R Programming

The wizard and the computer
An introduction to preprocessing corpora using R

Fabian Tomaschek
University of Tübingen, Germany

Corresponding author:

Fabian Tomaschek
Seminar für Sprachwissenschaft
Eberhard Karls University Tübingen
Wilhelmstrasse 19
Tübingen
e-mail: fabian.tomaschek@uni-tuebingen.de
Version: November 20, 2020

Abstract

The present text addresses readers without any programming skills or experience. It shows in (hopefully) comprehensible and short steps how to use the programming language R. The main focus of this text is to introduce basic functions as well as basic knowledge how to perform your programming task which you will need to preprocess corpora. Later in the text I also introduce how to simulate learning using an error-driven learning rule.

Disclaimer

This manuscript is work in progress. It is based on the knowledge and skills I acquired on my own during my academic journey. I am a phonetician turned psycho-linguist, not an educated programmer. This is why it is very probable that this manuscript is full of typos, errors and other problems. Things will go wrong. Possibly because you messed up, and most probably, because I messed up. If you find problems, I will be absolutely grateful, if you mark them in your PDF and send them to me. Of course, I will acknowledge this just below! Professional programmers will probably frown upon my approach. Well, that's life.

Acknowledgements to helping hands and minds

Knowledge is always a joint effort and we rarely are aware of where this knowledge comes from. Nevertheless, I would like to mention and thank some colleagues who, on my journey so far, have influenced my thoughts and my perspectives. In alphabetical order, these are Denis Arnold, R. Harald Baayen, Karen Beaman, Maja Linke, Jessie Nixon, Michael Ramscar, Jacolien van Rij, Motoki Saito, Tino Sering, Ben Tucker, Martijn Wieling. I have probably forgotten many more. If you are part of this group of people, please don't feel offended but simply let me know.

Acknowledgements to the readers

Many thanks go to Ryan Podlubny, Anna Jespersen, Mai Mayeg, Sandra Zorychta and Miriam Zwick, not only for working through the script but also for pointing me to typos, errors and other mistakes.

Contents

1	Why we are here.	
	Or rather, why you are here	7
1.1	How you probably ended up reading this manuscript	7
1.2	Why linguistic corpora	7
1.3	What is R?	7
1.4	The structure of this manuscript	8
1.5	Quantitative what?	8
1.6	Your task	9
1.7	How to work with this manuscript	9
1.8	Downloading the necessary data	9
1.9	Referencing this work	9
2	Let's start	11
2.1	First day at work	11
2.2	Assessing the work load	11
2.3	Excursion : Environment and Variables	13
2.4	R as a calculator	13
2.5	Assigning results to variables	15
2.6	Excursion : Saving data	17
3	Reading text files	19
3.1	Excursion : Setting the working directory	19
3.2	Excursion : Variable classes	19
3.3	Filenames	21
3.4	Excursion : Accessing vectors	23
3.5	Reading tables	23
3.6	Excursion : Creating own dataframes	26
3.7	Excursion : Creating sequences and repetitions	28
4	Extracting values from a dataframe	31
4.1	Excursion : How to design your programming task?	31
4.2	Assessing the task	32
4.3	Preparing our programing task	33
4.4	Excursion : Searching, finding, etc.	35
4.4.1	The == operator	35
4.4.2	The %in% operator	37
4.4.3	The grepl() function	38
4.4.4	The grep() function	39

4.4.5	The <code>which()</code> function	40
4.4.6	Inversing searching results	40
4.5	Extracting the appropriate values	41
5	Creating a large dataframe	47
5.1	Repetition	47
5.2	Excursion : Repetition in programming	47
5.3	Reading multiple files	49
5.4	Excursion : Binding dataframes	49
5.5	There is still a little problem to be solved	54
5.6	Finally we create a large dataframe	54
6	Adding speaker information to the dataframe	57
6.1	Recall your task	57
6.2	Adding speaker information	58
6.2.1	Checking the file	58
6.2.2	Reading the file	58
6.2.3	Merging information	61
6.3	Excursion : Debugging, part one	70
7	Adding more speaker information to the dataframe: Transforming information in columns	73
7.1	Problem 1: Deleting unnecessary characters	74
7.2	Problem 2: Replacing information	74
7.3	Problem 3: Extracting information out of strings	75
7.3.1	Possibility one: Splitting the string	75
7.3.2	Excursion : What are lists?	77
7.3.3	Splitting strings. And now the conclusion, part 1	79
7.3.4	Possibility two: Clipping or Trimming	82
7.3.5	Excursion : Installing and loading packages	82
7.3.6	Possibility two: Clipping or Trimming, the conclusion	83
7.3.7	Well now, how old are the speakers?	86
8	Adding lexical information: Counts	87
8.1	Word frequencies	88
8.1.1	Where do I get them from?	88
8.1.2	Count words in KEC	88
8.1.3	Excursion : Testing conditions using <code>if()</code>	90
8.1.4	Now the conclusion: Iterating across <code>W.count</code> with status information	93
9	Adding lexical information: Probabilities and Conditional Probabilities	97
9.1	Excursion : Probabilities	97
9.2	Excursion : Conditional Probabilities	98
9.3	Bigram counts	99
9.3.1	Creating bigrams	100
9.3.2	Bigram counts	102
9.3.3	Conditional probabilities	104

10 Preparation for Analysis	107
10.1 Excursion : Inspecting the distribution of your data	107
10.2 Excursion : Transformation of not normally distributed data	108
10.3 Excursion : Writing a function	112
10.4 Data distribution in our corpus	113
10.5 Dealing with problematic distributions	115
10.5.1 Log transform variables where necessary	115
10.5.2 Multimodal distribution of age	115
10.5.3 Distribution of Count.AB and Count.BC	116
10.6 Correlations between predictors	117
11 Performing pilot regressions and visualizations	121
11.1 Linear mixed-effects models	121
11.2 Visualizing linear regression	124
11.3 Well, that was it	124
12 Introduction to error-driven learning	127
12.1 Basics	128
12.1.1 Introduction	128
12.1.2 R installation	128
12.1.3 NDL installation	129
12.2 What is NDL? - Modelling learning	129
12.2.1 Modelling learning with the Rescorla-Wagner function	131
12.2.2 Condition A	131
12.2.3 Condition B	134
12.2.4 Condition C	135
12.2.5 Condition D	137
12.2.6 Condition E	140
12.2.7 Condition F	142
12.2.8 A "not so special" condition: Multiple Cue-to-Outcome connections	144
12.3 The math behind the Rescorla-Wagner learning	146
12.3.1 The Algorithm	146
12.3.2 Back to the not-so-special condition	147
12.4 Derive discriminative measures from weights	151
12.5 General description	151
12.6 Calculate discriminative measures	152
12.6.1 Code to calculate activation	152
12.6.2 Code to calculate activation diversity and prior	155

Chapter 1

Why we are here.

Or rather, why you are here

1.1 How you probably ended up reading this manuscript

Well, here we are. But why are we here? Sure, this is a deep philosophical question and many scholars and philosopher shave thought about it. But ‘why are we here’ actually boils down to: How did you end up holding this manuscript in your hands! Probably you became a student assistant at a linguistics lab and are asked to preprocess a corpus. Probably, you are working at a linguistics lab, and have to perform a corpus analysis. Probably, you want to write your PhD thesis and, well... have to query a corpus. Whatever your motivation is: if you answered one of the previous questions with yes, you are absolutely right to hold this manuscript in your hands.

1.2 Why linguistic corpora

Linguistic studies in recent years have changed from single participant studies to massive data analyses, in which the linguistic data of many many speakers is analyzed. In order to analyze such amounts of data, one needs first to preprocess it. Preprocessing data contains working steps such as reading your data from a source, adding information from other sources as well as creating subsets for later analysis. Thanks to computers these things can be easily handled by writing a script. This manuscript favors **R** for these tasks, as it enables the user to built on his or her preprocessing and carry out statistical analyses. Nevertheless, this manuscript teaches you how to generally set up and prepare your programing and preprocessing endeavors. In the end, you should have learned something about the logic of programing and how to cope with programing tasks. Hence, you should become a proficient programmer, independently of the programing language.

1.3 What is R?

R is a programing language which is used to statistically analyze experimental data. The language enables the programmer to write complex scripts in which he can repeat his/her analysis over and over again. In order to install R, you can download the latest version from <http://cran.r-project.org/>. It is supported for Windows, Mac and Linux platforms. There

are different ways to work with R. On windows platforms, starting R starts a console in which all calculations are performed. This is similar on Mac platforms. On Linux, R is run in a Terminal. In order to start it, you need to type "R" into your Terminal and hit the `jenterj` key. If you run R in a separate window, I always advise to write your analysis in a text editor, where you can store your program script. In order to run a script, you can copy+paste it into the R environment. However, you will benefit way more from typing the code by yourself. In this way, you learn "how to speak" the code, not only mimicking mine. Use a text editor which highlights R specific components such as functions, numbers, etc. On the other hand you can install **Rstudio** from <https://www.rstudio.com/>. The advantage of **Rstudio** is that it starts both, the R environment as well as an R specific editor. It furthermore supports a set of short cuts which enable a fast and easy interaction between the text editor and the R environment.

This introduction wants to teach you how to write programing scripts using R. If you have mastered this, you can go a step ahead and learn how to actually analyze data. For this, however, there have been already a lot of books written, such as Baayen (2008).

1.4 The structure of this manuscript

I usually skip this section, when reading a book. It is mostly boring because it only presents me information which I will read/experience anyways. If you happen to be such a person, too, stay here, because instead of summarizing each section at this place, I want to present our motivation for the structure of this manuscript to you. Our experience is that written introductions to programing are usually boring. And they assume that the programmer has no knowledge of programing at all. Although this is also the premise of this manuscript, I will push you right into the cold water. Why? Because I think that the presentation of pure facts does not allow you to learn to program ("*I ll, this is a variable of class X. In contrast to variables of the class Y, variables of the class X are (please fill in a word)*"). It is a well known fact that after a pure presentation, the reader/listener keeps not more than 20% of the presented knowledge. By contrast, hands-on sessions allow the reader/listener to experience the use of knowledge, to get to know "what happens if I push button X" and therefore experience the meaning of button X. In our case, buttons will be commands and functions, written in a script and parsed by the computer (hopefully correctly). Don't worry, I do not expect you to find out every detail on your own. Why should you then read this stupid introduction? I will, if I think it is necessary (which will be very often), take you on [Excursions](#) to discover the small but important details of R. If you think, I have missed something (which, again, I probably have), do not hesitate to contact me.

1.5 Quantitative what?

The present manuscript is tangent to Quantitative Linguistics or any quantitative science (Econometrics, or what else there might be. Use your favorite search engine to find more). Although I use Quantitative Linguistics as an example, I do not give here a thorough introduction to the subject and its theories. This is, really, beyond the skope of the present text.

1.6 Your task

Pushing you into the cold water, I will give you a *simple* task, like it is given on a everyday basis on a quantitative linguistics lab. For the accomplishment of your task, I will provide you with a working corpus, which is a subset of a database used in our lab. The task is: "Please preprocess the corpus". Easy, right? Well, let's get started.

1.7 How to work with this manuscript

I strongly advice you not to use this text as a bedtime reading. Rather, you should be working with R while you read this text. I present a lot of code in this manuscript. Use the code I present here. *Copy and paste* is a very powerful function which our computers provide to us. But copying and pasting does not help you in learning how to program. The better solution: copy the code by typing it yourself and see the results on your own. Read also the comments in the code, as they provide important information I often do not repeat in the text (by the way: comments in a script are marked with a hashtag #, as you will see). In addition, play with the code presented here, add and change the input presented here in order to get a sense of what you are doing. See programing as another language you learn, like French or Russian. Your communication partner is not another person, but the computer. And like all languages, programing has its own words and (syntactic) structures. Every function you will use is like a word in the real language. And unlike in a real language, function words have very often telling names, i.e. they tell you what they will do. But I will talk about this later. Finally, programing is like performing magic: you say (type) words into the command window, say abracadabra and simsalabim (or hit the enter button) and voila, the sesame opens (or the output of your process is displayed).

1.8 Downloading the necessary data

I provide you with a part of the Karl-Eberhards-Korpus (Arnold & Tomaschek, 2016), on whose annotations all the scripts in this manuscript are applied. You can download the data here:

- <https://osf.io/t5ku4/>

1.9 Referencing this work

I believe in open science and the exchange of knowledge. Part of this is bypassing big publishers. You might have noticed it, you did not have to purchase this introduction, but simply downloaded it for free. If you benefitted from this manuscript, I would be very happy if you could cite both, this manuscript and the corpus in your own word in the following way:

Tomaschek, F. (2020). *The wizard and the computer: An introduction to preprocessing corpora using R*. Retrieved from psyarxiv.com/jsv38

Denis Arnold and Fabian Tomaschek (2016): "*The Karl Eberhards Corpus of spontaneously spoken southern German in dialogues – audio and articulatory recordings*".

In: Draxler, Christoph; Kleber, Felicitas (eds.) (October 2016): Tagungsband der 12. Tagung Phonetik und Phonologie im deutschsprachigen Raum. P und P12. 12. Tagung Phonetik und Phonologie im deutschsprachigen Raum, 12. - 14. Oktober 2016, Ludwig-Maximilians-Universität München.

Chapter 2

Let's start

2.1 First day at work

Imagine the following situation: You got hired for a student assistant's job, either at the university or somewhere else. Now you are wondering: *"Why should this happen to me?"* Well, the truth is, it happens all the time. Ok, let's carry on. You get hired for your skills, which ever they might be (hopefully you have any, given the fact that you are enrolled at an university.... But, let's not get too deep into this...). The first day, you arrive at your new office and your boss says the following to you:

- Task

"Well, welcome to our lab. I'm happy that you are with us. So, I need you to preprocess this corpus of mine. You know, get everything in order, extract the usual measures. I would like you to use the great, fantastic and easy to learn statistical software R for this. Since you are new here, I'd say it will take you roughly two weeks. See you."

Well, congratulations. You probably are wondering: *"Does this situation actually ever happen?"* And my answer is: Yes, I have experienced it at least a dozen times. The problem is that R is not taught at Universities to students of Linguistics and related courses (at least at mine). Nor are students taught how to translate the wishes of their bosses and at some point of their career, their own, into actual code. When I was working on my master thesis, I had a friend who programed an entire experiment for me. Back then, he was a wizzard for me. Imagine what. Now I know how to do it myself. It only took me, well... seven years. Damn it, wrong information. No, just kidding.

2.2 Assessing the work load

Ok, before you start panicking, let's go back to your boss' assignment. He said something about getting everything in order and extracting the usual measures. Well, what does this mean. If you get an assignment and you don't understand what you should do, your first reaction should *always*, I mean like **always**, really **always** be to respond: *"Sorry, could*

you give me more details on that!". Details are the alpha and omega of every programming job because they define your plan. Before you start writing any script, make a plan, what you want to do. I don't know what kind of technique you use for creating plans and I won't teach you any of them. This is beyond the scope of my knowledge, partially because I make a plan while I go to bed and put it into practice next morning. And given my really, really not praiseworthy skills to draw, I OK, details. Let's say, I am your boss now. You should also think about precisising what kind of details you want to have. This might result in questions like the following:

- (1) What kind of corpus is this?
- (2) Where do I find it?
- (3) In what format is it stored?
- (4) Into what format should I translate the corpus?
- (5) Tell me more about 'the usual' measures? I don't know what you are talking about.

Now you probably say: *"Really, I should say, I don't know what my boss is talking about?"* Yes. In contrast to a famous movie quote, ignorance is not a bliss. Admitting that you don't know something is the first step to learn it. The following are the answers to your (possible) questions:

- (1) We start slow here. It is a corpus of spontaneously spoken language German. Two people were discussion a topic of their own interest.
- (2) I provide you with transcriptions of the discussions. Oh, wait. We already have **Annotations** for **Praat** of the transcriptions praat. I give you these to work on. They are in a standard text format.
- (3) I'll send you the directory.
- (4) Good question. You are a smart one. The final result, i.e. the result of your work, should be a **Dataframe** in R.
- (5) Well, you know: Time stamps. Word durations. Speaker information. Word frequencies. Conditional probabilities.

Unless you studied mathematics or had already an introduction to statistics, there is probably a quiet but scared voice in your head right now: *"Word frequencies? Conditional probabilities? What the hell is this?"*. Well, don't worry, we will get to that. And it is less complicated than the technical term suggests. But let us start easy here. Your first task will be to read the corpus by means of R. This will be handled in the following. But in order to do this, we need to give you a short introduction what R actually is and how to handle it. Furthermore, Praat is a very powerful software to process and analyze speech signals. We will deal with it in more detail in Section 4.2.

2.3 Excursion: Environment and Variables

Before you can start doing "the hard stuff", you need to be aware of how to handle some standard processes in the **R environment**. "Environment?" I hear the voice in your head asking. Depending on the operation system on your computer and your installation, your environment might look like that depicted in Figures 2.1 or 2.2.

Figure 2.1: *R Environment on Linux or Mac or Windows in a stand alone installation.*

Figure 2.1 shows you the R environment, the prompt, how you can see it when you install R as a stand alone installation. In this case, you need to install a good text editor which allows you to edit your scripts. There are plenty of softwares. The most important part is that your chosen one will highlight the structure of your script. This helps a lot during programming.

Figure 2.2 shows you how R looks like when you have installed Rstudio. The advantage of Rstudio, in contrast to a stand alone installation of R, is that it is equipped with a nice GUI which allows you an interaction between its inbuilt text editor and the R environment.

2.4 R as a calculator

OK, the following is rather trivial, but nevertheless important. R is an calculator. Though a complex one, still a calculator. And in the end, you will be running tons of calculations with R. So, the following code shows you the standard commands to perform calculations in R. Try the calculator out by typing the following commands into your R environment and hitting the **enter** key:

```
3+3 #addition

## [1] 6

3-3 #subraction
```

Figure 2.2: *R Environment in R studio* (<https://www.rstudio.com/wp-content/uploads/2014/04/rstudio-vcs.png>).

```
## [1] 0
```

```
3*3 #multiplication
```

```
## [1] 9
```

```
3/3 #division
```

```
## [1] 1
```

```
2^2 #power
```

```
## [1] 4
```

```
4^(1/2) #root
```

```
## [1] 2
```

```
# <- the hashtag comments sections and
# makes R ignore the code.
```

You can make the commands more complex:

```
3+3+3*2
## [1] 12

3+3+(3*2)
## [1] 12

3+(3+3)*2
## [1] 15
```

Remember that in more complex algebra, multiplication/division is applied prior to addition/subtraction. Therefore, brackets change the outcome of the calculation.

In contrast to your pocket calculator, R is able to store the results of your commands into variables. Variables are shells in which you can store almost everything you want. Imagine a situation in a kitchen. You want to bake a cake. The necessary ingredients, you need to make the dough for your cake, have to be mixed. The container for your ingredients is the bowl. In the same sense, the container for your results in R is the variable.

2.5 Assigning results to variables

```
first = 3+3+3*2
second = 3+3+(3*2)
third = 3+(3+3)*2
```

As you can see, storing the result of the command inhibits its displaying in the console. This is a big advantage. Imagine, you work with commands that produce a result with ten thousand entries. You don't want to display such a result every time. In order to see the content of a variable, you simply type the variable into your R environment and hit `enter`.

```
first
## [1] 12

second
## [1] 12

third
## [1] 15
```

A variable can of course contain more than one value (like your bowl can contain more than one egg). The number of values you can assign to a variable is limited by the working memory of your computer (**RAM** = **r**andom **a**ccess **m**emory). You should know how much RAM your computer has (if you don't know and don't know how to find it out... well, use your favorite search engine). You can compare RAM to the size of the table in your kitchen.

You always need to save space to shuffle your bowls around. If you pack it full with tons of bowls, well, working will be hard. The same with RAM. RAM and the size of your variables tremendously affect how fast your script runs through. Thus, keep an eye on the size of your variables and how large your RAM is. You can do this with a built-in function provided by R called: `object.size()` .

Built in functions are preprogramed and you can use them in order to perform some work. There are a lot of functions, which are very usefull and which you will learn while reading this manuscript. Functions are always followed by round brackets (). In order to apply the function to a given variable, you include the variable into the brackets, like in the following way:

```
object.size(first) # Raw display

## 48 bytes

print(object.size(first), unit='Mb') # Translate into different unit

## 0 Mb
```

`print()` is a very usefull function to display the content of a variable. You will see its variable use repeatedly. Now, the size of the variabe "first" is not large, but the size of variables can grow very fast. Be careful! Of course you can assign more than one value to a variable (otherwise its purpose would be very doubtful). A one-dimensional sequence of values is called a **Vector** . Assigning a vector to a variable is done by use of the function `c()` . `c()` is an abbreviation for 'concatenation'. As we mentioned above, functions have (to a certain extent) telling names.

```
fourth = c(1,2,3,4,5,6,7,8,9,10)
fourth

## [1] 1 2 3 4 5 6 7 8 9 10
```

Finally, you can apply the standart mathematical operations to an entire vector.

```
fourth+2

## [1] 3 4 5 6 7 8 9 10 11 12

fourth-2

## [1] -1 0 1 2 3 4 5 6 7 8

fourth*2

## [1] 2 4 6 8 10 12 14 16 18 20

fourth/2

## [1] 0.5 1.0 1.5 2.0 2.5 3.0 3.5 4.0 4.5 5.0
```

2.6 Excursion: Saving data

Although it is absolutely not necessary for you to save data at this point, we will introduce you to the way to do this right now. The reason for this is simple: you will progress through the script with your own pace, probably reading one chapter a day. You will want to be able to save your progress. Saving is performed by means of the `save()` function. Before you do so, you need to tell **R** where to store the file, or in other words, where your working directory is (it is, let's say, equivalent to telling your visitor, what room you want to go to before eating that delicious cake you baked in the morning). This is done with the `setwd()` function.

```
#DIR = 'YOURWORKINGDIRECTORY'
setwd(DIR) #set up your working directory
#saving one variable
save(fourth, file = 'fourth.rda')
#saving multiple variables
save(list=c('first', 'fourth'), file = 'fourth.rda')
```

Instead of providing the function with a mere filename, you can provide it also with the entire path to the directory where you want to store the file, e.g. `/home/fabian/myRdirectory/MyRdata.rda`. Remember that how exactly the path looks like depends on your system. The path shown before is a standard path for unix systems. In Windows systems the path looks something like `C:nnYOURPATHnnYOURFOLDER`. Retrieving your data is done by means of the `load()` function.

```
setwd(DIR)
#saving one item
load('fourth.rda')
#make the function show the loaded variables
load('fourth.rda', verbose=TRUE)

## Loading objects:
##   first
##   fourth
```

Again, you should give your stored data a telling name in order to remember what it contains. Everytime you will resume the study of R, you can load some of the variables we already worked on. It is up to you to find out, which ones.

Chapter 3

Reading text files

3.1 **Excursion:** Setting the working directory

Welcome back, you obviously have not thrown the manuscript into the nearest bin. Congratulations (For the true meaning of congratulations: <https://youtu.be/Is5BWcXApQ0?t=3m30s>). Now, we need to go back to our initial purpose: Preprocessing the corpus. I told you already that the format of the annotations are text files. Now, we need to read in these text files. There are plenty of possibilities to read in the text files. You will find out about them later. First, we need to concentrate on where your data is located on your computer, i.e. its **Directory** , aka **Working directory** . It is like asking your mother, father, friend, whoever, where the eggs are located for your cake. I usually check the location in my file manager and then set that directory as the working directory in R. The way this works is using the function `setwd()` in which you include the directory as a string. I usually first store the directory in a variable, and then give it to `setwd()`. You can find out your working directory using `getwd()` .

```
#DIR = 'REPLACE.THIS.STRING.WITH.YOUR.DIRECTORY.DELETE.HASHTAG'
setwd(DIR)
#getwd() #Obviously I will not show your mine.
```

3.2 **Excursion:** Variable classes

Above, you probably stumbled across the word **String** . Well, the truth is that a variable can contain different classes of values. You already learned about the most important class, the numeric class. We used in the above examples in Section 2.4. You can find out the class of the containing values using the function `class()` .

```
class(first)

## [1] "numeric"
```

`first` contains obviously numeric values. Variables can contain also strings. Strings are chains of signs (i.e. letters, numbers, and other signs) between apostrophies (") or quotation marks (").

```
first.string = 'This is a string and it contains a 4 (four) and a $ (Dollar sign)'  
class(first.string)  
## [1] "character"
```

Obviously, you can assign multiple strings to a variable.

```
second.string = c('first string', 'second string', 'third string')  
second.string  
## [1] "first string" "second string" "third string"
```

Remember that you can always assign a number in form of a string to variable. You need to know what class your values belong to, since R can calculate only with numeric values, not with **Characters**, i.e. strings. If you have a vector that contains only numbers, be its class numeric or string, you can switch between these using the functions `as.numeric()` and `as.character()`.

```
numbers = c(1,2,3,4,5)  
numbers  
## [1] 1 2 3 4 5  
  
class(numbers)  
## [1] "numeric"  
  
strings = as.character(numbers)  
strings  
## [1] "1" "2" "3" "4" "5"  
  
class(strings)  
## [1] "character"  
  
numbers.again = as.numeric(strings)  
numbers.again  
## [1] 1 2 3 4 5  
  
class(numbers.again)  
## [1] "numeric"
```

Numbers and strings can be distinguished insofar as strings are always displayed in quotation marks. If you concatenate a number (or numbers) and a string (even a single one) together and want to store it into a variable, the number will be changed per default into a string.

```
c(1,'a')

## [1] "1" "a"

c(1,2,3,4,'a')

## [1] "1" "2" "3" "4" "a"
```

3.3 Filenames

Alright, you have text data located on your computer, you know its location but now you need to read it in into the R environment. In order to read your data, you need to know the **File name** of your data. You can find out the name by checking it in your file manager. But would it not be great if there was a possibility to find the file names in R? Guess what, there is one: `list.files()` .

```
setwd(DIR)
filenames = list.files()
class(filenames)

## [1] "character"

filenames

## [1] "DF_textgrids_addedinformation2.rda"
## [2] "DF_textgrids_addedinformation3.rda"
## [3] "DF_textgrids_addedinformation4.rda"
## [4] "DF_textgrids_addedinformation.rda"
## [5] "DF_textgrids.rda"
## [6] "Distributions.pdf"
## [7] "rec_001_AS_id_002_CUT_20MIN_1.TextGrid"
## [8] "rec_001_AS_id_002_CUT_20MIN_2.TextGrid"
## [9] "rec_001_AS_id_002_CUT_20MIN_3.TextGrid"
## [10] "rec_001_IS_id_001_CUT_20MIN_1.TextGrid"
## [11] "rec_001_IS_id_001_CUT_20MIN_2.TextGrid"
## [12] "rec_001_IS_id_001_CUT_20MIN_3.TextGrid"
## [13] "rec_003_AS_id_006_CUT_10min_1.TextGrid"
## [14] "rec_003_AS_id_006_CUT_10min_2.TextGrid"
## [15] "rec_003_AS_id_006_CUT_10min_3.TextGrid"
## [16] "rec_003_AS_id_006_CUT_10min_4.TextGrid"
## [17] "rec_003_AS_id_006_CUT_10min_5.TextGrid"
## [18] "rec_003_AS_id_006_CUT_10min_6.TextGrid"
## [19] "rec_003_IS_id_005_CUT_10min_1.TextGrid"
## [20] "rec_003_IS_id_005_CUT_10min_2.TextGrid"
## [21] "rec_003_IS_id_005_CUT_10min_3.TextGrid"
## [22] "rec_003_IS_id_005_CUT_10min_4.TextGrid"
## [23] "rec_003_IS_id_005_CUT_10min_5.TextGrid"
## [24] "rec_003_IS_id_005_CUT_10min_6.TextGrid"
```

```
## [25] "rec_004_AS_id_008_CUT_10min_1.TextGrid"
## [26] "rec_004_AS_id_008_CUT_10min_2.TextGrid"
## [27] "rec_004_AS_id_008_CUT_10min_3.TextGrid"
## [28] "rec_004_AS_id_008_CUT_10min_4.TextGrid"
## [29] "rec_004_AS_id_008_CUT_10min_5.TextGrid"
## [30] "rec_004_AS_id_008_CUT_10min_6.TextGrid"
## [31] "rec_004_IS_id_007_CUT_10min_1.TextGrid"
## [32] "rec_004_IS_id_007_CUT_10min_2.TextGrid"
## [33] "rec_004_IS_id_007_CUT_10min_3.TextGrid"
## [34] "rec_004_IS_id_007_CUT_10min_4.TextGrid"
## [35] "rec_004_IS_id_007_CUT_10min_5.TextGrid"
## [36] "rec_004_IS_id_007_CUT_10min_6.TextGrid"
## [37] "rec_005_AS_id_010_CUT_10min_1.TextGrid"
## [38] "rec_005_AS_id_010_CUT_10min_2.TextGrid"
## [39] "rec_005_AS_id_010_CUT_10min_3.TextGrid"
## [40] "rec_005_AS_id_010_CUT_10min_4.TextGrid"
## [41] "rec_005_AS_id_010_CUT_10min_5.TextGrid"
## [42] "rec_005_AS_id_010_CUT_10min_6.TextGrid"
## [43] "rec_005_IS_id_009_CUT_10min_1.TextGrid"
## [44] "rec_005_IS_id_009_CUT_10min_2.TextGrid"
## [45] "rec_005_IS_id_009_CUT_10min_3.TextGrid"
## [46] "rec_005_IS_id_009_CUT_10min_4.TextGrid"
## [47] "rec_005_IS_id_009_CUT_10min_5.TextGrid"
## [48] "rec_005_IS_id_009_CUT_10min_6.TextGrid"
## [49] "rec_016_AS_id_032_1.TextGrid"
## [50] "rec_016_AS_id_032_2.TextGrid"
## [51] "rec_016_AS_id_032_3.TextGrid"
## [52] "rec_016_AS_id_032_4.TextGrid"
## [53] "rec_016_AS_id_032_5.TextGrid"
## [54] "rec_016_AS_id_032_6.TextGrid"
## [55] "rec_016_IS_id_031_1.TextGrid"
## [56] "rec_016_IS_id_031_2.TextGrid"
## [57] "rec_016_IS_id_031_3.TextGrid"
## [58] "rec_016_IS_id_031_4.TextGrid"
## [59] "rec_016_IS_id_031_5.TextGrid"
## [60] "rec_016_IS_id_031_6.TextGrid"
## [61] "SpeakerInformation.csv"
```

Well, isn't that a nice and long list. Would not it be great if we could simply check the first six entries of the list. Guess what, it is possible, using the function `head()` .

```
head(filenamees)

## [1] "DF_textgrids_addedinformation2.rda"
## [2] "DF_textgrids_addedinformation3.rda"
## [3] "DF_textgrids_addedinformation4.rda"
## [4] "DF_textgrids_addedinformation.rda"
## [5] "DF_textgrids.rda"
## [6] "Distributions.pdf"
```

We also can check out the last few lines of a vector or a dataframe by using the `tail()` .

```
tail(filenamees)

## [1] "rec_016_IS_id_031_2.TextGrid" "rec_016_IS_id_031_3.TextGrid"
## [3] "rec_016_IS_id_031_4.TextGrid" "rec_016_IS_id_031_5.TextGrid"
## [5] "rec_016_IS_id_031_6.TextGrid" "SpeakerInformation.csv"
```

3.4 Excursion: Accessing vectors

Let's simply chose the first entry of this list. But instead of copying the filename by using copy+paste, we access the vector by means of an **Index** . Indices indicate which entry of your data you want to access. Usually, indices are numeric. In order to access a vector, you type in the variable name and put square brackets `[]` after it. Accessing is of course independent of the class of values a variable contains.

```
index = 1
fn = filenamees[index]
print(fn)

## [1] "DF_textgrids_addedinformation2.rda"
```

Indices allow you not only to access ONE entry, but multiple.

```
filenamees[1] #accessing one entry directly

## [1] "DF_textgrids_addedinformation2.rda"

filenamees[c(1,3,5)] #accessing several entries

## [1] "DF_textgrids_addedinformation2.rda"
## [2] "DF_textgrids_addedinformation4.rda"
## [3] "DF_textgrids.rda"
```

Indexing a variable entry allows you to *read* its content, but also to change its content.

```
changeVector = c(1,2,3,4,5)
changeVector[3] = 99
changeVector

## [1] 1 2 99 4 5
```

3.5 Reading tables

Finally, you can read in our data. Of course there are a million ways to do this, depending on the format. But we will start with the easiest. Since your data is stored as a **text file**

, we will read it in as a **Table** . (By the way: text data can have many many endings, not only *.txt. In order to assess whether your data actually is stored as text data, simply open a file using your text editor).

Tables are instances of data sorted to rows and columns. This is done by means of the function `read.table()` and storing it into a new variable.

```
setwd(DIR)
filenames[1]

## [1] "DF_textgrids_addedinformation2.rda"

data = read.table(filenames[1])

## Warning in file(file, "rt"): cannot open file 'DF_textgrids_addedinformation2.rda':
## No such file or directory
## Error in file(file, "rt"): cannot open the connection
```

Well, there is a nice red **Error message** in the prompt (at least in the manuscript its red). We messed something up. The truth is that `read.table()` expects the number of columns to be equal in every row. We can tweak the reading by means of a **Parameter** . Parameters allow you to change the performance of a function. You can assess the possible parameters of a function (as well as its functionality) by typing a question mark ? before a function and hitting `enter` into the prompt. In this way, you access the function's **Manual** . By the way, the repeated use of `setwd()` in these examples is necessary only for the purpose of writing this manuscript. You need to set the working directory only once .

```
?read.table
```

In our case, we need to tell `read.table()` that there are columns missing in some rows and it needs to fill them with **NA** . NA represents "not available". It is an equivalent to NaN , which stands for "Not a Number". NAs are generally used to represent missing data. Importantly, you need to do this only if your dataframe contains missing columns. You can provide the filename to `read.table()` in two ways: Either you write provide the function with a proper string, as in the first example below. Or you provide it with an entry from a vector, as in the second example. The outcome is always the same. However, the second way to do this is the more flexible one, as you will see later.

```
setwd(DIR)
#use a string directly
data = read.table('rec_001_AS_id_002_CUT_20MIN_1.TextGrid', fill = NA)
#use an entry from a vector
filenames[1]

## [1] "DF_textgrids_addedinformation2.rda"

data = read.table(filenames[1], fill = NA)

## Warning in read.table(filenames[1], fill = NA): line 3 appears to contain embedded
nulls
```

```
## Warning in read.table(filenames[1], fill = NA): line 4 appears to contain embedded
nulls
## Warning in read.table(filenames[1], fill = NA): line 5 appears to contain embedded
nulls
## Warning in scan(file = file, what = what, sep = sep, quote = quote, dec = dec, :
embedded nul(s) found in input

class(data)

## [1] "data.frame"
```

The output of `read.table()` is a **Dataframe** as `class()` tells us. A dataframe is a two-dimensional variable structure with **Rows** and **Columns**. Entries in data frame run from top to bottom. The speciality of dataframes is that each column can contain values of different classes. We assess the **Variable structure** using the function `str()`.

```
class(data)

## [1] "data.frame"

str(data)

## 'data.frame': 887438 obs. of  1 variable:
 $ V1: Factor w/ 19671 levels "", "\"", "\"\"", "\"@", ...: 14977 18331 7...
```

`str()` tells us that "data" contains 887438 rows and 1 columns. Furthermore, it tells us that columns contain **Factors**. Factors are necessary for statistical analysis and we will not discuss factors until Section ??, which is why we need to make `read.table()` give us characters, setting the parameter **StringsAsFactors** to **FALSE**. Below you can see that the columns of "data" now contain character strings.

```
setwd(DIR)
data = read.table(filenames[1], fill = NA, stringsAsFactors=FALSE)

## Warning in read.table(filenames[1], fill = NA, stringsAsFactors = FALSE): line 3
appears to contain embedded nulls
## Warning in read.table(filenames[1], fill = NA, stringsAsFactors = FALSE): line 4
appears to contain embedded nulls
## Warning in read.table(filenames[1], fill = NA, stringsAsFactors = FALSE): line 5
appears to contain embedded nulls
## Warning in scan(file = file, what = what, sep = sep, quote = quote, dec = dec, :
embedded nul(s) found in input

class(data)

## [1] "data.frame"

str(data)

## 'data.frame': 887438 obs. of  1 variable:
## $ V1: chr  "RDX2" "X" "höchstens" "unmöglich" ...
```

3.6 Excursion: Creating own dataframes

As far as we know, all statistical analyses in R are performed on dataframes. So, let's create a dataframe, using the function `data.frame()`. In order to create a sequence of numbers, we use the function `seq()`.

```
DF1 = data.frame(Numbers = seq(1,10))
```

```
DF1
```

```
##      Numbers
## 1         1
## 2         2
## 3         3
## 4         4
## 5         5
## 6         6
## 7         7
## 8         8
## 9         9
## 10        10
```

You can see that the argument you assigned to the function `data.frame()` becomes the name of the first column. An important operator for the work with dataframes is the `$` (Dollar) operator. You need it, for the creation of new columns, for accessing columns in dataframes as well as the deletion of columns by assigning **NULL** to it.

Here, we attach a numeric vector as well as a string vector containing letters, created by means of the prestore variable **letters**.

```
DF1$NewNumbers = seq(11,20) #assigning new numbers and attaching columns
```

```
DF1$NewLetters = letters[1:10] #creating a column
```

```
DF1
```

```
##      Numbers NewNumbers NewLetters
## 1         1         11         a
## 2         2         12         b
## 3         3         13         c
## 4         4         14         d
## 5         5         15         e
## 6         6         16         f
## 7         7         17         g
## 8         8         18         h
## 9         9         19         i
## 10        10         20         j
```

```
str(DF1) #showing structure
```

```
## 'data.frame': 10 obs. of 3 variables:
## $ Numbers : int 1 2 3 4 5 6 7 8 9 10
## $ NewNumbers: int 11 12 13 14 15 16 17 18 19 20
## $ NewLetters: chr "a" "b" "c" "d" ...
```

```
DF1$NewLetters #accessing one column

## [1] "a" "b" "c" "d" "e" "f" "g" "h" "i" "j"

DF1$NewLetters = NULL #delet a column
DF1

##      Numbers NewNumbers
## 1         1         11
## 2         2         12
## 3         3         13
## 4         4         14
## 5         5         15
## 6         6         16
## 7         7         17
## 8         8         18
## 9         9         19
## 10        10         20
```

Of course, a dataframe can contain factors (which is, by the way, very important. Otherwise ANOVAS, linear models, linear mixed-effect regression, generalized additive models and so forth will not work).

It can easily happen that you need to rename one of your columns and rows, for whatever reasons that might be. This is done with the `colnames()` and `rownames()` functions. These functions allow you two things: 1. to inspect the columns' names as well as to change them and 2. to change them using a vector containing strings..

```
colnames(DF1) #inspecting the column names

## [1] "Numbers" "NewNumbers"

colnames(DF1) <- c('Numbers', 'Letters')
colnames(DF1)

## [1] "Numbers" "Letters"

str(DF1)

## 'data.frame': 10 obs. of 2 variables:
## $ Numbers: int 1 2 3 4 5 6 7 8 9 10
## $ Letters: int 11 12 13 14 15 16 17 18 19 20
```

Dataframes are also accessed by means of indices. Indices can be both, numeric or characters. The following code shows some examples.

```
## Accessing columns
DF1[, 'Letters'] #name of column

## [1] 11 12 13 14 15 16 17 18 19 20
```

```

DF1[,1] #number of column

## [1] 1 2 3 4 5 6 7 8 9 10

DF1[, 'A'] #not existing colum

## Error in '[.data.frame'(DF1, , "A)": undefined columns selected

## Accessing rows
DF1['1',] #name of row

##      Numbers Letters
## 1          1       11

DF1[1,] #number of row

##      Numbers Letters
## 1          1       11

DF1['asd',] #not existing row

##      Numbers Letters
## NA          NA       NA

## Writing into cells
DF1[1, 1] = 9999999
head(DF1)

##      Numbers Letters
## 1 9999999          11
## 2          2       12
## 3          3       13
## 4          4       14
## 5          5       15
## 6          6       16

```

3.7 Excursion: Creating sequences and repetitions

Typing a sequence of numbers every time is rather exhausting. This is why R has a set of function, which enable you to repeat the assignment procedure. The above sequence between 1 and 10 can easily be created by the function `seq()` (=sequence). Sequence has two parameters, `by` which defines the stepsize between your edges; and `length`, which defines the amount of items in your new sequence vector.

```

five = seq(from=1, to=10, by=1)
five

## [1] 1 2 3 4 5 6 7 8 9 10

```

```
six = seq(from=1,to=10,by=2)
six

## [1] 1 3 5 7 9

seven = seq(from=1,to=10,by=0.5)
seven

[1] 1.0 1.5 2.0 2.5 3.0 3.5 4.0 4.5 5.0 5.5 6.0 6.5...
## [15] 8.0 8.5 9.0 9.5 10.0

seven1 = seq(from=1,to=10,length=2)
seven1

## [1] 1 10

seven2 = seq(from=1,to=10,length=15)
seven2

[1] 1.000000 1.642857 2.285714 2.928571 3.571429 4.214286...
[8] 5.500000 6.142857 6.785714 7.428571 8.071429 8.714286...
## [15] 10.000000
```

The `seq()` function enables you to create sequences of numbers (guess what the abbreviation `seq()` stands for?). You obviously can not only choose the beginning and end number, but also in what steps the sequence has to be created.

If you simply want to create a sequence with identical entries, you can use the function `rep()` (=repeat).

```
eight = rep(1, times=5)
eight

## [1] 1 1 1 1 1
```

In case you want to repeat a vector, `rep()` has two possible parameters: `times` defines, how many times the entire vector is repeated; `each` defines, how many times you want to repeat every entry of your vector. Both parameters can also be used simultaneously.

```
rep(1:3, times=5)

## [1] 1 2 3 1 2 3 1 2 3 1 2 3 1 2 3

rep(1:3, each=5)

## [1] 1 1 1 1 1 2 2 2 2 2 3 3 3 3 3

rep(1:3, times=2, each=5)

## [1] 1 1 1 1 1 2 2 2 2 2 3 3 3 3 3 1 1 1 1 1 2 2 2 2 2 3 3 3 3 3
```

These three functions can be nested, i.e. used within each other. But beware, the nesting order affects the output significantly. Try out different orders of the nesting.

```
a = rep(seq(from=1,to=4), times=4)
a
## [1] 1 2 3 4 1 2 3 4 1 2 3 4 1 2 3 4
```

Chapter 4

Extracting values from a dataframe

4.1 **Excursion:** How to design your programming task?

Before we can carry on, we have to do another excursion. You know how to find your data, how to read it in, how to check its structure, etc. Congratulations. Writing a script means that you need to assess what steps have to be taken in order to achieve your aim. The problem is that you are used to do these things everyday, which is why it is hard to realize what steps a computer needs to take to achieve its (or rather your) aims. Programming is not voodoo. Programming is about giving the computer instructions as precise as possible what it has to do. Despite the superficial increasing complexity of computers nowadays – touchscreens, apps, all those buzzwords which your grandmother hates –, computers are simple machines. They need to be told step by step what they should execute. When I say step by step, I really mean *exactly each step*.

Lets use a metaphor to illustrate this. Imagine you want to borrow a book from a library. But you cannot go to the library yourself since you are lying sick in bed. Rather, a friend of yours will help you. You instruct your friend via telephone where he needs to go and what he needs to do. The problem is that once you convey wrong information, your friend will not wonder around the library but will simply stop doing whatever he or she is doing. Hence, you need to be precise.

That is what is happening in the programing world. You have to break down every substep of the process into a command. So, in order to get the book, you will need to provide the instructions in the following steps:

1. Go to the library. (~ start program)
2. Go to the shelf XY. (~ locate data)
3. Find the book with the code XYZ. (~ locate data)
4. Pick it up. (~ get data)
5. Return to the counter. (~ process data)
6. Have the book scanned. (~ still processing data)
7. Return to me and give me the book. (~ illustrate/plot/return data)

In real live we are used to cluster subprocesses into one single process. More importantly, we are not aware of these subprocesses any more, because we have performed them so many times. Try to remember when you did a complex task for the first time, e.g. borrowing a book, driving a car, baking a cake. Sometimes you skipped a step and the entire thing got

screwed up. The advantage here is that we can use our attention to other useful things – like watching TV or listening to the radio. The disadvantage is that you need to become aware of subprocesses again in order to become a good programmer.

Task

You have to become aware of everyday subprocesses. In order to do so, observe your own behavior for one entire day. Write into a protocol what you do. Break down your behavior into subprocesses. Your protocol should look something like that:

- Waking up
 - Open eyes
 - Sitting up
 - Putting one leg on the floor
 - Turning
 - Putting the other leg on the floor
 - etc.
- Breakfast
 - Going to the kitchen
 - Going to the bread's location
 - Picking the bread up
 - Going to the table
 - Putting the bread on the table
 - Going to the Nutela's location
 - etc.

If you repeat this several times, you will very fast get a feeling into how many subprocesses you can break down an entire process.

4.2 Assessing the task

We know about the structure of "data", but we still don't know anything about its content. However, it would be very, very unpleasant to check all of the 887438 rows. Instead, we only check the first 25 rows using `head()`. In contrast to above, we do not need to specify the name of the parameter which allows us to increase the number of preview columns.

```
head(data, 25)
```

```
##          V1
## 1      RDX2
## 2         X
## 3 höchstens
## 4 unmöglich
## 5 höchstens
## 6
## 7  \003das
```

```

## 8
## 9  \aversteh
## 10
## 11  \003ich
## 12
## 13 \005nicht
## 14
## 15  \003<P>
## 16
## 17  \002ja
## 18
## 19  \003hä
## 20
## 21  \003<P>
## 22
## 23 \004aber
## 24
## 25  \003<P>

```

Now, let me tell you a little bit about the format of the dataframe for those of you who have never worked with Praat before. Praat is used to analyze speech signals and allows the user to set boundaries within the signal, or in other words: annotate it. Figure 4.1 illustrates the software. You can see that there are two representations of the acoustic signal: the oscillogram and the spectrogram. These are not important to us. The x-axis represents time. More important are the tiers called "words-corr" and "words" (check right edge). "words" represents the raw annotations, "words-corr" represents the corrected annotations. In this sense, "words" is **Garbage**, i.e. material which we do not need for our analysis. You have to keep in mind that corpora are always full of garbage and you need to develop a strategy to deal with it.

You can see that the speech signal in Figure 4.1 contains the following entries: "iP", "hä", "was". "iP" represents a pause. "hä" is a hesitation marker, "was" is a German word. Now compare the figure with the output of `read.table()`. You can see the same entries, starting in line 17. Let us assume that all the files from Praat have the same structure. This means that from line 17 on the text file shows you the content of the annotation. Furthermore, you are provided with "xmin" and "xmax" entries. These represent the boundaries for the entries. Each entry is superheaded by the term "intervals", including its index. Timesteps for boundaries belong to the entry below them. For example, the first entry, iP, starts at 0 and ends at 18.285744. The second entry, "hä", starts at 18.285744 and ends at 18.489535. And so forth. Let us take a look above line 17. There you can find additional information. First, "xmin" equals to 0; "xmax" equals to 1200, i.e. the signal is 1200 seconds long (lines 3-4). Second, the tier "words-corr" also runs from 0 to 1200 seconds (lines 10-12). It contains 1749 entries (line 13). Using this knowledge, we will write a script, allowing us to extract each entry and additional information such as their time stamps.

4.3 Preparing our programing task

Ask yourself, if you had to do this by hand, i.e. find the word and the according time stamps, how would you do it?

Figure 4.1: *Illustration of the speech analysis software Praat.*

Task

Think about this for a moment and take some notes. Check to what extent your notes and mine are comparable.

Again, there are million ways to perform this task. The fact that I show you one does not mean that others are not as efficient as the one proposed. First, I will show you the one solution you probably came up with:

- (1) We need to find out, which lines belong to which tier, i.e. "words" and "words-corrected". Remember, we want to dismiss "words", as it is garbage. Once we have this information, we know our area of operations.
- (2) You start at entry number 1, i.e. "intervals [1]". You extract its start and end point and note its content, i.e. 'P;'.
- (3) You go to entry number 2, perform the same thing as in (2) and go to the next interval.
- (4) You repeat this over and over again until you encounter "words", where the garbage starts.

Am I right? This is how you would probably plan it. The truth is that this is the way to perform the task if you performed the task manually. But we don't do manually. We do computers! We skip the entire repetition process. Don't worry, we will come back to that in Section ???. But for now, it is easier to leave that away.

- (1) In column V1, we find out all the row numbers in which the entry "intervals" is located. Let us call these rows RW.
- (2) By definition, RW+1 always contains the start point of the interval.
- (3) By definition, RW+2 always contains the end point of the interval.
- (4) By definition, RW+3 always contains the entry of the interval.
- (5) Now, we need to find out, which lines belong to which tier, i.e. "words" and "words-corrected". Remember, we want to disregard "words", as it is garbage. Once we have this information, we can get rid of the superfluous information.

4.4 Excursion: Searching, finding, etc.

We encounter a new problem: How do we perform a **Searching** command in R? Well, there are a couple of ways to solve such a task. Let's take our own lives again as an example. We are looking for something almost everyday: lost keys, the jacket, glasses, a recipe etc. This is also the case in programing. Imagine, you have a dataframe which stores the results of an inquire where you measured the body size of men and women. Imagine, you would like to analyze the measurements of one gender alone. In this case, you need to extract all the measurements for either men or women from the dataset. What you don't want to do is to go by hand through the dataframe and look for all your gender markers. You want to have the computer done this work for you. There are several ways to perform a searching operation, each specific to a certain type to look for something. In the following we will present them.

4.4.1 The == operator

In case you are looking for a single item in a vector/matrix/array and one single item alone, you can use the == operator. The item you are looking for has to have an exact match in the vector you are searching in. The == operator can be used for all classes, i.e. numbers, strings and factors. The == operator will provide you with a **Logical vector** consisting of TRUE and FALSE's, TRUE representing a match and FALSE representing a non-match.

```
DF = c('m', 'm', 'f', 'f', 'f')
DF$gender == 'm'

## Error in DF$gender: $ operator is invalid for atomic vectors

DF

## [1] "m" "m" "f" "f" "f"
```

Using the logical vector which comes as a result of your searching operation, you can access your dataframe and use it to get a subset of your data. Importantly, the logical vector can be stored in a new variable. This is the case for all the vectors which are produced by your searching operation.

In the next step, we will create a dataframe using the function `data.frame()`. The new dataframe "DF" will contain 10 instances of 'm' and 'f', which we will generate using the function `rep()`. We will assign to each of these instances a random number using the randomization function `sample()`. See Section ?? for more details on that function. Using the logical vector provided by the `==`, you can access the dataframe. In this case, the logical vector functions like an index, where each TRUE is extracted, while each FALSE is dismissed.

```
#create a new dataframe with random number between 1 and 20.
DF = data.frame(gender = rep(c('m','f'),10), size = sample(20))
head(DF) #check how the dataframe looks like

##      gender size
## 1         m    19
## 2         f     7
## 3         m    15
## 4         f    14
## 5         m    17
## 6         f     3

DF$gender == 'm' #find all men

[1]  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE...
## [12] FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE

DF$gender == 'f' #find all women

[1] FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE...
## [12]  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE FALSE  TRUE

#use the logical vector to take a subset of the dataframe (in the y-dimension)
DF[DF$gender == 'f',]

##      gender size
## 2         f     7
## 4         f    14
## 6         f     3
## 8         f    10
## 10        f    13
## 12        f    20
## 14        f    16
## 16        f     8
## 18        f    18
## 20        f     5

where = DF$gender == 'm'
DF[where,]
```

```
##      gender size
## 1         m   19
## 3         m   15
## 5         m   17
## 7         m    4
## 9         m    1
## 11        m   12
## 13        m    2
## 15        m    6
## 17        m    9
## 19        m   11

#does not work, as the vector is applied to the wrong dimension
DF[,where]

## Error in '[.data.frame'(DF, , where):  undefined columns selected

#you can also look for a specific number
DF[DF$size == 18,]

##      gender size
## 18         f   18

#or for numbers larger/smaller than another number
head(DF[DF$size > 18,]) #larger

##      gender size
## 1         m   19
## 12        f   20

head(DF[DF$size < 18,]) #larger

##      gender size
## 2         f    7
## 3         m   15
## 4         f   14
## 5         m   17
## 6         f    3
## 7         m    4
```

4.4.2 The %in% operator

In case you have multiple items, you want to find in your dataframe (vector, string, etc), you can use the %in% operator. It is equivalent in its use to the == operator, but you can provide it with multiple items.

```
strings = c('aa','ab','bb','ac','aac')
DF = data.frame(factor = rep(strings,10), size = sample(length(strings)*10))
head(DF) #check how the dataframe looks like
```

```
##      factor size
## 1      aa   31
## 2      ab   14
## 3      bb   35
## 4      ac   47
## 5     aac   29
## 6      aa   46

DF$factor %in% c('aa', 'ab')

[1]  TRUE  TRUE FALSE FALSE FALSE  TRUE  TRUE FALSE FALSE FALSE...
[12]  TRUE FALSE FALSE FALSE  TRUE  TRUE FALSE FALSE FALSE  TRUE...
[23] FALSE FALSE FALSE  TRUE  TRUE FALSE FALSE FALSE  TRUE  TRUE...
[34] FALSE FALSE  TRUE  TRUE FALSE FALSE FALSE  TRUE  TRUE FALSE...
## [45] FALSE  TRUE  TRUE FALSE FALSE FALSE

DF[DF$factor %in% c('aa', 'ab'),]

##      factor size
## 1      aa   31
## 2      ab   14
## 6      aa   46
## 7      ab    9
## 11     aa   15
## 12     ab   48
## 16     aa    4
## 17     ab   28
## 21     aa   45
## 22     ab   27
## 26     aa   10
## 27     ab   16
## 31     aa    3
## 32     ab   36
## 36     aa   18
## 37     ab   38
## 41     aa   30
## 42     ab   40
## 46     aa   25
## 47     ab   37
```

4.4.3 The `grep1()` function

In case you want to have all entries, which contain one or several characters, you can use the `grep1()` function. This is especially useful in case you are looking for words but don't know all the endings of the words. E.g., by looking for the string 'alle', you will find 'alle', 'alles', 'aller'.

`grep1()` is part of a function family which uses regular expressions. If you want to know more about regular expressions, google it. At this point, we limit ourselves to the search of a single item.

```

head(DF) #check how the dataframe looks like

##      factor size
## 1      aa    31
## 2      ab    14
## 3      bb    35
## 4      ac    47
## 5     aac    29
## 6      aa    46

where = grepl('a', DF$factor)
print(where)

[1]  TRUE  TRUE FALSE  TRUE  TRUE  TRUE  TRUE FALSE  TRUE  TRUE...
[12]  TRUE FALSE  TRUE  TRUE  TRUE  TRUE  TRUE FALSE  TRUE  TRUE  TRUE...
[23] FALSE  TRUE  TRUE  TRUE  TRUE  TRUE FALSE  TRUE  TRUE  TRUE  TRUE...
[34]  TRUE  TRUE  TRUE  TRUE FALSE  TRUE  TRUE  TRUE  TRUE  TRUE  FALSE...
## [45]  TRUE  TRUE  TRUE FALSE  TRUE  TRUE

head(DF[where,])

##      factor size
## 1      aa    31
## 2      ab    14
## 4      ac    47
## 5     aac    29
## 6      aa    46
## 7      ab     9

```

You can find out the syntax of the function by typing in `?grepl` into your command window.

4.4.4 The `grep()` function

The `grep()` function has the same functionality like the `grepl()` function. It differs in its output.

Task

Construct a new dataframe (method shown above) with new strings, e.g. 'acb', 'acx', 'bcx', 'xcv' etc.

Search for subparts of the strings, e.g. 'ac', 'cx'. How does this affect the results? Use once the `grepl()` and once the `grep()` function. How do the outputs differ from each other?

`grepl()` and `grep()` use **Regular expressions** as inputs, a very usefull but complex representation of searching entries. There are a couple of strings which these two functions interpret as expressions, not as searching strings. In the following, we will show you those

which we use most commonly. Go to your favorite searching engine to find out more about regular expressions.

```
Strings = c('abc', 'abd', 'cdd', 'adz', 'xyz')
#find strings matching two inputs
where = grep('a|x', Strings)
Strings[ where ]

## [1] "abc" "abd" "adz" "xyz"

#find strings starting with a
where = grep('^a', Strings)
Strings[ where ]

## [1] "abc" "abd" "adz"

#find strings ending with d
where = grep('d$', Strings)
Strings[ where ]

## [1] "abd" "cdd"
```

A nice and very usefull functionality of `grep()` is that it can return immediately the values you were looking for, instead of the indices for the vector you provided. You simply need to set its parameter `value` to `TRUE`.

```
grep('d$', Strings, value = TRUE)

## [1] "abd" "cdd"
```

4.4.5 The `which()` function

The `which()` function shows you the indices for all `TRUE` values in a Boolean vector , i.e. a vector containing only `TRUE` or `FALSE`, and which are the output of `==` , `%in%` and `grepl()` . See Section ?? for more information. In this sense, it is similar to the `grep()` function.

```
which(grepl('d$', Strings))

## [1] 2 3
```

4.4.6 Inversing searching results

Finding a string or a number is an important step for appending or extracting information out of your corpus. However, sometimes you want to get rid of a certain item. For this, you need to inverse the results of your searching process. There are several possibilities to do this. The easiest is to negate the outcome of the `==` and `grepl()` commands using the operator `!` . However, this is only possible with Boolean vectors, which are discussed in Section ??.

```
Vector = c(1,2,3,4,5,6,7,8,9)
where = Vector == 2
Vector[where]

## [1] 2

Vector[!where] #negate

## [1] 1 3 4 5 6 7 8 9
```

If you want to inverse the searching results of `grep()` or `which()` , you need to do it with the minus operator `-` .

```
Vector = c(1,2,3,4,5,6,7,8,9)
where = grep(2, Vector)
Vector[where]

## [1] 2

Vector[-where] #negate

## [1] 1 3 4 5 6 7 8 9
```

4.5 Extracting the appropriate values

Do you remember what we wanted to do? If not, here's a repetition of our task:

- (1) In column V1, we find out all the row numbers in which the entry "intervals" is located. Let us call these rows "RW".
- (2) By definition, RW+1 always contains the start point of the interval.
- (3) By definition, RW+2 always contains the end point of the interval.
- (4) By definition, RW+3 always contains the entry of the interval.
- (5) Now, we need to find out, which lines belong to which tier, i.e. "words" and "words-corrected". Remember, we want to disregard "words", as it is garbage. Once we have this information, we can get rid of the superfluous information.
- (6)

Let's tackle the first task, finding out where "intervals" is located. We will use the `==` operator for this task because we need to find only one item in the column "V1" in the dataframe "data" . Remember, in order to access a column in a dataframe, we use the `$` (Dollar) operator. In this way, the column functions as a vector. Simultaneously, we will store the result of the search operation, a logical vector consisting of TRUE and FALSE, into the new variable "whereIntervals" .

```
whereIntervals = data$V1 == 'intervals'
head(whereIntervals)

## [1] FALSE FALSE FALSE FALSE FALSE FALSE
```

Although the logical vector is a very powerful vector, its shortcoming is that you cannot do numeric calculation with it. As we stated above, we need to add some numbers to this vector in order to access the entries of our interest. In order to do so, we need to find out the indices of the TRUE entries using the function `which()`.

```
whereIntervals.numeric = which(whereIntervals)
head(whereIntervals.numeric)

## integer(0)
```

Alright, now we have the indices of the term "intervals". Using the numeric vector "whereIntervals.numeric", we now can access our entries of interest in column "V3" (since they are located in that column). Before we do so, we create three new vectors, using our instructions from above:

```
where.start = whereIntervals.numeric+1
where.end = whereIntervals.numeric+2
where.word = whereIntervals.numeric+3
```

Finally, we can access the entries in question using "where.start", "where.end" and "where.word" as indices and simultaneously storing them in new variables "w.start", "w.end", "word". Make it a habit to check the results of your processes using `head()` after every new process you create! In this way you spare you a lot of time finding errors in your script by making sure that the code you wrote is soundDebugging.

```
w.start = data$V3[where.start]
w.end = data$V3[where.end]
word = data$V3[where.word]
head(w.start)

## NULL

head(w.end)

## NULL

head(word)

## NULL
```

As you can see, we found our entries in question. However, they are in quotes, which means that they are character strings. In order to use these values as numbers and calculate with them, we need to convert these strings into numeric values using the function `as.numeric()`. Otherwise calculations are not possible. We can apply `as.numeric()` directly to the output of our indexing operation from above. You can see that you do not need

to store the result of a process into a variable and apply a new process to the new variable. We create such **Nested commands** in order to save space – writing space in your script and RAM – as well as time, because every line of code takes time to process.

```
class(w.start)

## [1] "NULL"

w.start = as.numeric(data$V3[where.start])
w.end = as.numeric(data$V3[where.end])

head(w.start)

## numeric(0)

head(w.end)

## numeric(0)

class(w.end)

## [1] "numeric"
```

Using `head()` you can see that the output lost its quotes. Finally, we need to store these values into a.... what did the Boss say? Exactly, a dataframe.

Task

Store the values in the vectors "w.start", "w.end" and "word" in a dataframe with the columns "word", "w.start" and "w.end". Store the dataframe into the variable "DF.textgrid". Check its class and structure. Add a new column "wordduration" to "DF.textgrid", in which you store a vector containing the word durations.

Allright, your solution should look something like the following code.

```
#Solution 1#
#-----#
#create dataframe
DF.textgrid = data.frame(
word = word,
w.start = w.start,
w.end = w.end,
stringsAsFactors =FALSE)

#adding wordduration
DF.textgrid$wordduration = DF.textgrid$w.end - DF.textgrid$w.start
head(DF.textgrid)

## [1] w.start      w.end      wordduration
## <0 rows> (or 0-length row.names)

#Solution 2#
#-----#
DF.textgrid = data.frame(
word = word,
w.start = w.start,
w.end = w.end,
wordduration = w.end - w.start,
stringsAsFactors =FALSE)

class(DF.textgrid)

## [1] "data.frame"

str(DF.textgrid)

## 'data.frame': 0 obs. of  3 variables:
## $ w.start      : num
## $ w.end        : num
## $ wordduration: num

head(DF.textgrid)

## [1] w.start      w.end      wordduration
## <0 rows> (or 0-length row.names)
```

It is important to note that the *variables* "w.start", "w.end" and "word" are independent from the *columns* in the dataframe "DF.textgrid", even though they share names. We might have chosen different names in order to reduce the confusion. But, well, life ain't easy. Note that the parameter **stringsAsFactors** is set to FALSE. Simply take this for granted, we will talk about this later. Congratulations, you created your first dataframe on the basis of a text input. I am very proud of you. There is a final thing to do:

We need to find out, which entries belong to the tier "words", which we want to dismiss and which belong to "words-corrected". There are plenty solutions to that problem. One

would have been to find the area for "words" and "words-corrected" using the term "name" (do `head(data,25)` and check line 10. You can see that "name" indicates the name of the tier). In this way we would have dismissed all unnecessary columns right from the start. Another possibility is to work with the dataframe we already have.

The solution is as follows: We first find out which of the two tiers is listed first, and which is listed second. We need to automatize this in case the order is reversed in another data set. You don't want to open every file and check the order. In order to do this, we find the row numbers in column V1 of the original variable "data" in which "name" is located and use them as indices in V3 to extract the tier names. As we can see below, the first part contains the annotations for the tier "words-corr", the second part contains the annotations for the tier "words".

```
index = which(data$V1 == 'name')#get column number
Tiernames = data$V3[index]
Tiernames

## NULL
```

But how do we find where the first part ends and the second starts? The solution is to use time to our advantage. Remember that the duration of the signal is 1200 seconds. Since we have two tiers in this textgrid, we should have 1200 twice in the dataframe (and also in the original text file. Open and check it out). Also, there should be twice the start of the timing, i.e. 0.

```
index.start = which(DF.textgrid$w.start == 0)#get column number
index.start

## integer(0)
```

By default, everything between the first entry of "index.start" and the second entry of "index.start" belongs to . The second entry marks the beginning of the part which belongs to , which is why the row above NA is the end of part one. The following code translates this thoughts into an R routine. See that `seq()` now takes another parameter, "by" defines the step size of the sequence. Test this by including "2" instead of "1".

```
index.firstpart = seq(index.start[1], index.start[2]-1, by = 1)

## Error in seq.default(index.start[1], index.start[2] - 1, by = 1): 'from' must be
a finite number

DF.firstpart = DF.textgrid[index.firstpart,]

## Error in '[.data.frame'(DF.textgrid, index.firstpart, ): object 'index.firstpart'
not found

tail(DF.firstpart) # check last few entries

## Error in tail(DF.firstpart): object 'DF.firstpart' not found
```

Note that in the manuscript the output of `tail()` was rounded before the digits. However, in your prompt you should see at least three digits.

Finally, we can get gauge some basic statistics, e.g. large was our dataframe, how many rows did we lose, etc. using `dim()` , `nrow()` and `ncol()` .

```
#dimensions of old dataframe
dim(DF.textgrid)#get size of both dimensions

## [1] 0 3

ncol(DF.textgrid)#get number of columns

## [1] 3

nrow(DF.textgrid)#get number of rows

## [1] 0

#dimensions of new dataframe
dim(DF.firstpart)#get size of both dimensions

## Error in eval(expr, envir, enclos): object 'DF.firstpart' not found

ncol(DF.firstpart)#get number of columns

## Error in ncol(DF.firstpart): object 'DF.firstpart' not found

nrow(DF.firstpart)#get number of rows

## Error in nrow(DF.firstpart): object 'DF.firstpart' not found
```

Chapter 5

Creating a large dataframe

5.1 Repetition

Well, you did it! You extracted the necessary data from a text file lying around somewhere on your hard disk. Very good. Now we need to repeat that process for the remaining 60 files. Remember your friend in the library who helped you finding your book from section 4.1. It is as though you told him to get 10 instead of only one book from the library. We constantly repeat processes in real life: We put ten eggs into a bowl, instead of one; we clean ten forks, instead of one. These things have in common that they have a core element which is repeated. Regarding your programming endeavour, you need to repeat all those processes from above for every single file. We will look at this process in the next section.

5.2 **Excursion:** Repetition in programming

Take the baking example from above, where you want to bake a cake. You need to break 10 eggs (large cake!) into the bowl. The recipe you follow will simply state "Take the eggs and put them into a bowl". You will then take one egg, break it, put egg white and egg yolk into a bowl and throw the eggshell away. Then you will repeat the procedure until no eggs are left. What have you done? You repeated exactly the same behavior several times. The recipe did not have to state every of the ten repetitions, it simply assumed you were capable enough to perform them.

Repetition in R is accomplished by means of the `for()` function, telling the computer to repeat a specific command. We call this a **Loop** and the computer *loops* or *iterates* across the command.

The syntax of this function is a little bit different to that of all previous code because you need to take into account multiple things. Let me guide you through this. `for()` works similar to other functions, only that you need provide it with two parameters, V1 and V2: `for(V1 in V2)`. Let us start with V2 first. V2 is a vector in which you store the items across which `for()` has to loop. When the vector has to contain numbers, we create a sequence using the operator `:`, which is similar to `seq()` insofar as it creates a sequence of numbers differing by one. What then happens is that `for()` goes through V2 step by step and stores the current value in V1. `for()` creates an environment which, apart from one variable, never changes. This one variable is V1. The environment is defined by the curly braces `"{}"`. Let us look at the following example:


```

for (iloop in 1:3){
result = iloop+1
cat('In loop', iloop, 'we add one to', iloop, 'and the result is', result, '\n')
}

## In loop 1 we add one to 1 and the result is 2
## In loop 2 we add one to 2 and the result is 3
## In loop 3 we add one to 3 and the result is 4

```

You can see that V1 is represented by "iloop", whereas V2 is represented by "1:3". In this example, we tell the computer to loop across numbers from one to ten. In each loop we take hand the current loop value and hand it over to `iloop` by means of the `in` command. Within each loop, we take the value in the `iloop` variable, add one to it and store it in `result`. Afterwards we insert `iloop` and `result` into the `cat()` function in order to print the output in the command window (or here in the text).

Within loops, we need to use the `cat()` function since the output of variables is inhibited within a loop (why is this so, hmmm?). The nice thing about `for()` is that it does not have necessarily loop across numbers as in the example above. Rather, you can provide it with what ever you want: numbers, strings, factors, etc. Both possibilities, i.e. looping over numbers or over character strings, have their advantages and disadvantages.

```

stringvector = c('a', 'b', 'c', 'd')
for (iloop in stringvector){
cat('The current value of iloop is: ', iloop, '\n')
}

## The current value of iloop is:  a
## The current value of iloop is:  b
## The current value of iloop is:  c
## The current value of iloop is:  d

```

Here, V1 is still represented by "iloop", whereas we do not create a vector directly in V2 but hand it over the variable "stringvector". We strongly advise to use telling names for the looping variable. You can, for example, always use the letter 'i' in front of your telling-name-variable, in order to indicate that the variable contains the for-value (i for iterate).

By the way, the syntax of `cat()` is really easy. You can provide it with string sequences and variables, separated by commata and it will generate a big string sequence out of the input and simultaneously show it on your prompt. Very usefull.

The important thing to remember is that the `for()` environment works like the standard environment, apart from the fact that V1 changes. For example, you can test whether the content of V1 contains an item you are looking for:

```

stringvector = c('Live', 'long', 'and', 'prosper')
for (iloop in stringvector){
Vulcan.o = grepl('o', iloop)
cat('The value in iloop contains the letter "o":', Vulcan.o, '\n')
}

```

```
## The value in iloop contains the letter "o": FALSE
## The value in iloop contains the letter "o": TRUE
## The value in iloop contains the letter "o": FALSE
## The value in iloop contains the letter "o": TRUE
```

Finally, a piece of advice: Take care of your brackets and do not forget to close them in the correct order! This is one of the most common scripting errors which consumes lots of time and energy to debug.

5.3 Reading multiple files

Allright, let's wrap up what we know so far. You already can:

- localize files
- read in files
- search content in a vector
- create for loops

This is basically all it takes, to read in a corpus (at least the one you have here). Let us first design the programming task:

- 1) get file names, store in variable
- 2) iteratively access filenames
- 3) use filename to read in textgrid
- 4) bind new textgrid to existing dataframe
- 5) store large dataframe

5.4 Excursion: Binding dataframes

There are two things missing at this point, which are essential for writing our code. First, we need to find out how many entries the variable containing filenames has. This is done with the `length()` function.

```
V = seq(from = 1, to = 20)
length(V)

## [1] 20
```

The problem is binding two existing dataframes. This step is performed with the `rbind()` function, which stands for "row bind".

```

DF1 = data.frame(A = 1:3, B = 11:13)
DF2 = data.frame(A = 1:3, B = 11:13)
DF3 = data.frame() #empty DF
DF3 = rbind(DF1, DF2)
DF3

##   A  B
## 1 1 11
## 2 2 12
## 3 3 13
## 4 1 11
## 5 2 12
## 6 3 13

```

There are two premises for this function. First: the two dataframes in question need to have the same amount of columns. There are two possibilities to do this. You can either gauge the size of the dataframe using `dim()` . Or you can directly question the number of columns using `ncol()` . The equivalent for rows is `nrow()` .

```

dim(DF1)#size of dataframe

## [1] 3 2

ncol(DF1)#number of columns

## [1] 2

nrow(DF1)#number of rows

## [1] 3

```

It is important to know that `dim()` , `ncol()` and `nrow()` can be used for dataframes only, while `length()` works only with vectors.

```

dim(V)

## NULL

ncol(V)

## NULL

nrow(V)

## NULL

```

Hence, when we try to merge dataframes with different dimensions we get an error:

```
#create dataframes
DF1 = data.frame(A = 1:3, B = 11:13, C = letters[1:3])
DF2 = data.frame(A = 1:3, B = 11:13)
#try to bind ---- nope, not working
DF3 = data.frame() #empty DF
try({DF3 = rbind(DF1, DF2)})
DF3

## data frame with 0 columns and 0 rows
```

By the way, you probably saw function `try()` in the above code. This function comes in handy whenever you don't know whether your code will evaluate, especially later, when we use for-loops. Here, I used it, because otherwise my script to create this manuscript wouldn't evaluate and you would be able to read it. To see the error message, copy paste the above code into your R environment.

The second premise of `rbind()` is: The two dataframes to be bound need to have the same column names. Remember `colnames()` ? Exactly, the function we use to give a dataframe column names can also be used to inquire them.

```
#create dataframes
DF1 = data.frame(A = 1:3, B = 11:13)
DF2 = data.frame(C = 1:3, D = 11:13)
#assess names
colnames(DF1)

## [1] "A" "B"

colnames(DF2)

## [1] "C" "D"

#bind
DF3 = data.frame() #empty DF
try({DF3 = rbind(DF1, DF2)}, silent = FALSE)
DF3

## data frame with 0 columns and 0 rows
```

There are two solutions to this problem. Either you create new dataframes with the same column names. This is what we already have done above, which is why I will not repeat the process here again. Or, you give one of the dataframes the column names of the other one.

```
#create dataframes
DF1 = data.frame(A = 1:3, B = 11:13)
DF2 = data.frame(C = 1:3, D = 11:13)
#assess column names
colnames(DF1)
```

```
## [1] "A" "B"

colnames(DF2)

## [1] "C" "D"

#change column names
colnames(DF1) <- colnames(DF2) #isn't that cool?
#assess column names
colnames(DF1)

## [1] "C" "D"

colnames(DF2)

## [1] "C" "D"

DF3 = data.frame() #empty DF
DF3 = rbind(DF1, DF2)
DF3#it worked!

##      C  D
## 1 1 11
## 2 2 12
## 3 3 13
## 4 1 11
## 5 2 12
## 6 3 13
```

Ok. There one last thing you need to remember when binding dataframes. You need to make sure that the data of interest is in the same columns in both of your original dataframes. Also, the classes of your columns should match. Otherwise `rbind()` will mess up your data structure and the class of your columns.

```
#create dataframes
DF1 = data.frame(A = 1:3, B = letters[1:3])
DF2 = data.frame(A = letters[11:13], B = 11:13)
#assess structure of original DFs
str(DF1)

## 'data.frame': 3 obs. of 2 variables:
## $ A: int 1 2 3
## $ B: Factor w/ 3 levels "a","b","c": 1 2 3

str(DF2)

## 'data.frame': 3 obs. of 2 variables:
## $ A: Factor w/ 3 levels "k","l","m": 1 2 3
## $ B: int 11 12 13
```

```
#bind
DF3 = data.frame() #empty DF
DF3 = rbind(DF1, DF2)

## Warning in '[<-.factor'('*tmp*', ri, value = 11:13):  invalid factor level, NA generated

DF3#well, that is not nice

##      A      B
## 1 1      a
## 2 2      b
## 3 3      c
## 4 k <NA>
## 5 1 <NA>
## 6 m <NA>
```

Hence, we need to make sure that column A contains numbers, while column B contains letters. The easiest solution to this problem is to rename the column names in the second dataframe. `rbind()` is intelligent enough to bind the data by column names and not by their order.

```
colnames(DF2) = c('B', 'A')
DF3 = data.frame() #empty DF
DF3 = rbind(DF1, DF2)
DF3#well, that worked

##      A B
## 1   1 a
## 2   2 b
## 3   3 c
## 4 11 k
## 5 12 l
## 6 13 m
```

In summary, you to bind two dataframes, you need to take into account that

- 1) they have an equal number of columns
- 2) the columns in question have the same name
- 3) the data in question is in the same column

Task

Given the code you already know, write a script on your own which performs all the steps from above.

5.5 There is still a little problem to be solved

TODO: Write text about checking Tiernames

5.6 Finally we create a large dataframe

Finally, after 5 sections we arrived at the point where we can read in the entire directory full of textgrids into R and create a big large dataframe. The following code performs all the steps from above. We first find out the filenames using `list.files()`. Next, we create an empty dataframe using `data.frame()`, to which we will iteratively append the read textgrid. Finally, we create a `for()` loop in which we iterate across the filenames, read the according file, preprocess it in the way we have done in the previous section, and bind the result to the large dataframe using `rbind()`. Importantly, the size of "DFbig" will iteratively grow, gauged by means of `nrow()` nested in `print()`. Check the code line by line in order to understand what it actually does.

```
#DIR = 'REPLACE.THIS.STRING.WITH.YOUR.DIRECTORY'
setwd(DIR) #set working directory
filenames = list.files(pattern='TextGrid') #get filenames

DFbig = data.frame() #create new dataframe
nrow(DFbig)

## [1] 0

for (ifile in 1:length(filenames)){
  #get one filename. this is the point where the
  #result of the iteration is applied
  FN = filenames[ifile]
  #store dataframe temporarily
  tmp = read.table(FN, fill = NA, stringsAsFactors=FALSE)

  ## take only corrected words
  index = which(tmp$V1 == 'name')
  Tiernames = tmp$V3[index]
  take = c('words-corr', 'words_corr', 'words-korr', 'words_korr')
  index.start = index[Tiernames %in% take]
  DF.corrected = tmp[seq(index.start, nrow(tmp),1),]

  #find intervals
  whereIntervals = DF.corrected $V1 == 'intervals'
  whereIntervals.numeric = which(whereIntervals)
  where.start = whereIntervals.numeric+1
  where.end = whereIntervals.numeric+2
  where.word = whereIntervals.numeric+3

  #extract words and timestamps
  w.start = as.numeric(DF.corrected $V3[where.start])
```

```

w.end = as.numeric(DF.corrected $V3[where.end])
word = DF.corrected $V3[where.word]

#create dataframe
DF.tmp = data.frame(
  filename = FN, #save filename
  word = word, #save word tags
  w.start = w.start, #word start
  w.end = w.end, #word end
  wordduration = w.end - w.start,
  stringsAsFactors =FALSE)

#rbind to DFbig
DFbig = rbind(DFbig, DF.tmp)
cat(ifile, '. iteration:', nrow(DFbig), 'rows\n')
}

## 1 . iteration: 3569 rows
## 2 . iteration: 5707 rows
## 3 . iteration: 7814 rows
## 4 . iteration: 10731 rows
## 5 . iteration: 13475 rows
## 6 . iteration: 16030 rows
## 7 . iteration: 16791 rows
## 8 . iteration: 17510 rows
## 9 . iteration: 18328 rows
## 10 . iteration: 19681 rows
## 11 . iteration: 20896 rows
## 12 . iteration: 22402 rows
## 13 . iteration: 24180 rows
## 14 . iteration: 25981 rows
## 15 . iteration: 27469 rows
## 16 . iteration: 28553 rows
## 17 . iteration: 29922 rows
## 18 . iteration: 31070 rows
## 19 . iteration: 31954 rows
## 20 . iteration: 32931 rows
## 21 . iteration: 33960 rows
## 22 . iteration: 34628 rows
## 23 . iteration: 35097 rows
## 24 . iteration: 36331 rows
## 25 . iteration: 37375 rows
## 26 . iteration: 38364 rows
## 27 . iteration: 39317 rows
## 28 . iteration: 40356 rows
## 29 . iteration: 41671 rows
## 30 . iteration: 42300 rows
## 31 . iteration: 43159 rows
## 32 . iteration: 44189 rows
## 33 . iteration: 44991 rows

```



```
## 34 . iteration: 46053 rows
## 35 . iteration: 47135 rows
## 36 . iteration: 47897 rows
## 37 . iteration: 49441 rows
## 38 . iteration: 50803 rows
## 39 . iteration: 52248 rows
## 40 . iteration: 53416 rows
## 41 . iteration: 54611 rows
## 42 . iteration: 56112 rows
## 43 . iteration: 57115 rows
## 44 . iteration: 57945 rows
## 45 . iteration: 58688 rows
## 46 . iteration: 59377 rows
## 47 . iteration: 60200 rows
## 48 . iteration: 61285 rows
## 49 . iteration: 62853 rows
## 50 . iteration: 64592 rows
## 51 . iteration: 66265 rows
## 52 . iteration: 68075 rows
## 53 . iteration: 69671 rows
## 54 . iteration: 70958 rows

dim(DFbig)

## [1] 70958      5

#DO. NOT. FORGET. TO. SAVE!!!!
save(DFbig, file = 'DF_textgrids.rda')
```

Congratulations, you have preprocessed your first corpus (at least the one I know of). Now the fun starts.

Chapter 6

Adding speaker information to the dataframe

6.1 Recall your task

Let us recall what the task was and what we already can check off.

- (1) We start slow here. It is a corpus of spontaneously spoken language German. Two people were discussion a topic of their own interest.
→**OK, we got this**☺!
- (2) I provide you with transcriptions of the discussions. Oh, wait. We already have **Annotations** for **Praat** of the transcriptions praat. I give you these to work on. They are in a standard text format.
→**Got this, too!**☺!
- (3) I'll send you the directory.
→**Found it!**
- (4) Good question. You are a smart one. The final result, i.e. the result of your work, should be a **Dataframe** in R.
→**Nailed it**☺!
- (5) Well, you know: Time stamps. Word durations. Speaker information. Word frequencies. Conditional probabilities.
→**OH DEAR** ☹️

Ok, what points have we already tackled? We have time stamps, we have word durations. We are lacking speaker information and all that weird mathematical stuff. Everything what happens from now on builds on the knowledge you already gained. In a nutshell, all following processing steps consist of some kind of iteration process, during which information is searched and added to existing dataframes. What I will try to teach you here is how to write code which will combine information from different sources.

6.2 Adding speaker information

6.2.1 Checking the file

We first start with the simpler information: meta data about speakers. Before you read the file into R, go to the directory where you stored the file "SpeakerInformation.csv" containing speaker information and open it *manually* with a standard text editor. csv stands for "comma separated values" and is a common text format for tables. Make sure you open it with a text editor and not with an excel type of a program. **Make it a habit to open new files in a text editor to inspect them. Especially text files. This enables you to plan your coding efforts better.** Opening the file containing the speaker information you can see that the first line is different from all the following file. This file is the **Header** of the table, representing its column names. Columns are separated by spaces. Furthermore you can see that some information is written in quotes, representing strings, and other is not, representing numbers. We will deal with these differences once we have loaded the table. You also can see some *NAs* in the table. This indicates missing information about the speakers.

Note that the first line, the header, contains 6 columns, whereas in the table below there are 7 columns. The second line contains the following information, divided into columns:

- 1) "rec.001_AS.id.002"
- 2) 2
- 3) "männlich"
- 4) 1990
- 5) "28032014"
- 6) 190
- 7) 70

6.2.2 Reading the file

We will cope with this inconsistency in a second. First, we load the table using `read.csv()`. `read.csv()` is similar to `read.table()`, but is applicable to files ending in *.csv, a common text output of excel type of programs.

We set the **header** parameter to TRUE, so that the first line in the file is used as column names.

```
setwd(DIR)
SpeakerInfo = read.csv(
  file = 'SpeakerInformation.csv',
  stringsAsFactor=FALSE
)

head(SpeakerInfo)
```

```

      SubjID      Sex Birthyear RecordingDateDDMMYYYY...
rec_001_AS_id_002    2 männlich      1990      2803201...
rec_001_IS_id_001    1 männlich      1989      2803201...
rec_003_AS_id_006    6 weiblich      1991      1108201...
rec_003_IS_id_005    5 weiblich      1991      1108201...
rec_004_AS_id_008    8 männlich      1988      2708201...
rec_004_IS_id_007    7 weiblich      1985      2708201...
##                               WeightInKg
## rec_001_AS_id_002              70
## rec_001_IS_id_001              68
## rec_003_AS_id_006              55
## rec_003_IS_id_005              60
## rec_004_AS_id_008              86
## rec_004_IS_id_007              51

str(SpeakerInfo)

## 'data.frame': 72 obs. of  6 variables:
##  $ SubjID      : int  2 1 6 5 8 7 10 9 12 11 ...
##  $ Sex          : chr  "männlich" "männlich" "weiblich" ...
##  $ Birthyear    : int  1990 1989 1991 1991 1988 1985 198...
##  $ RecordingDateDDMMYYYY: int  28032014 28032014 11082014 110820...
##  $ SizeInCm     : int  190 180 167 170 179 159 182 171 1...
##  $ WeightInKg   : int  70 68 55 60 86 51 68 65 58 80 ...

```

Using `head()` and `str()` we can investigate the information in the file. Both functions tell you that the dataframe "SpeakerInfo" contains 6 columns. The reason for this is that `read.csv()` used the content of the first column in the file as row names. The row names now represent the corresponding file names for each subject. Next we need to know what each column contains:

- 1) **SubjID**: The speaker's identity number. Note that this number coincides with the last number in the file name (cf. row names).
- 2) **Sex**: The speaker's sex. This column contains characters.
- 3) **Birthyear**: His or her birthyear. This column contains integers.
- 4) **RecordingDateDDMMYYYY**: The recording date for the file. This column contains integers.
- 5) **SizeInCm**: The body size of the speaker in cm. This column contains integers.
- 6) **WeightInKg**: The speaker's weight in kg. This column contains integers.

Recall that the fifth column in the "SpeakerInformation.txt" file contained numeric values in quotes (if you scroll down you will see that it really contains only numeric values). `read.csv()` is smart enough to realize this and transforms the content between the quotes into numerics. However, we don't want to have this. The format of this information is two positions for days, two positions for months and four positions for years. However, transforming the values into integers, `read.csv()` deleted the first zero, reducing the number of characters of 8 in the file to 7.

```
SpeakerInfo[4:8,c('SubjID', 'RecordingDateDDMMYYYY')]
```

```
##              SubjID RecordingDateDDMMYYYY
## rec_003_IS_id_005      5          11082014
## rec_004_AS_id_008      8          27082014
## rec_004_IS_id_007      7          27082014
## rec_005_AS_id_010     10          1092014
## rec_005_IS_id_009      9          1092014
```

You can gauge the number of characters in "RecordingDateDDMMYYYY" using `nchar()`. We expect that there are always 8 characters across all rows. Thus we look whether all entries actually have 8 entries. If this is the case, the sum of the check should be equal to the number of rows in "SpeakerInfo".

```
HowManyPositions = nchar(SpeakerInfo$RecordingDateDDMMYYYY)
AllEight = HowManyPositions == 8
sum(AllEight)

## [1] 39

nrow(SpeakerInfo)

## [1] 72

sum(AllEight) == nrow(SpeakerInfo)

## [1] FALSE
```

Well, nope, that is not the case. This is very unfortunate, since we need the positional information to identify which position represents which piece of information (day, month, year) (Stupid creators of the corpus wouldn't come up with an easier storage format). We therefore need to tell `read.csv()` not to transform this column to integers. We do this using the parameters `as.is`, instructing the function not to change the class of the input, and `quote`, deleting the information of how strings are quoted. In this way we force `read.csv()` not to apply any changes to the values.

```
setwd(DIR)
SpeakerInfo = read.csv(
  file = 'SpeakerInformation.csv', #filename
  header = TRUE, #there is a header
  stringsAsFactor=FALSE, #no factors
  as.is = TRUE, #do not perform transformations
  quote = "" #no quotes
)

SpeakerInfo[4:8,c('SubjID', 'RecordingDateDDMMYYYY')]

##              SubjID RecordingDateDDMMYYYY
## "rec_003_IS_id_005"      5          "11082014"
```

```
## "rec_004_AS_id_008"      8      "27082014"
## "rec_004_IS_id_007"      7      "27082014"
## "rec_005_AS_id_010"    10      "01092014"
## "rec_005_IS_id_009"      9      "01092014"
```

The new output does not transform the date column into integers. However, it now contains the quotes from the file. This is a problem we can deal with later. The most important task is to add this information to our dataframe containing the TextGrids.

6.2.3 Merging information

You probably do not have the TextGrids in your working memory anymore. So, we first need to load them. This is standard procedure.

```
setwd(DIR)
load("DF_textgrids.rda", verbose = TRUE)

## Loading objects:
##   DFbig

head(DFbig,2)

##               filename word  w.start   w.end
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid <P> 0.00000 18.28574
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid  hä 18.28574 18.48954
##   wordduration
## 1   18.2857449
## 2    0.2037911
```

Next, we need to think about, how we will append the speaker information to the TextGrid dataframe. We will use a `for()` loop to perform this task.

Task

Before I show you my design, write a pen+pencil description of how you would perform this task manually.

Your description will probably look like one of the following two. Design 1:

- 1) Take row 1 from dataframe "DFbig".
- 2) Take filename.
- 3) Find appropriate speaker in dataframe "SpeakerInfo".
- 4) Copy information from "SpeakerInfo" and append to row 1 in dataframe "DFbig"
- 5) Go to row 2 and repeat the above steps.

The design above uses the large dataframe "DFbig" as a basis and performs the searching processes in the smaller dataframe "SpeakerInfo". "DFbig" has 70958 rows, "SpeakerInfo" has 72. If you use "DFbig" as the basic dataframe across which you iterate you will have more iterations than if you had done this with "SpeakerInfo". This is why there is another design for the task. Design 2:

- 1) Take row 1 from dataframe "SpeakerInfo".
- 2) Take filename.
- 3) Find the rows in dataframe "DFbig" which match the filename.
- 4) Copy information from "SpeakerInfo" and append to all respective rows in dataframe "DFbig"
- 5) Go to row 2 in "SpeakerInfo" and repeat the above steps.

We will write this code step by step and built up the code sequentially. First, need to make sure that we can easily access the filenames in "SpeakerInfo". Filenames are stored as row names in "SpeakerInfo". We therefore store it additionally in the column "filename". We also rename the row names by numbers, in order to decrease the size of the output.

```
#change row names
SpeakerInfo$filename= rownames(SpeakerInfo)
rownames(SpeakerInfo) = 1:nrow(SpeakerInfo)
head(SpeakerInfo)
```

	SubjID	Sex	Birthyear	RecordingDateDDMMYYYY	SizeInCm	Wei...
1	2	"männlich"	1990	"28032014"	190	...
2	1	"männlich"	1989	"28032014"	180	...
3	6	"weiblich"	1991	"11082014"	167	...
4	5	"weiblich"	1991	"11082014"	170	...
5	8	"männlich"	1988	"27082014"	179	...
6	7	"weiblich"	1985	"27082014"	159	...

```
##          filename
## 1 "rec_001_AS_id_002"
## 2 "rec_001_IS_id_001"
## 3 "rec_003_AS_id_006"
## 4 "rec_003_IS_id_005"
## 5 "rec_004_AS_id_008"
## 6 "rec_004_IS_id_007"
```

Then, we write the frame for our process, the `for()` -loop:

```
#write frame for-loop
for (irow in 1:nrow(SpeakerInfo)){
}
```

Next, we use `irow` to access the rows in "SpeakerInfo" and store it in a new variable, "info".

```
#access SpeakerInfo iteratively
for (irow in 1:nrow(SpeakerInfo)){
  info = SpeakerInfo[irow,]#access rows in SpeakerInfo
}
```

I usually check the output of my code. I do this for every single line I write in order to minimize the probability of bugs. For this, I save a dummy index in the variable across `for()` iterates and use it to perform the code I've written:

```
#checking by dummy coding
irow = 1
info = SpeakerInfo[irow,]
info

  SubjID      Sex Birthyear RecordingDateDDMMYYYY SizeInCm Wei...
1      2 "männlich"    1990      "28032014"      190    ...
##                filename
## 1 "rec_001_AS_id_002"
```

Well, that works nicely. Next step: Using "filename" we find the rows in "DFbig" which belong to the subject by matching it to "filename" in "DFbig". However, there are two problems. Compare the filename from "DFbig" with the one with "SpeakerInfo":

```
#Matching process using filename
info$filename

## [1] "\"rec_001_AS_id_002\""

DFbig$filename[1]

## [1] "rec_001_AS_id_002_CUT_20MIN_1.TextGrid"
```

First, the content of "info\$filename" is flanked by quotes, which is not the case in the filenames from "DFbig". Second, the filename from "DFbig" has an extension attached to it, namely "CUT_20MIN_1.TextGrid". This extension was either information from processing (CUT), indicates the duration of the recording (20MIN), indicates the recording number "1" and the file extension ".TextGrid". We don't need neither of this information, as this does not inform us about the speaker.

Whatever the reason for a mismatch, we need to change the filenames so that they match again. Concretely, we need to erase the quotes from the first case and the extension of the

filename in the second case. We will do this using the function `gsub()`, which also uses **Regular expressions**. As all functions in R, `gsub()` can be applied to a single entry as well as to an entire vector. It takes as arguments first the string which has to be changed, then the result and last the variable or string in which the change has to be performed. I usually apply a changing function to a single entry so that I can see the output of the changing process.

```
# Substitution and deletion in strings
one = gsub("'", "", SpeakerInfo$filename[1])
one

## [1] "rec_001_AS_id_002"

two = gsub("_CUT_20MIN_1.TextGrid", "", DFbig$filename[1])
two

## [1] "rec_001_AS_id_002"

one == two #check whether they match

## [1] TRUE
```

Now that the result looks the way we want it, we can apply it to the entire vector and store it in the original column. Processes as changing names or renaming, etc., have to be performed before the `for()` loop. Otherwise they will be performed in each iteration, unnecessarily increasing processing time. It might be that the extension in "DFbig" varies between the filenames. In order to find all possible variations, we need to check all the possible filenames. We can do this by applying `unique()` to the vector. `unique()` reduces the data to unique items.

```
#check extents of needed changes
unique(DFbig$filename)

## [1] "rec_001_AS_id_002_CUT_20MIN_1.TextGrid"
## [2] "rec_001_AS_id_002_CUT_20MIN_2.TextGrid"
## [3] "rec_001_AS_id_002_CUT_20MIN_3.TextGrid"
## [4] "rec_001_IS_id_001_CUT_20MIN_1.TextGrid"
## [5] "rec_001_IS_id_001_CUT_20MIN_2.TextGrid"
## [6] "rec_001_IS_id_001_CUT_20MIN_3.TextGrid"
## [7] "rec_003_AS_id_006_CUT_10min_1.TextGrid"
## [8] "rec_003_AS_id_006_CUT_10min_2.TextGrid"
## [9] "rec_003_AS_id_006_CUT_10min_3.TextGrid"
## [10] "rec_003_AS_id_006_CUT_10min_4.TextGrid"
## [11] "rec_003_AS_id_006_CUT_10min_5.TextGrid"
## [12] "rec_003_AS_id_006_CUT_10min_6.TextGrid"
## [13] "rec_003_IS_id_005_CUT_10min_1.TextGrid"
## [14] "rec_003_IS_id_005_CUT_10min_2.TextGrid"
## [15] "rec_003_IS_id_005_CUT_10min_3.TextGrid"
## [16] "rec_003_IS_id_005_CUT_10min_4.TextGrid"
## [17] "rec_003_IS_id_005_CUT_10min_5.TextGrid"
```

```
## [18] "rec_003_IS_id_005_CUT_10min_6.TextGrid"
## [19] "rec_004_AS_id_008_CUT_10min_1.TextGrid"
## [20] "rec_004_AS_id_008_CUT_10min_2.TextGrid"
## [21] "rec_004_AS_id_008_CUT_10min_3.TextGrid"
## [22] "rec_004_AS_id_008_CUT_10min_4.TextGrid"
## [23] "rec_004_AS_id_008_CUT_10min_5.TextGrid"
## [24] "rec_004_AS_id_008_CUT_10min_6.TextGrid"
## [25] "rec_004_IS_id_007_CUT_10min_1.TextGrid"
## [26] "rec_004_IS_id_007_CUT_10min_2.TextGrid"
## [27] "rec_004_IS_id_007_CUT_10min_3.TextGrid"
## [28] "rec_004_IS_id_007_CUT_10min_4.TextGrid"
## [29] "rec_004_IS_id_007_CUT_10min_5.TextGrid"
## [30] "rec_004_IS_id_007_CUT_10min_6.TextGrid"
## [31] "rec_005_AS_id_010_CUT_10min_1.TextGrid"
## [32] "rec_005_AS_id_010_CUT_10min_2.TextGrid"
## [33] "rec_005_AS_id_010_CUT_10min_3.TextGrid"
## [34] "rec_005_AS_id_010_CUT_10min_4.TextGrid"
## [35] "rec_005_AS_id_010_CUT_10min_5.TextGrid"
## [36] "rec_005_AS_id_010_CUT_10min_6.TextGrid"
## [37] "rec_005_IS_id_009_CUT_10min_1.TextGrid"
## [38] "rec_005_IS_id_009_CUT_10min_2.TextGrid"
## [39] "rec_005_IS_id_009_CUT_10min_3.TextGrid"
## [40] "rec_005_IS_id_009_CUT_10min_4.TextGrid"
## [41] "rec_005_IS_id_009_CUT_10min_5.TextGrid"
## [42] "rec_005_IS_id_009_CUT_10min_6.TextGrid"
## [43] "rec_016_AS_id_032_1.TextGrid"
## [44] "rec_016_AS_id_032_2.TextGrid"
## [45] "rec_016_AS_id_032_3.TextGrid"
## [46] "rec_016_AS_id_032_4.TextGrid"
## [47] "rec_016_AS_id_032_5.TextGrid"
## [48] "rec_016_AS_id_032_6.TextGrid"
## [49] "rec_016_IS_id_031_1.TextGrid"
## [50] "rec_016_IS_id_031_2.TextGrid"
## [51] "rec_016_IS_id_031_3.TextGrid"
## [52] "rec_016_IS_id_031_4.TextGrid"
## [53] "rec_016_IS_id_031_5.TextGrid"
## [54] "rec_016_IS_id_031_6.TextGrid"
```

It can be seen that there are multiple variations of the extension. We therefore need to repeatedly erase all the extensions. We will split the different parts of the extension in order to capture all possible combinations between the subparts of the extensions.

```
# apply changes
SpeakerInfo$filename = gsub("'", "", SpeakerInfo$filename)#replace by empty
DFbig$filename.orig = DFbig$filename#store the original fn.
DFbig$filename = gsub("_CUT_", "", DFbig$filename)
DFbig$filename = gsub("20MIN", "", DFbig$filename)
DFbig$filename = gsub("10min", "", DFbig$filename)
DFbig$filename = gsub("_1.TextGrid", "", DFbig$filename)
DFbig$filename = gsub("_2.TextGrid", "", DFbig$filename)
```

```
DFbig$filename = gsub("_3.TextGrid", "", DFbig$filename)
DFbig$filename = gsub("_4.TextGrid", "", DFbig$filename)
DFbig$filename = gsub("_5.TextGrid", "", DFbig$filename)
DFbig$filename = gsub("_6.TextGrid", "", DFbig$filename)
```

Task

We need to get rid of copy+pasting behavior during our writing scripts, as this burns our time. Also, copy+pasting increases the probability of typos. Write a script, in which the substitution using `gsub()` is performed iteratively in a `for()` loop.

We need to check how our commands performend. Make it a habbit to check how your new script affects your data. In this way, you reduce the necessity to debug (we will talk about this later).

```
unique(DFbig$filename)

## [1] "rec_001_AS_id_002" "rec_001_IS_id_001" "rec_003_AS_id_006"
## [4] "rec_003_IS_id_005" "rec_004_AS_id_008" "rec_004_IS_id_007"
## [7] "rec_005_AS_id_010" "rec_005_IS_id_009" "rec_016_AS_id_032"
## [10] "rec_016_IS_id_031"

#well, that worked nicely!
```

Now that we have dealt with the matching problem, we can look for the appropriate rows using `==`. Using the new vector "whichrows" as indices, we access "DFbig" and check its content. Again, we perform the process for one entry,

```
# perform searching and matching
irow = 1 #dummy coding
info = SpeakerInfo[irow,] #extract information
whichrows = DFbig$filename == info$filename #find rows
head(DFbig[whichrows,])

##          filename word  w.start   w.end wordduration
## 1 rec_001_AS_id_002 <P>  0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002  hä 18.28574 18.48954    0.2037911
## 3 rec_001_AS_id_002 was 18.48954 18.73572    0.2461807
## 4 rec_001_AS_id_002 <P> 18.73572 19.04996    0.3142447
## 5 rec_001_AS_id_002  hä 19.04996 19.54812    0.4981559
## 6 rec_001_AS_id_002 <P> 19.54812 21.07000    1.5218827
##
##          filename.orig
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 3 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 4 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 5 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 6 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
```

then include it into the `for()` loop. Using "whichrows" we can access "DFbig" and copy the new information from "info". For this we setup dummy columns with `NA` entries. We can monitor the progress of the iteration using the function `cat()`, which concatenates strings and variables, which need to be separated by commas. In order break the line, you can include the line break command `'\n'` at the end of the concatenated string.

```
#Set up empty columns
DFbig$Sex = NA
DFbig$Birthyear = NA
DFbig$SizeInCm = NA
DFbig$WeightInKg = NA
DFbig$RecordingDateDDMMYYYY = NA

#for-loop
for (irow in 1:nrow(SpeakerInfo)){

  cat('Processed Speaker ', irow, '/', nrow(SpeakerInfo), '\n')

  info = SpeakerInfo[irow,]#access rows in SpeakerInfo
  whichrows = DFbig$filename == info$filename #find rows

  #attach information
  DFbig$Sex[whichrows] = info$Sex
  DFbig$Birthyear[whichrows] = info$Birthyear
  DFbig$SizeInCm[whichrows] = info$SizeInCm
  DFbig$WeightInKg[whichrows] = info$WeightInKg
  DFbig$RecordingDateDDMMYYYY[whichrows] = info$RecordingDateDDMMYYYY
}

## Processed Speaker 1 / 72
## Processed Speaker 2 / 72
## Processed Speaker 3 / 72
## Processed Speaker 4 / 72
## Processed Speaker 5 / 72
## Processed Speaker 6 / 72
## Processed Speaker 7 / 72
## Processed Speaker 8 / 72
## Processed Speaker 9 / 72
## Processed Speaker 10 / 72
## Processed Speaker 11 / 72
## Processed Speaker 12 / 72
## Processed Speaker 13 / 72
## Processed Speaker 14 / 72
## Processed Speaker 15 / 72
## Processed Speaker 16 / 72
## Processed Speaker 17 / 72
## Processed Speaker 18 / 72
## Processed Speaker 19 / 72
## Processed Speaker 20 / 72
```

```
## Processed Speaker 21 / 72
## Processed Speaker 22 / 72
## Processed Speaker 23 / 72
## Processed Speaker 24 / 72
## Processed Speaker 25 / 72
## Processed Speaker 26 / 72
## Processed Speaker 27 / 72
## Processed Speaker 28 / 72
## Processed Speaker 29 / 72
## Processed Speaker 30 / 72
## Processed Speaker 31 / 72
## Processed Speaker 32 / 72
## Processed Speaker 33 / 72
## Processed Speaker 34 / 72
## Processed Speaker 35 / 72
## Processed Speaker 36 / 72
## Processed Speaker 37 / 72
## Processed Speaker 38 / 72
## Processed Speaker 39 / 72
## Processed Speaker 40 / 72
## Processed Speaker 41 / 72
## Processed Speaker 42 / 72
## Processed Speaker 43 / 72
## Processed Speaker 44 / 72
## Processed Speaker 45 / 72
## Processed Speaker 46 / 72
## Processed Speaker 47 / 72
## Processed Speaker 48 / 72
## Processed Speaker 49 / 72
## Processed Speaker 50 / 72
## Processed Speaker 51 / 72
## Processed Speaker 52 / 72
## Processed Speaker 53 / 72
## Processed Speaker 54 / 72
## Processed Speaker 55 / 72
## Processed Speaker 56 / 72
## Processed Speaker 57 / 72
## Processed Speaker 58 / 72
## Processed Speaker 59 / 72
## Processed Speaker 60 / 72
## Processed Speaker 61 / 72
## Processed Speaker 62 / 72
## Processed Speaker 63 / 72
## Processed Speaker 64 / 72
## Processed Speaker 65 / 72
## Processed Speaker 66 / 72
## Processed Speaker 67 / 72
## Processed Speaker 68 / 72
## Processed Speaker 69 / 72
## Processed Speaker 70 / 72
```

```
## Processed Speaker 71 / 72
## Processed Speaker 72 / 72
```

After the process has been performed, we check the result.

```
#Check result of for-loop
head(DFbig)

##          filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002 <P> 0.00000 18.28574    18.2857449
## 2 rec_001_AS_id_002 hä 18.28574 18.48954    0.2037911
## 3 rec_001_AS_id_002 was 18.48954 18.73572    0.2461807
## 4 rec_001_AS_id_002 <P> 18.73572 19.04996    0.3142447
## 5 rec_001_AS_id_002 hä 19.04996 19.54812    0.4981559
## 6 rec_001_AS_id_002 <P> 19.54812 21.07000    1.5218827
##          filename.orig      Sex Birthyear Si...
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"      1990 ...
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"      1990 ...
## 3 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"      1990 ...
## 4 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"      1990 ...
## 5 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"      1990 ...
## 6 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"      1990 ...
##      WeightInKg RecordingDateDDMMYYYY
## 1           70      "28032014"
## 2           70      "28032014"
## 3           70      "28032014"
## 4           70      "28032014"
## 5           70      "28032014"
## 6           70      "28032014"
```

We also check, if all rows contain the necessary information or if something is missing. We do this using two new functions, `is.na()`, which checks whether there are any **NA** in the vector, and `sum`, which sums the numbers in a vector (which is basically an addition function). We hope that the result of this calculation is zero, as the new information should erase all the previous NAs. For this, we use the function `complete.cases()`, which provides a **FALSE**, whenever there is missing information in the dataframe in the form of **NA** or **NaN**, otherwise a **TRUE**.

```
complete.cases(NA)

## [1] FALSE

complete.cases(NaN)

## [1] FALSE

complete.cases(1)

## [1] TRUE

complete.cases('a')

## [1] TRUE
```

We thus need to inverse the result using the negation operator ! :

```
# check deletion of NAs
nrow(DFbig)

## [1] 70958

sum(!complete.cases(DFbig))

## [1] 9673
```

Well, it is clearly not the case that all the NAs we set up above have been erased. This opens a completely new topic, namely debugging, namely: finding the error which made this happend. But before we debug our code, we first, of course, store our results so far

```
setwd(DIR)
save(DFbig, file = 'DF_textgrids_addedinformation.rda')
```

Task

Think about it: What could possibly be a reason for the missing information?

6.3 Excursion: Debugging, part one

Ok, our script does not perform the way we want it to perform. We either quit our job or start what good programmers do: **Debugging** . Actually, it is said that debugging requires almost 30% of your programming time. I don't know whether this is true, but you need to know how to locate the source of your problem and how to solve it. Alright, Sherlock, buckle up, this will be an interesting ride. We already have a hint, how to find our problem: the **NA** s.

First, we need to find the columns responsible for the problem. We can do this using the function `summary` , which provides a short summary of what is happening in the dataframe.

```
summary(DFbig)
```

filename	word	w.start	w.en...
Length:70958	Length:70958	Min. : 0.0	Min. :...
Class :character	Class :character	1st Qu.: 170.3	1st Qu.:...
Mode :character	Mode :character	Median : 341.2	Median :...
		Mean : 370.5	Mean :...
		3rd Qu.: 512.1	3rd Qu.:...
		Max. :1240.7	Max. :...
			...
wordduration	filename.orig	Sex	Birt...
Min. : 0.00005	Length:70958	Length:70958	Min. :...
1st Qu.: 0.13906	Class :character	Class :character	1st Qu.:...
Median : 0.22825	Mode :character	Mode :character	Median:...
Mean : 0.52718			Mean :...

```

3rd Qu.: 0.40000
Max.    :65.55813
3rd Qu...
Max.    ...
NA's    ...

##      SizeInCm      WeightInKg      RecordingDateDDMMYYYY
## Min.      :159    Min.      :51.0    Length:70958
## 1st Qu.:168    1st Qu.:60.0    Class :character
## Median :170    Median :65.0    Mode  :character
## Mean    :173    Mean     :64.2
## 3rd Qu.:180    3rd Qu.:68.0
## Max.    :190    Max.     :86.0
##

```

Well, it seems as though column "Birthyear" is the culprit. Next, we need to find the rows, in which there are still NAs left.

```

#store rows
whereNA = is.na(DFbig$Birthyear)
sum(whereNA)

## [1] 9673

```

Now, one possible reason for why there is missing data in "DFbig" is that eventhough we took great care to match the filenames in "SpeakerInfo" and those in "DFbig", there are still some mismatches. In order to find for which filenames there is missing data, we need to extract the appropriate filenames using "whereNA" and then reduce the amount of duplicated data using `unique()` (since we dont want to look at 9673 entries).

```

NoInfoHere = unique(DFbig$filename[whereNA])
NoInfoHere

## [1] "rec_016_IS_id_031"

```

Well, this is great, we did a good job above. It is only 1 filename which causes the problems. We first need to rule out the possibility that the filename "rec_016_IS_id_031" may be even missing in "SpeakerInfo". We first search "NoInfoThere" in the vector containing the filenames in "SpeakerInfo" using `==`, which gives us a Boolean vector and then sum the vector up using `sum()`. If the sum is zero, the name is not in "SpeakerInfo" (and we have a problem). If the sum is larger than zero, well, then it is in there.

```

IsItThere = SpeakerInfo$filename == NoInfoHere
head(IsItThere)

## [1] FALSE FALSE FALSE FALSE FALSE FALSE

sum(IsItThere)

## [1] 1

```

Nice, the result is larger than one. Hence, it is there. Another problem down, still one to go. Why does this filename result in an NA? Another possibility is that the information

about the birth year was not in "SpeakerInfo" in the first place. For this we need to extract the appropriate row in "SpeakerInfo" using the index in "IsItThere" and then check what information it contains.

```
str(SpeakerInfo[IsItThere,])

## 'data.frame': 1 obs. of 7 variables:
## $ SubjID : int 31
## $ Sex : chr "\"weiblich\""
## $ Birthyear : int NA
## $ RecordingDateDDMMYYYY: chr "\"11122014\""
## $ SizeInCm : int 168
## $ WeightInKg : int 65
## $ filename : chr "rec_016_IS_id_031"
```

Ok, the mystery is solved. "SpeakerInfo" does not contain any birth year information for speaker "rec_016_IS_id_031". The experimenter messed up and he or she did not check the speaker's input. This is unfortunate, but there is nothing we can do.

Chapter 7

Adding more speaker information to the dataframe: Transforming information in columns

So far, so good. Looking at the `head()` of "DFbig", it becomes clear that there is still work to do.

```
setwd(DIR)
load('DF_textgrids_addedinformation.rda', verbose = TRUE)

## Loading objects:
##   DFbig

head(DFbig)

##           filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002 <P>  0.00000 18.28574    18.2857449
## 2 rec_001_AS_id_002  hä 18.28574 18.48954     0.2037911
## 3 rec_001_AS_id_002 was 18.48954 18.73572     0.2461807
## 4 rec_001_AS_id_002 <P> 18.73572 19.04996     0.3142447
## 5 rec_001_AS_id_002  hä 19.04996 19.54812     0.4981559
## 6 rec_001_AS_id_002 <P> 19.54812 21.07000     1.5218827
##           filename.orig    Sex Birthyear Si...
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"    1990    ...
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"    1990    ...
## 3 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"    1990    ...
## 4 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"    1990    ...
## 5 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"    1990    ...
## 6 rec_001_AS_id_002_CUT_20MIN_1.TextGrid "männlich"    1990    ...
##   WeightInKg RecordingDateDDMMYYYY
## 1          70          "28032014"
## 2          70          "28032014"
## 3          70          "28032014"
## 4          70          "28032014"
## 5          70          "28032014"
## 6          70          "28032014"
```

- 1) The 'Sex' column contains quotes.
- 2) Unless you are a speaker of German (or any related language), you probably do not understand the German adjectives in that column.
- 3) The recording date is in a weird string format.
- 4) We still lack the age of the speaker at the time of recording.

7.1 Problem 1: Deleting unnecessary characters

Let us tackle problem 1: The quotes. We have seen in the previous section that characters can easily be removed from a string using `gsub()` .

```
unique(DFbig$Sex)

## [1] "\"männlich\"" "\"weiblich\""

DFbig$Sex = gsub('"', "", DFbig$Sex)
unique(DFbig$Sex)

## [1] "männlich" "weiblich"
```

7.2 Problem 2: Replacing information

Next stop: The German adjectives. First we need to know what "Sex" actually contains, then we can change its content:

```
unique(DFbig$Sex)

## [1] "männlich" "weiblich"
```

In German, *männlich* means "male" and *weiblich* means female. What we want to do is to change the columns content to "male", if it was previously *männlich* and to "female" otherwise (makes sense, doesn't it?). R provides a function which does exactly such a thing: `ifelse()` , which needs three arguments. The first argument in `ifelse()` is the test, usually a procedure resulting in a Boolean vector; the second argument contains the entry, which should be inserted if the Boolean vector contains TRUE; the third argument contains the entry, which should be inserted if the Boolean vector contains FALSE.

```
DFbig$Sex2 = ifelse(DFbig$Sex == 'weiblich', 'female', 'male')
head(DFbig, 2)

##           filename word  w.start   w.end wordduration
## 1 rec_001_AS_id_002  <P> 0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002   hä 18.28574 18.48954    0.2037911
##           filename.orig      Sex Birthyear Size...
```

```

1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
##   RecordingDateDDMMYYYY Sex2
## 1              "28032014" male
## 2              "28032014" male

```

7.3 Problem 3: Extracting information out of strings

Next problem: The format of the recording date. This is a little bit tricky. The vector contains strings, not numbers (which is also indicated by the quotes. You might wonder, why this is a problem. Well, the creators of this great corpus (among others me), did not bother to provide you with the age of the speakers but only with their birthyear and the date of recording. Thus, we need to calculate the age on our own (one of many, many, many calculations).

```

class(DFbig$RecordingDateDDMMYYYY)

## [1] "character"

```

Since the information is a string, we cannot perform any calculations with the date. We first must transform the date into something numeric and, more specific, readable as "days", "months" and "years". Fortunately, the column name tells us that the structure of the information is as follows: two positions for days (DD), two positions for months (MM) and four positions four years (YYYY). The easiest thing to do would be to directly access the string and extract the appropriate numbers (as you would do with a vector). Unfortunately, this is not possible in R. But there are possibilities to work around that. I will present you two of these.

7.3.1 Possibility one: Splitting the string

Splitting strings is actually the more tedious solution as it incorporates several sub-steps. Nevertheless, I will present it here, because this section teaches you a very valuable lesson, namely, how to split strings and how to work with the splitted result. There are a couple of in-built functions in R which facilitate the work with strings (but really, only a couple). One of them is the function `strsplit()`, which needs two arguments. Argument one is a vector of strings (which can contain one or many entries); argument two is the character at which the string has to be split. The following example illustrates, how `strsplit()` works, using the 'X' as the character at which the string has to be split.

```

Ex1 = 'ThisXisXanXexample'
Res  = strsplit(Ex1, split = 'X')
Res

## [[1]]
## [1] "This"      "is"        "an"        "example"

```

`strsplit()` can also be applied to a whole vector.

```

Ex2 = c('ThisXisXexampleXone', 'ThisXisXexampleXtwo')
Res  = strsplit(Ex2, split = 'X')
Res

## [[1]]
## [1] "This"      "is"        "example"   "one"
##
## [[2]]
## [1] "This"      "is"        "example"   "two"

class(Res)

## [1] "list"

```

You might have noticed that the output of "Res" differs a little bit from what we are used to. Every entry from "Ex1" is preceded by a number on a double square bracket, e.g. `[[1]]`. `class()` tells us that "Res" is a **List**. This is a variable class we have not encountered before. Before we carry on, we need to do another excursion.

7.3.2 Excursion: What are lists?

Lists allow you to combine various variables into one variable. The advantage of lists is that, in contrast to dataframes, the size and also the class of each entry can vary. Lists are created in a similar way like dataframes, using `list()`. Above, the entries were indexed by numbers. However, list entries can also have names.

```
L1 = list(1, 2, 3)
L1

## [[1]]
## [1] 1
##
## [[2]]
## [1] 2
##
## [[3]]
## [1] 3

L2 = list(a=c(1,2), b='b1', c=3)
L2

## $a
## [1] 1 2
##
## $b
## [1] "b1"
##
## $c
## [1] 3
```

The structure of lists can be assessed using `length()`, `str()` and `names()`.

```
length(L1)

## [1] 3

length(L2)

## [1] 3

str(L1)

## List of 3
## $ : num 1
## $ : num 2
## $ : num 3

str(L2)
```

```
## List of 3
## $ a: num [1:2] 1 2
## $ b: chr "b1"
## $ c: num 3

names(L1) #L1 does not have names

## NULL

names(L2)

## [1] "a" "b" "c"
```

The different naming possibilities allow you to access lists in various forms. Lists are accessed using the double squared brackets: `[[string'/indexnumber]]`.

```
L1[[1]]

## [1] 1
```

7.3.3 Splitting strings. And now the conclusion, part 1

Above, we discovered that we can split strings into their single items. The result is a list. Thus, we first `strsplit()` the strings in column "Birthyear" into their single entries, with an empty entry as the splitting parameter, which then splits all characters, and then extract the necessary positions in a `for()` loop.

```
RecInfo = strsplit(DFbig$RecordingDateDDMMYYYY, split = "")
head(RecInfo,2)

## [[1]]
## [1] "\" "2" "8" "0" "3" "2" "0" "1" "4" "\"
##
## [[2]]
## [1] "\" "2" "8" "0" "3" "2" "0" "1" "4" "\"

length(RecInfo)

## [1] 70958

nrow(DFbig)

## [1] 70958
```

Using `head()`, we can see that every row became now a new entry in the list. Every entry is a single vector in which, with the number of entries is equal to the number of characters in a "Birthyear" row.

The length of the list is equal to the number of rows of `DFbig`. What we can do is, using `for()`, to iterate across the dataframe "SpeakerInfo" line by line and the list entry by entry. Then we extract the appropriate numbers by accessing the character vector and store them to "RecDay" (first two), "RecMonth" (next two) and "RecYear" (last four entries).

```
## Test process
iline = 1 #dummy coding
RecInfo[[iline]]

## [1] "\" "2" "8" "0" "3" "2" "0" "1" "4" "\"

length(RecInfo[[iline]])#gauge number of entries

## [1] 10

DFbig[iline,]

##          filename word w.start    w.end wordduration
## 1 rec_001_AS_id_002 <P>         0 18.28574      18.28574
##                               filename.orig      Sex Birthyear Size...
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
##   RecordingDateDDMMYYYY Sex2
## 1          "28032014" male
```


As you can see above, the number of entries in the date vector is 10. Recall that the date actually had 8. The reason for this is that the date, since it was included as a string and left as it is, is surrounded by parentheses. We need to take this into account when indexing the date vector "RecDate".

```
## Test process
iline = 1 #dummy coding
RecDate = RecInfo[[iline]]
RecDate[2:3] #DD

## [1] "2" "8"

RecDate[4:5] #MM

## [1] "0" "3"

RecDate[6:9] #YYYY

## [1] "2" "0" "1" "4"
```

Note that the entries in "RecDate" still are characters. More importantly, they are single independent strings. We need to find a way to do two things:

- 1) Collapse the characters so that 2, 8 becomes "28".
- 2) Change the string "28" to the number 28.

We will perform the two outline processes using our dummy before we implement the script into a `for()` loop. Let us tackle 1 from above. Let us tackle collapsing first. Collapsing strings ... (I call it here collapsing because "concatenation" is already reserved for whatever `c()` does. I don't know another appropriate term. If you know one, let me know!). OK, collapsing strings can be done using the `paste()` function with its parameter `collapse`. You can provide `collapse` whatever string you want. But in order to merge two strings without anything inbetween, you need to provide it with an empty string.

```
## Test process
iline = 1 #dummy coding
RecDate = RecInfo[[iline]]
paste(RecDate[2:3], collapse="") #DD

## [1] "28"

paste(RecDate[4:5], collapse="") #MM

## [1] "03"

paste(RecDate[6:9], collapse="") #YYYY

## [1] "2014"
```

Well, that looks nice, doesn't it? Next we need to transform them into numbers. Otherwise we won't be able to calculate how old participants were. Or have you ever succeeded in subtracting a number from a string? Yes? Really? Try it out!

```
'a'-1

## Error in "a" - 1: non-numeric argument to binary operator
```

Next, we change the strings into numbers using `as.numeric()` .

```
## Test process
iline = 1 #dummy coding
RecDate = RecInfo[[iline]]
#do it for days
DD= paste(RecDate[2:3], collapse="") #DD
class(DD)

## [1] "character"

DD = as.numeric(DD)#rewrite DD
class(DD)

## [1] "numeric"

#do it for the rest
MM = as.numeric(paste(RecDate[4:5], collapse="")) #MM
YY = as.numeric(paste(RecDate[6:9], collapse="")) #YYYY
MM

## [1] 3

YY

## [1] 2014
```

Having understood how the process works, we need to use it to include the date, the month and the year into our original dataframe "DFbig". The solution here is to use a `for()` loop, which iterates across the entire dataframe "DFbig" and simultaneously attaches the newly created values. Since both, "RecDate" and "DFbig" have the same number of rows, it does not matter on whose bases we perform the process. Although the two processes below (a and b) could be performed simultaneously (by directly storing the result of `as.numeric()` to "DFbig"), I like to keep such things apart as they allow me inspect the result of my code in a more clear way.

```
#Prepare DFbig
DFbig$RecDD = NA
DFbig$RecMM = NA
DFbig$RecYY = NA
```

```

for(iline in 1:nrow(DFbig)){

  # a) get splitted date and create numbers
  RecDate = RecInfo[[iline]]
  DD= as.numeric(paste(RecDate[2:3], collapse=""))
  MM = as.numeric(paste(RecDate[4:5], collapse=""))
  YY = as.numeric(paste(RecDate[6:9], collapse=""))

  # b) attach values to DFbig
  DFbig$RecDD[iline] = DD
  DFbig$RecMM[iline] = MM
  DFbig$RecYY[iline] = YY

}#frame

## Final test
head(DFbig,2)

##           filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002  <P>  0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002   hä 18.28574 18.48954    0.2037911
##           filename.orig    Sex Birthyear Size...
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich    1990    ...
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich    1990    ...
##   RecordingDateDDMMYYYY Sex2 RecDD RecMM RecYY
## 1           "28032014" male    28     3  2014
## 2           "28032014" male    28     3  2014

```

That was the tedious solution, which you will have to perform from time to time. The next section shows you an easier solution.

7.3.4 Possibility two: Clipping or Trimming

Trimming or **Clipping** means that you cut away a certain part of your string. Unfortunately, there is no in-built function in R which performs this (or at least, none that I know of) because R is not such a good processing tool for strings.

Wouldn't it be great if there was somewhere a function to do this? Maybe someone else already made the effort and wrote one.

7.3.5 Excursion: Installing and loading packages

Guess what: There is indeed! There is a large online database (called **Repository**) where dozens of functions are stored in the format of a **Package** (also sometimes called **Library**). Packages contain scripts, functions and data other people created and which they decided to share with the world. The package which contains the scripts we need for our purpose is called "stringr". In order to load a package, you first need to install it. There are two possibilities to do this:

- Download and installation directly from the internet.
- Installation from a local hard drive location.

Download and installation directly from the internet needs a stable internet connection (one might assume that this is absolutely logical, but the variability in human logic is insane). In order to do so, you need to use the `install.packages()` function, to which you provide the package name in the form of a string. If you type the following code, usually you will be asked to select a **Mirror**, i.e. the location from which to download your package. Select the one which is nearest to your location.

```
install.packages('stringr')
```

It might happen that the download cannot be performed or is aborted, or R tells you that the package cannot be found in the mirror you provided. One reason for this might be that the package you need is not accessible under the prespecified repositories in your R installation. But this is not the end of the world, because you can specify the repository from which R should download the package.

```
install.packages('stringr', repos="http://mirrors.softliste.de/cran/")
```

The `repos` parameter is also the one to use in case you downloaded the package from somewhere in the internet (e.g. github) and you need to install it from your local harddrive. For this, of course, you need to find out where you saved the package.

```
install.packages('stringr', repos="ReplaceByYourDirectory")
```

Now that the package is successfully installed on your computer, it does not mean that you can "simply" use it in R. No, you need to load the package into your R environment. This is done using the `library()` function. Do not forget to do this at the beginning of every session. Otherwise some of your code will not work.

```
library('stringr')  
library(stringr) #exactly the same thing!
```

7.3.6 Possibility two: Clipping or Trimming, the conclusion

In the "stringr" package, there is a function called `str_sub()`, which allows us to do exactly what we want. It takes the parameters `start` and `end` in which you specify how many positions from left or right you want to clip.

```
#Test str_sub()  
String = 'abcde'  
str_sub(String, start = 2) #start at b  
  
## Error in str_sub(String, start = 2): could not find function "str_sub"
```

Well, what happened there? Of course, we forgot to load the library! OK, new attempt. Parameters `start` and `end` not only specify at which positions you want to start/end clipping, but also whether you want to start from left or from right.

```
library(stringr)
#Test str_sub()
String = 'abcde'
str_sub(String, start = 2)#start 2 positions from left

## [1] "bcde"

str_sub(String, start = -2)#start 2 positions from right

## [1] "de"

str_sub(String, end = 2)#end 2 positions from left

## [1] "ab"

str_sub(String, end = -2)#end 2 positions from right

## [1] "abcd"
```

I usually do something like the code above, whenever I start to use `str_sub()` , since I never, never can remember how to specify where to start and end clipping. Let's get back to our problem, the extraction of days, months and years out of "RecordingDateDDMMYYYY" (You probably already forgot about this, didn't you?). We need to find out, how to specify `start` and `end` in `str_sub()` in order to get the exact positions.

```
library(stringr)
Test = DFbig$RecordingDateDDMMYYYY[1]
Test#there are still those stupid quotes

## [1] "\"28032014\""

#Delete quotes
DFbig$RecordingDateDDMMYYYY = gsub('"', '', DFbig$RecordingDateDDMMYYYY)#
Test = DFbig$RecordingDateDDMMYYYY[1]
Test#Nice

## [1] "28032014"

str_sub(Test, start = 1, end = 2)#Great.

## [1] "28"

str_sub(Test, start = 3, end = 4)#Yeah!!

## [1] "03"

str_sub(Test, start = 5)#Strike!!!

## [1] "2014"
```

The nice thing about functions is that they are applicable to an entire vector, independently how long it is (OK, not independently, of course there are some restrictions like RAM, your life time or the patience of your boy or girlfriend. But R is very capable in handling long vectors). Hence, in the next step we will extract days, months and years out of "RecordingDateDDMMYYYY".

```
library(stringr)
Days = str_sub(DFbig$RecordingDateDDMMYYYY , start = 1, end = 2)
Months = str_sub(DFbig$RecordingDateDDMMYYYY , start = 3, end = 4)
Years = str_sub(DFbig$RecordingDateDDMMYYYY , start = 5)
class(Days)

## [1] "character"
```

We will save some lines of code and transform the output simultaneously into numbers.

```
library(stringr)
Days = as.numeric(str_sub(DFbig$RecordingDateDDMMYYYY , start = 1, end = 2))
Months = as.numeric(str_sub(DFbig$RecordingDateDDMMYYYY , start = 3, end = 4))
Years = as.numeric(str_sub(DFbig$RecordingDateDDMMYYYY , start = 5))
class(Days)

## [1] "numeric"
```

The only thing which is left is: We need to append that values as new columns to "DFbig". Again, in order to save lines, we do this directly, without storing the result of the extraction+changing process in separate variables. Finally, we check the result of the process using `head()`.

```
library(stringr)
DFbig$RecDD = as.numeric(str_sub(DFbig$RecordingDateDDMMYYYY , start = 1, end = 2))
DFbig$RecMM = as.numeric(str_sub(DFbig$RecordingDateDDMMYYYY , start = 3, end = 4))
DFbig$RecYY = as.numeric(str_sub(DFbig$RecordingDateDDMMYYYY , start = 5))
head(DFbig,2)

##           filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002  <P>  0.00000 18.28574    18.2857449
## 2 rec_001_AS_id_002   hä 18.28574 18.48954     0.2037911
##           filename.orig    Sex Birthyear Size...
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich    1990    ...
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich    1990    ...
##   RecordingDateDDMMYYYY Sex2 RecDD RecMM RecYY
## 1           28032014 male    28    3  2014
## 2           28032014 male    28    3  2014
```

Above, we used only three lines of code to perform (compare that to the code above)! What we have done here is called **Optimization**. But this is the topic of section ??, where I will talk more about how to optimize your code to make it faster.

7.3.7 Well now, how old are the speakers?

```
DFbig$Age = DFbig$RecYY - DFbig$Birthyear
unique(DFbig$Age)

## [1] 24 25 23 26 29 21 NA

#DO #NOT #FORGET #TO #SAVE
setwd(DIR)
save(DFbig, file = 'DF_textgrids_addedinformation.rda')
```

Chapter 8

Adding lexical information: Counts

Well, this is it. You know all the tools to clean and preprocess a corpus. You can go out into the world of corpus processing and show what you can do. Ok, just little bit joking. The truth is, there is much more, much more complex techniques. But all boils down to "search and add" processes. However, these can be rather complex, depending on what you want to add. We will tackle that complexity step by step. First we need to recall what information we need to add to our dataframe.

- (1) Time stamps. 😊
- (2) Word durations. 😊
- (3) Speaker information. 😊
- (4) Word frequencies. 😞
- (5) Conditional probabilities. 😞

We have two tasks left before we can give the dataframe to your boss (OK, to me!). Use `head()` to check what you have already accomplished. This makes you feel good since you see how the dataframe grows in size.

```
setwd(DIR)
load('DF_textgrids_addedinformation.rda', verbose = TRUE)

## Loading objects:
##   DFbig

head(DFbig)

##           filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002 <P>  0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002  hä  18.28574 18.48954    0.2037911
## 3 rec_001_AS_id_002 was  18.48954 18.73572    0.2461807
## 4 rec_001_AS_id_002 <P>  18.73572 19.04996    0.3142447
## 5 rec_001_AS_id_002  hä  19.04996 19.54812    0.4981559
## 6 rec_001_AS_id_002 <P>  19.54812 21.07000    1.5218827
##           filename.orig    Sex Birthyear Size...
```


1	rec_001_AS_id_002_CUT_20MIN_1.TextGrid	männlich	1990	...			
2	rec_001_AS_id_002_CUT_20MIN_1.TextGrid	männlich	1990	...			
3	rec_001_AS_id_002_CUT_20MIN_1.TextGrid	männlich	1990	...			
4	rec_001_AS_id_002_CUT_20MIN_1.TextGrid	männlich	1990	...			
5	rec_001_AS_id_002_CUT_20MIN_1.TextGrid	männlich	1990	...			
6	rec_001_AS_id_002_CUT_20MIN_1.TextGrid	männlich	1990	...			
##	WeightInKg	RecordingDateDDMMYYYY	Sex2	RecDD	RecMM	RecYY	Age
## 1	70	28032014	male	28	3	2014	24
## 2	70	28032014	male	28	3	2014	24
## 3	70	28032014	male	28	3	2014	24
## 4	70	28032014	male	28	3	2014	24
## 5	70	28032014	male	28	3	2014	24
## 6	70	28032014	male	28	3	2014	24

8.1 Word frequencies

8.1.1 Where do I get them from?

Word frequencies, also called **Frequencies of occurrences** or **Counts** are nothing but counts of individual word tokens. In other words, how often does *I* occur in a language, *you*, *he*, etc. The first question which should pop up in your mind is: "*What kind of a corpus should I use?*". There is no easy answer to this question. I even doubt that there is any. It starts at the decision: Written or spoken. Fact is that due to their nature written corpora are tremendously larger than spoken corpora. Nowadays it is easier to use an automatic grabber which searches the internet and copy+pastes every file onto your computer than to record a corpus of (spontaneously) spoken language, transcribe and annotate it. It took almost four years to collect the data for the corpus you are using in this manuscript. Furthermore distributions vary between spoken and written language. One example: The German first singular personal pronoun "ja" (Engl. I) is the most often word in the KEC, while in the written corpus SdeWac it is the word "der" (Engl. the, mask.). The reason for this is obvious: In a dialog environment "ja" is used as a dialogue marker. Texts are written in the third person, thus "der". You can imagine that this will affect your analysis and your results. In addition, the genre of the written text will affect your results: A collection of classic books will have a different frequency distribution than a collection of e.g. newspaper articles, which are written in a very distanced and technical tone, or movie subtitles, which focus on the transmission of emotions. There are a couple of scientists figuring out, which corpus to use. If you are uncertain, use several corpora and test their effects in your analyses.

8.1.2 Count words in KEC

Here, we will focus on the counts from the corpus at hand, the KEC. There is an easy way to calculate number of occurrences of an item, namely the function `table()`. The column names of the results represent the counted item, the entries represent the counts.

```
table(rep(11:20,1:10))

##
## 11 12 13 14 15 16 17 18 19 20
##  1  2  3  4  5  6  7  8  9 10
```

In our corpus, we want to count how often a certain word occurs. We apply `table()` , and then immediately sort the results to a decreasing order using `sort()` .

```
W.counts = sort(table(DFbig$word), decreasing = TRUE)
head(W.counts)

##
##  <P>    ja    ich    das    so    und
## 10451  2510  2155  1814  1516  1455
```

Nice, we have counted how often a word occurs, now we need to attach the results to our original dataframe "DFbig". We will use a `for()` loop for this. It will iterate across "W.counts", going to the `iwordth` entry in "W.counts", extracting the column name, looking for the rows, where the column name matches the word in "DFbig" and attaching the `iwordth` count from "W.counts". As always, we try whether this plan works using dummy coding!

```
iword = 1 #index
count = W.counts[iword] #access
count#show

##  <P>
## 10451

which.word = colnames(count)#take colname
which.word#show

## NULL

where = DFbig$word == which.word#find rows
head(where)

## logical(0)
```

Well, this did not work. The reason for this is simple but not trivial. "W.counts" has only one dimension, since it is a vector. And by definition, vectors do not have any column names. R is a little bit picky about this. It is a pain, but you will get used to it. Or will create a new programming language. In order to extract the name from "count", we need to use the function `names()` , without "col" or "row".

```
iword = 1 #index
count = W.counts[iword] #access
count#show

##  <P>
## 10451
```

```

which.word = names(count)#take colname
which.word#show

## [1] "<P>"

where = DFbig$word == which.word#find rows
head(where)

## [1] TRUE FALSE FALSE TRUE FALSE TRUE

```

Good, this worked. Now we include this piece of code into a `for()` loop, in which "iword" will iterate between 1 and the length of "W.counts".

```

DFbig$W.counts = NA #create dummy column

for(iword in 1:length(W.counts)){#iterate
  count = W.counts[iword]#access counts
  which.word = names(count)#take colname
  where = DFbig$word == which.word#find rows in DFbig
  DFbig$W.counts[where] = count
}#end for loop

```

Since there are 6017 entries in "W.counts", you can expect that this process will take a while. This is why we want to have a status update about the process shown in our prompt. However, we don't want an update every iteration, but every 1000th iteration. This is a little bit more complex. For this, we need to use two new functions, namely the operator `%%`, which calculates modulo and `if()` which tests conditions.

8.1.3 Excursion: Testing conditions using `if()`

`if()` is one of the most important functions for programming. It is actually surprising that we start to use it so late in the manuscript. I bet that when you open another programming introduction it will appear on the first 50 pages. Well, remember, this is not just another programming introduction. The input to `if()` is a single entry, namely either TRUE or FALSE. Whenever `if()` is provided a TRUE, it will evaluate the code within defined environment, created by the curly brackets. Whenever it is provided a FALSE, nothing will happen. However, you need to create a command which evaluates to a one entry item containing TRUE or FALSE.

```

V1 = 1:5
V1 == 1

## [1] TRUE FALSE FALSE FALSE FALSE

if(V1 == 1){cat('Done')}

## Warning in if (V1 == 1) {: the condition has length > 1 and only the first element
will be used

## Done

```

If you provide an entire vector to `if()` , it will by default take only the first element, as the warning message indicates (READ THE WARNING MESSAGES!). This means that you actually do not have full control of what is happening. You thus need to reduce the scope of testing.

```
V1 = 1:5
V1 == 1

## [1] TRUE FALSE FALSE FALSE FALSE

if(V1[1] == 1){print('Done')}

## [1] "Done"

if(V1[2] == 1){print('Done')}#nothing happens
```

However, `if()` provides an escape mechanism for cases, where the input does not evaluate to TRUE. This is done by means of `else` . The crucial part of `else` is that it has to appear exactly in the same row as the closing bracket of `if()` , followed by an opening bracket . Again, count the number opening and closing brackets so that you don't forget to close your environments.

```
V1 = 1:5
V1 == 1

## [1] TRUE FALSE FALSE FALSE FALSE

if(V1[1] == 1){print('Done')}

## [1] "Done"

if(V1[2] == 1){
  print('Done')
} else {
  print('None')
}

## [1] "None"
```

However, remember that you do not have to use necessarily `else` . If you don't want anything else to happen, if `if()` is not provided a TRUE, then you simply leave it out.

```
V1 = 1:5
V1 == 1

## [1] TRUE FALSE FALSE FALSE FALSE

if(V1[1] == 1){print('Done')}

## [1] "Done"

if(V1[2] == 1){print('Done')}
```

Again, the input to `if()` are all commands which evaluate to a single entry item containing TRUE or FALSE. This means that it can be TRUE or FALSE itself, but also tests for larger or smaller numbers.

```
if(6 > 5){
  print('larger five')
} else {
  print('smaller five')
}

## [1] "larger five"

#####
if(4 > 5){
  print('larger five')
} else {
  print('smaller five')
}

## [1] "smaller five"
```

Also, you can look whether a vector "V.test" contains any items from vector "V.find". This is done with `%in%`. As you probably do not recall, the output of `%in%` is a Boolean vector "B.vector", i.e. it contains multiple TRUE and FALSE entries.

```
V.test = 1:10
V.find = c(4,3,11)
B.vector = V.test %in% V.find
B.vector

## [1] FALSE FALSE TRUE TRUE FALSE FALSE FALSE FALSE FALSE
```

What you want to know is whether any items from "V.find" is present in "V.test". This means that whenever there is one TRUE in "B.vector", your task is fulfilled and you can be happy. A simple method to query whether there is a TRUE in a Boolean vector is to `sum()` it up.

```
sum(B.vector)

## [1] 2
```

Thus, whenever the `sum()` of the Boolean vector is larger than 0, there is a TRUE, which means that "V.test" contains any items from "V.find". Thus in the next step, you need to test, whether the sum is larger than 0.

```
sum(B.vector)>0

## [1] TRUE
```

This command can be provided to `if()` and tested:

```
V.test = 1:10
V.find = c(4,3,11)
B.vector = V.test %in% V.find

if(sum(B.vector)>0){
print('V.test contains stuff from V.find.')
}

## [1] "V.test contains stuff from V.find."
```

Don't worry, we will use all of the possibilities from above repeatedly in the manuscript. You do not necessarily need to remember all of them!

8.1.4 Now the conclusion: Iterating across W.count with status information

Using **Modulo %%** , we can query whether "iword" matches 1000. The reason for this is that whenever it does not, the output of modulo will be larger than 0.

Modulo shows you the "Rest" after a division.

```
3 %% 2 #test 1

## [1] 1

33 %% 2 #test 2

## [1] 1

33 %% 1000

## [1] 33

1000 %% 1000

## [1] 0

2000 %% 1000

## [1] 0

1000 %% 1000 == 0

## [1] TRUE
```

In the following, I include a `cat()` command, whenever "iword" equals any number which can fully be divided by 1000, thus modulo equals zero. Otherwise nothing happens.

```
DFbig$W.counts = NA #create dummy column
```

```

for(iword in 1:length(W.counts)){#iterate

  if(iword %% 1000 == 0){cat('Processed', iword, 'of', length(W.counts), 'items\n')}#print

  count = W.counts[iword]#access counts
  which.word = names(count)#take colname
  where = DFbig$word == which.word#find rows in DFbig
  DFbig$W.counts[where] = count
}#end for loop

## Processed 1000 of 6017 items
## Processed 2000 of 6017 items
## Processed 3000 of 6017 items
## Processed 4000 of 6017 items
## Processed 5000 of 6017 items
## Processed 6000 of 6017 items

head(DFbig)

##          filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002  <P>  0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002   hä 18.28574 18.48954   0.2037911
## 3 rec_001_AS_id_002  was 18.48954 18.73572   0.2461807
## 4 rec_001_AS_id_002  <P> 18.73572 19.04996   0.3142447
## 5 rec_001_AS_id_002   hä 19.04996 19.54812   0.4981559
## 6 rec_001_AS_id_002  <P> 19.54812 21.07000   1.5218827

          filename.orig      Sex Birthyear Size...
1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990    ...
2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990    ...
3 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990    ...
4 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990    ...
5 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990    ...
6 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990    ...
  WeightInKg RecordingDateDDMMYYYY Sex2 RecDD RecMM RecYY Age W....
1          70          28032014 male   28    3  2014  24    ...
2          70          28032014 male   28    3  2014  24    ...
3          70          28032014 male   28    3  2014  24    ...
4          70          28032014 male   28    3  2014  24    ...
5          70          28032014 male   28    3  2014  24    ...
6          70          28032014 male   28    3  2014  24    ...

```

Isn't that cool? If you want, you can try out the following code in your prompt. I did not evaluate it in the manuscript because, well... you will see.

```

DFbig$W.counts = NA #create dummy column

for(iword in 1:length(W.counts)){#iterate

  if(iword %% 1000 == 0){
    cat('\nProcessed', iword, 'of', length(W.counts), 'items\n')
  }
}

```

```
} else {  
  cat('.'.')  
}  
  
count = W.counts[iword]#acce counts  
which.word = names(count)#take colname  
where = DFbig$word == which.word#find rows in DFbig  
DFbig$W.counts[where] = count  
}  
#end for loop  
  
setwd(DIR)  
save(DFbig, file = 'DF_textgrids_addedinformation2.rda')
```


Chapter 9

Adding lexical information: Probabilities and Conditional Probabilities

Ah, here comes the spooky mathematical word which scared almost all of you (at least those without a mathematical background. Believe me, it totally freaked me out the first time I saw it). Let us first look at the mathematics behind transitional probabilities. I promise, this will be the only time you have to deal with formulas and math in this manuscript. And the math behind transitional probabilities is really easy.

9.1 **Excursion:** Probabilities

Let us start with an easy question: What is a probability and how do we calculate it? Imagine we have the following items in our vector "V.items".

```
V.items = rep(letters[1:10], 1:10)
V.items

[1] "a" "b" "b" "c" "c" "c" "d" "d" "d" "d" "e" "e" "e" "e" "e"...
[18] "f" "f" "f" "f" "g" "g" "g" "g" "g" "g" "g" "h" "h" "h" "h"...
[35] "h" "h" "i" "i" "i" "i" "i" "i" "i" "i" "i" "j" "j" "j" "j"...
## [52] "j" "j" "j" "j"
```

Whenever I randomly select a number out of vector "V.items", how probable is it that I will encounter a "a" or a "b" and so forth. Probabilities are easily calculated, namely by dividing the count of each item (found using `table()`) by the total number of items in our vector "V.items". As a result, probabilities always range between 0 (= no probability) and 1 (=absolute certainty).

```
#calculate number of occurrences
V.count = table(V.items)
V.count

## V.items
```

```
##   a   b   c   d   e   f   g   h   i   j
##   1   2   3   4   5   6   7   8   9  10

#calculate total number of items
TotalNumberOfItems = sum(V.count) # Version 1
TotalNumberOfItems

## [1] 55

TotalNumberOfItems = length(V.items) # Version 2
TotalNumberOfItems

## [1] 55

#Calculate probability of items
P.items = V.count/TotalNumberOfItems
round(P.items,2)

## V.items
##      a      b      c      d      e      f      g      h      i      j
## 0.02 0.04 0.05 0.07 0.09 0.11 0.13 0.15 0.16 0.18
```

In the example above, I used `round()` to round the output of "P.items" to the second digit. "P.items" tells you that you will encounter an "a" with a probability of 0.0181818, a "b" with a probability of 0.0363636, and so forth. That was easy wasn't it?

Now we can calculate the probabilities, also called relative frequencies, to encounter a certain word in our KEC corpus.

```
DFbig$P.word = DFbig$W.counts / sum(DFbig$W.counts)
```

Probabilities have the advantage that they make calculations comparable when values stem from different corpora with a different total amount of items. In mathematical terms, a probability P_i for an item W_i in a vector with N items with $i = 1...N$, is expressed by equation ??:

$$P_i = \frac{C_i}{\sum_{i=1}^n C_i} \quad (9.1)$$

Compare equation ?? with the code above where we calculated the probabilities for the words in "DFbig". You can map the parts of the equation one to one on the code above.

9.2 Excursion: Conditional Probabilities

Conditional probabilities tell you the probability of an item (or event) given some other items (or events). Let me first present you some real life examples for this. For example, the probability that you will carry a cup of coffee when you visit a coffee shop is very high. Hence, the question is, how high is the probability of you carrying a cup of coffee given that you visited a coffee shop? Let us start with the math right from the start and then translate it into R code.

Quantitative linguistics are interested in conditional probabilities of a word W_i given the preceding word W_{i-1} or following word W_{i+1} .

The conditional probability of a word W_i given word W_{i-1} , $P(W_i|W_{i-1})$, is expressed by equation 9.2:

$$P(W_i|W_{i-1}) = \frac{P(W_{i-1} \cap W_i)}{P(W_{i-1})} \quad (9.2)$$

where $P(W_{i-1} \cap W_i)$ denotes the joint probability of W_{i-1} and W_i . The inversed conditional probability of a word W_i given word W_{i+1} , $P(W_i|W_{i+1})$, is expressed by equation 9.3:

$$P(W_i|W_{i+1}) = \frac{P(W_i \cap W_{i+1})}{P(W_{i+1})} \quad (9.3)$$

where $P(W_i \cap W_{i+1})$ denotes the joint probability of W_i and W_{i+1} . The inversed conditional probability is called *inversed* because you calculate the probability of an event given some kind of future events. Basically you want to predict the presence given some event in the future. Physically this makes absolute no sense since we cannot see in the future (Unless you watched too many Star Trek episodes where Kirk travels back in time. Have you actually thought about the fact that the crew always travels back in time but never into the future?). Cognitively, however, it makes sense if you accept the idea that the brain is constantly predicting and planning ahead. However, I don't want to deal with cognitive premises here. I only want to focus on those measures which are necessary and useful for cognitive science and quantitative linguistics.

OK, enough math, now we need to translate the equations 9.2 and 9.3 into code. Look at the equations: In order to calculate conditional probabilities, we first need to have the joint probabilities of two words (preceding + current and current + following). In order to calculate a probability for an item, we first need to find the count for this item (cf. equation ??).

9.3 Bigram counts

A sequence of two words AB , or two items whatever they may be, are called a **Bigram**. Accordingly, a sequence of three words or items ABC is called a **Trigram**, etc. Looking at "DFbig", we see that we do not have any "Bigrams" in our dataframe. We first need to create it.

```
setwd(DIR)
load('DF_textgrids_addedinformation2.rda', verbose = TRUE)

## Loading objects:
##   DFbig

head(DFbig)

##           filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002  <P>  0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002   hä 18.28574 18.48954    0.2037911
## 3 rec_001_AS_id_002  was 18.48954 18.73572    0.2461807
```

Table 9.1: Illustration of bigrams including preceding (AB) and following (BC) entry.

Target	BG_AB	BG_BC
1	NA 1	1 2
2	1 2	2 3
3	2 3	3 4
4	3 4	4 5
5	4 5	5 NA

```
## 4 rec_001_AS_id_002 <P> 18.73572 19.04996 0.3142447
## 5 rec_001_AS_id_002 hä 19.04996 19.54812 0.4981559
## 6 rec_001_AS_id_002 <P> 19.54812 21.07000 1.5218827

      filename.orig      Sex Birthyear Size...
1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
3 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
4 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
5 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
6 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich      1990      ...
WeightInKg RecordingDateDDMMYYYY Sex2 RecDD RecMM RecYY Age W...
1          70          28032014 male    28     3  2014  24  ...
2          70          28032014 male    28     3  2014  24  ...
3          70          28032014 male    28     3  2014  24  ...
4          70          28032014 male    28     3  2014  24  ...
5          70          28032014 male    28     3  2014  24  ...
6          70          28032014 male    28     3  2014  24  ...
```

Let us first think what we want to do here for a second.

- 1) We need to find the preceding words and the following words and their according probabilities for every participant.
- 2) We need to create "Bigrams" using the preceding and the following words.
- 3) We need to calculate their frequency of occurrence.
- 4) Using these counts we then can calculate local conditional probabilities.

9.3.1 Creating bigrams

Finding the preceding and the following word is easy, as for a random word with the index i , the preceding has the index $i - 1$ and the following the index $i + 1$. Table 9.1 illustrates how bigrams based on different directions with the i^{th} word as a target look like (and also, how they can be constructed).

The first entry in the AB bigram contains a NA, because it takes into account "nothing" before the first entry. The same holds true for the last entry in the BC bigram, as the last entry is also followed by "nothing".

We will create word bigrams for every annotation file, as we cannot "bridge" across the file end. This means that we will iterate across files using `for()`. In the following, we denote the word and the according count preceding our target word with "Word.A", and the one following the target word with "Word.C". This is based on the notation of a sequence "ABC", where "B" is the target word. In every iteration we will first extract the rows for

every filename and store them in a temporary variable. Then we create two new variables "A" and "C" on the basis of the "Target" variable which contains the words in the correct sequence. "A" contains the words shifted one position to the front, "B" contains the words shifted one position to the back. In the following code, "NA" from above is replaced by the hashtag "#".

Then, we create the bigrams by concatenating "Word.A" and "Target" or "Target and "Word.B" by the help of `paste()`. Simultaneously, we shift the word counts, so that we know the frequencies of occurrence for "Word.A" and "Word.B".

Finally, we will `rbind()` the temporary variable to a larger variable, "DFbig2". In copying the code, do not forget to test the code between the `for()` brackets using a dummy variable, e.g. `ifn = 1`. In addition, we create two columns containing the counts from the preceding and the following word. We use the same technique to do this as we use to create the shifted word sequences.

```
## create bigrams by filename!

#extract filenames
FN = unique(DFbig$filename)
#create dummy variable
DFbig2 = data.frame()

for (ifn in FN){
  #extract rows for filename
  tmp = DFbig[DFbig$filename == ifn,]

  #store "target words"
  Target = tmp$word

  #shift word location and
  #create preceding/following words
  tmp$Word.A= c('#', Target[1:(nrow(tmp)-1)])
  tmp$Word.C= c(Target[2:nrow(tmp)], '#')

  #concatenate bigrams
  tmp$Bigram.AB = paste(tmp$Word.A, Target)
  tmp$Bigram.BC = paste(Target, tmp$Word.C)

  #shift word location and
  #create preceding/following word counts
  Target = tmp$W.counts
  tmp$Count.A= c(NA, Target[1:(nrow(tmp)-1)])
  tmp$Count.C= c(Target[2:nrow(tmp)], NA)

  #rbind
  DFbig2 = rbind(DFbig2, tmp)
}

head(DFbig2,2)
```

```
##          filename word  w.start    w.end wordduration
## 1 rec_001_AS_id_002  <P>  0.00000 18.28574   18.2857449
## 2 rec_001_AS_id_002   hä 18.28574 18.48954    0.2037911
##          filename.orig    Sex Birthyear Size...
## 1 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich    1990    ...
## 2 rec_001_AS_id_002_CUT_20MIN_1.TextGrid männlich    1990    ...
##      WeightInKg RecordingDateDDMMYYYY Sex2 RecDD RecMM RecYY Age W....
## 1          70          28032014 male    28    3  2014  24    ...
## 2          70          28032014 male    28    3  2014  24    ...
##      Word.A Word.C Bigram.AB Bigram.BC Count.A Count.C
## 1      #      hä      # <P>    <P> hä      NA      18
## 2    <P>     was    <P> hä     hä was    10451    487

tail(DFbig2,2)

##          filename word  w.start    w.end wordduration
## 70957 rec_016_IS_id_031   ja 607.5059 607.6423    0.1363792
## 70958 rec_016_IS_id_031 klar 607.6423 607.9800    0.3377041
##          filename.orig    Sex Birthyear SizeInCm W...
## 70957 rec_016_IS_id_031_6.TextGrid weiblich    NA    168    ...
## 70958 rec_016_IS_id_031_6.TextGrid weiblich    NA    168    ...
##      RecordingDateDDMMYYYY Sex2 RecDD RecMM RecYY Age W.count...
## 70957          11122014 female    11    12  2014  NA    251...
## 70958          11122014 female    11    12  2014  NA     7...
##      Word.C Bigram.AB Bigram.BC Count.A Count.C
## 70957   klar    <P> ja    ja klar    10451    73
## 70958      #    ja klar    klar #     2510    NA
```

9.3.2 Bigram counts

In the next step, we need to append calculate the bigram counts to "DFbig2". For this, we will again use function `table()`. Simultaneously, we transform the result into a dataframe for an easier useability.

```
#Bigram counts
AB.count = data.frame(table(DFbig2$Bigram.AB))
BC.count = data.frame(table(DFbig2$Bigram.BC))
head(AB.count)

##      Var1 Freq
## 1      13
## 2    a a    1
## 3  Aaa <P>    1
## 4    A Aber    1
## 5  ääh Alex    1
## 6 aaha mhm    1

head(BC.count)
```

```
##      Var1 Freq
## 1      13
## 2     a a    1
## 3   Aaa <P>   1
## 4    A Aber   1
## 5   ääh Alex   1
## 6  aaha mhm   1
```

Finally, we need to append the bigram counts to "DFbig2". In preceding chapters, we performed the appending by means of a `for()` loop. Let me be frank here: I did this for the purpose of training. There is a way more efficient method to merge information from two different dataframes. Namely, the function `merge()` performs it equally well. `merge()` takes two dataframes and merges them according to a matching information in a defined row, e.g. "BG_AB". This means that both dataframes need to contain a row with the same name. This is why we need to change the column names in "AB.count" and "BC.count". We store the result in a new variable "DFbig3", in case anything goes wrong (which it sometimes does!).

```
#change column names for AB
colnames(AB.count) = c('Bigram.AB', 'Count.AB')

#merge two dataframes
DFbig3 = merge(AB.count, DFbig2, by = 'Bigram.AB')

#change column names for BC
colnames(BC.count) = c('Bigram.BC', 'Count.BC')

#merge two dataframes
DFbig3 = merge(BC.count, DFbig3, by = 'Bigram.BC')
head(DFbig3,2)
```

	Bigram.BC	Count.BC	Bigram.AB	Count.AB	filename	word	...
1		13	<P>	4	rec_004_AS_id_008	4...	
2		13		13	rec_004_AS_id_008	...	

```

w.end wordduration                                filename.orig ...
1 477.2023      0.2975419 rec_004_AS_id_008_CUT_10min_6.TextGrid m...
2  57.9957      0.3168713 rec_004_AS_id_008_CUT_10min_6.TextGrid m...
  Birthyear SizeInCm WeightInKg RecordingDateDDMMYYYY Sex2 RecDD...
1    1988      179      86          27082014 male    27...
2    1988      179      86          27082014 male    27...
##   RecYY Age W.counts Word.A Word.C Count.A Count.C
## 1  2014  26      36    <P>    10451      36
## 2  2014  26      36      36      36
```

```
dim(DFbig2)#original DF

## [1] 70958    23

dim(DFbig3)#new DF. Everything alright.

## [1] 70958    25
```


Notice two things: First, `merge()` appends the new columns at the beginning. Second, the order of the dataframe is mixed. You can check this easily by typing `head(DFbig2)` into your prompt. This does not generally pose a problem. However, we are used to see things in the correct order, which is why we correct that using `order()` which takes vectors as inputs and provides you with an integer vector. We will order "DFbig3" first according to the filename, then according to the beginning of the words, which are chronologically. We will use the new vector "ord" as indices to create an ordered dataframe.

```
#assess correct order

ord = order(DFbig3$filename, DFbig3$w.start)
head(ord)

## [1] 44873 48415 51521 70327 11536 36165

DFbig3 = DFbig3[ord,]
head(DFbig3,3)
```

	Bigram.BC	Count.BC	Bigram.AB	Count.AB	filename	wo...
44873	<P>	4	<P> <P>	36	rec_001_AS_id_002	<...
48415	<P> hä	15	# <P>	10	rec_001_AS_id_002	<...
51521	<P> ja	1458	<P> <P>	36	rec_001_AS_id_002	<...

```
##          w.end wordduration          filename.orig
## 44873  2.25000          2.25000 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 48415 18.28574          18.28574 rec_001_AS_id_002_CUT_20MIN_1.TextGrid
## 51521  1.98000          1.98000 rec_001_AS_id_002_CUT_20MIN_2.TextGrid
```

	Sex	Birthyear	SizeInCm	WeightInKg	RecordingDateDDMMYY...
44873	männlich	1990	190	70	280320...
48415	männlich	1990	190	70	280320...
51521	männlich	1990	190	70	280320...

```
##          RecDD RecMM RecYY Age W.counts Word.A Word.C Count.A Count.C
## 44873      28    3  2014  24   10451   <P>      10451    36
## 48415      28    3  2014  24   10451    #    hä      NA     18
## 51521      28    3  2014  24   10451   <P>    ja    10451   2510
```

9.3.3 Conditional probabilities

Finally, we can calculate the conditional probabilities in our corpus, namely, $P(B|A)$ and $P(B|C)$, i.e the conditional probability that our target word occurs given the preceding or the following word. We will abbreviate "conditional probabilities" with "cndP".

```
#P(B|A)
DFbig3$cndP.B_A = DFbig3$Count.AB / DFbig3$Count.A

#P(B|C)
DFbig3$cndP.B_C = DFbig3$Count.BC / DFbig3$Count.C
```

```
## save  
setwd(DIR)  
save(DFbig3, file = 'DF_textgrids_addedinformation3.rda')
```


Chapter 10

Preparation for Analysis

You might have already guessed it, but science is about analyzing data. All those experiments, recordings, data collections, extractions out of corpora are for naught if you do not analyze your data. There are plenty ways for analysis, ANOVAs, ANCOVAS, linear mixed-effects regression, non-linear regression, Bayesian statistics, and so forth. For some of those methods the data must fulfill some prerequisites.

10.1 **Excursion:** Inspecting the distribution of your data

One of the premisses of linear and non-linear regression is that the data is normally distributed. In a few moments, I will show you, what "normally distributed" means. However, investigating distributions (and many things more) needs a figure. Therefore we must first learn how to visualize data, or in programmer terms, how to **Plot** things. Plotting data is really easy. It is one of R's main purposes, since science would not be science without visualization. The old stupid saying "A picture says more than a thousand words" is here valid more than anywhere else. People believe figures more than naked numbers, even when they represent exactly the same thing. Moreover, visualization allows you to understand and check the data before you do any statistical analysis. Thus: Plot before analyze!

In the following I will show you how to simply plot data in R. So far, we will plot distributions. More "complex" plots will be considered in sections ??.

We will start with the visualization of data with a **Normal distribution** . Let us create a simple toy dataframe using `rnorm()` , a function which allows you to generate random numbers which are normally distributed with a specified **Standard deviation** around a specified **Mean** . **(STOP: Do you know what a "standard deviation" is? Do you know what "mean" is? This concerns especially not-native English speakers. And no, it does not mean "nasty"!!!. If you answered one of these questions with "no" → ask your favorite search engine! No, not later. Now!).** We can check the properties of the values using `mean()` and, calculating the mean and the standard deviation, respectively. **Have you already checked the meaning of "mean" and "standard deviation"? Yes? Well, carry on, soldier!**

```
NumberOfValues = 100
V.randvalues = rnorm(NumberOfValues, mean = 0, sd = 1)
mean(V.randvalues)
```

```
## [1] 0.225527

sd(V.randvalues)

## [1] 1.024868
```

Note that the mean is *near* zero, the standard deviation *near* one, exactly as we have specified them in `rnorm()` 's parameters (Surprise!). Note further that the values in the manuscript will deviate from the values in your R session, because `rnorm()` creates random values. Next, we want to know how the data is distributed. The **Distribution** of the data informs you about how often a certain value occurs in your variable. For this, we use the `density()` function, which calculates the distribution, and feed the result directly into the `plot()` function.

Another possibility to illustrate the data distribution is the **histogram**, using function `hist()`. It differs from the density plot insofar that it illustrates the number of values in a certain bin. In contrast to `density()` you don't have to call `plot()` in order to see the result. The result can be seen in Figure 10.2.

10.2 Excursion: Transformation of not normally distributed data

For an easier grip to the topic we will start with a toy example for which we will create our own data. For this, we will create a dataframe with three kinds of values: normally distributed and two kinds of skewed distributions. The latter will serve as "prototypes" of not normally distributed data.

```
#crate data frame with differently
#distributed values
DF = data.frame(Normdist = rnorm(1000, mean = 10, sd = 1))
DF$NotNormdist1= exp(DF$Normdist)
DF$NotNormdist2= -1000/DF$Normdist
```

In Figure 10.3, top row, you can see the three different distributions. The plot was created by opening a plotting environment which can contain six Figures, using the `( par()` function. In the `par()` function, there is the `( mfrow)` parameter, which takes a two entry vector, specifying the number of rows and the number of columns in your plot. Next, we apply `attach()` to our dataframe "DF". This makes all its columns available for use without the need to use the precursor "DF\$". Next we plot the three different kinds of distributions.

The top left plot in Figure 10.3 illustrates a normally distributed values as we have already seen above. Normally distributed data has to fulfill plenty properties. Most of them are mathematical in nature and serve as statistical tests for normal distributions such as the Shapiro-Wilk test for normality (function `shapiro.test()`). (I really don't want to go too deep into the mathematics behind normal distributions. But you are free to ask your favorite search engine to ask for more information!)

For visual inspection, it is important that the data is distributed symmetrically around a center, namely the mean. Second, it has to have a bell-like shape. See the top left Figure as a prototype of normally distributed data.

Figure 10.1: *Density plot for "V.randvalues".*

```
plot(density(V.randvalues),
     main = 'Density plot',
     ylab = 'Values in your variable',
     xlab = 'Relative amount given the total number of values')
```

The top central plot in Figure 10.3 illustrates a skewed distribution with a long right tail. A long tail indicates strong outliers, which we want to avoid in our analysis. Check the numbers on the x-axis. They are in a **Scientific notation**, i.e. very large numbers are abbreviated by indicating by how many zeros they are followed (and when they are negative, how many decimal places the number contains). This means "1e+05" equals "100000", "2e+05" equals "200000", "1e-1" equals "0.1", etc. You can see that most of the values are distributed between zero and 100.000. Then some of them increase up to 500.000. This indicates exponentially distributed values (which they indeed are, because we transformed them above using the `exp()` function). Exponentially distributed values, or values with a long right tail, can be transformed into normally distributed values using the `log()` function. The result can be seen in the bottom central plot in Figure 10.3. Taking the log of data reduces the difference between the values especially in the upper ranges. As a result, the

Figure 10.2: *Histogram for a normal distribution of random values. X-axis represents sorted unique values in analyzed variable; Y-axis represents the density at the respective x-point.*

```
hist(V.randvalues,  
main = 'Histogram of\nnormally distributed random values',  
ylab='Density',  
xlab = 'Value bins')
```

Figure 10.3: Histogram for a normal distribution of random values. *X*-axis represents sorted unique values in analyzed variable; *Y*-axis represents the density at the respective *x*-point.

```
attach(DF) #make the content available without the need to use \f

par(mfrow=c(2,3))
plot(density(Normdist), main = 'normal distribution')
plot(density(NotNormdist1), main = 'skewed distribution with a right tail.')
plot(density(NotNormdist2), main = 'skewed distribution with a left tail.')

# solution: Transformation!
frame()#empty figure
plot(density(log(NotNormdist1)), main = 'log transformed data')
plot(density(-1000/NotNormdist2), main = 'inverse transformed data')
```

```
detach(DF) #remove from search path
```

data becomes normally distributed.

The top left plot in Figure 10.3 illustrates a skewed distribution with a (long) left tail. Here you can ignore the values on the x-axis. Such a distribution does not occur very often. But if it does, an inverse transformation can change it into a normal distribution. The result can be seen in the bottom left plot.

Thus, you saw two main possibilities of transformation.

10.3 Excursion: Writing a function

In the next section, we want to investigate the distribution of the values in our variables and how they look like when we transform them. We could write code for every column and every transformation. But we don't have to, since this involves lots of copy and pasting, and therefore, lots of possibilities for errors. What we want to do is to create an environment, which will perform this process iteratively without being prone to errors. This is exactly what a **Function** does. The following lines show a template for functions. The syntax is as follows: You start with a `FunctionName` to which you store your code, defining it as a "function". In the round brackets, you define the input variables you want to provide to your function. Using the curly brackets, you open the environment in which you provide the code for your function. Finally, you use the `return()` command to return the variable of your interest to the "normal" workspace. You can provide **Scalars** to the function (one entry vectors), as well as entire vectors. One important property of functions is that all the variables within the function are not returned to the surface workspace. They remain hidden. Only the "returned" variable becomes visible and storable.

```
FunctionName = function(INPUT){
  OUTPUT <- Do something to (INPUT)
  return(OUTPUT)
}
```

Next, we write a simple example function, which adds the content of two variables, "a" and "b" and returns the outcome of the command stored in "c".

```
MyAddFunction <- function(a,b){
  c = a+b
  return(c)
}

#adding scalars
MyAddFunction(1,2)

## [1] 3

#adding vectors
MyAddFunction(c(1,2,3),c(1,2,3))

## [1] 2 4 6

#store result in variable
D = MyAddFunction(1,2)

#a,b,c not visible in workspace
a

## NULL

b

## Error in eval(expr, envir, enclos): object 'b' not found

c

## function (...) .Primitive("c")
```

In the following, we write a function which will take our dataframe "DF" and a specific column of our dataframe "T" and plot the densities for the raw values, the log-transformed values and the inversely transformed values.

```
MyDensityFunction = function (DF, T){
  par(mfcol=c(1,3))

  Y = DF[,T]
  plot(density(Y,na.rm=TRUE), xlab = T, main = 'raw')

  X = log(Y)
  plot(density(X,na.rm=TRUE), xlab = T, main = 'log')

  X<- -1000/Y;
  plot(density(X,na.rm=TRUE), xlab = T, main = '-1000/Y')
}
```

10.4 Data distribution in our corpus

First, we need to load our data.

```
setwd(DIR)
load('DF_textgrids_addedinformation3.rda', verbose = TRUE)

## Loading objects:
##   DFbig3
```

We first need to do a little bit of preprocessing, before we can plot the distributions. Recall that we are interested in the distributions of counts for words. That means that we can exclude all entries that contain as target words pauses "¡P¡". We look for pauses using the `==` operator, providing a Boolean vector. Also, we want to get rid of all the NAs we got due to the inclusion of speaker information. This can be performed using the `complete.cases()` function, which provides you with a Boolean vector in which TRUE stands for a row in which there are no NAs, and FALSE stands for a row, in which there is an NA in any of the columns. We can combine both Boolean vectors using the `( )` operator, creating the Boolean vector "take", whose content we will use as indices to extract the rows of interest from "DFbig3" and store them in "DFbig4".

```
isNoNaN = complete.cases(DFbig3)
sum(!isNoNaN)

## [1] 9691

take = !DFbig3$word == '<P>' & isNoNaN
sum(take)

## [1] 52452

DFbig4 = DFbig3[take,]
```

Using our custom made density function from above, we want to plot the distributions for all the variables which might be necessary for our analysis. These are the following:

```
VariableOfInterest = c("wordduration", "Age", "Count.A",
"W.counts", "Count.C", "Count.AB", "Count.BC",
"cndP.B_A", "cndP.B_c")
```

Next we check, whether our list in "VariableOfInterest" actually has the correct column names from "DFbig4". We do this using the %in% operator.

```
VariableOfInterest[!VariableOfInterest %in% colnames(DFbig4)]

## [1] "cndP.B_c"
```

Oops, there is a type. "cndP.B_c" should actually be "cndP.B_C".

```
VariableOfInterest = c("wordduration", "Age",
"Count.A", "Count.C", "W.counts", "Count.AB", "Count.BC",
"cndP.B_A", "cndP.B_C")
VariableOfInterest[!VariableOfInterest %in% colnames(DFbig4)]

## character(0)

length(VariableOfInterest)

## [1] 9
```

Allright, no typos there. We can carry on. Since there are so many figures we want to plot, we want to store them onto our hard drive, let's say, in a PDF. This is what we do with the function pdf(), whose parameters specify the filename, the width and the height of the figure in inches. This opens a plotting device, which has to be closed using dev.off() after we have finished plotting. We will iterate across all the entries in "VariableOfInterest" in a for() loop, extract one column name, give it to MyDensityFunction() and plot it.

```
setwd(DIR)
pdf('Distributions.pdf', width = 9, height = 3) #height/width in inches
for (itake in VariableOfInterest){
  MyDensityFunction(DFbig4, itake)
}
dev.off()

## pdf
## 2
```

Now go to your specified directors in "DIR" and open the "Distributions.pdf" file. You can see that the file contains 13 pages, as many variables, as we stored in "VariableOfInterest". You can see that each page contains three plots in one row, as we have specified in MyDensityFunction() using par(). We will go through each of the pages and discuss them.

- 1) **Wordduration:** The raw values are highly skewed, logging them puts them to a distribution resembling a normal distribution. Thus, we need to log them.
- 1) **Age:** No transformation provides any good here. Even though the raw values are not normally distributed, we will keep them as they are. (We will discuss another possible transformation later, in Section ??).
- 1) **Count.A and Count.C:** These show a bimodal distribution in the raw data, which remains after log-transformation. A short inquiry which takes all the rows with Count.A > 8000 will show you that the reason for that spike are the pauses <P> in the preceding word. The same holds true for Count.C for the following word. The only solution to this bimodal distribution would be to exclude these from the analysis IF we wanted to use these predictors in the following analysis. We don't want to, but keep them for the sake of a later example (since I know something you don't know yet).
- 1) **W.counts:** The distribution of the raw values has a long right tail. Although logging the values helps to remedy that, the transformed distribution is far away from a normal one. But that's life, folks. Reality is not uniform. We have to take what we have.
- 1) **Count.AB and Count.BC:** The raw distribution informs us that there is a spike near zero (probably a count of one) and that there is a long right tail. Logging these values changes the distribution, but the values can be considered to be something like "bimodal": lots of ones, and everything else. We will deal with that in a second.
- 1) **cndP.B_C and cndP.B_C:** The logged values resemble something near a normal distribution. Nice!

10.5 Dealing with problematic distributions

10.5.1 Log transform variables where necessary

We use the preamble "L." to indicate that values are logged.

```
DFbig4$L.wordduration = log(DFbig4$wordduration)
DFbig4$L.W.counts = log(DFbig4$W.counts)
DFbig4$L.cndP.B_A = log(DFbig4$cndP.B_A)
DFbig4$L.cndP.B_C = log(DFbig4$cndP.B_C)
```

10.5.2 Multimodal distribution of age

The distribution of age is a problematic one, because we have almost a distribution into six categories and two of the categories (below 22 and above 28) have larger distances from the main group. In order to remedy that shortcoming, we can **rank** the values and making the distances between the values uniform. The following code illustrates how to do that. `as.factor()` creates factors of the integers, which have internally a numeric and ordered

coding. Applying sequentially `as.numeric()` to the factors transforms them into ordered numbers. We store that into the column "R.Age", with "R." representing "ranked". Then we include the distribution of age into our former `Distributions.pdf` plot.

```
DFbig4$R.Age = as.numeric(as.factor(DFbig4$Age))
head(DFbig4$R.Age)

## [1] 3 3 3 3 3 3

## Plot
pdf('Distributions.pdf', width = 9, height = 3) #height/width in inches
  for (itake in VariableOfInterest){
    MyDensityFunction(DFbig4, itake)
  }

par(mfrow=c(1,3))

plot(density(DFbig4$R.Age), main = 'Ranked Age')

dev.off()

## pdf
## 2
```

10.5.3 Distribution of Count.AB and Count.BC

The distributions of "Count.AB" and "Count.BC" indicate that there are lots of counts of ones, and less representations of other values. This results in a strongly skewed distribution, independently of the transformation. One possible solution for this is to transform such a distribution into a categorical predictor: "ones" vs. "largerThanOnes". We will perform this with the `ifelse()` function. We immediately transform these into factors (which are necessary for the statistical analysis) using `as.factor()`. We check the distribution using `table()`.

```
DFbig4$Count.AB.cat = ifelse(DFbig4$Count.AB == 1, "ones", "largerOnes")
DFbig4$Count.BC.cat = ifelse(DFbig4$Count.BC == 1, "ones", "largerOnes")
DFbig4$L.Count.BC.cat = log(DFbig4$Count.BC)

attach(DFbig4)
table(Count.AB.cat)

## Count.AB.cat
## largerOnes      ones
##      33699      18753

table(Count.BC.cat)
```

```
## Count.BC.cat
## largerOnes      ones
##      33245      19207

detach(DFbig4)
```

You can see that the relation of "ones" to "largerOnes" is in both cases almost 3:2.

10.6 Correlations between predictors

If you were thinking, "now I am finished, now I can give my Boss the preprocessed corpus", well, I must disappoint you. There are still two steps to be performed, before we can start with our analysis. The first one is a check whether there exist correlations between our numerical predictors. Correlations between multiple variables in your analysis is also called **Collinearity**. Correlation is denoted with the letter R . **Predictors**? Well, yes, predictors. So far, I have not talked about predictors because we have not talked about analysis yet. But every preprocessing, every analysis is (or should be at least), driven by a hypothesis, that, given a thing X the thing Y will change. In our case, the thing X is the predictor, the thing Y is the dependent variable. More concretely, in a more linguistic approach, a very interesting thing is how word durations change given certain lexical predictors, such as word frequency, transitional probability, etc. However, the premise of linear regression is that the predictors do not correlate with each other. There are plenty mathematical reasons for this, and also fierce fights among scientists how to deal with that problem, which I don't want to bore you with. If you are indeed interested in that topic, you can check ?? and the references therein.

Also, there are plenty solutions for that problem. One possibility is to exclude a predictor which correlates strongly with another one. Correlation values range between -1 and 1. -1 indicates a perfect negative correlation, i.e. high values of X indicate low values of Y ; 1 indicates a perfect positive correlation, i.e. high values of X indicate high values of Y .

We will do this in the following. We first need to calculate a correlation matrix between all our numeric variables of interest using the `cor()` function. There are basically two methods the function provides. The default is the Pearson's product-moment correlation, which calculates the correlation between the raw values; the Spearman correlation on the other hand calculates a rank wise correlation, i.e. first the values are ranked and then correlated. The advantage of this method is that the distribution of the data matters less, as outliers are reincluded into the bulge. We use `round()` to round the values to the second digit so that the output saves space. Recall, that in a correlation matrix rows equal columns and the numbers are mirrored along the diagonal. We will use `upper.tri()` to replace one half of the matrix with empty cells in order to increase readability. The content of variable is displayed in Table ??.

```
#setting variables of interest
VariableOfInterest = c("L.worddduration", "L.W.counts",
"L.cndP.B_A", "L.cndP.B_C", "R.Age", 'Count.C', 'Count.A')

#checking against colnames
```

Table 10.1: Correlation table for the numeric predictors in our dataframe.

	L.wordduration	L.W.counts	L.cndP.B_A	L.cndP.B_C	R.Age	Count.C	Count.A
L.wordduration							
L.W.counts	-0.51						
L.cndP.B_A	-0.2	0.41					
L.cndP.B_C	-0.38	0.43	0.14				
R.Age	0.05	0	-0.01	-0.01			
Count.C	0.22	0	0.02	-0.73	0.01		
Count.A	-0.01	0.11	-0.61	0.04	0.01	0.07	

```
VariableOfInterest[!VariableOfInterest %in% colnames(DFbig4)]

## character(0)

#no types
RR = round(cor(DFbig4[,VariableOfInterest], method = 'spearman'), 2)
RRorig = RR
RR[upper.tri(RR, diag=TRUE)] = ''
```

Inspecting the first column of Table ??, we can see that there are correlations with variable strength between "L.wordduration" and other predictors. The strongest is with "L.W.counts", the weakest ist with "Count.A", which could have been expected.

According to baayen:2008 a threshold where correlations become harmful in a model is ± 0.7 , as then one predictor explains almost 50% of the other predictor's variance ($R = 0.7$, $R^2 = \text{variance explained} = 0.49$). We lower that threshold to 0.6. Inspecting the entire Table ??, we find two correlations higher than an absolute of 0.6: "Count.C" with "L.cndP.B_C" ($R = -0.73$) and "Count.A" with "L.cndP.B_A" ($R = -0.61$) (now you know why I wanted to keep the counts in the pool!).

We need to do the following thing: We select the two correlating predictors and inspect their correlation with the dependent variable, in our case, "L.wordduration", and eliminate the one with the lower correlation from further analysis. Here, the less well correlated predictors are "Count.A" ($R_{\text{Count.A}} = -0.01$ vs. $R_{\text{L.cndP.B_A}} = -0.2$) and "Count.C" ($R_{\text{Count.C}} = 0.22$ vs. $R_{\text{L.cndP.B_C}} = -0.38$).

There are two additional predictors which we did not test for correlation, because they are categorical "Count.AB.cat" and "Count.BC.cat". However, we can perform two simple linear models to inspect whether first, one will predict the other and second whether they both predict any kind of numerical predictor from above. We immediately perform a summary on the model using `summary()`.

```
#model 1
summary(lm(L.cndP.B_A~Count.AB.cat, data = DFbig4))

##
## Call:
## lm(formula = L.cndP.B_A ~ Count.AB.cat, data = DFbig4)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -5.4804 -1.3710  0.0878  1.4715  3.8944
```

```
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)   -3.89445    0.01114  -349.51  < 2e-16
## Count.AB.catones  0.12038    0.01864    6.46 1.06e-10
##
## Residual standard error: 2.045 on 52450 degrees of freedom
## Multiple R-squared:  0.000795, Adjusted R-squared:  0.000776
## F-statistic: 41.73 on 1 and 52450 DF,  p-value: 1.056e-10
```

Indeed, the model above informs us that the difference between the levels "largerOnes" and "ones" is significant, as can be seen in the coefficients table. The same is true in the following as well.

```
#model 2
summary(lm(L.cndP.B_C~Count.AB.cat, data = DFbig4))

##
## Call:
## lm(formula = L.cndP.B_C ~ Count.AB.cat, data = DFbig4)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -5.6896 -1.5166 -0.0136  1.5724  4.4156
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)   -3.56484    0.01209  -294.96  <2e-16
## Count.AB.catones -0.85077    0.02021   -42.09  <2e-16
##
## Residual standard error: 2.219 on 52450 degrees of freedom
## Multiple R-squared:  0.03267, Adjusted R-squared:  0.03266
## F-statistic: 1772 on 1 and 52450 DF,  p-value: < 2.2e-16
```

And in the following as well.

```
#model 2
summary(lm(L.cndP.B_A~Count.BC.cat, data = DFbig4))

##
## Call:
## lm(formula = L.cndP.B_A ~ Count.BC.cat, data = DFbig4)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -5.6870 -1.3671  0.0121  1.4397  4.3430
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)   -3.56741    0.01103  -323.30  <2e-16
```



```
## Count.BC.catones -0.77558    0.01823  -42.53   <2e-16
##
## Residual standard error: 2.012 on 52450 degrees of freedom
## Multiple R-squared:  0.03334, Adjusted R-squared:  0.03332
## F-statistic: 1809 on 1 and 52450 DF,  p-value: < 2.2e-16
```

Surprise, in the following too. This means that we can exclude both categorical predictors from our analysis, since they strongly are predictive of the value distributions in the numeric variables.

```
#model 2
summary(lm(L.cndP.B_C~Count.BC.cat, data = DFbig4))

##
## Call:
## lm(formula = L.cndP.B_C ~ Count.BC.cat, data = DFbig4)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -5.328 -1.547  0.016  1.624  3.926
##
## Coefficients:
##              Estimate Std. Error  t value Pr(>|t|)
## (Intercept)   -3.83594    0.01237  -310.112   <2e-16
## Count.BC.catones -0.09030    0.02044   -4.418    1e-05
##
## Residual standard error: 2.255 on 52450 degrees of freedom
## Multiple R-squared:  0.0003719, Adjusted R-squared:  0.0003529
## F-statistic: 19.52 on 1 and 52450 DF,  p-value: 1e-05
```

```
setwd(DIR)
save(DFbig4, file = 'DF_textgrids_addedinformation4.rda')
```

Chapter 11

Performing pilot regressions and visualizations

```
setwd(DIR)
load('DF_textgrids_addedinformation4.rda', verbose = TRUE)

## Loading objects:
##   DFbig4
```

This chapter shows you how to perform a pilot regression analysis using linear mixed-effects models. The focus is on two points: first how to interpret the output of the model and second, how to visualize the model. It is not an introduction to model criticism, or model building, which are complex topics on their own. Again, if you are interested in statistics you are invited to read [baayen:2008](#).

11.1 Linear mixed-effects models

For doing linear mixed-effect regression, you first need to install the third-party package `lme4` using the following code:

```
install.packages('lme4')
```

The function we need for model is `lmer()`. The function's syntax is presented in the next code section:

```
model = lmer(DependentVariable ~ Predictors +(1|RandomIntercepts), data = data)
```

Let me first explain what I mean by the above code:

- **DependentVariable:** This is the variable whose behavior you want to predict depending on the predictors.
- **Predictors:** Here come the variables which you want to use to predict the dependent variable's behavior. These variables can be continuous (aka. numeric) or categorical (aka factors with levels).

- `RandomIntercepts`: This is where the "mixed" comes from in "mixed-effects". You provide the model information about where the data comes from. Usually we provide information about the participants or speakers (in our case represented by "filename") and the items we use (represented by "word").

Next, we perform a simple analysis. The question is: How does word duration correlate with our predictors? As a categorical predictor we will use "Sex2"

```
library(lme4) #load package

model = lmer(L.wordduration ~
  L.W.counts
  + L.cndP.B_A
  + L.cndP.B_C
  + R.Age
  + Sex2
  +(1|filename)
  +(1|word), data = DFbig4)

print(summary(model), corr = FALSE)

## Linear mixed model fit by REML ['lmerMod']
## Formula: L.wordduration ~ L.W.counts + L.cndP.B_A + L.cndP.B_C + ...
##      Sex2 + (1 | filename) + (1 | word)
## Data: DFbig4
##
## REML criterion at convergence: 67941.1
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -12.3665  -0.5701  -0.0242   0.5306   8.0799
##
## Random effects:
## Groups Name Variance Std.Dev.
## word (Intercept) 0.091461 0.30243
## filename (Intercept) 0.006306 0.07941
## Residual 0.196523 0.44331
## Number of obs: 52452, groups: word, 5552; filename, 9
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) -1.061057  0.065099 -16.30
## L.W.counts -0.161297  0.004215 -38.27
## L.cndP.B_A  0.003453  0.001133  3.05
## L.cndP.B_C -0.046567  0.001081 -43.06
## R.Age 0.023055  0.019019  1.21
## Sex2male -0.007852  0.060126 -0.13
```

Above you see the summary of the table, calculated by providing "model" to function `summary()` (Yes, `summary` is a versatile tool in R). We piped the summary through `print()`

in order to suppress the output of the correlation matrix between the predictors. Let me go through that summary from top to bottom. First you get the name of the model "Linear mixed model fit ...". Next you see the formula you used in the model and the data used (DFbig4). You also get an REML criterion, this Random Effects Maximum likelihood criterion informs you about the goodness of fit. Next we get a presentation of the residuals. They are roughly normally distributed (i.e. roughly the same number for min and max, zero in the for the median). Then we have the summary table for the random effects, for "word" and "filename". These inform you how large the variance is in that group and what their standard deviation is. Finally we have the summary for the fixed effects (our holy predictors).

- You first see "(Intercept)", well, this row serve two purposes: first, it IS the intercept for numeric predictors, such as "L.W.counts". You remember, the intercept indicates at what point the linear vector *intercepts* the y-axis. Here, the intercept estimated to be -1.06 (Column Estimate). Since we logged word durations, we can back transform this value into raw values by $\exp(-1.06) = 0.346$ ms. The standard error (SE) informs us about the variability of the estimate, the t-value (t) informs us about how precise the estimate is (remember: $t = \text{Estimate}/\text{SE}$). The consensus here is that everything larger than absolute 2 indicates that the estimate is significantly not zero (that's important, remember that!). This is the case for the intercept.
- Next line, "L.W.counts" we have an estimate of -0.16. This is the **Slope** of the covariate. In other words, how fast does word duration increase/decrease with one unit of "L.W.counts". The negative value informs us that the slope is negative, i.e. a decrease in word durations is associated with an increase in "L.W.counts". The t-value informs us that the estimate is significant.
- Next line, "L.cndP.B_A", also a numeric predictor, this is why the estimate is a slope, is positive. Small, but positive and the t-value informs us that it is significant. Thus, longer word durations are associated with higher conditional probabilities of our target word given the preceding word.
- Next line, "L.cndP.B_C", also a numeric predictor. The slope is negative, the t-value way beyond absolute 2. This indicates that word durations become shorter, the higher the conditional probability of the target word given the following word.
- Next line, "R.Age", also a numeric predictor, also a slope. However, the t-value is below 2. The correlation between word durations and age can be regarded as not significant.
- Final line, "Sex2male": Well, "Sex2" is a factorial predictor with 2 levels: "male" and "female". What this line indicates the difference between "female", which is implicitly coded in the "(Intercept)", and "male". In other words: how much do you need to add to the intercept (=female), in order to get to the group of males (=Sex2male). In our case, the answer is zero, as the t-value informs us that the difference is not significant.

11.2 Visualizing linear regression

There is a fine possibility to visualize the results of a linear regression analysis and it comes in the form of the `visreg` package. The code is provided below, we perform the visualization exemplarily for "L.W.counts". The plotting function `visreg()` not only plots the estimate but also the data points.

Figure 11.1: *Visualization of the linear mixed-effects model.*

```
library(visreg)
visreg(model, "L.W.counts",
main = "Association between word duration and word counts",
ylab = "log(Word duration)", xlab = "log(Word counts)")
```

11.3 Well, that was it

These are the basic steps to preprocess a corpus. I hope you had fun so far. In the near future, I will add some sections about how to boost the performance of your code (two words:

parallel processing, yeah!) and also how to plot nicer figures. But for give me time for that. Thanks for reading. And thanks for the fish!

Chapter 12

Introduction to error-driven learning

```
NDLnetwork <- function(DFw, cueYpos = 1, ocYpos = 3, lwdM = 15, col = c('red',  
'blue'), lty = c(1, 1), lwd = 1, cex = 1, rad = 0.5){  
  cues = rownames(DFw)  
  outcomes = colnames(DFw)  
  
  wooutcomes = 1:length(outcomes)  
  wocues = 1:length(cues)  
  
  x = c()  
  for (iout in wooutcomes){  
    for (icues in wocues){  
      x = c(x, icues, iout)  
    }  
  }  
  
  rel = data.frame(matrix(x, ncol = 2, byrow = TRUE), stringsAsFactors = FALSE)  
  colnames(rel) = c('cueX', 'oc')  
  rel$ocX = NA  
  rel$cueY = cueYpos  
  rel$ocY = ocYpos  
  
  newX = seq(from = 2, to = length(cues)-1, length = length(outcomes))  
  
  for (iline in 1:nrow(rel)){  
    rel$w[iline] = DFw.weight[rel$cueX[iline], rel$oc[iline]]  
    rel$ocX[iline] = newX[rel$oc[iline]]  
  }  
  
  XLIM = c(0, max(c(rel$cueX, rel$ocX))+1)  
  YLIM = c(cueYpos-0.3, ocYpos+0.3)  
  
  plot(0, 0, pch = '.', xlim = XLIM, ylim = YLIM, xaxt = 'n', yaxt = 'n', ylab = '', xlab = '')  
  arrows(x0 = rel$cueX, y0 = rel$cueY, x1 = rel$ocX, y1 = rel$ocY, col = ifelse(rel$w > 0,  
col[1], col[2]), lty = ifelse(rel$w > 0,  
lty[1], lty[2]), lwd = abs(rel$w)*lwdM*lwd, length = 0)
```



```

par(new=TRUE)
symbols(rel$ocX,rel$ocY,circles =rep(rad,nrow(rel)), bg = 'grey', inches = FALSE, xlim =XLIM
text(rel$ocX,rel$ocY,labels = outcomes[rel$oc], cex = cex)#write outcomes

par(new=TRUE)
symbols(rel$cueX,rel$cueY,circles =rep(rad,nrow(rel)), bg = 'grey', inches = FALSE, xlim =XLIM
text(rel$cueX,rel$cueY,labels = cues[rel$cueX], cex = cex)#write cues

}

#####

QQ = function (DF, T){
  par(mfcol=c(2,3))

  Y = DF[,T]

  qqnorm(Y, main = 'raw', xlab = T, pch = '.'); qqline(Y)
  plot(density(Y,na.rm=TRUE), xlab = T, main = 'raw')

  if(min(Y)<0){X<-log( Y+ abs(min(Y))+0.01)} else {X = log(Y)}

  qqnorm(X, main = 'log', xlab = T, pch = '.'); qqline(X)
  plot(density(X,na.rm=TRUE), xlab = T, main = 'log')

  X<- -1000/(Y);
  qqnorm(X, main = '-1000/Y', xlab = T, pch = '.'); qqline(X)
  plot(density(X,na.rm=TRUE), xlab = T, main = '-1000/Y')
}

```

12.1 Basics

12.1.1 Introduction

Error-driven learning is a computational formalization of discriminative learning (Ramscar et al., 2010, 2013), which is based on a long tradition of research in animal learning (e.g. Kamin, 1968; Pavlov, 2006). Error-driven learning as presented here uses the Rescorla-Wagner learning rule (Rescorla, 1988; Rescorla & Wagner, 1972) implemented in the Naive Discriminative Learning Package (NDL) in R (Arppe et al., 2018). An overview of discriminative learning can be found here (Hoppe et al., 2020); a review of learning theories and their relation to discriminative learning can be found in (Nixon, 2020).

12.1.2 R installation

In order to install R, you can download the latest version from <http://cran.r-project.org/>. It is supported for Windows, Mac and Linux platforms. There are different ways to work with R. On windows platforms, starting R starts a console in which all calculations are performed.

This is similar on Mac platforms. On Linux, R is run in a Terminal. In order to start it, you need to type "R" into your Terminal and hit the `enter` key. If you run R in a separate window, I always advise to write your analysis in a text editor, where you can store your program script. In order to run a script, you can copy+paste it into the R environment. Use a text editor which highlights R specific components such as functions, numbers, etc.

On the other hand you can install **Rstudio** from <https://www.rstudio.com/>. The advantage of **Rstudio** is that it starts both, the R environment as well as an R specific editor. It furthermore supports a set of short cuts which enable a fast and easy interaction between the text editor and the R environment.

12.1.3 NDL installation

There are currently two NDL packages, one available over CRAN, the other still in development. We will come back later to the differences between these two. In order to install NDL1 (the original NDL package), you need to execute the following code in the R prompt.

If you happen to have access to NDL2, you can install it with the following code. NDL2 is currently only available for Linux machines. It is possible that you need to install some additional libraries on your linux system before you can install NDL2. One such library might be the header files from the International Components for Unicode library. For this, open a terminal and install (`libc-uc-dev`) with your common installation command.

Congratulations, you have installed NDL.

12.2 What is NDL? - Modelling learning

NDL models learning by calculating the association strength between **cues** and **outcomes**. Association strengths are represented in a cue-outcome weight matrix. But, what are cues, what are outcomes, what is association strength? You remember Pavlov's dog, don't you? That's the poor dog tortured for science. When presented with food, the dog started to drool. For a certain amount of time, food was presented together with the ring of a bell. After a while, when only the bell was rang, the dog started to drool, expecting food, although no food was presented (Pavlov, 2006). What happened? The dog learned to associate the bell with food. As a result it started to drool in expectation of food when presented with the bell, although food itself was not presented. The bell became a cue for the outcome: food. This is classical conditioning. NDL is capable to model the association strength between the cue – bell – and its outcome – food –.

In the following we will encounter five cue-to-outcome conditions illustrating how the association strength is affected by the respective structure (Figure 12.1).

- A) In the first condition, two outcomes are present. Both outcomes have an equivalent number of independent, non-shared cues associated to them.
- B) In the second condition, three outcomes are present. All the three have an equivalent number of independent cues associated to them.
- C) In the third condition, two outcomes are present. In contrast to the previous conditions, they have a different number of cues associated to them.

Figure 12.1: *Five possible cue-to-outcome conditions. Letters represent outcomes; colored circles around them represent their associated cues with size representing number of cues. See text for discussion of the different conditions.*

- D) In the fourth condition, two outcomes are present. Both outcomes share a set of cues, but also have independent cues associated to them.
- E) In the fifth condition, two outcomes share a set of cues. In one outcome the shared cues are more often than in the other.
- F) In the sixth condition, two outcomes are present. However, outcome B has a set of cues which is a complete subset of cues from outcome A.

We will look at the development of association strength as a function of presentation events without any mathematical details. We will handle these once we have dealt with these five conditions.

12.2.1 Modelling learning with the Rescorla-Wagner function

12.2.2 Condition A

In order to model learning, we need to create learning events. The way the algorithm is implemented in NDL is that it needs at least two cues and two outcomes. We create learning events in R by means of a data frame, in which every row represents one learning event. Cues are stored in the column **Cues**, Outcomes are stored in the column **{Outcomes}**. *bell to food* occurs 1000 times, as does *goodbye to leave*. The important thing here is that these two events alternate. In addition we need to provide the information, how often this event occurred by means of the column **Frequency**. We set the frequency for each event to 1. We will deal with frequencies later. For the classical conditioning, we associate the bell with food. As a second outcome, we choose "leaving" and saying "goodbye" as its cue.

```
### Create Learning Events ###
DF.learn = data.frame(
  Cues=rep(c('bell', 'goodbye'), 1000),
  Outcomes=rep(c('food', 'leave'), 1000),
  Frequency=1,
  stringsAsFactors=FALSE)
head(DF.learn,3)
```

```
##      Cues Outcomes Frequency
## 1    bell      food          1
## 2 goodbye    leave          1
## 3    bell      food          1
```

Now we can run the Rescorla-Wagner algorithm with the learning set using the function `.RescorlaWagner`. The way the function works is that it traces the association weight between one cue and one outcome during the entire learning. Hence, you 1) need to specify your data frame containing the learning events; 2) need to specify how often RescorlaWagner iterates through the data frame (`nruns=1`); 3) you need to tell the function which cue it has to trace while learning the association (`traceCue='bell'`) as well as which outcome to trace (`traceOutcome='food'`); 4) finally, you want to tell the function that it should not randomize the order of your events (`random=FALSE`).

Running the function calculates the association strength between the cue **bell** and the outcome {**food**}. As a side note: sometimes we call association strength also connection strength or connection weight. These terms are interchangeable.

```
### run Rescorla-Wagner###
library(ndl)

## Legend
# DF.learn, #learning dataframe
# nruns = 1, #how often should the learning set be repeated
# traceCue='bell', #which cue should be traced
# traceOutcome="food", #which outcome should be traced
# random = FALSE #do not randomize

DF.learn.rw_food <- RescorlaWagner(
```

```

DF.learn,
nruns = 1,
traceCue='bell',
traceOutcome="food",
random = FALSE
)
class(DF.learn.rw_food)

## [1] "RescorlaWagner"

DF.learn.rw_leave <- RescorlaWagner(
DF.learn,
nruns = 1,
traceCue='bell',
traceOutcome="leave",
random = FALSE
)

### Plot ###
par(mfrow=c(1,2)) #set up number of plotting rows and columns
plot(DF.learn.rw_food, ylimit= c(0,1), lwd = 2, main = 'bell --> food')
plot(DF.learn.rw_leave, ylimit= c(0,1), lwd = 2, main = 'bell --> leave')

```

Figure 12.2: *Modelling the association strength between bell and food using the . See text for a detailed description.*

The result of this learning is illustrated in Figure 12.2, including the necessary code to plot the associations. The x-axis represents the presented events from 1 to 2000 (as each two cue-to-outcome events is presented 1000 times). The y-axis represents the association strength (or connection strength or connection weight). The black curve represents the development of the association strength.

Simultaneously, no association between **bell** and {leave} is learned, as these two never occur together. Of course, **goodbye** and {leave} are associated in the same way as **bell** and {food}, which is why it is not plotted here. But you are welcome to check the development of this association on your own.

You can see that the association strength between cue and outcome increases, until it asymptotes at 1 roughly after 1000 events. The poor dog had to be tortured about 1000 times to fully understand the connection between bell and food. You can see that 1 is represented by a horizontal dashed line; this is the maximal learnability between this cue-to-outcome connection, i.e. how well this connection can be learned. You can affect the development of the connection strength by means of three parameters in the function.

- 1) **lambda** represents the learning rate. It is set by default to 1. Increasing **lambda** will increase the rate with which an association is learned. An increase in **lambda** from 1 to 1.5 is illustrated in Figure 12.3,a. You can see that the curve only asymptotes at 1.5, in contrast to Figure 12.2.
- 2) **alpha** and **beta1** represent the salience of the cue and outcome, respectively. Their default values are 0.1. The maximal salience is a product of **alpha** and **beta1**. Increasing/decreasing them represents higher/lower salience, thus faster/slower learning. An increase in **alpha** to 0.5 is illustrated in Figure 12.3,b. You can see that the curve asymptotes faster than in Figure 12.2.

Figure 12.3: *Increasing the learning rate **lambda** increases the maximal learnability of an association; increasing cue salience **alpha** changes the speed of learning.*

12.2.3 Condition B

In condition B we consider the case where three outcomes have the same number of associated cues. For this, we create a learning set like the one from above, but add a third outcome. What happens now is that the probability of the **bell** to **food** event is lower: while in our first example it occurred every second event, now it occurs every third event. The simple effect is that, compared to the first learning model in Figure 12.2, the association is learned slower (Figure 12.4). Below is the code to create the learning set with a larger number of outcomes and associated cues.

```
### Run model with three outcomes ###
nrepetitions = 1000
Cues = rep(c('bell', 'beep', 'goodbye'), nrepetitions)
Outcomes = rep(c('food', 'food', 'leave'), nrepetitions)
DF.learn3 = data.frame(Cues = Cues, Outcomes = Outcomes, Frequency = 1, stringsAsFactors =
```

Figure 12.4: *Learning the association between **bell** and {**food**} in a learning set, in which this event occurs less often in contrast to the learning set used for Figure 12.2. You can see that the association is learned slower.*

12.2.4 Condition C

In condition C we investigate how the total association strength per cue changes, when an outcome is associated with a larger number of cues. As you can see below, if multiple cues are associated simultaneously to an outcome in one learning event, these cues are separated by an underscore ' '.

```
### Run model with two outcomes ###
nrepetitions = 1000
Cues = rep(c('bell_laugh_reddoor', 'goodbye'), nrepetitions)
Outcomes = rep(c('food', 'leave'), nrepetitions)
DF.learn4 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
  Frequency = 1,
  stringsAsFactors = FALSE)

# learn associations
RW_bell = RescorlaWagner(DF.learn4, traceCue='bell',
  traceOutcome="food", random = FALSE)
RW_laugh = RescorlaWagner(DF.learn4, traceCue='laugh',
  traceOutcome="food", random = FALSE)
RW_door = RescorlaWagner(DF.learn4, traceCue='reddoor',
  traceOutcome="food", random = FALSE)

### Plot ###

par(mfrow=c(1,3))
plot(RW_bell, ylimit= c(0,1), lwd = 2, main = 'bell --> food')
plot(RW_laugh, ylimit= c(0,1), lwd = 2, main = 'laugh --> food')
plot(RW_door, ylimit= c(0,1), lwd = 2, main = 'reddoor --> food')
```

Figure 12.5: *Increasing the amount of associated cues for an outcome reduces the maximal learnability of each cue. Only when their weights are summed up, they provide the maximal association strength possible.*

Figure 12.5 illustrates how the association strength changes, once more cues are associated with one outcome. You can see that the maximal association strength for each of these cues

is roughly 0.33, i.e. one third of the maximal association strength in Figure 12.2. In order to get an as high association strength between the three cues **bell**, **laugh**, **reddoor** in the present condition as in condition A, you need to sum up the weights across all the three cues. In this sense, the maximal learnability of an association depends on the number of cues associated to an outcome (remember: parameterized by λ in function).

The question now arises, what kind of effects does the smaller association strength between each of the three cues and their outcome {food} have. Remember that by modelling an association strength between cue and outcome, we are trying to predict behavior. The dog responds to food by producing saliva. Comparing conditions A and C, our model predicts that when presented only with the **bell**, the dog will produce more saliva in condition A than in condition C, as the association strength in condition C is smaller. In order to obtain the same amount of saliva in condition C as in condition A, we need to present all the three cues together.

12.2.5 Condition D

Figure 12.6: *Two situations in which the color red expresses two different semantics, i.e. it is a cue for two different outcomes*

Conditions A to C are very unique insofar as we rarely encounter outcomes which are associated to cues that are not associated to other outcomes. Quite the contrary, multiple cues share multiple outcomes, as illustrated in Figure 12.1,D. For example, the color red is a completely different cue for the two semantics illustrated in Figure 12.6. I leave it to you to ponder about the two different semantics. We therefore will create in the following code a situation in which the two outcomes {food} and {leave} have a cue in common, **bye**. We will train the function in two `for()` loops, the outer loop iterating across the outcomes to trace, the inner loop iterating across the cues to trace. We will store the six final models (= three cues times two outcomes) in a list. A list is a variable whose inputs can be of different classes and different sizes.

```
### Run model with two outcomes ###
nrepetitions = 1000
Cues = rep(c('bell_bye', 'bye_beep'), nrepetitions)
Outcomes = rep(c('food', 'leave'), nrepetitions)
DF.learn5 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
  Frequency = 1,
  stringsAsFactors = FALSE)

traceOutcomes = c('food', 'leave')
traceCues = c('bell', 'bye', 'beep')

Models = list()
z = 0 # count loops

# learn associatio
for (ioutcome in traceOutcomes){ #outer loop
  for (icue in traceCues){ #inner loop
    z = z+1 #count loops
    Models[[z]] <- RescorlaWagner(DF.learn5,
      traceCue=icue, traceOutcome=ioutcome, random = FALSE)
  }#frame inner loop
```

```

}#frame outer loop

length(Models)#check number of entries in list

## [1] 6

```

The following code plots the different cue-to-outcome combinations. In a loop, the Rescorla-Wagner object is extracted from the list `Models` and stored in the variable `Mod`, which is then used for plotting. As you can see, the RescorlaWagner object has entries `traceCue` and `traceOutcome`, indicating which cue and outcome were traced. There are several differences in association strength between Figure 12.2 and Figure 12.7.

```

par(mfrow=c(2,3))
for (imod in 1:length(Models)){
  Mod = Models[[imod]] #access list of models
  plot(Mod, ylimit= c(-1,1), lwd = 2,
       main = paste(Mod$traceCue, '-->', Mod$traceOutcome) )
  abline(h=0,col='grey') # plot a horizontal line at zero
}

```

Figure 12.7: *Development of association strength from the three cues `bell`, `bye`, `beep` to `{food}` and `{leave}`. In contrast to Figure 12.2, the association strength for the unencountered cues decreases. See text for more details.*

- The maximal **bell**→**food** association strength is lower in Figure 12.7 than in Figure 12.2. This is because {**food**} has more outcomes which share the maximal association strength.
- However, the maximal association strength to {**food**} is not distributed equally across its two associated cue. Rather, the association strength **bye**→**food** is lower than **bell**→**food**. The reason for this is that **bye** is also associated to {**leave**}, to which it has the same association strength. Simultaneously, you can see that **bye** has the same association strength to both {**food**} and {**leave**}. This means that when presented with **bye**, the dog expects both outcomes to the same extent. In conclusion, in addition to distributing the association strength across the cues associated to one outcome, the association strength is affected by whether these cues are also associated to other outcomes.
- In contrast to Figure 12.2, the **beep**→**food** association strength is not flat at zero, but developed a negative association strength. A negative association strength means that a cue is dissociated from an outcome; in other words: when presented with the **beep**, the dog absolutely does not expect {**food**} (but expects you to {**leave**}, as you can see in the lower row).

You are probably wondering where this difference comes from for the **beep**→**food** association, as in both models the **beep** cue does not co-occur with {**food**} as an outcome. At this point, suffice it to say that they become associated indirectly by means of **bye**, which creates a link between **beep** and {**food**}. We will deal with this once we discuss Rescorla-Wagner's formalization of learning in section 12.3.

12.2.6 Condition E

In contrast to condition D, the shared cue **bye** occurs more often for {food}. In addition, it occurs independently from its co-cue **bell**. Figure 12.8 illustrates the development of the association strength for these cues. In contrast to Figure 12.7, **bell**→food loses its informativity. This seems completely counter-intuitive, as one would have expected that **bell** remains a valuable cue. But due to its higher occurrence, **bye**→food becomes the main cue, while it also loses informativity for {leave}.

```
### Run model with two outcomes ###
nrepetitions = 1000
Cues = rep(c('bell_bye','bye', 'bye_beep'),nrepetitions)
Outcomes = rep(c('food','food', 'leave'), nrepetitions)
DF.learn5 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
  Frequency = 1,
  stringsAsFactors = FALSE)

traceOutcomes = c('food', 'leave')
traceCues = c('bell', 'bye', 'beep')

Models = list()
z = 0 # count loops

# learn associatio
for (ioutcome in traceOutcomes){ #outer loop
  for (icue in traceCues){ #inner loop
    z = z+1 #count loops
    Models[[z]] <- RescorlaWagner(DF.learn5,
      traceCue=icue, traceOutcome=ioutcome, random = FALSE)
  }#frame inner loop
}#frame outer loop

length(Models)#check number of entries in list

## [1] 6
```

```

par(mfrow=c(2,3))
for (imod in 1:length(Models)){
  Mod = Models[[imod]] #access list of models
  plot(Mod, ylim=c(-1,1), lwd = 2,
       main = paste(Mod$traceCue, '-->', Mod$traceOutcome) )
  abline(h=0,col='grey') # plot a horizontal line at zero
}

```

Figure 12.8: *Development of association strength from the three cues **bell**, **bye**, **beep** to {food} and {leave}. In contrast to Figure 12.2, the association strength for the unencountered cues decreases. See text for more details.*

12.2.7 Condition F

The final condition F represents a situation in which one outcome is associated to cues which are a subset of a cue set associated to another outcome. The following code creates this cue-to-outcome structure and trains the three different cue-to-outcome associations.

```
### Run model with two outcomes ###
nrepetitions = 1000
Cues = rep(c('bell_beep', 'beep'), nrepetitions)
Outcomes = rep(c('food', 'leave'), nrepetitions)
DF.learn5 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
  Frequency = 1,
  stringsAsFactors = FALSE)

traceOutcomes = c('food', 'leave')
traceCues = c('bell', 'beep')

Models = list()
z = 0 # count loops

# learn association
for (ioutcome in traceOutcomes){ #outer loop
  for (icue in traceCues){ #inner loop
    z = z+1 #count loops
    Models[[z]] <- RescorlaWagner(DF.learn5,
      traceCue=icue, traceOutcome=ioutcome, random = FALSE)
  }#frame inner loop
}#frame outer loop
```

```

par(mfrow=c(2,2))
for (imod in 1:length(Models)){
  Mod = Models[[imod]] #access list of models
  plot(Mod, ylimit= c(-1,1), lwd = 2,
        main = paste(Mod$traceCue, '-->', Mod$traceOutcome) )
  abline(h=0,col='grey') # plot a horizontal line at zero
}

```

Figure 12.9: *Development of association strength from the three cues **bell**, **bye**, **beep** to {food} and {leave}.*

12.2.8 A "not so special" condition: Multiple Cue-to-Outcome connections

Conditions D to F nicely illustrate how the association strength is affected by different cue-to-outcome combinations. Another condition not dealt with so far is one, where multiple outcomes share multiple cues. In the following code we construct a situation with three outcomes, {food}, {leave} and {come}. Two outcomes have always one independent cue and one shared cue. In contrast to previous simulations we are interested how the association strength of the unique cues, i.e. the one not shared, develops.

```
### Run model with two outcomes ###
nrepetitions = 500
Cues = rep(c('bell_bye', 'bye_beep', 'hello_beep'), nrepetitions)
Outcomes = rep(c('food', 'leave', 'come'), nrepetitions)
DF.learn7 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
  Frequency = 1,
  stringsAsFactors = FALSE)

traceOutcomes = c('food', 'leave', 'come')
traceCues = c('bell', 'bye', 'beep', 'hello')

Models = list()
z = 0 # count loops

# learn association
for (icue in traceCues){
  for (ioutcome in traceOutcomes){

    z = z+1 #count loops

    NAME = paste(icue, ioutcome)
    Models[[NAME]] <- RescorlaWagner(DF.learn7,
      traceCue=icue, traceOutcome=ioutcome, random = FALSE)
  }#frame inner loop
}#frame outer loop
```

As expected, **bell** and **hello** develop positive association strength to {food} and {come}, respectively (Figure 12.10). Unexpectedly, **bell** develops a positive association strength for {come}, as does **hello** for {food}. As we have seen in Conditions D to F, they should have developed a negative association strength, as they never occur with these outcomes. We will be able to answer this question, once we have inspected the algorithm of the Rescorla-Wagner learning rule. (Regard this condition as a cliffhanger!).

```

par(mfrow=c(4,3))

for (imod in 1:length(Models)){
  Mod = Models[[imod]] #access list of models
  plot(Mod, ylimit= c(-1,1), lwd = 2,
       main = paste(Mod$traceCue, '-->', Mod$traceOutcome) )
  abline(h=0,col='grey') # plot a horizontal line at zero
}

```

Figure 12.10: Development of association strength from unique cues to outcomes which share

12.3 The math behind the Rescorla-Wagner learning

12.3.1 The Algorithm

In the previous section we have shown you how the combination between cues and outcomes affects the development of their association strength. Some of the effects matched your intuition, some probably did not. It is time to investigate the mathematics behind Rescorla-Wagner learning (?). The learning is formalized in Equation (12.1) and is implemented in the `RescorlaWagner()` function:

$$w_i^{t+1} = w_i^t + \Delta w_i^t \quad \Delta w_i^t = \begin{cases} \text{a) } 0 & \text{if ABSENT}(C_i, t) \\ \text{b) } \alpha_i \beta \left(\lambda - \sum_{\text{PRESENT}(C_k, t)} w_j \right) & \text{if PRESENT}(C_i, t) \ \& \ \text{PRESENT}(O, t) \\ \text{c) } \alpha_i \beta \left(0 - \sum_{\text{PRESENT}(C_k, t)} w_j \right) & \text{if PRESENT}(C_i, t) \ \& \ \text{ABSENT}(O, t) \end{cases} \quad (12.1)$$

Equation 12.1 requires that all possible cues C_i and all possible outcomes O in your learning set are assessed right from the start. In this way it takes into account cue competition. In Formula 12.1, t iterates across all of your learning events which are specified in your learning set. In every t event, it checks whether the cue C which you want to trace and the outcome O you want to trace are present and adjusts the association strength w_i^t between them. w_i^t is adjusted by means of Δw_i^t . The result is used in the next event, w_i^{t+1} . Δw_i^t is calculated in the following way:

- **Condition a)** If the cue in question and the outcome in question are not present during the event, Δw_i^t is assigned zero, i.e. no adjustment happens.
- **Condition b)** If the cue in question and the outcome in question are present during the event, Δw_i^t is calculated by subtracting the sum across all weights present in the learning event from λ (e.g. the weights for **bell** and **bye** for {food} in Condition E).
 λ is the maximum learnability of the association between cues and outcomes. It is set by default to 1. The result is weighted by the product of α , representing the salience of the present cue, and β , representing the salience of the situation in which the outcome can be found. α and β are set by default to 0.01. We have found in multiple experiments that when their product equals to 0.001, the data is fit best, resulting in best predictions and best classifications.
- **Condition c)** If the outcome in question is not present during the event, Δw_i^t is calculated by subtracting the sum across all weights present in the learning event from 0. Again, the result is multiplied by α and β .

Again, λ represents the maximum learnability, i.e. the maximum association strength between an outcome and its cues; α represents the salience of the cue. β represents the salience of the situation, in which the outcome does or does not occur. Let us check this in an example.

12.3.2 Back to the not-so-special condition

Given the algorithm in Equation 12.1, we can try to understand what is happening in the not-so-special condition, in which cues which were never co-occurred with an outcome surprisingly developed a positive association strength. The easiest way to do this is to investigate the math properly at different stages of the learning. We will investigate the cue-to-outcome association strengths after 10, 50 and 100 presentations of every outcome in order to assess exactly, why the positive strength develops. For this, we need to use a new function, `RW.iter()`. Simultaneously, you can see how the Rescorla-Wagner learning algorithm is implemented into the R code.

```

RW.iter = function (cuesOutcomes, Lambda = 1, Alpha = 0.1, Beta = 0.1) {
  allCues = unique(unlist(strsplit(cuesOutcomes$Cues, "_")))
  allOutcomes = unique(unlist(strsplit(cuesOutcomes$Outcomes, "_")))
  M.weights = matrix(0, length(allCues), length(allOutcomes))
  rownames(M.weights)=sort(allCues)
  colnames(M.weights)=sort(allOutcomes)
  theCues = strsplit(cuesOutcomes$Cues, "_")
  theOutcomes = strsplit(cuesOutcomes$Outcomes, "_")
  rate=Alpha*Beta
  knownOutcomes = NULL
  for (i in 1:nrow(cuesOutcomes)) {
    currentCues = theCues[[i]]
    currentOutcome = theOutcomes[[i]]
    toAdd = currentOutcome[!is.element(currentOutcome, knownOutcomes)]
    knownOutcomes = c(knownOutcomes, toAdd)
    Vtotal = colSums(M.weights[currentCues,currentOutcome,drop=FALSE])
    M.weights[currentCues,currentOutcome] =
      t(t(M.weights[currentCues,currentOutcome])
        + (Lambda-Vtotal)*rate)
    otherOutcomes = knownOutcomes[!is.element(knownOutcomes, currentOutcome)]
    if (!is.null(otherOutcomes)){
      Vtotal = colSums(M.weights[currentCues,otherOutcomes,drop=FALSE])
      M.weights[currentCues,otherOutcomes] = t(t(M.weights[currentCues,otherOutcomes])
        + (0-Vtotal)*rate)
    }#end if
  }#end for
  return(M.weights) }

```

Instead of tracing a single cue-to-outcome combination, as has been done before, function `RW.iter()` calculates the association strength iteratively for all cue-to-outcome combinations. Its output is a *weight matrix*, which contains the cue-to-outcome association strengths for all cue-to-outcome combinations as it is obtained at the end of the learning set. In the next step we will calculate the weight matrices for learning set with 30 learning events, in which every outcome has been presented 10 times.

```
#Take first 30 time steps (10 presentations * 3 outcomes)
M.weights_30 = RW.iter(DF.learn7[1:30,])
```

Remember that the Rescorla-Wagner Algorithm iterates through the entire learning set. Hence, by having `RW.iter()` learning the first 30, first 150 and first 500 learning events, we can inspect what will happen in the next time step. After 30 learning events, i.e. 10 presentations of each cue-to-outcome constellation, the association strength between every constellation shown in the following table, as is the next learning event:

```
M.weights_30

##               come           food           leave
## beep    0.0875596999 -0.0046945143  0.086869555
## bell    0.0001080719  0.0916100101 -0.004050310
## bye    -0.0039422384  0.0867231418  0.087706113
## hello   0.0916100101  0.0001923539 -0.004886868

DF.learn7[31,]

##           Cues Outcomes Frequency
## 31 bell_bye      food           1
```

In learning event 31, the cues **bell bye** are simultaneously present with the outcome food. Δw_i^t for {food} is calculated by $\alpha * \beta * (\lambda = 1 - (\text{weight for bell} = 0.09161 + \text{weight for bye} = 0.0867231)) = \alpha * \beta * 0.8216668 = 8.2166685 \times 10^{-5}$.

However, we are interested in the adaptation for **hello to food**. For this, we need to go to an event where **hello** is present as a cue. This happens to be event 33, in which {come} is also the outcome. For a precise illustration, we rerun `RW.iter()` with 32 learning events.

```
M.weights_32 = RW.iter(DF.learn7[1:32,])
DF.learn7[33,]

##           Cues Outcomes Frequency      Cues2
## 33 hello_beep      come           1 hello beep

M.weights_32

##               come           food           leave
## beep    0.0867231418 -0.0055969673  0.095132164
## bell    0.0001464135  0.0998266786 -0.004886868
## bye    -0.0047404547  0.0940373573  0.095132164
## hello   0.0916100101  0.0001923539 -0.004886868
```

Now remember that whenever cues are present, but outcomes are absent, Δw_i^t equals the difference between zero and the sum of the weights from that cue to those outcomes (Equation 12.1, c).

We therefore need to obtain the weights between the cues **hello beep** and the outcome {food} at the learning event 32. This is 1.9235394×10^{-4} for **hello** and -0.0046945 for **beep**. These two have to be summed and subtracted from zero. Note that the association weight for **beep** is negative, whereas the other is positive. The sum of these two equals -0.0045022. Since we have to subtract sum of weights from zero, we end up adding them to zero.

As a result, Δw_i^t has a positive value for the combination **hello to food**, eventhough they have never occurred together. The reason why **hello** has a positive weight to {food} is simple: because **beep** occurred in the learning order before **hello** occurred, it obtained negative weight to {food}, making the sum of weights negative. Moreover, **beep** occurs more frequently than **hello**. Likewise, **bell** has a positive association strength for {come}, eventhough these two never cooccurred together.

12.4 Derive discriminative measures from weights

12.5 General description

Table 12.1: The table illustrates a two-layer cue-to-outcome network with k cues $\mathcal{C}$ and n outcomes $\mathcal{O}$. We illustrate the calculation of d-measures for the word form $\{\text{dogs}\}$ as an outcome $\mathcal{O}_j$, located in the second column, and its associated cue set $\mathcal{CS}_\Omega = \text{ld dO Og gz zb}$, located in rows 3 to 7. Each $\mathcal{C}_i$ in the network is associated with each $\mathcal{O}_j$ by a weight $\mathcal{W}_{(i,j)}$, representing their connection strength, where $i = 1, 2, \dots, k$ and $j = 1, 2, \dots, n$. See Section ?? for explanation about the measures derived from this network.

$\mathcal{C}_{(1)}$	$\mathcal{O}_{(1)}$ $\mathcal{W}_{(1,1)}$	$\{\text{dogs}\}_{(2)}$ $\mathcal{W}_{(1,2)}$	...	$\mathcal{O}_{(n)}$ $\mathcal{W}_{(1,n)}$
$\mathcal{C}_{(2)}$	$\mathcal{W}_{(2,1)}$	$\mathcal{W}_{(2,2)}$	...	$\mathcal{W}_{(2,n)}$
ld	$\mathcal{W}_{(3,1)}$	$\mathcal{W}_{(3,2)}$	...	$\mathcal{W}_{(3,n)}$
dO	$\mathcal{W}_{(4,1)}$	$\mathcal{W}_{(4,2)}$	...	$\mathcal{W}_{(4,n)}$
Og	$\mathcal{W}_{(5,1)}$	$\mathcal{W}_{(5,2)}$	...	$\mathcal{W}_{(5,n)}$
gz	$\mathcal{W}_{(6,1)}$	$\mathcal{W}_{(6,2)}$	...	$\mathcal{W}_{(6,n)}$
zb	$\mathcal{W}_{(7,1)}$	$\mathcal{W}_{(7,2)}$	...	$\mathcal{W}_{(7,n)}$
...	...	...	...	...
$\mathcal{C}_{(k)}$	$\mathcal{W}_{(k,1)}$	$\mathcal{W}_{(k,2)}$	...	$\mathcal{W}_{(k,n)}$
$\mathcal{A}$	$\mathcal{A}_{(1)}$	$\mathcal{A}_{(2)}$	...	$\mathcal{A}_{(n)}$

The examples we have dealt so far represent very artificial set ups, eventough you can recreate in your laboratory by means of very controlled experiments: You *can* train your dog to react to three cues, you *can* train the participants of your experiment in a Wug-test to react to three cues. But life is much more complicated. If you want to have a fair representation of the real life, you need to increase the amount of cues per outcome to a hundred or even a thousand. Also, you need to increase the amount of outcomes appropriately, since, as you can imagine, there are a lot of outcomes in real life and they share a lot of cues (remember the example with the color red and cue competition).

The weight matrix represents the association strength between every cue and every outcome in your learning set. It is very powerful as you can investigate single association strengths. But how about when you are presented with multiple cues: ten cues, a hundred, one thousand. There you cannot check and compare every single cue (you can, but it can become very, very tedious and I wonder what you could possibly learn from it).

After training, we can ask three different questions to the network: First, given the presentation of a cue set $\mathcal{CS}_\Omega$ to the trained network, how strongly does it activate the associated outcome $\mathcal{O}_j$?

Second, the presentation of a cue set $\mathcal{CS}_\Omega$ to the trained network, to how many outcomes $\mathcal{O}$ is it related? Or, in other words, how strongly entrenched is the cue set $\mathcal{CS}_\Omega$ in the network?

Third, to how many cues is an outcome $\mathcal{O}_j$ related? Or, in other words, how strongly is an outcome $\mathcal{O}_j$ entrenched in the network?

We defined three types of discriminative measures to answer these questions, as explained in the following.

- **Activations** represent the total association strength between a cue set $\mathcal{CS}_\Omega$ associated with a specific outcome (e.g. *ld dO Og gz zb* to $\{\text{dogs}\}$) and all outcomes $\mathcal{CS}_{1\dots n}$. It is calculated by

$$\mathcal{A}_j = \sum_1^m \mathcal{W}_{mj} \quad (12.2)$$

where $\mathcal{A}_j$ is the activation vector $\mathcal{A}$ (cf. bottom dark gray row in Table 12.1); $\mathcal{W}_{mj}$ represents the weights between a cue set and its associated outcome; $j = 1, 2, \dots, n$ ranges across all outcomes in the network and $m = 1, 2, \dots, M$ ranges across all cues within a cue set $\mathcal{CS}_\Omega$ with the size M . The specific association strength between an outcome $\mathcal{O}_j$ and its associated cue set $\mathcal{CS}_j$ is called *activation*.

For example, in Table 12.1, $\mathcal{CS}_\Omega$ for $\mathcal{O}_j = \{\text{dogs}\}$ is **ld dO Og gz zb**. The respective weights, marked in red, in the grey highlighted column, are $\mathcal{W}_{3,2}$, $\mathcal{W}_{4,2}$, $\mathcal{W}_{5,2}$, $\mathcal{W}_{6,2}$, $\mathcal{W}_{7,2}$, which are summed to $\mathcal{A}_2$.

- **Activation Diversity** $\mathcal{D}_\Omega$ is a measure for a cue set $\mathcal{CS}_\Omega$ which is associated to $\mathcal{O}_j$. It is defined by the sum of the absolute activations $\mathcal{A}_{j=1\dots n}$ across the activation vector $\mathcal{A}$ (cf. Table 12.1, bottom dark gray row):

$$\mathcal{D}_\Omega = \sum_1^j |\mathcal{A}_j| \quad (12.3)$$

Independently of whether activations in $\mathcal{A}$ are positive or negative, $\mathcal{D}_\Omega$ represents a measure of variability of $\mathcal{A}_j$ in $\mathcal{A}$ around zero, given that $\mathcal{A}_j \in \{-\lambda : \lambda\}$ and $\bar{\mathcal{A}} \sim 0$. We interpret that value as a measure of entrenchment or lexicality, with higher $\mathcal{D}_\Omega$ associated with higher entrenchment. The reasoning is that a cue set which is highly lexical creates lots of activity in the network; one that is not lexical creates little activity in the network. In other words, the higher $\mathcal{D}_\Omega$, the more activity $\mathcal{CS}_\Omega$ creates across the network, the deeper are cues entrenched in the network, the more lexical they are.

- **Prior** $\mathcal{P}_j$ is a measure for an outcome $\mathcal{O}_j$. It is defined as the sum of the absolute weights (1-norm) from all cues $\mathcal{C}_{i=1\dots k}$ to one outcome $\mathcal{O}_j$. This can be regarded as the weight vector for the outcome $\mathcal{O}_j$, as illustrated by the gray highlighted column for $\{\text{dogs}\}$ in Table 12.1).

$$\mathcal{P}_j = \sum_1^i |\mathcal{W}_{(i,j)}| \quad (12.4)$$

Independently of whether weights in the weight vector for the outcome $\mathcal{O}_j$ are positive or negative, $\mathcal{P}_j$ represents a measure of variability of $\mathcal{W}_{(i,j)}$ around zero, given that $\mathcal{W}_{(i,j)} \in \{-\lambda : \lambda\}$ and $\bar{\mathcal{W}} \sim 0$. Similar to cue lexicality, we interpret that value as a measure of entrenchment, with higher $\mathcal{P}_j$ associated with higher entrenchment. In other words, the higher $\mathcal{P}_j$, the higher the prior activation of $\mathcal{O}_j$ in the network.

12.6 Calculate discriminative measures

12.6.1 Code to calculate activation

The following code calculates the for `{come}`, `{food}` and `{leave}` using the cues set up in `DF.learn7`. However, we do not want to do this for all events but only for the unique cue-to-outcome combinations. This is why we reduce the data frame to the respective rows using `unique()`. We call this the **test data frame**. To store activations, we **initialize** a new column, **Activation**. We fill the column with to be able to include numbers. Furthermore, given that Rescorla-Wagner can be used as a **Classifier**, we initialize the column **found** to store the classified cues. We fill the column with empty to be able to include strings.

```
DF.test1 = unique(DF.learn7)
DF.test1$Activations = NA
DF.test1$Found = ''
```

Next, we iterate across all events in *DF.test1* in a `for()` loop. In each iteration, we store the event in `tmp`. In the following we do this for the first iteration so we can inspect each step. To calculate the activation, we access the columns with the and the them using the `strsplit()` at the positions with an underscore. Note that `strsplit()` provides you with a `,` which is why we need to use `unlist()` to create a `.` We use this vector `CUES` to access the rows of matrix *M.weights_30*, which we store in `WEIGHTS`. According to formula 12.3 we need to sum the weights. Given that `WEIGHTS` is a matrix, we need to use the function `colSums()` which sums across the rows in each column. We store the result in the variable `ACT`, which is a vector of activations. We immediately sort `ACT` in a decreasing order using `sort()` . We use the content in the column `Outcomes` from `tmp` to acces `ACT` and store it in the respective row in *DF.test1*. Given that `ACT` is already sorted in decreasing order, the with the highest activation is regarded to be the classified item by the network. Using `names()` , we can get the entry names of the vector. We acces the first entry and store it in `found`. As you can see, the network.

```
irow = 1
#extract event
tmp = DF.test1[irow,]
#extract cues and split
CUES = unlist(strsplit(tmp$Cues, split = '_'))
CUES

## [1] "bell" "bye"

#extract weights
WEIGHTS = M.weights_30[CUES,]
WEIGHTS

##           come      food      leave
## bell  0.0001080719 0.09161001 -0.00405031
## bye   -0.0039422384 0.08672314  0.08770611

#calculate activations
ACT = colSums(WEIGHTS)
ACT

##           come      food      leave
## -0.003834167  0.178333152  0.083655803

#sort act
ACT = sort(ACT, decreasing = TRUE)
ACT

##           food      leave      come
##  0.178333152  0.083655803 -0.003834167
```

```
#access the outcome in the event
DF.test1$Activations[irow] = ACT[tmp$Outcomes]
DF.test1$Found[irow] = names(ACT)[1]
# result
DF.test1
```

##	Cues	Outcomes	Frequency	Cues2	Activations	Found
## 1	bell_bye	food	1	bell bye	0.1783332	food
## 2	bye_beep	leave	1	bye beep	NA	
## 3	hello_beep	come	1	hello beep	NA	

The following piece of code includes all the above into a `for()` loop. Usurprisingly for this simple example, the network classified all outcomes correctly given their cues.

```
for(irow in 1:nrow(DF.test1)){
#extract event
tmp = DF.test1[irow,]
#extract cues and split
CUES = unlist(strsplit(tmp$Cues, split = '_'))
#extract weights
WEIGHTS = M.weights_30[CUES,]
WEIGHTS
#calculate activations
ACT = colSums(WEIGHTS)
#sort act
ACT = sort(ACT, decreasing = TRUE)
#access the outcome in the event
DF.test1$Activations[irow] = ACT[tmp$Outcomes]
DF.test1$Found[irow] = names(ACT)[1]
}

# result
DF.test1
```

##	Cues	Outcomes	Frequency	Cues2	Activations	Found
## 1	bell_bye	food	1	bell bye	0.1783332	food
## 2	bye_beep	leave	1	bye beep	0.1745757	leave
## 3	hello_beep	come	1	hello beep	0.1791697	come

Testing can also be done on **unseen data**, i.e. cue-to-outcome combinations which the network has not yet encountered. This is done in the following. where we set up new cue-to-outcome combinations.

```
#new test data
Cues = c('bell_beep', 'hello_bye', 'bye_beep')
Outcomes = c('food', 'leave', 'come')
DF.test2 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
```

```

Activations = NA,
Found = '',
stringsAsFactors = FALSE)

#run
for(irow in 1:nrow(DF.test2)){
  #extract event
  tmp = DF.test1[irow,]
  #extract cues and split
  CUES = unlist(strsplit(tmp$Cues, split = '_'))
  #extract weights
  WEIGHTS = M.weights_30[CUES,]
  WEIGHTS
  #calculate activations
  ACT = colSums(WEIGHTS)
  #sort act
  ACT = sort(ACT, decreasing = TRUE)
  #access the outcome in the event
  DF.test2$Activations[irow] = ACT[tmp$Outcomes]
  DF.test2$Found[irow] = names(ACT)[1]
}

# result
DF.test2

##           Cues Outcomes Activations Found
## 1 bell_beep      food   0.1783332  food
## 2 hello_bye     leave   0.1745757  leave
## 3  bye_beep      come   0.1791697   come

```

12.6.2 Code to calculate activation diversity and prior

As stated above, is ‘simply’ the sum of absolute activations. A is the sum of absolute weights associated with an outcome. The code below uses the code from above and incorporates the calculation.

```

#new test data
Cues = c('bell_beep', 'hello_bye', 'bye_beep')
Outcomes = c('food', 'leave', 'come')
DF.test2 = data.frame(
  Cues = Cues,
  Outcomes = Outcomes,
  Activations = NA,
  Found = '',
  stringsAsFactors = FALSE)

```

```

#run
for(irow in 1:nrow(DF.test2)){
  #extract event
  tmp = DF.test1[irow,]
  #extract cues and split
  CUES = unlist(strsplit(tmp$Cues, split = '_'))
  #extract weights
  WEIGHTS = M.weights_30[CUES,]
  WEIGHTS
  #calculate activations
  ACT = colSums(WEIGHTS)
  #sort act
  ACT = sort(ACT, decreasing = TRUE)
  #access the outcome in the event
  DF.test2$Activations[irow] = ACT[tmp$Outcomes]
  DF.test2$Found[irow] = names(ACT)[1]
  DF.test2$ActivationDiversity = sum(abs(ACT)) # activation diversity
  DF.test2$Prior = sum(abs(M.weights_30[,tmp$Outcomes]))
}

# result
DF.test2

##           Cues Outcomes Activations Found ActivationDiversity Prior
## 1 bell_beep    food    0.1783332  food          0.2656546 0.18322
## 2 hello_bye   leave    0.1745757  leave          0.2656546 0.18322
## 3 bye_beep    come     0.1791697   come          0.2656546 0.18322

```

Bibliography

- Arnold, D. & F. Tomaschek. 2016. The Karl Eberhards Corpus of spontaneously spoken Southern German in dialogues - audio and articulatory recordings. In Christoph Draxler & Felicitas Kleber (eds.), *Tagungsband der 12. Tagung Phonetik und Phonologie im deutschsprachigen Raum. P und P12. 12. Tagung Phonetik und Phonologie im deutschsprachigen Raum, 12. - 14. Oktober 2016*, 9–11. Ludwig-Maximilians-Universität München. <https://ids-pub.bsz-bw.de/frontdoor/index/index/year/2017/docId/5944>.
- Arppe, Antti, Peter Hendrix, Petar Milin, R. Harald Baayen, Tino Sering & Cyrus Shaoul. 2018. ndl: Naive Discriminative Learning. <https://CRAN.R-project.org/package=ndl>.
- Baayen, R. H. 2008. *Analysing linguistic data: A practical introduction to statistics*. Cambridge, MA: Cambridge University Press.
- Hoppe, Dorothée B., Petra Hendriks, Michael Ramscar & Jacolien van Rij. 2020. An Exploration of Error-Driven Learning in Simple Two-Layer Networks From a Discriminative Learning Perspective. Tech. rep. PsyArXiv. doi:10.31234/osf.io/py5kd. <https://psyarxiv.com/py5kd/>.
- Kamin, Leon J. 1968. Attention-like processes in classical conditioning. In M. R. Jones (ed.), *Miami symposium on the prediction of behavior*, 9–31. Miami: Miami University Press.
- Nixon, Jessie S. 2020. Of mice and men: Speech sound acquisition as discriminative learning from prediction error, not just statistical tracking. *Cognition* 197. 104081. doi:<https://doi.org/10.1016/j.cognition.2019.104081>. <http://www.sciencedirect.com/science/article/pii/S0010027719302549>.
- Pavlov, I. P. 2006. *Conditioned reflexes*. London: Oxford.
- Ramscar, M., Dye, M. & Klein, J. 2013. Children Value Informativity Over Logic in Word Learning -. *Psychological Science* 24(6). 1017–1023. doi:<https://doi.org/10.1177/0956797612460691>. https://journals.sagepub.com/doi/full/10.1177/0956797612460691?casa_token=CzQq7-JFh9gAAAAA%3A0ITl9tpijaWk36ivjJN7MmI-J2LPtHEGamHg4mUR3UbJkGdq-ZIN5-6sbhfdCpc3Wc-W2F8LI6M-qA.
- Ramscar, M., D. Yarlett, M. Dye, K. Denny & K. Thorpe. 2010. The Effects of Feature-Label-Order and their implications for symbolic learning. *Cognitive Science* 34(6). 909–957.
- Rescorla, R. 1988. Pavlovian Conditioning - It's Not What You Think It Is. *American Psychologist* 43(3). 151–160.
- Rescorla, R. & A. Wagner. 1972. A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W.F. Prokasy (eds.), *Classical conditioning II: Current research and theory*, 64–69. Appleton Century Crofts, New York.

Index

, 40, 70
(, 108, 113
-, 41
:, 47
==, 35–37, 40, 41, 66, 71, 113
Activation Diversity, 155
Activation, 152
NA, 152
Prior, 155
\$, 26, 41
%%, 90, 93
%in%, 37, 40, 92, 114
as.character(), 20
as.factor(), 115, 116
as.numeric(), 20, 42, 81, 116
attach(), 108
by, 45
c(), 16, 80
cat(), 48, 67, 93
class(), 19, 25, 76
colSums(), 153
colnames(), 27, 51
complete.cases(), 69, 113
cor(), 117
cues, 153
data.frame(), 26, 36, 54
density(), 108
dev.off(), 114
dim(), 46, 50
else, 91
exp(), 109
for(), 47, 48, 54, 61, 63, 64, 66, 67, 79–81, 89,
90, 100, 101, 103, 114, 137, 153, 154
getwd(), 19
grep(), 39–41
grepl(), 38–40, 48
gsub(), 64, 66, 74
head(), 22, 32, 42, 43, 59, 73, 79, 85, 87
hist(), 108
if(), 90–92
ifelse(), 74, 116
install.packages(), 83
in, 48
is.na(), 69
length(), 49, 50, 77
library(), 83
list(), 77
list.files(), 21, 54
list, 153
lmer(), 121
load(), 17
log(), 109
mean(), 107
merge(), 103, 104
names(), 77, 89, 153
nchar(), 60
ncol(), 46, 50
nrow(), 46, 50, 54
object.size(), 16
order(), 104
outcome, 153
par(), 108, 114
paste(), 80, 101
pdf(), 114
plot(), 108
print(), 16, 54, 123
rbind(), 49, 51–54, 101
read.csv(), 58–60
read.table(), 24, 25, 33, 58
rep(), 29, 36
return(), 112
rnorm(), 107, 108
round(), 98, 117
rownames(), 27
sample(), 36
save(), 17
seq(), 26, 28, 29, 45, 47

setwd(), 17, 19, 24
shapiro.test(), 108
sort(), 89, 153
str(), 25–27, 59, 77
str_sub(), 83, 84
strings, 152
strsplit(), 153
strsplit(), 75, 79
sum(), 71, 92
summary(), 118, 122
summary, 70
sum, 69
table(), 88, 89, 97, 102, 116
tail(), 23, 45
temporary variable, 153
try(), 51
unique(), 64, 71, 152
unlist(), 153
upper.tri(), 117
vector, 153
visreg(), 124
which(), 40–42

Accessing columns in dataframes, 26, 41
 Accessing lists in a *for()* loop, 79
 Accessing vectors, 23
 Accessing vectors in a *for()* loop, 79
 Activation Diversity $\mathcal{D}_\Omega$, 152
 Activations, 151
 Annotations, 12, 57
as.is, 60
 Attaching new columns to dataframes, 26

Bigram, 99
 Boolean, 40
 Boolean vector, 71, 74, 92
 Brackets, round, 15
 Brackets, square, 23

Characters, 20
 Class, 52
 Classes, 19
 Classifier, 152
 Clipping, 82
 Collapse, 80
 Collinearity, 117
 Column, 41
 Columns, 25, 26, 44
 Conditional probability, 104
 Correlation, 117
 Counts, 88, 93

Counts, bigrams, 99
 cues, 129
 Data transformation, 108
 Dataframe, 12, 25, 32, 35, 36, 41, 43, 50, 57
 Dataframe, accessing, 27
 Dataframes, append, 49
 Dataframes, bind, 49
 Dataframes, deleting columns, 26
 Debugging, 70
 Density, 108
 Directory, 19
 Distribution, 108
 Error message, 24
 Factors, 25
 FALSE, 25, 35, 36, 41, 69, 74, 90
 File name, 21
 Frequencies of occurrences, 88
 Function, 112
 Functions, 16
 Garbage, 33
 Header, 58
 header, 58
 histogram, 108
 Index, 23, 36, 42
 initialize, 152
 Intercept, 123
 letters, 26
 Library, 82
 List, 76, 77, 79
 Logical vector, 35, 36, 41
 Loop, 47
 Manual, 24
 Mean, 107
 Mirror, 83
 Modulo, 93
 NA, 24, 58, 67, 69, 70
 NaN, 24, 69
 Nested commands, 43, 54
 Normal distribution, 107
 NULL, 26
 Number of columns in plot, 108
 Number of rows in plot, 108
 Optimization, 85
 outcomes, 129
 Outliers, 109

- Package, 82, 83
- Package, installation, 82
- Package, loading, 82
- Parameter, 24, 28, 32, 40
- Plot, 107
- Plotting, 107
- Praat, 12, 33, 34, 57
- Predictors, 117
- Prior $\mathcal{P}_j$, 152
- Probabilities, 97
- Probabilities, conditional, 98

- quote, 60

- R as calculator, 13
- R environment, 13
- R installation, 7
- RAM, 15
- Randomization, 36
- rank, 115
- Read tables, 23
- Regular expressions, 39, 64
- Repetition, 28, 36, 47
- Repository, 82
- Rows, 25
- Rows, names, 59
- RStudio, 13
- Rstudio, 8

- Scalars, 112
- Scientific notation, 109
- Searching, 35
- Searching, inverse, 40
- Searching, negate, 40
- Sequence, 28
- Slope, 123
- Split, 75, 79
- Standard deviation, 107
- String, 19, 25
- Strings, 19, 42
- Strings, split, 75, 79
- StringsAsFactors, 25
- stringsAsFactors, 44

- Table, 24
- Task preparation, 33
- test data frame, 152
- text file, 24
- Token, 88
- Transformation, 108
- Trigram, 99
- Trimming, 82
- TRUE, 35, 36, 40, 41, 69, 74, 90

- unseen data, 154

- Variable structure, 25
- Variables, 15, 44
- Vector, 16, 23, 41, 47, 79, 85
- Visualization, 107

- Word frequencies, 88
- Working directory, 19, 24