

Classification

Maryclare Griffin

2024-10-08

This material is based on Chapters 2 and 4 of Introduction to Statistical Learning (ISL) and parts of Chapter 4 of Elements of Statistical Learning (ESL). We will tend to follow ISL more closely, and look to ESL for occasional additional higher level material. There may be some slight changes to the notation.

For the past several lectures, we have been focused on the supervised learning problem where we observe inputs \mathbf{X} and a quantitative output \mathbf{y} and assuming

$$Y = f(X) + \epsilon,$$

where f is a fixed but unknown function of $X = X_1, \dots, X_p$ and ϵ is a mean zero random error term which is independent of X .

In practice, we often encounter a **qualitative** or **categorical** output \mathbf{y} , with elements that each take on one of M distinct values. Sometimes these distinct values are also called **classes**, and the problem of predicting a future output value or class is referred to as **classification**.

If we want to predict future values of the output, it no longer makes sense to think about minimizing

$$E[(Y - \hat{f}(X))^2]$$

when Y is qualitative. Rather, our goal will be to obtain predictions \hat{Y} based on estimated functions of the inputs $\hat{f}(X)$ that correctly identify the value the observed output is most likely to take on, i.e. we want to obtain predictions \hat{Y} from $\hat{f}(X)$ that satisfy $\hat{Y} = Y$ as often as possible. We refer to this as predictions that minimize the **error rate** and we refer to the estimated function $\hat{f}(X)$ mapping from $\hat{f}(X)$ to \hat{Y} as a **classifier**.

Let $I(y_i \neq \hat{y}_i)$ refer to an indicator function that is equal to 0 if $y_i = \hat{y}_i$, i.e. if the prediction for the i -th training data point matches the observed value of the output, and 1 otherwise. We can quantify the performance of a classifier on the training data via the **training error rate**,

$$\sum_{i=1}^n I(y_i \neq \hat{y}_i).$$

Analogously to the setting where the output is quantitative, we can also quantify the performance of a classifier on test data that were not used to fit the classifier. Letting (x_0, y_0) refer to an arbitrary data point that has not been seen before, we define the **test error rate** as

$$I(y_0 \neq \hat{y}_0). \tag{1}$$

Again, we can imagine averaging this over all possible observations that we haven't seen before but may in the future. The best classifier will minimize the average test error rate over test data that were not used to fit the classifier.

So far the definitions of the error rate, training error rate, and test error rate are a bit awkward, because they do not explicitly depend on the $\hat{f}(X)$, rather they depend on $\hat{f}(X)$ implicitly through \hat{Y} . Relatedly, thinking of the output as the sum of an unknown function of the inputs and a mean zero error is no longer meaningful - what is $E[Y|X] = f(X)$ if Y is qualitative? It doesn't make sense!

Accordingly, we will make the relationship between $\hat{f}(X)$ and \hat{Y} more explicit. It can be proven that (1) is minimized on average by a classifier that predicts the class that is most likely given the inputs, i.e. if $v_j = 1, \dots, v_M$ represents all of the possible values that Y can take on and if the probability that Y takes on value v_j given inputs X is known $\Pr(Y = v_j|X)$, this classifier assigns predictions \hat{y}_0 according to

$$\hat{y}_0 = \operatorname{argmax}_j \Pr(Y = v_j|x_0),$$

Note that when $M = 2$, this corresponds to predicting whichever class has probability $\Pr(Y = v_j|x_0) > 0.5$.

This suggests defining

$$f_j(X) = \Pr(Y = v_j|X).$$

We now have a subscript j associated with $f(X)$ that reflects the fact that we need to define these probabilities for every possible value of Y .

Note that because $j = 1, \dots, M$ indexes all of the possible values that the output Y could take on, $\sum_{j=1}^M \Pr(Y = v_j) = \sum_{j=1}^M f_j(X) = 1$. This means that given $M - 1$ probabilities, we can always reconstruct the remaining probability. For this reason, we will sometimes model for qualitative output only specify $\Pr(Y = v_j) = f_j(X)$ and estimate $f_j(X)$ for $M - 1$ values of j , most frequently for $j = 2, \dots, M$ or $j = 1, \dots, M - 1$.

Just as in the regression setting where we could decompose the performance of an estimated classifier into reducible and irreducible, we can do something similar in the classification setting. The equivalent to irreducible error in the classification setting is the **Bayes error rate**,

$$1 - E[\max_j f_j(X)], \tag{2}$$

where the expectation is taken with respect to X . This describes the error rate even if the true probabilities $f_j(X)$ were known. The Bayes error rate is equal to 0 when the **Bayes decision boundary**, which describes the set of values of X for which all possible values of the outcome are equally likely perfectly separates the outputs. This corresponds to the setting where $\max_j \Pr(Y = v_j|X) = 1$ for all X and the outputs are deterministic functions of the inputs. In real life (and furthermore in this *statistical* learning class in which we are studying random outputs), the outputs are rarely deterministic functions of the inputs and the Bayes error rate is rarely 0.

Now we will introduce our first classifier which is very closely related to based on K -nearest neighbors regression, **K-nearest neighbors classification**. Like its regression counterpart, KNN classification is one of the simplest non-parametric methods for classification.

Let $\mathcal{N}_i^{(K)}$ refer to the set of K indices of observed values of the predictor that are closest to x_i . Then the KNN estimate of $f_j(x_i)$ is

$$\hat{f}_j(x_i) = \frac{1}{K} \sum_{k \in \mathcal{N}_i^{(K)}} I(y_k = v_j)$$

and the KNN classifier assigns the prediction $\hat{y}_i = v_j$ where v_j is the value of the output associated with the highest estimated probability $\hat{f}_j(x_i)$

Again, the value of K is chosen by the user and determines the bias and variance of the estimate of f . Smaller K correspond to less biased but more variable estimates, whereas larger K correspond to (potentially) more biased and less variable estimates.

Just as in the regression setting, the performance of KNN classifiers depends on how K is chosen and if K is chosen to be too small, KNN classifiers can overfit the data. Furthermore, the performance of KNN classifiers also deteriorates rapidly as p , the number of predictors and/or dimension of X , increases. This is because it is harder to find neighbors in high dimensions. This is related to the idea of the **curse of dimensionality**.

This leads us to parametric classifiers. In what follows we will mainly consider the **binary** setting When Y is **binary** with $M = 2$ levels, we will refer to the levels of Y as equal to $v_1 = 0$ or $v_2 = 1$, specify

$\Pr(Y = 1) = f_2(X)$, and drop the subscript on $f_2(X)$, letting $f(X) = f_2(X)$. We can recover $f_1(X)$ from $f(X)$ according to $f_1(X) = 1 - f(X)$. Note that the KNN classifier's estimate of $f(X)$ in this case is identical to the estimate of $f(X)$ obtained from KNN regression.

Similarly, the simplest parametric classifier is obtained by using linear regression to estimate $f(X)$. We assume $f(X) = \beta_0 + \beta_1 X_1 + \dots + \beta_p X_p$ and estimate $\beta_0, \beta_1, \dots, \beta_p$ by finding the values $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$ that minimize the least squares criterion

$$\sum_{i=1}^n (y_i - b_0 - x_{i1}b_1 - \dots - x_{ip}b_p)^2,$$

with respect to b_0, b_1, \dots, b_p . Then predictions \hat{Y} can be obtained by setting $\hat{Y} = I(\hat{\beta}_0 + \hat{\beta}_1 X_1 + \dots + \hat{\beta}_p X_p > 0.5)$. As in the regression case, there is only a unique minimizer of the least squares criterion when $p < n$, i.e. we have more observations than predictors, and when the predictors $\mathbf{x}_1, \dots, \mathbf{x}_p$ aren't too correlated with each other. We emphasize that this is only reasonable in the binary setting with $M = 2$ when the output is defined to take on values 0 and 1.

There are one obvious issue with using linear regression for classification in this way - it can produce nonsensical estimated probabilities that are less than 0 or greater than 1! This leads us to **logistic regression**, which is a special case of a **generalized linear model**. A generalized linear model is made up of:

- A choice of **link function** $g(\cdot)$ that relates the conditional mean of an output $E[Y|X]$ to a linear function of the inputs $\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p$, satisfying $E[Y|X] = g(\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p)$;
- An assumed conditional distribution of the output given values of the input.

The parameters of a generalized linear model, including $\beta_0, \beta_1, \dots, \beta_p$, can then be estimated using maximum likelihood.

Logistic regression uses the link function

$$\begin{aligned} g(\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p) &= \frac{\exp\{\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p\}}{1 + \exp\{\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p\}} \\ &= \frac{1}{1 + \exp\{-\beta_0 - \beta_1 X_1 - \dots - \beta_p X_p\}}, \end{aligned}$$

which maps the real valued linear function $\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p$ to a number between 0 and 1. This is sometimes called the **logistic** or **sigmoid** function. Logistic regression also assumes that the outputs Y are Bernoulli random variables conditional on the inputs X ,

$$Y|X \stackrel{\text{indep.}}{\sim} \text{Bernoulli} \left(\frac{1}{1 + \exp\{-\beta_0 - \beta_1 X_1 - \dots - \beta_p X_p\}} \right).$$

When we use a nonlinear model, we can still interpret nonzero β_j for $j > 0$ as indicating the presence of a relationship between the output and the input, but the interpretation of individual β_j for $j > 0$ changes. The interpretation will now describe how the log odds of $Y = 1$ change with changes in individual inputs X_j , where the log odds are

$$\log \left(\frac{\Pr(Y = 1|X)}{1 - \Pr(Y = 1|X)} \right).$$

Personally, I don't find log odds very interpretable beyond the fact that higher log odds correspond to $Y = 1$ being more likely and lower log odds correspond to $Y = 0$ being more likely. Nonetheless, the interpretation of any β_j for $j > 0$ will be the expected change in log odds of $Y = 1$ associated with a one unit change in X_j , holding all other predictors constant. Sometimes, people point to this awkward interpretation of the regression coefficients as a reason not to use logistic regression but that is beyond the scope of this class.

Estimates of the logistic regression coefficients $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$ can then be obtained by maximizing the log likelihood

$$\begin{aligned} \sum_{i=1}^n \log(p(y_i|x_i)) &= \sum_{i=1}^n y_i \log\left(\frac{1}{1 + \exp\{-b_0 - b_1 X_1 - \dots - b_p X_p\}}\right) + \\ &\quad (1 - y_i) \log\left(1 - \frac{1}{1 + \exp\{-b_0 - b_1 X_1 - \dots - b_p X_p\}}\right) \\ &= \sum_{i=1}^n y_i \log\left(\frac{1}{1 + \exp\{-b_0 - b_1 X_1 - \dots - b_p X_p\}}\right) + \\ &\quad (1 - y_i) \log\left(\frac{\exp\{-b_0 - b_1 X_1 - \dots - b_p X_p\}}{1 + \exp\{-b_0 - b_1 X_1 - \dots - b_p X_p\}}\right) \end{aligned}$$

with respect to b_0, b_1, \dots, b_p . This is a nonlinear function of b_0, b_1, \dots, b_p , and we will rely on statistical software to maximize it and find the maximum likelihood estimator $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$.

This may seem new, we have already been using maximum likelihood in this class without explicitly saying it! The least squares estimator is the maximum likelihood estimator obtained by assuming that the output is conditionally normal given the inputs with mean $\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p$ and constant variance!

A great think about many commonly used generalized linear models broadly and also for logistic regression is that it has been shown that the asymptotic distribution of the estimated regression coefficients $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$ as $n \rightarrow \infty$ can be derived to be normal and centered at the truth $\beta_0, \beta_1, \dots, \beta_p$ with a variance-covariance matrix that can be estimated from the data. This means that we can construct hypotheses tests of null hypotheses that any individual β_j is equal to a certain value or that subsets of coefficients are equal to a certain value and construct confidence intervals for the estimated regression coefficients.

Logistic regression is also straightforward to extend to qualitative outputs with $M > 2$ levels. This extension is called **multinomial logistic regression**. Let β_{kj} for $j = 0, 1, \dots, p$ refer to the regression coefficient associated with the probability that the output Y is equal to the k -th level v_k for the intercept (if $j = 0$) or predictor X_j (if $j > 0$). Multinomial logistic regression assumes that the probability that the probability that the output Y is equal to the k -th level v_k is

$$\frac{\exp\{\beta_{k0} + \beta_{k1} X_1 + \dots + \beta_{kp} X_p\}}{\sum_{l=1}^m \exp\{\beta_{l0} + \beta_{l1} X_1 + \dots + \beta_{lp} X_p\}}.$$

Going back to the notation we introduced when we first started talking about qualitative outputs, this is $f_k(X)$ or $\Pr(Y = v_k|X)$. Having assumed this link function, multinomial logistic regression then assumes that the outputs Y are multinomial random variables conditional on the inputs X with size 1,

$Y|X \stackrel{indep.}{\sim}$

$$\text{Multinomial}\left(1, \left(\frac{\exp\{\beta_{10} + \beta_{11} X_1 + \dots + \beta_{1p} X_p\}}{\sum_{l=1}^m \exp\{\beta_{l0} + \beta_{l1} X_1 + \dots + \beta_{lp} X_p\}}\right), \dots, \left(\frac{\exp\{\beta_{M0} + \beta_{M1} X_1 + \dots + \beta_{Mp} X_p\}}{\sum_{l=1}^m \exp\{\beta_{l0} + \beta_{l1} X_1 + \dots + \beta_{lp} X_p\}}\right)\right).$$

Interpretation of the estimated regression coefficients is related to the log odds ratio of level k versus level l ,

$$\log\left(\frac{\Pr(Y = v_k|X)}{\Pr(Y = v_l|X)}\right) = (\beta_{k0} - \beta_{l0}) + (\beta_{k1} - \beta_{l1}) X_1 + \dots + (\beta_{kp} - \beta_{lp}) X_p.$$

Note that we just see the differences of pairs of regression coefficients for different levels or classes in this expression. This is related to the fact that the multinomial probabilities must sum to 1, which means that the regression coefficients associated with one level/class are a deterministic function of the rest, and reflects the fact that the linear logistic regression model can only identify each regression coefficient up to a constant. To address this in practice, it is common to choose a **reference** or **baseline** level and set the associated regression coefficients to 0. Often the last level, level M , is chosen as a baseline, meaning that $\beta_{M0} = \beta_{M1} = \dots = \beta_{Mp} = 0$ is assumed. Then the remaining regression coefficients β_{kj} for $k = 1, \dots, M - 1$

and $j > 0$ can be interpreted as the expected change in log odds of $Y = v_k$ relative to the reference level $Y = v_M$ associated with a one unit change in X_j , holding all other predictors constant.

Alternatively, the **softmax** coding, which does not explicitly constrain the regression coefficients can be used. However in this case it is crucial that inference be based on contrasts of the regression coefficients $\beta_{kj} - \beta_{lj}$ to be meaningful.

The probabilities $\Pr(Y = v_k|X)$ are invariant to the coding used.

As with logistic regression for binary outputs, the parameters $\hat{\beta}_{k0}, \hat{\beta}_{k1}, \dots, \hat{\beta}_{kp}$ are estimated using maximum likelihood and there is theory that supports the construction of normal-based confidence intervals and hypothesis tests of the value(s) of the regression coefficients.

Now we will consider alternative **generative** models for classification that:

- Are more stable than linear or logistic regression based methods when classes or levels of the output are “well separated” by the inputs;
- Can be more accurate if the distribution of the inputs conditional on the output is approximately normal;
- Can also easily be used for qualitative outputs with $M > 2$ levels.

These models all build from Bayes theorem,

$$\begin{aligned} \Pr(A|B) &= \frac{\Pr(A, B)}{\Pr(B)} \\ &= \Pr(B|A) \left(\frac{\Pr(A)}{\Pr(B)} \right) \\ &= \frac{\Pr(B|A)\Pr(A)}{\sum_{A'} \Pr(B|A')\Pr(A')}. \end{aligned}$$

In the specific context of classification, the event A corresponds to the event that $Y = v_k$ and the event B corresponds to the event that $X = x$. A generative model specifies a model for the inputs conditional on the value v_k of the output Y that defines

$$\Pr(B|A) = \Pr(X = x|Y = v_k),$$

denoted (with some abuse of notation) as $f_k(X)$, and a model for the **prior** probability that the output Y takes on value v_k ,

$$\Pr(A) = \Pr(Y = v_k),$$

denoted by π_k , which we can estimate using the proportion of values of the output in the training data equal to the k -th value v_k .

The term generative describes that such models are based on how the observed values of the inputs are generated given the level or class of the output.

Probabilities used for classification are given by

$$\Pr(Y = v_k|X) = \frac{\pi_k f_k(X)}{\sum_{l=1}^M \pi_l f_l(X)}.$$

We will talk about three different models, which differ in their assumed distribution of the inputs given the class of the output:

- **Linear discriminant analysis (LDA)**, which assumes that inputs are conditionally normally distributed given $Y = v_k$ with a mean μ_k that depends on the value of the output but a common variance Σ that is common across all possible values of the output v_k

- **Quadratic discriminant analysis (QDA)**, which assumes that inputs are conditionally normally distributed given $Y = v_k$ with a mean μ_k and variance Σ_k that both depend on the value of the output
- **Naive Bayes**

LDA assumes

$$X|Y = v_k \sim \text{normal}(\mu_k, \Sigma)$$

for $k = 1, \dots, M$. Letting \mathcal{W}_k refer to the set of n_k indices for training observations with output equal to v_k , i.e. $\mathcal{W}_k = \{i : Y_i = v_k\}$, estimates $\hat{\mu}_k$ and $\hat{\Sigma}$ of μ_k and Σ can be obtained according to

$$\begin{aligned}\hat{\mu}_k &= \frac{1}{n_k} \sum_{i \in \mathcal{W}_k} x_i \text{ for } k = 1, \dots, M \\ \hat{\Sigma} &= \frac{1}{n - M} \sum_{k=1}^M \sum_{i \in \mathcal{W}_k} (x_i - \hat{\mu}_k)(x_i - \hat{\mu}_k)^\top.\end{aligned}$$

QDA assumes

$$X|Y = v_k \sim \text{normal}(\mu_k, \Sigma_k)$$

for $k = 1, \dots, M$. Estimates $\hat{\mu}_k$ of μ_k are the same as in LDA, and estimates $\hat{\Sigma}$ of Σ can be obtained according to

$$\begin{aligned}\hat{\mu}_k &= \frac{1}{n_k} \sum_{i \in \mathcal{W}_k} x_i \text{ for } k = 1, \dots, M \\ \hat{\Sigma}_k &= \frac{1}{n_k - 1} \sum_{i \in \mathcal{W}_k} (x_i - \hat{\mu}_k)(x_i - \hat{\mu}_k)^\top.\end{aligned}$$

For LDA and QDA, $f_k(X)$ is a multivariate normal density function, which yields probabilities of the form

$$\begin{aligned}\Pr(Y = v_k|X) &= \frac{\frac{\pi_k}{\sqrt{2\pi^p} \sqrt{|\Sigma|_k}} \exp\{-\frac{1}{2}(x - \mu_k)^\top \Sigma_k^{-1}(x - \mu_k)\}}{\sum_{l=1}^m \frac{\pi_l}{\sqrt{2\pi^p} \sqrt{|\Sigma|_l}} \exp\{-\frac{1}{2}(x - \mu_l)^\top \Sigma_l^{-1}(x - \mu_l)\}} \\ &= \frac{\frac{\pi_k}{\sqrt{|\Sigma|_k}} \exp\{-\frac{1}{2}(x - \mu_k)^\top \Sigma_k^{-1}(x - \mu_k)\}}{\sum_{l=1}^m \frac{\pi_l}{\sqrt{|\Sigma|_l}} \exp\{-\frac{1}{2}(x - \mu_l)^\top \Sigma_l^{-1}(x - \mu_l)\}},\end{aligned}$$

which can be estimated by setting $\Sigma_k = \Sigma$ for LDA and plugging in the estimates $\hat{\pi}_k$, $\hat{\mu}_k$, and $\hat{\Sigma}_k$ (or $\hat{\Sigma}$ for LDA) described above.

Note that when we use LDA, we can simplify further,

$$\begin{aligned}\Pr(Y = v_k|X) &= \frac{\frac{\pi_k}{\sqrt{|\Sigma|}} \exp\{-\frac{1}{2}(x - \mu_k)^\top \Sigma^{-1}(x - \mu_k)\}}{\sum_{l=1}^m \frac{\pi_l}{\sqrt{|\Sigma|}} \exp\{-\frac{1}{2}(x - \mu_l)^\top \Sigma^{-1}(x - \mu_l)\}} \\ &= \frac{\pi_k \exp\{x^\top \Sigma^{-1} \mu_k - \frac{1}{2} \mu_k^\top \Sigma^{-1} \mu_k\}}{\sum_{l=1}^m \pi_l \exp\{x^\top \Sigma^{-1} \mu_l - \frac{1}{2} \mu_l^\top \Sigma^{-1} \mu_l\}}.\end{aligned}$$

In practice, instead of computing these normalized probabilities we may recognize that the ordering of the probabilities $\Pr(Y = v_k|X)$ is the same as the ordering of the numerators of $\Pr(Y = v_k|X)$. Furthermore, applying a log transformation retains the ordering while offering more computational stability.

We can also rearrange terms to relate LDA and QDA to the softmax parameterization of multinomial logistic regression,

$$\Pr(Y = v_k | X) = \frac{(\exp\{\log(\pi_k) - \frac{1}{2}(\mu_k^\top \Sigma_k^{-1} \mu_k + \log(|\Sigma_k|))\}) \exp\{x^\top \Sigma_k^{-1} \mu_k - \frac{1}{2} \text{tr}(\Sigma_k^{-1} x x^\top)\}}{\sum_{l=1}^M (\exp\{\log(\pi_l) - \frac{1}{2}(\mu_l^\top \Sigma_l^{-1} \mu_l + \log(|\Sigma_l|))\}) \exp\{x^\top \Sigma_l^{-1} \mu_l - \frac{1}{2} \text{tr}(\Sigma_l^{-1} x x^\top)\}},$$

where we can recognize $\log(\pi_k) - \frac{1}{2}(\mu_k^\top \Sigma_k^{-1} \mu_k + \log(|\Sigma_k|))$ as the intercept, elements of $\Sigma_k^{-1} \mu_k$ as regression coefficients for linear functions of the predictors, and elements of Σ_k^{-1} as regression coefficients for quadratic functions of the predictors and their two-way interactions. With some additional algebra, we could alternatively relate LDA and QDA to the baseline parametrization of multinomial logistic regression.

Note that elements of Σ_k^{-1} are constrained to ensure that Σ_k^{-1} is a symmetric positive definite matrix. When we use LDA, as opposed to QDA, the two-way interactions cancel and vanish.

Note that although we can recognize LDA and QDA as multinomial logistic regression models, fitted LDA and QDA models may not be identical to the multinomial logistic regression models obtained by using the same predictors because the LDA and QDA coefficients are not obtained by maximizing the multinomial logistic regression likelihood.

Naive Bayes addresses a major limitation of LDA and QDA, computational challenges that arise for LDA and QDA when p is large, specifically the challenges that arise when estimating one or more $p \times p$ covariance matrices, by assuming independence of individual predictors X_1, \dots, X_p .

The assumption of independence of individual predictors is rarely something that we actually believe, thus **naive** in the name of the method, but is often practically useful. It also can address the limitation of needing to assume normality of the predictors X .

Assuming independence of individual predictors allows the multivariate density of the predictors $f_k(X)$ to be factorized into p simpler univariate densities $f_{kj}(X_j)$,

$$f_k(X) = \prod_{j=1}^p f_{kj}(X_j).$$

The univariate densities $f_{kj}(X_j)$ can be modeled and estimated flexibly, without assuming that the distribution of each predictor X_j is independent and do not require estimation of any $p \times p$ covariance matrices. There are many ways to model $f_{kj}(X_j)$ and we will not discuss all of them here. Two of the simplest ways are:

- If X_j is continuous, assuming $X_j | Y = v_k \sim \text{normal}(\mu_{kj}, \sigma_{kj}^2)$;
- If X_j is qualitative with relatively few levels L , estimating $f_{kj}(X_j)$ using a histogram.

A more flexible but also more advanced approach may non-parametrically estimate $f_{kj}(X_j)$, e.g. by estimating $f_{kj}(X_j)$ by smoothing a histogram of X_j as done by a **kernel density estimator**.

The probabilities used for classification in Naive Bayes are:

$$\Pr(Y = v_k | X) = \frac{\pi_k \prod_{j=1}^p f_{kj}(X_j)}{\sum_{l=1}^M \pi_l \prod_{j=1}^p f_{lj}(X_j)}.$$

Rewriting this as before to resemble the softmax parametrization of multinomial logistic regression, we have

$$\Pr(Y = v_k | X) = \frac{\exp\left\{\log(\pi_k) + \sum_{j=1}^p \log(f_{kj}(X_j))\right\}}{\sum_{l=1}^M \exp\left\{\log(\pi_l) + \sum_{j=1}^p \log(f_{lj}(X_j))\right\}},$$

where we can recognize $\log(\pi_k)$ as an intercept and $\log(f_{kj}(X_j))$ as potentially nonlinear functions of each predictor. This corresponds to a **generalized additive model**, which we will learn about in future weeks.

To summarize, we have learned about the following types of models for classification: * Linear Regression Models * Logistic Regression Models * Linear Discriminant Analysis * Quadratic Discriminant Analysis * Naive Bayes

All of these models take different approaches to specifying a functional form for and estimating the parameters of the unknown probabilities

$$\Pr(Y = v_k|X).$$

There is no single choice that is always best regardless of the data.

One last point to consider is whether or not we may sometimes want to adjust the threshold used to convert estimated probabilities to predicted values of the output. In the context of binary classification with $M = 2$, we have been using the threshold of 0.5 to obtain predicted values, i.e. we predict $\hat{Y} = 1$ when $\hat{\Pr}(Y = v_k|X) > 0.5$ and $\hat{Y} = 0$ otherwise. This is known to minimize average classification error, treating false positives (predicting $\hat{Y} = 1$ when the true $Y = 0$) and false negatives (predicting $\hat{Y} = 0$ when the true $Y = 1$) equivalently. In real life, there may be reason to weight false positives and false negatives differently. When this is the case, it can make sense to choose an alternative cutoff c and then predict $\hat{Y} = 1$ when $\hat{\Pr}(Y = v_k|X) > c$ and $\hat{Y} = 0$ otherwise. Each cutoff c is associated with a corresponding false positive rate and false negative rate, and the “best” value of c can be chosen by a practitioner either to reflect different costs of false positives versus false negatives or to ensure that the number of false positives does not exceed some prespecified number. We won’t talk about this in detail but it’s something I want you to be aware of.