

Chapter 8

Credit Scoring

Credit scoring provides a statistical assessment of a borrower’s creditworthiness that helps financial institutions in making decisions on loan applications. In this chapter, we review the uses of discriminant analysis, and binomial logistic regression, with application to credit scoring. We also cover the properties of receiver operating characteristics curves.

8.1 Discriminant Analysis	195
8.2 Decision Rule	199
8.3 Logistic Regression	202
8.4 ROC Curve	207
Exercises	212


8.1 Discriminant Analysis

Consider a set Ω of credit applicants which is partitioned as

$$\Omega = B \cup G$$

into a subset G of “good” (or solvent) applicants, and a subset B of “bad” (or insolvent) applicants, with $B \cap G = \emptyset$. Applicants are selected at random within the set Ω according to a probability distribution \mathbb{P} , so that

$$\mathbb{P}(B) + \mathbb{P}(G) = 1.$$



Here, $\mathbb{P}(B)$ represents the probability that an applicant chosen at random will default, while $\mathbb{P}(G)$ represents the probability that a randomly selected applicant is solvent. In addition, each applicant $\omega \in \Omega$ is assigned a real-valued rating (or score) $X(\omega)$ via a *random variable**

$$\begin{aligned} X : \Omega &\longrightarrow \mathbb{R} \\ \omega &\longmapsto X(\omega) \end{aligned}$$

with probability density function $f_X : \mathbb{R} \rightarrow \mathbb{R}_+$.

Definition 8.1. 1) *The function*

$$\begin{aligned} \mathbb{R} &\longrightarrow [0, 1] \\ x &\longmapsto \mathbb{P}(B \mid X = x) \end{aligned}$$

is respectively called the probability default curve.

2) *The function*

$$\begin{aligned} \mathbb{R} &\longrightarrow [0, 1] \\ x &\longmapsto \mathbb{P}(G \mid X = x) \end{aligned}$$

is called the probability acceptance curve.

Denoting by $f_X(x \mid B)$, resp. $f_X(x \mid G)$, the probability density function of X given B , resp. X given G , we have

$$f_X(x) = f_X(x \mid G)\mathbb{P}(G) + f_X(x \mid B)\mathbb{P}(B), \quad x \in \mathbb{R},$$

and the Bayes formula yields

$$\mathbb{P}(B \mid X = x) = f_X(x \mid B) \frac{\mathbb{P}(B)}{f_X(x)}$$

* See (MOE and UCLES 2016, page 14 lines 4-5) and (MOE and UCLES 2020, page 19 lines 4-5).

$$= \frac{\mathbb{P}(B)f_X(x | B)}{\mathbb{P}(G)f_X(x | G) + \mathbb{P}(B)f_X(x | B)}, \quad (8.1)$$

and

$$\mathbb{P}(G | X = x) = \frac{\mathbb{P}(G)f_X(x | G)}{\mathbb{P}(G)f_X(x | G) + \mathbb{P}(B)f_X(x | B)}.$$

Definition 8.2. 1) The True Positive Rate (TPR) is given by the tail distribution function

$$\bar{F}_G(x) := \mathbb{P}(X > x | G) = \int_x^\infty f_X(y | G)dy, \quad x \in \mathbb{R}.$$

2) The False Positive Rate (FPR) is given by the tail distribution function

$$\bar{F}_B(x) := \mathbb{P}(X > x | B) = \int_x^\infty f_X(y | B)dy, \quad x \in \mathbb{R}.$$

Example. In case X is Gaussian distributed given $\{G, B\}$ with the conditional densities

$$f_X(x | G) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu_G)^2/(2\sigma^2)} \quad (8.2)$$

and

$$f_X(x | B) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu_B)^2/(2\sigma^2)}, \quad (8.3)$$

with $\mu_B < \mu_G$, we have

$$f_X(x) = \frac{\mathbb{P}(G)}{\sqrt{2\pi}\sigma} e^{-(x-\mu_G)^2/(2\sigma^2)} + \frac{\mathbb{P}(B)}{\sqrt{2\pi}\sigma} e^{-(x-\mu_B)^2/(2\sigma^2)}, \quad x \in \mathbb{R},$$

and by (8.1) we find the probability default curve

$$\begin{aligned} \mathbb{P}(B | X = x) &= \frac{\mathbb{P}(B)f_X(x | B)}{\mathbb{P}(G)f_X(x | G) + \mathbb{P}(B)f_X(x | B)} \\ &= \frac{\mathbb{P}(B)e^{-(x-\mu_B)^2/(2\sigma^2)}}{\mathbb{P}(G)e^{-(x-\mu_G)^2/(2\sigma^2)} + \mathbb{P}(B)e^{-(x-\mu_B)^2/(2\sigma^2)}} \\ &= \frac{1}{1 + e^{\alpha + \beta x}}, \quad x \in \mathbb{R}, \end{aligned} \quad (8.4)$$

where we let

$$\beta := \frac{\mu_G - \mu_B}{\sigma^2} > 0$$

and

$$\alpha := -\beta \frac{\mu_G + \mu_B}{2} + \log \frac{\mathbb{P}(G)}{\mathbb{P}(B)}.$$




Fig. 8.1: Probability default curve $x \mapsto \mathbb{P}(B | X = x) = 1/(1 + e^{a+x})$.

Similarly, the probability acceptance curve is given by

$$\begin{aligned}
 \mathbb{P}(G | X = x) &= \frac{\mathbb{P}(G)f_X(x | G)}{\mathbb{P}(G)f_X(x | G) + (1 - \mathbb{P}(G))f_X(x | B)} \\
 &= \frac{\mathbb{P}(G)e^{-(x-\mu_G)^2/(2\sigma^2)}}{\mathbb{P}(G)e^{-(x-\mu_G)^2/(2\sigma^2)} + \mathbb{P}(B)e^{-(x-\mu_B)^2/(2\sigma^2)}} \\
 &= \frac{1}{1 + e^{-\alpha - \beta x}} \\
 &= \frac{e^{\alpha + \beta x}}{1 + e^{\alpha + \beta x}}, \quad x \in \mathbb{R}. \tag{8.5}
 \end{aligned}$$




Fig. 8.2: Probability acceptance curve $x \mapsto \mathbb{P}(G | X = x) = e^{a+x}/(1 + e^{a+x})$.

8.2 Decision Rule

We decide to accept the applicants whose score $X(\omega)$ belongs to a decision (or acceptance) set $\mathcal{A} \subset \mathbb{R}$, and to reject those whose score $X(\omega)$ belongs to the rejection set $\mathcal{A}^c = \mathbb{R} \setminus \mathcal{A}$ which is the complement of \mathcal{A} in \mathbb{R} .

Definition 8.3. *Let*

- $D(x)$ represents the cost incurred by the default of an applicant with score $x \in \mathcal{A}$, and
- $L(x)$ represents the loss (or missed earnings) incurred by the rejection of an applicant with score $x \in \mathcal{A}^c$.

The cost associated to this decision rule is defined as

$$\underbrace{D(X)\mathbb{1}_{\{X \in \mathcal{A}\} \cap B}}_{\text{Cost of accepting a "bad" applicant}} + \underbrace{L(X)\mathbb{1}_{\{X \in \mathcal{A}^c\} \cap G}}_{\text{Loss from rejecting a "good" applicant}}$$

Theorem 8.4. *The optimal acceptance set $\mathcal{A}^* \subset \mathbb{R}$ that minimizes the expected cost*

$$\mathbb{E} [D(X)\mathbb{1}_{\{X \in \mathcal{A}\} \cap B} + L(X)\mathbb{1}_{\{X \in \mathcal{A}^c\} \cap G}]$$

is given by

$$\mathcal{A}^* = \left\{ x \in \mathbb{R} : \lambda(x) \geq \frac{D(x) \mathbb{P}(B)}{L(x) \mathbb{P}(G)} \right\},$$

where $\lambda(x)$ is the likelihood ratio

$$\lambda(x) := \frac{f_X(x | G)}{f_X(x | B)} = \frac{\mathbb{P}(G | X = x) \mathbb{P}(B)}{\mathbb{P}(B | X = x) \mathbb{P}(G)}, \quad x \in \mathbb{R}.$$

Proof. The expected cost corresponding to an acceptance set $\mathcal{A} \subset \mathbb{R}$ can be written as

$$\begin{aligned} & \mathbb{E} [D(X)\mathbb{1}_{\{X \in \mathcal{A}\} \cap B} + L(X)\mathbb{1}_{\{X \in \mathcal{A}^c\} \cap G}] \\ &= \int_{\mathcal{A}} D(x) \mathbb{P}(B \cap \{X \in dx\}) + \int_{\mathcal{A}^c} L(x) \mathbb{P}(G \cap \{X \in dx\}) \\ &= \int_{\mathcal{A}} D(x) \mathbb{P}(B | X = x) f_X(x) dx + \int_{\mathcal{A}^c} L(x) \mathbb{P}(G | X = x) f_X(x) dx \\ &= \int_{-\infty}^{\infty} (\mathbb{1}_{\mathcal{A}}(x) D(x) \mathbb{P}(B | X = x) + \mathbb{1}_{\mathcal{A}^c}(x) L(x) \mathbb{P}(G | X = x)) f_X(x) dx. \end{aligned}$$

The expected cost can be minimized pointwise by finding the set \mathcal{A} that minimizes the conditional expected cost

$$x \mapsto \mathbb{E} [D(X)\mathbb{1}_{\{X \in \mathcal{A}\} \cap B} + L(X)\mathbb{1}_{\{X \in \mathcal{A}^c\} \cap G} | X = x]$$

$$= \mathbb{1}_{\mathcal{A}}(x)D(x)\mathbb{P}(B | X = x) + \mathbb{1}_{\mathcal{A}^c}(x)L(x)\mathbb{P}(G | X = x)$$

given that $X = x$. For this, we note that

$$\begin{aligned} & \mathbb{1}_{\mathcal{A}}(x)D(x)\mathbb{P}(B | X = x) + \mathbb{1}_{\mathcal{A}^c}(x)L(x)\mathbb{P}(G | X = x) \\ & \geq \min(D(x)\mathbb{P}(B | X = x), L(x)\mathbb{P}(G | X = x)) \\ & = D(x)\mathbb{P}(B | X = x)\mathbb{1}_{\{D(x)\mathbb{P}(B|X=x) \leq L(x)\mathbb{P}(G|X=x)\}} \\ & \quad + L(x)\mathbb{P}(G | X = x)\mathbb{1}_{\{L(x)\mathbb{P}(G|X=x) < D(x)\mathbb{P}(B|X=x)\}} \\ & = D(x)\mathbb{P}(B | X = x)\mathbb{1}_{\mathcal{A}^*}(x) + L(x)\mathbb{P}(G | X = x)\mathbb{1}_{(\mathcal{A}^*)^c}(x), \end{aligned}$$

where the set \mathcal{A}^* which achieves equality in the above inequality is given by

$$\begin{aligned} \mathcal{A}^* & := \{D(x)\mathbb{P}(B | X = x) \leq L(x)\mathbb{P}(G | X = x)\} \\ & = \left\{ x \in \mathbb{R} : \frac{D(x)}{L(x)} \leq \frac{\mathbb{P}(G | X = x)}{\mathbb{P}(B | X = x)} \right\}. \end{aligned}$$

In addition, the optimal acceptance set \mathcal{A}^* can be rewritten in terms of the *likelihood ratio* function

$$\lambda(x) := \frac{f_X(x | G)}{f_X(x | B)} = \frac{\mathbb{P}(G | X = x) \mathbb{P}(B)}{\mathbb{P}(B | X = x) \mathbb{P}(G)}, \quad x \in \mathbb{R},$$

as

$$\mathcal{A}^* = \left\{ x \in \mathbb{R} : \lambda(x) \geq \frac{D(x) \mathbb{P}(B)}{L(x) \mathbb{P}(G)} \right\}.$$

□

For simplicity, in the sequel we assume that $D = D(x)$ and $L = L(x)$ are constant in $x \in \mathbb{R}$, in which case we have

$$\mathcal{A}^* = \left\{ x \in \mathbb{R} : \lambda(x) \geq \frac{D \mathbb{P}(B)}{L \mathbb{P}(G)} \right\}.$$

Proposition 8.5. *In the Gaussian example (8.2)-(8.3) with $\mu_B < \mu_G$, the optimal acceptance set \mathcal{A}^* is given by*

$$\mathcal{A}^* = \left[\frac{\mu_G + \mu_B}{2} + \frac{1}{\beta} \log \left(\frac{D \mathbb{P}(B)}{L \mathbb{P}(G)} \right), \infty \right), \quad (8.6)$$

where

$$\beta := \frac{\mu_G - \mu_B}{\sigma^2} > 0,$$

under the condition

$$\mathbb{E}[X | B] = \mu_B < \mu_G = \mathbb{E}[X | G].$$

Proof. In the Gaussian example (8.2)-(8.3), the likelihood ratio is given by

$$\begin{aligned} \lambda(x) &= \frac{f_X(x | G)}{f_X(x | B)} \\ &= e^{-(x-\mu_G)^2/(2\sigma^2)+(x-\mu_B)^2/(2\sigma^2)} \\ &= e^{-(\mu_G^2-\mu_B^2-2x(\mu_G-\mu_B))/(2\sigma^2)} \\ &= e^{\beta x - (\mu_G^2-\mu_B^2)/(2\sigma^2)}, \quad x \in \mathbb{R}, \end{aligned}$$

hence the condition

$$\lambda(x) = e^{\beta x - (\mu_G^2-\mu_B^2)/(2\sigma^2)} \geq \frac{D \mathbb{P}(B)}{L \mathbb{P}(G)}$$

is equivalent to

$$\begin{aligned} x &\geq \frac{\mu_G^2 - \mu_B^2}{2\beta\sigma^2} + \frac{1}{\beta} \log\left(\frac{D \mathbb{P}(B)}{L \mathbb{P}(G)}\right) \\ &\geq \frac{\mu_G^2 - \mu_B^2}{2(\mu_G - \mu_B)} + \frac{1}{\beta} \log\left(\frac{D \mathbb{P}(B)}{L \mathbb{P}(G)}\right) \\ &\geq \frac{\mu_G + \mu_B}{2} + \frac{1}{\beta} \log\left(\frac{D \mathbb{P}(B)}{L \mathbb{P}(G)}\right) \\ &=: x^*, \end{aligned}$$

hence $\mathcal{A}^* = [x^*, \infty)$, provided that

$$\beta := \frac{\mu_G - \mu_B}{\sigma^2} > 0,$$

which yields (8.6). □

We note that the optimal boundary point x^* satisfies the relation

$$\lambda(x^*) = \frac{\mathbb{P}(X = x^* | G)}{\mathbb{P}(X = x^* | B)} = \frac{\mathbb{P}(G | X = x^*) \mathbb{P}(B)}{\mathbb{P}(B | X = x^*) \mathbb{P}(G)} = \frac{D \mathbb{P}(B)}{L \mathbb{P}(G)},$$

i.e.

$$\frac{\mathbb{P}(G | X = x^*)}{\mathbb{P}(B | X = x^*)} = \frac{D}{L}. \quad (8.7)$$

Figure 8.3 illustrates the optimal decision rule by taking $L = D = \$1$ and using the default and acceptance curves (8.4)-(8.5).

Fig. 8.3: Animated graph of optimal decision rule.*

In Figure 8.3, the conditional expected cost function

$$x \mapsto \mathbb{1}_{\mathcal{A}}(x)D(x)\mathbb{P}(B \mid X = x) + \mathbb{1}_{\mathcal{A}^c}(x)L(x)\mathbb{P}(G \mid X = x)$$

is represented by the purple curve for a given set \mathcal{A} . Its uniform minimum over different threshold values is obtained for \mathcal{A}^* of the form $\mathcal{A}^* = [x^*, \infty)$ where x^* lies at the intersection of the curves $x \mapsto D(x)\mathbb{P}(B \mid X = x)$ and $x \mapsto L(x)\mathbb{P}(G \mid X = x)$ as in (8.7).

The acceptance rate, or probability that an applicant is accepted according to the rule \mathcal{A} , is given by

$$\begin{aligned} \mathbb{P}(X \in \mathcal{A}) &= \mathbb{P}(\{X \in \mathcal{A}\} \cap G) + \mathbb{P}(\{X \in \mathcal{A}\} \cap B) \\ &= \mathbb{P}(X \in \mathcal{A} \mid G)\mathbb{P}(G) + \mathbb{P}(X \in \mathcal{A} \mid B)\mathbb{P}(B), \end{aligned}$$

where

$$\mathbb{P}(\{X \in \mathcal{A}\} \cap B) = \mathbb{P}(X \in \mathcal{A} \mid B)\mathbb{P}(B)$$

is the default rate, or probability that an applicant accepted according to the rule \mathcal{A} will default.

We can also minimize the default rate $\mathbb{P}(\{X \in \mathcal{A}\} \cap B)$ subject to a given acceptance rate $\mathbb{P}(X \in \mathcal{A}) = a$.

8.3 Logistic Regression

In this section, we address the problem of constructing the random score variable $X : \Omega \rightarrow \mathbb{R}$ in a concrete setting. For this, consider a set of m

* The animation works in Acrobat Reader on the entire pdf file.

financial criteria or indicators $(x_{i,j})_{j=1,2,\dots,m}$ applying to each of n credit applicants $i = 1, 2, \dots, n$.

```
1 install.packages("caret")
2 library(caret); data(GermanCredit); head(GermanCredit)
```

j	Age	ForeignWorker	Property.RealEstate	Housing.Own	CreditHistory	Class
$x_{i,j}$	$x_{i,1} = 67$	$x_{i,2} = 1$	$x_{i,3} = 0$	$x_{i,4} = 1$	$x_{i,5} = 0$	$c_i = 1$

The credit scoring class (“good” or “bad”) of applicant $n^o i$ is denoted by $c_i \in \{0, 1\}$ depending on his status, *i.e.* $c_i = 1$ for “good” applicants and $c_i = 0$ for “bad” applicants.

Linear regression

The score z_i of a given credit applicant in row $n^o i$ is modeled as

$$z_i = F \left(\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j} \right), \quad i = 1, 2, \dots, n.$$

where $(\gamma_j)_{j=0,1,\dots,m}$ is a family of linear coefficients, where $p_i := 1 - z_i$ represents the probability that applicant $n^o i$ may default. In a linear regression model we would take $F(z) := z$, hence the system of equations

$$c_i = \gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}, \quad i = 1, 2, \dots, n,$$

would be used to estimate the coefficients $(\gamma_j)_{j=0,1,\dots,m}$.

Binomial logistic regression

The shortcoming of linear models is that $F(z_i)$, which is assumed to represent a probability value, may exit the interval $[0, 1]$. Logistic regression models address this issue by replacing $F(x) = x$ with the logistic CDF F_L defined as

$$F_L(x) := \frac{e^x}{1 + e^x}, \quad x \in \mathbb{R}.$$




Fig. 8.4: Logistic CDF $x \mapsto F_L(x) = e^x / (1 + e^x)$.

We model the probability $p_i = \mathbb{P}(i \in G)$ that applicant $n^o i$ is rated “good” as

$$p_i = \mathbb{P}(i \in G) = \mathbb{P}\left(G \mid X = \gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}\right) = F_L\left(\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}\right)$$

as in (8.5), $1 \leq i \leq n$. The probability of sampling a given $\{0, 1\}$ -valued sequence $(c_i)_{1 \leq i \leq n}$ of applicant classifications is

$$\prod_{\substack{1 \leq i \leq n \\ c_i=0}} \mathbb{P}(i \in B) \prod_{\substack{1 \leq i \leq n \\ c_i=1}} \mathbb{P}(i \in G) = \prod_{i=1}^n ((\mathbb{P}(i \in B))^{1-c_i} (\mathbb{P}(i \in G))^{c_i}).$$

In order to estimate the sequence of coefficients $(\gamma_j)_{j=0,1,\dots,m}$, we aim at maximizing the log-likelihood ratio

$$\begin{aligned} \log L(\beta|x) &:= \log \prod_{i=1}^n ((\mathbb{P}(i \in B))^{1-c_i} (\mathbb{P}(i \in G))^{c_i}) \\ &= \log \prod_{i=1}^n \left(\left(\bar{F}_L\left(\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}\right) \right)^{c_i} \left(F_L\left(\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}\right) \right)^{1-c_i} \right) \\ &= \log \prod_{i=1}^n \left(\left(\frac{e^{\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}}}{1 + e^{\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}}} \right)^{c_i} \left(\frac{1}{1 + e^{\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}}} \right)^{1-c_i} \right) \\ &= \sum_{i=1}^n c_i \log \frac{e^{\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}}}{1 + e^{\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}}} + \sum_{i=1}^n (1 - c_i) \log \frac{1}{1 + e^{\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}}} \end{aligned}$$

over $\beta = (\gamma_j)_{j=0,1,\dots,m}$. The default probabilities are then given from

$$1 - p_i = \bar{F}_L \left(\gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j} \right), \quad i = 1, \dots, n,$$

and the *logit*

$$\bar{F}_L^{-1}(p_i) = -\log \frac{p_i}{1 - p_i} = \gamma_0 + \sum_{j=1}^m \gamma_j x_{i,j}, \quad i = 1, \dots, n,$$

or log-odds, represents probabilities on a logit scale.

```
1 ggplot(GermanCredit, aes(x = Class)) + geom_bar(aes(y = (.count.)/sum(.count.)), fill=c(
  "red", "darkgreen")) + labs(y = "prob.") + theme_bw()
2 Train <- createDataPartition(GermanCredit$Class, p=0.6, list=FALSE)
   training <- GermanCredit[ Train, ]; testing <- GermanCredit[ -Train, ]
```

The data is randomly split into a training set and a testing set using the `createDataPartition` command in **R**. The training set is used to fit the data in a generalized linear model using the `glm()` command. The testing set is then used to estimate the corresponding default probabilities.

```
1 mod.glm<-glm(Class ~ Age + ForeignWorker + Property.RealEstate + Housing.Own +
  CreditHistory.Critical, data=training, family="binomial")
   head(testing$Class); head(predict(mod.glm, newdata=testing, type="response"))
3 testing <- cbind(rownames(testing), testing); colnames(testing)[1] <- "ID"
   testing$Score5<-predict(mod.glm, newdata=testing, type="response")
5 testsamp <- head(testing,100); colnames(testsamp) <- make.unique(names(testsamp))
   ggplot(testsamp, aes(x=ID, y=Score5, fill=Class)) + geom_bar(stat="identity") +
   scale_fill_manual(values=c("red", "darkgreen")) +
   scale_y_continuous(limits=c(0,1), expand = c(0,0)) + theme_bw(base_size = 18) +
   xlab(NULL) + theme(axis.text.x = element_blank(), aspect.ratio=0.5) +
   geom_hline(yintercept = 0.75, lwd=1.6)
```





Fig. 8.5: Logistic regression output on 5 criteria.

Based on the above 100 samples, the next  code identifies the True Positive Rate (TPR) = $34/77 = 44.16\%$, and the False Positive Rate (FPR) = $8/23 = 34.78\%$, at the threshold $p = 0.75$.

```

1 cat("TPR5=",length(testsamp$Score5[length(testsamp$Score5>0.75 & testsamp$Class=="Good")
  /length(testsamp$Score5[testsamp$Class=="Good"]),'\n')
2 cat("FPR5=",length(testsamp$Score5[length(testsamp$Score5>0.75 & testsamp$Class=="Bad")
  /length(testsamp$Score5[testsamp$Class=="Bad"]),'\n')
pred <- ifelse(testsamp$Score5 > 0.75, "Good", "Bad")
4 confusionMatrix(factor(noquote(pred)),factor(testsamp$Class))

```

```

1 mod.glm<-glm(Class ~ Duration + Amount + InstallmentRatePercentage + ResidenceDuration
  + Age + NumberExistingCredits + NumberPeopleMaintenance + Telephone +
  ForeignWorker + CheckingAccountStatus.lt.0 + CheckingAccountStatus.0.to.200 +
  CheckingAccountStatus.gt.200 + CheckingAccountStatus.none +
  CreditHistory.NoCredit.AllPaid + CreditHistory.ThisBank.AllPaid +
  CreditHistory.PaidDuly + CreditHistory.Delay + CreditHistory.Critical +
  Purpose.NewCar + Purpose.UsedCar + Purpose.Furniture.Equipment +
  Purpose.Radio.Television + Purpose.DomesticAppliance + Purpose.Repairs +
  Purpose.Education + Purpose.Vacation + Purpose.Retaining + Purpose.Business +
  Purpose.Other + SavingsAccountBonds.lt.100 + SavingsAccountBonds.100.to.500 +
  SavingsAccountBonds.500.to.1000 + SavingsAccountBonds.gt.1000 +
  SavingsAccountBonds.Unknown + EmploymentDuration.lt.1 +
  EmploymentDuration.1.to.4 + EmploymentDuration.4.to.7 + EmploymentDuration.gt.7
  + EmploymentDuration.Unemployed + Personal.Male.Divorced.Separated +
  Personal.Female.NotSingle + Personal.Male.Single + Personal.Male.Married.Widowed +
  Personal.Female.Single + OtherDebtorsGuarantors.None +
  OtherDebtorsGuarantors.CoApplicant + OtherDebtorsGuarantors.Guarantor +
  Property.RealEstate + Property.Insurance + Property.CarOther + Property.Unknown +
  OtherInstallmentPlans.Bank + OtherInstallmentPlans.Stores +
  OtherInstallmentPlans.None + Housing.Rent + Housing.Own + Housing.ForFree +
  Job.UnemployedUnskilled + Job.UnskilledResident + Job.SkilledEmployee +
  Job.Management.SelfEmp.HighlyQualified, data=training, family="binomial")
2 testsamp$Score61<-predict(mod.glm, newdata=testing, type="response")
testsamp <- head(testing,100);colnames(testsamp) <- make.unique(names(testsamp))
4 ggplot(testsamp, aes(x=ID, y=Score61, fill=Class)) + geom_bar(stat="identity") +
  scale_fill_manual(values=c("red","darkgreen")) +
  scale_y_continuous(limits=c(0,1),expand = c(0,0)) + theme_bw(base_size = 18) +
  xlab(NULL) + theme(axis.text.x = element_blank(),aspect.ratio=0.5) +
  geom_hline(yintercept = 0.75,lwd=1.6)

```




Fig. 8.6: Logistic regression output on 61 criteria.

In comparison with Figure 8.5, the 100 samples in Figure 8.6 above yield a higher True Positive Rate (TPR) = $48/77 = 62.34\%$ and a lower False Positive Rate (FPR) = $5/23 = 21.74\%$ at the level $p = 0.75$. In other words, the count of true positive samples has increased from 34 to 48, and the count of false positive samples has decreased from 6 to 5.

```

1  cat('TPR61=',length(testsamp$Score61[testsamp$Score61>0.75 &
2  testsamp$Class=="Good"])/length(testsamp$Score61[testamp$Class=="Good"]),'\n')
3  cat('FPR61=',length(testsamp$Score61[testsamp$Score61>0.75 &
4  testsamp$Class=="Bad"])/length(testsamp$Score61[testamp$Class=="Bad"]),'\n')
pred <- ifelse(testsamp$Score61 > 0.75, "Good", "Bad")
confusionMatrix(factor(noquote(pred)),factor(testsamp$Class))

```

8.4 ROC Curve

The ROC curve is a plot of the True Positive Rate values

$$x \mapsto \bar{F}_G(x)$$

against the False Positive Rate function

$$x \mapsto \bar{F}_B(x).$$

Definition 8.6. *The Receiver Operating Characteristic (ROC) curve is the function of the threshold $p \in [0, 1]$ defined as*

$$\begin{aligned}
 [0, 1] &\longrightarrow [0, 1] \\
 p &\longmapsto \text{ROC}(p) := \bar{F}_G(\bar{F}_B^{-1}(p)),
 \end{aligned}$$

where \bar{F}_B^{-1} denotes the inverse of the tail distribution function \bar{F}_B .

The construction of Definition 8.6 proceeds in two steps.

- a) Starting from a given FPR level p , compute the associated threshold x^* as $x^* := \bar{F}_B^{-1}(p)$, i.e.

$$p = \bar{F}_B(x^*) = \mathbb{P}(X > x^* | B).$$

- b) From the threshold level, x^* estimate the corresponding TPR as

$$\mathbb{P}(X > x^* | G) = \bar{F}_G(x^*) = \bar{F}_G(\bar{F}_B^{-1}(p)) = \text{ROC}(p).$$

Proposition 8.7. *The ROC function can be rewritten as the integral*

$$\text{ROC}(p) = \int_0^p \lambda(\bar{F}_B^{-1}(q)) dq, \quad 0 \leq p \leq 1,$$

where the likelihood ratio $\lambda(x)$ is given by

$$\lambda(x) = \frac{f_X(x | G)}{f_X(x | B)}, \quad x \in \mathbb{R}.$$

Proof. The slope of the ROC curve at the point $p \in [0, 1]$ is given by

$$\begin{aligned} \frac{d}{dp} \bar{F}_G(\bar{F}_B^{-1}(p)) &= \bar{F}'_G(\bar{F}_B^{-1}(p)) \frac{d}{dp} \bar{F}_B^{-1}(p) \\ &= \frac{\bar{F}'_G(\bar{F}_B^{-1}(p))}{\bar{F}'_B(\bar{F}_B^{-1}(p))} \\ &= \lambda(\bar{F}_B^{-1}(p)), \end{aligned}$$

hence we have

$$\bar{F}_G(\bar{F}_B^{-1}(p)) = \int_0^p \lambda(\bar{F}_B^{-1}(q)) dq, \quad 0 \leq p \leq 1.$$

□

When X is Gaussian distributed given $\{G, B\}$ with the conditional densities

$$f_X(x | G) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu_G)^2/(2\sigma^2)}$$


and

$$f_X(x | B) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu_B)^2/(2\sigma^2)},$$

the likelihood ration is given by

$$\lambda(x) = \frac{f_X(x | G)}{f_X(x | B)} = e^{\beta x - (\mu_G^2 - \mu_B^2)/(2\sigma^2)}, \quad x \in \mathbb{R},$$

with $\beta := (\mu_G - \mu_B)/\sigma^2$.

The ROC curve in this Gaussian setting can be computed using the following  code.

```

1 FG <- function(x, mug, sigma) {1 - pnorm((x - mug) / sigma)}
2 FBINV <- function(x, mub, sigma) {mub + sigma * qnorm(1 - x)}
3 ROC <- function(x, mub, mug, sigma) {
4   sapply(x, function(x_val) ifelse(x_val == 0, 0, FG(FBINV(x_val, mub, sigma), mug,
      sigma)))}

```

Figure 8.7 presents three samples of ROC curves in the Gaussian example with successively $(\mu_B, \mu_G) = (1, 4)$, $(\mu_B, \mu_G) = (1, 2)$, and $(\mu_B, \mu_G) = (1, 1)$.




Fig. 8.7: Gaussian ROC curves.

We check that the classification is better when $\mu_B \ll \mu_G$. The ROC Curve shows the performance of the classification procedure, which is quantified by the Area Under The Curve (AUC). A perfect classification would correspond to a single point with coordinates $(0, 1)$ and AUC equal to 1, which corresponds to FPR=0% false negatives and TPR=100% true positives. The closer the AUC is to 1, the better the classification is performing.

On the other hand, a completely random guess would correspond to a point on the diagonal line. Points above the diagonal represent good classification results (better than random), while points below the line represent poor results (worse than random) ([Wikipedia](https://en.wikipedia.org/wiki/ROC_curve)).

```

1 TPR5<-length(testing$Score5[testing$Score5>0.75 &
  testing$Class=="Good"])/length(testing$Score5[testing$Class=="Good"])
2 FPR5<-length(testing$Score5[testing$Score5>0.75 &
  testing$Class=="Bad"])/length(testing$Score5[testing$Class=="Bad"])
3 cat("TPR5=",TPR5,"\n"); cat("FPR5=",FPR5,"\n")
4 TPR61<-length(testing$Score61[testing$Score61>0.75 &
  testing$Class=="Good"])/length(testing$Score61[testing$Class=="Good"])
5 FPR61<-length(testing$Score61[testing$Score61>0.75 &
  testing$Class=="Bad"])/length(testing$Score61[testing$Class=="Bad"])
6 cat("TPR61=",TPR61,"\n"); cat("FPR61=",FPR61,"\n")

```

The above True Positive Rates (TPR) and False Positive Rates (FPR) based on 5 and 61 criteria at the level $p = 0.75$ are recomputed in the above `R` code on the whole 400 samples, and plotted on the next ROC graphs of Figure 8.8.

```

1 install.packages("ROCR"); library(ROCR)
2 pred5<-prediction(as.numeric(testing$Score5),as.numeric(testing$Class))
3 perf5 <- performance(pred5,"tpr","fpr")
4 dev.new(width=16,height=7); par(mar=c(4.5,4.5,2,2))
5 plot(perf5,col="purple",lwd=3, xaxs="i", yaxs="i",cex.lab=2,las=1)
6 segments(FPR5,0,FPR5,TPR5, col="purple", lwd =2)
7 segments(0,TPR5,FPR5,TPR5, col="purple", lwd =2)
8 pred61<-prediction(as.numeric(testing$Score61),as.numeric(testing$Class))
9 perf61 <- performance(pred61,"tpr","fpr"); par(new=TRUE)
10 plot(perf61,col="blue",lwd=3,main="", ann=FALSE, xaxs="i", yaxs="i")
11 legend("bottomright", legend=c("61 criteria","5 criteria"),col=c("blue","purple"), lwd=3, cex=3)
12 segments(0,0,1,1, col="black", lwd =3)
13 segments(FPR61,0,FPR61,TPR61, col="blue", lwd =2)
14 segments(0,TPR61,FPR61,TPR61, col="blue", lwd =2)

```

The ROC graphs in the next Figure 8.8 confirm the improvement in classification reached when switching from 5 to 61 criteria.




Fig. 8.8: ROC curves based on 5 criteria and 61 criteria.

Using a neural network in R

Note: Restarting R is *needed* before running the following code.

```

1 install.packages("neuralnet");
2 library(neuralnet);library(caret);data(GermanCredit);library(ROCR);
3 Train <- createDataPartition(GermanCredit$Class, p=0.6, list=FALSE)
4 training <- GermanCredit[ Train, ];testing <- GermanCredit[ -Train, ]
5 nn <- neuralnet(Class ~ Age + ForeignWorker + Property.RealEstate + Housing.Own +
6   CreditHistory.Critical, data=training, hidden=c(3,1), linear.output=FALSE,
7   threshold=0.05)
8 nn$result.matrix; plot(nn,col.intercept = "blue")
9 temp_test <- subset(testing, select = c("Age", "ForeignWorker", "Property.RealEstate",
10   "Housing.Own", "CreditHistory.Critical"))
11 head(temp_test); nn.results <- compute(nn, temp_test)
12 results <- data.frame(actual = testing$Class, prediction = nn.results$net.result)
13 head(results); pred <- ifelse(results[,3] > 0.75, "Good", "Bad")
14 confusionMatrix(factor(results$actual),factor(pred))
15 pred2<-prediction(as.numeric(results[,3]),as.numeric(results$actual))
16 perf <- performance(pred2,"tpr","fpr"); plot(perf,col="purple",lwd=3, xaxs = "i", yaxs = "i")

```

Improved performance may be achieved by rescaling the binary variables to $\{-1, 1\}$.

```

1 training$ForeignWorker <- ifelse(training$ForeignWorker == 0, -1, 1)
2 training$Property.RealEstate <- ifelse(training$Property.RealEstate == 0, -1, 1)
3 training$Housing.Own <- ifelse(training$Housing.Own == 0, -1, 1)
4 training$CreditHistory.Critical <- ifelse(training$CreditHistory.Critical == 0, -1, 1)
5 nn2 <- neuralnet(Class ~ Age + ForeignWorker + Property.RealEstate + Housing.Own +
6   CreditHistory.Critical, data=training, hidden=c(3,1), linear.output=FALSE,
7   threshold=0.05)
8 nn2.results <- compute(nn2, temp_test)
9 results2 <- data.frame(actual = testing$Class, prediction = nn2.results$net.result)
10 pred3 <- ifelse(results2[,3] > 0.75, "Good", "Bad")
11 confusionMatrix(factor(results2$actual),factor(pred3))
12 pred4<-prediction(as.numeric(results2[,3]),as.numeric(results$actual))
13 dev.new(width=16,height=7); par(mar = c(4.5,4.5,2,2))
14 plot(perf,col="purple",lwd=3, xaxs = "i", yaxs = "i",cex.lab=2,las=1)
15 perf3 <- performance(pred4,"tpr","fpr"); par(new=TRUE)
16 plot(perf3,col="blue",lwd=3,main="", ann=FALSE, xaxs="i", yaxs="i")
17 legend("bottomright", legend=c("Rescaled","Non rescaled"),col=c("blue","purple"), lwd=3,
18   cex=3)

```


Using a neural network in Python

Download the corresponding [IPython notebook](#) that can be run [here](#) using this [data file](#).

Using random forests in Python

Download the corresponding [IPython notebook](#) that can be run [here](#) using this [data file](#).

Using XGBoost in 

Download the corresponding  [code](#).

Exercises

Exercise 8.1 Consider a set Ω of applicants decomposed into the partition $\Omega = G \cup B$, where each applicant ω is assigned a score $X(\omega)$ which is exponentially distributed given $\{G, B\}$ with the conditional densities

$$f_X(x | G) = \lambda_G e^{-\lambda_G x} \mathbb{1}_{[0, \infty)}(x)$$

and

$$f_X(x | B) = \lambda_B e^{-\lambda_B x} \mathbb{1}_{[0, \infty)}(x),$$

where $\lambda_B > \lambda_G > 0$.

- Compute the conditional expected values $\mathbb{E}[X | G]$ and $\mathbb{E}[X | B]$.
- Compute the *probability default curve*

$$x \mapsto \mathbb{P}(B | X = x) = \frac{\mathbb{P}(B) f_X(x | B)}{\mathbb{P}(G) f_X(x | G) + \mathbb{P}(B) f_X(x | B)}$$

in terms of the likelihood ratio $\lambda(x) := f_G(x) / f_B(x)$, $x \in \mathbb{R}$.

- Determine the acceptance set

$$\mathcal{A} := \{x \in \mathbb{R} : DP(B | X = x) \leq LP(G | X = x)\},$$

where

- L represents the loss incurred by the rejection of an applicant, and
- D represents the loss incurred by the default of an applicant.

Exercise 8.2 Consider a set Ω of applicants decomposed into the partition $\Omega = G \cup B$, where each applicant ω is assigned a uniformly distributed score $X(\omega)$ given $\{G, B\}$, with the conditional densities

$$f_X(x | G) = \frac{1}{\lambda_G} \mathbb{1}_{[0, \lambda_G]}(x) dx$$

and

$$f_X(x | B) = \frac{1}{\lambda_B} \mathbb{1}_{[0, \lambda_B]}(x) dx,$$

where $0 < \lambda_B < \lambda_G$.

a) Compute the *probability default curve*

$$x \mapsto \mathbb{P}(B \mid X = x) = \frac{\mathbb{P}(B)f_X(x \mid B)}{\mathbb{P}(G)f_X(x \mid G) + \mathbb{P}(B)f_X(x \mid B)}.$$

b) Compute the conditional expected values $\mathbb{E}[X \mid G]$ and $\mathbb{E}[X \mid B]$.

c) Determine the acceptance set

$$A := \{x \in \mathbb{R}_+ : D\mathbb{P}(B \mid X = x) \leq L\mathbb{P}(G \mid X = x)\},$$

where

- L represents the missed earnings incurred by the rejection of applicant,
- D represents the loss incurred by the default of an applicant.

Exercise 8.3

a) Compute the ROC function $x \mapsto \overline{F}_G(\overline{F}_B^{-1}(x))$ of Exercise 8.1, $x \in [0, 1]$, and draw its graph for sample values of λ_G, λ_B .

b) Compute the ROC function $x \mapsto \overline{F}_G(\overline{F}_B^{-1}(x))$ of Exercise 8.2, $x \in [0, 1]$, and draw its graph for sample values of λ_G, λ_B .

Exercise 8.4 Consider a set Ω of customers decomposed as the partition $\Omega = G \cup B$, where each customer ω is assigned a uniformly distributed score $X(\omega)$ given $\{G, B\}$, with the conditional densities

$$f_X(x \mid G) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu_G)^2/(2\sigma^2)}$$

and

$$f_X(x \mid B) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu_B)^2/(2\sigma^2)},$$

with $\mu_B < \mu_G$.

a) Compute the *probability default curve*

$$x \mapsto \mathbb{P}(B \mid X = x) := \frac{\mathbb{P}(B)f_X(x \mid B)}{\mathbb{P}(G)f_X(x \mid G) + \mathbb{P}(B)f_X(x \mid B)}.$$

b) Compute the likelihood ratio $\lambda(x)$ defined as $\lambda(x) := f_X(x \mid G)/f_X(x \mid B)$, $x \in \mathbb{R}$.

c) Letting

- $L(x) := e^{-ax}$ represent the missed earnings incurred by rejecting an applicant with score $x \in \mathbb{R}$, with $a > 0$,
- $D(x) := e^{bx}$ represent the loss incurred by the default of an applicant with score $x \in \mathbb{R}$, with $b > 0$,

determine the acceptance set

$$\mathcal{A} := \{x \in \mathbb{R}_+ : D(x)\mathbb{P}(B | X = x) \leq L(x)\mathbb{P}(G | X = x)\}.$$