CS5560: Practice Problems

January 4, 2022

Week 2

Topics: Definition & Examples of Exponential Family Models

1. Exercise 9.2 in [Murphy(2012)].
2. Exercise 8.1(a), Proof of theorems 8.1,8.3 in [Koller and Friedman(2009)].

Week 3

Topics: Exponential Family of Discriminative Models and examples.

1. Exercise 8.4 in [Murphy(2012)].
2. Consider a (exp family) discriminative model with $\mathcal{Y} = \{0, 1, 2, \ldots\}$ and $\phi(y) \equiv y$. Simplify the expression for the likelihood function (as much as possible). This model is known as the Poisson Regression model, which is widely used in bio-statistical applications (e.g., y is number of diseases or number of genomic reads etc.).

Week 4

Topics: Method of Moments, Sufficient Statistic, and Maximum Likelihood Estimation

1. Exercises 3.1, 3.6, 3.8, 3.11 (a)-(b), (proof of) theorem4.1.1 in [Murphy(2012)].
2. Exercises 17.1-17.5, 8.1(b), 8.3, 8.4, 8.6 in [Koller and Friedman(2009)].

Week 5

Topics: MLE — Convexity and Gradient, Importance Sampling, MH algorithm

1. Write down MLE optimization problem of (univariate) Gaussian in terms of the usual mean, and variance. Show that this objective is not a convex function. Whereas, from lectures we know that, the same objective is convex in terms of the natural parameters!

2. Let the target likelihood be a Dirichlet[1] with number of categories as $k \geq 2$ and concentration parameters $\alpha_i > 0$. Now estimate the mean vector, $\hat{\mu}$, of this Dirichlet using the Importance Sampling Technique with the proposal likelihood as uniform in $\Delta_k \equiv \left\{x \in \mathbb{R}^k \mid x \geq 0, \mathbf{1}^\top x = 1\right\}$ and m samples from the proposal. Plot the error in this estimate, $\|\hat{\mu} - \mu\|$, vs m for a fixed k. Here μ is the exact mean vector whose i^{th} entry is given by $\mu_i = \frac{\alpha_i}{\sum_{i=1}^k \alpha_i}$. Also, plot the error in this estimate, $\|\hat{\mu} - \mu\|$, vs k for a fixed m. For both these plots fix some values of α_is. Report your observations from these plots.

Week 6

Topics: MCMC methods: MH and Gibbs sampling

1. Answer question 4 titled "MCMC 2" in `https://www.stat.purdue.edu/~varao/STAT545/midterm15_2.pdf`.

2. Repeat question 2 in week 5 using MH and (conditional) proposal as Gaussian with mean as the previous iterate. Repeat the same with Gibbs sampling and (joint) proposal as Gaussian. Using all the plots obtained, summarize your observations about the performance of these methods.

Week 7

Topics: Training Discriminative Models, Inference algorithms for clustering, classification and regression.

1. Study entire sections 3.5, 4.2 (except 3.5.1.2, 3.5.2, 3.5.4, and other starred subsections) from [Murphy(2012)].

2. Exercises 4.7, 4.19, 4.20, 4.21, 4.22, 4.23, 7.2-7.6, 7.9, 8.1, 8.3, 8.4 in [Murphy(2012)].

3. Write code (in your favourite language) that solves a generic regression problem over Euclidean input spaces. In particular, use an exponential family discriminative model, with input-feature-map as identity (assume Euclidean input vectors). Your code must handle any user-specified output-feature-map (recall that in regression, the outputs/labels are real-valued). Training is done using MCLE and label inferred for a new input is the posterior mean estimated using Gibbs sampling. Compare

[1] `https://en.wikipedia.org/wiki/Dirichlet_distribution`.

accuracy the model with your output-feature-map vs Linear Regression (i.e., quadratic output feature map; refer section in [Murphy(2012)]) on few benchmark datasets: `https://archive.ics.uci.edu/ml/datasets.php?format=mat&task=reg&att=num&area=&numAtt=less10&numIns=100to1000&type=mvar&sort=nameUp&view=table`. You can now play with output-feature-map and see how it effects the accuracy.

Week 8

Topics: Inference in Supervised Learning with/without Missing Data, Bayesian Models, Computing posterior

1. Read sections 3.3-3.3.3, 3.4-3.4.3, 3.5.1.2, 4.6.3-4.6.3.3, 7.6-7.6.1 in [Murphy(2012)].
2. Exercises 3.4, 3.7, 3.9, 3.10a, 3.11c-d, 3.20, 4.11, 5.9 in [Murphy(2012)].

Week 9

Topics: Conjugate prior, posterior predictive/generative (Bayesian Averaging), MAP estimation (hybrid models)

1. Read sections 3.3.4, 3.4.4, 3.5.1.2, 4.6.3.4, 4.6.3.6, 5.2.1, 7.5, 7.6, 8.3.6, 8.4, 9.2.5, 9.3 from [Murphy(2012)].
2. Exercises 9.1, 8.6, 8.5, 7,8(b), 7.9, 7.10.

Week 10

Topics: Model Selection, Maximum Marginal Likelihood, Hierarchical Bayes Models (self-study), Cross-Validation (self-study)

1. Exercises 5.1, 5.7, 5.8a-d in [Murphy(2012)].
2. Read sections 5.3, 5.5, 5.6 (except starred sections) in [Murphy(2012)].

Week 11

Topics: Bayes nets, Local Markov conditions, chain rule, I-maps

1. Exercises 3.1-3.4, 3.6-3.7, 3.15, in [Koller and Friedman(2009)].
2. Exercises 10.1, 10.5, 10.6, in [Murphy(2012)].

References

[Koller and Friedman(2009)] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques.* MIT Press, 2009.

[Murphy(2012)] Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective.* The MIT Press, 2012. ISBN 0262018020, 9780262018029.