{y})}[def_i]$, and $\mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{\eta} | \mathbf{y})}[home]$.

1. [4 points] Find the joint likelihood $p(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\eta})$.
2. [4 points] Write down the Metropolis-Hastings algorithm for sampling from posterior $p(\boldsymbol{\theta}, \boldsymbol{\eta} | \mathbf{y})$, and derive the acceptance function for a proposal distribution of your choice (e.g. isotropic Gaussian).

3. [12 points] Implement the Metropolis-Hastings algorithm to inference the posterior distribution. The data can be found from http://cs.cmu.edu/~pradeepr/708/resources/premier_league_2013_2014.dat, which contains a 380×4 matrix. The first column is the number of goals y_{g0} scored by the home team, the second column is the number of goals y_{g1} scored by the away team, the third column is the index for the home team $h(g)$, and the fourth column is the index for the away team $a(g)$.
- Use an isotropic Gaussian proposal distribution $\mathcal{N}(0, \sigma^2 I)$ and use 0.1 as the starting point.
 - Run the MCMC chain for 5000 steps to burn in and then collect 5000 samples with t steps in between (i.e., run M-H for $5000t$ steps and collect only each t -th sample). This is called thinning, which reduces the autocorrelation of the MCMC samples introduced by the Markovian process. The parameter sets are $\sigma = 0.005, 0.05, 0.5$, and $t = 1, 5, 20, 50$.
 - Plot the trace plot of the burn in phase and the MCMC samples for the latent variable *home* using proposal distributions with different σ and t .
 - Estimate the rejection ratio for each parameter setting, report your results in a table.
 - Comment on the results. Which parameter setting worked the best for the algorithm?
 - Use the results from the optimal parameter setting:
 - (a) plot the posterior histogram of variable *home* from the MCMC samples.
 - (b) plot the estimated attacking strength $\mathbb{E}_{p(\theta, \eta | \mathbf{y})}[att_i]$ against the estimated defending strength $\mathbb{E}_{p(\theta, \eta | \mathbf{y})}[def_i]$ for each the team in one scatter plot. Please make sure to identify the team index of each point on your scatter plot using the index to team mapping in Table 1.
4. [0 points] Despite what the data says, conclude that Manchester United is the best team in the Premier League!

Note: You are free to use Python or MATLAB for your implementation. You are NOT allowed to use any existing implementations of Metropolis-Hastings in this problem. Please include all the required results (figures + tables) in your writeup PDF submission, as well as submit your code to Gradescope separately.

References

- G. Baio and M. Blangiardo. Bayesian hierarchical model for the prediction of football results. *Journal of Applied Statistics*, 37(2):253–264, 2010.
- X. Wang and A. McCallum. Topics over time: A non-markov continuous-time model of topical trends. 2006.