



AGH

Wydział Applied Mathematics

Project for the course Data Science

The Selection of a Model for German Credit Applications

Philippe De Brouwer

05 October, 2021

Abstract

We use the data of German Credit applications to build a credit model. We use two approaches: neural networks and logistic regression and conclude that the neural network is not preferable, because it over-fits too much and is black box. For the logistic regression we also build a challenger model that has one variable more and conclude that this model is the best so far.

Contents

1	Introduction	3
2	The Data	3
2.1	Introduction	3
2.2	The Data Quality	4
3	Data Wrangling	4
3.1	Categorical Variables	4
3.1.1	The information Value	4
3.2	The Continous Variables	5
3.2.1	Decide which Continuous Variable to Use	5
3.3	Data Binning	7
3.3.1	The Categorical Variables	7
3.3.2	The Continuous variables	9
4	The Logistic Regression	9
5	The performance of the Model	11
6	Validation of the Model	13
6.1	Monte Carlo Cross Validation	13
7	The Challenger Models	13
7.1	Neural Network	13
7.2	Another logistic regression: logistic 2	14
8	Conclusion	17
	Bibliography	19

1 Introduction

This is a template that can be used for the report of the project that goes with the course *Data Science* of **Philippe De Brouwer** for AGH. It is obviously not perfect, nor really finished. We invite you to use this as a starting point, and rethink all decisions yourself. Eventually we hope that you can build a better model.

Best is to take the models that is selected here (“logistic 2’’) as a reference and take it as a challenge to build a better model. The quality of the model is both the performance and the robustness (“not over-fit”). Ideally it is better in at least one of those two areas while at least as good in the other criterion.

You are free to take another view anything. In particular we suggest you to consider:

1. the selection of variables,
2. interaction variables (matrix variables),
3. the choice of model (logistic, probit, decision tree, neural network, SVM, etc.), or eventually improve the models chosen (e.g. how could you make the neural network to be less over-fitted?),
4. of course you can also choose to consider other performance criteria and even use other cross validation methods, such as *k-fold cross validation*, and
5. the documentation, are the better ways of presenting it, what to present, etc.

Good luck!

The data is believed to be public domain.

The main reference is De Brouwer (2020), and more in particular (De Brouwer 2020, pt. 1 to 6 with elements of part 8), the slides presented in the classroom, and the code that we walked through together.

2 The Data

2.1 Introduction

The data consists of 1000 rows and 22 columns. In Table 1 we list all column names.

Table 1: The column names of the data.

Column Names
default
account_check_status
duration_in_month
credit_history
purpose
credit_amount
savings
present_emp_since
installment_as_income_perc
personal_status_sex
other_debtors
present_res_since
property
age

Column Names
other_installment_plans
housing
credits_this_bank
job
people_under_maintenance
telephone
foreign_worker
isGood

The breakdown of the dependent variable “default” is as presented in Table 2.

Table 2: The Data consists of a high number of defaults. We assume that this is because the data has been prepared by dropping randomly “good” customers.

Var1	Freq
0	700
1	300

2.2 The Data Quality

From the 1000 observations, there are 0 rows with missing values. Some of the values have one of the categorical variables that indicates the lack of information. We refer to the following sections where we will show how these “NA” values will be treated.

No suspicious outliers were found.

3 Data Wrangling

3.1 Categorical Variables

3.1.1 The information Value

The first step is selecting the variables that we ultimately want to use in our model. First we will consider the categorical variables and study the information value.

Table 3: The table of all information values for each categorical variable ordered in decreasing order. We will work with the ones that have an information value above 0.1.

varName	IV
account_check_status	0
credit_history	0
purpose	0
savings	0
present_emp_since	0

varName	IV
personal_status_sex	0
other_debtors	0
property	0
other_installment_plans	0
housing	0
job	0
telephone	0
foreign_worker	0

3.2 The Continuous Variables

The continuous variables are presented in Table 4. Also their content follows from the naming convention.

Table 4: List of the continuous scale variables in the data.

Continuous Scale Variables
duration_in_month
credit_amount
installment_as_income_perc
present_res_since
age
credits_this_bank
people_under_maintenance

3.2.1 Decide which Continuous Variable to Use

Note that we have introduced a variable *isGood* which is one for good credits and zero for bad credits.

In order to select continuous scale variables we will investigate the correlation and study the plot of the average behaviour in function of the relevant variables in order to detect non-linear dependencies. Later (in next section) we will consider interactions between the different variables (relevant for the logistic regression for example).

Based on the correlation we can certainly consider *duration_in_month*, *credit_amount*, *installment_as_income_perc* (which is usually a strong candidate, however in this data the dependency is not remarkable: -0.0724039), and *age*.

We also note the high correlation between *installment_as_income_perc*, *duration_in_month* and *credit_amount* and make note that we might need to combined them in order to avoid co-linearity if we would decide to keep both variables in a regression model.

We used a loess model and histogram to assess what binning is optimal for these variables. Here we present the loess with superimposed confidence level for the variables where the correlation with *isGood* is above 0.05%.

We already knew that the variable *installment_as_income_perc* had a reasonably low correlation with the dependent variable. From Figure 2 we learn also that the shape of the relationships is rather flat. Therefore we decide in a first approach not to use this variable (the models “logistic 1”

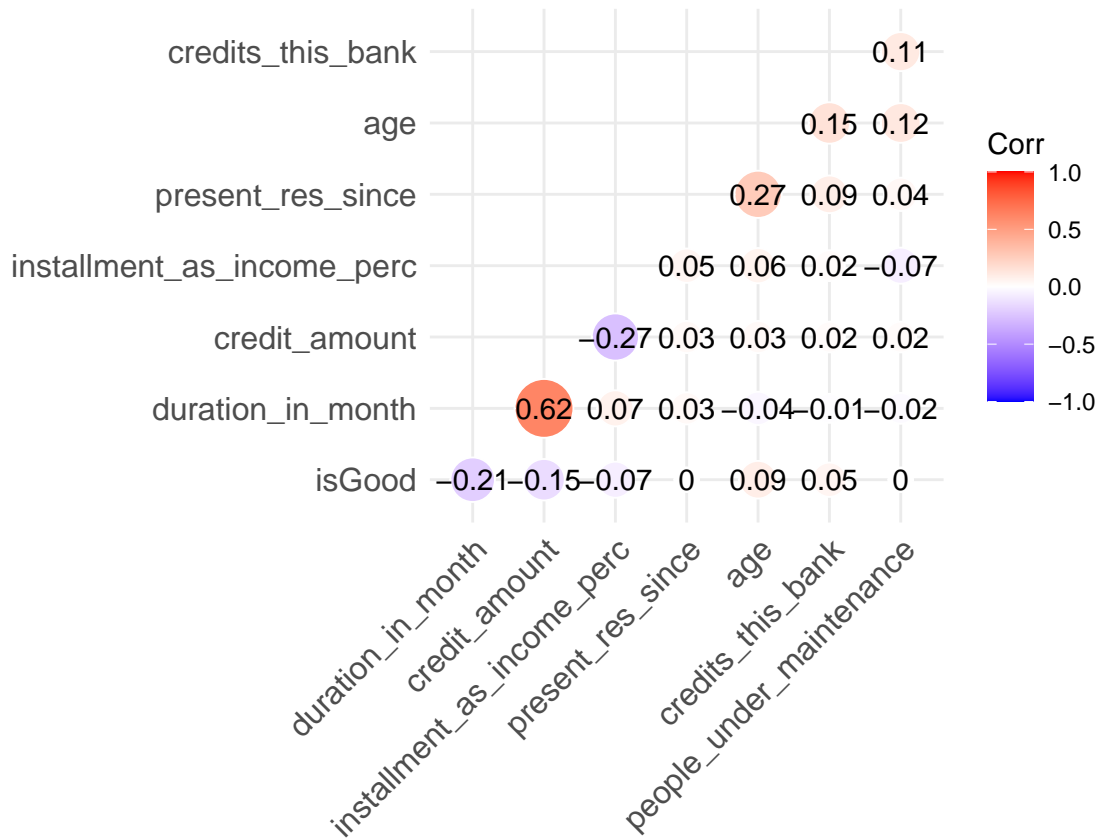


Figure 1: The correlation matrix for all continuous scale variables. Note that *isGood* is the binary variable that equals 1 for all successful credits and 0 for credits that defaulted.

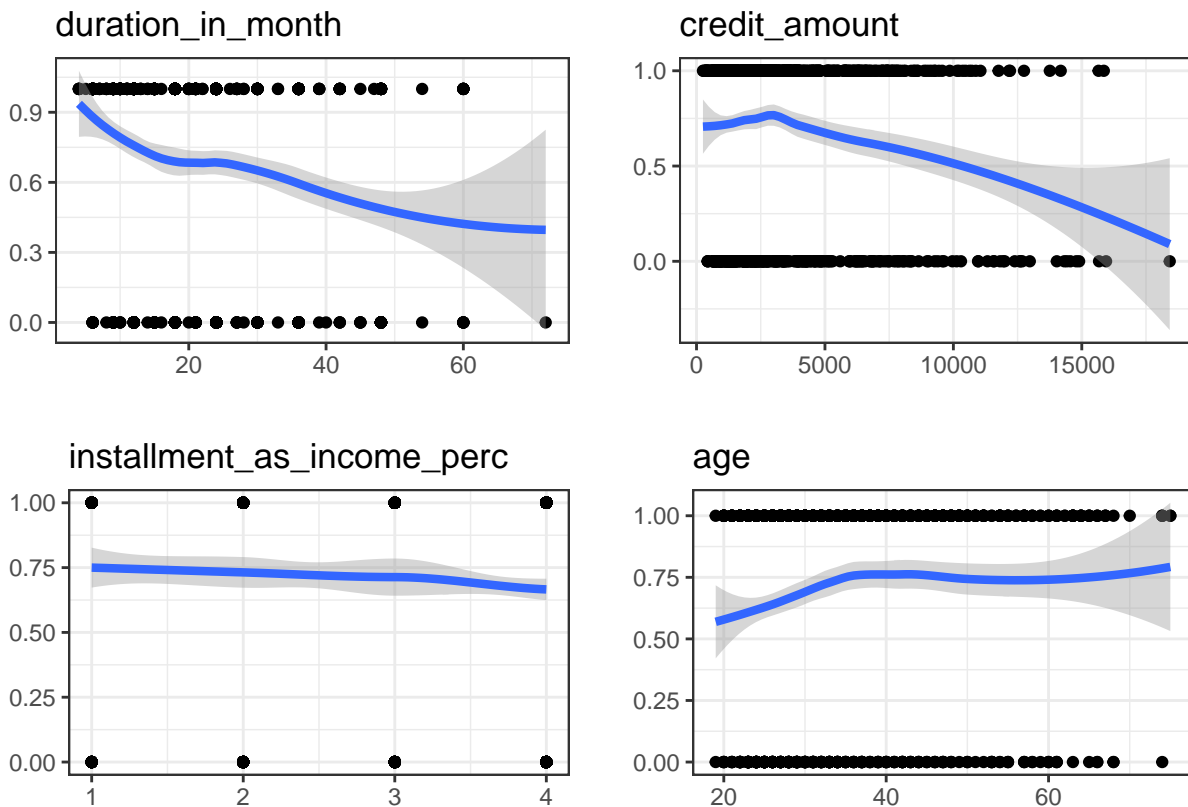


Figure 2: The loss estimations and its confidence level for the variables with correlation above.

and “neuralnet” will be built without this parameter) and build a challenger models that includes this variable (“logistic 2”).

3.3 Data Binning

3.3.1 The Categorical Variables

The existing binning is probably directly taken from the production system. This means that we should at least check if it is optimal for a logistic regression and its cross validation. Ideally we want to avoid bins that are too small and we also want to make sure that the relationships observed are realistic, logical and consistent.

We decided to make sure all bins have at least 200 observations in them and will take bins that have similar WOE together.

3.3.1.1 Account_check_status To illustrate the thought process, we display the information for *account_check_status*. First we study the WOE of the existing bins (see Figure 3) and their detailed information about histogram, weight of evidence and information value (see Table 5).




Figure 3: The WOE for the variable *account_check_status*.

Table 5: The WOE and IV for the variable *account_check_status*.

<i>account_check_status</i>	n	nbrBad	nbrGood	pctGood	WOE	IV
< 0 DM	274	135	139	0.51	-0.82	0.21
0 <= ... < 200 DM	269	105	164	0.61	-0.40	0.05
>= 200 DM / salary assignments for	63	14	49	0.78	0.41	0.01

account_check_status	n	nbrBad	nbrGood	pctGood	WOE	IV
no checking account	394	46	348	0.88	1.18	0.40

Then we decide on which bins to combine. In this case, we decide to take the second and third bin together. This is done by the following code:

```
df$acc_chk_sts <- if_else(d$account_check_status == '< 0 DM', 'neg',
  if_else((d$account_check_status == '0 <= ... < 200 DM') |
    (d$account_check_status == '>= 200 DM / salary assignments for at least 1 year'),
    'pos',
    if_else(d$account_check_status == 'no checking account', 'none', 'ERROR'))
# Check for mistakes (they should appear as ERROR or an undefined category):
df$acc_chk_sts %>% table
```

```
## .
## neg none pos
## 274 394 332
```

3.3.1.2 Other variables We work with similar logic through all other variables and make the following transformations:

```
### credit_history
df$credit_hist <- if_else((d$credit_history == 'no credits taken/ all credits paid back duly' |
  (d$credit_history == 'all credits at this bank paid back duly' |
    'clear',
    if_else((d$credit_history == 'existing credits paid back duly till now') |
      (d$credit_history == 'delay in paying off in the past'), 'paid',
      if_else(d$credit_history == 'critical account/ other credits existing (no)') |
        'crit', 'ERROR'))))

### savings
df$savings <- if_else((d$savings == '... < 100 DM') | (d$savings == '100 <= ... < 500 DM') |
  if_else((d$savings == 'unknown/ no savings account') | (d$savings == '500 <= ... < 1000 DM'),
    'small', 'ERROR') |
  if_else(d$savings == '.. >= 1000 DM ', 'large', 'ERROR'))

### Property
df$property <- if_else(d$property == 'unknown / no property', 'none',
  if_else((d$property == 'if not A121/A122 : car or other, not in attribute 6' |
    (d$property == 'if not A121 : building society savings agreement/ life insurance' |
      'movbl',
      if_else(d$property == 'real estate', 'real', 'ERROR'))))

### purpose
df$purpose <- if_else((d$purpose == '(vacation - does not exist?)' |
  (d$purpose == 'furniture/equipment') |
  (d$purpose == 'car (new)') |
  (d$purpose == 'car (used)'), 'friv',
  if_else((d$purpose == 'education') |
    (d$purpose == 'other'), 'other', 'ERROR'))
```



```

        (d$purpose == 'business') |
        (d$purpose == 'repairs') |
        (d$purpose == 'radio/television')
        , 'invest',
    if_else((d$purpose == 'domestic appliances') |
            (d$purpose == 'car (used)') |
            (d$purpose == 'retraining')
            , 'useful', 'ERROR'))

```

Finally we check the information value of the new binning:

variable	InformationValue	Predictability
acc_chk_sts	0.6345968	Highly Predictive
credit_hist	0.2918291	Highly Predictive
savings	0.1908894	Highly Predictive
property	0.1126339	Highly Predictive
purpose	0.1524439	Highly Predictive

3.3.2 The Continuous variables

Based on a similar reasoning¹, but also making sure that we capture any non-linear behaviour, we split the continuous variables as follows:

```

#### duration_in_month
df$duration <- if_else(d$duration_in_month <= 12, 'L',
                      if_else(d$duration_in_month <= 24 & d$duration_in_month > 12, 'M',
                              if_else(d$duration_in_month > 24, 'H', 'ERROR')))

### credit_amount
df$credit_amount <- if_else(d$credit_amount <= 2500, 'L',
                           if_else(d$credit_amount <= 5000 & d$credit_amount > 2500, 'M',
                                   if_else(d$credit_amount > 5000, 'H', 'ERROR')))

### age
df$age <- if_else(d$age <= 30, 'L',
                 if_else(d$age > 30, 'H', 'ERROR'))

```

4 The Logistic Regression

The first model that we will study is a logistic regression. We name it for future reference “logistic 1.”

```

# Define the formula:
frm <- isGood ~ acc_chk_sts + credit_hist + savings + property +
          purpose + duration + credit_amount + age +
          savings * age + purpose * credit_amount

```

¹We want to end up with bins that are not too small and we want behaviour in that one bin to be as similar as possible.

```

# Fit the model:
m <- glm(formula = frm, data = df, family = "binomial")

# Investigate the model:
summary(m)

##
## Call:
## glm(formula = frm, family = "binomial", data = df)
##
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -2.8866  -0.8189   0.4397   0.7532   1.8493
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)    -1.11864    0.59072  -1.894  0.05826 .
## acc_chk_stsnone  1.61515    0.21375   7.556 4.15e-14 ***
## acc_chk_stspos   0.49996    0.18391   2.719  0.00656 **
## credit_histcrit  1.59628    0.30375   5.255 1.48e-07 ***
## credit_histpaid  0.84587    0.26744   3.163  0.00156 **
## savingsother    0.19850    0.49063   0.405  0.68578
## savingsssmall  -0.66076    0.41477  -1.593  0.11115
## propertynone   -0.42292    0.23075  -1.833  0.06683 .
## propertyreal    0.27113    0.19690   1.377  0.16851
## purposeinvest   0.17096    0.44221   0.387  0.69904
## purposeuseful   0.59843    0.40826   1.466  0.14270
## durationL      1.20145    0.26489   4.536 5.74e-06 ***
## durationM      0.61142    0.22419   2.727  0.00639 **
## credit_amountL -0.34256    0.37784  -0.907  0.36461
## credit_amountM  0.03419    0.44872   0.076  0.93926
## ageL          -0.48965    0.57961  -0.845  0.39823
## savingsother:ageL  0.06200    0.72922   0.085  0.93224
## savingsssmall:ageL  0.06987    0.60705   0.115  0.90837
## purposeinvest:credit_amountL  0.38957    0.51715   0.753  0.45127
## purposeuseful:credit_amountL  0.40192    0.48748   0.824  0.40967
## purposeinvest:credit_amountM  0.51364    0.59006   0.870  0.38403
## purposeuseful:credit_amountM  0.63575    0.59408   1.070  0.28455
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 1221.73  on 999  degrees of freedom
## Residual deviance:  969.86  on 978  degrees of freedom
## AIC: 1013.9
##
## Number of Fisher Scoring iterations: 5

```

We notice that none of the interactions is statistically significant and decide to leave them out in

order to make the model more robust. Then we fit the model again in the following code:

```
frm <- isGood ~ acc_chk_sts + credit_hist + savings + property +
  purpose + duration + credit_amount + age
m <- glm(formula = frm, data = df, family = "binomial")
summary(m)
```

```
##
## Call:
## glm(formula = frm, family = "binomial", data = df)
##
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -2.8724  -0.7963   0.4363   0.7562   1.8765
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)  -1.38987    0.46805  -2.969  0.00298 **
## acc_chk_stsnone  1.62853    0.21266   7.658 1.89e-14 ***
## acc_chk_stspos   0.51819    0.18281   2.835  0.00459 **
## credit_histcrit  1.58537    0.29885   5.305 1.13e-07 ***
## credit_histpaid  0.84956    0.26297   3.231  0.00124 **
## savingsother    0.21178    0.36509   0.580  0.56185
## savingsssmall  -0.63239    0.30271  -2.089  0.03670 *
## propertynone   -0.39933    0.22927  -1.742  0.08155 .
## propertyreal    0.27157    0.19643   1.383  0.16681
## purposeinvest   0.50159    0.20073   2.499  0.01246 *
## purposeuseful   0.95534    0.19943   4.790 1.66e-06 ***
## durationL      1.18845    0.26254   4.527 5.99e-06 ***
## durationM      0.60622    0.22223   2.728  0.00638 **
## credit_amountL -0.07055    0.25710  -0.274  0.78378
## credit_amountM  0.43522    0.24878   1.749  0.08022 .
## ageL           -0.41904    0.16488  -2.541  0.01104 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 1221.73  on 999  degrees of freedom
## Residual deviance:  971.25  on 984  degrees of freedom
## AIC: 1003.3
##
## Number of Fisher Scoring iterations: 5
```

5 The performance of the Model

We will check the performance of the model with the aid of the ROC curve.

The model has an AUC of 0.7973381, and a KS of 0.4604762. These findings are not bad and we will proceed now to calculate the optimal cutoff. This cutoff is as follows:

KS Plot




Figure 4: The ROC (receiver operating curve) for our model

ROC Curve




Figure 5: The lift of the model (bottom): the cumulative percentage of responders (ones) captured by the model

```
## The optimal cutoff if the false positives cost 10 times more than the false negatives:
##           [,1]
## sensitivity 0.3471429
## specificity 0.9566667
## cutoff      0.8926366
```

We can also use the package *InformationValue* to calculate the optimal cutoff:

```
optimalCutoff(actuals = df$isGood, predictedScores = predScores, optimiseFor = "Zeros")
## [1] 0.9925079
optimalCutoff(actuals = df$isGood, predictedScores = predScores, optimiseFor = "Both")
## [1] 0.6825079
optimalCutoff(actuals = df$isGood, predictedScores = predScores, optimiseFor = "Ones")
## [1] 0.1625079
optimalCutoff(actuals = df$isGood, predictedScores = predScores, optimiseFor = "misclassification")
## [1] 0.5525079
```

6 Validation of the Model

To validate the model we will use the cross validation method, and opt for the Monte Carlo Cross Validation. [^While we could opt for the *k-fold* validation (with *k* not too high), we choose the Monte Carlo methods because it allows to draw more random selections.]

6.1 Monte Carlo Cross Validation

We use the Monte Carlo Cross Validation with a test dataset that spans 30% of our observations and 70% in the training data-set. We will draw 200 times a training data-set of 0.7 and study the AUC on the testing data-set.

The results are in Figure 6. We notice that the model is over-fit, but even for the test data-set the performance is acceptable. All variables retained have at least one of the categories with a *p-value* that is smaller than 0.01. We could consider to leave out the categories that have a higher *p-value* to make the model more robust. We choose, however, not to do that because all coefficients are somehow logical.

The median of the observed values for the AUC of the test data is 0.7751938, the average is 0.7745367 with a standard deviation of 0.0243891.

7 The Challenger Models

We will build two challenger models.

7.1 Neural Network

As a challenger model, we use a neural network on the same binned data and with three hidden layers with resp 16, 8, and 4 hidden neurons.




Figure 6: The histogram for the AUC of the randomised data-sets with the Monte Carlo cross validation for the first logistic regression model.

First we fit the model on all the data:

```
## [1] "AUC: 0.831480952380953"
```

```
## [1] "KS: 0.56952380952381"
```

Now we perform the Monte Carlo Cross Validation for the neural network. The results are in Figure 8.

7.2 Another logistic regression: logistic 2

We use also the installments as percentage of income, the correlation was small but it is usual a good variable to use for consumer lending. We also decide the parameter as it is. The variable is a number between 1 and 4.

```
summary(d$installment_as_income_perc)
```

```
##   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##  1.000  2.000  3.000  2.973  4.000  4.000
```

The variable is encoded as a number between 1 and 4. We believe this is not too different from the other binary variables, and will fit the model with the variable unchanged.

```
d2 <- cbind(df, installment_as_income_perc = d$installment_as_income_perc)
```

```
frm2 <- isGood ~ acc_chk_sts + credit_hist + savings + property +
  purpose + duration + credit_amount + age + installment_as_income_perc
```




Figure 7: The ROC for the neural network with three hidden layers of 14, and 7 neurons respectively.




Figure 8: The results of the cross validation for the neural network. We notice that the neural network with only three layers does not significantly overfit and hence is a valid model.

```
m2 <- glm(formula = frm2, data = d2, family = "binomial")
summary(m2)
```

```
##
## Call:
## glm(formula = frm2, family = "binomial", data = d2)
##
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -2.8433  -0.7887   0.4450   0.7453   1.9900
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)    -0.87463    0.50690  -1.725 0.084449 .
## acc_chk_stsnone  1.62590    0.21374   7.607 2.80e-14 ***
## acc_chk_stspos   0.48315    0.18404   2.625 0.008658 **
## credit_histcrit  1.60182    0.29956   5.347 8.93e-08 ***
## credit_histpaid  0.84192    0.26318   3.199 0.001379 **
## savingsother    0.25345    0.36741   0.690 0.490298
## savingssmall   -0.62136    0.30494  -2.038 0.041582 *
## propertynone   -0.33794    0.23097  -1.463 0.143436
## propertyreal    0.24173    0.19774   1.222 0.221536
## purposeinvest   0.50731    0.20147   2.518 0.011803 *
## purposeuseful   1.00217    0.20118   4.981 6.31e-07 ***
## durationL       1.03128    0.26975   3.823 0.000132 ***
## durationM       0.51989    0.22447   2.316 0.020555 *
## credit_amountL  0.22399    0.28098   0.797 0.425344
## credit_amountM  0.55479    0.25345   2.189 0.028599 *
## ageL            -0.42528    0.16543  -2.571 0.010145 *
## installment_as_income_perc -0.20986    0.08029  -2.614 0.008950 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 1221.73  on 999  degrees of freedom
## Residual deviance:  964.29  on 983  degrees of freedom
## AIC: 998.29
##
## Number of Fisher Scoring iterations: 5
## [1] "AUC: 0.800790476190476"
## [1] "KS: 0.463333333333333"
```

The coefficient of *installment_as_income_perc* is significant (the p-value is 0.0089499). This is encouraging.

We will now subject this model to the same Monte Carlo cross validation and compare the results.

```
cv_mc2 <- crossv_mc(d2, n = nRuns, test = 1 - pctTrain)
mods2 <- map(cv_mc2$train, ~ glm(frm2, data = .))
```





Figure 9: The ROC for the logistic challenger logistic regression

```

AUCs2_test <- numeric(0)
AUCs2_train <- numeric(0)
for(k in 1:nRuns) {
  AUCs2_test[k] <- fAUC(mods2[[k]], as.data.frame(cv_mc2$test[[k]]))
  AUCs2_train[k] <- fAUC(mods2[[k]], as.data.frame(cv_mc2$train[[k]]))
}

allAUCs <- rbind(tibble(model = "train data", AUC = AUCs2_train),
                tibble(model = "test data", AUC = AUCs2_test))
p1 <- ggplot(allAUCs, aes(AUC, fill = model, colour = model)) +
  geom_density(alpha=0.5)
p2 <- ggplot(allAUCs, aes(AUC, fill = model, colour = model)) +
  stat_ecdf()
grid.arrange(p1, p2, ncol = 1)

```

The median of the observed values for the AUC of the test data is 0.7811891, the average is 0.7787291 with a standard deviation of 0.0254239.

8 Conclusion

Comparing the two logistic regressions it seems that the challenger model that uses the additional variable (debt installments as a percentage of income), but also prone to over-fitting. The neural network performs significantly better on the training data, but is even more over-fit: it performs worse on the testing data. For that reason we have to reject this particular neural network.

In the following table we summarise these results:




Figure 10: The distribution of the AUC for the challenger model “logistic 2.”




Figure 11: The kernel density for the observed areas under the curve (top) and the cumulative probability density functions (bottom) for the base and challenger models. All AUCs shown are for the test data only.

Model	Quality	Mean AUC on Test Data	AUC on All Data
logistic 1	slightly over-fit	0.7745367	0.7973381
logistic 2	slightly over-fit	0.7787291	0.8007905
neural network	significantly over-fit	0.7235076	0.831481

The decisions of the logistic regression are transparent in that sense that they can be explained to the customer. The explainability is important for the bank (who seeks assurance that he or she does the right thing, the customer and eventual courts). The neural networks seems to over-fit more – so we would suggest to look into simpler neural networks with less neurons and/or layers.

The overall best models that we have is “logistic 2” and we recommend this for implementation.

Bibliography

De Brouwer, Philippe J.S. 2020. *The Big r-Book: From Data Science to Learning Machines and Big Data*. New York: John Wiley & Sons, Ltd. [https://doi.org/https://doi.org/10.1002/9781119632757](https://doi.org/10.1002/9781119632757).