


Last 2 classes

- Beta-Bernoulli model (for learning 'p' from $\text{Ber}(p)$ data)
- Decision theory - MAP estimator, posterior prediction, credible intervals
- Bayes Nets - Graphical 'language' for complex probabilistic models (DAG which represents independence rel's in your model)
- Conditional independence \equiv d-separation

Today

- Dirichlet priors
- Naive Bayes classifier
- Gaussian models

$$\left(\begin{array}{l} \text{are } X_A \perp\!\!\!\perp X_B \mid X_C \\ X_A = \{X_i \mid i \in A\} \end{array} \right)$$



(Bishop, Chapter 8, sec 1.2(!))

multiclass data (clustering / classification)

- data $D = \{X_1, X_2, \dots, X_n\} \in \{1, 2, \dots, K\}^n$
← count of word i in data ← number of words
eg - set of english words
- for $i \in [K]$, data D contains N_i copies of type i
- model \mathcal{M} : X_i generated i.i.d. from $Multinomial(\theta_1, \theta_2, \dots, \theta_K)$ distn

Eg - email filtering, classes = {spam, not spam}
sentiment analysis datasets \equiv 'bag of words' repⁿ of an email



← Each word is of type k w.p. $\frac{\theta_k}{\sum \theta_i}$ iid

$Mult(n, \theta_1, \theta_2, \dots, \theta_K)$ NAIVE

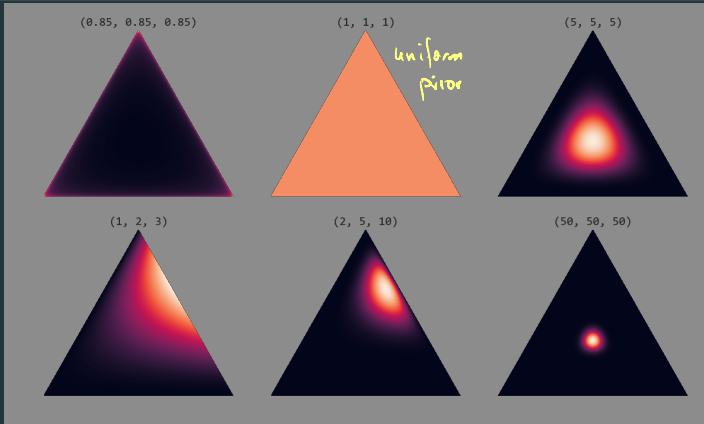
Eg - If $K=2$, then $Bin(n, p)$
ie, $\mathcal{A} = \{0, 1\}$ = $Mult(n, p, 1-p)$

$$\mathcal{L}(N_1, N_2, \dots, N_K | \theta) = \prod_{i=1}^K \theta_i^{N_i} \quad \text{Multinomial Likelihood}$$

the Dirichlet distribution

Dirichlet distribution

- $x \in \{x_i \in [0, 1], \sum_{i=1}^K x_i = 1\}$, parameters: $\Theta = (\alpha_1, \alpha_2, \dots, \alpha_K)$
 - pdf: $p(x) \propto \prod_{i=1}^K x_i^{\alpha_i - 1}$
- K-simplex*
hyperparameters $\alpha_i \geq 0$



the Dirichlet distribution

Dirichlet distribution

- $x \in \{x_i \in [0, 1], \sum_{i=1}^K x_i = 1\}$, parameters: $\Theta = (\alpha_1, \alpha_2, \dots, \alpha_K)$
- denote $\alpha = \{\alpha_i\}_{i=1}^K$; Dirichlet pdf

$$p(x) = \frac{1}{Z(\alpha)} \prod_{i=1}^K x_i^{\alpha_i - 1}$$

$\Gamma(\alpha_1 + \alpha_2 + \dots + \alpha_K) / \Gamma(\alpha_1) \Gamma(\alpha_2) \dots \Gamma(\alpha_K)$

- normalizing constant: $\frac{1}{Z(\alpha)} = \frac{\Gamma(\sum_{i=1}^K \alpha_i)}{\prod_{i=1}^K \Gamma(\alpha_i)}$
(Dirichlet norm const)

• Recall - $\Gamma(\alpha + 1) = \alpha \Gamma(\alpha)$, $\Gamma(1) = 1$
 $\Gamma(k) = (k-1)!$ if $k \in \mathbb{N}$

multiclass data and Dirichlet priors

- for $i \in [K]$, data D contains N_i copies of type i
- model \mathcal{M} : X_i generated i.i.d. from $Multinomial(\theta_1, \theta_2, \dots, \theta_K)$ distn

Dirichlet-Multinomial model

- prior parameters: $\Theta_0 = (\alpha_1, \alpha_2, \dots, \alpha_K) \in \mathbb{R}_+^K$ (hyperparameters)
- Dirichlet prior: $Dir(\alpha_1, \alpha_2, \dots, \alpha_K) \sim p(\theta) \propto \prod_{i=1}^K \theta_i^{\alpha_i - 1}$
- likelihood: $p(D|\theta) = \prod_{i=1}^K \theta_i^{N_i}$
- posterior: $p(\theta|D) \sim Dir(\alpha_1 + N_1, \alpha_2 + N_2, \dots, \alpha_K + N_K)$
- marginal likelihood: let $m = \sum_{i=1}^K \alpha_i$ posterior = $\frac{\text{prior} \times \text{likelihood}}{\text{marginal likelihood}}$


$$p(D) = \frac{\Gamma(m)}{\Gamma(n+m)} \prod_{i=1}^K \frac{\Gamma(N_i + \alpha_i)}{\Gamma(\alpha_i)}$$

Spam filtering

- Consider only spam emails: D_1, D_2, \dots, D_n
- $D_i \equiv (N_1^i, N_2^i, \dots, N_k^i) \equiv$ word counts



$D_1 = (0, 4, 9)$
→
(learning)



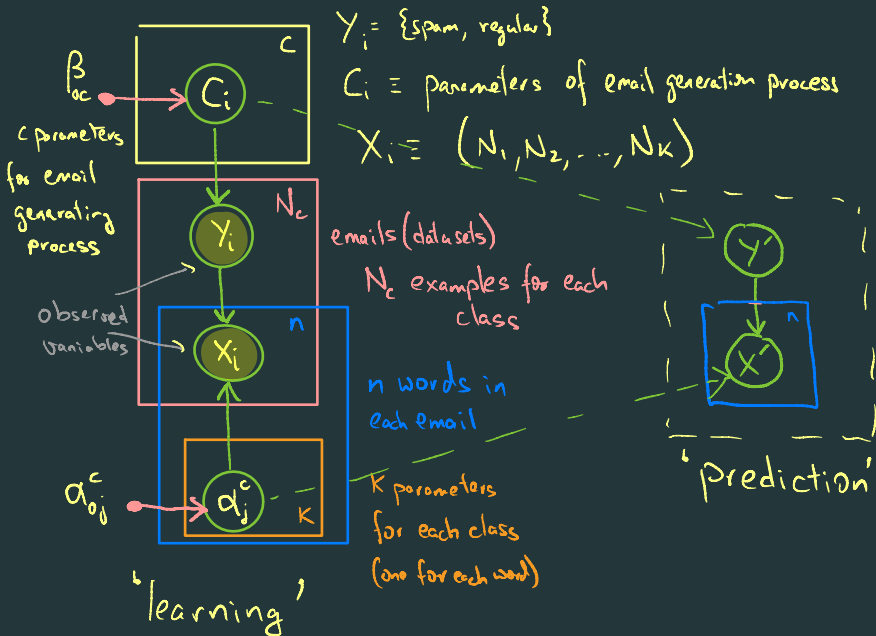
- Now if we get a new email D' , what is $P[D' | \text{spam}]$

$$\begin{aligned} &\propto \prod_{i=1}^k (\alpha_i^s)^{N_i'} \\ &= \frac{1}{Z(1,5,10)} (1)^{N_1'} (5)^{N_2'} (10)^{N_3'} \end{aligned}$$

Naive Bayes classifier

- For spam - take data $D_1^s, D_2^s, \dots, D_{n_s}^s$
 - learn posterior $\text{Dir}(\alpha_1^s, \alpha_2^s, \dots, \alpha_k^s)$
- For non-spam (regular) - take data $D_1^r, D_2^r, \dots, D_{n_r}^r$
 - learn posterior $\text{Dir}(\alpha_1^r, \alpha_2^r, \dots, \alpha_k^r)$
- Given new email D' is it regular or spam?
 - $= (N_1, N_2, \dots, N_k)$
 - compute $P(D' | \text{spam}) = \frac{\prod_{i=1}^k (\alpha_i^s)^{N_i}}{Z(\{\alpha_i^s\})}$, $P(D' | \text{regular}) = \frac{\prod_{i=1}^k (\alpha_i^r)^{N_i}}{Z(\{\alpha_i^r\})}$
 - $P(\text{spam} | D') \propto P(D' | \text{spam}) \underbrace{p(\text{spam})}_{?}$, $P(\text{regular} | D') \propto P(D' | \text{reg}) \underbrace{p(\text{reg})}_{?}$
need model for class label ?

Naive Bayes model for spam classification



generative models for continuous data