Common Instructions: You must submit answers via Google Classroom's assignment submission interface before the appropriate deadline. The format of the submission file is `pdf` (for e.g., scan your hand-written answers using mobile app and make a pdf from it). Please write legibly and concisely. Always simplify your answers as much as possible. For coding/simulations you are free to use any programming language.

Week1

1. Exercises 2.1, 2.4, 2.7, 2.10, 2.16, 2.17 in Kevin Murphy's textbook.

* Exercises 27, 37, 38 in Chapter 2 in Sheldon Ross's book[1].

Week1a

1. In section 2 in Murphy's book many common discrete (sec 2.3), and common continuous (sec 2.4) distributions are presented. For each of these cases implement the parameter estimation algorithm using (any variant of) method of moments. In each case, your code must take as input the training data and output the parameters of the distribution. Also, repeat this exercise for Multivariate Gaussian (sec 2.5.2) and the Dirichlet distribution (sec 2.5.4). You are free to use any numerical equation solver or numerical optimization package. Finally, you need to submit the following (in the same pdf):

   (a) For each univariate distribution plot the estimated parameter values vs. the number of training datapoints. Plot all parameters in the same graph (different colors/markers for each parameter). One graph per distribution. The details of your training data will be communicated separately.

   (b) For each multivariate distribution (Gaussian and Dirichlet), compute the (joint) likelihood of the training dataset[2] with the estimated parameter values. Plot likelihood of training data vs. the number of training datapoints.

   (c) Note your observations (if any) from the plots.

   Please do NOT submit anything else than the above list.

Week2

1. Exercises 3.6, 3.8(a), 3.11(a), 3.11(b), 4.1, 4.2, 4.3, 4.4.

[1] We follow the 10^{th} edition of Sheldon Ross's "Introduction to Probability Models".

[2] From iid assumption, it follows that the (joint) likelihood of the training data is product of likelihoods of each training datapoint (datum).

* Consider the random variable corresponding to the MLE[3]. Analogous to the Central Limit Theorem, can something be said about the asymptotic distribution of the MLE? Provide details.

Week3

1. Implement the parameter estimation algorithm for the Gaussian model based regression model we studied in lectures. Your code should take in as input a regression training dataset (with n_1 input variables and n_2 output variables) and build the regression function that predicts the output for any input. Report accuracy wrt. each output variable separately with your dataset.

Week4

1. Exercise 7.3, 7.4, 7.5, 7.6.

* Exercise 7.7.

Week5

1. Implement the Gaussian Discriminant (Bayes classifier;eqn.4.33) and Naive bayes classifier (use Gaussian model for each input feature/attribute). Compare the accuracy[4] of these two classifiers with that of Logistic Regression[5] on your dataset.

2. Exercise 4.19, 4.21, 4.22.

* Exercise 4.20, 8.6, 8.7.

Week6

1. Exercises 3.2, 3.4, 3.7, 3.10, 3.11, 3.15, 4.14.

* Consider the discriminative predictive model based on exponential family introduced in lectures, with output variables sufficient statistics given by $\psi(\cdot)$ and the canonical parameter corresponding to $\psi_i(y)$ as $w_i^\top \phi(x)$. Show that this model is equivalent to (is reparametrized version of) the generative predictive (posterior) model introduced in lecture with sufficient statistics as $\phi(x) \otimes \psi(y)$. Note that, because of this equivalence, it

[3] Because samples are random, the estimate is also random.

[4] Accuracy is percentage of correctly classified test data samples.

[5] Download code from `https://www.csie.ntu.edu.tw/~cjlin/liblinear/` and use "train -s 0 -B 1 -C trainingsetfile modelfile" command and use "predict testfile modelfile outputfile".

makes sense to call $\phi(\cdot)$ as the sufficient statistics of inputs even in the discriminative model!

* Exercise 4.11.

Week7

1. Repeat the technical derivations in sections 3.5.1.2-3.5.2. Clearly write down the final simplified expressions for MLE, MAP, Bayesian Bayes classifier. Contrast them and write down your comments.
2. Repeat the technical derivation in section 7.6.3.1, while giving details of non-trivial steps that may be left as exercise in the book.
3. Read section 8.4 carefully.
4. Exercises 4.19, 7.8.

* Exercises 5.1, 5.7, 5.9, 7.10.

Week8

1. Exercises 5.8a-5.8d, 6.1, 6.2a.
2. Read section 5.3.1, 5.5.1.

* Exercise 5.7

Week9

1. Implement the following two algorithms for solving the MLE problem associated with the mixture of Gaussians model.
   (a) Gradient descent algorithm (using Armijo's rule[6] for step-size)
   (b) EM algorithm

   Submit a single graph consisting of two plots: Value of negative log-likelihood of the training data vs. iterations with each of the above two implementations. Use the same data as in "Week5" (You must now ignore the outputs/labels in the training data and use only the inputs).

[6] `https://en.wikipedia.org/wiki/Wolfe_conditions`

Week10

1. Exercises 17.3, 17.4 in your textbook.

* Write code for (unsupervised) training of HMM and Viterbi algorithm. Verify your code on the three toy datasets in `https://www.cs.princeton.edu/courses/archive/fall06/cos402/hw/hw5/code.zip`. The description of these datasets is here: `https://www.cs.princeton.edu/courses/archive/fall06/cos402/hw/hw5/hw5.html`. Submit a small paragraph summarizing your findings from these simulations.