Why the tidyverse?

Amelia McNamara @AmeliaMN
www.amelia.mn

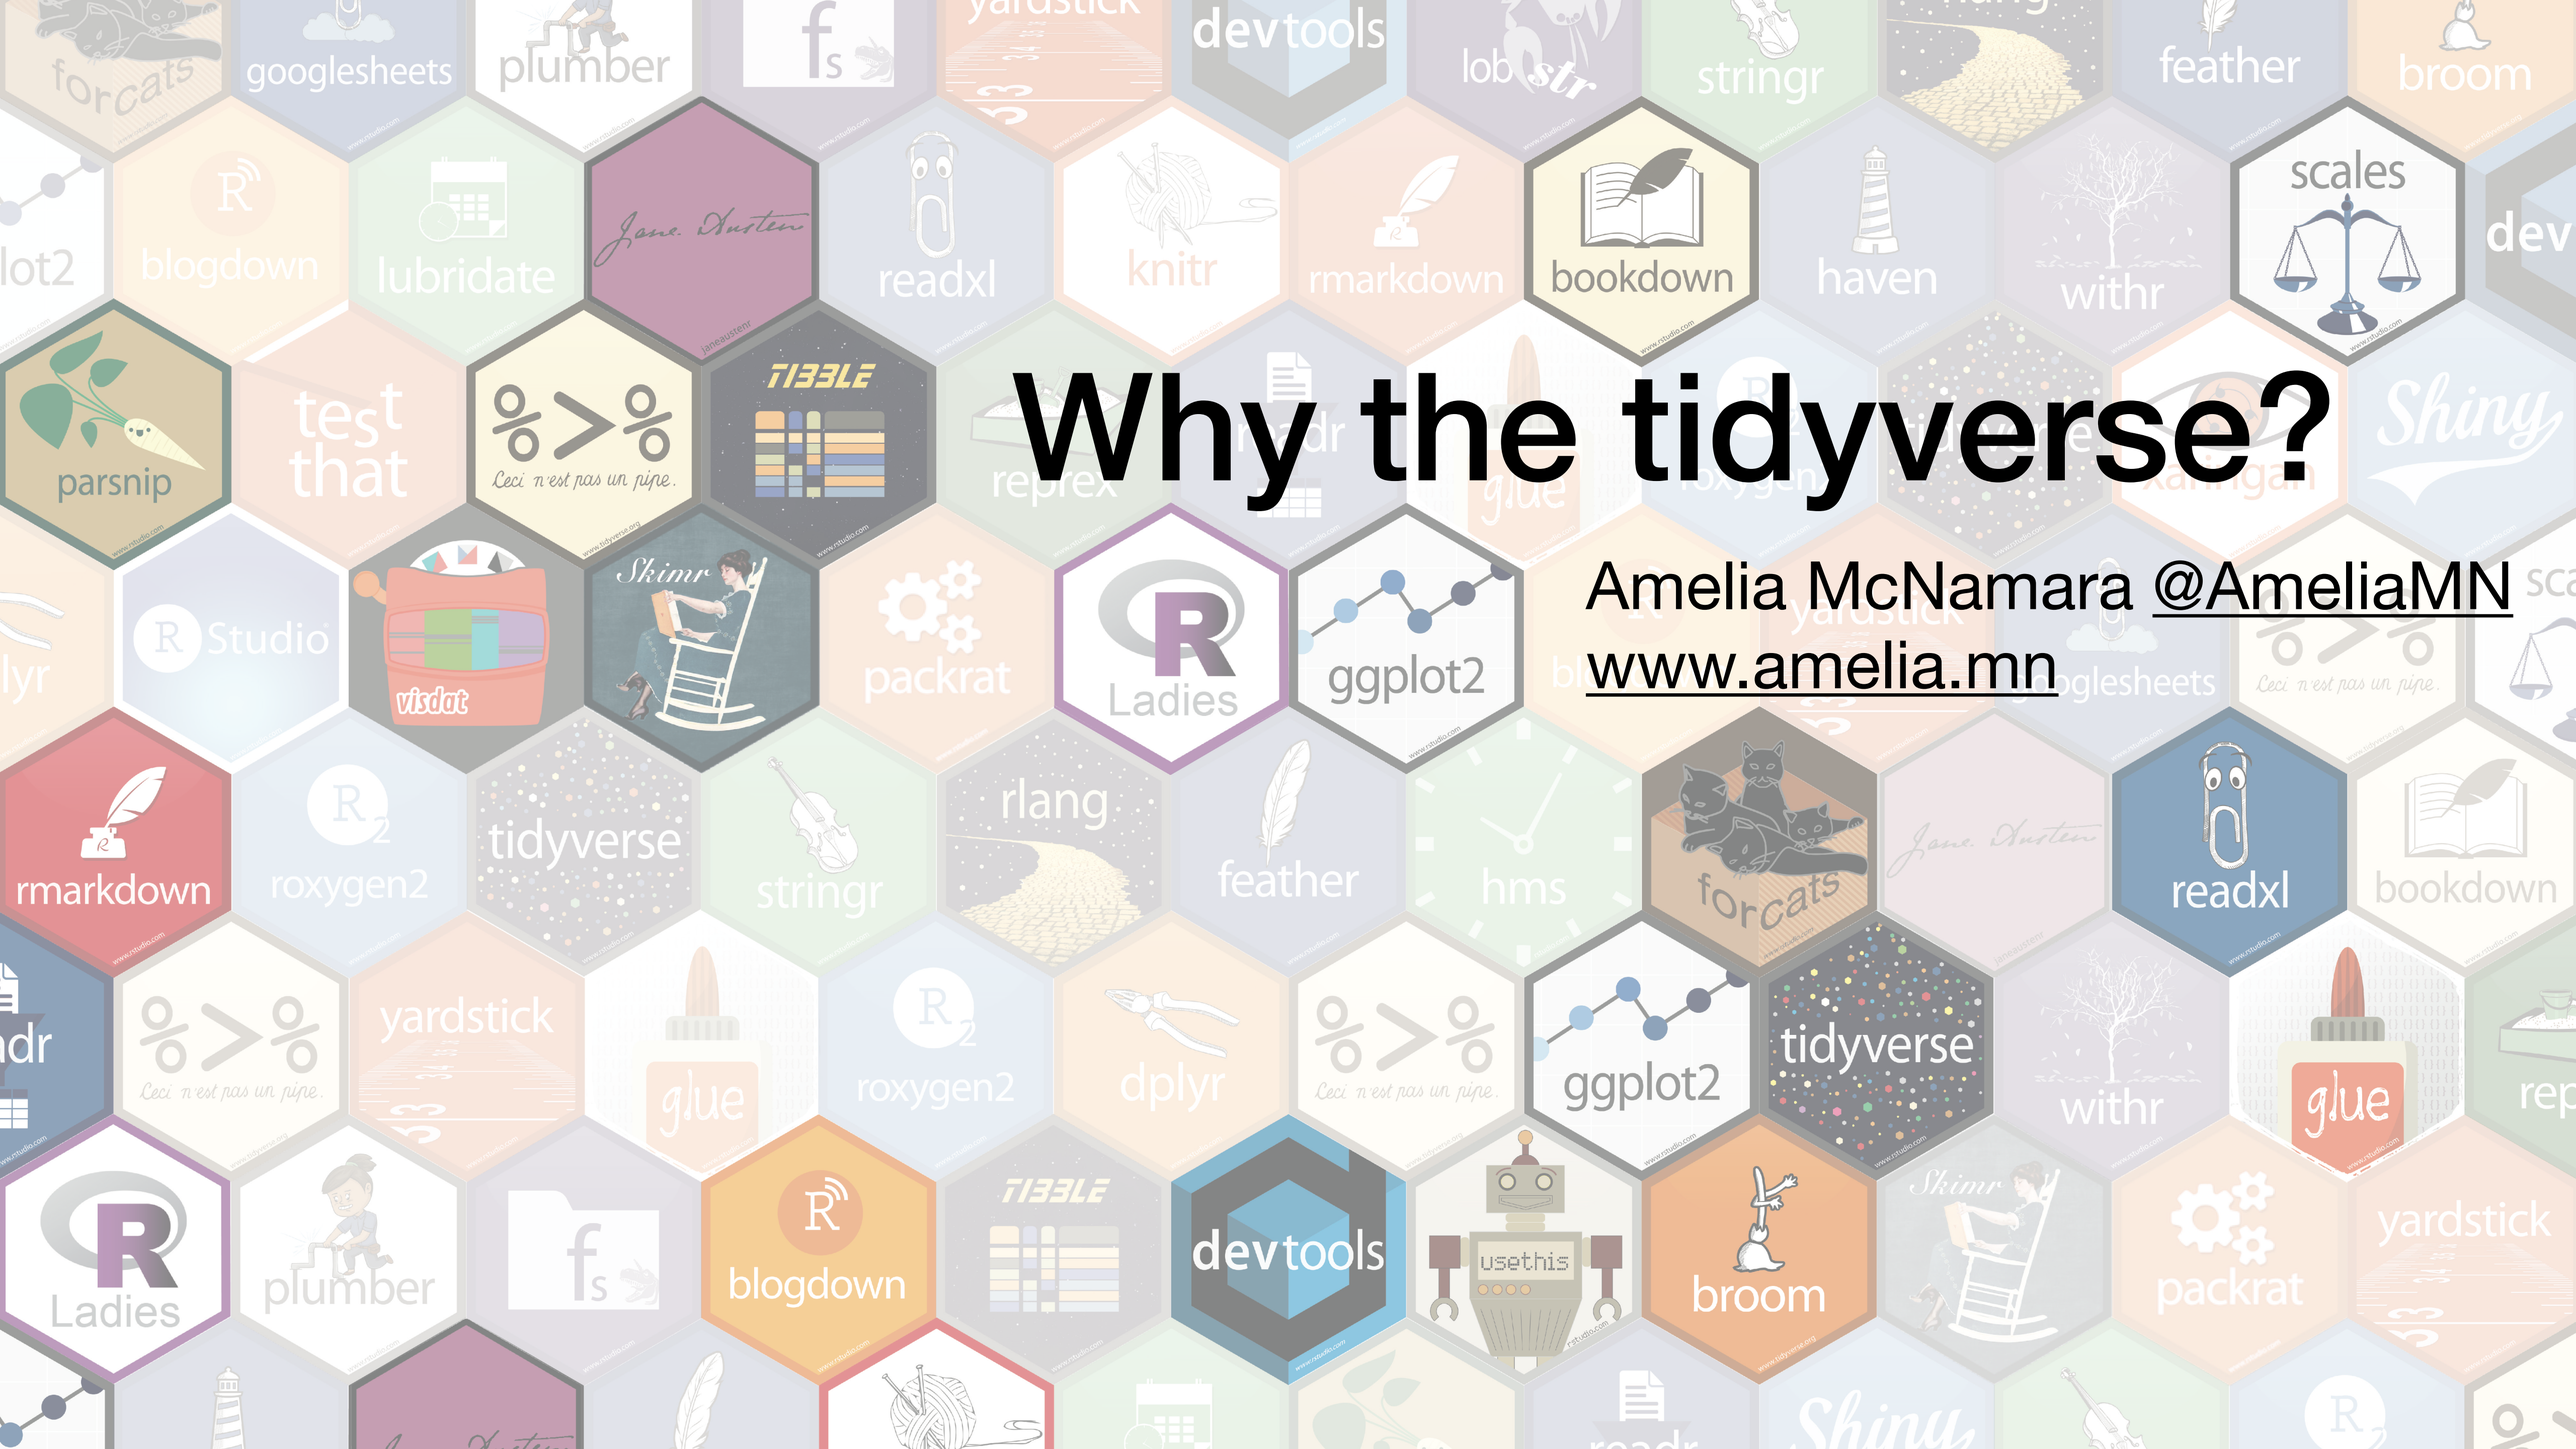

What is the tidyverse?

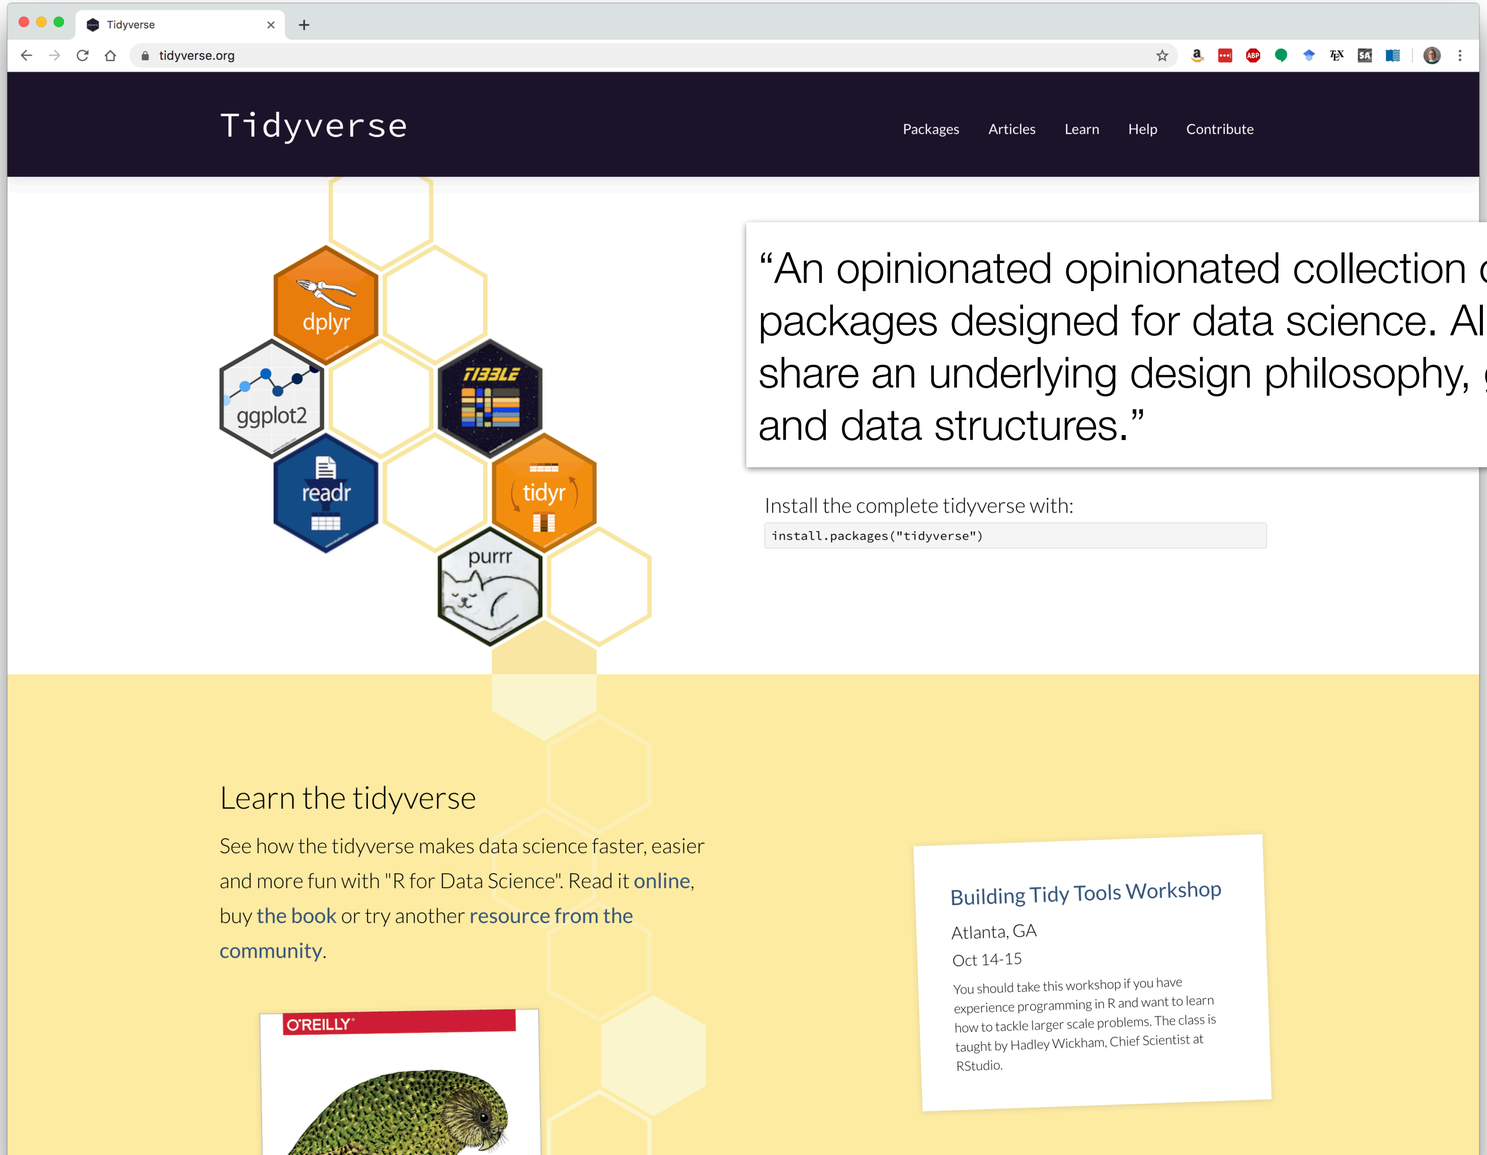

Core tidyverse

ggplot2

ggplot2 is a system for declaratively creating graphics, based on The Grammar of Graphics. You provide the data, tell ggplot2 how to map variables to aesthetics, what graphical primitives to use, and it takes care of the details. Learn more ...

dplyr

dplyr provides a grammar of data manipulation, providing a consistent set of verbs that solve the most common data manipulation challenges. Learn more ...

tidyr

tidyr provides a set of functions that help you get to tidy data. Tidy data is data with a consistent form: in brief, every variable goes in a column, and every column is a variable. Learn more ...

readr

readr provides a fast and friendly way to read rectangular data (like csv, tsv, and fwf). It is designed to flexibly parse many types of data found in the wild, while still cleanly failing when data unexpectedly changes. Learn more ...

purrr

purrr enhances R's functional programming (FP) toolkit by providing a complete and consistent set of tools for working with functions and vectors. Once you master the basic concepts, purrr allows you to replace many for loops with code that is easier to write and more expressive. Learn more ...

tibble

tibble is a modern re-imagining of the data frame, keeping what time has proven to be effective, and throwing out what it has not. Tibbles are data.frames that are lazy and surly: they do less and complain more forcing you to confront problems earlier, typically leading to cleaner, more expressive code. Learn more ...

stringr

stringr provides a cohesive set of functions designed to make working with strings as easy as possible. It is built on top of stringi, which uses the ICU C library to provide fast, correct implementations of common string manipulations. Learn more ...

forcats

forcats provides a suite of useful tools that solve common problems with factors. R uses factors to handle categorical variables, variables that have a fixed and known set of possible values. Learn more ...

More

Packages

As well as the core tidyverse, installing this package also installs a selection of other packages that you're likely to use frequently, but probably not in every analysis. This includes packages for:

- Working with specific types of vectors:
    - hms, for times.
    - lubridate, for date/times.
- Importing other types of data:
    - feather, for sharing with Python and other languages.
    - haven, for SPSS, SAS and Stata files.
    - httr, for web apis.
    - jsonlite for JSON.
    - readxl, for `.xls` and `.xlsx` files.
    - rvest, for web scraping.
    - xml2, for XML.
- Modelling
    - modelr, for modelling within a pipeline
    - broom, for turning models into tidy data

Even more

tidytext (Julia Silge and David Robinson)

skimr (Elin Waring, Michael Quinn, Amelia McNamara, Eduardo Ariño de la Rubia, Hao Zhu, Shannon Ellis)

(fun rOpensci connection)

Tidyverse philosophy/manifesto

"There are four basic principles to a tidy API:

- Reuse existing data structures.
- Compose simple functions with the pipe.
- Embrace functional programming.
- Design for humans."

Pure, predictable, pipeable

https://tidyverse.tidyverse.org/articles/manifesto.html

History of Programming Languages

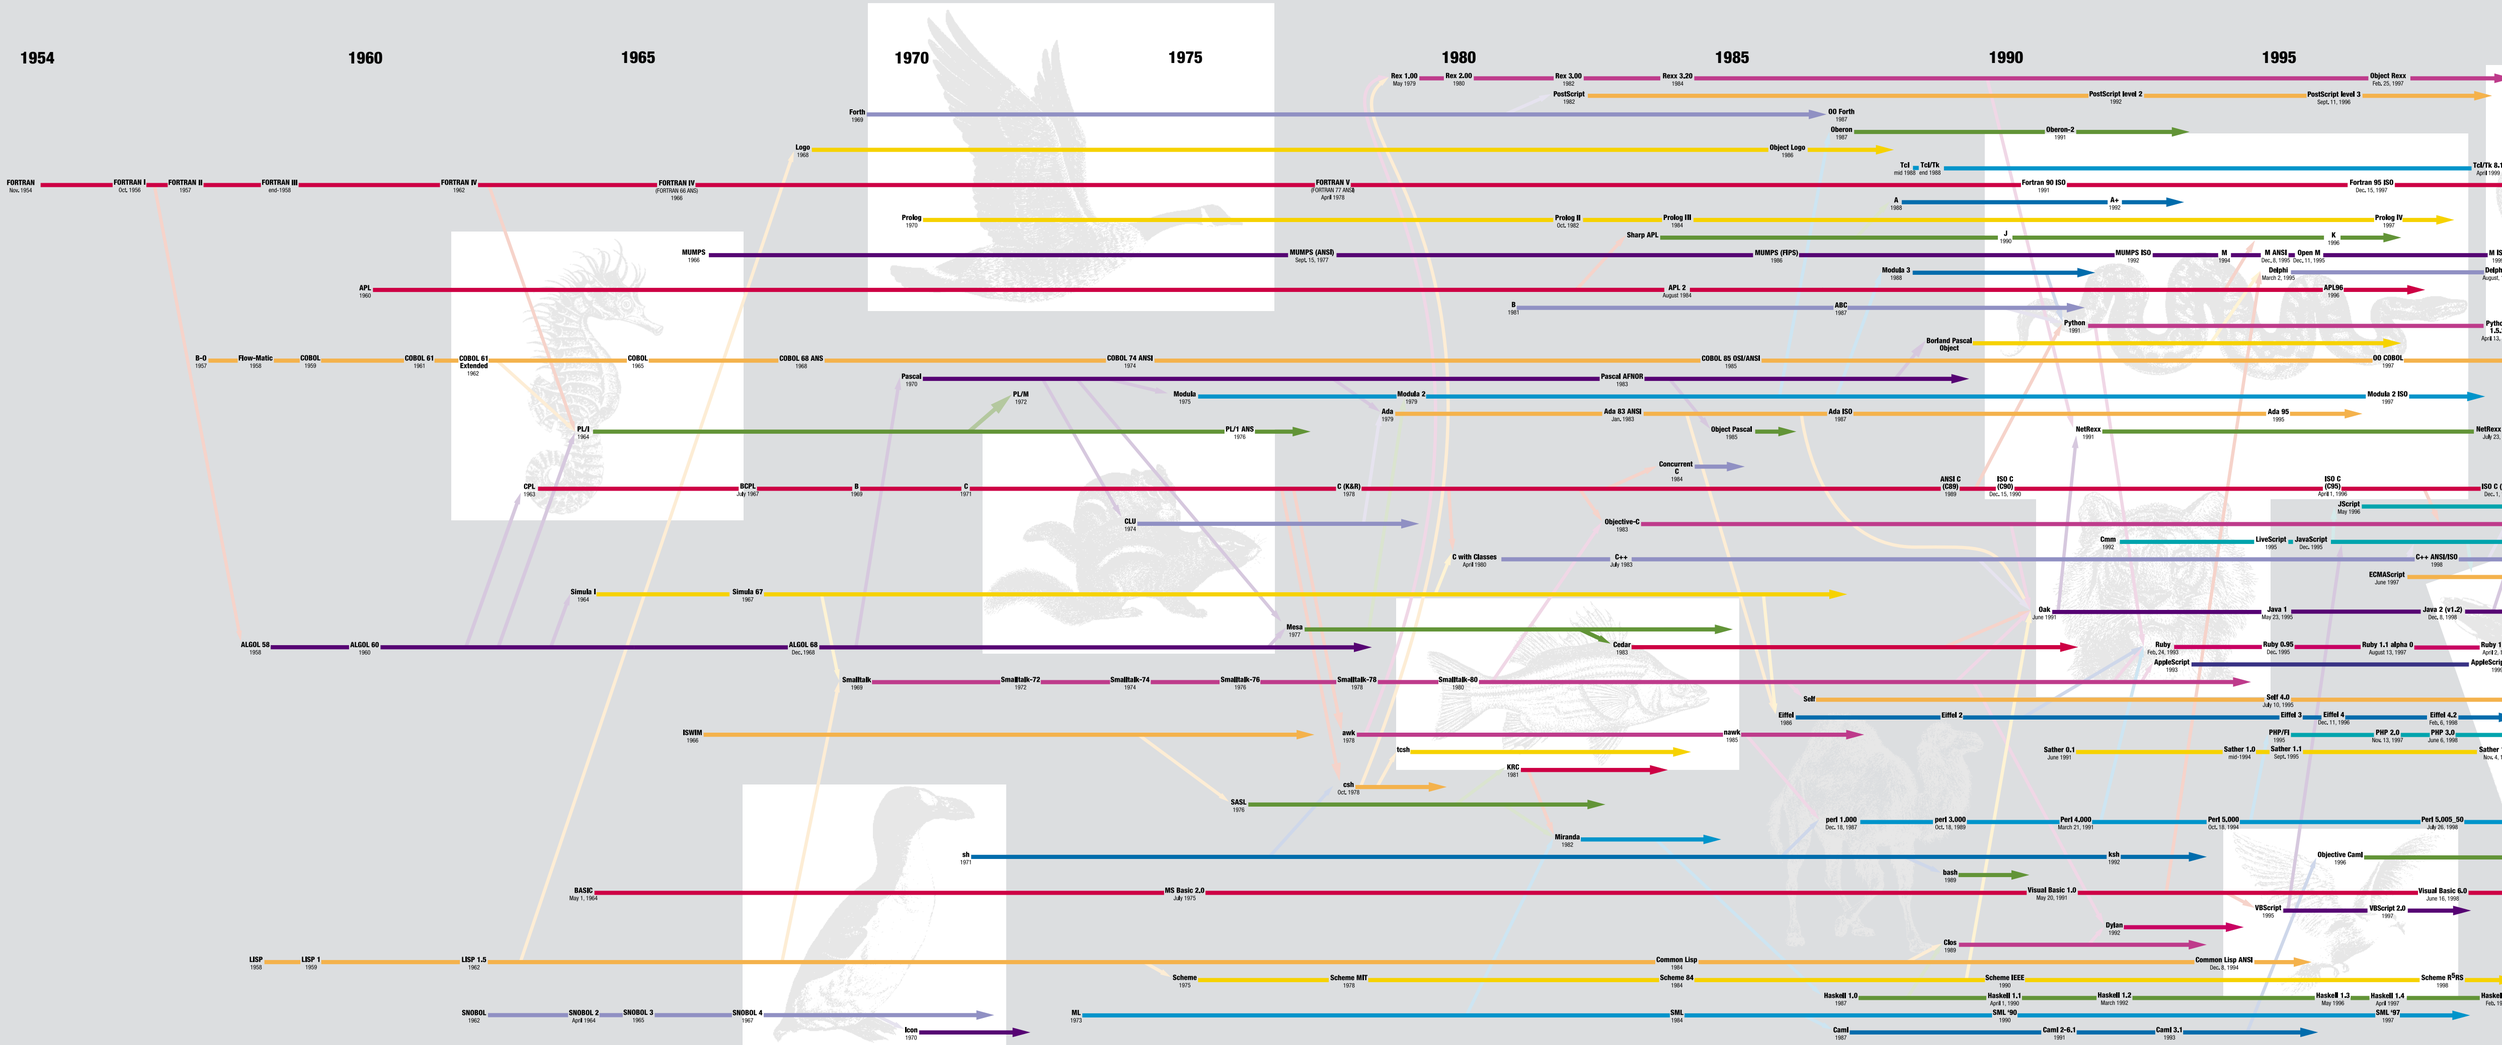

www.oreilly.com

For more than half of the fifty years computer programmers have been writing code, O'Reilly has provided developers with comprehensive, in-depth technical information. We've kept pace with rapidly changing technologies as new languages have emerged, developed, and matured. Whether you want to learn something new or need answers to tough technical questions, you'll find what you need in O'Reilly books and on the O'Reilly Network.

This timeline includes fifty of the more than 2500 documented programming languages. It is based on an original diagram created by Éric Lévénez (*www.levenez.com*), augmented with suggestions from O'Reilly authors, friends, and conference attendees.

For information and discussion on this poster, go to *www.oreilly.com/go/languageposter*.

Languages

O'REILLY®

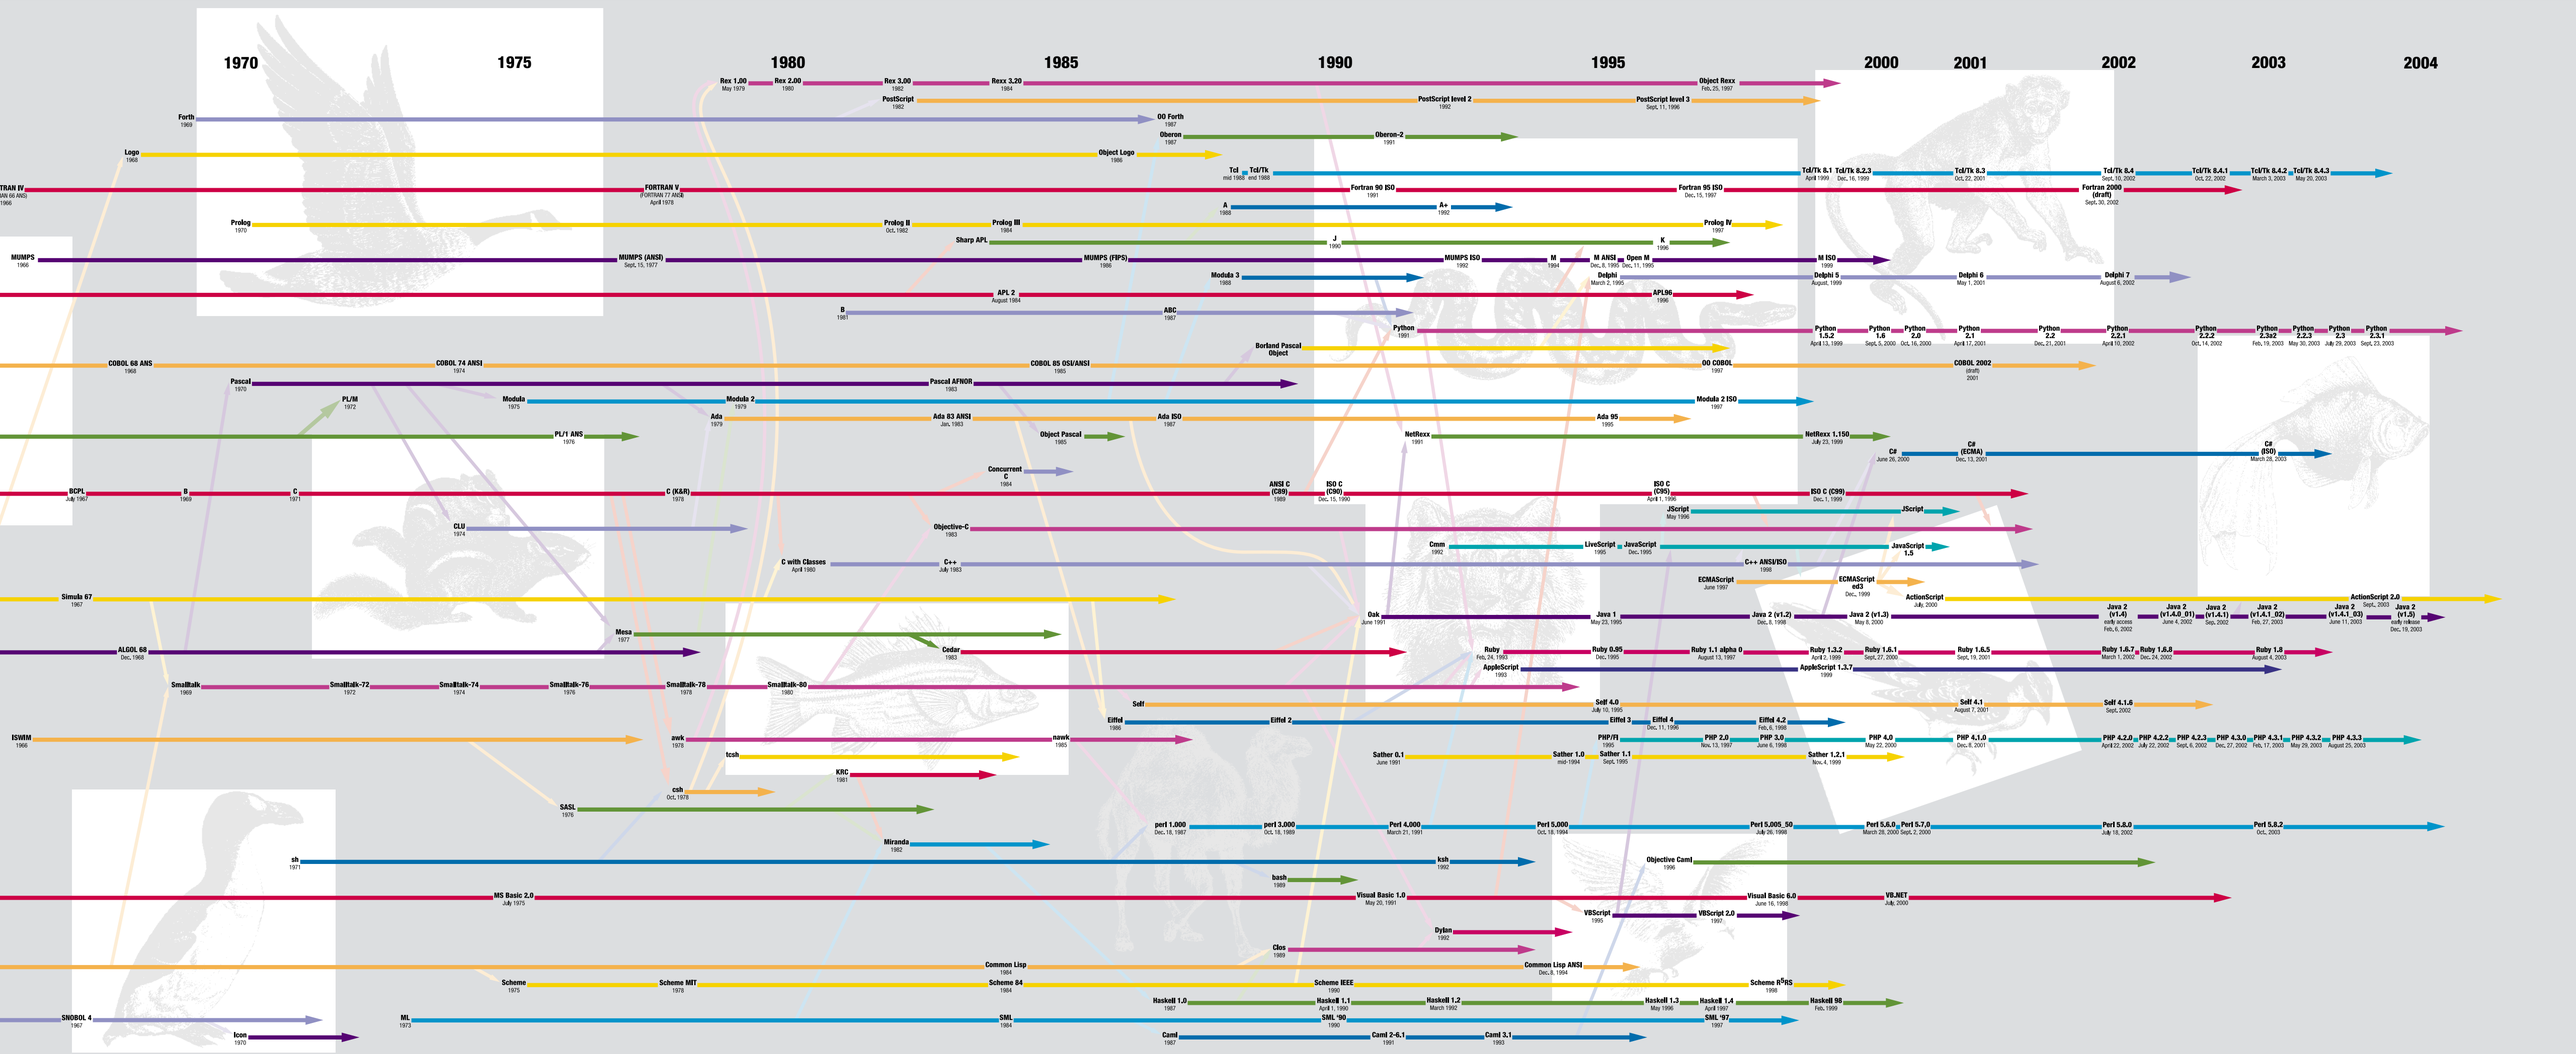

This timeline includes fifty of the more than 2500 documented programming languages. It is based on an original diagram created by Éric Lévénez (*www.levenez.com*), augmented with suggestions from O'Reilly authors, friends, and conference attendees.

For information and discussion on this poster, go to *www.oreilly.com/go/languageposter*.

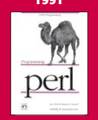

FORTRAN

Developed in the 1950s, still exists today

```
! sum.f90
! Performs summations using in a loop using EXIT statement
! Saves input information and the summation in a data file

program summation
implicit none
integer :: sum, a

print*, "This program performs summations. Enter 0 to stop."
open(unit=10, file="SumData.DAT")

sum = 0

do
 print*, "Add:"
 read*, a
 if (a == 0) then
  exit
 else
  sum = sum + a
 end if
 write(10,*) a
end do
```

Wikipedia

S

Written by John Chambers

Originally developed in 1975 as an alternative to FORTRAN for statistical computing

“Books”

@rabaath

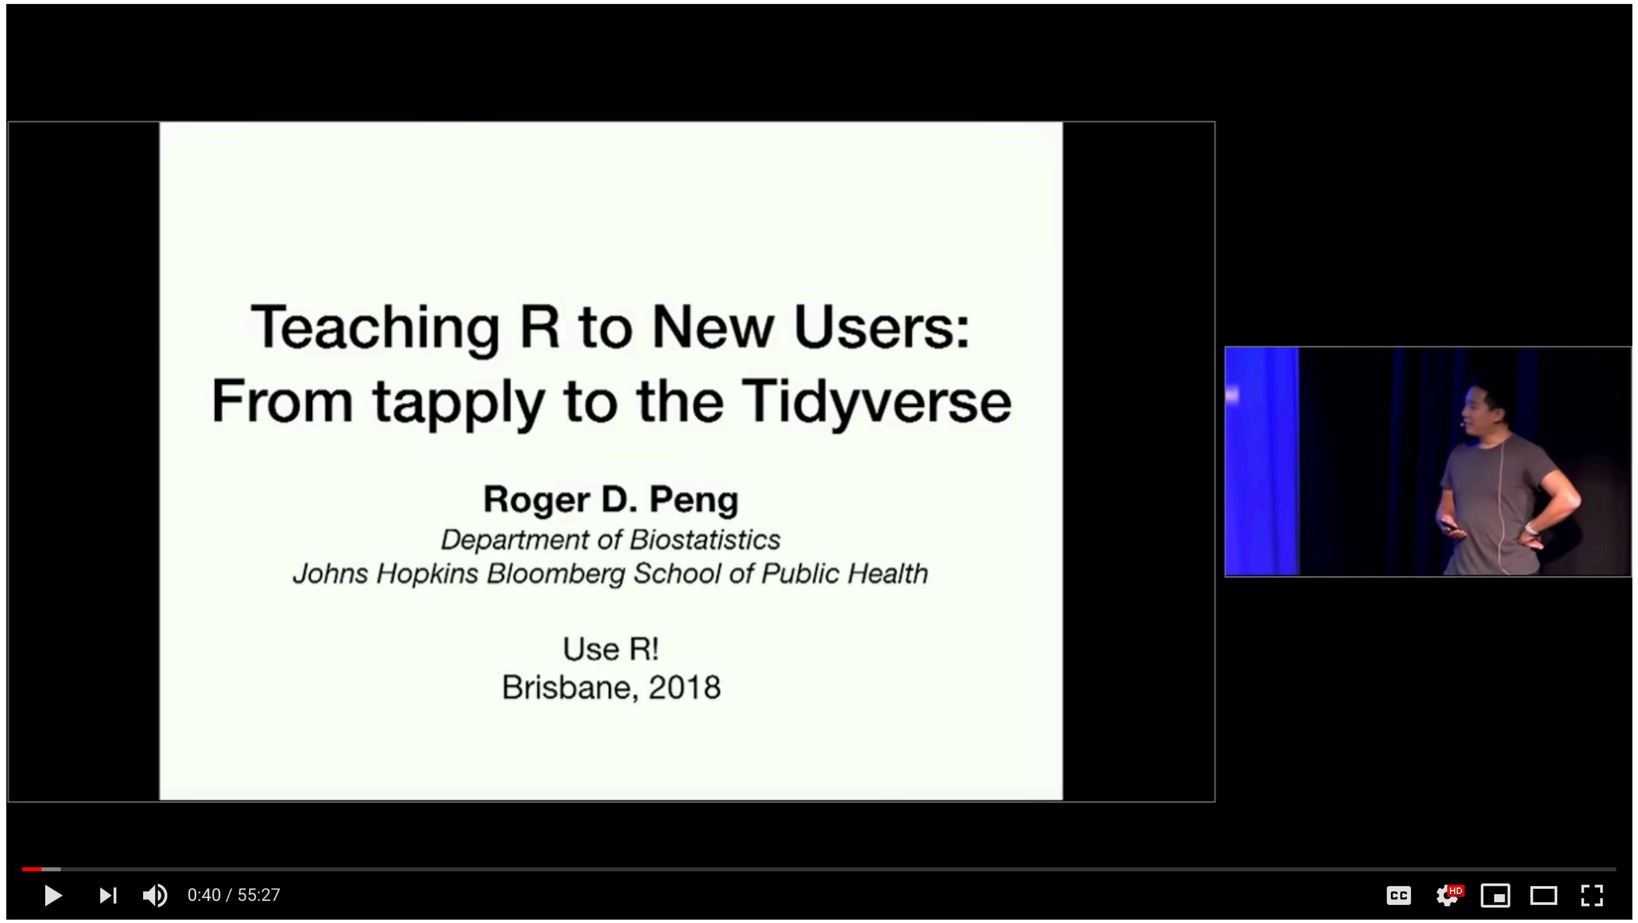

Roger Peng's useR 2018 keynote,
Teaching R to New Users: From tapply to Tidyverse (video, blogpost)

"The ambiguity [of the S language] is real and goes to a key objective: we wanted users to be able to begin in an **interactive environment**, where they did not consciously think of themselves as programming. Then as their needs became clearer and their sophistication increased, they should be able to **slide gradually into programming**, when the language and system aspects would become more important."

- John Chambers,"Stages in the Evolution of S."
Quoted in Roger Peng's keynote

McNamara, Amelia Ahlers. (2015).

Bridging the Gap Between Tools for Learning and for Doing Statistics.

http://bit.ly/BridgingTheToolGap

or, two arXiv pre-prints:

http://bit.ly/ModernStatComputing

http://bit.ly/StateOfStatCompEd

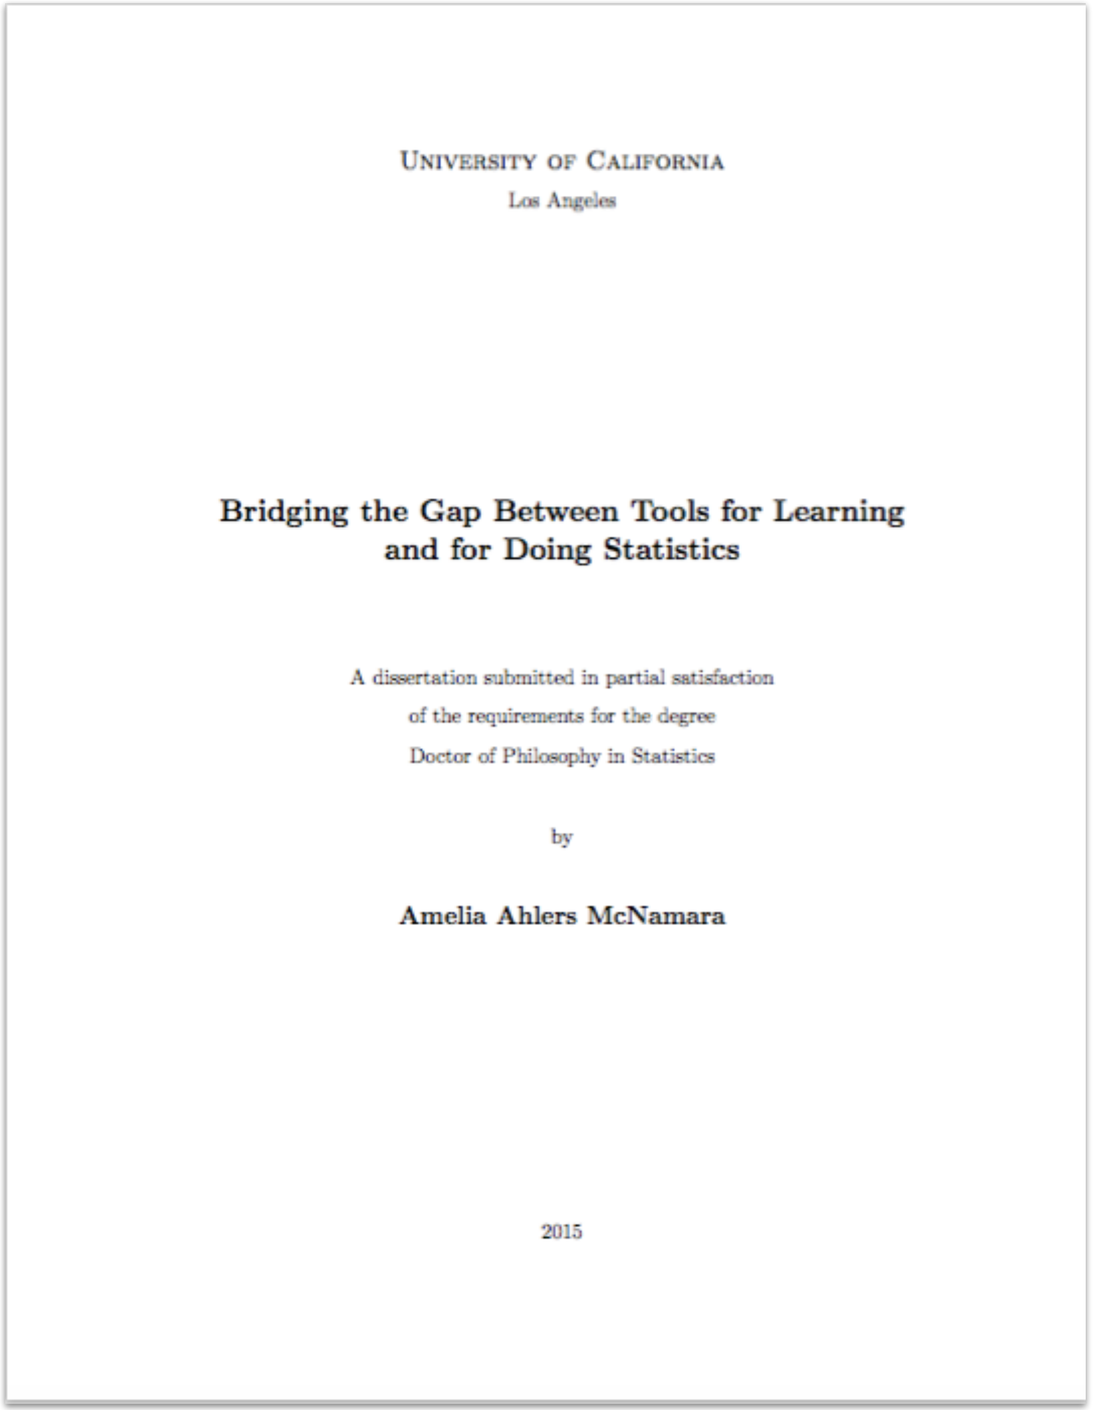

UNIVERSITY OF CALIFORNIA
Los Angeles

Bridging the Gap Between Tools for Learning and for Doing Statistics

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Statistics

by

Amelia Ahlers McNamara

2015

http://bit.ly/KeyAttributesPreprint

THE AMERICAN STATISTICIAN
2018, VOL. 0, NO. 0, 1–10: Statistical Computing and Graphics
https://doi.org/10.1080/00031305.2018.1482784

Key Attributes of a Modern Statistical Computing Tool

Amelia McNamara

Statistical and Data Sciences, Smith College, Northampton, MA

ABSTRACT

In the 1990s, statisticians began thinking in a principled way about how computation could better support the learning and doing of statistics. Since then, the pace of software development has accelerated, advancements in computing and data science have moved the goalposts, and it is time to reassess. Software continues to be developed to help do and learn statistics, but there is little critical evaluation of the resulting tools, and no accepted framework with which to critique them. This article presents a set of attributes necessary for a modern statistical computing tool. The framework was designed to be broadly applicable to both novice and expert users, with a particular focus on making more supportive statistical computing environments. A modern statistical computing tool should be accessible, provide easy entry, privilege data as a first-order object, support exploratory and confirmatory analysis, allow for flexible plot creation, support randomization, be interactive, include inherent documentation, support narrative, publishing, and reproducibility, and be flexible to extensions. Ideally, all these attributes could be incorporated into one tool, supporting users at all levels, but a more reasonable goal is for tools designed for novices and professionals to "reach across the gap," taking inspiration from each others' strengths.

ARTICLE HISTORY
Received September 2016
Revised May 2018

KEYWORDS
Bootstrap; Data visualization; Exploratory data analysis; Randomization; Reproducibility; Software design; Software evaluation

1. Introduction

Tools shape the way we see the world, and statistical computing tools are starting to blur, and we believe this lowers the barrier

http://bit.ly/KeyAttributesPreprint

Table 1. Summary of attributes.

1. Accessibility
2. Easy entry for novice users
3. Data as a first-order persistent object
4. Support for a cycle of exploratory and confirmatory analysis
5. Flexible plot creation
6. Support for randomization throughout
7. Interactivity at every level
8. Inherent documentation
9. Simple support for narrative, publishing, and reproducibility
10. Flexibility to build extensions

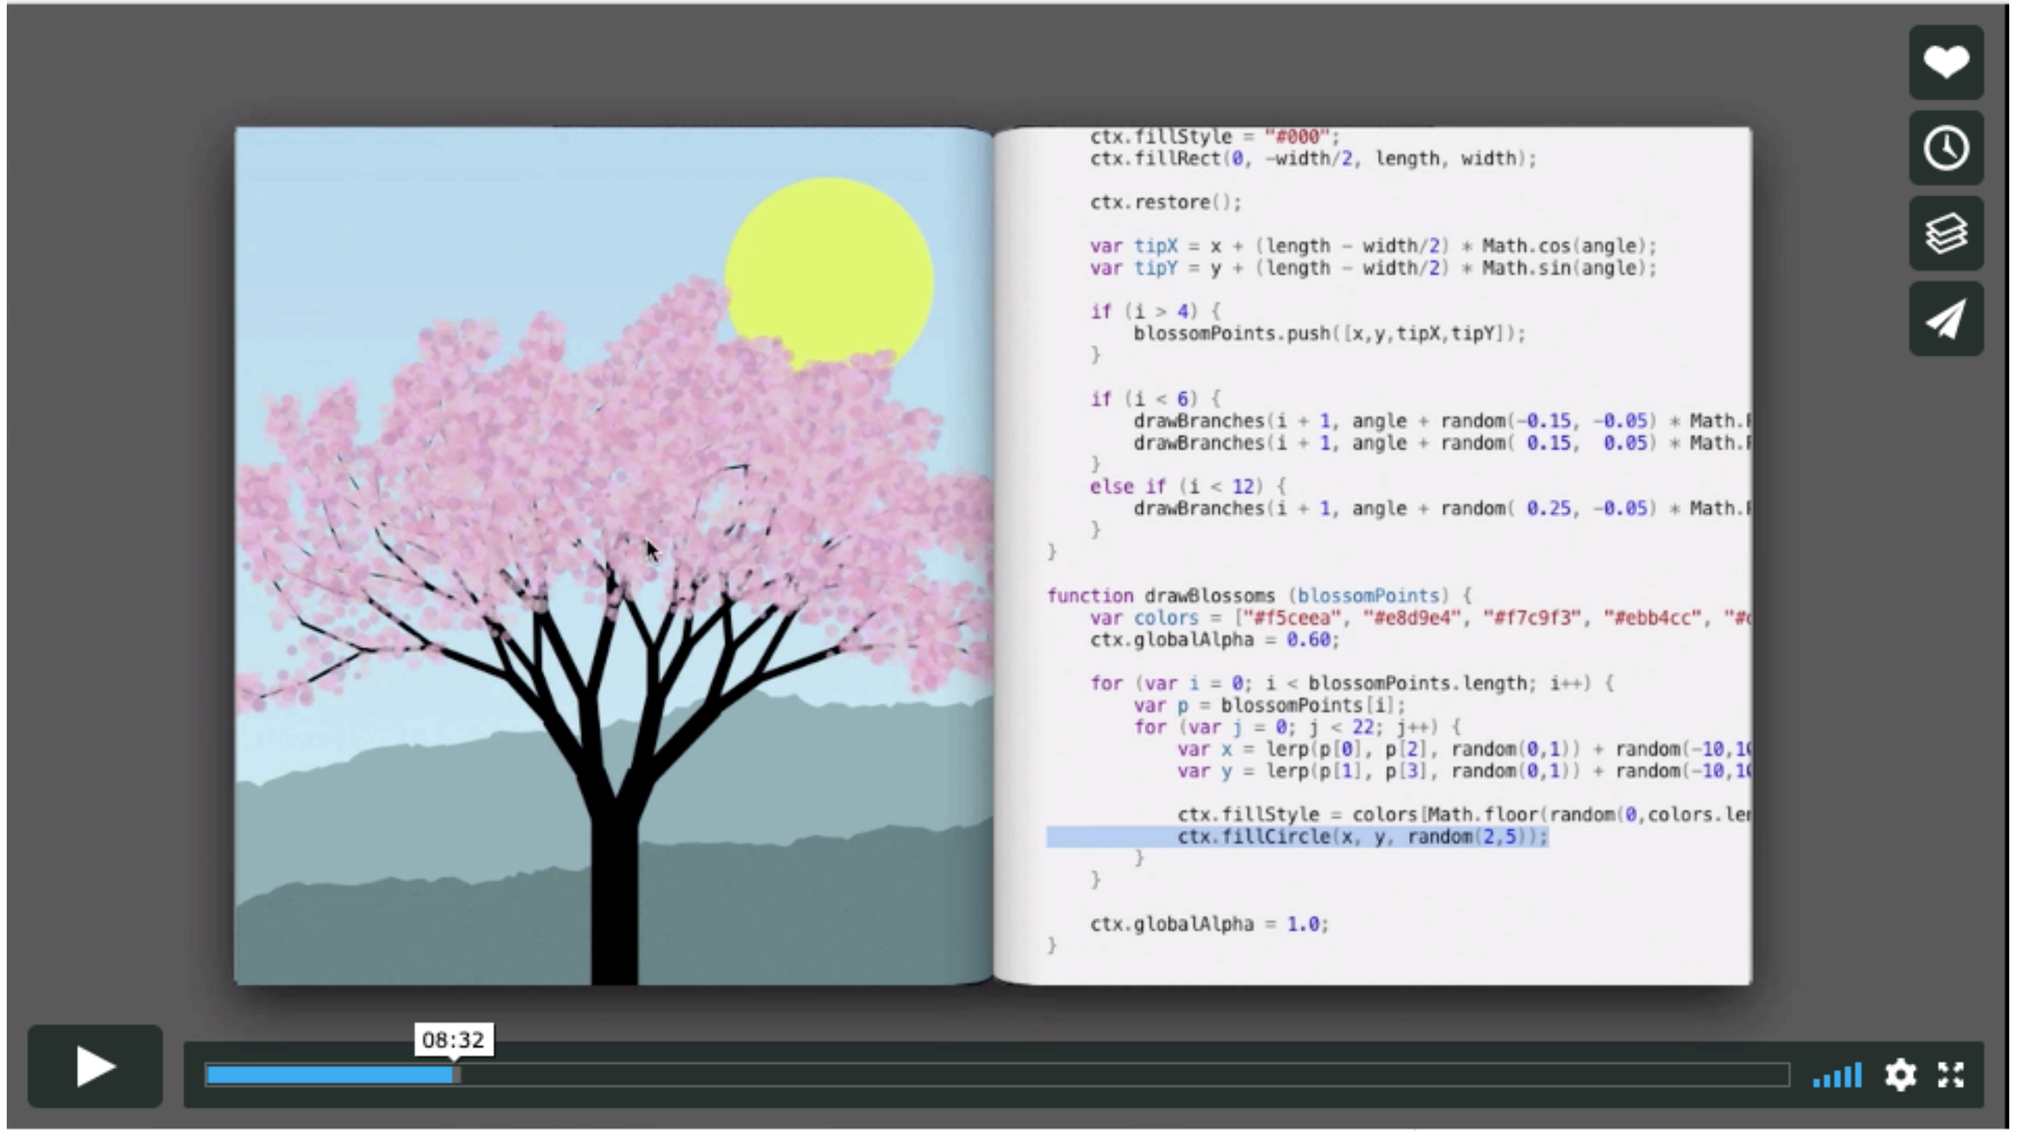

Bret Victor - Inventing on Principle (video)

R Syntax Comparison : : **CHEAT SHEET**

Dollar sign syntax

```
goal(data$x, data$y)
```

SUMMARY STATISTICS:

one continuous variable:
```
mean(mtcars$mpg)
```

one categorical variable:
```
table(mtcars$cyl)
```

two categorical variables:
```
table(mtcars$cyl, mtcars$am)
```

one continuous, one categorical:
```
mean(mtcars$mpg[mtcars$cyl==4])
mean(mtcars$mpg[mtcars$cyl==6])
mean(mtcars$mpg[mtcars$cyl==8])
```

PLOTTING:

one continuous variable:
```
hist(mtcars$disp)

boxplot(mtcars$disp)
```

one categorical variable:
```
barplot(table(mtcars$cyl))
```

two continuous variables:
```
plot(mtcars$disp, mtcars$mpg)
```

two categorical variables:
```
mosaicplot(table(mtcars$am, mtcars$cyl))
```

one continuous, one categorical:
```
histogram(mtcars$disp[mtcars$cyl==4])
histogram(mtcars$disp[mtcars$cyl==6])
histogram(mtcars$disp[mtcars$cyl==8])

boxplot(mtcars$disp[mtcars$cyl==4])
boxplot(mtcars$disp[mtcars$cyl==6])
boxplot(mtcars$disp[mtcars$cyl==8])
```

WRANGLING:

subsetting:
```
mtcars[mtcars$mpg>30, ]
```

making a new variable:
```
mtcars$efficient[mtcars$mpg>30] <- TRUE
mtcars$efficient[mtcars$mpg<30] <- FALSE
```

Formula syntax

```
goal(y~x|z, data=data, group=w)
```

SUMMARY STATISTICS:

one continuous variable:
```
mosaic::mean(~mpg, data=mtcars)
```

one categorical variable:
```
mosaic::tally(~cyl, data=mtcars)
```

two categorical variables:
```
mosaic::tally(cyl~am, data=mtcars)
```

one continuous, one categorical:
```
mosaic::mean(mpg~cyl, data=mtcars)
```

tilde

PLOTTING:

one continuous variable:
```
lattice::histogram(~disp, data=mtcars)

lattice::bwplot(~disp, data=mtcars)
```

one categorical variable:
```
mosaic::bargraph(~cyl, data=mtcars)
```

two continuous variables:
```
lattice::xyplot(mpg~disp, data=mtcars)
```

two categorical variables:
```
mosaic::bargraph(~am, data=mtcars, group=cyl)
```

one continuous, one categorical:
```
lattice::histogram(~disp|cyl, data=mtcars)

lattice::bwplot(cyl~disp, data=mtcars)
```

The variety of R syntaxes give you many ways to "say" the same thing

read **across** the cheatsheet to see how different syntaxes approach the same problem

Tidyverse syntax

```
data %>% goal(x)
```

SUMMARY STATISTICS:

one continuous variable:
```
mtcars %>% dplyr::summarize(mean(mpg))
```

one categorical variable:
```
mtcars %>% dplyr::group_by(cyl) %>%
dplyr::summarize(n())
```

the pipe

two categorical variables:
```
mtcars %>% dplyr::group_by(cyl, am) %>%
dplyr::summarize(n())
```

one continuous, one categorical:
```
mtcars %>%  dplyr::group_by(cyl) %>%
  dplyr::summarize(mean(mpg))
```

PLOTTING:

one continuous variable:
```
ggplot2::qplot(x=mpg, data=mtcars, geom = "histogram")

ggplot2::qplot(y=disp, x=1, data=mtcars, geom="boxplot")
```

one categorical variable:
```
ggplot2::qplot(x=cyl, data=mtcars, geom="bar")
```

two continuous variables:
```
ggplot2::qplot(x=disp, y=mpg, data=mtcars, geom="point")
```

two categorical variables:
```
ggplot2::qplot(x=factor(cyl), data=mtcars, geom="bar") +
       facet_grid(.~am)
```

one continuous, one categorical:
```
ggplot2::qplot(x=disp, data=mtcars, geom = "histogram") +
       facet_grid(.~cyl)

ggplot2::qplot(y=disp, x=factor(cyl), data=mtcars,
              geom="boxplot")
```

WRANGLING:

subsetting:
```
mtcars %>% dplyr::filter(mpg>30)
```

making a new variable:
```
mtcars <- mtcars %>%
dplyr::mutate(efficient = if_else(mpg>30,  TRUE, FALSE))
```

SMITH COLLEGE

https://www.rstudio.com/resources/cheatsheets/

Safety

Open Wrangling categorical data in R -- Amelia McNamara and Nicholas... #26 Opened by jennybc

AmeliaMN commented on May 4, 2016

I'm just coming off of final student projects, so I'm thinking about things that might be useful to new data practitioners in R. Some ideas

1. A comparison of different ways to express the same action using different syntaxes. Probably I would focus on subsetting in different ways (rows/columns). For example, `mtcars %>% select(wt)` versus `mtcars[,6]` versus `mtcars["wt"]` or `mtcars %>% filter(mpg>30)` versus `mtcars[mtcars$mpg>30,]` Other than subsetting, I could also look at ways to create new variables, e.g. `mtcars %>% mutate(ratio = gear/carb)` versus `mtcars$ratio <- mtcars$gear/mtcars$carb` This one might be too simplistic and/or too related to #8.

2. Explanation of factors and how to recode them. I might need to talk to **@hadley** about best practices here, because my current solutions are a bit hacky and I often get warning messages. There are a few different factor issues I/my students often run into.

   a. Starting with the simplest: you want to change the formatting of the factor labels so they all start with a capital letter. When doing this, it is so easy to accidentally ruin your data, so you need a little EDA workflow: look at the `summary()` of the factor and note the numbers in each category, then try your level changes, then look at the `summary()` again.

   c. Another problem is reordering factor levels-- maybe because you want ggplot2 to show them in a particular order, or because there is some inherent order to your levels. Again, I often do `SummaryStats <- SummaryStats %>% mutate(Treatment = factor(Treatment, levels=c("Control", "E25", "E50", "E100")))` and ruin everything before I remember it's actually `SummaryStats <- SummaryStats %>% mutate(treatment = factor(treatment, levels=levels(treatment)[c(1,3,4,2)]))`

   b. Even easier to mess up is when you have a categorical variable with 10+ categories and want to condense down to 3-4. Again, this is where my hack often runs into errors.

None yet

Milestone

No milestone

Notifications

Unsubscribe

You're receiving notifications because you were mentioned.

7 participants

Preprint

View 14 tweets

NOT PEER-REVIEWED

"PeerJ Preprints" is a venue for early communication or feedback before peer review. Data may be preliminary.

Learn more about preprints or browse peer-reviewed articles instead.

Wrangling categorical data in R

Research article | Computer Education | Data Science | Scientific Computing and Simulation | Social Computing

Me

@askdrstats

Amelia McNamara[1], Nicholas J Horton[2]

August 30, 2017

http://bit.ly/WranglingCats

Highlighted in Practical Data Science for Stats

Author and article information

Abstract

Data wrangling is a critical foundation of data science, and wrangling of categorical data is an important component of this process. However,

Journal
The American Statistician ›
Volume 72, 2018 - Issue 1: Special Issue on Data Science

Article

Wrangling Categorical Data in R

Amelia McNamara & Nicholas J. Horton

Pages 97-104 | Received 01 May 2017, Accepted author version posted online: 27 Jul 2017, Published online: 27 Jul 2017

 https://doi.org/10.1080/00031305.2017.1356375

Amelia McNamara[a*] http://orcid.org/0000-0003-4916-2433 & Nicholas J. Horton[b] http://orcid.org/0000-0003-3332-4311

[a] Program in Statistical and Data Sciences, Smith College, Northampton, MA

[b] Department of Mathematics and Statistics, Amherst College, Amherst, MA

CONTACT Amelia McNamara *amcnamara@smith.edu* Program in Statistical and Data Sciences,

```
> badApproach <- GSS$OpinionOfIncome
> summary(badApproach)
    Above average           Average     Below average        Don't know Far above average
              483              1118               666                21                65
Far below average         No answer              NA's
              179                 6                 2
> levels(badApproach) <- c("Far above average", "Above average",
+                          "Average", "Below Average", "Far below average",
+                          "Don't know", "No answer")
> summary(badApproach)
Far above average     Above average           Average     Below Average Far below average
              483              1118               666                21                65
       Don't know         No answer              NA's
              179                 6                 2
>
```

```
> summary(GSS$BaseOpinionOfIncome)
    Above average           Average     Below average        Don't know Far above average
              483              1118               666                21                65
Far below average        No answer              NA's
              179                 6                 2
> GSS$BaseOpinionOfIncome <-
+   factor(GSS$BaseOpinionOfIncome,
+          levels = c("Far above average", "Above average", "Average ", "Below Average",
+                     "Far below average", "Don't know", "No answer"))
> summary(GSS$BaseOpinionOfIncome)
Far above average     Above average           Average     Below Average Far below average
               65               483                 0                 0               179
      Don't know        No answer              NA's
               21                 6              1786
>
```

```
> library(forcats)
> summary(GSS$OpinionOfIncome)
    Above average           Average     Below average        Don't know Far above average
              483              1118               666                21                65
Far below average        No answer              NA's
              179                 6                 2
> GSS <- GSS %>%
+   mutate(tidyOpinionOfIncome =
+            fct_relevel(OpinionOfIncome,
+                        "Far above average",
+                        "Above average",
+                        "Average",
+                        "Below average",
+                        "Far below average"))
> summary(GSS$tidyOpinionOfIncome)
Far above average     Above average           Average     Below average Far below average
               65               483              1118               666               179
       Don't know        No answer              NA's
               21                 6                 2
>
```

Amelia McNamara,
"Working with categorical data in R without losing your mind" (video, slides)

Community

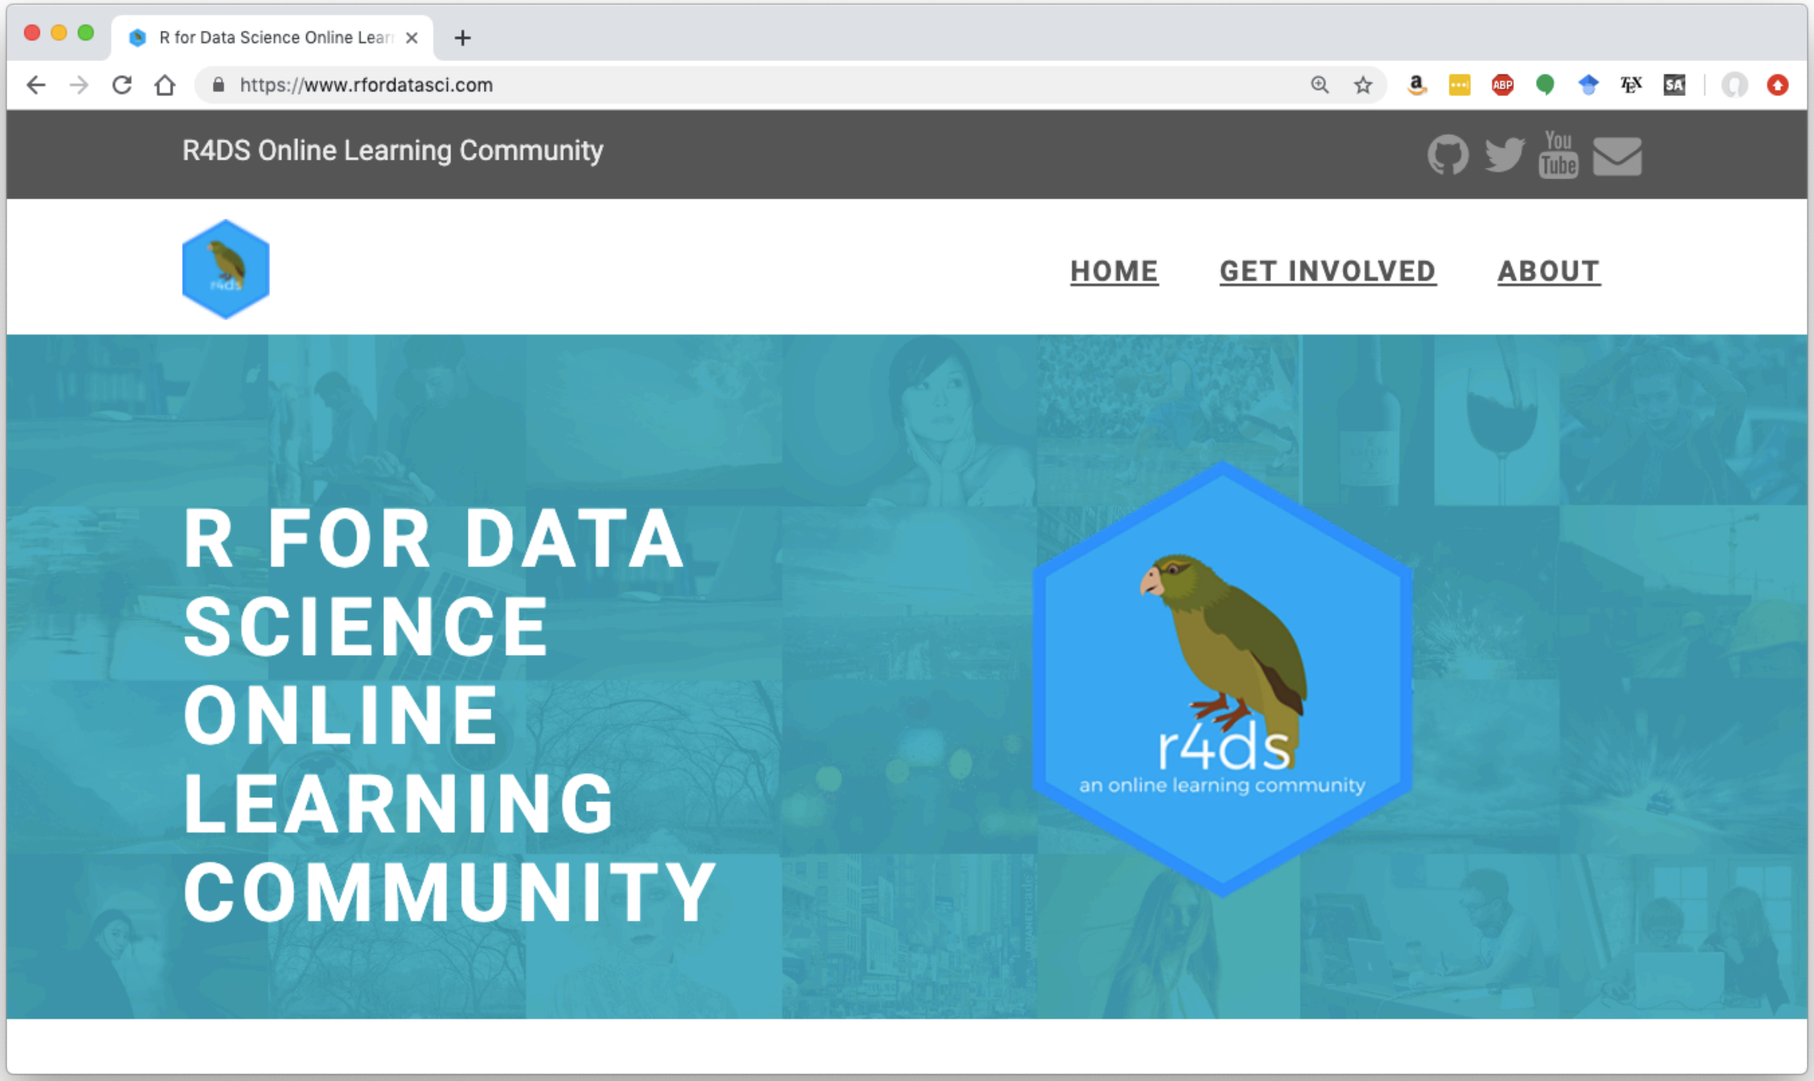

R for Data Science Online Learning Community
(r4ds book, community)

R Syntax Comparison : : **CHEAT SHEET**

Dollar sign syntax

```
goal(data$x, data$y)
```

SUMMARY STATISTICS:

one continuous variable:
```
mean(mtcars$mpg)
```

one categorical variable:
```
table(mtcars$cyl)
```

two categorical variables:
```
table(mtcars$cyl, mtcars$am)
```

one continuous, one categorical:
```
mean(mtcars$mpg[mtcars$cyl==4])
mean(mtcars$mpg[mtcars$cyl==6])
mean(mtcars$mpg[mtcars$cyl==8])
```

PLOTTING:

one continuous variable:
```
hist(mtcars$disp)

boxplot(mtcars$disp)
```

one categorical variable:
```
barplot(table(mtcars$cyl))
```

two continuous variables:
```
plot(mtcars$disp, mtcars$mpg)
```

two categorical variables:
```
mosaicplot(table(mtcars$am, mtcars$cyl))
```

one continuous, one categorical:
```
histogram(mtcars$disp[mtcars$cyl==4])
histogram(mtcars$disp[mtcars$cyl==6])
histogram(mtcars$disp[mtcars$cyl==8])

boxplot(mtcars$disp[mtcars$cyl==4])
boxplot(mtcars$disp[mtcars$cyl==6])
boxplot(mtcars$disp[mtcars$cyl==8])
```

WRANGLING:

subsetting:
```
mtcars[mtcars$mpg>30, ]
```

making a new variable:
```
mtcars$efficient[mtcars$mpg>30] <- TRUE
mtcars$efficient[mtcars$mpg<30] <- FALSE
```

Formula syntax

```
goal(y~x|z, data=data, group=w)
```

SUMMARY STATISTICS:

one continuous variable:
```
mosaic::mean(~mpg, data=mtcars)
```

one categorical variable:
```
mosaic::tally(~cyl, data=mtcars)
```

two categorical variables:
```
mosaic::tally(cyl~am, data=mtcars)
```

one continuous, one categorical:
```
mosaic::mean(mpg~cyl, data=mtcars)
```

tilde

PLOTTING:

one continuous variable:
```
lattice::histogram(~disp, data=mtcars)

lattice::bwplot(~disp, data=mtcars)
```

one categorical variable:
```
mosaic::bargraph(~cyl, data=mtcars)
```

two continuous variables:
```
lattice::xyplot(mpg~disp, data=mtcars)
```

two categorical variables:
```
mosaic::bargraph(~am, data=mtcars, group=cyl)
```

one continuous, one categorical:
```
lattice::histogram(~disp|cyl, data=mtcars)

lattice::bwplot(cyl~disp, data=mtcars)
```

The variety of R syntaxes give you many ways to "say" the same thing

read **across** the cheatsheet to see how different syntaxes approach the same problem

Tidyverse syntax

```
data %>% goal(x)
```

SUMMARY STATISTICS:

one continuous variable:
```
mtcars %>% dplyr::summarize(mean(mpg))
```

one categorical variable:
```
mtcars %>% dplyr::group_by(cyl) %>%
dplyr::summarize(n())
```

the pipe

two categorical variables:
```
mtcars %>% dplyr::group_by(cyl, am) %>%
dplyr::summarize(n())
```

one continuous, one categorical:
```
mtcars %>%  dplyr::group_by(cyl) %>%
  dplyr::summarize(mean(mpg))
```

PLOTTING:

one continuous variable:
```
ggplot2::qplot(x=mpg, data=mtcars, geom = "histogram")

ggplot2::qplot(y=disp, x=1, data=mtcars, geom="boxplot")
```

one categorical variable:
```
ggplot2::qplot(x=cyl, data=mtcars, geom="bar")
```

two continuous variables:
```
ggplot2::qplot(x=disp, y=mpg, data=mtcars, geom="point")
```

two categorical variables:
```
ggplot2::qplot(x=factor(cyl), data=mtcars, geom="bar") +
       facet_grid(.~am)
```

one continuous, one categorical:
```
ggplot2::qplot(x=disp, data=mtcars, geom = "histogram") +
       facet_grid(.~cyl)

ggplot2::qplot(y=disp, x=factor(cyl), data=mtcars,
              geom="boxplot")
```

WRANGLING:

subsetting:
```
mtcars %>% dplyr::filter(mpg>30)
```

making a new variable:
```
mtcars <- mtcars %>%
dplyr::mutate(efficient = if_else(mpg>30,  TRUE, FALSE))
```

SMITH COLLEGE

https://www.rstudio.com/resources/cheatsheets/

Learnability?

A Randomized Controlled Trial on the Wild Wild West of Scientific Computing with Student Learners

Timothy Rafalski, P. Merlin Uesbeck, Cristina Panks-Meloney, Patrick Daleiden, William Allee, Amelia McNamara, Andreas Stefik

tl;dr they (we?) didn't find a difference between base, formula, and tidyverse syntaxes in completion time or number of errors (qualitative or interpreter)

https://dl.acm.org/citation.cfm?id=3291279.3339421

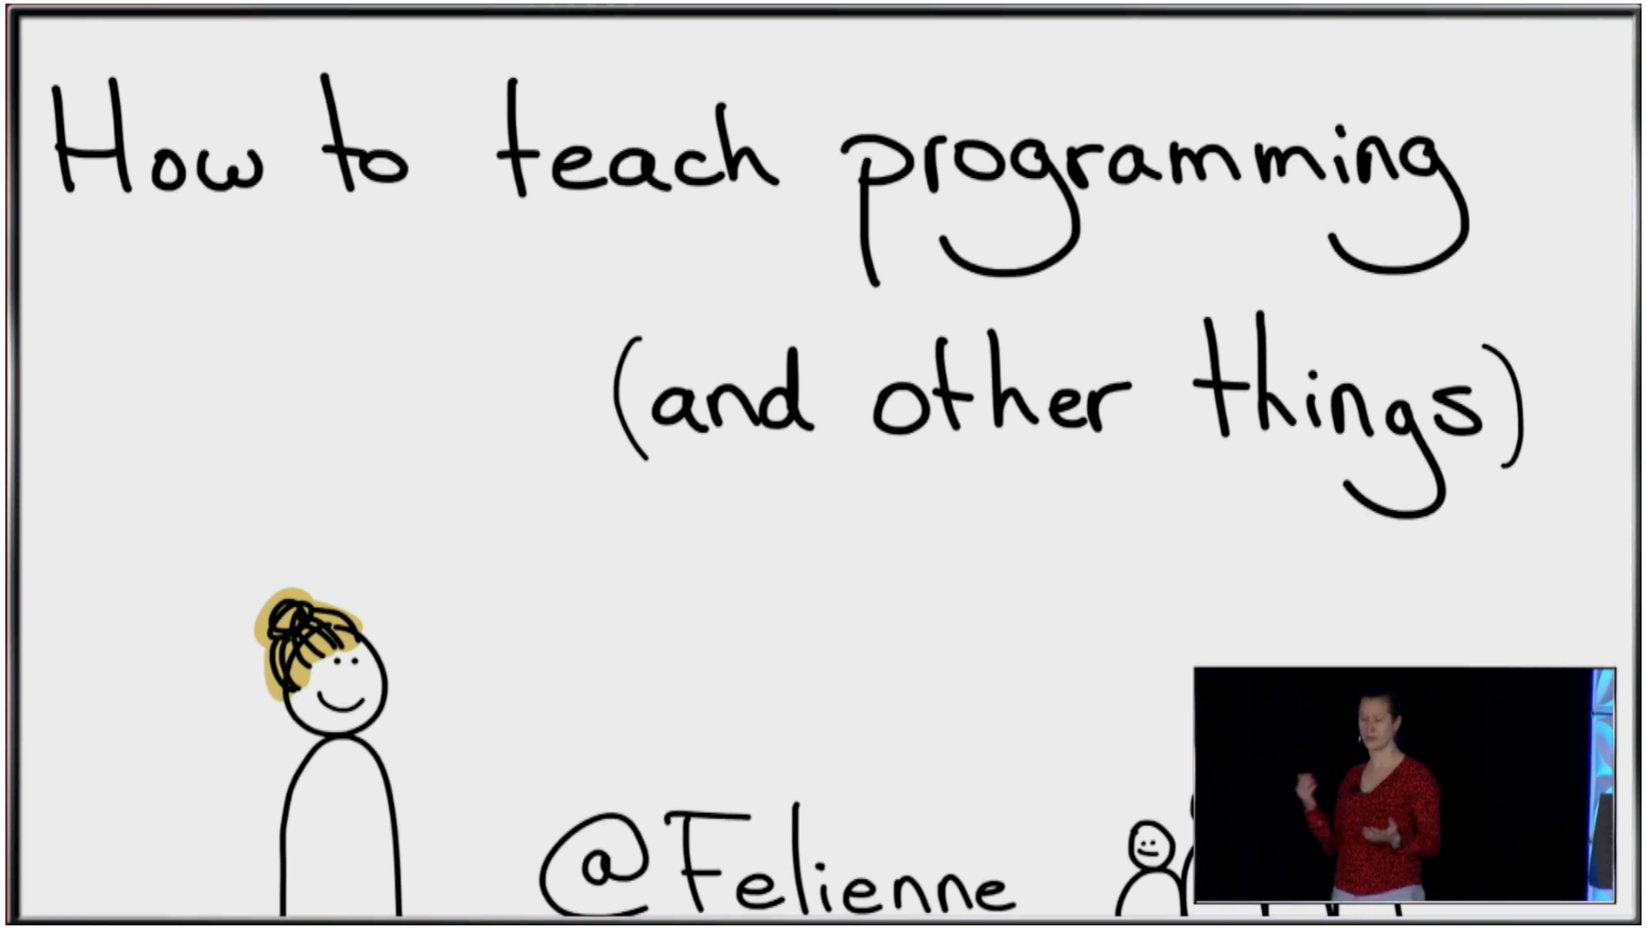

Felienne's rstudio::conf 2019 keynote,
Explicit Direct Instruction in Programming Education (video)

THE CARPENTRIES

software carpentry

DATA CARPENTRY

library carpentry

Teach the most useful thing first

Reduce cognitive load

Don’t touch the learner’s keyboard

Verbalize what you are doing

Universal Design

“Universal Design is the design and composition of an environment so that it can be accessed, understood and used to the greatest extent possible by all people regardless of their age, size, ability or disability.”

- Centre for Excellence in Universal Design

tl;dr R is best, but could be better

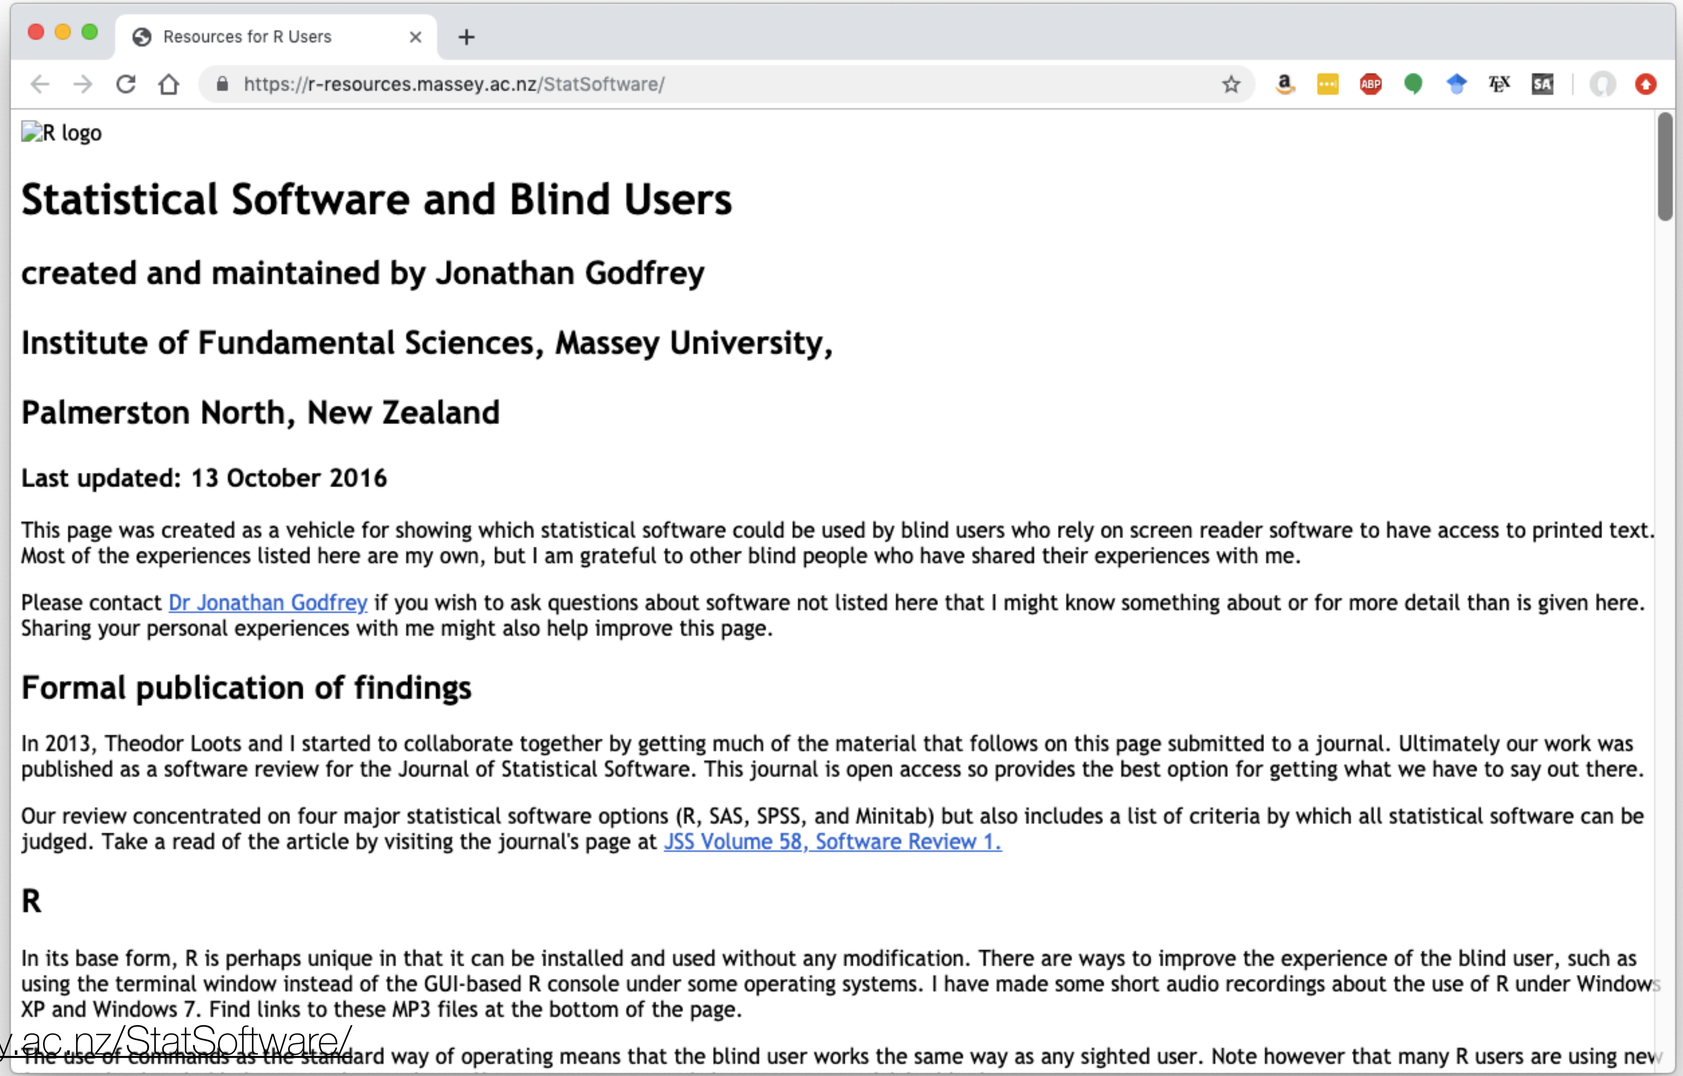

https://r-resources.massey.ac.nz/StatSoftware/

Emily Shea - "Perl Out Loud" (video)

“Due to this disability, I cannot type or write by hand. Many people have asked me about the stack that enables me to be productive in spite of this limitation. I hope this information is helpful both for people with more severe limitations, and for programmers with mild repetitive stress injuries who can benefit from reducing their keyboard use.”

–Naomi Saphra, "What Does a Coder Do If They Can't Type?"

“[…] by the time I was 20 I had developed a chronic wrist injury that I still have decades later. Maybe I should have walked away from computers, but by then I was hooked! That was a dark time, but eventually I struck upon a solution… I could dictate my code to someone else who could be my hands.”

–Ian Gilman, "Wrists & Apprentices"

Pair programming

https://www.youtube.com/watch?v=vgkahOzFH2Q

Readability

Summary statistics three ways

base

```
> mean(mtcars$mpg[mtcars$cyl==4])
[1] 26.66364
> mean(mtcars$mpg[mtcars$cyl==6])
[1] 19.74286
> mean(mtcars$mpg[mtcars$cyl==8])
[1] 15.1
```

mosaic

```
> mean(mpg~cyl, data=mtcars)
       4        6        8
26.66364 19.74286 15.10000
```

dplyr

```
> mtcars %>%
+   group_by(cyl) %>%
+   summarize(mean(mpg))
# A tibble: 3 x 2
    cyl `mean(mpg)`
  <dbl>       <dbl>
1     4        26.7
2     6        19.7
3     8        15.1
```

Scatterplot three ways

base

```
plot(mtcars$disp, mtcars$mpg)
```

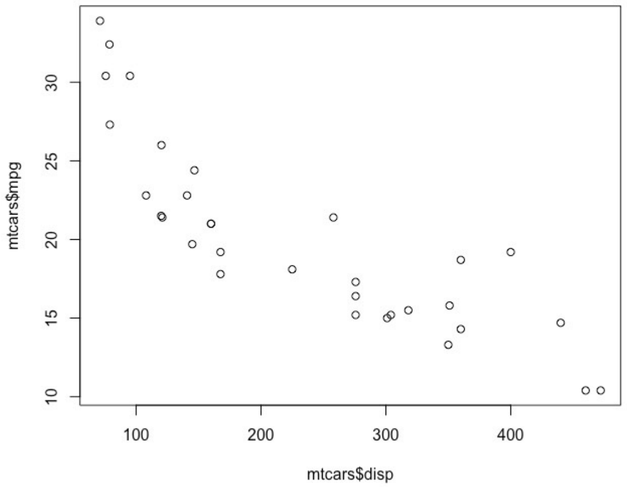

lattice

```
xyplot(mpg~disp, data=mtcars)
```

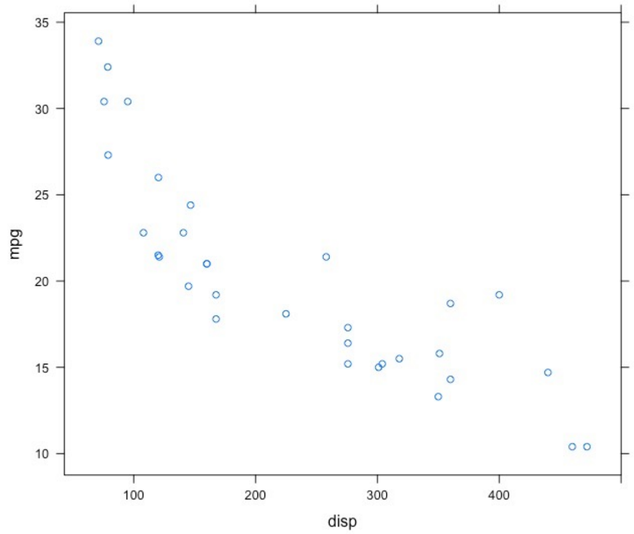

ggplot2

```
qplot(x=disp, y=mpg,
    data=mtcars, geom="point")
```

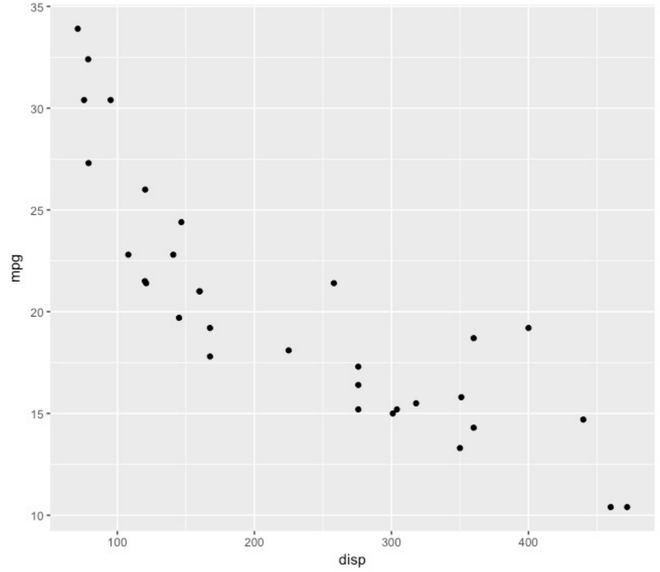

Sets of scatterplots three ways

```
xyplot(mpg ~ wt | as.factor(cyl), data = mtcars)

par(mfrow = c(1, 3))
plot(mtcars$wt[mtcars$cyl == 4], mtcars$mpg[mtcars$cyl == 4])
plot(mtcars$wt[mtcars$cyl == 6], mtcars$mpg[mtcars$cyl == 6])
plot(mtcars$wt[mtcars$cyl == 8], mtcars$mpg[mtcars$cyl == 8])


ggplot(mtcars, aes(x=wt, y = mpg)) + geom_point() +
 facet_grid(~cyl)
```

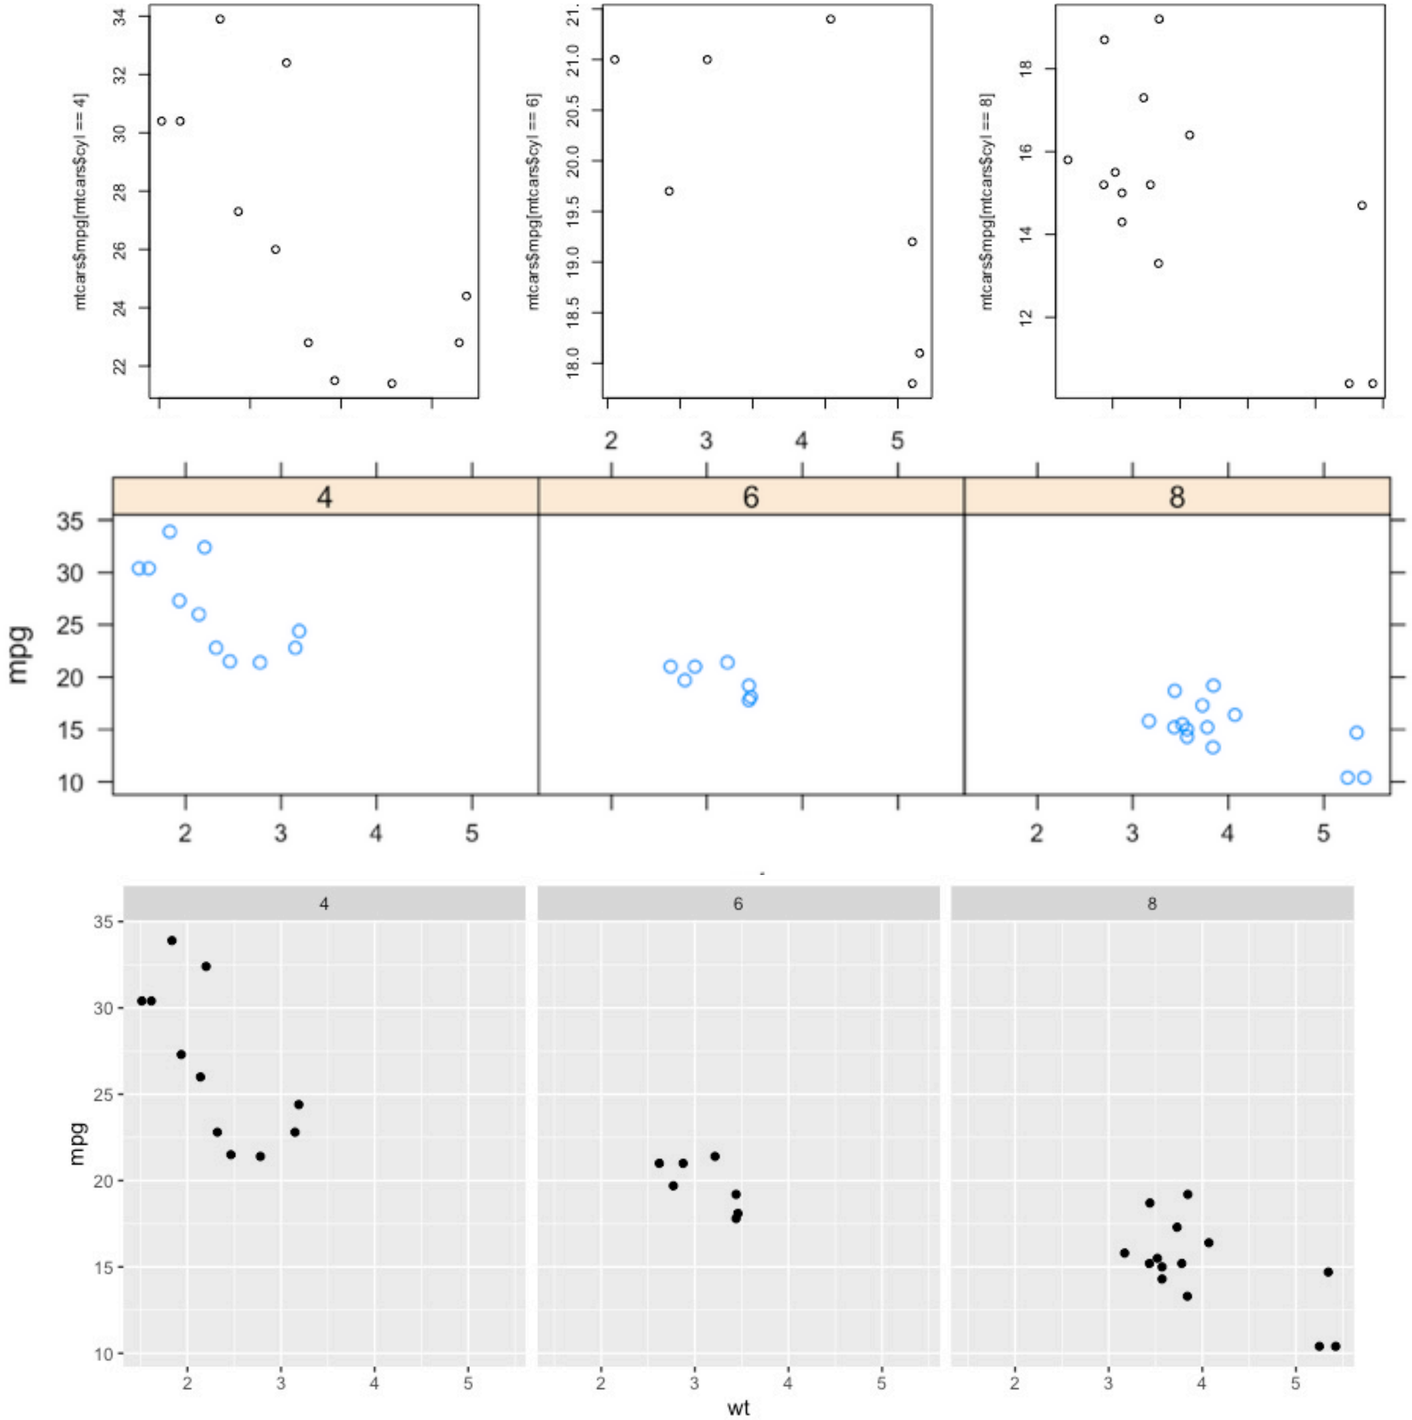
mtcars$mpg[mtcars$cyl == 4]
mtcars$mpg[mtcars$cyl == 6]
mtcars$mpg[mtcars$cyl == 8]
4
6
8
mpg
wt

“Less volume, more creativity.”

–Mike McCarthy / `mosaic` package philosophy

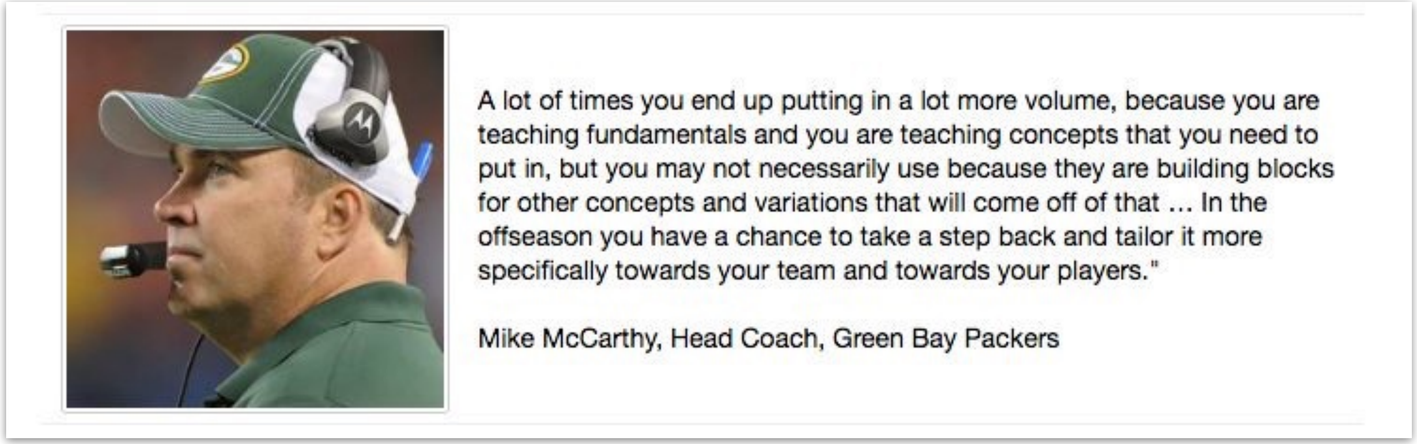

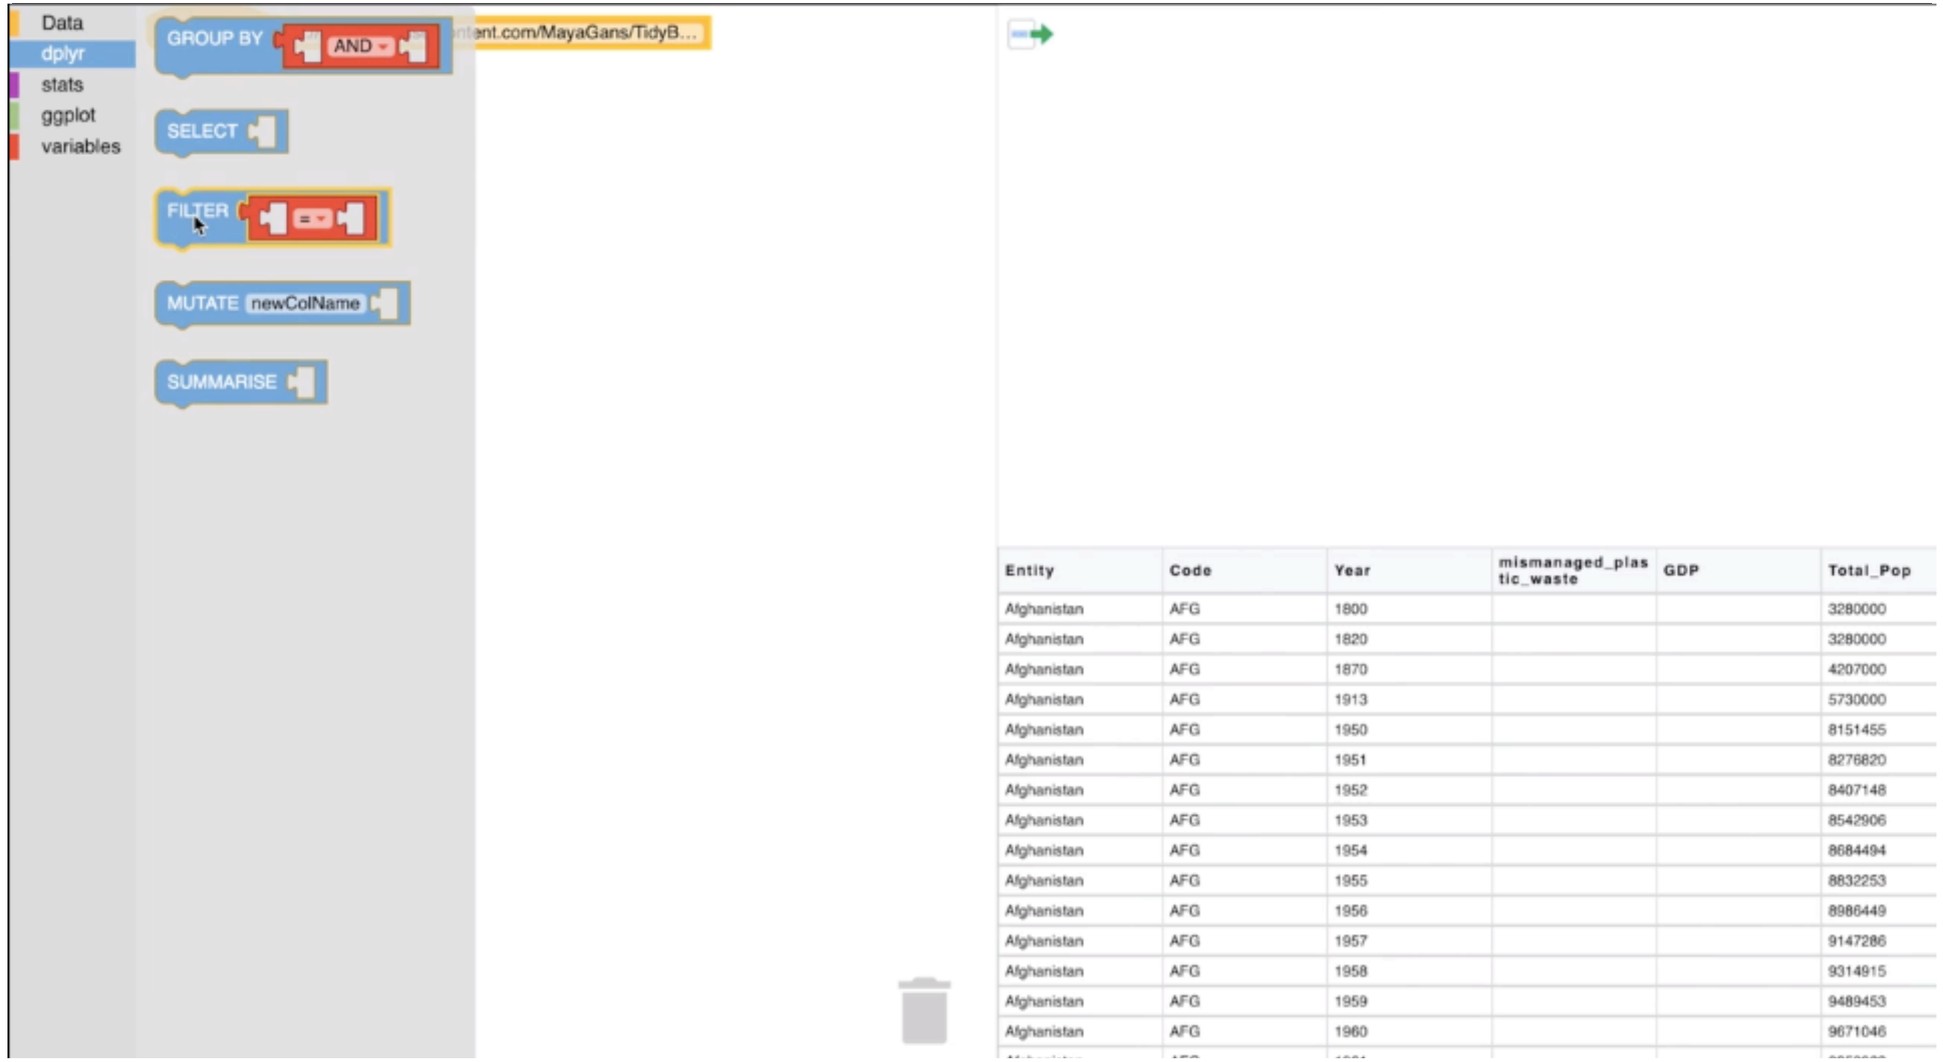

Entity	Code	Year	mismanaged_plastic_waste	GDP	Total_Pop
Afghanistan	AFG	1800			3280000
Afghanistan	AFG	1820			3280000
Afghanistan	AFG	1870			4207000
Afghanistan	AFG	1913			5730000
Afghanistan	AFG	1950			8151455
Afghanistan	AFG	1951			8276820
Afghanistan	AFG	1952			8407148
Afghanistan	AFG	1953			8542906
Afghanistan	AFG	1954			8684494
Afghanistan	AFG	1955			8832253
Afghanistan	AFG	1956			8986449
Afghanistan	AFG	1957			9147286
Afghanistan	AFG	1958			9314915
Afghanistan	AFG	1959			9489453
Afghanistan	AFG	1960			9671046

Maya Gans, tidyblocks demo (video)

data8 exercise 1 (python)

```
# Count how many times the names Jim, Tom, and Huck appear in each chapter.

counts = Table().with_columns([
        'Jim', np.char.count(huck_finn_chapters, 'Jim'),
        'Tom', np.char.count(huck_finn_chapters, 'Tom'),
        'Huck', np.char.count(huck_finn_chapters, 'Huck')
    ])

# Plot the cumulative counts:
# how many times in Chapter 1, how many times in Chapters 1 and 2, and so
on.

cum_counts = counts.cumsum().with_column('Chapter', np.arange(1, 44, 1))
cum_counts.plot(column_for_xticks=3)
plots.title('Cumulative Number of Times Each Name Appears', y=1.08);
```

https://predictablynoisy.com/data8-template/textbook/01/3/1/literary-characters/

Exercise ported to R

```
# Count how many times the names Jim, Tom, and Huck appear in each chapter.

counts = huck_finn_chapters %>%
  filter(word %in% c('jim', 'tom', 'huck')) %>%
  group_by(chapter, word) %>%
  summarize(count = n())

# Plot the cumulative counts:
# how many times in Chapter 1, how many times in Chapters 1 and 2, and so on.

cum_counts = counts %>%
  group_by(word) %>%
  mutate(cumulativesum = cumsum(count))

ggplot(cum_counts) +
  geom_line(aes(x=chapter, y = cumulativesum, color = word)) +
  ggtitle("Cumulative Number of Times Each Name Appears")
```

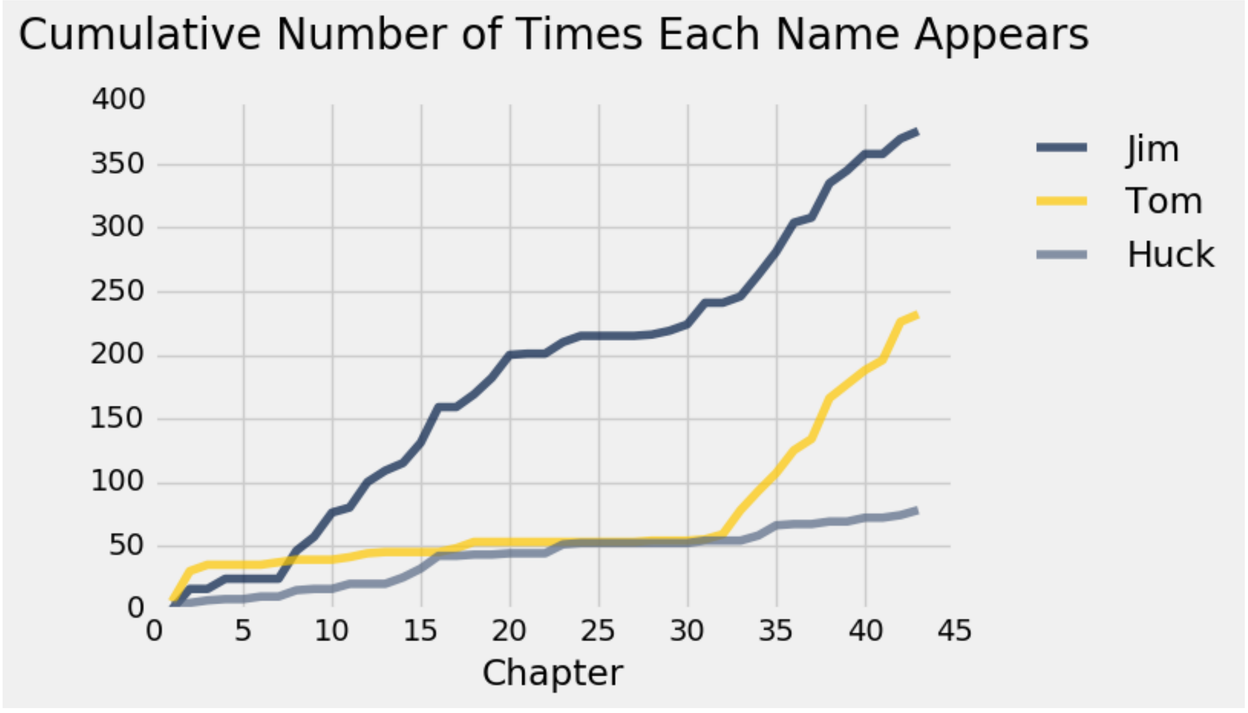
Cumulative Number of Times Each Name Appears
400
350
300
250
200
150
100
50
0
0
5
10
15
20
25
30
35
40
45
Chapter
Jim
Tom
Huck

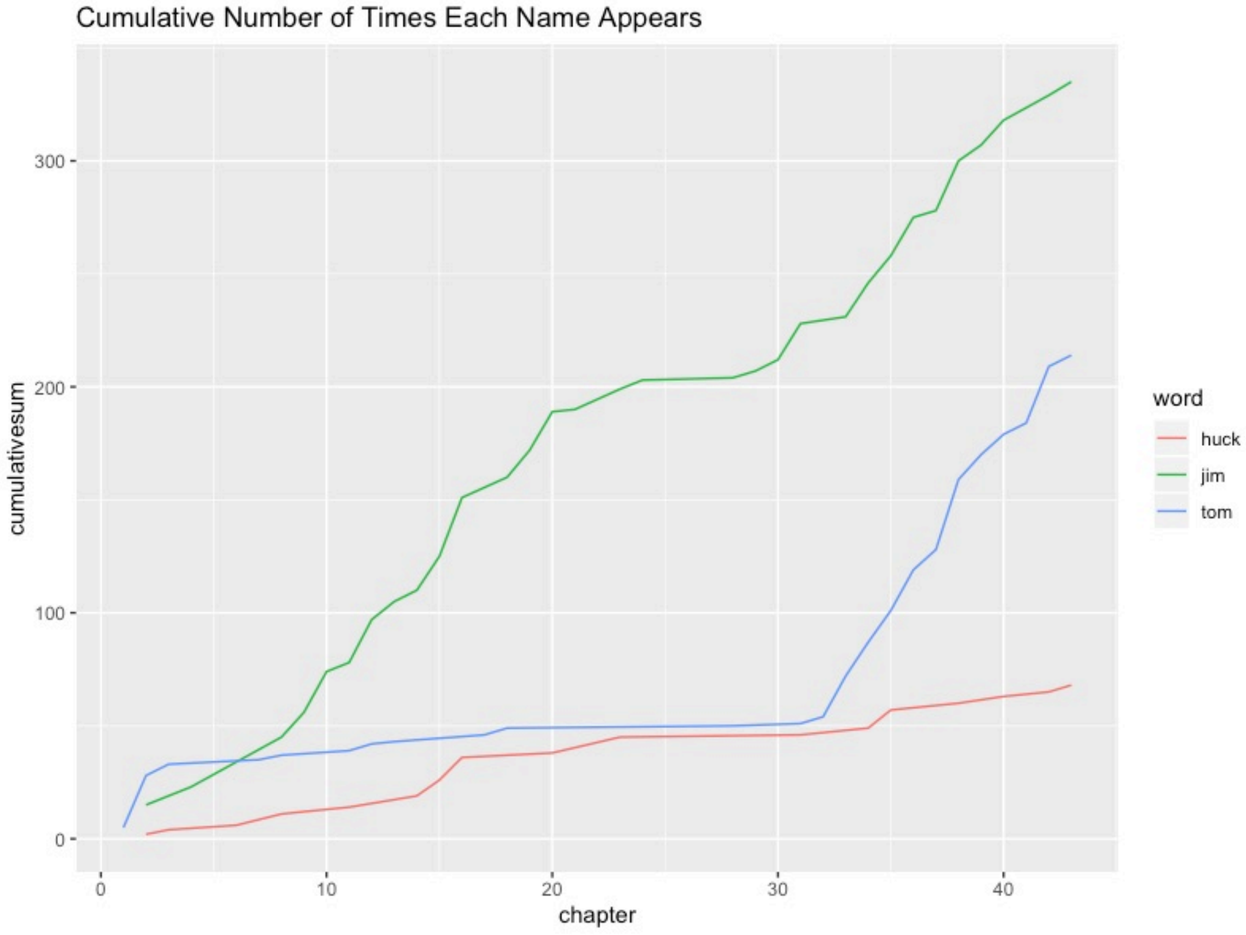
Cumulative Number of Times Each Name Appears
300
200
100
0
cumulativesum
0
10
20
30
40
chapter
word
huck
jim
tom

Amelia McNamara
@AmeliaMN

"Code as language" -@hadleywickham #icots10 😍😍

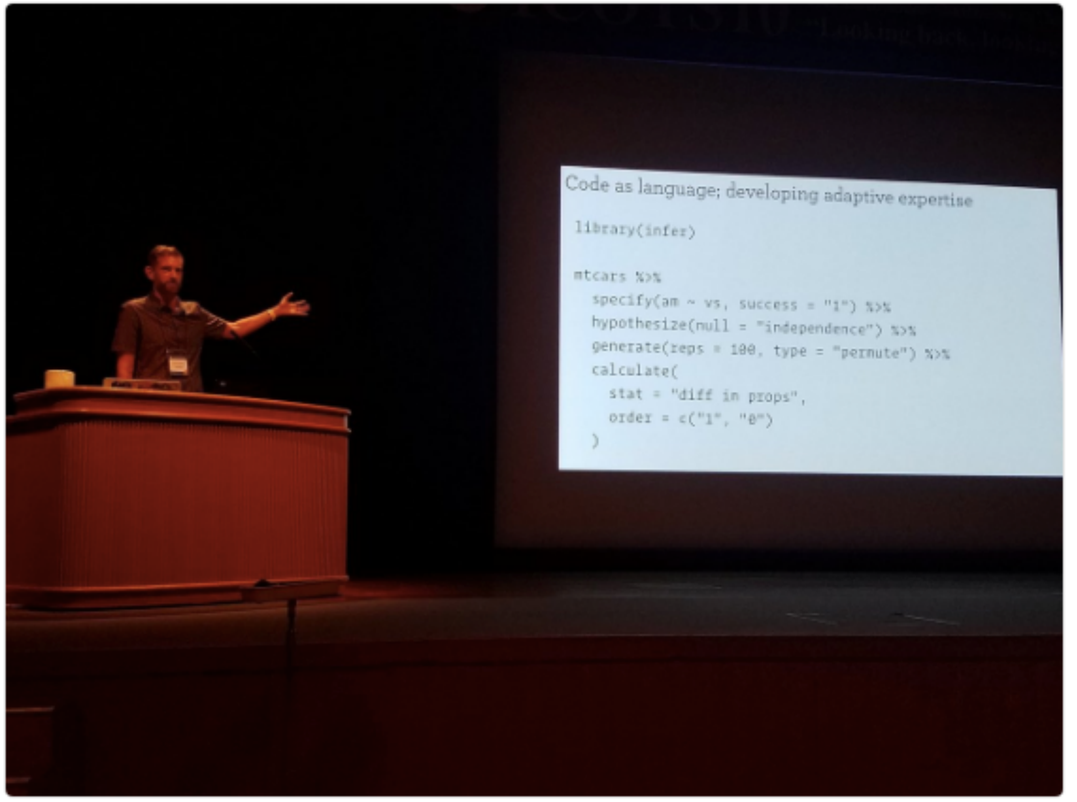

12:26 AM - 10 Jul 2018

30 Retweets **129** Likes

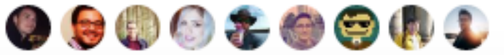

Lera Boroditsky, "How language shapes the way we think" (video)

“It’s a language for data analysis. And if you think that the language is a little incoherent, a little confusing, a bit of a maze, well then all I have to say is welcome to data analysis.”

– Roger Peng, useR keynote

Evaluating the Design of the R Language: Objects and Functions For Data Analysis

Floreal Morandat, Brandon Hill, Leo Osvald and Jan Vitek. ECOOP'12 Proceedings of the 26th European conference on Object-Oriented Programming. 2012. https://dl.acm.org/citation.cfm?id=2367172

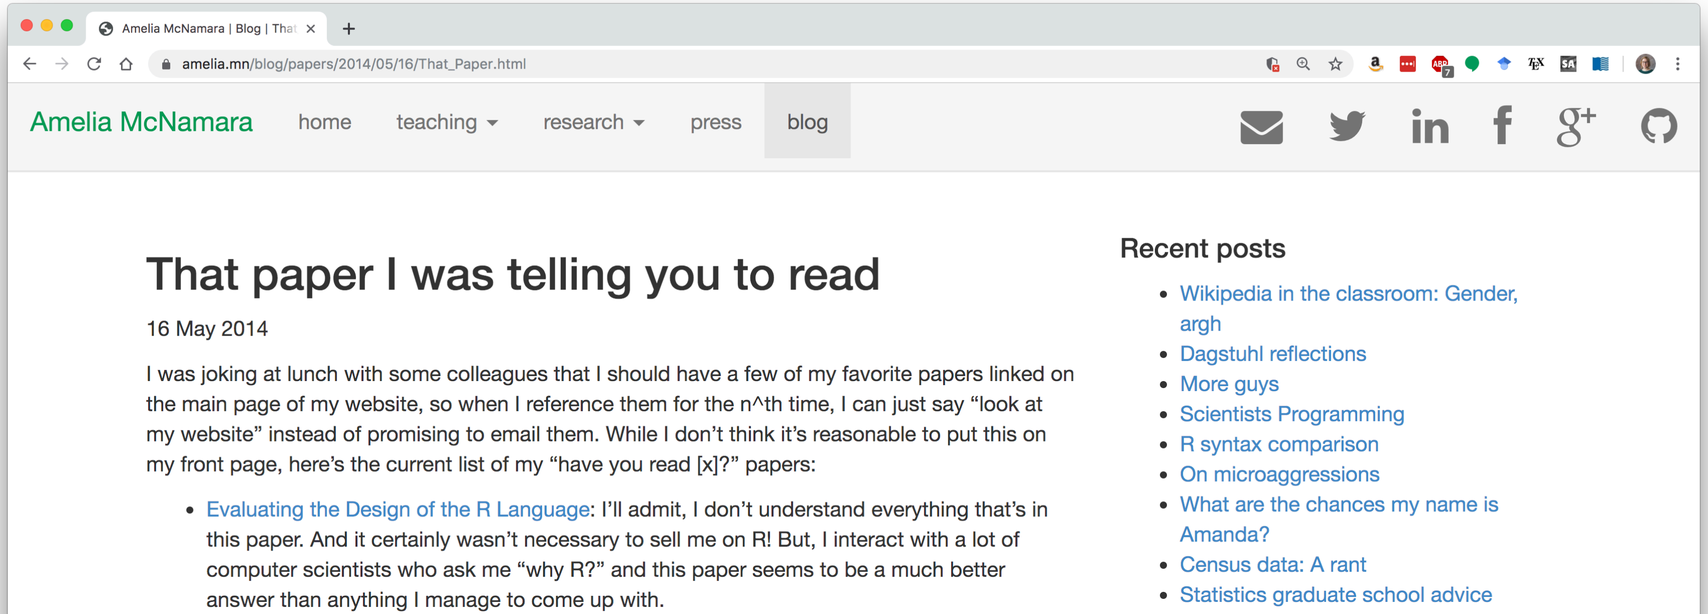

“We will be remiss in our duty to our students if we do not see that they learn to use the computer more easily, flexibly, and thoroughly than we ever have; we will be remiss in our duties to ourselves if we do not try to improve and broaden our own uses.”

–John Tukey,
The Technical Tools of Statistics
talking about the class of 1970

Takeaways

Design for humans

Universal Design is good for everyone

If you're teaching code, read it out loud

Tidyverse syntax seems more readable

Data analysis is sometimes incoherent!

Our goal is to make it more coherent over time


Thank you!

An aside— the dress

Fabric design is on GitHub
https://github.com/ameliamn/hexfabric

Fabric for purchase is on Spoonflower
https://www.spoonflower.com/profiles/ameliamn

I don’t take dress orders. Making this one was challenging enough!