

Machine Learning and Statistical Learning

Logistic regression

Ewen Gallic
ewen.gallic@gmail.com

MASTER in Economics - Track EBDS - 2nd Year



Let us turn to another category of supervised learning methods: **classification**.

In the regression case, the response variable y was real-valued.

We consider now the case in which it is **qualitative** (or **categorical**).

In this course, we will cover three classifiers:

- the logistic regression (this document)
- the k-nearest neighbors (next week)
- Support vector machines (next week).

Note that trees, random forests, bagging and boosting will be covered in Pierre Michel's course.

Some references

- Berk (2008) . Statistical learning from a regression perspective, volume 14. Springer.
- James et al. (2013) . An introduction to statistical learning, volume 112. Springer.
- Scott Long (1997) . Regression models for categorical and limited dependent variables. Advanced quantitative techniques in the social sciences.

1. Introduction

Introduction

In classification problems, we want to assign a class to a **quantitative** response.

We saw in the introduction of the course that it may be a good idea to estimate a probability for each of the categories and then to assign the class, based on the mode, for example.

Why not fitting a linear regression

In the case in which the response variable is categorical, the linear regression may not be appropriate.

Let us consider a simple case in which we are trying to predict the occupation of individuals, based on their characteristics. Let us suppose there are three classes for the response variable:

The response variable could be:

$$y = \begin{cases} 1 & \text{blue-collar jobs} \\ 2 & \text{white-collar jobs} \\ 3 & \text{professional jobs} \end{cases}$$

Why not fitting a linear regression

Technically, it is possible to fit a linear model using this response variable. But this implies :

- an **ordering** on the outcome, putting white-collar jobs between blue-collar and professional jobs
- the difference between blue-collar and white-collar is the same as the difference between white-collar and professional.

We could code the response variable differently, and the results would be completely different.

Why not fitting a linear regression

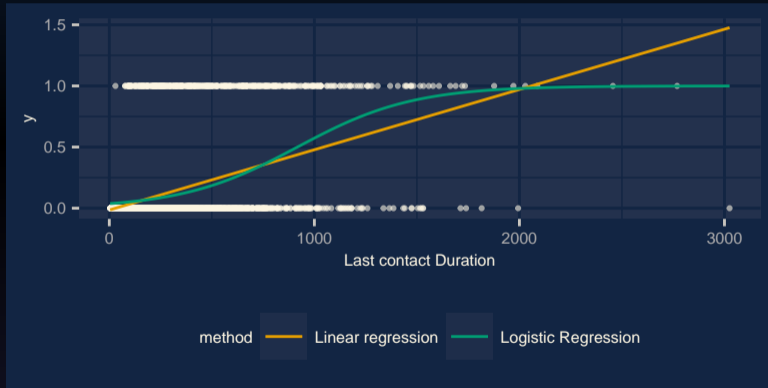
However, if the response variable is categorical, but the variables **can be ordered**, and **if we think that the gap between each category is similar**, a linear regression can be envisaged. For example, if the response variable is some age class:

- [18-25]
- [26-35]
- [36-45]
- [46-55]

If the **response variable is binary**, fitting a linear regression is less problematic. If our response variable is coded using 0 and 1 values:

- as we saw in the previous chapter, flipping the coding of the variable does not change the prediction
- but the prediction may lie outside the $[0, 1]$ interval, making them hard to interpret as probabilities.

Why not fitting a linear regression



2. Logistic regression

Dataset

Let us consider some data on which we will build some examples: Bank Marketing Data Set

Source : Moro, Cortez, and Rita (2014) A Data-Driven Approach to Predict the Success of Bank Telemarketing. Decision Support Systems, Elsevier, 62:22-31.

To download the data: <https://archive.ics.uci.edu/ml/datasets/Bank+Marketing>

Dataset

The data concerns results from direct marketing campaigns (phone calls) of a Portuguese banking institution.

The aim is to predict whether the client will subscribe from a term deposit (variable y).

The original dataset has 45,211 training observations. We will use the provided random subset using only 10% of observations here.

Dataset

- bank client data:
 - ▶ age (numeric)
 - ▶ job : type of job (categorical: 'admin.', 'blue-collar', 'entrepreneur', 'housemaid', 'management', 'retired', 'self-employed', 'services', 'student', 'technician', 'unemployed', 'unknown')
 - ▶ marital : marital status (categorical: 'divorced', 'married', 'single', 'unknown'; note: 'divorced' means divorced or widowed)
 - ▶ education (categorical: 'basic.4y', 'basic.6y', 'basic.9y', 'high.school', 'illiterate', 'professional.course', 'university.degree', 'unknown')
 - ▶ default: has credit in default? (categorical: 'no', 'yes', 'unknown')
 - ▶ balance: average yearly balance, in euros (numeric)
 - ▶ housing: has housing loan? (categorical: 'no', 'yes', 'unknown')
 - ▶ loan: has personal loan? (categorical: 'no', 'yes', 'unknown')

Dataset

- related with the last contact of the current campaign:
 - ▶ contact: contact communication type (categorical: 'cellular', 'telephone')
 - ▶ day: last contact day of the month (numeric: 1 to 31)
 - ▶ month: last contact month of year (categorical: 'jan', 'feb', 'mar', ..., 'nov', 'dec')
 - ▶ duration: last contact duration, in seconds (numeric). Important note: this attribute highly affects the output target (e.g., if duration=0 then y='no'). Yet, the duration is not known before a call is performed. Also, after the end of the call y is obviously known. Thus, this input should only be included for benchmark purposes and should be discarded if the intention is to have a realistic predictive model.
- other attributes:
 - ▶ campaign: number of contacts performed during this campaign and for this client (numeric, includes last contact)
 - ▶ pdays: number of days that passed by after the client was last contacted from a previous campaign (numeric; 999 means client was not previously contacted)
 - ▶ previous: number of contacts performed before this campaign and for this client (numeric)
 - ▶ poutcome: outcome of the previous marketing campaign (categorical: 'failure', 'nonexistent', 'success')

Dataset

The response variable is coded as follows:

$$y = \begin{cases} 1 & \text{if the client has subscribed to a term deposit} \\ 0 & \text{if the client has not subscribed to a term deposit} \end{cases}$$

4000 (88.48%) clients have not subscribed while 521 (11.52%) have.

Logistic regression

Let us consider a simple example in which we try to model the class of the response variable y using the variable `balance` as the sole predictor (which corresponds to the average yearly balance, in euros).

The logistic regression models the **probability** that the client will subscribe to a long term deposit or not. The probability of subscribing given `balance` writes:

$$\mathbb{P}(y = \text{yes} \mid \text{balance})$$

We can rely on this probability to assign a class (yes or no) to the observation.

- For example, we can assign the class `yes` for all observations where $\mathbb{P}(y = \text{yes} \mid \text{balance}) > 0$.
- But we can also select a **different threshold**.

2.1 Model

2.1.1 The Linear Probability Model

A little detour

Before turning to the logistic model, let us a little detour to see what happens when we fit a linear model to the relationship between our predictor \mathbf{x} and our response variable y .

The structural model writes, for $i \in 1, \dots, n$:

$$y_i = \mathbf{x}_i \beta + \varepsilon_i,$$

where ε_i is an error term with zero mean and variance σ^2 .

If the predictor is a binary outcome, the model becomes:

$$y_i = \beta_0 + \beta_1 x_i + \varepsilon_i.$$

The conditional expectation of y given x is:

$$\mathbb{E}(y \mid x) = \beta_0 + \beta_1 x.$$

Linear regression model for binary outcome

As we consider the case in which y is a binary variable, the **unconditional expectation** of y is the probability that the event occurs:

$$\mathbb{E}(y_i) = [1 \times \mathbb{P}(y_i = 1)] + [0 \times \mathbb{P}(y_i = 0)] = \mathbb{P}(y_i = 1).$$

For the regression model:

$$\mathbb{E}(y_i | \mathbf{x}_i) = [1 \times \mathbb{P}(y_i = 1 | \mathbf{x}_i)] + [0 \times \mathbb{P}(y_i = 0 | \mathbf{x}_i)] = \mathbb{P}(y_i = 1 | \mathbf{x}_i).$$

The expected value of y given \mathbf{x} is the probability that $y = 1$ given \mathbf{x} . We therefore have:

$$\mathbb{P}(y_i = 1 | \mathbf{x}_i) = \mathbf{x}_i \beta$$

Problems with the linear regression model and binary outcomes

Multiple problems arise when fitting a linear regression model to a binary outcome response.

First, let us consider the issue regarding **heteroscedasticity**.

If y is a binary outcome random variable, with mean μ , then:

$$\mathbb{E}(Y) = 0 \cdot \mathbb{P}(Y = 0) + 1 \cdot \mathbb{P}(Y = 1) \Rightarrow \mathbb{P}(Y = 1) = \mu.$$

As $\mathbb{E}(Y^2) = 0^2 \cdot \mathbb{P}(Y = 0) + 1^2 \cdot \mathbb{P}(Y = 1) = \mathbb{P}(Y = 1)$, the variance of Y writes:

$$\text{Var}(Y) = \mathbb{E}(Y^2) - \mathbb{E}(Y)^2 = \mu(1 - \mu).$$

Problems with the linear regression model and binary outcomes

We know that $\mathbb{E}(y | \mathbf{x}) = \mathbf{x}\beta$, hence the conditional variance of y depends on \mathbf{x} :

$$\text{Var}(y | \mathbf{x}) = \mathbb{P}(y = 1 | \mathbf{x}) [1 - \mathbb{P}(y = 1 | \mathbf{x})] = \mathbf{x}\beta(1 - \mathbf{x}\beta).$$

Hence, the variance of the errors depends on the values of \mathbf{x} and is not constant.

As a consequence, the least squares estimate of β is inefficient and the standard errors are biased.

Problems with the linear regression model and binary outcomes

Let us now turn to the **normality assumption**.

As the response variable only takes values between 0 and 1, the error at point x_0 must be equal to:

$$\begin{cases} 1 - \mathbb{E}(y \mid x_0) & \text{if } y = 1 \\ 0 - \mathbb{E}(y \mid x_0) & \text{if } y = 0 \end{cases}$$

Hence, the errors cannot be normally distributed, which leads to biased estimates.

Another problem arises, as we saw previously, that of **nonsensical predictions**: some predictions may be negative or greater than one.

This is problematic, as we are predicting a probability.

2.1.2 A latent model for binary variables

A latent model for binary variables

To overcome the mentioned issues, $\mathbb{P}(y \mid \mathbf{x})$ must be modeled using a function that gives outputs between 0 and 1 for all values of \mathbf{x} .

Let us suppose there is an **unobserved** (or **latent**) variable y^* ranging from $-\infty$ to $+\infty$ that generates the observed values of y :

- the large values of y^* are observed as 1
- the small values of y^* are observed as 0.

The latent variable

For example, let us consider the clients who subscribe to a term deposit (the observed response variable y). This variable can be observed in two states:

- the client has subscribed
- the client has not subscribed.

Some clients may be closed to the decision of subscribing, other may be very far from it. In both cases, we observe $y = 0$.

The idea behind the **latent** y^* is that there exists an underlying **propensity to subscribe** that generates the observed state.

We **cannot observe** y^* , but at some point, a **change in y^* results in a change in what we observe**:

- for example, as the balance increases, the propensity to subscribe to a term deposit increases too.

The latent variable

The latent y^* is assumed to be linearly related to the observed predictors:

$$y_i^* = \mathbf{x}_i \beta + \varepsilon_i.$$

The observed variable y writes:

$$y_i = \begin{cases} 1 & \text{if } y_i^* > \tau \\ 0 & \text{if } y_i^* \leq \tau \end{cases} \quad (1)$$

where τ is called the **threshold** (or **cutpoint**), assumed to be equal to zero for now.

Estimating the model: intuition

The dependant variable is not observed. The model cannot be estimated using least squares.

We will see in a next section how to estimate the model using **maximum likelihood** (ML).

The ML estimation requires some assumptions about the distribution of the errors ε .

For the **logit model**, we assume that the **distribution of the errors is logistic**.

We also assume that:

- $\mathbb{E}(\varepsilon \mid \mathbf{x}) = 0$
- $\mathbb{V}ar(\varepsilon \mid \mathbf{x}) = \pi^2/3$

Note: if we assume the errors are $\mathcal{N}(0, 1)$, it results in the probit model.

Standard logistic distribution

The probability distribution function of the standard logistic distribution writes:

$$\lambda(\varepsilon) = \frac{\exp(\varepsilon)}{[1 + \exp(\varepsilon)]^2}.$$

The cumulative distribution function writes:

$$\Lambda(\varepsilon) = \frac{\exp(\varepsilon)}{1 + \exp(\varepsilon)}$$

Standard logistic distribution

The logistic distribution can be rescaled to have a unit variance, resulting in the **standardized logistic distribution**.

The probability distribution becomes:

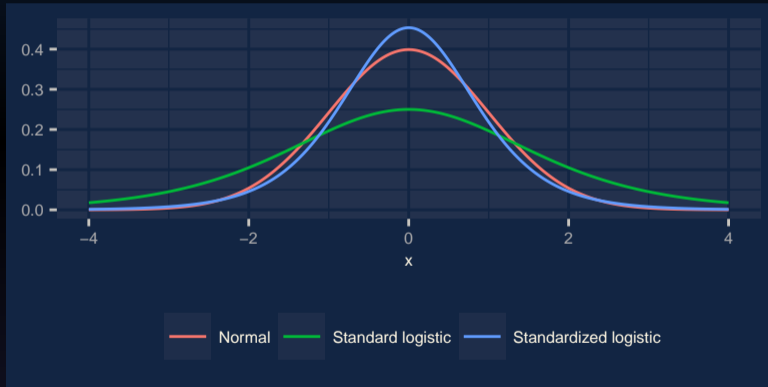
$$\lambda^S(\varepsilon) = \frac{\gamma \exp(\gamma\varepsilon)}{[1 + \exp(\gamma\varepsilon)]^2}.$$

The cumulative distribution function writes:

$$\Lambda^S(\varepsilon) = \frac{\exp(\gamma\varepsilon)}{1 + \exp(\gamma\varepsilon)},$$

where $\gamma = \pi/\sqrt{3}$

Standard logistic distribution



Estimating the model: intuition

Assuming the distribution of the errors is useful to compute the probability of $y = 1$ for a given \mathbf{x} .

Indeed:

$$\begin{aligned}\mathbb{P}(y = 1 \mid \mathbf{x}) &= \mathbb{P}(y^* > 0 \mid \mathbf{x}) \\ &= \mathbb{P}(\mathbf{x}\beta + \varepsilon > 0 \mid \mathbf{x}) \\ &= \mathbb{P}(\varepsilon > -\mathbf{x}\beta \mid \mathbf{x}) \\ &= \mathbb{P}(\varepsilon \leq \mathbf{x}\beta \mid \mathbf{x}) \quad (\text{symmetric cdf}) \\ &= F(\mathbf{x}\beta) = \Lambda(\mathbf{x}\beta).\end{aligned}$$

Hence the probability of observing an event given \mathbf{x} is the cumulative density evaluated at $\mathbf{x}\beta$.

Deriving the model without appealing to a latent variable

Another way of deriving the logit model without appealing to a latent variable is to specify a nonlinear model relating the predictors to the probability of an event.

In a first step, the probability of an event is **transformed into the odds**:

$$\frac{\mathbb{P}(y = 1 \mid \mathbf{x})}{\mathbb{P}(y = 0 \mid \mathbf{x})} = \frac{\mathbb{P}(y = 1 \mid \mathbf{x})}{1 - \mathbb{P}(y = 1 \mid \mathbf{x})}$$

The **odds** are used to evaluate how often an event happens relative to how often it does not.

The odds range from 0 (when $\mathbb{P}(y = 1 \mid \mathbf{x}) = 0$) to ∞ (when $\mathbb{P}(y = 1 \mid \mathbf{x}) = 1$).

Deriving the model without appealing to a latent variable

The log of the odds, as known as the **logit** ranges from $-\infty$ to ∞ .

This suggests a model that is linear in the logit:

$$\ln \left[\frac{\mathbb{P}(y = 1 \mid \mathbf{x})}{1 - \mathbb{P}(y = 1 \mid \mathbf{x})} \right] = \mathbf{x}\beta \quad (2)$$

Hence, a marginal change in \mathbf{x} changes the log odds by the coefficient β associated with \mathbf{x} .

We can easily show that Eq. 2 equivalent to the **logistic function**:

$$\mathbb{P}(y = 1 \mid \mathbf{x}) = \frac{\exp(\mathbf{x}\beta)}{1 + \exp(\mathbf{x}\beta)}. \quad (3)$$

Deriving the model without appealing to a latent variable

From Eq. 3, it can easily be shown that:

$$\frac{\mathbb{P}(y = 1 \mid \mathbf{x})}{1 - \mathbb{P}(y = 1 \mid \mathbf{x})} = \exp(\mathbf{x}\beta) \quad (4)$$

Hence, we see that a marginal increase in \mathbf{x} can be equivalently interpreted as follows: it multiplies the odds by the coefficient β associated with \mathbf{x}

2.2 Estimation

Estimation

In the case of a binary response variable and the logit model, we want to estimate the coefficients β_0 and β_1 from the logistic function:

$$\mathbb{P}(y = 1 \mid \mathbf{x}) = \frac{\exp(\beta_0 + \beta_1 \mathbf{x})}{1 + \exp(\beta_0 + \beta_1 \mathbf{x})}.$$

First, let us define p as the probability of observing the value of y :

$$p_i = \begin{cases} \mathbb{P}(y_i = 1 \mid \mathbf{x}_i) & \text{if } y_i = 1 \text{ is observed} \\ 1 - \mathbb{P}(y_i = 1 \mid \mathbf{x}_i) & \text{if } y_i = 0 \text{ is observed} \end{cases} \quad (5)$$

Assuming the observations are independent, the likelihood equation writes:

$$\mathcal{L}(\beta \mid y, \mathbf{x}) = \prod_{i=1}^n p_i \quad (6)$$

Estimation

Hence, combining Eq. 5 and Eq. 6:

$$\mathcal{L}(\beta | y, \mathbf{x}) = \prod_{y=1} \mathbb{P}(y_i = 1 | \mathbf{x}_i) \prod_{y=0} [1 - \mathbb{P}(y_i = 1 | \mathbf{x}_i)] \quad (7)$$

We have seen that $\mathbb{P}(y_i = 1 | \mathbf{x}_i) = F(\mathbf{x}_i\beta)$, *i.e.*, $\mathbb{P}(y_i = 1 | \mathbf{x}_i) = \Lambda(\mathbf{x}_i\beta)$ in the case of the logit model.

Hence, the likelihood equations also writes:

$$\mathcal{L}(\beta | y, \mathbf{x}) = \prod_{y=1} \Lambda(\mathbf{x}_i\beta) \prod_{y=0} [1 - \Lambda(\mathbf{x}_i\beta)] \quad (8)$$

Estimation

The log likelihood thus writes:

$$\ell(\beta \mid y, \mathbf{x}) := \ln \mathcal{L}(\beta \mid y, \mathbf{x}) = \sum_{y=1} \ln \Lambda(\mathbf{x}_i \beta) + \sum_{y=0} \ln [1 - \Lambda(\mathbf{x}_i \beta)] \quad (9)$$

$$= \sum_{i=1}^n \mathbf{1}_{(y_i=1)} \ln \Lambda(\mathbf{x}_i \beta) + (1 - \mathbf{1}_{y_i=1}) \ln [1 - \Lambda(\mathbf{x}_i \beta)] \quad (10)$$

The estimates of β are chosen to **maximize** the likelihood function.

Effect of a variable on the odds

Let us get back to our clients data.

```
::: .cell layout-align="center"
```

```
fig.showtext='false'
```

	Model 1
(Intercept)	-3.25593456*** (0.08457673)
duration	0.00354955*** (0.00017136)
AIC	2705.75264185
BIC	2718.58561882
Log Likelihood	-1350.87632092
Deviance	2701.75264185
Num. obs.	4521
*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$	

Table 1: Statistical models

```
:::
```

We see that $\beta_1 = 0.00354955$:

- a unit increase in the duration of the last call is associated with an increase in the probability of subscribing to a term deposit
- a unit increase in the duration of the last call is associated with an increase in the log odds of the response by 0.00354955.

2.2.1 Numerical methods

Numerical methods

Let us look at how to use numerical methods to estimate the coefficients.

We will look at different ways of doing so, using iterative solutions.

First, we need to begin with an initial guess that we will denote β_0 : the **start values**.

At each iteration, we will try to improve the previous guess by adding a vector ζ_n of adjustments:

$$\beta_{n+1} = \beta_n + \zeta_n.$$

The iterations stops as soon as a convergence criterion is met:

- roughly, when the gradient of the log likelihood gets close to 0
- or when the estimates do not change from one step to the next.

Sometimes, there is no convergence, and we do not get the ML estimates.

Numerical methods

ζ_n can be expressed as:

$$\zeta_n = D_n \gamma_n,$$

- γ_n : **gradient vector**:

- ▶ $\gamma_n = \frac{\partial \ln \mathcal{L}}{\partial \beta_n}$

- ▶ it indicates the direction of the change in the log likelihood for a change in the parameters

- D_n : **direction matrix**

- ▶ it reflects the curvature of the likelihood function (how rapidly the gradient is changing)

The Newton-Raphson Method

The rate of change in the slope of $\ln \mathcal{L}$ is indicated by the second derivatives, contained in the Hessian matrix.

In the case of a single predictor, we have two parameters to estimate, β_0 and β_1 . The Hessian matrix is thus:

$$\frac{\partial^2 \ln \mathcal{L}}{\partial \beta \partial \beta^\top} = \begin{pmatrix} \frac{\partial^2 \ln \mathcal{L}}{\partial \beta_0 \partial \beta_0} & \frac{\partial^2 \ln \mathcal{L}}{\partial \beta_0 \partial \beta_1} \\ \frac{\partial^2 \ln \mathcal{L}}{\partial \beta_1 \partial \beta_0} & \frac{\partial^2 \ln \mathcal{L}}{\partial \beta_1 \partial \beta_1} \end{pmatrix}$$

If $\frac{\partial^2 \ln \mathcal{L}}{\partial \beta_0 \partial \beta_0}$ is large relative to $\frac{\partial^2 \ln \mathcal{L}}{\partial \beta_1 \partial \beta_1}$:

- the gradient is changing more rapidly as β_0 changes than β_1 changes.

In that case, smaller adjustments to the estimates of β_0 would be indicated.

The Newton-Raphson Method

Recall that:

$$\ell(\beta | y, \mathbf{x}) = \sum_{i=1}^n \mathbf{1}_{(y_i=1)} \ln \Lambda(\mathbf{x}_i \beta) + (1 - \mathbf{1}_{(y_i=1)}) \ln [1 - \Lambda(\mathbf{x}_i \beta)]$$

Hence, it can easily be shown that, for all $j = 0, 1, \dots, p$:

$$\frac{\partial \ell(\beta | y, \mathbf{x})}{\partial \beta_j} = \sum_{i=1}^n \mathbf{x}_{ij} [\mathbf{1}_{(y_i=1)} - \Lambda(\mathbf{x}_i \beta)]$$

In matrix form:

$$\frac{\partial \ell(\beta | y, \mathbf{x})}{\partial \beta} = \sum_{i=1}^n \mathbf{x}_i [\mathbf{1}_{(y_i=1)} - \Lambda(\mathbf{x}_i \beta)]$$

The Newton-Raphson Method

The Newton-Raphson algorithm requires the Hessian matrix:

$$\frac{\partial^2 \ell(\beta \mid y, \mathbf{x})}{\partial \beta \partial \beta^\top} = - \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top [\Lambda(\mathbf{x}_i \beta)(1 - \Lambda(\mathbf{x}_i \beta))]$$

The following equation is used by the Newton-Raphson algorithm at each iteration:

$$\beta_{n+1} = \beta_n - \left(\frac{\partial^2 \ln \mathcal{L}}{\partial \beta_n \partial \beta_n^\top} \right)^{-1} \frac{\partial \ln \mathcal{L}}{\partial \beta_n}$$

The Newton-Raphson Method

To illustrate the iterative process of the Newton-Raphson algorithm, let us go through an example, using matrix notations.

- y : column vector of the response variable, of dimension $n \times 1$
- X : matrix of predictors, including the constant, of dimension $n \times (p + 1)$
- p_n : vector of fitted probabilities at iteration n , of dimension $n \times 1$
- Ω : diagonal matrix of weights, of dimension $n \times n$
 - ▶ the i th element of the diagonal is: $\Lambda(X_i\beta_{n-1})(1 - \Lambda(X_i\beta_{n-1}))$

The Newton-Raphson Method

Using these notations, we can write the gradient and the Hessian as follows:

$$\frac{\partial \ell(\beta \mid y, \mathbf{x})}{\partial \beta} = X^\top (y - p_{n-1})$$
$$\frac{\partial^2 \ell(\beta \mid y, \mathbf{x})}{\partial \beta \partial \beta^\top} = -X^\top \Omega_{n-1} X$$

The Newton-Raphson Method

At each step:

$$\begin{aligned}\beta_n &= \beta_{n-1} + (X^\top \Omega X)^{-1} X^\top (y - p_{n-1}) \\ &= \underbrace{(X^\top \Omega X)^{-1} (X^\top \Omega X)}_{I_n} \beta_{n-1} + (X^\top \Omega X)^{-1} X^\top (y - p_{n-1}) \\ &= (X^\top \Omega X)^{-1} X^\top \Omega [X \beta_{n-1} + \Omega^{-1} X^\top (y - p_{n-1})] \\ &= (X^\top \Omega X)^{-1} X^\top \Omega z\end{aligned}$$

where $z = [X \beta_{n-1} + \Omega^{-1} X^\top (y - p_{n-1})]$.

Written this way, z can be viewed as the response of the model, X the predictors and β_n the coefficients of a weighted least squares regression.

2.3 Predictions

Predictions

Is is easy to use the estimates to perform some predictions.

For example, in our example of clients subscribing to term deposit:

$$y_i = \beta_0 + \beta_1 \text{duration}_i + \varepsilon_i$$

We estimated that $\hat{\beta}_0 = -3.25593456$ and $\hat{\beta}_1 = 0.00354955$.

Hence, the probability of subscribing to a term deposit for an individual whose last call duration was 250 seconds is given by:

$$\begin{aligned}\hat{\mathbb{P}}(y = 1 \mid \text{duration} = 250) &= \frac{\exp(\hat{\beta}_0 + \hat{\beta}_1 \times 250)}{1 + \exp(\hat{\beta}_0 + \hat{\beta}_1 \times 250)} \\ &= \frac{\exp(-3.25593456 + 0.00354955 \times 250)}{1 + \exp(-3.25593456 + 0.00354955 \times 250)} \\ &= 0.0856028\end{aligned}$$

2.4 Interpretation

2.4.1 The effect of the Parameters

Binary Regression Model

Let us consider the binary response model with a single predictor \mathbf{x} . The model writes:

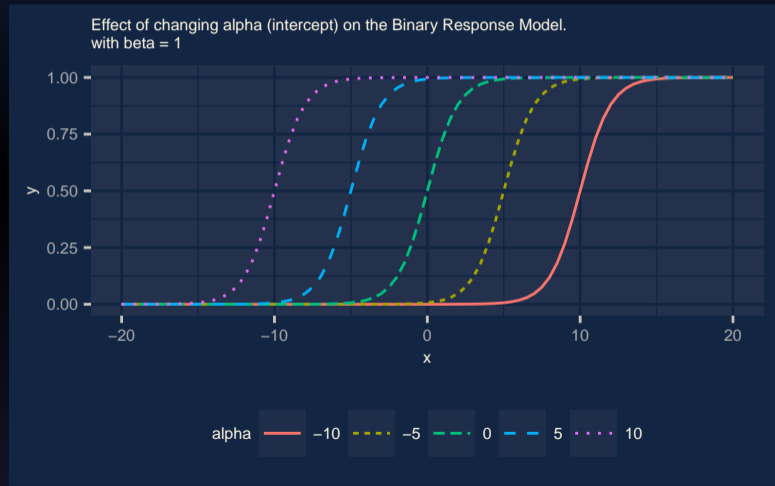
$$\mathbb{P}(y = 1 \mid \mathbf{x}) = F(\alpha + \beta\mathbf{x}),$$

i.e., for a logit model:

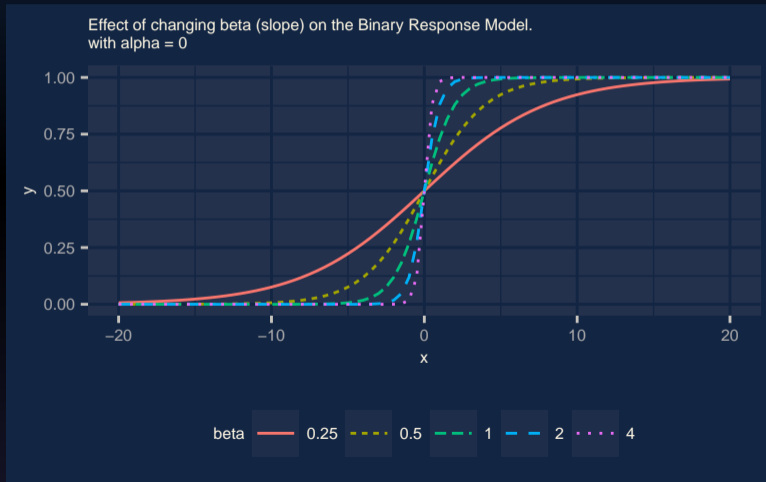
$$\mathbb{P}(y = 1 \mid \mathbf{x}) = \frac{\exp(\alpha + \beta\mathbf{x})}{1 + \exp(\alpha + \beta\mathbf{x})}$$

Let us vary α and β to observe how it affects the probability.

Varying β



Varying α



3. Binary Response Model with Multiple Predictors

Let us now consider multiple regressors (only two, for illustrative purposes):

$$\mathbb{P}(y = 1 \mid \mathbf{x}) = F(\alpha + \beta_1 x_1 + \beta_2 x_2),$$

i.e., for a logit model:

$$\mathbb{P}(y = 1 \mid \mathbf{x}) = \frac{\exp(\alpha + \beta_1 x_1 + \beta_2 x_2)}{1 + \exp(\alpha + \beta_1 x_1 + \beta_2 x_2)}$$

Let us consider the following values:

- $\alpha = 1$
- $\beta_1 = 1$
- $\beta_2 = .75$

Let us vary α and β_1 and β_2 to observe how it affects the probability.

Varying α



Figure 1: Effects of changing α on the Logit Model $\mathbb{P}(y = 1 \mid x_1, x_2)$, with $\beta_1 = 1$ and $\beta_2 = .75$

Varying β_1



Figure 2: Effects of changing α on the Logit Model $\mathbb{P}(y = 1 \mid x_1, x_2)$, with $\alpha = 1$ and $\beta_2 = .75$

Varying β_2



Figure 3: Effects of changing α on the Logit Model $\mathbb{P}(y = 1 \mid x_1, x_2)$, with $\alpha = 1$ and $\beta_1 = 1$

Using Predicted Probabilities

To know the effect of a covariate on the probability of the event, we can look at how this probability changes when varying a covariate.

However, where there are more than two predictors, the surface response cannot be plotted as it was done previously when facing a single covariate.

In such cases, the interpretation depends on whether or not **the relationship between the response and the predictors can be considered as linear**. To that end, we need to compute the **range of probabilities**.

Using Predicted Probabilities

The predicted probability of an event given some values for the predictors \mathbf{x} for the i th individual writes:

$$\hat{\mathbb{P}}(y = 1 \mid \mathbf{x}_i) = F(\mathbf{x}_i \hat{\beta})$$

The minimum and maximum probabilities **in the sample** write, respectively:

$$\min \hat{\mathbb{P}}(y = 1 \mid \mathbf{x}_i) = \min_i F(\mathbf{x}_i \hat{\beta})$$

$$\max \hat{\mathbb{P}}(y = 1 \mid \mathbf{x}_i) = \max_i F(\mathbf{x}_i \hat{\beta})$$

- If the range of probabilities ($\max \hat{\mathbb{P}}(y = 1 \mid \mathbf{x}_i) - \min \hat{\mathbb{P}}(y = 1 \mid \mathbf{x}_i)$) is **between 0.2 and 0.8**, the relationship between \mathbf{x}_i and y can be considered as linear:
 - ▶ thus, the marginal effect of \mathbf{x}_i can be obtained using simple measures
- Otherwise, other methods need to be used.

The effect of Each Variable on the Predicted Probability

To assess the effect of a (numerical) variable on the predicted probability, we can look at how the predicted probability changes as the variable varies **from its minimum to its maximum value**.

To do so, for the k th variable:

- the other variables are set to their average value ($\bar{\mathbf{x}}_{(-k)}$)
- we compute the predicted probability when \mathbf{x}_k is at its maximum value ($\hat{\mathbb{P}}(y = 1 \mid \bar{\mathbf{x}}_{(-k)}, \max \mathbf{x}_k)$)
- we compute the predicted probability when \mathbf{x}_k is at its minimum value ($\hat{\mathbb{P}}(y = 1 \mid \bar{\mathbf{x}}_{(-k)}, \min \mathbf{x}_k)$)

The **predicted change in the probability** as \mathbf{x}_k varies from its min to its max is then computed as:

$$\hat{\mathbb{P}}(y = 1 \mid \bar{\mathbf{x}}_{(-k)}, \max \mathbf{x}_k) - \hat{\mathbb{P}}(y = 1 \mid \bar{\mathbf{x}}_{(-k)}, \min \mathbf{x}_k)$$

The effect of Each Variable on the Predicted Probability

If, among the $(-k)$ covariates, some are non numerical:

- it is possible to set these variables to the mode of the distribution
- or different combinations of the levels can be envisaged.

Discrete Change

Let us now consider the **partial change in y** .

Let $\mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_k)$ be the probability of an event given some values for $\mathbf{x}_{(-k)}$ and a specific value for \mathbf{x}_k .

Let $\mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_{k+\delta})$ be the probability of an event after a variation of δ in \mathbf{x}_k , keeping all other variables unchanged.

The **discrete change** in the probability for a change of δ in x_k is:

$$\mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_{k+\delta}) - \mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_k)$$

Discrete Change

$$\mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_{k+\delta}) - \mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_k)$$

Following a variation of δ in \mathbf{x}_k , keeping all the other variables unchanged, the predicted probability of an event is changed by $\mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_{k+\delta}) - \mathbb{P}(y = 1 \mid \mathbf{x}_{(-k)}, \mathbf{x}_k)$.

Some usual values are picked for δ :

- a **(centered) unit change**: increasing $\bar{\mathbf{x}}_k$ to $\bar{\mathbf{x}}_k + 1$
- a **standard deviation change**:



$$\mathbb{P}\left(y = 1 \mid \bar{\mathbf{x}}_{(-k)}, \bar{\mathbf{x}}_k + \frac{\text{std}(\mathbf{x}_k)}{2}\right) - \mathbb{P}\left(y = 1 \mid \bar{\mathbf{x}}_{(-k)}, \bar{\mathbf{x}}_k - \frac{\text{std}(\mathbf{x}_k)}{2}\right)$$

- a **change from 0 to 1 for dummy variables**

4. Assessing the quality of classification

Assessing the quality of classification

When fitting a classifier on sample data and testing it on a test sample, we can compare the predictions with the *observed values* and compute some metrics to assess goodness of fit.

4.1 Confusion tables

Confusion tables

It is possible to use **confusion tables** to compare the predictions of the fitted model to the actual classes. These tables cross-tabulate the **observed classes** against the classes that the **classifier assigns**.

Let us consider that we trained a model on a train dataset to be able to predict a categorical response variable with two classes: “yes” and “no”.

		<i>Predicted class</i>		Model error
		no predicted	yes predicted	
<i>True class</i>	no	a	b	$b/(a + b)$
	yes	c	d	$c/(c + d)$
Use error		$c/(a + c)$	$b/(b + d)$	Overall error $= \frac{(b+c)}{(a+b+c+d)}$

Table 2: A confusion table.

Confusion tables

We can create these tables for both predictions made on the training and the testing samples.

Four kinds of performance assessment can be made from confusion tables.

1. Looking at the overall proportion of cases **incorrectly classified**

		Predicted class		Model error
		no predicted	yes predicted	
True class	no	a	b	$b/(a + b)$
	yes	c	d	$c/(c + d)$
Use error		$c/(a + c)$	$b/(b + d)$	Overall error $= \frac{(b+c)}{(a+b+c+d)}$

Confusion tables

2. Sometimes, we may want to be **more accurate for one class than for another** (e.g., if we are trying to detect a cancer). Looking at **false positive** and **false negative**.

		Predicted class		Model error
		no predicted	yes predicted	
True class	no	a	b	$b/(a+b)$
	yes	c	d	$c/(c+d)$
Use error		$c/(a+c)$	$b/(b+d)$	Overall error $= \frac{(b+c)}{(a+b+c+d)}$

Confusion tables

3. The **column proportions** help evaluate how useful the classifier results are likely to be if put to work: what happens when forecasting.

		<i>Predicted class</i>		
		no predicted	yes predicted	
<i>True class</i>	no	a	b	$b/(a + b)$
	yes	c	d	$c/(c + d)$
Use error		$c/(a + c)$	$b/(b + d)$	Overall error $= \frac{(b+c)}{(a+b+c+d)}$

Here the number of false positive is b .

Confusion tables

4. The **ratio of the number of false negative to the number of false positives** shows how the results are trading one kind of error for the other:
 - ▶ here c/b : if b is 5 times larger than c , there are 5 false positives for every false negative: the classifier produces results in which false negatives are five times more important than false positives.

Example

Let us go back to our logistic regression classification, on the data of client who subscribe or not to a term deposit. We model the probability of subscribing using the information on duration, education, campaign.

We obtain the following confusion matrix:

Table 3: Confusion matrix for the logistic regression.

y	No	Yes
No	3940	60
Yes	434	87

The **overall error** (we only consider a training sample here) is $\frac{434+60}{4,521} = 0.1092679$.

Example

Table 4: Confusion matrix for the logistic regression.

y	No	Yes
No	3940	60
Yes	434	87

- **false positive**: the classifier predicted that the client would subscribe but he actually did not

▶ $\frac{60}{3,940+60} = 0.015$

- **false negative**: the classifier predicted that the client would not subscribe but he actually did

▶ $\frac{434}{434+87} = 0.8330134$

We notice that the classifier is more accurate to classify the class "No" (when clients do not subscribe).

Example

Table 5: Confusion matrix for the logistic regression.

y	No	Yes
No	3940	60
Yes	434	87

Let us look at the ratio of the number of false negative to the number of false positives: $\frac{434}{60} = 7.233333$.

Hence, for every false positive (predicting that the client will subscribe), there are 7.23 false negative (predicting that the client will not subscribe when in fact he would).

4.2 The ROC curve

Costs

In the point mentioned in the previous slide, there is an underlying idea of the existence of **some costs**.

Recall that c is the number of false negative in our table, and b is the number of false positives.

The **off-diagonal cells** of the confusion table give us the number of **false positive** and **false negatives**.

The ratio of the cells c/b tells us that for every false positive, there are c/b false negatives:

- one false positive is “worth” c/b false negatives
- the cost is c/b times greater: **it is c/b times more costly to misclassify a false positive than to misclassify a false negative.**

Costs

So here, in the example of the clients who subscribe or not to a term deposit, for every false positive (predicting that the client will subscribe), there are 7.23 false negative (predicting that the client will not subscribe when in fact he would):

- this may not be acceptable for a business industry...
- we may want to be able to sacrifice some accuracy in predicting false positive in favor of better predict false negative.

Threshold of 0.5

So far, to assign a class to an observation x_0 in the case of a binary categorical response variable, we have used a threshold of 0.5:

- predicting default class if $\mathbb{P}(Y = 1 \mid X = x_0) \geq 0.5$ for a given observation $X = x_0$
- predicting alternative class otherwise

I.e., this corresponds to assign an observation x_0 the class k for which $\mathbb{P}(Y = k \mid X = x_0)$ is the highest.

Reminder

		<i>Predicted class</i>		Model error
		-	+	
<i>True class</i>	-	TN	FP	$FPR = FP / (TN + FP)$
	+	FN	TP	$FNR = FN / (FN + TP)$
Use error		$FOR = FN / (TN + FN)$	$FDR = FP / (FP + TP)$	Overall error $ACC = \frac{(FP + FN)}{(TN + FP + FN + TP)}$

Table 6: A confusion table.

Changing the threshold value

Now, let us consider the bank data again and our logistic regression.

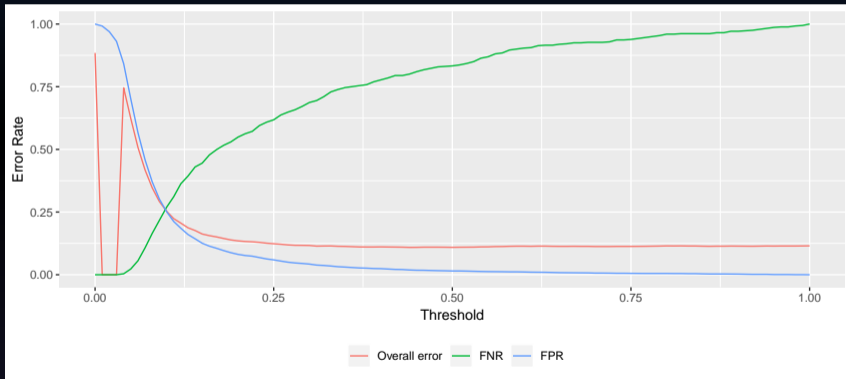


Figure 4: Varying the threshold on the logistic regression of the bank example. The red line represents the overall error. The blue line represents the fraction of misclassified clients among the non-subscribers. The green line represents the fraction of misclassified clients among the subscribers.

Changing the threshold value

In the previous graph, we can note that:

- a threshold of 0.5 minimizes the overall error rate
- reducing the value of the threshold:
 - ▶ diminishes the error rate among the individuals who subscribed ($FNR = FN/(FN + TP)$)
 - ▶ but increases the error rate among the individuals who did not ($FPR = FP/(FP + TN)$).

The ROC curve

For a binary categorical response variable, it is possible to draw a **ROC curve** which allows us to display the two types of error for different thresholds.

ROC means *receiver operating characteristics*.

The ROC curve is obtained by plotting the true positive rate against false positive rate, at different values of the threshold.

The **area under the ROC curve (AUC)** allows us to assess the overall performance of a classifier:

- a value of 0.5 corresponds to the AUC for a random classifier
- the closer the AUC is to 1, the better

This is therefore a metric that can be used to select between different classifiers...

The ROC curve

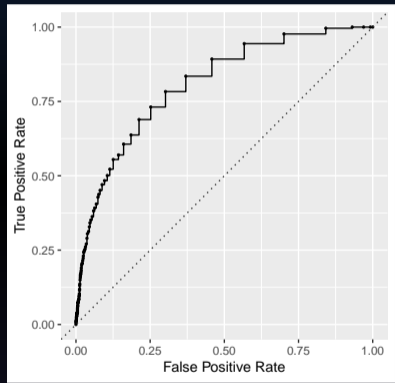


Figure 5: ROC curve for the logistic regression on the bank data.

In this example, the AUC is 0.8221.

5. References

- Berk, Richard A. 2008. *Statistical Learning from a Regression Perspective*. Vol. 14. Springer.
- James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 112. Springer.
- Moro, Sérgio, Paulo Cortez, and Paulo Rita. 2014. "A Data-Driven Approach to Predict the Success of Bank Telemarketing." *Decision Support Systems* 62: 22–31.
- Scott Long, J. 1997. "Regression Models for Categorical and Limited Dependent Variables." *Advanced Quantitative Techniques in the Social Sciences* 7.