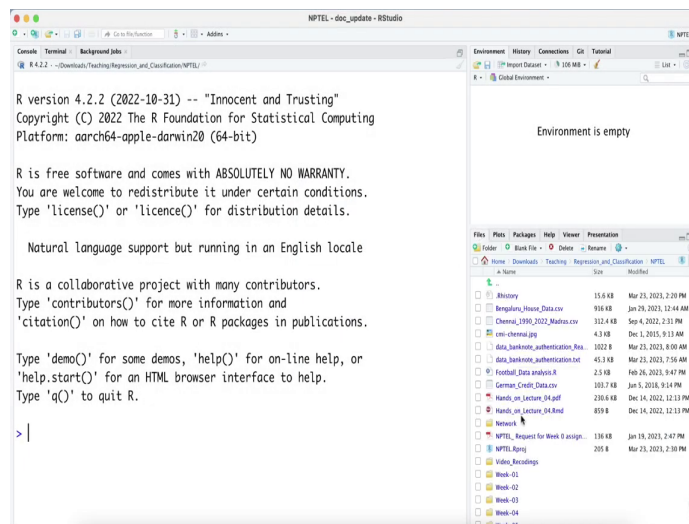


Predictive Analytics - Regression and Classification
Prof. Sourish Das
Department of Mathematics
Chennai Mathematical Institute

Lecture - 58
Hands on with R: Classify fake bank note with GLM

Hello all. Welcome to the part B of lecture 19. In this part we will do Hands on with R in this hands on video. We are going to do work on another data analysis where let me just first open the R.

(Refer Slide Time: 00:44)



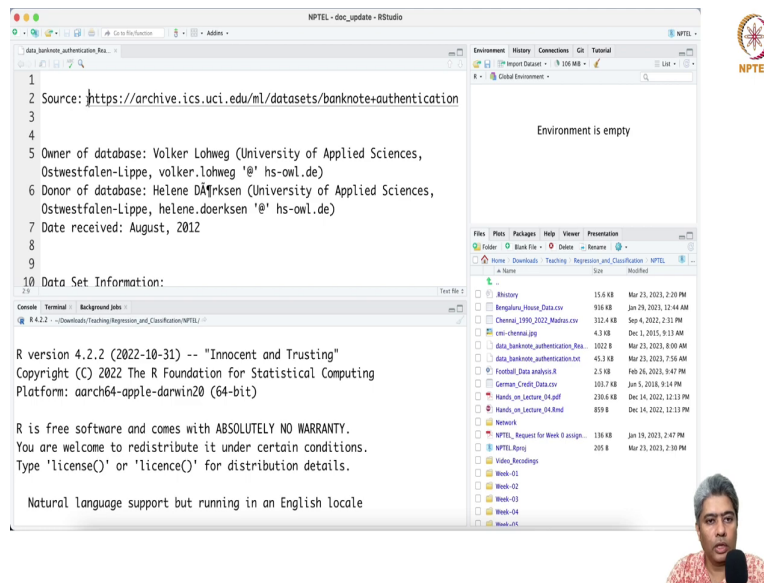
The screenshot shows the RStudio interface. The console on the left displays the following text:

```
R version 4.2.2 (2022-10-31) -- "Innocent and Trusting"  
Copyright (C) 2022 The R Foundation for Statistical Computing  
Platform: aarch64-apple-darwin20 (64-bit)  
  
R is free software and comes with ABSOLUTELY NO WARRANTY.  
You are welcome to redistribute it under certain conditions.  
Type 'license()' or 'licence()' for distribution details.  
  
Natural language support but running in an English locale  
  
R is a collaborative project with many contributors.  
Type 'contributors()' for more information and  
'citation()' on how to cite R or R packages in publications.  
  
Type 'demo()' for some demos, 'help()' for on-line help, or  
'help.start()' for an HTML browser interface to help.  
Type 'q()' to quit R.  
  
> |
```

The file explorer on the right shows a list of files and folders, including 'data_banknote_authentication.Rna', 'data_banknote_authentication.txt', 'Football_Data_analysis.R', 'German_Credit_Data.csv', 'Hands_on_Lecture_04.pdf', 'Hands_on_Lecture_04.Rmd', 'Network', 'NPTEL_Request for Week 0 assign.', 'NPTEL.Rproj', 'Video_Recordings', 'Week-01', 'Week-02', 'Week-03', 'Week-04', and 'Week-05'.



(Refer Slide Time: 00:58)



The screenshot shows the RStudio interface. The main editor window contains a script with the following content:

```
1  
2 Source: https://archive.ics.uci.edu/ml/datasets/banknote+authentication  
3  
4  
5 Owner of database: Volker Lohweg (University of Applied Sciences,  
6 Ostwestfalen-Lippe, volker.lohweg '@' hs-owl.de)  
7 Donor of database: Helene Dörksen (University of Applied Sciences,  
8 Ostwestfalen-Lippe, helene.doerksen '@' hs-owl.de)  
9 Date received: August, 2012  
10 Data Set Information:
```

The console window shows the R version and platform information:

```
R version 4.2.2 (2022-10-31) -- "Innocent and Trusting"  
Copyright (C) 2022 The R Foundation for Statistical Computing  
Platform: aarch64-apple-darwin20 (64-bit)  
  
R is free software and comes with ABSOLUTELY NO WARRANTY.  
You are welcome to redistribute it under certain conditions.  
Type 'license()' or 'licence()' for distribution details.  
  
Natural language support but running in an English locale
```

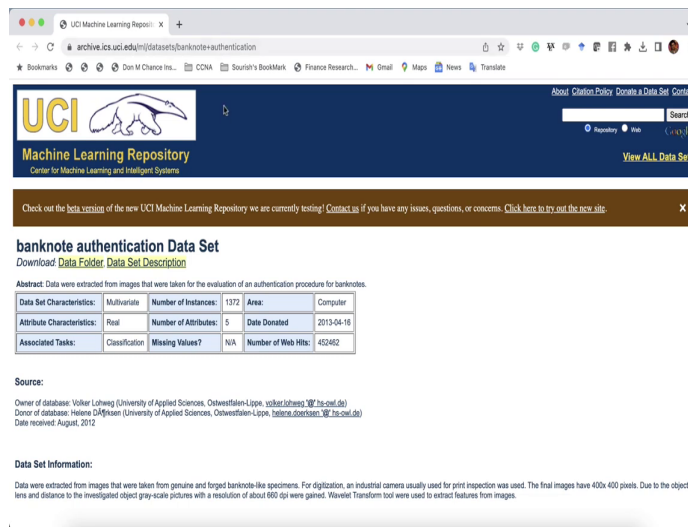
The file explorer on the right shows a list of files in the current project directory:

File Name	Size	Modified
..		
..history	15.6 KB	Mar 23, 2023, 2:20 PM
Bengaluru_House_Data.csv	910 KB	Jan 29, 2023, 12:44 AM
Chemical_1910_2012_Matlab.csv	312.4 KB	Sun 4, 2012, 2:31 PM
uci-chemml.zip	4.3 KB	Dec 1, 2015, 9:13 AM
data_banknote_authentication.Rna	1022 B	Mar 23, 2023, 8:00 AM
data_banknote_authentication.txt	45.3 KB	Mar 23, 2023, 7:56 AM
Football_Data_analysis.R	2.5 KB	Feb 26, 2023, 9:47 PM
German_Credit_Data.csv	103.7 KB	Jan 5, 2018, 9:14 PM
Handwritten_Lecture_04.pdf	238.5 KB	Dec 14, 2022, 12:13 PM
Handwritten_Lecture_04.Rmd	859 B	Dec 14, 2022, 12:13 PM
Network		
NPTEL_Request_for_Week_0_assign.	130 KB	Jan 19, 2023, 2:47 PM
NPTEL_Survey	205 B	Mar 23, 2023, 2:30 PM
Video_Recordings		
Week-01		
Week-02		
Week-03		
Week-04		
Week-05		

The NPTEL logo is visible in the top right corner of the RStudio window.

So, what I am going to do here we have the from the UCI machine learning repository. If we just go there and let me just open my net, yeah ok.

(Refer Slide Time: 01:18)



UCI Machine Learning Repository

Machine Learning Repository
Center for Machine Learning and Intelligent Systems

banknote authentication Data Set

Download: [Data Folder](#) [Data Set Description](#)

Abstract: Data were extracted from images that were taken for the evaluation of an authentication procedure for banknotes.

Data Set Characteristics:	Multivariate	Number of Instances:	1372	Area:	Computer
Attribute Characteristics:	Real	Number of Attributes:	5	Date Donated:	2013-04-16
Associated Tasks:	Classification	Missing Values?	N/A	Number of Web Hits:	452482

Source:

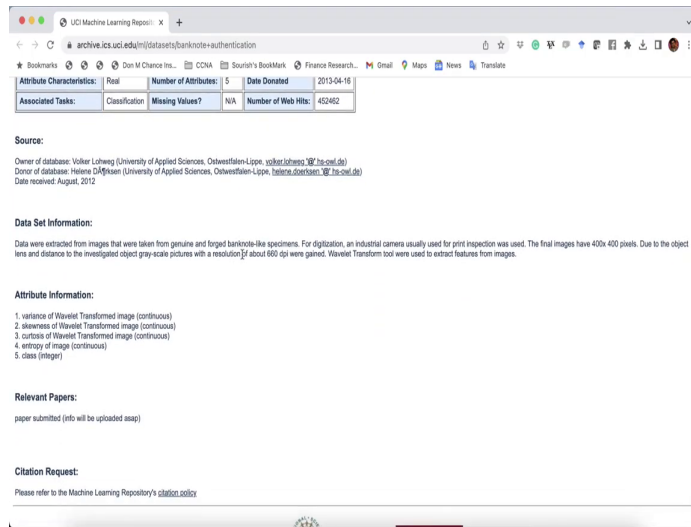
Owner of database: Volker Lohweg (University of Applied Sciences, Ostwestfalen-Lippe, volker.lohweg@fh-oil.de)
Donor of database: Helene Döpfner (University of Applied Sciences, Ostwestfalen-Lippe, helene.doepner@fh-oil.de)
Date received: August, 2012

Data Set Information:

Data were extracted from images that were taken from genuine and forged banknote-like specimens. For digitization, an industrial camera usually used for print inspection was used. The final images have 400x400 pixels. Due to the object lens and distance to the investigated object gray-scale pictures with a resolution of about 600 dpi were gained. Wavelet Transform tool were used to extract features from images.



(Refer Slide Time: 01:29)



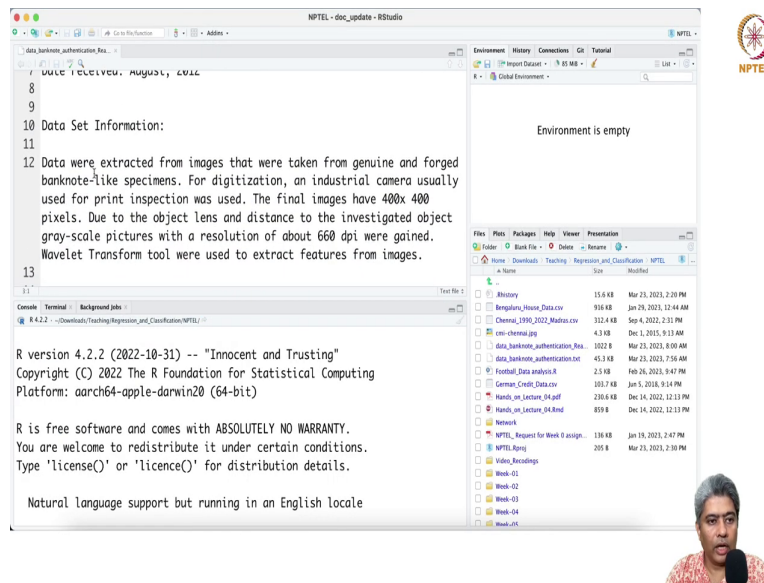
The screenshot shows a web browser window with the URL `archive.ics.uc.edu/ml/datasets/banknote-authentication`. The browser's address bar and tabs are visible at the top. Below the browser window, there is a table with the following data:

Attribute Characteristics:	Real	Number of Attributes:	5	Date Donated:	2013-04-16
Associated Tasks:	Classification	Missing Values?	N/A	Number of Web Hits:	452482

Below the table, there are sections for 'Sources', 'Data Set Information', 'Attribute Information', 'Relevant Papers', and 'Citation Request'. The 'Sources' section lists the owner and donor of the database. The 'Data Set Information' section describes the data extraction process. The 'Attribute Information' section lists five attributes: variance of Wavelet Transformed image, skewness of Wavelet Transformed image, contrast of Wavelet Transformed image, entropy of image, and class. The 'Relevant Papers' section mentions a submitted paper. The 'Citation Request' section refers to the repository's citation policy.



(Refer Slide Time: 01:34)



So, this is the UCI machine learning data repository. This is called bank note authentication data set. So, this is from Denmark. It says basically data were extracted from images that were taken from genuine and forged bank not like specimens. For digitization and industrial camera usually used for print inspection was used.

The final images have a 400 by 400 pixel due to object lens distance to the investigated object grayscale pictures with resolution of about 600 dpi where gained. Wavelet Transformation tool were used to extract the features from the images.

(Refer Slide Time: 02:15)

Wavelet Transform tool were used to extract features from images.

```
13  
14  
15 Attribute Information:  
16  
17 1. variance of Wavelet Transformed image (continuous)  
18 2. skewness of Wavelet Transformed image (continuous)  
19 3. curtosis of Wavelet Transformed image (continuous)  
20 4. entropy of image (continuous)  
21 5. class (integer)  
22  
23  
24
```

R version 4.2.2 (2022-10-31) -- "Innocent and Trusting"
Copyright (C) 2022 The R Foundation for Statistical Computing
Platform: aarch64-apple-darwin20 (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

Natural language support but running in an English locale

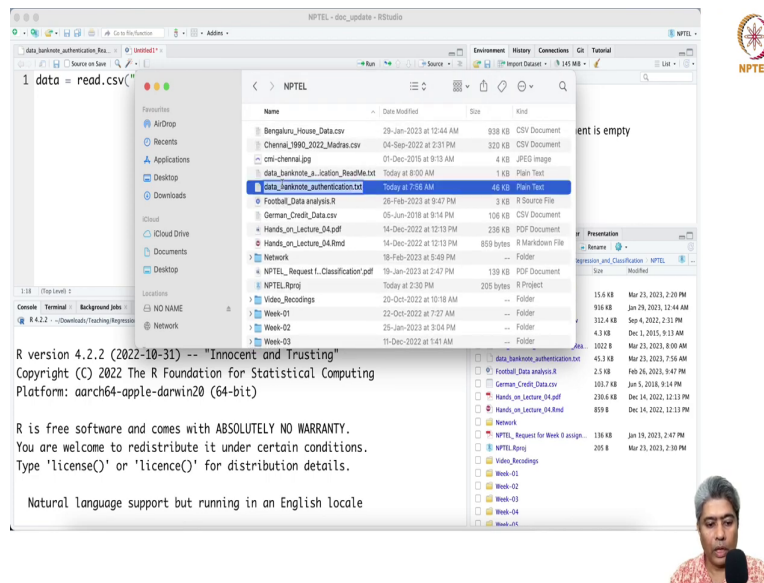
Environment: Global Environment

Files: Packages: Help: Viewer: Presentation

File Name	Size	Modified
..		
..Rhistory	15.6 KB	Mar 23, 2023, 2:20 PM
..Bengaluru_House_Data.csv	910 KB	Jan 29, 2023, 12:44 AM
..Cleaning_1910_2022_Matlab.csv	312.4 KB	Sun 4, 2022, 2:37 PM
..cnc-chemical.jpg	4.3 KB	Dec 1, 2015, 9:13 AM
..data_banknote_authentication.Rna	1022 B	Mar 23, 2023, 8:00 AM
..data_banknote_authentication.Rna	45.3 KB	Mar 23, 2023, 7:56 AM
..Football_Data_analysis.R	2.5 KB	Feb 26, 2023, 9:47 PM
..German_Credit_Data.csv	103.7 KB	Jan 5, 2018, 9:14 PM
..Handwritten_Letter_04.pdf	238.5 KB	Dec 14, 2022, 12:13 PM
..Handwritten_Letter_04.Rmd	859 B	Dec 14, 2022, 12:13 PM
..Network		
..NPTEL_Request_for_Week_0_assign.	130 KB	Jan 19, 2023, 2:47 PM
..NPTEL.Rproj	209 B	Mar 23, 2023, 2:30 PM
..Video_Recordings		
..Week-01		
..Week-02		
..Week-03		
..Week-04		
..Week-05		

And these variants of wavelet transformations, skewness of the wavelet transformations, curtosis of the wavelet transformations, entropy of images and the classification manually done classification 0 or 1 that was done.

(Refer Slide Time: 02:40)



The screenshot shows the RStudio interface. The top-left pane displays the code editor with the command `data = read.csv()`. The top-right pane shows a file explorer for the 'NPTEL' directory, listing various files and folders. The bottom-left pane shows the R console output, which includes the R version (4.2.2), copyright information, and a warning about natural language support.

R version 4.2.2 (2022-10-31) -- "Innocent and Trusting"
Copyright (C) 2022 The R Foundation for Statistical Computing
Platform: aarch64-apple-darwin20 (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

Natural language support but running in an English locale

The file explorer shows the following files and folders:

Name	Date Modified	Size	Kind
Bengaluru_House_Data.csv	29-Jan-2023 at 12:44 AM	938 KB	CSV Document
Chennai_1990_2022_Madras.csv	04-Sep-2022 at 8:31 PM	320 KB	CSV Document
cn-channels.jpg	01-Dec-2016 at 8:19 AM	4 KB	JPEG Image
data_banknote_authentication.txt	Today at 8:00 AM	1 KB	Plain Text
data_banknote_authentication.txt	Today at 7:58 AM	46 KB	Plain Text
Football_Data_analysis.R	28-Feb-2023 at 9:47 PM	3 KB	R Source File
German_Credit_Data.csv	05-Jun-2018 at 9:14 PM	106 KB	CSV Document
Hands_on_Lecture_04.pdf	14-Dec-2022 at 12:13 PM	238 KB	PDF Document
Hands_on_Lecture_04.Rmd	14-Dec-2022 at 12:13 PM	859 bytes	R Markdown File
Network	18-Feb-2023 at 9:49 PM	-	Folder
NPTEL_Request_1_Classification.pdf	19-Jan-2023 at 2:47 PM	139 KB	PDF Document
NPTEL_Rproj	Today at 3:30 PM	205 bytes	R Project
Video_Recordings	20-Oct-2022 at 10:18 AM	-	Folder
Week-01	22-Oct-2022 at 7:27 AM	-	Folder
Week-02	25-Jan-2023 at 3:04 PM	-	Folder
Week-03	11-Dec-2022 at 1:41 AM	-	Folder

So, now let us start analysis. So, data equal to we will do read dot csv read dot csv. Let me just go quickly to the bank note and all downloads and classification. So, from so, this is the data set name. I have shared this data with in the swarm portal.

(Refer Slide Time: 03:25)

The screenshot shows the RStudio interface with the following components:

- Source Editor:** Contains the R code: `data = read.csv("data_banknote_authentication.txt"); header=F`
- Environment:** Shows the variable `data` as a data frame with 1371 observations and 5 variables.
- Files Panel:** Lists files in the current directory, including `data_banknote_authentication.txt` and `data_banknote_authentication.csv`.
- Terminal:** Displays the R help text for `read.csv()`, including instructions on using `contributor()`, `demo()`, `help()`, `help.start()`, and `q()`.



(Refer Slide Time: 03:28)

The screenshot displays the RStudio interface. The main window shows a data frame with 1371 observations and 5 variables. The console shows the following code and output:

```
> data = read.csv("data_banknote_authentication.txt")
> View(data)
>
```

The console also displays the following text:

```
R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.

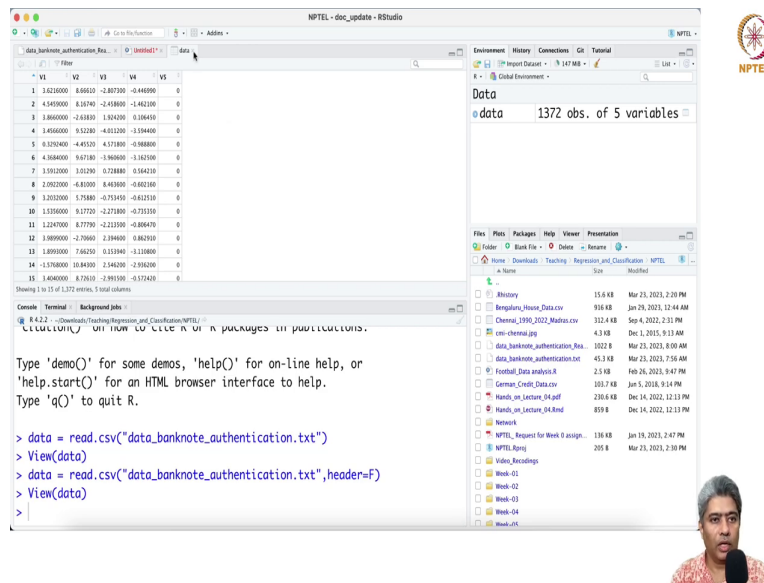
Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.
```

The file browser on the right shows the following files:

File Name	Type	Modified
..		
..history	15.6 KB	Mar 23, 2023, 2:20 PM
Bengaluru_House_Data.csv	910 KB	Jan 29, 2023, 12:44 AM
Chemical_1990_2022_Mat.csv	312.4 KB	Sep 4, 2022, 2:37 PM
csi-chemical.jpg	4.3 KB	Dec 1, 2015, 9:13 AM
data_banknote_authentication.Rsa	1022 B	Mar 23, 2023, 8:00 AM
data_banknote_authentication.txt	45.3 KB	Mar 23, 2023, 7:56 AM
Football_Data_analyis.R	2.5 KB	Feb 26, 2023, 9:47 PM
German_Credit_Data.csv	103.7 KB	Jan 5, 2018, 9:14 PM
Hands_on_Lecture_04.pdf	238.5 KB	Dec 14, 2022, 12:13 PM
Hands_on_Lecture_04.Rmd	859 B	Dec 14, 2022, 12:13 PM
Network		
NPTEL_Request_for_Week_0_assign.	130 KB	Jan 19, 2023, 2:47 PM
NPTEL.Rproj	205 B	Mar 23, 2023, 2:30 PM
Video_Recordings		
Week-01		
Week-02		
Week-03		
Week-04		
Week-05		

So, if we just read this ok. So, it does not have any header because looks like this cannot be the name. So, it does not have any header.

(Refer Slide Time: 03:56)



The screenshot shows the RStudio interface with the following components:

- Environment:** Shows a data frame named 'data' with 1372 observations and 5 variables.
- Files:** A file explorer showing the current directory with various files and folders.
- Console:** Contains the following R code and output:

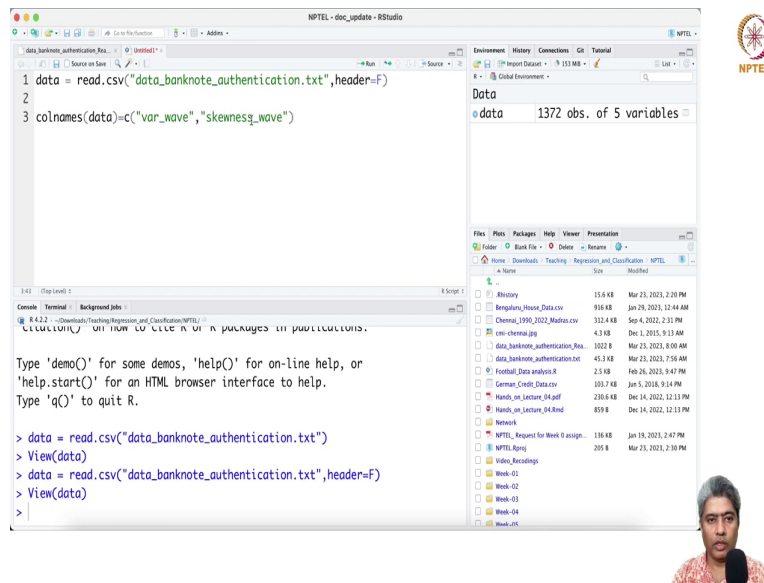
```
> data = read.csv("data_banknote_authentication.txt")
> View(data)
> data = read.csv("data_banknote_authentication.txt", header=F)
> View(data)
>
```
- Terminal:** Shows the R version (R 4.2.2) and the current working directory (C:/Users/.../Documents/Teaching/Regression and Classification/NPTEL).

The data frame 'data' has the following structure:

V1	V2	V3	V4	V5
1	1.6120000	0.69600	-2.807300	-0.446900
2	1.5450000	0.16740	-2.451800	-1.462100
3	1.8660000	-2.51830	1.524200	0.106450
4	1.4560000	0.52280	-4.011200	-1.594400
5	1.1704000	-0.41520	-4.071800	-0.506800
6	4.1040000	-0.42130	-1.962600	-1.162100
7	1.5912000	1.01290	0.728880	0.564210
8	2.0020000	-0.41000	0.483600	-0.602160
9	1.2012000	1.75880	-0.751450	-0.612310
10	1.1510000	0.17720	-2.271800	-0.731510
11	1.1247000	0.77790	-2.211000	-0.884470
12	1.0095000	-0.20660	2.194600	0.862910
13	1.8930000	7.66250	0.151940	-1.110800
14	-1.1768000	10.84300	2.546200	-2.936200
15	1.4040000	0.72610	-2.981500	-0.577420

So, we have to say header equal to header equal to false header equal to false and now we have the data correctly read, alright.

(Refer Slide Time: 04:04)



The screenshot shows the RStudio interface with the following content:

```
1 data = read.csv("data_banknote_authentication.txt",header=F)
2
3 colnames(data)=c("var_wave", "skewnes_wave")
```

Environment: data (1372 obs. of 5 variables)

```
> data = read.csv("data_banknote_authentication.txt")
> View(data)
> data = read.csv("data_banknote_authentication.txt",header=F)
> View(data)
>
```

Console output:

```
R 4.2.2 - Checkpoint (Teaching Regression and Classification)
C:\Users\user> R CMD SHLIB "C:\Program Files\R\R422\bin\i386\Rcpp.dll"
Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.

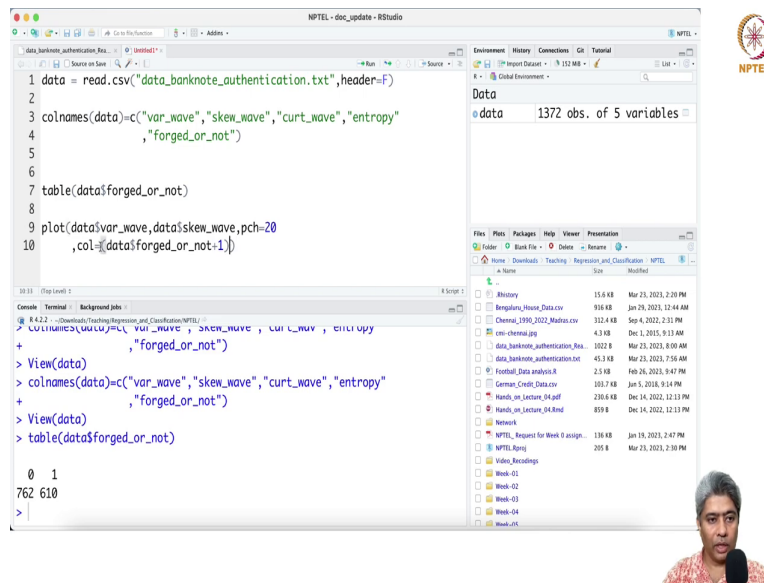
> data = read.csv("data_banknote_authentication.txt")
> View(data)
> data = read.csv("data_banknote_authentication.txt",header=F)
> View(data)
>
```

Files list:

File Name	Type	Modified
..		
..History	15.6 KB	Mar 23, 2023, 2:20 PM
Bengaluru_House_Data.csv	910 KB	Jan 29, 2023, 12:44 AM
Chemical_1990_2022_Matlab.csv	312.4 KB	Sun 4, 2023, 2:31 PM
csi-chemical.jpg	4.3 KB	Dec 1, 2015, 9:13 AM
data_banknote_authentication.R	45.3 KB	Mar 23, 2023, 8:00 AM
data_banknote_authentication.txt	1022 B	Mar 23, 2023, 7:56 AM
Football_Data_analyis.R	2.5 KB	Feb 26, 2023, 9:47 PM
German_Credit_Data.csv	103.7 KB	Jan 5, 2018, 9:14 PM
Hands_on_Lecture_04.pdf	238.5 KB	Dec 14, 2022, 12:13 PM
Hands_on_Lecture_04.Rmd	859 B	Dec 14, 2022, 12:13 PM
Network		
NPTEL_Request_for_Week_0_assign.	130 KB	Jan 19, 2023, 2:47 PM
NPTEL.Rproj	205 B	Mar 23, 2023, 2:30 PM
Video_Recordings		
Week-01		
Week-02		
Week-03		
Week-04		
Week-05		

Now, what we have to do? We have to put the colnames data equal to c. First, let us take these name variants of wavelet. So, first one will be var of wavelet, second will be the skewness, curtosis and entropy, ok.

(Refer Slide Time: 04:45)



The screenshot shows the RStudio interface with the following code in the script editor:

```
1 data = read.csv("data_banknote_authentication.txt",header=F)
2
3 colnames(data)=c("var_wave", "skew_wave", "curt_wave", "entropy"
4                 , "forged_or_not")
5
6
7 table(data$forged_or_not)
8
9 plot(data$var_wave,data$skew_wave,pch=20
10      ,col=(data$forged_or_not+1))
```

The console output shows the execution of these commands:

```
> View(data)
> colnames(data)=c("var_wave", "skew_wave", "curt_wave", "entropy"
+                 , "forged_or_not")
> View(data)
> table(data$forged_or_not)
 0  1
762 610
> |
```

The Environment pane on the right shows the 'data' object with 1372 observations and 5 variables. The Files pane on the bottom right shows a list of files in the current directory.

Skewness of wavelet maybe I will just say skew. I do not want to put very large name. Skewness, curtosis of wavelets and “entropy”; and finally, last one is the class which is forged or not, right. So, let me just run this. So, now we have the these names are being given ok. So, maybe I will just say wave. Yeah. And now I think this is fine, right, ok.

Next, I will just have a look that how many forged cases and how many original note were there. So, out of 13,722 13,000, no, sorry, 1372 observations, 610 of them are forged and 762 of them are not forged. They are true original bank note. So, what I am going to do?

I am going to first plot data, dollar, variance of wave, comma, data, dollar skewness of wavelength maybe pch equal to 20. Now, what I am going to do is something you will I think

(Refer Slide Time: 07:15)

The screenshot shows an RStudio session with the following code in the editor:

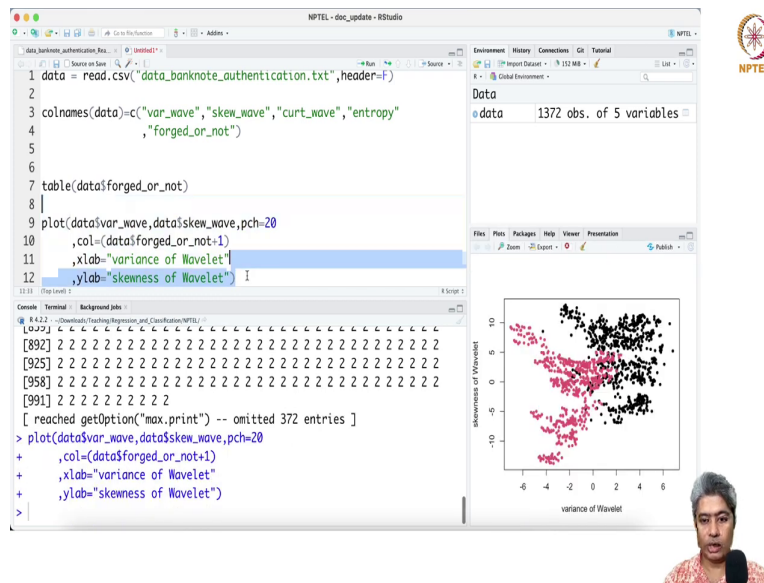
```
1 data = read.csv("data_banknote_authentication.txt", header=F)
2
3 colnames(data) = c("var_wave", "skew_wave", "curt_wave", "entropy",
4                   "forged_or_not")
5
6
7 table(data$forged_or_not)
8
9 plot(data$var_wave, data$skew_wave, pch=20
10      , col=data$forged_or_not+1)
```

The console output shows a matrix of 1s and 2s, where the first column contains 1s and the second column contains 2s. The rows are labeled with IDs: [529], [562], [595], [628], [661], [694], [727], [760], [793], [826], [859].

The RStudio interface also shows the Environment pane with a 'data' object containing 1372 observations and 5 variables. The Files pane shows the project structure.

Now, what happens if I add 1 to it? What is happening? You actually is very simple. I am just adding either 1 or 0. It was 0 or 1. So, all 0s have become 1 and 1s have become 2. Now, as a color, 1 color equal to 1 means it is it will be black and color equal to 2 means it will be red, ok.

(Refer Slide Time: 07:30)




The screenshot displays the RStudio interface. The script editor contains the following R code:

```
1 data = read.csv("data_banknote_authentication.txt", header=F)
2
3 colnames(data)=c("var_wave", "skew_wave", "curt_wave", "entropy",
4                 "forged_or_not")
5
6
7 table(data$forged_or_not)
8
9 plot(data$var_wave, data$skew_wave, pch=20
10      , col=(data$forged_or_not+1)
11      , xlab="variance of Wavelet"
12      , ylab="skewness of Wavelet")
```

The console shows the output of the code, including the distribution of the 'forged_or_not' variable and the execution of the plot command. The plot window shows a scatter plot with 'variance of Wavelet' on the x-axis and 'skewness of Wavelet' on the y-axis. The data points are colored based on the 'forged_or_not' variable, with pink points representing forged banknotes and black points representing authentic ones.

So, this is a this is I am taking making forged or not as the color that I am going to use xlab is let me just take simply “variance of Wavelet”, ok. And ylab, ylab equal to this is simply “skewness of Wavelet” transformed image, ok.

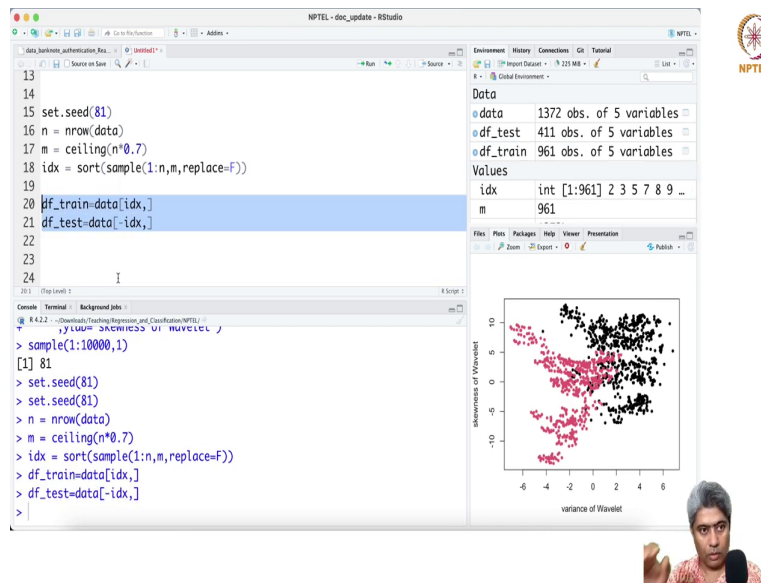
(Refer Slide Time: 08:06)



Let me just make this plot. Yeah, let me just zoom it. So, this is the data. So, now we see this data is very peculiar. The data has a very peculiar shape. You can see there is a very peculiar shape and there is bit of a overlap between the two kind of things. So, one is was forged, which was marked as red or these are like these are the specimens or points which have which represents forged note and the black ones are the one, which represent the original note.

So, there is we see there is a bit of a overlapping is happening, but mostly they are kind of well separated with slight overlapping between the 2. So, can we use wavelet variance and skewness of the wavelet these two feature to do the classification.

(Refer Slide Time: 09:19)



The screenshot displays the RStudio interface. The script editor contains the following R code:

```
13
14
15 set.seed(81)
16 n = nrow(data)
17 m = ceiling(n*0.7)
18 idx = sort(sample(1:n,m,replace=F))
19
20 df_train=data[idx,]
21 df_test=data[-idx,]
22
23
24
```

The console shows the execution of the code:

```
R 4.2.2 ~RStudio: Training Regression and Classification with Wavelet
> sample(1:10000,1)
[1] 81
> set.seed(81)
> set.seed(81)
> n = nrow(data)
> m = ceiling(n*0.7)
> idx = sort(sample(1:n,m,replace=F))
> df_train=data[idx,]
> df_test=data[-idx,]
>
```



The Environment pane on the right shows the following objects:

- data: 1372 obs. of 5 variables
- df_test: 411 obs. of 5 variables
- df_train: 961 obs. of 5 variables

The Values pane shows the values for 'idx' and 'm':

```
idx int [1:961] 2 3 5 7 8 9 ...
m    961
```

A scatter plot is visible in the bottom right, with 'variance of Wavelet' on the x-axis and 'skewness of Wavelet' on the y-axis. The plot shows two distinct clusters of points: one in pink and one in black.



So, what we will do we will first we will set a seed say sample first we need a set a seed. So, we need a number 1 is to say 1000. So, we will take 81. So, 81 as sitting our seed set dot seed ok, alright. And then n is the nrow of data and m is the ceiling of n times 0.7.

So, I am taking 70 percent for training data set and rest of them has test idx equal to sort, sample, 1, is to n comma m comma replace equal to false, ok. And then so, let me just have this. So, this idx is 961 cases have come and then now df train equal to data idx comma and df test equal to test equal to minus idx. So, let me just run the whole thing, yeah.

(Refer Slide Time: 11:10)

The screenshot shows the RStudio interface. The script editor contains the following R code:

```
17 m = ceiling(n*0.7)
18 idx = sort(sample(1:n,m,replace=F))
19
20 df_train=data[idx,]
21 df_test=data[-idx,]
22
23 fit1 = glm(forged_or_not~var_wave+skew_wave, data= df_train
24           ,family = binomial(link="logit"))
25
26 summary(fit1)
27
28
```

The console shows the execution of the code:

```
> set.seed(81)
> n = nrow(data)
> m = ceiling(n*0.7)
> idx = sort(sample(1:n,m,replace=F))
>
> df_train=data[idx,]
> df_test=data[-idx,]
> fit1 = glm(forged_or_not~var_wave+skew_wave, data= df_train
+           ,family = binomial(link="logit"))
+
>
```

The Environment pane shows the following objects:

df_test	411 obs. of 5 variables
df_train	961 obs. of 5 variables
fit1	List of 30

The Values pane shows the following values:

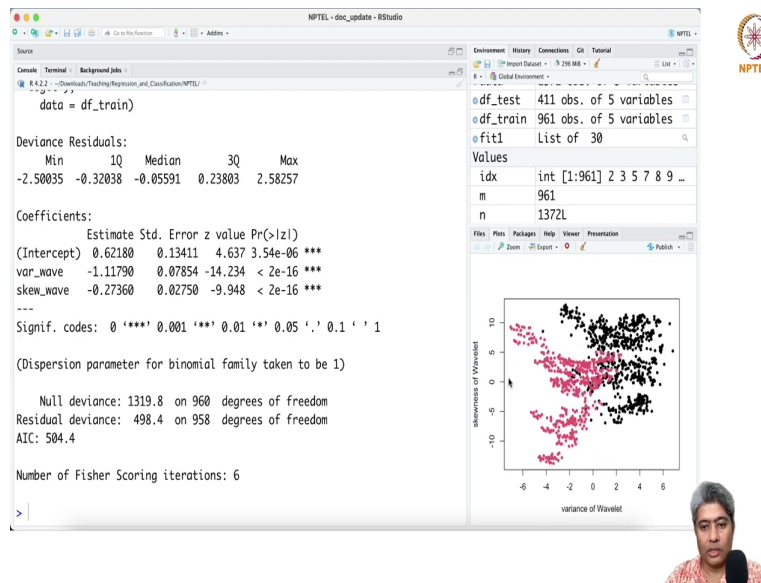
idx	int [1:961] 2 3 5 7 8 9 ...
m	961
n	1372L

The plot shows a scatter plot of 'skewness of Wavelet' (y-axis, ranging from -1.0 to 1.0) versus 'variance of Wavelet' (x-axis, ranging from -6 to 6). The data points are colored in pink and black, showing a clear separation between the two groups.

So, first thing I am going to do I am going to fit a simple model with wavelet variance of wavelet and the skewness of wavelet as the you know feature of simple logistic regression fit. First model will going to be glm, ok glm. So, what we have is forged or not is as this comma variance of wavelet plus skewness of wavelet.

These two data equal to data equal to df train and then family equal to binomial link equal to “logit”, alright. Now, this is what I am going to do. I am going to run this model. Now, if you run this summary fit1.

(Refer Slide Time: 12:34)



The screenshot displays the RStudio interface. The console on the left shows the execution of a linear regression model. The data is split into training and testing sets. The model coefficients are highly significant, indicating a strong relationship between the variables. The scatter plot on the right visualizes the relationship between the variance of the wavelet and its skewness, showing a clear negative correlation.

```
data = df_train)

Deviance Residuals:
    Min       1Q   Median       3Q      Max
-2.50035  -0.32038  -0.05591   0.23803   2.58257

Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)  0.62180    0.13411   4.637 3.54e-06 ***
var_wave    -1.11790    0.07854 -14.234 < 2e-16 ***
skew_wave    -0.27360    0.02750  -9.948 < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 1319.8 on 960 degrees of freedom
Residual deviance: 498.4 on 958 degrees of freedom
AIC: 504.4

Number of Fisher Scoring iterations: 6
```

Variable	Obs	Variables
df_test	411	5
df_train	961	5
fit1	List of 30	

Values

Variable	Value
idx	int [1:961] 2 3 5 7 8 9 ...
m	961
n	1372L

Skewness of Wavelet vs Variance of Wavelet

So, what you have it is saying that both variance of wavelet and the skewness of the wavelet is going to have a very strong accuracy sorry strong effect on the both variance of wavelet and the skewness of the wavelet is going to have a on the whether note will be forged or not, ok.

(Refer Slide Time: 13:08)

The screenshot shows the RStudio interface. The script editor contains the following R code:

```
18 idx = sort(sample(1:n,m,replace=F))
19
20 df_train=data[idx,]
21 df_test=data[-idx,]
22
23 fit1 = glm(forged_or_not~var_wave+skew_wave, data= df_train
24           , family = binomial(link="logit"))
25
26 summary(fit1)
27
28 df_test$prob1=predict(fit1,newdata = df_test,type="response")
29
```

The Environment pane shows the following objects:

Object	Details
df_test	411 obs. of 6 variables
df_train	961 obs. of 5 variables
fit1	List of 30

The Values pane shows the following values for the fit1 object:

Variable	Value
idx	int [1:961] 2 3 5 7 8 9 ...
m	961
n	1372L

The Console pane shows the following output:

```
> df_test$prob1=predict(fit1,newdata = df_test,type="response")
>
```

The plot shows the relationship between the variance of Wavelet (x-axis, ranging from -6 to 6) and the skewness of Wavelet (y-axis, ranging from -1.0 to 1.0). The data points are colored in pink and black, showing a clear separation between the two groups.

So, let us take that as and then now what we are going to do we are going to make the prediction on the test data set, ok. df test df test dollar probability1 equal to predict fit1, newdata equal to df test and then type equal to “response”, ok.

(Refer Slide Time: 13:54)

The screenshot shows the RStudio interface. The main editor displays a data frame with columns: var_wave, skew_wave, curt_wave, entropy, kurtosis, and prob1. The environment pane on the right shows objects: df_test (411 obs. of 6 variables), df_train (961 obs. of 5 variables), and fit1 (List of 30). The console shows the following code and output:

```
> df_test$prob1 = predict(fit1, newdata = df_test, type = "response")
> View(df_test)
>
```

The scatter plot in the bottom right shows 'skewness of Wavelet' on the y-axis (ranging from -1.0 to 1.0) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The data points are colored in pink and black, showing a clear separation between the two groups.



So, now if you see this, alright. So, now you can see that these values have come 0 and all this probability of 0.24.

(Refer Slide Time: 14:17)

The screenshot displays the RStudio interface. The top-left pane shows a data table with columns: `var_wave`, `skewness`, `var_wave`, `entropy`, `forged_or_not`, and `prob1`. The console shows the following R code and output:

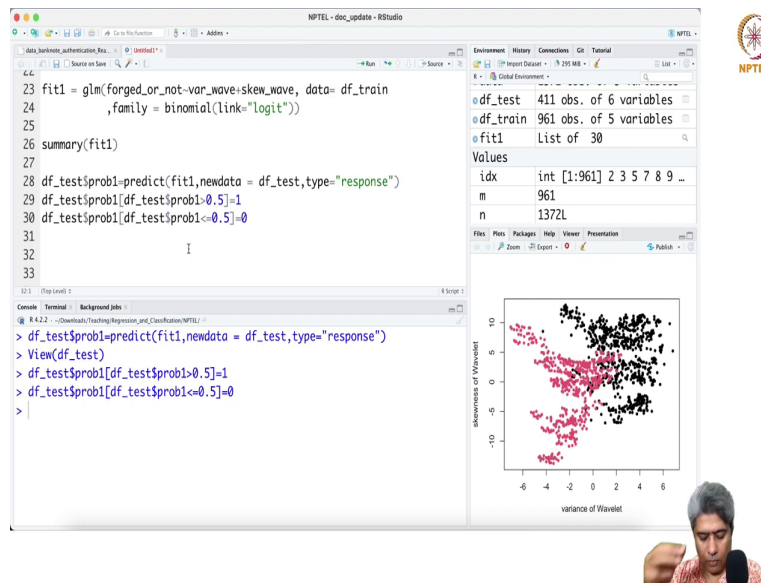
```
> df_test$prob1=predict(fit1,newdata = df_test,type="response")
> View(df_test)
> |
```

The top-right pane shows the environment with objects: `df_test` (411 obs. of 6 variables), `df_train` (961 obs. of 5 variables), and `fit1` (List of 30). The bottom-right pane shows a scatter plot of `skewness of Wavelet` (y-axis, -1.0 to 1.0) versus `variance of Wavelet` (x-axis, -6 to 6). The plot shows two clusters of points: one in pink and one in black.



So, let us see if it is. So, mostly they are close to 0 some 0.1, 0.2 and if we just go down to 1 and you will see that they are reasonably close to 0.9, 0.8, 0.9, 0.8, right. Here is a one point where you got a you know 0.26, but the probability is less, but you know actually it is a forged note. So, there will be looks like there could be some misclassification possible. So, what we are going to do? We are going to make the prediction. These are the probability of this particular node to be forged note.

(Refer Slide Time: 14:51)



The screenshot displays the RStudio interface. The main editor window contains the following R code:

```
23 fit1 = glm(forged_or_not~var_wave+skew_wave, data= df_train
24           , family = binomial(link="logit"))
25
26 summary(fit1)
27
28 df_test$prob1=predict(fit1,newdata = df_test,type="response")
29 df_test$prob1[df_test$prob1>0.5]=1
30 df_test$prob1[df_test$prob1<=0.5]=0
31
32
33
```

The Environment pane on the right shows the following objects:

- df_test: 411 obs. of 6 variables
- df_train: 961 obs. of 5 variables
- fit1: List of 30

The Values pane shows the following data:

idx	int	[1:961]	2	3	5	7	8	9	...
m									961
n									1372L

The Console window shows the following commands and output:

```
> df_test$prob1=predict(fit1,newdata = df_test,type="response")
> View(df_test)
> df_test$prob1[df_test$prob1>0.5]=1
> df_test$prob1[df_test$prob1<=0.5]=0
>
```

A scatter plot is visible in the bottom right corner, with 'variance of Wavelet' on the x-axis (ranging from -6 to 6) and 'skewness of Wavelet' on the y-axis (ranging from -1.0 to 1.0). The plot shows two clusters of points: one in pink and one in black.

The NPTEL logo is visible in the top right corner of the RStudio window.

Now, I am going to make the prediction df testdollarprob1 equal to df testdollarprob1. If it is 0 greater than 0.5, we will call it 1. And if it is less than equal to 0.5, we will call it 0. So; that means, not forged ok, alright.

(Refer Slide Time: 15:32)

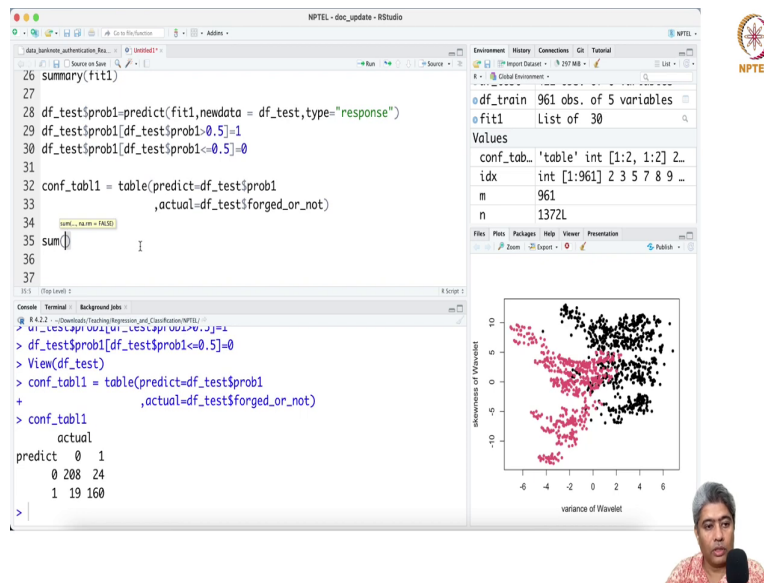
The screenshot shows the RStudio interface. The top-left pane displays a data table with columns: `var_wave`, `skew_wave`, `curr_wave`, `entropy`, `forpnl_cr_pct`, and `prob1`. The top-right pane shows environment variables: `df_test` (411 obs. of 6 variables), `df_train` (961 obs. of 5 variables), and `fit1` (List of 30). The bottom-right pane shows a scatter plot of `skewness of Wavelet` (y-axis, -1.0 to 1.0) versus `variance of Wavelet` (x-axis, -6 to 6). The plot shows two clusters of points: one in pink and one in black. The bottom-right corner features a small video feed of a person speaking.

```
df_test 411 obs. of 6 variables
df_train 961 obs. of 5 variables
fit1 List of 30
Values
idx int [1:961] 2 3 5 7 8 9 ...
m 961
n 1372L
```

```
> df_test$prob1=predict(fit1,newdata = df_test,type="response")
> View(df_test)
> df_test$prob1[df_test$prob1>0.5]=1
> df_test$prob1[df_test$prob1<=0.5]=0
> View(df_test)
>
```

Now, what we are going to do? So, now, if I just look into the df test. So, now, they are all. So, here are there are some mis-classification we can see, but mostly they are agreeing.

(Refer Slide Time: 15:45)



The screenshot shows an RStudio interface with the following elements:

- Source Editor:** Contains R code for model fitting and evaluation:

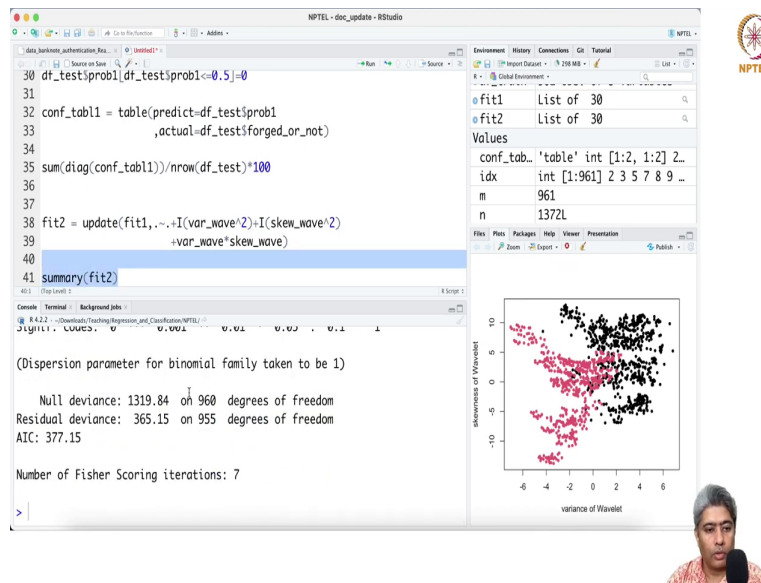
```
summary(fit1)
df_test$prob1=predict(fit1,newdata = df_test,type="response")
df_test$prob1[df_test$prob1<=0.5]=1
df_test$prob1[df_test$prob1<=0.5]=0
conf_tab1 = table(predict=df_test$prob1
                  ,actual=df_test$forged_or_not)
sum()
```
- Environment:** Shows objects: df_train (961 obs. of 5 variables), fit1 (List of 30), and Values (conf_tab1, idx, m, n).
- Console:** Shows the execution of the code and the resulting confusion matrix:

```
> df_test$prob1[df_test$prob1<=0.5]=0
> View(df_test)
> conf_tab1 = table(predict=df_test$prob1
+                  ,actual=df_test$forged_or_not)
> conf_tab1
      actual
predict 0  1
      0 208 24
      1  19 160
```
- Plot:** A scatter plot titled 'skewness of Wavelet' vs 'variance of Wavelet'. The x-axis ranges from -6 to 6, and the y-axis ranges from -10 to 10. Data points are colored black and pink.
- NPTEL Logo:** Located in the top right corner.
- Speaker:** A small video inset of a man speaking is visible in the bottom right corner.

So, now we are going to I am going to calculate a confusion table, ok. So, table predict equal to df testdollarprob1 comma actual equal to df test dollar forged or not, ok. Now, if I just run the confusion table. So, this is my confusion table actually predicted 0 and actually 0 was 208 cases and in is the test data set. Similarly, predicted 1 and actually it was 1 was 160 cases.

There are 24 cases where it was predicted it is not forged, but it is actually forged and there are 19 cases where it was predicted forged, but it was not actually forged. So, there are possibility of and let us just check the accuracy.

(Refer Slide Time: 17:06)



The screenshot shows the RStudio interface with the following content:

```
30 df_test$probl<-df_test$probl<-0.5]-0
31
32 conf_tab1 = table(predict=df_test$probl
33                   ,actual=df_test$forged_or_not)
34
35 sum(diag(conf_tab1))/nrow(df_test)*100
36
37
38 fit2 = update(fit1,.-I(var_wave^2)+I(skew_wave^2)
39              +var_wave*skew_wave)
40
41 summary(fit2)
```

Console output:

```
(Dispersion parameter for binomial family taken to be 1)
Null deviance: 1319.84 on 960 degrees of freedom
Residual deviance: 365.15 on 955 degrees of freedom
AIC: 377.15

Number of Fisher Scoring iterations: 7
```

Environment pane:

fit1	List of 30
fit2	List of 30

Values pane:

conf_tab1	'table' int [1:2, 1:2] 2...
idx	int [1:961] 2 3 5 7 8 9 ...
m	961
n	1372L

Scatter plot: skewness of Wavelet vs variance of Wavelet. The plot shows two clusters of points: one in red (lower variance and skewness) and one in black (higher variance and skewness).

So, first I am going to take the diagonal of the confidence table. These are the correct predictions all correct prediction take the sum of that; sum of that. These are the all 368 are the total correct prediction and total number of all test is. Actually, just I can take nrow of df test is 411 of them there are 411 of them if I just make 100 out of this so, 89.53. So, this is my accuracy 89.53.

Now, if I make, I will make a second model fit2. What I will do is essentially I will just update fit 1 with same model plus few more engineered feature. So, I will just take variance of wave and square them plus a skewness of wave lengths and square them and finally, plus variance of wave times skewness of wave, ok. If I just run this and summary let me run the summary of fit2.

(Refer Slide Time: 19:20)

The screenshot shows the RStudio interface. The script editor contains the following code:

```
30 df_test$probl[df_test$probl<=0.5]=0
31
32 conf_tab1 = table(predict=df_test$probl
33                   ,actual=df_test$forged_or_not)
34
35 sum(diag(conf_tab1))/nrow(df_test)*100
36
37
38 fit2 = update(fit1, ~.~I(var_wave^2)+I(skew_wave^2)
39              +var_wave*skew_wave)
40
41 summary(fit2)
```

The console output shows the coefficients for the model fit2:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.564866	0.204163	7.665	1.79e-14 ***
var_wave	-1.876305	0.176926	-10.605	< 2e-16 ***
skew_wave	-0.060625	0.038243	-1.585	0.1129
I(var_wave^2)	0.057639	0.029567	1.949	0.0512 .
I(skew_wave^2)	-0.056032	0.007548	-7.424	1.14e-13 ***
var_wave:skew_wave	0.142255	0.026317	5.405	6.47e-08 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The environment pane shows fit1 and fit2 as lists of length 30. The values pane shows the structure of the confusion matrix. The plot pane shows a scatter plot of 'skewness of Wavelet' (y-axis, -10 to 10) versus 'variance of Wavelet' (x-axis, -6 to 6). The plot shows two clusters of points: one in red (skewness < 0) and one in black (skewness > 0).

So, looks like variance of wavelets and skewness of wavelet ok was. So, initially variance of wavelet ok if I just run the summary of fit1 we see that both wavelet both variance and skewness of the wavelet is as effect, but if I run on the forged or not whereas, if I run summary of fit2 where, it shows basically variance of wavelet does have a effect skewness does not, but the quadratic effect of skewness does have a effect and their interaction does have a effect.

So, this is very interesting phenomena, but let us see what is the whether adding this feature help me increasing the out of the sample accuracy or not. So, that will be my that will be the real test case.

(Refer Slide Time: 20:33)

The screenshot shows an RStudio session. The script editor contains the following code:

```
36
37
38 fit2 = update(fit1, ~.-I(var_wave^2)+I(skew_wave^2)
39               +var_wave*skew_wave)
40
41 summary(fit2)
42
43
44 df_test$prob2=predict(fit2,newdata = df_test,type="response")
45 df_test$prob2[df_test$prob2>0.5]=1
46 df_test$prob2[df_test$prob2<=0.5]=0
47
```

The console shows the following output:

```
Residual deviance: 365.15 on 955 degrees of freedom
AIC: 377.15

Number of Fisher Scoring iterations: 7

> df_test$prob2=predict(fit2,newdata = df_test,type="response")
> View(df_test)
> df_test$prob2[df_test$prob2>0.5]=1
> df_test$prob2[df_test$prob2<=0.5]=0
>
```

The Environment pane shows:

- df_test: 411 obs. of 7 variables
- df_train: 961 obs. of 5 variables
- fit1: List of 30
- fit2: List of 30

The Values pane shows:

```
conf_tab_ 'table' int [1:2, 1:2] 2_
idx       int [1:961] 2 3 5 7 8 9 ...
```

The plot shows 'skewness of Wavelet' on the y-axis (ranging from -1.0 to 1.0) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The data points are colored black and pink, showing a clear separation between the two groups.

So, what I will do? I will just take this guy copy and paste it here, but here instead of probability1 I have to take probability2 and then instead of fit2 I will just do fit1 I will just do fit2 it will add a new column it should add a new column the probability2, ok.

(Refer Slide Time: 20:47)

The screenshot displays the RStudio interface. The top-left pane shows a data table with columns: `var_wave`, `show_wave`, `cut_wave`, `entropy`, `target_or_not`, `prob1`, and `prob2`. The top-right pane shows environment variables: `df_test` (411 obs. of 7 variables), `df_train` (961 obs. of 5 variables), `fit1` (List of 30), and `fit2` (List of 30). The bottom-left pane (Console) shows the following R code and output:

```
AIC: 377.15

Number of Fisher Scoring iterations: 7

> df_test$prob2=predict(fit2,newdata = df_test,type="response")
> View(df_test)
> df_test$prob2[df_test$prob2>0.5]=1
> df_test$prob2[df_test$prob2<=0.5]=0
> View(df_test)
>
```

The bottom-right pane shows a scatter plot of `showness of Wavelet` (y-axis, ranging from -1.0 to 1.0) versus `variance of Wavelet` (x-axis, ranging from -6 to 6). The plot contains two clusters of points: one in pink and one in black. A small inset image of a man speaking is visible in the bottom right corner of the RStudio window.

And then on the probability 2 I will make the prediction if it is more than 5, 0.5 I will say 1 otherwise I will say it is 0. So, now they are making 0 1, ok. Here is a misclassification; here was misclassification it was corrected here is a misclassification here is misclassification.

(Refer Slide Time: 21:26)

The screenshot shows an RStudio session with the following code in the editor:

```
43  
44 df_test$prob2 = predict(fit2, newdata = df_test, type = "response")  
45 df_test$prob2[df_test$prob2 > 0.5] = 1  
46 df_test$prob2[df_test$prob2 <= 0.5] = 0  
47  
48  
49 conf_tabl2 = table(predict = df_test$prob2  
50 , actual = df_test$forged_or_not)  
51  
52  
53 sum(diag(conf_tabl2)) / nrow(df_test) * 100  
54
```

The console output shows the execution of the code:

```
> view(w_1_1_1_1)  
> conf_tabl2 = table(predict = df_test$prob2  
+ , actual = df_test$forged_or_not)  
> conf_tabl2  
      actual  
predict 0  1  
      0 212 18  
      1  15 166  
> sum(diag(conf_tabl2)) / nrow(df_test) * 100  
[1] 91.9708  
>
```



The environment pane shows:

- df_train: 961 obs. of 5 variables
- fit1: List of 30
- fit2: List of 30

The Values pane shows:

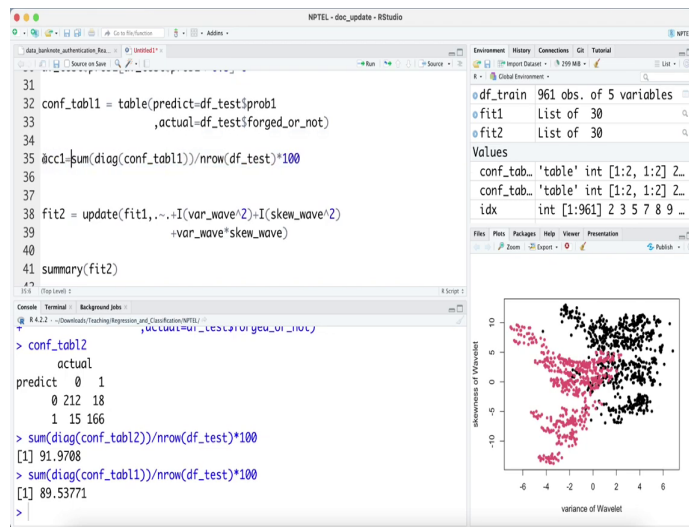
```
conf_tabl_ 'table' int [1:2, 1:2] 2_...  
conf_tabl_ 'table' int [1:2, 1:2] 2_...  
idx        int [1:961] 2 3 5 7 8 9 ...
```

The plot shows 'skewness of Wavelet' on the y-axis (ranging from -10 to 10) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The data points are colored in pink and black, showing a clear separation between the two classes.



So, we will see and we can just calculate a confusion table simply second confusion table with probability2 and what is the situation of this is the second confusion table for this with the second model and if we just compute the accuracy simple accuracy this is 91.97. So, 92 percent accuracy we are getting whereas, for from this first model we are getting 89.5 percent.

(Refer Slide Time: 22:02)



The screenshot displays the RStudio interface with the following content:

```
31
32 conf_tab1 = table(predict=df_test$prob1
33                   ,actual=df_test$forged_or_not)
34
35 acc1 = sum(diag(conf_tab1))/nrow(df_test)*100
36
37
38 fit2 = update(fit1,.-I(var_wave^2)+I(skew_wave^2)
39              ,+var_wave*skew_wave)
40
41 summary(fit2)
```

Environment: Global Environment

- df_train: 961 obs. of 5 variables
- fit1: List of 30
- fit2: List of 30

Values:

- conf_tab_ 'table' int [1:2, 1:2] 2_
- conf_tab_ 'table' int [1:2, 1:2] 2_
- idx: int [1:961] 2 3 5 7 8 9 _

Console:

```
> conf_tab1
      actual
predict 0  1
       0 212 18
       1  15 166
> sum(diag(conf_tab2))/nrow(df_test)*100
[1] 91.9708
> sum(diag(conf_tab1))/nrow(df_test)*100
[1] 89.53771
>
```

Figure: A scatter plot showing the relationship between the variance of Wavelet (x-axis, ranging from -6 to 6) and the skewness of Wavelet (y-axis, ranging from -1.0 to 1.0). The data points are colored in black and pink, showing a positive correlation between the two variables.



(Refer Slide Time: 22:23)

The screenshot displays the RStudio interface. The script editor contains the following R code:

```
49 conf_tabl2 = table(predict=df_test$prob2
50                   ,actual=df_test$forged_or_not)
51
52
53 acc2=sum(diag(conf_tabl2))/nrow(df_test)*100
54
55 acc2-acc1
56
57 ## visualisation to understand the effect of feature engineering
58
59 test_data = dum_var = data.frame(matrix(NA,nrow=1,ncol=2))
60
```



The console shows the following output:

```
> sum(diag(conf_tabl2))/nrow(df_test)*100
[1] 91.9708
> sum(diag(conf_tabl1))/nrow(df_test)*100
[1] 89.53771
> acc1=sum(diag(conf_tabl1))/nrow(df_test)*100
> acc2=sum(diag(conf_tabl2))/nrow(df_test)*100
> acc2-acc1
[1] 2.43309
> test_data = dum_var = data.frame(matrix(NA,nrow=1,ncol=2))
>
```

The Environment pane shows the following variables:

Variable	Class	Value
acc1	numeric	89.5377128953771
acc2	numeric	91.970802919708
conf_tabl1	matrix	int [1:2, 1:2] 2...
conf_tabl2	matrix	int [1:2, 1:2] 2...
idx	matrix	int [1:961] 2 3 5 7 8 9 ...

The plot pane shows a scatter plot with 'variance of Wavelet' on the x-axis (ranging from -6 to 6) and 'skewness of Wavelet' on the y-axis (ranging from -10 to 10). The plot contains two clusters of points: one in pink and one in black.



So, nearly 2.5 percent gain that we are getting, ok. So, this could be accuracy1 and this could be accuracy2 and accuracy2 minus accuracy1. So, 2.4 percent gain that we are getting because of the because of engineering proper feature engineering. Now, let us see how the geometry of the predictive space has changed using because of adding the feature engineer engineered feature.

So, I will do some visualisation to understand the effect of feature engineering, ok. Let us try to understand that. So, first I am going to do test data I am going to define a test data or dummy variable and dummy variable data dot frame and with that I will going to see matrix NA comma nrow equal to 1 ncol equal to 1 or sorry ncol equal to 2.

(Refer Slide Time: 23:50)

The screenshot displays the RStudio interface. The source editor shows the following R code:

```
54  
55 acc2=acc1  
56  
57 ## visualisation to understand the effect of feature engineering  
58  
59 test_data = dum_var = data.frame(matrix(NA,nrow=1,ncol=2))  
60 colnames(test_data)=colnames(dum_var)=c("var_wave", "skew_wave")  
61  
62  
63 var_wave=seq(-7,7,by=0.1)  
64  
65
```

The console shows the execution of the following commands:

```
> acc1=sum(diag(conf_tab1))/nrow(df_test)*100  
> acc2=sum(diag(conf_tab2))/nrow(df_test)*100  
> acc2-acc1  
[1] 2.43309  
> test_data = dum_var = data.frame(matrix(NA,nrow=1,ncol=2))  
> View(test_data)  
> View(dum_var)  
> colnames(test_data)=colnames(dum_var)=c("var_wave", "skew_wave")  
> var_wave=seq(-7,7,by=0.1)  
>
```

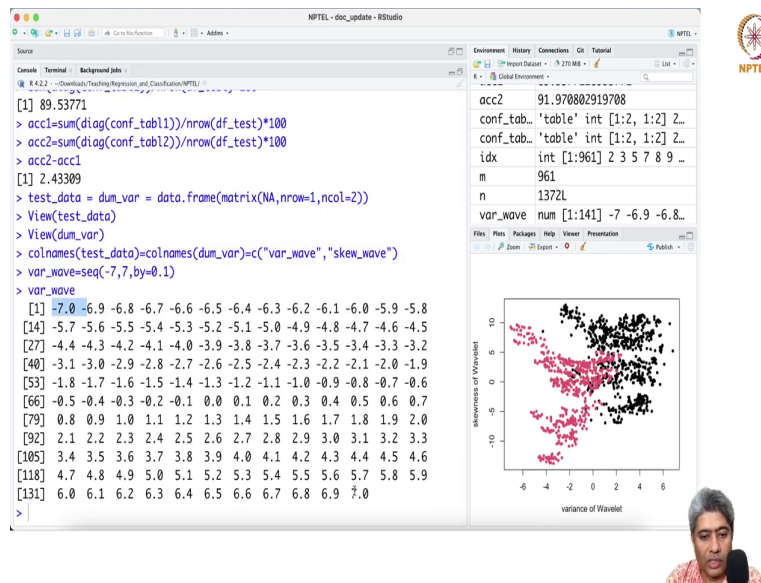
The environment pane shows the following variables:

acc2	91.970802919708
conf_tab...	'table' int [1:2, 1:2] 2..
conf_tab...	'table' int [1:2, 1:2] 2..
idx	int [1:961] 2 3 5 7 8 9 ..
m	961
n	1372L
var_wave	num [1:141] -7 -6.9 -6.8..

The scatter plot shows 'skewness of Wavelet' on the y-axis (ranging from -10 to 10) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The plot contains two clusters of points: one in pink and one in black.

So, test data and this is the I am just creating a simply place holder, ok. Just a place holder and then I am going to add a colnames of a test data equal to colnames equal to dummy variable equals to this 2, ok. These are the then variance of wave, alright. So, variance now what I am going to do? So, variance of wavelet as a range of between see it is between minus 7 and 7 and skewness of wavelet is somewhere between maybe minus 12 to 12, ok. So, I am going to create this variance of wave define minus 7 to 7 by 0.1.

(Refer Slide Time: 25:12)



The screenshot shows an RStudio session with the following R code and output:

```
[1] 89.53771
> acc1=sum(diag(conf_tab1))/nrow(df_test)*100
> acc2=sum(diag(conf_tab2))/nrow(df_test)*100
> acc2-acc1
[1] 2.43309
> test_data = dum_var = data.frame(matrix(NA,nrow=1,ncol=2))
> View(test_data)
> View(dum_var)
> colnames(test_data)=colnames(dum_var)=c("var_wave", "skew_wave")
> var_wave=seq(-7,7,by=0.1)
> var_wave
[1] -7.0 -6.9 -6.8 -6.7 -6.6 -6.5 -6.4 -6.3 -6.2 -6.1 -6.0 -5.9 -5.8
[14] -5.7 -5.6 -5.5 -5.4 -5.3 -5.2 -5.1 -5.0 -4.9 -4.8 -4.7 -4.6 -4.5
[27] -4.4 -4.3 -4.2 -4.1 -4.0 -3.9 -3.8 -3.7 -3.6 -3.5 -3.4 -3.3 -3.2
[40] -3.1 -3.0 -2.9 -2.8 -2.7 -2.6 -2.5 -2.4 -2.3 -2.2 -2.1 -2.0 -1.9
[53] -1.8 -1.7 -1.6 -1.5 -1.4 -1.3 -1.2 -1.1 -1.0 -0.9 -0.8 -0.7 -0.6
[66] -0.5 -0.4 -0.3 -0.2 -0.1 0.0 0.1 0.2 0.3 0.4 0.5 0.6 0.7
[79] 0.8 0.9 1.0 1.1 1.2 1.3 1.4 1.5 1.6 1.7 1.8 1.9 2.0
[92] 2.1 2.2 2.3 2.4 2.5 2.6 2.7 2.8 2.9 3.0 3.1 3.2 3.3
[105] 3.4 3.5 3.6 3.7 3.8 3.9 4.0 4.1 4.2 4.3 4.4 4.5 4.6
[118] 4.7 4.8 4.9 5.0 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9
[131] 6.0 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 6.9 7.0
>
```

The environment pane shows the following variables:

```
acc2      91.970802919708
conf_tab... 'table' int [1:2, 1:2] 2...
conf_tab... 'table' int [1:2, 1:2] 2...
idx       int [1:961] 2 3 5 7 8 9 ...
m         961
n         1372L
var_wave  num [1:141] -7 -6.9 -6.8...
```

The plot shows 'skewness of Wavelet' on the y-axis (ranging from -10 to 10) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The data points are colored in a gradient from pink to black, showing a clear upward trend where higher variance corresponds to higher skewness.

So, basically if I just do that now you can see that it just creating values between minus 7 to 7 with a difference of 0.1 ok, it just created these values.

(Refer Slide Time: 25:28)

The screenshot shows the RStudio interface with the following code in the editor:

```
54  
55 acc2-acc1  
56  
57 ## visualisation to understand the effect of feature engineering  
58  
59 test_data = dum_var = data.frame(matrix(NA,nrow=1,ncol=2))  
60 colnames(test_data)=colnames(dum_var)=c("var_wave", "skew_wave")  
61  
62  
63 var_wave=seq(-7,7,by=0.1)  
64 skew_wave=seq(-11,11,by=0.1)  
65
```

The console shows the execution of the last line of code:

```
> skew_wave=seq(-11,11,by=0.1)  
>
```

The Environment pane on the right shows the following objects:

conf_tab_..	'table' int [1:2, 1:2] 2..
conf_tab_..	'table' int [1:2, 1:2] 2..
idx	int [1:961] 2 3 5 7 8 9 ..
m	961
n	1372L
skew_wave	num [1:221] -11 -10.9 -1..
var_wave	num [1:141] -7 -6.9 -6.8..

The plot on the right is a scatter plot with 'variance of Wavelet' on the x-axis (ranging from -6 to 6) and 'skewness of Wavelet' on the y-axis (ranging from -10 to 10). The data points are colored in black and pink, showing a clear separation between the two groups.

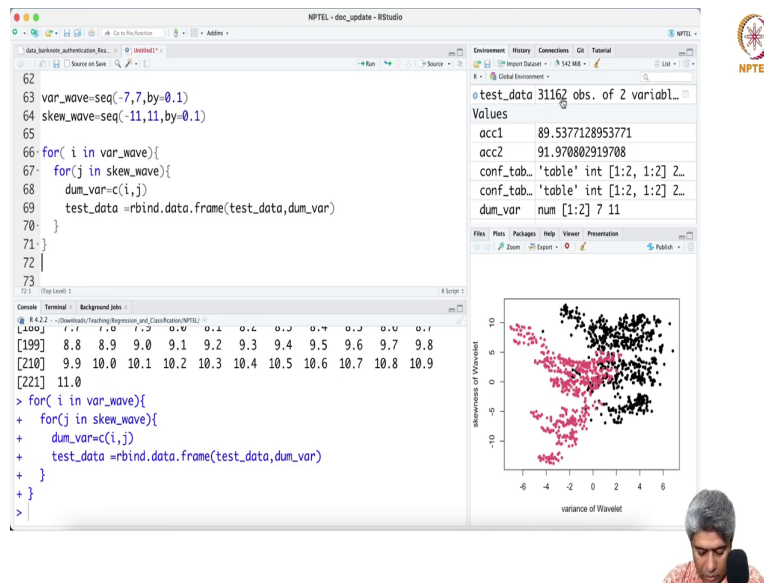


(Refer Slide Time: 25:44)

The screenshot shows the RStudio interface. The console on the left displays a grid of values for the variable 'skew_wave' across 221 iterations, ranging from -11.0 to 11.0 in increments of 0.1. The environment pane on the right shows the current workspace with variables like 'conf_tab...', 'idx', 'm', 'n', 'skew_wave', and 'var_wave'. A scatter plot in the bottom right pane shows 'skewness of Wavelet' on the y-axis (ranging from -10 to 10) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The plot contains two clusters of points: one in pink and one in black. The NPTEL logo is visible in the top right corner of the RStudio window.

Similarly, if I just say skewness of wave is going to take sequence of value between minus 11 to 11 by 0.1 difference this is also going to create values between minus 11 and 11 with a difference of 0.1, ok equal distance of 0.1. So, I am getting grid I just created grid.

(Refer Slide Time: 26:01)



The screenshot displays the RStudio interface. The script editor on the left contains the following R code:

```
62
63 var_wave=seq(-7,7,by=0.1)
64 skew_wave=seq(-11,11,by=0.1)
65
66 for( i in var_wave){
67   for(j in skew_wave){
68     dum_var=c(i,j)
69     test_data=rbind.data.frame(test_data,dum_var)
70   }
71 }
72 [
73
74 ]
```

The console on the bottom left shows the execution of the code, including the creation of the joint grid:

```
> for( i in var_wave){
+   for(j in skew_wave){
+     dum_var=c(i,j)
+     test_data=rbind.data.frame(test_data,dum_var)
+   }
+ }
```

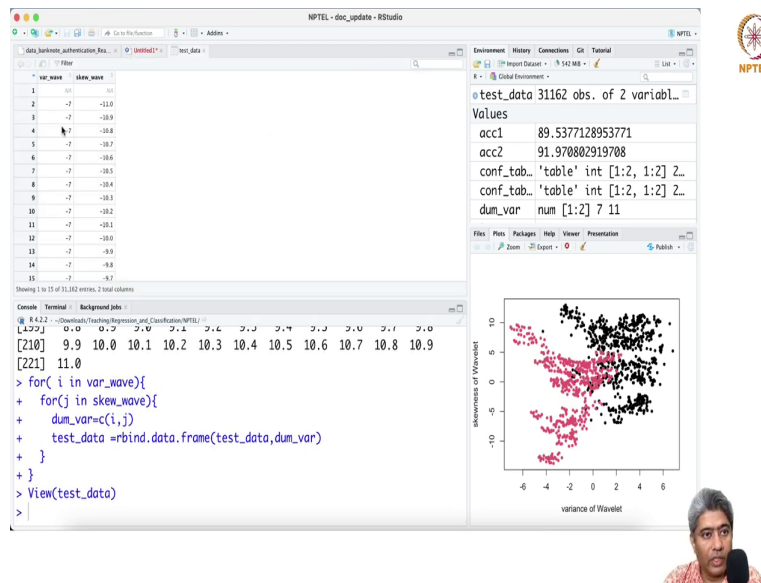
The Environment pane on the right shows the variable `test_data` with 31162 observations of 2 variables. The Values pane displays the following data:

Variable	Value
acc1	89.5377128953771
acc2	91.970802919708
conf_tab_	'table' int [1:2, 1:2] 2_
conf_tab_	'table' int [1:2, 1:2] 2_
dum_var	num [1:2] 7 11

The plot on the right is a scatter plot with 'variance of Wavelet' on the x-axis (ranging from -6 to 6) and 'skewness of Wavelet' on the y-axis (ranging from -10 to 10). The data points are colored in black and pink, forming a dense, roughly circular cluster centered around the origin.

Now, I am going to create a joint grid; joint grid how that is how. So, for i in 1 i in variance of wave for j in skewness of wave and then dum variable equal to c of whatever i and j values are there in the thing and then what I am going to do? Test data equal to rbind data dot frame test data comma dummy variable, alright.

(Refer Slide Time: 27:13)

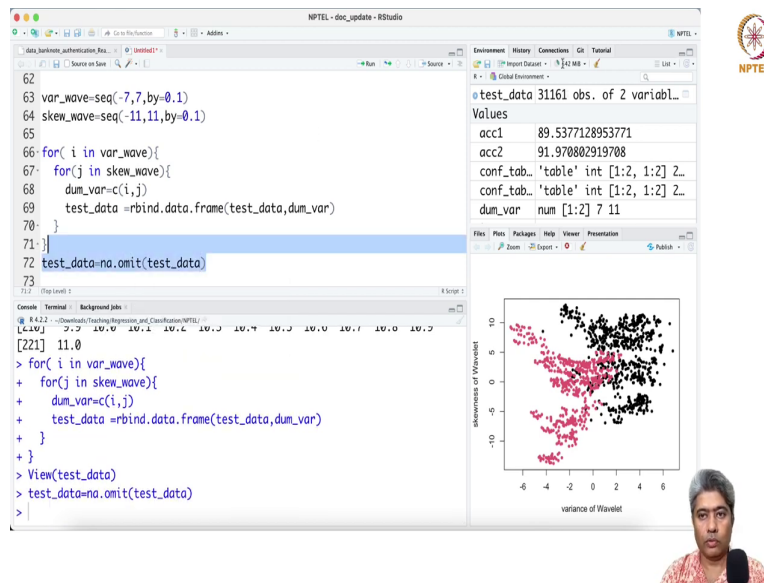


The screenshot displays the RStudio interface. The top-left pane shows a data frame with columns 'var_wave' and 'skew_wave'. The top-right pane shows the environment with 'test_data' containing 31162 observations of 2 variables. The bottom-right pane shows a scatter plot of 'skewness of Wavelet' (y-axis, ranging from -1.0 to 1.0) versus 'variance of Wavelet' (x-axis, ranging from -6 to 6). The plot shows two distinct clusters of points: one in red and one in black. The console window shows the following R code:

```
> for(i in var_wave){  
+   for(j in skew_wave){  
+     dum_var=c(i,j)  
+     test_data =rbind.data.frame(test_data,dum_var)  
+   }  
+ }  
> View(test_data)  
>
```

So, let me just run it, ok. So, it has created test data with 31,162 observations if I just create that you will see that it has created all a complete grid of two values.

(Refer Slide Time: 27:27)



The screenshot shows the RStudio interface. The main editor contains the following R code:

```
62
63 var_wave=seq(-7,7,by=0.1)
64 skew_wave=seq(-11,11,by=0.1)
65
66 for( i in var_wave){
67   for(j in skew_wave){
68     dum_var=c(i,j)
69     test_data =rbind.data.frame(test_data,dum_var)
70   }
71 }
72 test_data=na.omit(test_data)
73
```

The console shows the execution of the code, with the following output:

```
[221] 11.0
> for( i in var_wave){
+   for(j in skew_wave){
+     dum_var=c(i,j)
+     test_data =rbind.data.frame(test_data,dum_var)
+   }
+ }
> View(test_data)
> test_data=na.omit(test_data)
>
```

The Environment pane shows the following variables:

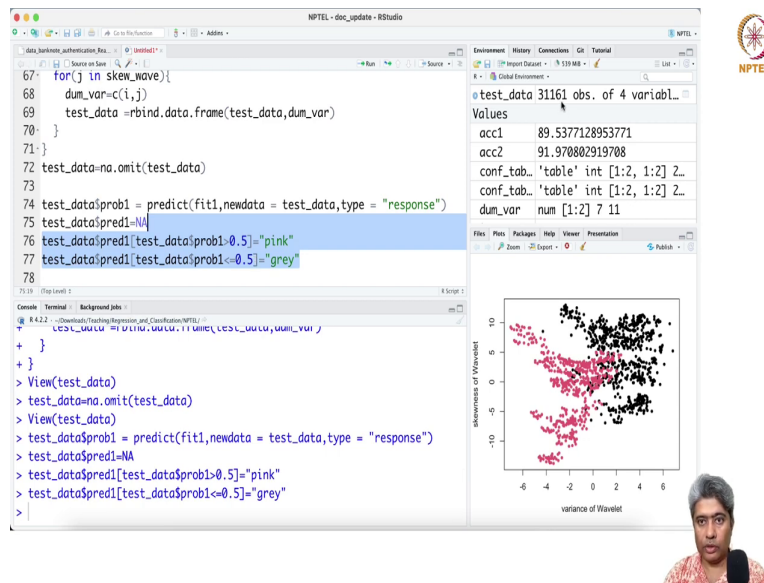
Variable	Value
acc1	89.5377128953771
acc2	91.970802919708
conf_tab_	'table' int [1:2, 1:2] 2_
conf_tab_	'table' int [1:2, 1:2] 2_
dum_var	num [1:2] 7 11

The plot pane shows a scatter plot of 'skewness of Wavelet' (y-axis, ranging from -10 to 10) versus 'variance of Wavelet' (x-axis, ranging from -6 to 6). The plot displays two clusters of points: one in black and one in pink.

The NPTEL logo is visible in the top right corner of the RStudio window.

And now what I am going to do I am going to create a let me just put a na dot omit dot omit. So, there will be the first row will be deleted and now you can see all the values are being created.

(Refer Slide Time: 27:48)



The screenshot displays the RStudio interface. The script editor on the left contains the following R code:

```
67: for(j in skew_wave){
68:   dum_var=c(i,j)
69:   test_data = rbind.data.frame(test_data,dum_var)
70: }
71: }
72: test_data=na.omit(test_data)
73:
74: test_data$prob1 = predict(fit1,newdata = test_data,type = "response")
75: test_data$pred1=NA
76: test_data$pred1[test_data$prob1>0.5]="pink"
77: test_data$pred1[test_data$prob1<=0.5]="grey"
78:
```

The console on the bottom left shows the execution of these commands:

```
R 4.2.2 -- Command: Training Regression and Classification NPTEL
> test_data = rbind.data.frame(test_data,dum_var)
+ }
+ }
+ }
> View(test_data)
> test_data=na.omit(test_data)
> View(test_data)
> test_data$prob1 = predict(fit1,newdata = test_data,type = "response")
> test_data$pred1=NA
> test_data$pred1[test_data$prob1>0.5]="pink"
> test_data$pred1[test_data$prob1<=0.5]="grey"
>
```

The environment pane on the right shows the following variables:

Variable	Class	Length
acc1	num	1
acc2	num	1
conf_tab_	'table' int	[1:2, 1:2] 2_
conf_tab_	'table' int	[1:2, 1:2] 2_
dum_var	num	[1:2] 7 11

The plot window on the right shows a scatter plot of 'skewness of Wavelet' (y-axis, ranging from -10 to 10) versus 'variance of Wavelet' (x-axis, ranging from -6 to 6). The data points are colored based on the prediction: pink points represent 'forges' (probability > 0.5) and grey points represent 'not forges' (probability ≤ 0.5). The plot shows a clear separation between the two groups, with pink points clustered in the lower-left and lower-right regions, and grey points clustered in the upper-left and upper-right regions.

Now, essentially what I am going to do is I am going to create for these all possible values I am going to create probability1 equal to predict fit from the 1st model newdata equal to test data and type equal to “response”, ok.

Now, what I am going to do is test data from the test data I am going to do a prediction and here is the whole thing the prediction is basically probability1 if it is greater than probability1 greater than 0.5 then the prediction will be “pink” basically it is a forged; that means, its forged, ok and if it is less than equal to 5 we will use grey; that means, it is not “grey” color I am just going to make this things, ok. Now, if you just go there you see pink and grey.

(Refer Slide Time: 29:37)

The screenshot shows the RStudio interface with the following components:

- Environment:** Shows 'test_data' with 31161 observations and 4 variables.
- Values:** Lists variables: acc1 (89.537128953771), acc2 (91.970802919708), conf_tab_ (table' int [1:2, 1:2] 2...), and dum_var (num [1:2] 7 11).
- Code Editor:** Contains R code for data viewing and prediction:

```
> View(test_data)
> test_data-na.omit(test_data)
> View(test_data)
> test_data$probi = predict(fit1,newdata = test_data,type = "response")
> test_data$pred1=NA
> test_data$pred1[test_data$probi>0.5]="pink"
> test_data$pred1[test_data$probi<=0.5]="grey"
> View(test_data)
```
- Console:** Shows the execution of the above code.
- Scatter Plot:** A plot of 'skewness of Wavelet' vs 'variance of Wavelet'. The y-axis ranges from -1.0 to 1.0, and the x-axis ranges from -6 to 6. Data points are colored pink and grey, showing a clear separation between the two groups.



(Refer Slide Time: 29:44)

The screenshot shows an RStudio session with the following code in the editor:

```
75 test_data$pred1=NA
76 test_data$pred1[test_data$prob1<=0.5]="pink"
77 test_data$pred1[test_data$prob1<0.5]="grey"
78
79
80 plot(test_data$var_wave, test_data$skew_wave
81       ,pch=20, col=test_data$pred1
82       ,xlab="variance of Wavelet"
83       ,ylab="skewness of Wavelet")
84 points(df_test$var_wave, df_test$skew_wave
85        ,pch=20
86        ,col=(df_test$forged_or_not+1))
```

The console shows the execution of these commands. The environment pane on the right displays the structure of the 'test_data' object:


```
test_data 31161 obs. of 4 variables
Values
acc1      89.5377128953771
acc2     91.970802919708
conf_tab_ 'table' int [1:2, 1:2] 2...
conf_tab_ 'table' int [1:2, 1:2] 2...
dum_var   num [1:2] 7 11
```

The plot shows 'skewness of Wavelet' on the y-axis (ranging from -10 to 10) and 'variance of Wavelet' on the x-axis (ranging from -8 to 6). Points are colored pink or grey based on the 'pred1' variable. A pink region is visible on the left side of the plot, and a grey region is on the right.

So, now what I am going to do? I am going to make a plot I am going to make a plot that test data dollar variance of wave test datadollarskewness of wave and then pch equal to 20 color equal to test data dollar my prediction color whatever the prediction color. Now from the above plot so, I can I will just take the lab labels xlabel and ylabel, ok.

So, that is the my predicted area. So, later the one which so, any point that will fall in this area will be called forged note and any points or any notes whose variance of wavelet and the skewness of the wavelet falls in this area will be called not forged or original note, ok. And then what I will do? I will I can just put the points essentially df testdollarvariance of wave comma df testdollarskewness of wave pch equal to 20 and color equal to df test dollar the forged or not plus 1.

(Refer Slide Time: 32:04)



So, I have to just take this and then alright. So, if I just run this. So, now, if you see so, these are the points. So, let me just. So, clearly you can yeah I think this is now probably better right. Now, this looks good.

So, what we are seeing here that these are the points you can see there are some misclassification point the points which were definitely where will be misclassified and then there are like you know forged note, but they will be called as a original note and then there will be notes which are original, but they will be called as forged.

So, there is a effectively bit of a difficult this is sort of a difficult zone, ok. This is definitely a bit of a difficult zone, alright. Now, what I am going to do I am going to produce the same

plot, but for the second model remember that second model were second model were developed using the engineered features with quadratic terms, alright.

(Refer Slide Time: 33:44)

The screenshot shows an RStudio session with the following code in the script editor:

```
84 points(df_test$var_wave,df_test$skew_wave
85 ,pch=20
86 ,col=(df_test$forged_or_not+1))
87
88
89
90 test_data$prob2 = predict(fit2,newdata = test_data,type = "response")
91 test_data$pred2=NA
92 test_data$pred2[test_data$prob2>0.5]="pink"
93 test_data$pred2[test_data$prob2<=0.5]="grey"
94
95
```

The console shows the execution of the following commands:

```
> points(df_test$var_wave,df_test$skew_wave
+ ,pch=20
+ ,col=(df_test$forged_or_not+1))
> test_data$prob2 = predict(fit2,newdata = test_data,type = "response")
> test_data$pred2=NA
> test_data$pred2[test_data$prob2>0.5]="pink"
> test_data$pred2[test_data$prob2<=0.5]="grey"
>
```

The environment pane shows the following variables:

Variable	Class	Dimensions
acc1	num	1x1
acc2	num	1x1
conf_tab_	'table'	int [1:2, 1:2] 2_
conf_tab_	'table'	int [1:2, 1:2] 2_
dum_var	num	[1:2] 7 11

The plot shows 'skewness of Wavelet' on the y-axis (ranging from -10 to 10) and 'variance of Wavelet' on the x-axis (ranging from -6 to 6). The plot contains two sets of points: pink points (representing probability2 > 0.5) and grey points (representing probability2 <= 0.5). A decision boundary is visible, separating the pink and grey regions.

So, I am going to copy this thing. So, probability2 with second model probability2 probability2 prediction2 probability2 prediction2 probability2 let me just run this.

(Refer Slide Time: 34:01)

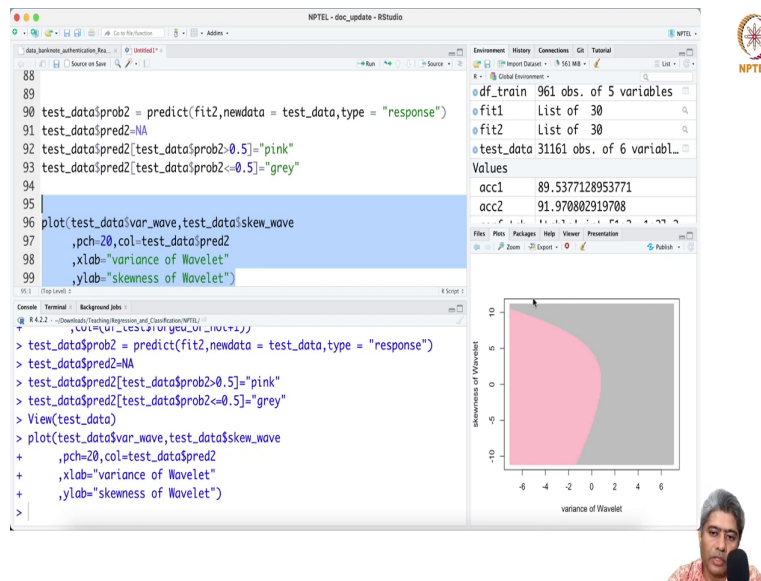
The screenshot displays the RStudio interface. The top-left pane shows a data table with columns: var_wave, skew_wave, prob1, prob2, and pred2. The top-right pane shows environment variables: df_train (961 obs. of 5 variables), fit1 (list of 30), fit2 (list of 30), and test_data (31161 obs. of 6 variables). The bottom-right pane shows accuracy metrics: acc1 (89.3377128953771) and acc2 (91.970802919708). The bottom-left pane shows the console with the following code:

```
R 4.2.2 ~ RStudio: Training Regression and Classification with Wavelet  
+ ,xlab="variance of Wavelet"  
+ ,ylab="skewness of Wavelet")  
> points(df_test$var_wave, df_test$skew_wave  
+ ,pch=20  
+ ,col=(df_test$forged_or_not+1))  
> test_data$pred2 = predict(fit2, newdata = test_data, type = "response")  
> test_data$pred2=NA  
> test_data$pred2[test_data$prob2>0.5]="pink"  
> test_data$pred2[test_data$prob2<=0.5]="grey"  
> View(test_data)  
>
```

The bottom-right plot is a scatter plot with 'variance of Wavelet' on the x-axis (ranging from -6 to 6) and 'skewness of Wavelet' on the y-axis (ranging from -1.0 to 1.0). The plot shows two distinct clusters of points: pink points (representing forged data) and grey points (representing non-forged data), separated by a diagonal decision boundary.

So, now if I have the test data, I have now got the second probability and their corresponding predictions as well.

(Refer Slide Time: 34:27)



The screenshot shows the RStudio interface with the following code in the editor:

```
88
89
90 test_data$pred2 = predict(fit2, newdata = test_data, type = "response")
91 test_data$pred2[NA]
92 test_data$pred2[test_data$prob2 > 0.5] = "pink"
93 test_data$pred2[test_data$prob2 <= 0.5] = "grey"
94
95
96 plot(test_data$var_wave, test_data$skew_wave
97       , pch=20, col=test_data$pred2
98       , xlab="variance of Wavelet"
99       , ylab="skewness of Wavelet")
```

The console shows the execution of the code:

```
> test_data$pred2 = predict(fit2, newdata = test_data, type = "response")
> test_data$pred2[NA]
> test_data$pred2[test_data$prob2 > 0.5] = "pink"
> test_data$pred2[test_data$prob2 <= 0.5] = "grey"
> View(test_data)
> plot(test_data$var_wave, test_data$skew_wave
+       , pch=20, col=test_data$pred2
+       , xlab="variance of Wavelet"
+       , ylab="skewness of Wavelet")
>
```

The Environment pane shows the following objects:

- df_train: 961 obs. of 5 variables
- fit1: List of 30
- fit2: List of 30
- test_data: 31161 obs. of 6 variables

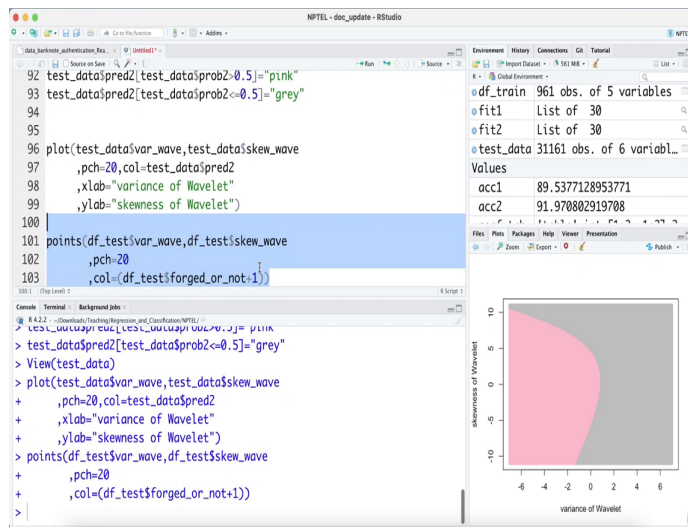
The Values pane shows the following values:

acc1	acc2
89.5377128953771	91.970802919708

The plot shows the skewness of Wavelet on the y-axis (ranging from -0.5 to 0.5) and the variance of Wavelet on the x-axis (ranging from -6 to 6). The data points are colored pink and grey, forming a quadratic geometry.

So, now if I just run this, but the color will be instead of prediction1 the prediction2. Now if I just do that. So, now, you can see this is a completely different geometry a quadratic geometry sort of coming up nicely, ok sort of a ellipse or kind of thing behavior is coming up, which is a very interesting and now if we just put the points here, right.

(Refer Slide Time: 35:00)



The screenshot shows the RStudio interface with the following R code in the script editor:

```
92 test_data$pred2[test_data$prob2<=0.5]="pink"  
93 test_data$pred2[test_data$prob2<=0.5]="grey"  
94  
95  
96 plot(test_data$var_wave, test_data$skew_wave  
97       ,pch=20, col=test_data$pred2  
98       ,xlab="variance of Wavelet"  
99       ,ylab="skewness of Wavelet")  
100  
101 points(df_test$var_wave, df_test$skew_wave  
102        ,pch=20  
103        ,col=(df_test$forged_or_not+1))
```

The console shows the execution of the code:

```
4.4.2 - Overview: Training Regression and Classification Models  
> test_data$pred2[test_data$prob2<=0.5]="pink"  
> test_data$pred2[test_data$prob2<=0.5]="grey"  
> View(test_data)  
> plot(test_data$var_wave, test_data$skew_wave  
+       ,pch=20, col=test_data$pred2  
+       ,xlab="variance of Wavelet"  
+       ,ylab="skewness of Wavelet")  
> points(df_test$var_wave, df_test$skew_wave  
+       ,pch=20  
+       ,col=(df_test$forged_or_not+1))  
>
```

The Environment pane on the right shows the following objects:

- df_train: 961 obs. of 5 variables
- fit1: List of 30
- fit2: List of 30
- test_data: 31161 obs. of 6 variables

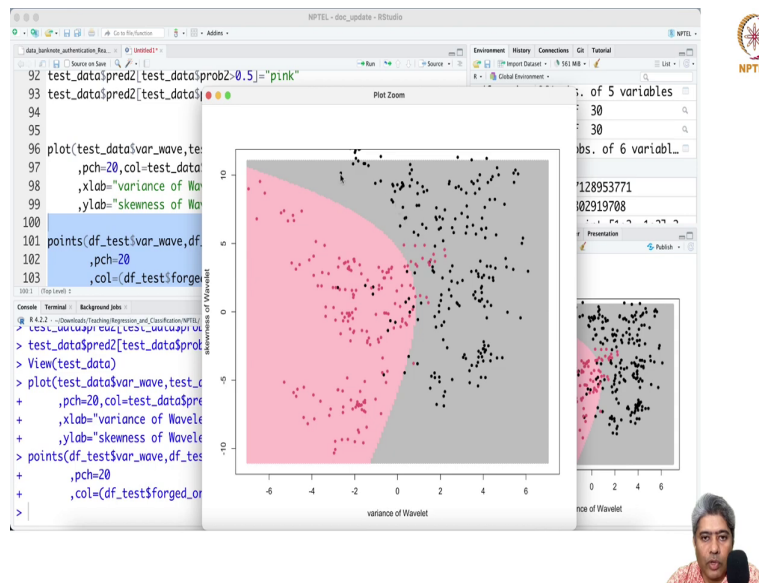
The Values pane shows the following data:

acc1	acc2
89.5377128953771	91.970802919708

The plot shows the skewness of Wavelet (y-axis, ranging from -10 to 10) versus the variance of Wavelet (x-axis, ranging from -6 to 6). The plot is divided into two regions: a pink region for $\text{prob2} \leq 0.5$ and a grey region for $\text{prob2} > 0.5$. The pink region is on the left, and the grey region is on the right.

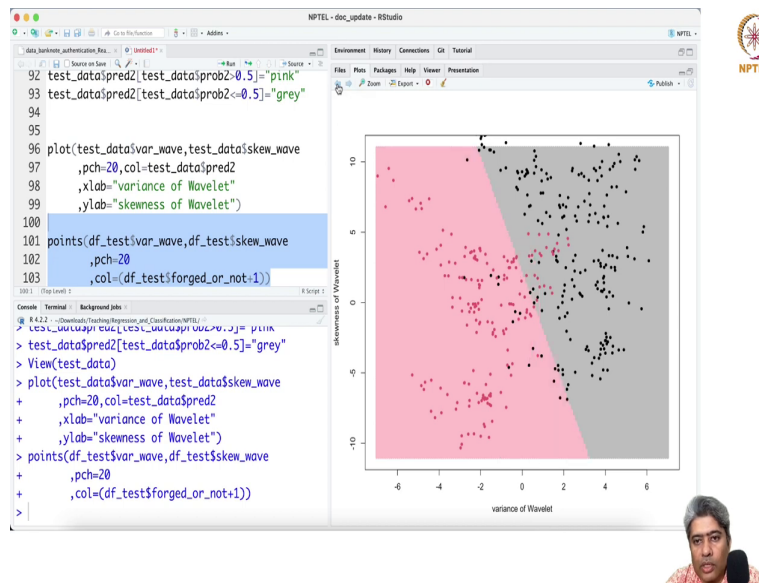


(Refer Slide Time: 35:06)



So, put the points here then we can see that you know there is a it is trying to capture these points whereas. So, let me just you know little bit increase this guy. So, it will be let, yeah.

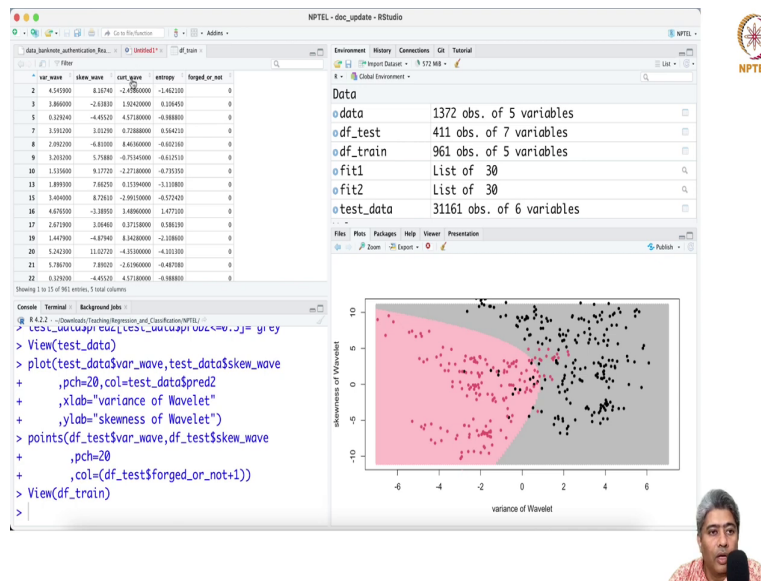
(Refer Slide Time: 35:30)



So, this is the quadratic behavior that we are seeing here and whereas, if you just this is the simple without any feature engineering, we will get a simple linear you know discriminator whereas, here we are getting a sort of a nice non-parametric quadratic discriminator. Here you can see these points are getting misclassified right, these points are getting misclassified whereas, here they are not getting misclassified this point is getting misclassified.

But it is better than too many and similarly I think if this point is getting misclassified here whereas, sorry about that, this is better this point are not getting misclassified. So, few points which are nicely not getting misclassified and trying to capture this and overall amount of the sample accuracy is going up.

(Refer Slide Time: 36:51)





The screenshot displays the RStudio interface. The top-left pane shows a data frame with columns: `skew_wave`, `var_wave`, `curt_wave`, `entropy`, and `forged_or_not`. The top-right pane shows the 'Data' tab with the following summary:

- `data`: 1372 obs. of 5 variables
- `df_test`: 411 obs. of 7 variables
- `df_train`: 961 obs. of 5 variables
- `fit1`: List of 30
- `fit2`: List of 30
- `test_data`: 31161 obs. of 6 variables

The bottom-left pane shows the R console with the following code:

```
R 4.2.2 - Cheatham Training Repository and Classification NPTEL  
> View(test_data)  
> plot(test_data$var_wave, test_data$skew_wave  
+       , pch=20, col=test_data$pred2  
+       , xlab="variance of Wavelet"  
+       , ylab="skewness of Wavelet")  
> points(df_test$var_wave, df_test$skew_wave  
+        , pch=20  
+        , col=(df_test$forged_or_not+1))  
> View(df_train)
```

The bottom-right pane shows a scatter plot of 'skewness of Wavelet' (y-axis, ranging from -10 to 10) versus 'variance of Wavelet' (x-axis, ranging from -6 to 6). The plot contains two sets of data points: pink points (representing the test set) and black points (representing the training set). A pink shaded region is visible on the left side of the plot, corresponding to the 'pred2' variable in the code.



So, feature engineering typically helps remember that I have not tried the cubic and the other higher order transformations, in addition we have few more data for example, we have curtosis and entropy, which I do not have I have not used. So, my recommendation is you guys use that and check whether it helps you to improve the accuracy.

So, I will stop here and in the next video, I will try to do predictive analysis with a new data set with a new problem. I hope you are enjoying these hands-on real data analysis. Take care, bye.