# Notation 

Kevin P. Murphy

## Last updated October 25, 2006

## 1 General

- $x^{*} \in \arg \max _{x} f(x)$ means $x^{*}$ is the value of $x$ that maximizes the function $f$, i.e., $f\left(x^{*}\right)=\max _{x} f(x)$. Note that there may be multiple global maxima, in which case we break ties randomly.
- Indicator function: $I(e)=1$ if event $e$ is true, $I(e)=0$ otherwise.
- Delta function: $\delta(x)=1$ if $e=0$ and $\delta(x)=0$ otherwise.
- Sometimes probability mass functions (for discrete random variables) are written $P(X)$, and probability density functions (for continuous random variables) are written $p(X)$. We will use $p()$ for both.
- Usually we write random variables as capital letters and values of random variables as lower case, e.g., $p(X=x)$ is the probability $X$ has value $x$. However, we do not follow this convention very closely.
- If $X$ is distributed according to distribution $f$ with parameters $\theta$, i.e., $p(X)=f(X \mid \theta)$, then we write $X \sim f(\theta)$.
- We will often write probability distributions up to a constant of proportionality, $p(x \mid \theta) \propto f(x, \theta)$. This normalization constant is often denoted $1 / Z(\theta)$, where $Z$ is called the partition function.
- Vectors are usually column vectors. $T$ denotes transpose, so $x^{T}$ is a row vector. Sometimes we will write vectors in bold, e.g., $\mathbf{x}$, or as $\vec{x}$, but usually we will just write $x$. Matrices will usually be written as capital letters, $X$. However, using this convention we will cannot distinguish matrices from scalar (or vector) random variables. It should be clear from context.
- We use the following matlab notation: $1: n$ denotes the sequence of integers $\{1,2, \ldots, n\}$ and $X(i, j, k)$ is element $i, j, k$ of some matrix, where $i, j, k$ could each be a sequence of indices.


## 2 Data and variables

- $X_{n i}$ is the $i$ 'th component/ feature / variable of data case $n$, for $i=1: D$, where $D$ is the number of features/ variables, and $n=1: N$, where $N$ is the number of training samples. (In general, $D$ may depend on $n$ if each example has a different size, but we will rarely write $D_{n}$.) If there is a single training/ test sample, we just write $X_{i}$ for the $i$ 'th variable.
- $X=X(1: N, 1: D)$ is the design matrix. The $n$ 'th row is the $n$ 'th example $X_{n}^{T}$ (since each example $X_{n}$ is a column vector); the $i$ 'th column of $X$ is all values of the $i$ 'th feature. We also write this as $\mathcal{D}=\left\{x_{n}\right\}_{n=1}^{N}$, which is more general notation, since it does not assume that all examples have the same number of features (e.g., a document may contain sentences of different lengths, so we would use $\mathcal{D}$ rather than $X$ ).
- In supervised learning problems, there is a distinguished output variable $y_{n}$, so $\mathcal{D}=\left\{x_{n}, y_{n}\right\}$. In classification, $y_{n} \in\{1, \ldots, C\}$, where $C$ is the number of classes. In regression, $y_{n} \in \mathbb{R}$.
- If $X_{n i}$ is a scalar, then $X_{n i} \in \mathbb{R}$ or $X_{n i} \in \mathbb{R}^{+}$. If it is a vector, then $X_{n i} \in \mathbb{R}^{K}$, where $X_{n i k}$ is the $k$ 'th component of the $i$ 'th variable for $k=1: K$, where $K$ is the dimensionality of each variable. (In general, $K$ may depend on $i$, but we rarely write $K_{i}$ ).
- If $X_{n i}$ is binary, then $X_{n i} \in\{0,1\}$. If $X_{n i}$ is categorical, then $X_{n i} \in\{1, \ldots, K\}$, where $K$ is the number of states of variable $i$. (In general, $K$ may depend on $i$, but we rarely write $K_{i}$ ). We write $X_{n i}=k$ if the $i$ 'th variable is in state $k$, where $k \in 1: K$. Sometimes you will see a 1 -of- $K$ encoding, where $X_{n i} \in\{0,1\}^{K}$, where $X_{n i k}=I\left(X_{n i}=k\right)$. We also use $j$ to index states, mostly of variables that are "parents" of $X_{i}$.


## 3 Bernoullis and multinomials

- We define the Bernoulli distribution $X \sim B e(\theta)$ for $X \in\{0,1\}$ by

$$
\begin{equation*}
B e(X \mid \theta)=\theta^{X}(1-\theta)^{1-X} \tag{1}
\end{equation*}
$$

We denote the minimal sufficient statistics for a Bernoulli distribution by the number of heads and tails: $N_{1}=$ $\sum_{n} I\left(X_{n}=1\right), N_{0}=\sum_{n} I\left(X_{n}=0\right)$. Alternatively, we can use $N_{1}$ and $N=N_{1}+N_{0}$.

- We define the multinomial distribution $X \sim M u(\theta)$ for $X \in\{1, \ldots, K\}$ by

$$
\begin{equation*}
M u(X \mid \theta)=\prod_{j=1}^{K} \theta_{j}^{I(X=j)} \tag{2}
\end{equation*}
$$

Put another way, $p(X=j \mid \theta)=\theta_{j}$. We denote the sufficient statistics for a multinomial distribution by $N_{j}=\sum_{n} I\left(X_{n}=j\right)$.

- We define the Beta distribution $\theta \sim \operatorname{Beta}\left(\alpha_{1}, \alpha_{0}\right)$ for $\theta \in[0,1]$ by

$$
\begin{equation*}
\operatorname{Beta}\left(\theta \mid \alpha_{1}, \alpha_{0}\right)=\frac{\Gamma\left(\alpha_{1}+\alpha_{0}\right)}{\Gamma\left(\alpha_{1}\right) \Gamma\left(\alpha_{0}\right)} \theta^{\alpha_{1}-1}(1-\theta)^{\alpha_{0}-1} \tag{3}
\end{equation*}
$$

where $\Gamma(x)$ is the gamma function. Here $\alpha_{0}, \alpha_{1} \in \mathbb{R}^{+}$are called hyper parameters (pseudo counts) and $\alpha=\alpha_{0}+\alpha_{1}$ is the equivalent sample size (strenght) of the prior.

- We define the Dirichlet distribution $\theta \sim \operatorname{Dir}\left(\alpha_{1}, \ldots, \alpha_{K}\right)$ for $\theta \in[0,1]^{K}$ by

$$
\begin{equation*}
\operatorname{Dir}\left(\theta \mid \alpha_{1}, \ldots, \alpha_{K}\right) \propto \prod_{j=1}^{K} \theta_{j}^{\alpha_{j}-1} \tag{4}
\end{equation*}
$$

Here $\alpha_{j} \in \mathbb{R}^{+}$are called hyper parameters (pseudo counts) and $\alpha=\sum_{j} \alpha_{j}$ is the equivalent sample size.

- We define the likelihood as

$$
\begin{equation*}
L(\theta)=p(\mathcal{D} \mid \theta) \tag{5}
\end{equation*}
$$

and the log-likelihood as

$$
\begin{equation*}
\ell(\theta)=\log p(\mathcal{D} \mid \theta) \tag{6}
\end{equation*}
$$

- We denote the maximum likelihood estimate by

$$
\begin{equation*}
\hat{\theta}^{M L} \in \arg \max _{\theta} p(\mathcal{D} \mid \theta) \tag{7}
\end{equation*}
$$

We denote the maximum a posterior estimate by

$$
\begin{equation*}
\hat{\theta}^{M A P} \in \arg \max _{\theta} p(\mathcal{D} \mid \theta) p(\theta) \tag{8}
\end{equation*}
$$

We denote the posterior mean estimate by

$$
\begin{equation*}
\hat{\theta}^{\text {mean }}=E[\theta \mid \mathcal{D}] \tag{9}
\end{equation*}
$$

## 4 Naive Bayes classifier

- The 1 d Gaussian density is denoted $\mathcal{N}(x \mid \mu, \sigma)$.
- In a generative classifer, the class prior is usually denoted $p(Y=c)=\pi_{c}$, if we assume $Y$ has a multinomial distribution.
- In the naive Bayes model, we have

$$
\begin{equation*}
p(x \mid y=c)=\prod_{i=1}^{D} p\left(x_{i} \mid y=c\right) \tag{10}
\end{equation*}
$$

In the case of $K$-ary features, we have

$$
\begin{equation*}
p\left(x_{i} \mid y=c\right)=\prod_{k} \theta_{i c k}^{I\left(X_{i}=k\right)} \tag{11}
\end{equation*}
$$

where $\theta_{i c k}=P\left(X_{i}=k \mid Y=c\right)$. The sufficient statistics are $N_{i c k}$, which is the number of times $X_{i}=k$ amongst those training cases where $Y=c$. In the case of binary features, we have

$$
\begin{equation*}
p\left(x_{i} \mid y=c\right)=\theta_{i c}^{I\left(X_{i}=1\right)}\left(1-\theta_{i c}\right)^{I\left(X_{i}=0\right)} \tag{12}
\end{equation*}
$$

where $\theta_{i c}=P\left(X_{i}=1 \mid Y=c\right)$. The sufficient statistics are $N_{i c 1}$, the number of times $X_{i}=1$ amongst cases where $Y=c$, and $N_{i c}=N_{c}$, the number of times $X_{i}=0$ or $X_{i}=1$ in cases where $Y=c$.

## 5 Markov chains

- The transition matrix is $T_{j k}^{i}=p\left(X_{i}=k \mid X_{i-1}=j\right)$, which is independent of $i$ if the chain is stationary. The sufficient statistics to estimate this are the observed number of $j \rightarrow k$ transitions: $N_{j k}=\sum_{n=1}^{N} \sum_{i=2}^{D} I\left(X_{n i}=\right.$ $k, X_{n i-1}=j$ ). There is no $i$ index since we assume the parameters are shared (tied) across time.
- The initial state distribution is $\pi_{k}^{1}=p\left(X_{1}=k\right)$.
- The stationary distribution is $\pi$ which satisfies $\pi T=\pi$ (if we treat $\pi$ as a row vector).


## 6 Information theory

- The entropy of a random variable $X \in 1: K$ with discrete distribution $p$ is denoted by

$$
\begin{equation*}
H(p)=H(X)=-\sum_{k=1}^{K} p(X=k) \log _{2} p(X=k)=-\sum_{k} p_{k} \log p_{k} \tag{13}
\end{equation*}
$$

The joint entropy is denoted $H(X, Y)$ and the conditional entropy as $H(X \mid Y)$. The mutual information is denoted $I(X, Y)$ (often written as $I(X ; Y)$ ). The Kullback-Leibler divergence between two distributions is denoted $K L(p \| q)$.

