

Lecture 21: Gaussian Mixture Model

April 2020

Lecturer: Steven Wu

Scribe: Steven Wu

Maximum Likelihood Estimation

Now we will switch over to a different unsupervised learning problem that aims to model the underlying probability distribution P from which the observed examples are drawn. As we discussed before in the lecture on logistic regression, there is a general principle called *maximum likelihood estimation* (MLE):

- Pick a set of probability models for the data: $\mathcal{P} := \{p_\theta : \theta \in \Theta\}$.
- Given samples x_1, \dots, x_n , pick the model that maximized the likelihood

$$\max_{\theta \in \Theta} L(\theta) = \max_{\theta \in \Theta} \ln \prod_{i=1}^n p_\theta(x_i) = \max_{\theta \in \Theta} \sum_{i=1}^n \ln p_\theta(x_i)$$

Example 0.1 (Coin flips). Heads: $x_i = 1$ and tails $x_i = 0$. The Bernoulli distribution has the parameter—the bias or the probability of heads $\theta \in [0, 1]$. We can write

$$p_\theta(x_i) = \theta^{x_i} (1 - \theta)^{1-x_i}$$

Let $H = \sum_i x_i$ and $T = \sum_i (1 - x_i)$ be the number of heads and tails.

$$L(\theta) = \sum_{i=1}^n (x_i \ln \theta + (1 - x_i) \ln(1 - \theta)) = H \ln \theta + T \ln(1 - \theta).$$

By using the first-order condition, we derive that the solution is

$$\theta = \frac{H}{T + H} = \frac{H}{n}.$$

Gaussian Mixture Model

Gaussian mixture model (GMM) is the following generative model:

- Draw a latent class Y such that $\Pr[Y = j] = \pi_j$
- Then draw X conditioned on Y : $X | Y = j \sim \mathcal{N}(\mu_j, \Sigma_j)$.




Figure 1: Statistician Karl Pearson wanted to understand the distribution of the ratio between forehead breadth and body length for crabs. He fit a mixture of two Gaussians. Figure due to Peter Macdonald.

The parameter $\theta = ((\pi_1, \mu_1, \Sigma_1), \dots, (\pi_k, \mu_k, \Sigma_k))$ and the probability density at each point x is

$$p_\theta(x) = \sum_{j=1}^k p_{\mu_j, \Sigma_j}(x) \pi_j$$

where p_{μ_j, Σ_j} denotes the multivariate Gaussian density function:

$$p_{\mu_j, \Sigma_j}(x) = \frac{1}{\sqrt{(2\pi)^d \det(\Sigma_j)}} \exp\left(-\frac{1}{2}(x - \mu_j)^\top \Sigma_j^{-1} (x - \mu_j)\right)$$

The MLE problem is the maximization of

$$L(\theta) = \sum_{i=1}^n \ln \left(\sum_{j=1}^k p_{\mu_j, \Sigma_j}(x_i) \pi_j \right) = \sum_{i=1}^n \ln \left[\sum_{j=1}^k \frac{\pi_j}{\sqrt{(2\pi)^d \det(\Sigma_j)}} \exp\left(-\frac{1}{2}(x_i - \mu_j)^\top \Sigma_j^{-1} (x_i - \mu_j)\right) \right]$$

Unlike the MLE problem for coin flips, we cannot obtain a closed-form solution here. In fact, MLE for GMM is known to be NP-hard, but we will introduce a well-known heuristic in this lecture.

Expectation and Maximization

We will introduce the method of *expectation and maximization* (EM) for solving the MLE problem for GMM. We will introduce a set of auxiliary variables in matrix form $R \in \mathbb{R}^{n \times k} := R \in [0, 1]^{n \times k} : R \mathbf{1}_k = \mathbf{1}_n$, such that each R_{ij} that defines the probability that each example x_i to the j -th Gaussian distribution. Let us define augmented likelihood as

$$L(\theta, R) = \sum_{i=1}^n \sum_{j=1}^k R_{ij} \ln \left(\sum_{j=1}^k \frac{p_\theta(x_i, y_i = j)}{R_{ij}} \right)$$

Note that we can write the original likelihood function as:

$$\begin{aligned} L(\theta) &= \sum_{i=1}^n \ln(p_\theta(x_i)) \\ &= \sum_{i=1}^n 1 \ln(p_\theta(x_i)) \\ &= \sum_{i=1}^n \sum_{j=1}^k p_\theta(y_i = j | x_i) \ln(p_\theta(x_i)) \\ &= \sum_{i=1}^n \sum_{j=1}^k p_\theta(y_i = j | x_i) \ln \left(\frac{p_\theta(x_i, y_i = j)}{p_\theta(y_i = j | x_i)} \right) \quad (\text{Bayes rule: } p_\theta(x) p_\theta(y|x) = p_\theta(x, y)) \end{aligned}$$

This means, if we set $R_{ij} = p_\theta(y_i = j | x_i)$, then $L(\theta, R) = L(\theta)$.

The EM method performs the following alternating optimization over θ and R to maximize augmented likelihood function $L(\theta, R)$. the algorithm first initialize $(\pi_0)_j = 1/k$, $(\Sigma_0)_j = I$, and $(\mu_0)_j$ randomly. Then over iterations $t = 1, \dots, T$:

- **E-step:** set $(R_t)_{ij} := p_{\theta_{t-1}}(y_i = j | x_i)$. This means

$$(R_t)_{ij} = p_{\theta_{t-1}}(y_i = j | x_i) = \frac{p_{\theta_{t-1}}(y_i = j, x_i)}{p_{\theta_{t-1}}(x_i)} = \frac{\pi_j p_{\mu_j, \Sigma_j}(x_i)}{\sum_{l=1}^k \pi_l p_{\mu_l, \Sigma_l}(x_i)}$$

(We omit the subscript t on the righthmost expression.)

- **M-step:** set $\theta_t = \arg \max_{\theta \in \Theta} L(\theta; R_t)$

$$\pi_j = \frac{\sum_{i=1}^n R_{ij}}{\sum_{i=1}^n \sum_{l=1}^k R_{il}} = \frac{\sum_{i=1}^n R_{ij}}{n} \quad (1)$$

$$\mu_j = \frac{\sum_{i=1}^n R_{ij} x_i}{\sum_{i=1}^n R_{ij}} = \frac{\sum_{i=1}^n R_{ij} x_i}{n \pi_j} \quad (2)$$

$$\Sigma_j = \frac{\sum_{i=1}^n R_{ij} (x_i - \mu_j)(x_i - \mu_j)^\top}{n \pi_j} \quad (3)$$

(We omit all the subscripts t above.)

By using first-order condition and also Lagrange duality, one can show that the choices in (1), (2) and (3) solve the problem of $\arg \max_{\theta \in \Theta} L(\theta; R_t)$. We will leave it as an exercise.

Theorem 0.2. Let $(R_0, \theta_0) \in \mathbb{R}^{n \times k} \times \Theta$ be initialized arbitrarily. Let (R_t, θ_t) be given by EM:

$$(R_t)_{ij} := p_{\theta_{t-1}}(y = j | x_i) \quad \theta_t := \arg \max_{\theta \in \Theta} L(\theta; R_t)$$

Then for all t ,

$$L(\theta_t; R_t) \leq \max_{R \in \mathbb{R}^{n \times k}} L(\theta_t; R) = L(\theta_t; R_{t+1}) = L(\theta_t) \leq L(\theta_{t+1}; R_{t+1}) \quad (4)$$

In particular, this implies $L(\theta_t) \leq L(\theta_{t+1})$.

Proof. Let us first prove the easy steps in (4) from left to right.

- First, $L(\theta_t; R_t) \leq \max_{R \in \mathbb{R}^{n \times k}} L(\theta_t; R)$ follows from maximization over R .
- $L(\theta_t; R_{t+1}) = L(\theta_t)$ follows from the definition of R_{t+1} and augmented likelihood.
- $L(\theta_t) \leq L(\theta_{t+1}; R_{t+1})$ follows from maximization over θ .

Now we just need to show $\max_{R \in \mathbb{R}^{n \times k}} L(\theta_t; R) = L(\theta_t; R_{t+1})$. We will rely on a useful tool from convex analysis called the *Jensen's inequality*: for any concave function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, any a_1, \dots, a_k , and any weights $\lambda_1, \dots, \lambda_k \geq 0$ such that $\sum_{j=1}^k \lambda_j = 1$, the following inequality holds

$$\sum_{j=1}^m \lambda_j f(a_j) \leq f\left(\sum_{j=1}^m \lambda_j a_j\right)$$

Using this tool, we can bound the augmented likelihood as follows

$$\begin{aligned} L(\theta_t, R) &= \sum_{i=1}^n \sum_{j=1}^k R_{ij} \ln \frac{p_{\theta_t}(x_i, y_i = j)}{R_{ij}} \\ &\leq \sum_{i=1}^n \ln \left(\sum_{j=1}^k R_{ij} \frac{p_{\theta_t}(x_i, y_i = j)}{R_{ij}} \right) && \text{(Jensen's inequality)} \\ &\leq \sum_{i=1}^n \ln \left(\sum_{j=1}^k p_{\theta_t}(x_i, y_i = j) \right) \\ &\leq \sum_{i=1}^n \ln (p_{\theta_t}(x_i)) = L(\theta_t) = L(\theta_t, R_{t+1}) \end{aligned}$$

This means $L(\theta_t, R) \leq L(\theta_t, R_{t+1})$ for any R , and so $\max_R L(\theta_t, R) = L(\theta_t, R_{t+1})$. \square

Choosing the number k . The number k is another hyperparameter. We can follow the same approach in supervised learning, and tune it with a validation set. As we increase k , we will gradually increase the log-likelihood on the training set, but the log-likelihood on the validation set will stop increasing at some point.

k -Means Clustering A related unsupervised learning method is k -means clustering. We won't go into details in this course. k -means is another alternating optimization method that aims to minimize the following k -means objective

$$\phi(\mu_1, \dots, \mu_k) = \sum_{i=1}^n \min_j \|x_i - \mu_j\|^2$$

The method introduces an "hard" assignment matrix $A \in \{0, 1\}^{n \times k}$, and alternatively optimizes A and μ_j 's:

- For each x_i , define $\mu(x_i)$ to be a closest center:

$$\|x_i - \mu(x_i)\| = \min_j \|x_i - \mu_j\|$$

- For each i , set $A_{ij} = \mathbf{1}[\mu(x_i) = \mu_j]$.

k -means can also provide initialization for the EM method.