

CPSC 540 Assignment 3 (due February 22 at midnight)

The assignment instructions are the same as for the previous assignment.

1. Name(s):
2. Student ID(s):

1 Discrete and Gaussian Variables

1.1 MLE for General Discrete Distribution

Consider a density estimation task, where we have two variables ($d = 2$) that can each take one of k discrete values. For example, we could have

$$X = \begin{bmatrix} 1 & 3 \\ 4 & 2 \\ k & 3 \\ 1 & k-1 \end{bmatrix}.$$

The likelihood for example x^i under a general discrete distribution would be

$$p(x_1^i, x_2^i | \Theta) = \theta_{x_1^i, x_2^i},$$

where θ_{c_1, c_2} gives the probability of x_1 being in state c_1 and x_2 being in state c_2 , for all the k^2 combinations of the two variables. In order for this to define a valid probability, we need all elements θ_{c_1, c_2} to be non-negative and they must sum to one, $\sum_{c_1=1}^k \sum_{c_2=1}^k \theta_{c_1, c_2} = 1$.

1. Given n training examples, [derive the MLE for the \$k^2\$ elements of \$\Theta\$](#) .
2. If we had separate parameter θ_{c_1} and θ_{c_2} for each variables, a reasonable choice of a prior would be a product of Dirichlet distributions,

$$p(\theta_{c_1}, \theta_{c_2}) \propto \theta_{c_1}^{\alpha_{c_1}-1} \theta_{c_2}^{\alpha_{c_2}-1}.$$

For the general discrete distribution, a prior encoding the same assumptions would be

$$p(\theta_{c_1, c_2}) \propto \theta_{c_1, c_2}^{\alpha_{c_1} + \alpha_{c_2} - 2}.$$

[Derive the MAP estimate under this prior](#) (assuming we use k^2 variables to parameterize Θ)

3. We often use discrete distributions as parts of more-complicated distributions (like mixture models and graphical models). In these cases we often need to fit a weighted NLL for the form

$$f(\Theta) = - \sum_{i=1}^n v_i \log p(x_1^i, x_2^i | \Theta),$$

with n non-negative weights. [What is the MLE for the \$k^2\$ elements of \$\Theta\$ under this weighting.](#)

Hint: it may be convenient to write the discrete likelihood for an example i in the form

$$p(x^i | \Theta) = \prod_{c \in [k]^2} \theta_c^{\mathbb{I}[x^i=c]},$$

where c is a vector containing (c_1, c_2) , $[x^i = c]$ evaluates to 1 if all elements are equal, and $[k]^2$ is all ordered pairs (c_1, c_2) . You can use a Lagrangian to enforce the sum-to-1 constraint on the log-likelihood, and you may find it convenient to define $n_c = \sum_{i=1}^n \mathcal{I}[x^i = c]$.

1.2 Gaussian Self-Conjugacy and Posterior Convergence Rate

Consider n IID samples x^i distributed according to a Gaussian with mean μ and covariance $\sigma^2 I$,

$$x^i \sim \mathcal{N}(\mu, \sigma^2 I).$$

Assume that μ itself is distributed according to a Gaussian

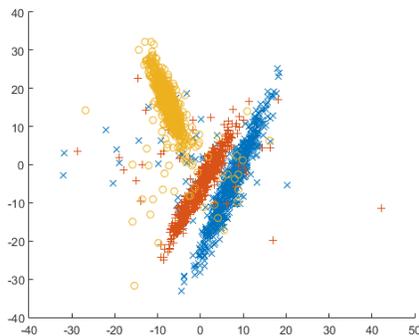
$$\mu \sim \mathcal{N}(\mu_0, \Sigma_0),$$

with mean μ_0 and (positive-definite) covariance Σ_0 . In this setting, the posterior for μ also follows a Gaussian distribution.¹

1. Derive the form of the posterior distribution, $p(\mu | X, \sigma^2, \mu_0, \Sigma_0)$.
Hint: the posterior is a product of Gaussian densities.
2. To measure the speed at which the posterior converges to μ , we can consider the variance of the posterior. Consider a version of the previous question where $d = 1$: how many examples n (in O -notation) do we need before we have $|\Sigma| < \epsilon$ (for the posterior variance Σ).

1.3 Generative Classifiers with Gaussian Assumption

Consider the 3-class classification dataset in this image:



In this dataset, we have 2 features and each colour represents one of the classes. Note that the classes are highly-structured: the colours each roughly follow a Gaussian distribution plus some noisy samples.

Since we have an idea of what the features look like for each class, we might consider classifying inputs x using a *generative classifier*. In particular, we are going to use Bayes rule to write

$$p(y^i = c | x^i, \Theta) = \frac{p(x^i | y^i = c, \Theta) \cdot p(y^i = c | \Theta)}{p(x^i | \Theta)},$$

where Θ represents the parameters of our model. To classify a new example \tilde{x}^i , generative classifiers would use

$$\hat{y}^i = \arg \max_{y \in \{1, 2, \dots, k\}} p(\tilde{x}^i | y^i = c, \Theta) p(y^i = c | \Theta),$$

¹We say that the Gaussian distribution is the ‘conjugate prior’ for the Gaussian mean parameter (we’ll formally discuss conjugate priors later in the course). Another reason the Gaussian distribution is important is that is the only (non-trivial) continuous distribution that has this “self-conjugacy” property.

where in our case the total number of classes k is 3.² Modeling $p(y^i = c \mid \Theta)$ is easy: we can just use a k -state categorical distribution,

$$p(y^i = c \mid \Theta) = \theta_c,$$

where θ_c is a single parameter for class c . The maximum likelihood estimate of θ_c is given by n_c/n , the number of times we have $y^i = c$ (which we've called n_c) divided by the total number of data points n .

Modeling $p(x^i \mid y^i = c, \Theta)$ is the hard part: we need to know the *probability of seeing the feature vector x^i given that we are in class c* . This corresponds to solving a density estimation problem for each of the k possible classes. To make the density estimation problem tractable, we'll assume that the distribution of x^i given that $y^i = c$ is given by a $\mathcal{N}(\mu_c, \Sigma_c)$ Gaussian distribution for a class-specific μ_c and Σ_c ,

$$p(x^i \mid y^i = c, \Theta) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma_c|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x^i - \mu_c)^T \Sigma_c^{-1} (x^i - \mu_c)\right).$$

Since we are distinguishing between the probability under k different Gaussians to make our classification, this is called *Gaussian discriminant analysis* (GDA). In the special case where we have a constant $\Sigma_c = \Sigma$ across all classes it is known as *linear discriminant analysis* (LDA) since it leads to a linear classifier between any two classes (while the region of space assigned to each class forms a convex polyhedron as in k -means clustering and softmax classification). Another common restriction on the Σ_c is that they are diagonal matrices, since this only requires $O(d)$ parameters instead of $O(d^2)$ (corresponding to assuming that the features are independent univariate Gaussians given the class label). Given a dataset $\mathcal{D} = \{(x^i, y^i)\}_{i=1}^n$, where $x^i \in \mathbb{R}^d$ and $y^i \in \{1, \dots, k\}$, the maximum likelihood estimate (MLE) for the μ_c and Σ_c in the GDA model is the solution to

$$\arg \max_{\mu_1, \mu_2, \dots, \mu_k, \Sigma_1, \Sigma_2, \dots, \Sigma_k} \prod_{i=1}^n p(x^i \mid y^i, \mu_{y^i}, \Sigma_{y^i}).$$

This means that the negative log-likelihood will be equal to

$$\begin{aligned} -\log p(X \mid y, \Theta) &= -\sum_{i=1}^n \log p(x^i, y^i \mid \mu_{y^i}, \Sigma_{y^i}) \\ &= \sum_{i=1}^n \frac{1}{2} (x^i - \mu_{y^i})^T \Sigma_{y^i}^{-1} (x^i - \mu_{y^i}) + \frac{1}{2} \sum_{i=1}^n \log |\Sigma_{y^i}| + \text{const.} \end{aligned}$$

In class we derived the MLE for this model under the assumption that we use full covariance matrices and that each class has its own covariance.

1. [Derive the MLE for the GDA model under the assumption of *common diagonal covariance* matrices, \$\Sigma_c = D\$ \(\$d\$ parameters\).](#) (Each class will have its own mean μ_c .)
2. [Derive the MLE for the GDA model under the assumption of *individual scale-identity* matrices, \$\Sigma_c = \sigma_c^2 I\$ \(\$k\$ parameters\).](#)
3. When you run `example_generative` it loads a variant of the dataset in the figure that has 12 features and 10 classes. This data has been split up into a training and test set, and the code fits a k -nearest neighbour classifier to the training set then reports the accuracy on the test data (around $\sim 63\%$ test error). The k -nearest neighbour model does poorly here since it doesn't take into account the Gaussian-like structure in feature space for each class label. Write a function `gda` that fits a GDA model to this dataset (using individual full covariance matrices). [Hand in the function and report the test set accuracy.](#)
4. In this question we would like to replace the Gaussian distribution of the previous problem with the more robust multivariate-t distribution so that it isn't influenced as much by the noisy data.

²The denominator $p(\tilde{x}^i \mid \Theta)$ is irrelevant to the classification since it is the same for all y .

Unlike the previous case, we don't have a closed-form solution for the parameters. However, if you run `example_student` it generates random noisy data and fits a multivariate-t model. By using the `studentT` model, write a new function `tda` that implements a generative model that is based on the multivariate-t distribution instead of the Gaussian distribution. [Report the test accuracy with this model.](#)

Hints: you may be able to substantially simplify the notation in the MLE derivations if you use the notation $\sum_{i \in y_c}$ to mean the sum over all values i where $y^i = c$. Similarly, you can use n_c to denote the number of cases where $y_i = c$, so that we have $\sum_{i \in y_c} 1 = n_c$. Note that the determinant of a diagonal matrix is the product of the diagonal entries, and the inverse of a diagonal matrix is a diagonal matrix with the reciprocals of the original matrix along the diagonal.

For the implementation you can use the result from class regarding the MLE of a general multivariate Gaussian. At test time for GDA, you may find it more numerically reasonable to compare log probabilities rather than probabilities of different classes, and you may find it helpful to use the `logdet` function to compute the log-determinant in a more numerically-stable way than taking the log of the determinant. (Also, don't forget to center at training and test time.)

For the last question, you may find it helpful to define an empty array that can be filled with k `DensityModel` objects using:

```
subModel = Array{DensityModel}(undef,k)
```

2 Mixture Models and Expectation Maximization

2.1 Categorical Mixture Model

Consider a density estimation with examples $x^i \in \{1, 2, \dots, k\}^d$ representing a set of d categorical variables. In this setting, a natural way to model an individual variable x_j^i would be with a categorical distribution,

$$p(x_j^i = c \mid \theta_{j,1}, \theta_{j,2}, \dots, \theta_{j,k}) = \theta_{j,c},$$

so the joint distribution would be

$$p(x^i \mid \theta_{j,1}, \theta_{j,2}, \dots, \theta_{j,k}) = \prod_{j=1}^d \theta_{j,x_j^i},$$

where all $\theta_{j,c} \geq 0$ and $\sum_{c=1}^k \theta_{j,c} = 1$ for each feature j . However, if we assume this structure for each variable then the variables would be independent. One way to model dependent count variables would be with a mixture of m independent categorical distributions,

$$p(x^i \mid \Theta) = \sum_{c=1}^m p(z^i = c \mid \Theta^t) p(x^i \mid z^i = c, \Theta^t) = \sum_{c=1}^m \pi_c \prod_{j=1}^d \theta_{j,x_j^i}^c,$$

where:

- π_c is the probability that the examples comes from mixture c , $p(z^i = c \mid \Theta) = \pi_c$.
- $\theta_{j,c'}^c$ is the probability that x_j^i is c' for examples from mixture c , $p(x_j^i = c' \mid z^i = c, \Theta)$.
- We use Θ as the set containing all the π_c and $\theta_{j,c'}^c$ values.

[Derive the EM update for this mixture model](#) (treating the z^i as missing values).

Hint: a lot of the work has been done for you in the EM notes on the course webpage.

2.2 Gaussian Mixture Model Implementation

The script *example_gaussian* fits a Gaussian distribution to a dataset that is not uni-modal, and giving a bad fit. Implement the EM for fitting a mixture of Gaussians to this data. [Hand in your code and the updated plot when using a mixture of 3 Gaussians.](#)

Hint: you may want to start by implementing the PDF (“predict”) function; you can then use this with the monotonicity of EM to debug your implementation.

It is possible that Σ_k^{t+1} may not be positive-definite, if you encounter this the standard fixes are to remove such clusters or use a MAP estimate for Σ where you add a small multiple of the identity matrix to the estimate.

3 Very-Short Answer Questions

Give a short and concise 1-sentence answer to the below questions.

1. If we have n samples of d discrete features that can each take k values, what is the cost of computing the MLE if we use a product of independent distributions? (Try to find the minimum dependence on n , d , and k .)
2. Describe a setting where it would make sense to use the general discrete distribution for density estimation with binary features.
3. What is the relationship between using a product of independent Gaussians and using a multivariate Gaussian.
4. Suppose we run the graphical LASSO method and it returns a tri-diagonal precision matrix. What would the graph look like?
5. Suppose we have a lot of extreme outliers in our dataset. Why is this less of a problem for a mixture of Gaussians than if we use a single Gaussian?
6. If we used GDA with where each class has its own diagonal covariance matrix, would this give linear or quadratic decision boundaries? What about if the covariance matrices are multiples of the identity matrix?
7. What is an advantage and a disadvantage of the EM algorithm over the imputation approach to handling MAR variables
8. What is an advantage of the Epanechnikov kernel over the Gaussian kernel.
9. What is an advantage and a disadvantage of parameter tying?
10. What is the difference between computing marginals and computing the stationary distribution of a Markov chain.

4 Project Proposal

For the final part of this assignment, you must [submit a project proposal](#) for your course project. The proposal should be a maximum of 2 pages (and 1 page or half of a page is ok if you can describe your plan concisely). The proposal should be written for me and the TAs, so you don’t need to introduce any ML background but you will need to introduce non-ML topics. The projects must be done in groups of 2-3. If you are doing your assignment in a group that is different from your project group, only 1 group member should include the proposal as part of their submission (we’ll do the merge across assignments, and this means that assignments could have multiple proposals). Please state clearly who is involved with each project proposal.

There is quite a bit of flexibility in terms of the type of project you do, as I believe there are many ways that people can make valuable contributions to research. However, note that ultimately the project will have three parts:

1. A very short paper review summarizing the pros and cons of a particular paper on the topic (due with Assignment 4).
2. A short literature review summarizing at least 10 papers on a particular topic (due with Assignment 5).
3. A final report containing at most 6 pages of text (the actual document can be longer due to figures, tables, references, and proofs) that emphasizes a particular “contribution” (i.e., what doing the project has added to the world).

The reason for this, even though it’s strange for some possible projects, is that this is the standard way that results are communicated to the research community.

The three main ingredients of the project proposal are:

1. What problem you are focusing on.
2. What you plan to do.
3. What will be the “contribution”.

Also, note that for the course project that negative results (i.e., we tried something that we thought we would work in a particular setting but it didn’t work) are acceptable (and often unavoidable).

Here are some standard project “templates” that you might want to follow:

- **Application bake-off:** you pick a specific application (from your research, personal interests, or maybe from Kaggle) or a small number of related applications, and try out a bunch of techniques (e.g., random forests vs. logistic regression vs. generative models). In this case, the contribution would be showing that some methods work better than others for this specific application (or your contribution could be that everything works equally well/badly).
- **New application:** you pick an application where ML methods where people aren’t using ML, and you test out whether ML methods are effective for the task. In this case, the contribution would be knowing whether ML is suitable for the task.
- **Scaling up:** you pick a specific machine learning technique, and you try to figure out how to make it run faster or on larger datasets (for example, how do we apply kernel methods when n is very large). In this case, the contribution would be the new technique and an evaluation of its performance, or could be a comparison of different ways to address the problem.
- **Improving performance:** you pick a specific machine learning technique, and try to extend it in some way to improve its performance (for example, how can we efficiently use non-linearity within graphical models). In this case, the contribution would be the new technique and an evaluation of its performance.
- **Generalization to new setting:** you pick a specific machine learning technique, and try to extend it to a new setting (for example, making a graphical-model version of random forests). In this case, the contribution would be the new technique and an evaluation of its performance, or could be a comparison of different ways to address the problem.
- **Perspective paper:** you pick a specific topic in ML, read a larger number of papers on the topic, then write a report summarizing what has been done on the topic and what are the most promising directions of future work. In this case, the contribution would be your summary of the relationships between the existing works, and your insights about where the field is going.

- **Coding project:** you pick a specific method or set of methods (like independent component analysis), and build an implementation of them. In this case, the contribution could be the implementation itself or a comparison of different ways to solve the problem.
- **Theory:** you pick a theoretical topic (like the variance of cross-validation or the convergence of proximal stochastic gradient in the non-convex setting), read what has been done about it, and try to prove a new result (usually by relaxing existing assumptions or adding new assumptions). The contribution could be a new analysis of an existing method, or why some approaches to analyzing the method will not work.

The above are just suggestions, and many projects will mix several of these templates together, but if you are having trouble getting going then it's best to stick with one of the above templates. Also note that the above includes topics not covered in the course (like random forests), so there is flexibility in the topic, but the topic should be closely-related to ML.

This question is mandatory but will not be formally marked: it's just a sanity check that you have at least one project idea that fits within the scope of 540 course project, and it's an excuse for you to allocate some time to thinking about the project. Also, there is flexibility in the choice of project topics even after the proposal: if you want to explore different topics you can ultimately choose to do a project that is unrelated to the one in your proposal/paper-review/literature-review, although it will likely be easier to do all 4 parts on the same topic.