

Hidden Variables, the EM Algorithm, and Mixtures of Gaussians

Computer Vision

CS 143, Brown

James Hays

Today's Class

- Examples of Missing Data Problems
 - Detecting outliers
- Background
 - Maximum Likelihood Estimation
 - Probabilistic Inference
- Dealing with “Hidden” Variables
 - EM algorithm, Mixture of Gaussians
 - Hard EM

Missing Data Problems: Outliers

You want to train an algorithm to predict whether a photograph is attractive. You collect annotations from Mechanical Turk. Some annotators try to give accurate ratings, but others answer randomly.

Challenge: Determine which people to trust and the average rating by accurate annotators.



Annotator
Ratings

10
8
9
2
8

Missing Data Problems: Object Discovery

You have a collection of images and have extracted regions from them. Each is represented by a histogram of “visual words”.


Challenge: Discover frequently occurring object categories, without pre-trained appearance models.



Missing Data Problems: Segmentation

You are given an image and want to assign foreground/background pixels.

Challenge: Segment the image into figure and ground without knowing what the foreground looks like in advance.




Missing Data Problems: Segmentation

Challenge: Segment the image into figure and ground without knowing what the foreground looks like in advance.


Three steps:

1. If we had labels, how could we model the appearance of foreground and background?
2. Once we have modeled the fg/bg appearance, how do we compute the likelihood that a pixel is foreground?
3. How can we get both labels and appearance models at once?



Maximum Likelihood Estimation

1. If we had labels, how could we model the appearance of foreground and background?



Maximum Likelihood Estimation

data \rightarrow $\mathbf{x} = \{x_1 \dots x_N\}$ parameters

$$\hat{\theta} = \operatorname{argmax}_{\theta} p(\mathbf{x} | \theta)$$
$$\hat{\theta} = \operatorname{argmax}_{\theta} \prod_n p(x_n | \theta)$$

Maximum Likelihood Estimation

$$\mathbf{x} = \{x_1 \dots x_N\}$$

$$\hat{\theta} = \operatorname{argmax}_{\theta} p(\mathbf{x} | \theta)$$

$$\hat{\theta} = \operatorname{argmax}_{\theta} \prod_n p(x_n | \theta)$$

Gaussian Distribution

$$p(x_n | \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x_n - \mu)^2}{2\sigma^2}\right)$$

Maximum Likelihood Estimation

$$\mathbf{x} = \{x_1 \dots x_N\}$$

$$\hat{\theta} = \operatorname{argmax}_{\theta} p(\mathbf{x} | \theta)$$

$$\hat{\theta} = \operatorname{argmax}_{\theta} \prod_n p(x_n | \theta)$$

Gaussian Distribution

$$p(x_n | \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x_n - \mu)^2}{2\sigma^2}\right)$$


$$\hat{\mu} = \frac{1}{N} \sum_n x_n \quad \hat{\sigma}^2 = \frac{1}{N} \sum_n (x_n - \hat{\mu})^2$$

Example: MLE

Parameters used to Generate

fg: $\mu=0.6$, $\sigma=0.1$

bg: $\mu=0.4$, $\sigma=0.1$



im



labels

```
>> mu_fg = mean(im(labels))  
      mu_fg = 0.6012
```

```
>> sigma_fg = sqrt(mean((im(labels)-mu_fg).^2))  
      sigma_fg = 0.1007
```


```
>> mu_bg = mean(im(~labels))  
      mu_bg = 0.4007
```

```
>> sigma_bg = sqrt(mean((im(~labels)-mu_bg).^2))  
      sigma_bg = 0.1007
```

```
>> pfg = mean(labels(:));
```

Probabilistic Inference

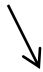
2. Once we have modeled the fg/bg appearance, how do we compute the likelihood that a pixel is foreground?



Probabilistic Inference

Compute the likelihood that a particular model generated a sample

component or label


$$p(z_n = m \mid x_n, \theta)$$

Probabilistic Inference

Compute the likelihood that a particular model generated a sample

component or label

$$p(z_n = m | x_n, \theta) = \frac{p(z_n = m, x_n | \theta_m)}{p(x_n | \theta)}$$

Probabilistic Inference

Compute the likelihood that a particular model generated a sample

component or label

$$p(z_n = m | x_n, \theta) = \frac{p(z_n = m, x_n | \theta_m)}{p(z_n | \theta)}$$
$$= \frac{p(z_n = m, x_n | \theta_m)}{\sum_k p(z_n = k, x_n | \theta_k)}$$

Probabilistic Inference

Compute the likelihood that a particular model generated a sample


component or label

$$p(z_n = m | x_n, \theta) = \frac{p(z_n = m, x_n | \theta_m)}{p(x_n | \theta)}$$

$$= \frac{p(z_n = m, x_n | \theta_m)}{\sum_k p(z_n = k, x_n | \theta_k)}$$

$$= \frac{p(x_n | z_n = m, \theta_m) p(z_n = m | \theta_m)}{\sum_k p(x_n | z_n = k, \theta_k) p(z_n = k | \theta_k)}$$

Example: Inference




im

Learned Parameters

fg: $\mu=0.6$, $\sigma=0.1$

bg: $\mu=0.4$, $\sigma=0.1$

```
>> pfg = 0.5;  
>> px_fg = normpdf(im, mu_fg, sigma_fg);  
>> px_bg = normpdf(im, mu_bg, sigma_bg);  
>> pfg_x = px_fg*pfg ./ (px_fg*pfg + px_bg*(1-pfg));
```







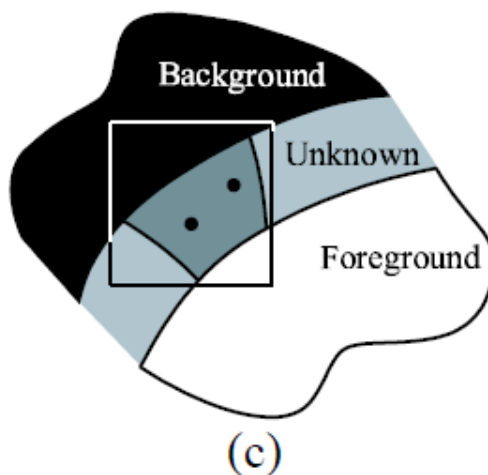

$p(\text{fg} \mid \text{im})$

Mixture of Gaussian* Example: Matting


Knockout

Ruzon-Tomasi

Bayesian




Mixture of Gaussian* Example: Matting



Dealing with Hidden Variables


3. How can we get both labels and appearance models at once?



Segmentation with Mixture of Gaussians

Pixels come from one of several Gaussian components

- We don't know which pixels come from which components
- We don't know the parameters for the components



Simple solution

1. Initialize parameters
2. Compute the probability of each hidden variable given the current parameters
3. Compute new parameters for each model, weighted by likelihood of hidden variables
4. Repeat 2-3 until convergence

Mixture of Gaussians: Simple Solution

1. Initialize parameters

2. Compute likelihood of hidden variables for current parameters

$$\alpha_{nm} = p(z_n = m | x_n, \boldsymbol{\mu}^{(t)}, \boldsymbol{\sigma}^{2(t)}, \boldsymbol{\pi}^{(t)})$$

3. Estimate new parameters for each model, weighted by likelihood

$$\hat{\boldsymbol{\mu}}_m^{(t+1)} = \frac{1}{\sum_n \alpha_{nm}} \sum_n \alpha_{nm} x_n \quad \hat{\boldsymbol{\sigma}}_m^{2(t+1)} = \frac{1}{\sum_n \alpha_{nm}} \sum_n \alpha_{nm} (x_n - \hat{\boldsymbol{\mu}}_m)^2 \quad \hat{\boldsymbol{\pi}}_m^{(t+1)} = \frac{\sum_n \alpha_{nm}}{N}$$

Expectation Maximization (EM) Algorithm

$$\text{Goal: } \hat{\theta} = \operatorname{argmax}_{\theta} \log \left(\sum_{\mathbf{z}} p(\mathbf{x}, \mathbf{z} | \theta) \right)$$


1. E-step: compute

$$\mathbb{E}_{\mathbf{z} | \mathbf{x}, \theta^{(t)}} \left[\log p(\mathbf{x}, \mathbf{z} | \theta) \right] = \sum_{\mathbf{z}} \log p(\mathbf{x}, \mathbf{z} | \theta) p(\mathbf{z} | \mathbf{x}, \theta^{(t)})$$

2. M-step: solve


$$\theta^{(t+1)} = \operatorname{argmax}_{\theta} \sum_{\mathbf{z}} \log p(\mathbf{x}, \mathbf{z} | \theta) p(\mathbf{z} | \mathbf{x}, \theta^{(t)})$$

Gaussian Mixture Example: Start




Advance apologies: in Black and White this example will be incomprehensible


After first iteration




After 2nd
iteration




After 3rd
iteration




After 4th
iteration




After 5th
iteration




After 6th
iteration




After 20th
iteration



Some Bio Assay data



GMM clustering of the assay data



Resulting Density Estimator



Mixture of Gaussian demos

- <http://www.cs.cmu.edu/~alad/em/>
- <http://lcn.epfl.ch/tutorial/english/gaussian/html/>

“Hard EM”

- Same as EM except compute z^* as most likely values for hidden variables
- K-means is an example
- Advantages
 - Simpler: can be applied when cannot derive EM
 - Sometimes works better if you want to make hard predictions at the end
- But
 - Generally, pdf parameters are not as accurate as EM

Missing Data Problems: Outliers

You want to train an algorithm to predict whether a photograph is attractive. You collect annotations from Mechanical Turk. Some annotators try to give accurate ratings, but others answer randomly.

Challenge: Determine which people to trust and the average rating by accurate annotators.



Annotator
Ratings

10
8
9
2
8


Missing Data Problems: Object Discovery

You have a collection of images and have extracted regions from them. Each is represented by a histogram of “visual words”.

Challenge: Discover frequently occurring object categories, without pre-trained appearance models.



What's wrong with this prediction?



$P(\text{foreground} \mid \text{image})$

Next class

- MRFs and Graph-cut Segmentation

