Data Science & Machine Learning

Dan S. Reznik
www.dat-sci.com
December, 2019

Agenda

- Morning: 8:30~12:00
  - The Business Impact of Data Science and Machine Learning
  - What is Machine Learning
  - What is Data Science
  - Toolbox: R, Python, etc.
- Afternoon: Hands-on w/ R: 13:00~17:30
  - Set up DS environment on the cloud
    - GitHub, Amazon Machine Instance
  - Data Wrangle
    - Option1: Google Forms
    - Option2: Open Data: data.gov
- Takeaways

AI vs Traditional Businesses

- Physical => Digital
- Product, Service => Cloud Platform
- Oil => Data
- Manuf. & Ownership => Sharing economy
- Propietary Software => Open Frameworks
- Algorithms => Leverage diverse datasets

5,862 views | Apr 5, 2019, 03:59pm

AI Is Destroying Traditional Business Thinking

Tom Davenport Contributor ⓘ

Enterprise & Cloud

Business is being transformed by three trends

Big Data

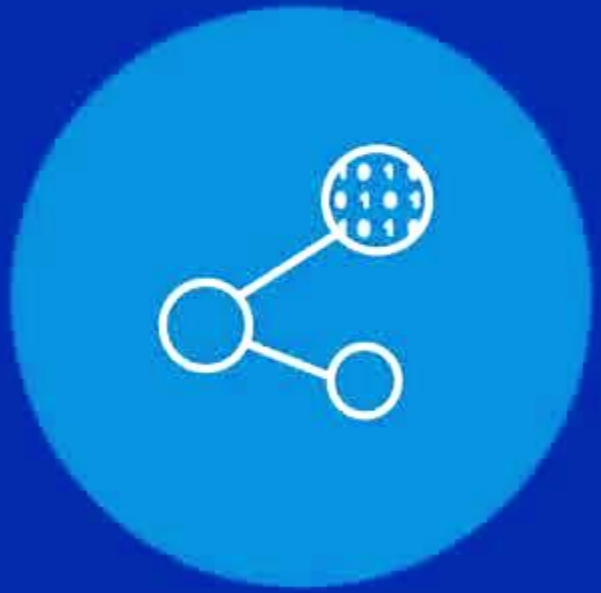

Cloud

Intelligence

ECONOMIC FORUM

COMMITTED TO IMPROVING THE STATE OF THE WORLD

The Jobs Landscape in 2022

emerging roles, global change by 2022

Top 10 Emerging

1. Data Analysts and Scientists
2. AI and Machine Learning Specialists
3. General and Operations Managers
4. Software and Applications Developers and Analysts
5. Sales and Marketing Professionals
6. Big Data Specialists
7. Digital Transformation Specialists
8. New Technology Specialists
9. Organisational Development Specialists
10. Information Technology Services

Key Technologies Will Deliver Change

Q. In your opinion, which three of these technologies have the most potential to change your organization over the next five years?

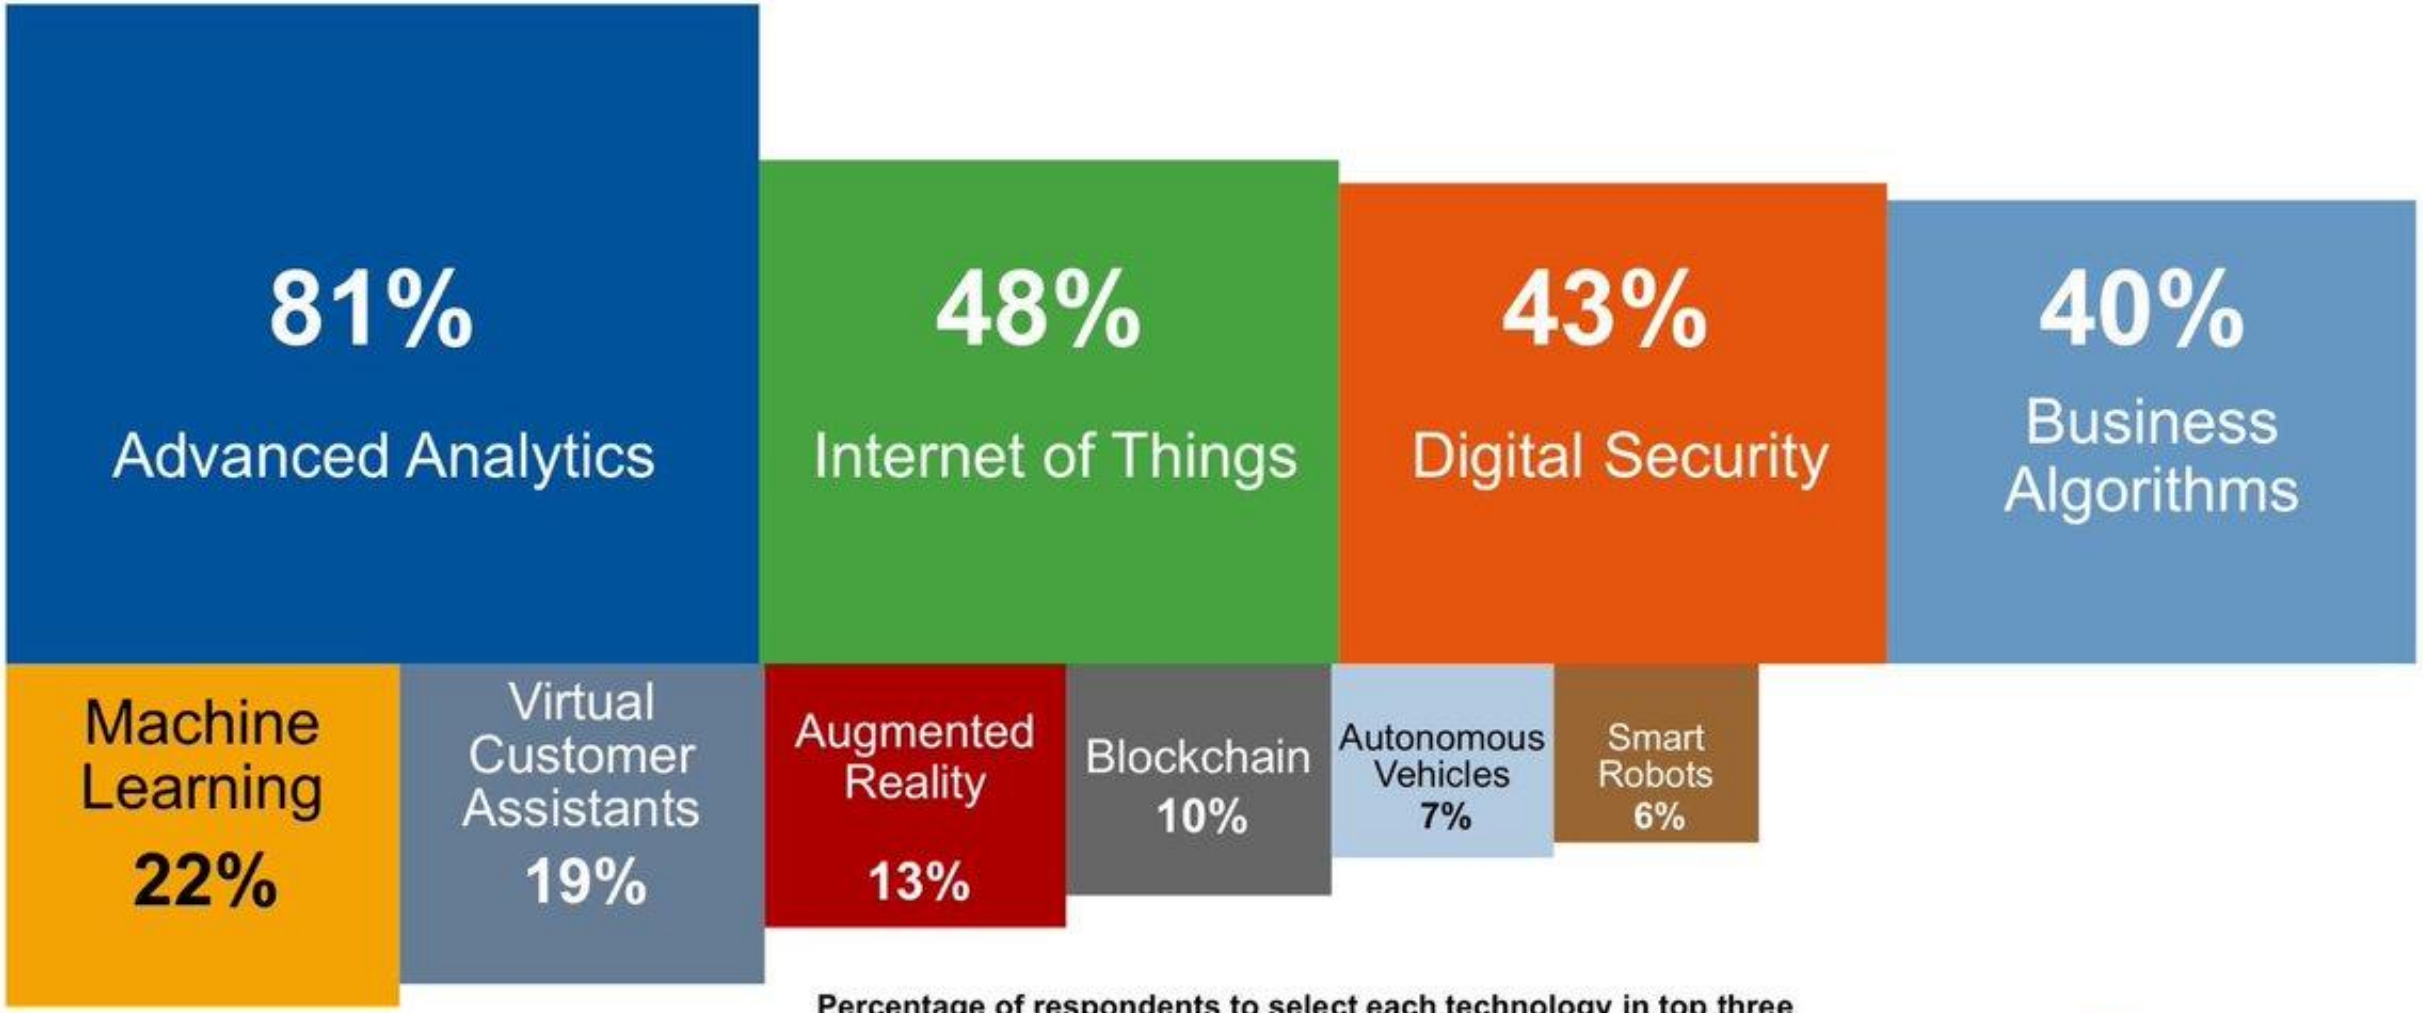

Percentage of respondents to select each technology in top three

#GartnerSYM

Gartner

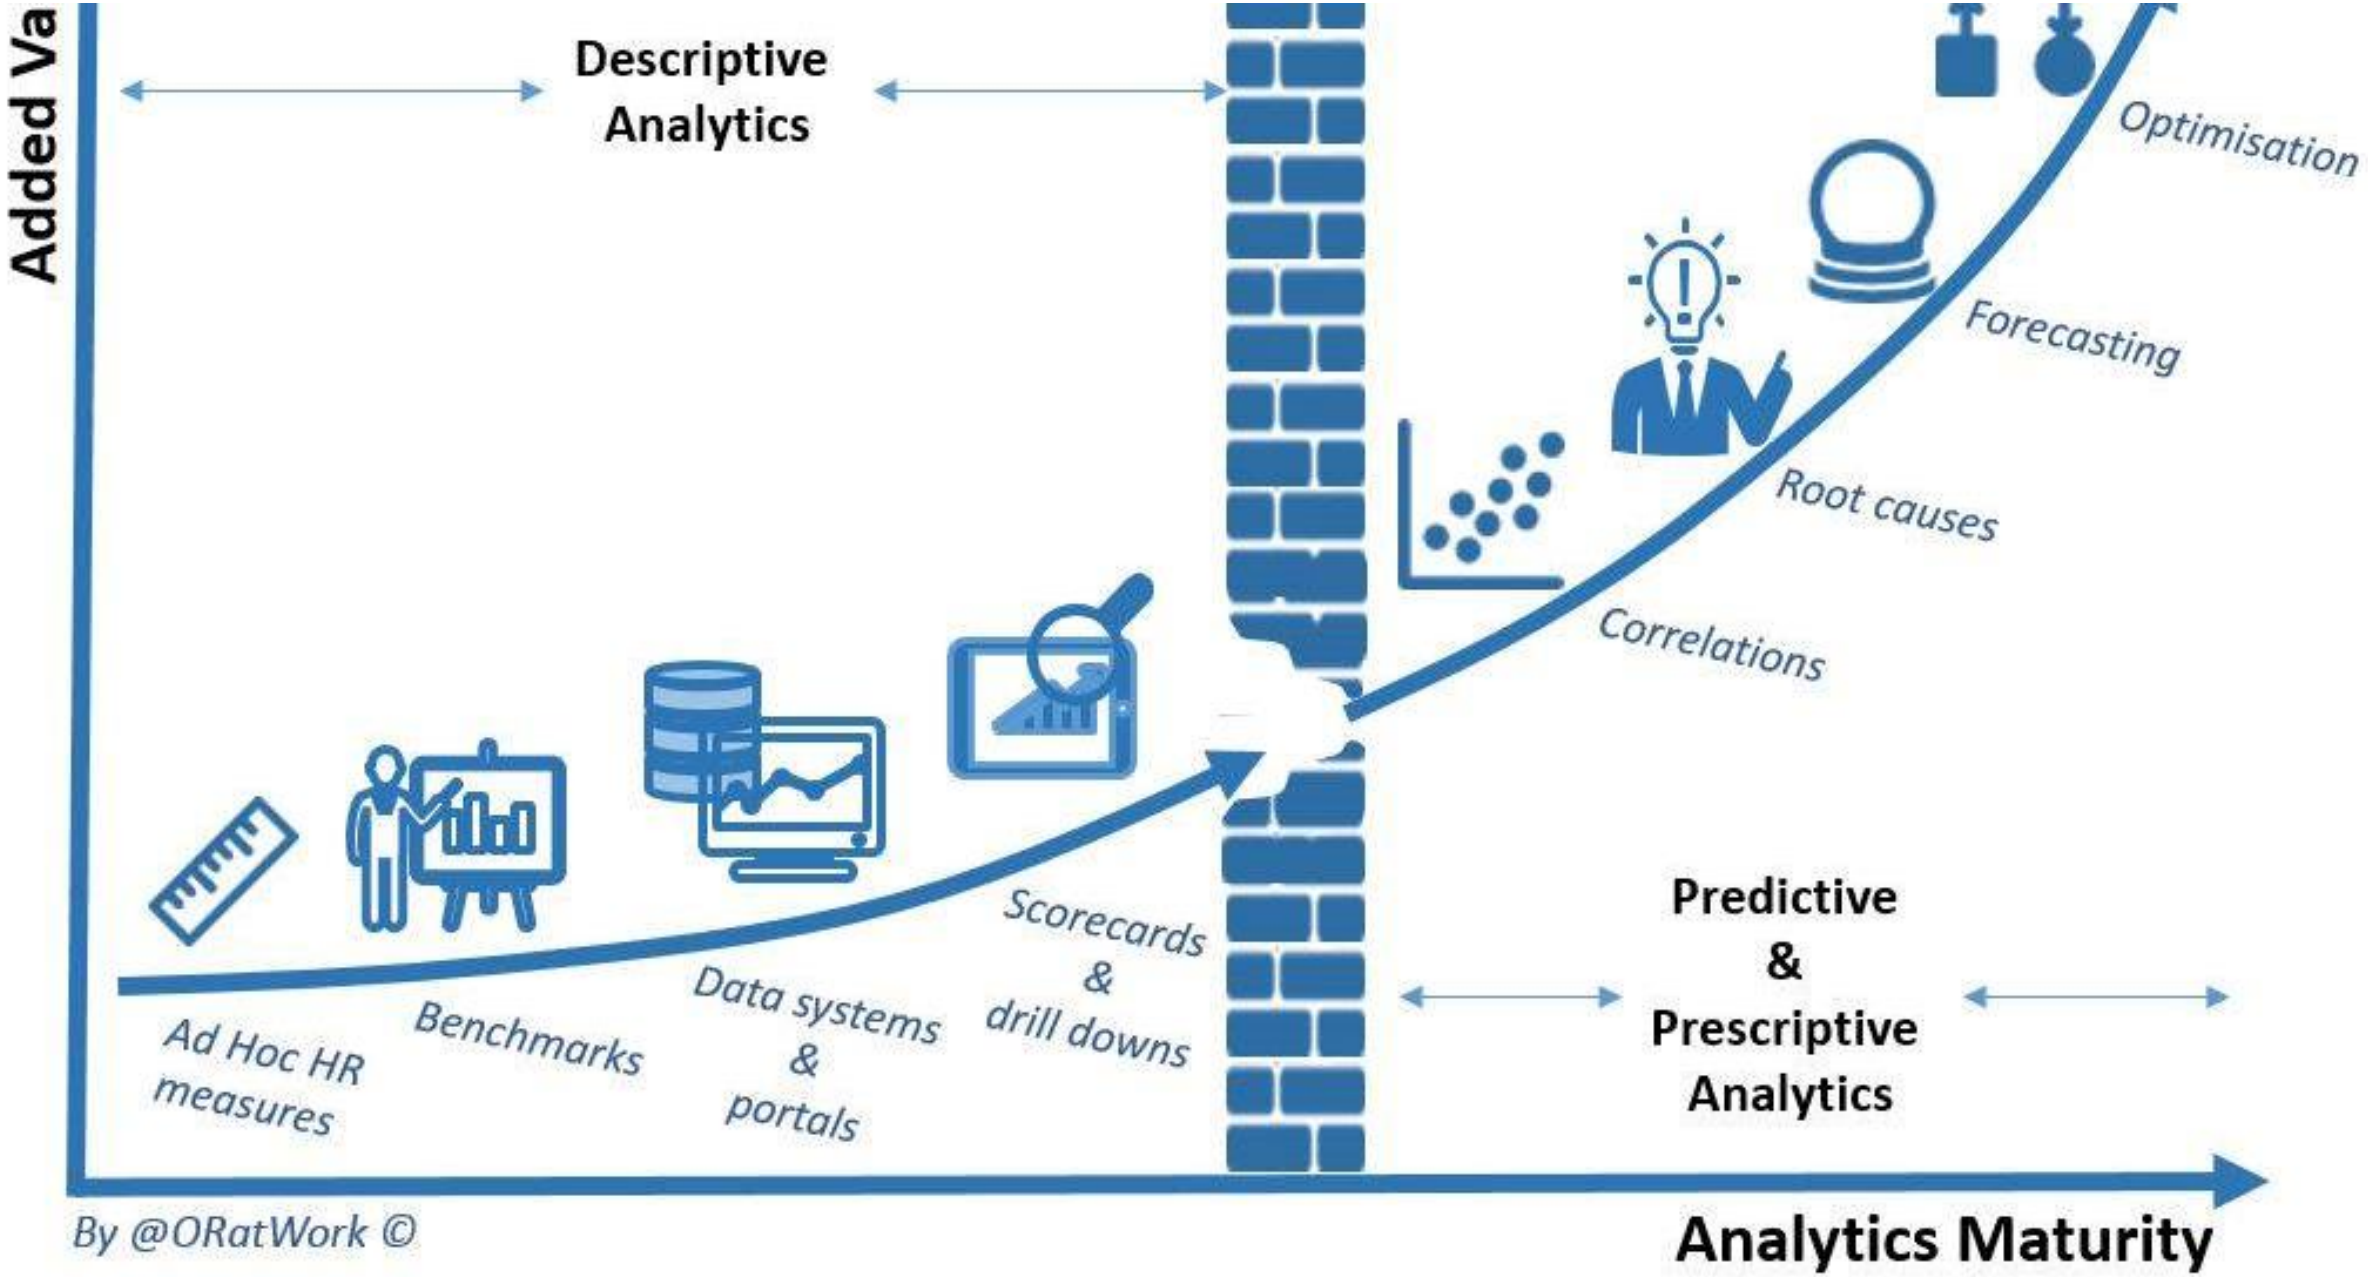
Added Va
Descriptive Analytics
Optimisation
Forecasting
Root causes
Correlations
Ad Hoc HR measures
Benchmarks
Data systems & portals
Scorecards & drill downs
Predictive & Prescriptive Analytics
By @ORatWork ©
Analytics Maturity


26,472
Likes
13,911
Subscribers
6,524
Followers
5,093
Circled By
45,322
Followers
1,765
Followers
Page Load Time (GA)
Redirection Time (GA)
Server Response Time (GA)
Response Time (Pingdom)
Load Time
Avg Response Time
Dec 1, 12:00 PM
Dec 2, 10:00 AM
11.75%
vs $15,069 (prev.)
Session Traffic
280,430
29.44%
vs 216,646 (prev.)
Wistia Video Stats
Metrics Driven Change Management
563
Play Count
75%
Engagement
49%
Play Rate
Salesforce Accounts by Country
Invalid state information provided for 646 of 13124 accounts.
Quickbooks & Salesforce: CAC (Last 30 Days)
$15,085
CAC Value
-29.99%
vs $21,548 prev. 30 days
CallRail Today's Average Call Duration
Company: The Klip Factory
10m:8s
Based on 19 calls
Leads (This Month)
10,753
0
8,961 (Last Month)
12,903
Alexa Page Views per User (Last 30 days)
amazingklips.com
goodklips.com
klipstore.com
20
10
0
Apr 1
Apr 6
Apr 11
Apr 16
Apr 21
Apr 26
Facebook Demographics
Male
Female
Unspecified
People
20,000
15,000
10,000
5,000
0
13-17
18-24
25-34
35-44
45-54
55-64
65+

atlas

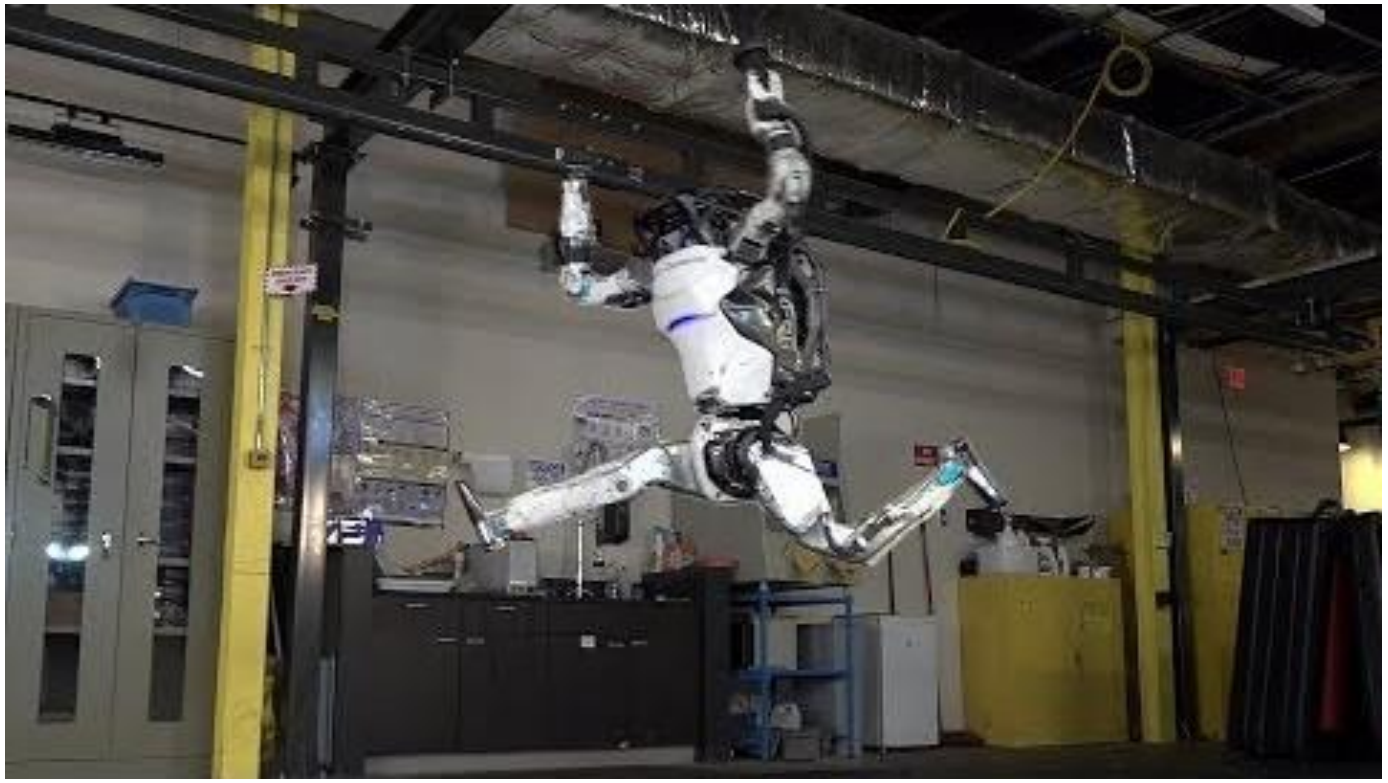


TESLA
STORE &
DELIVERY
CENTER


١:٢١ ص
Zain JO
"Open app store"
tap to edit
It doesn't look like you have an app named 'app store'. If you'd like, I can help you look for it on the App Store.
App Store

In terms of general intelligence, we're not even close to a rat.

Yann LeCun
Facebook AI Director

Sources of Data-Driven Value

REDUCE COSTS

INCREASE REVENUE

MANAGE RISK

Value from Data

Training Predictive Models

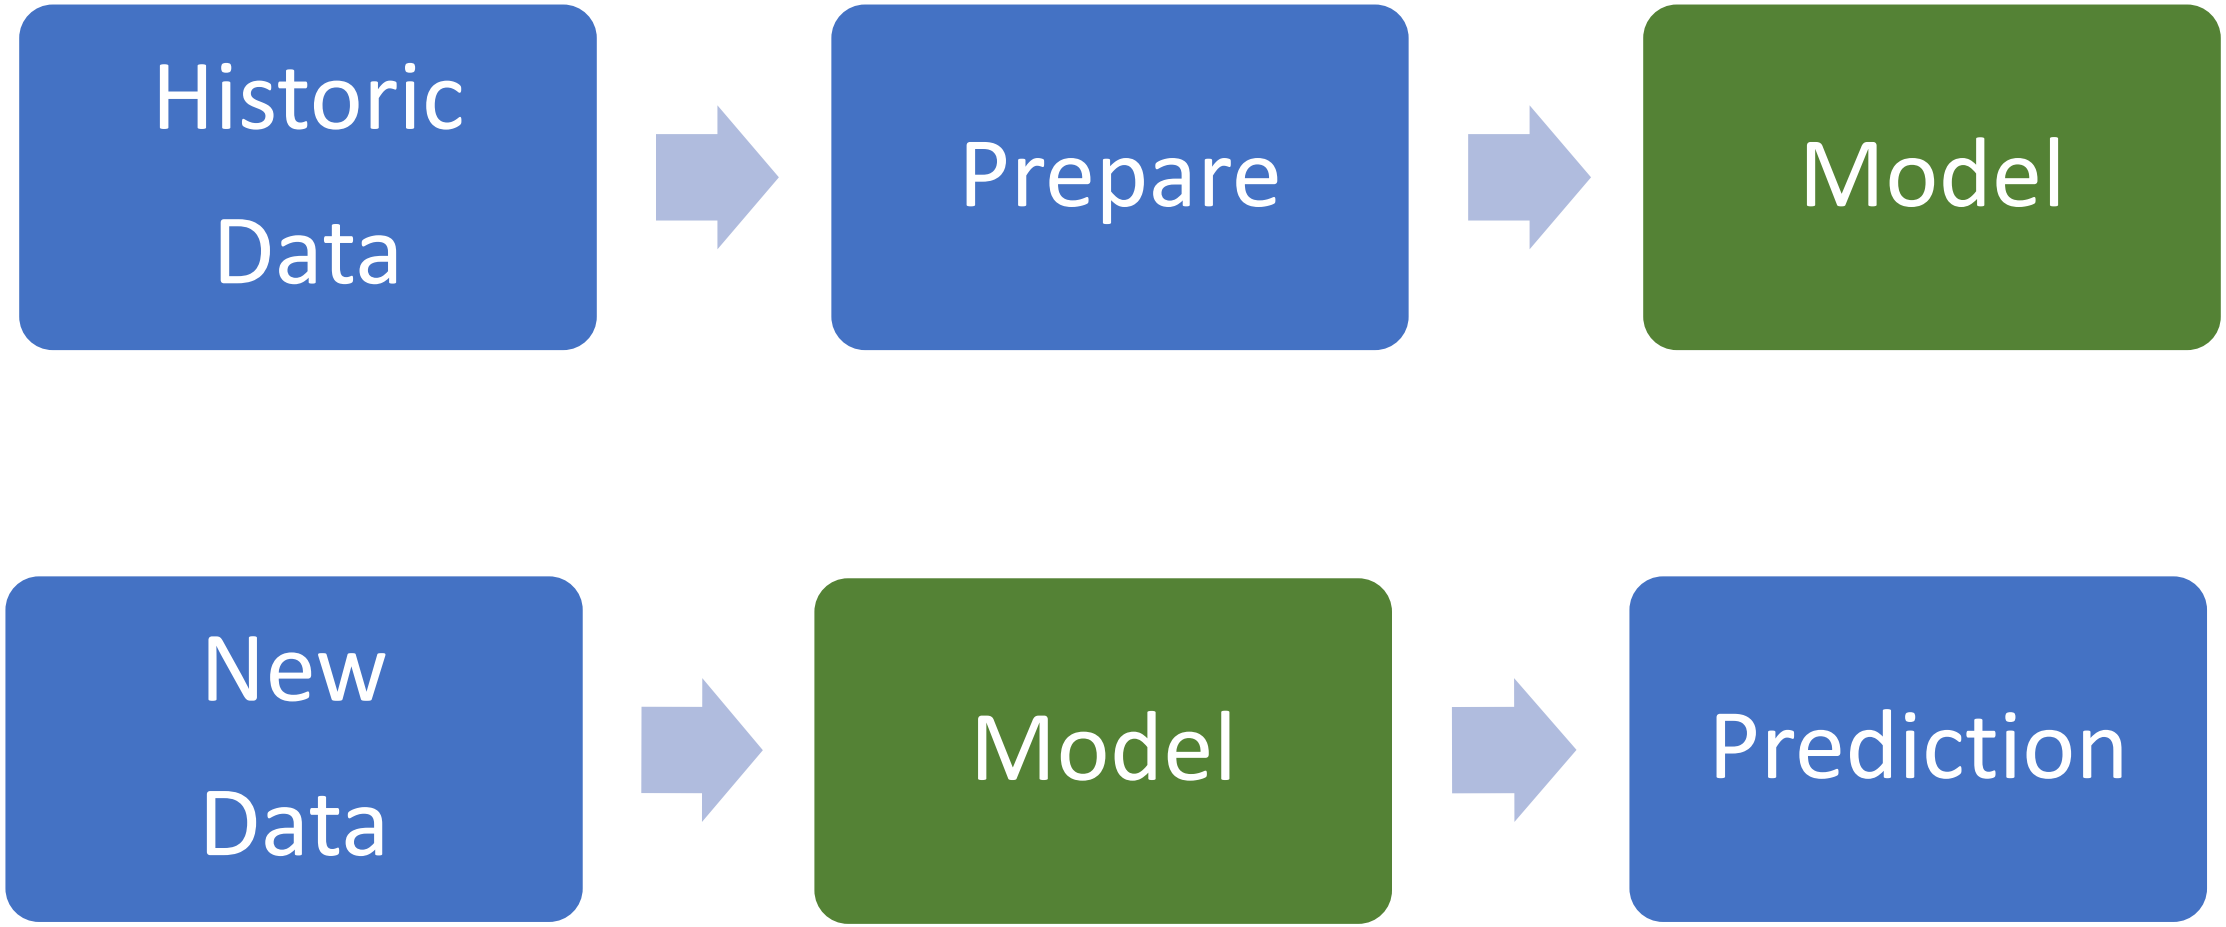

BADIR


1 Business question
2 Analysis plan
3 Data collection
4 Insights
5 Recommendations
Data Science, Technical Track
Decision Science, Business Track

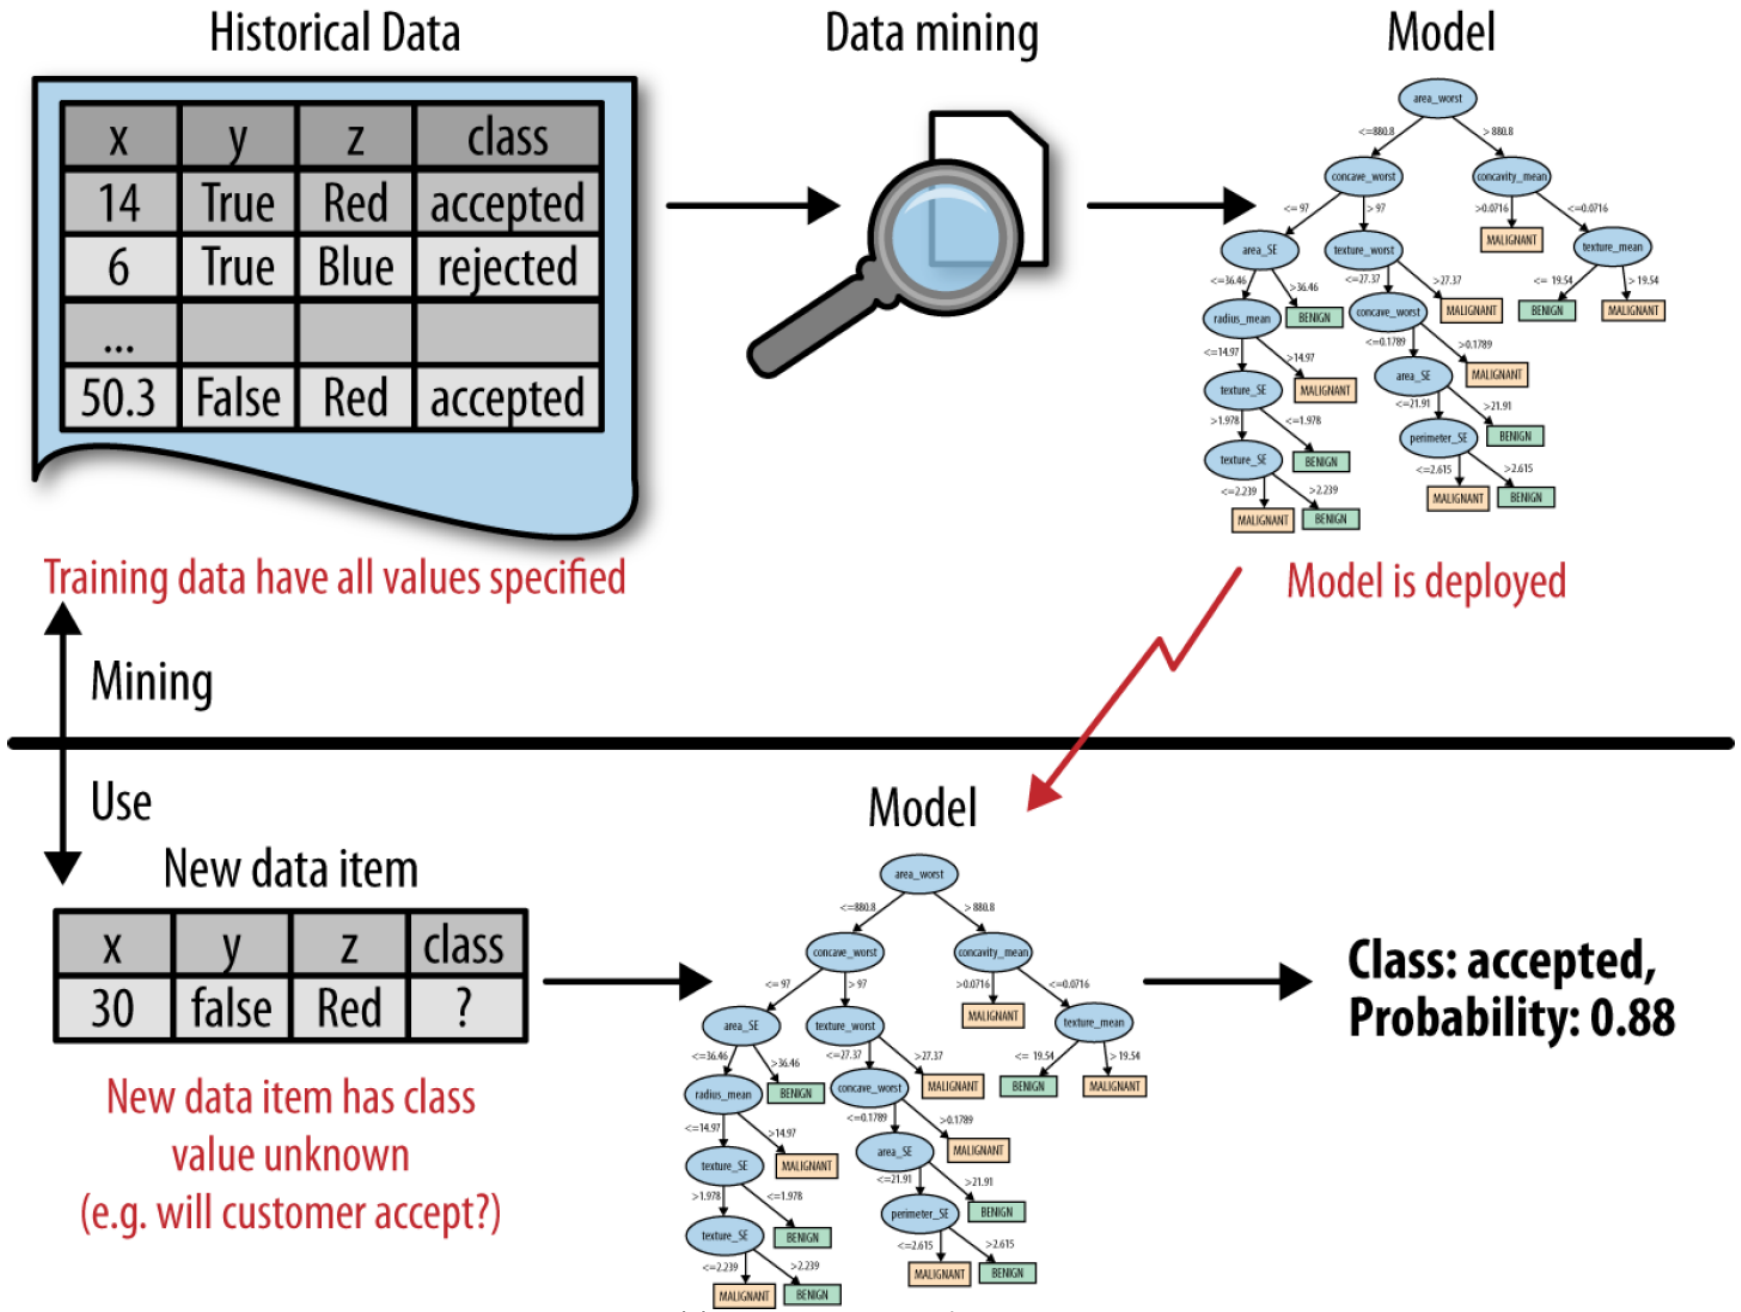
Historical Data
Data mining
Model
x y z class
14 True Red accepted
6 True Blue rejected
... 50.3 False Red accepted
Training data have all values specified
Model is deployed
Mining
Use
New data item
x y z class
30 false Red ?
New data item has class value unknown (e.g. will customer accept?)
Model
Class: accepted, Probability: 0.88

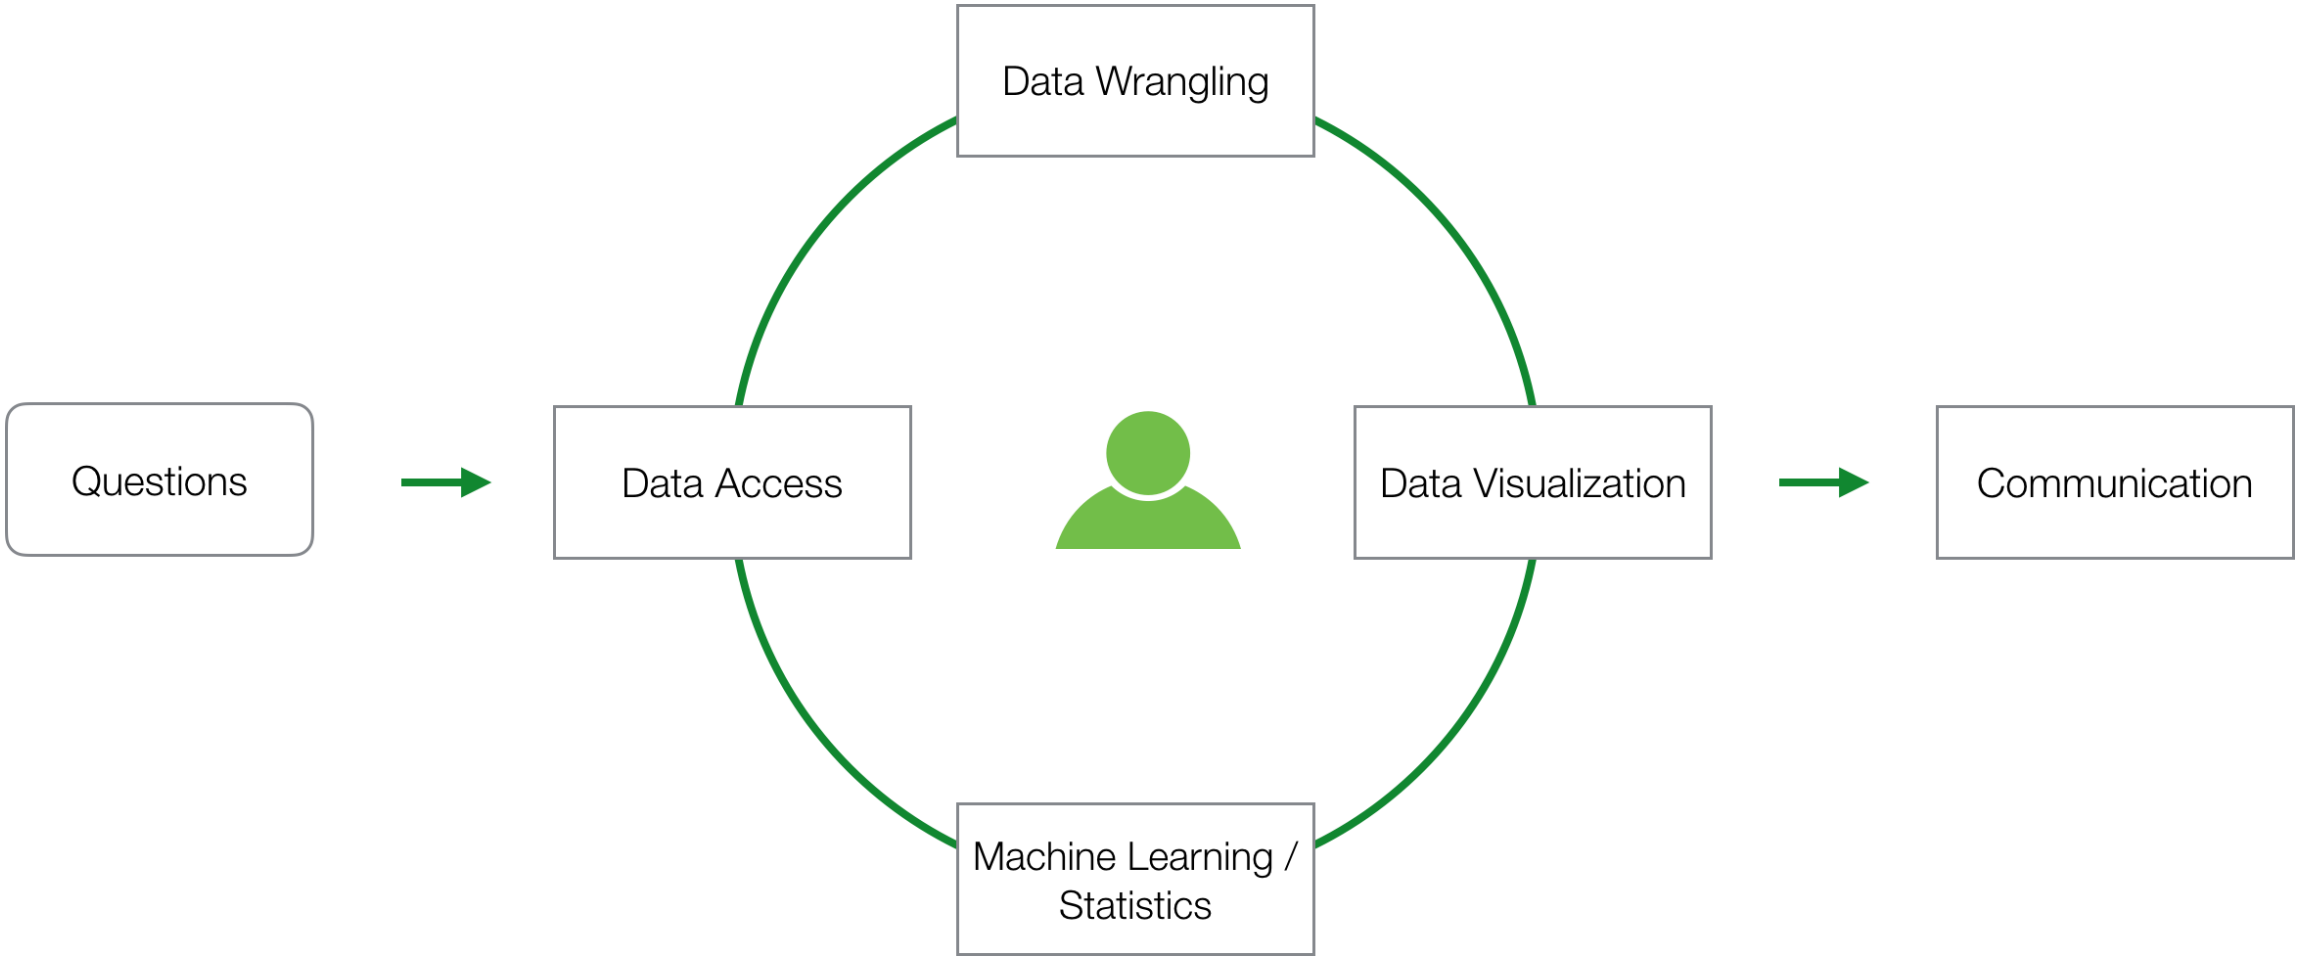
Data Wrangling
Questions
Data Access
Data Visualization
Communication
Machine Learning / Statistics

Merck's Mantis
Data Democratization

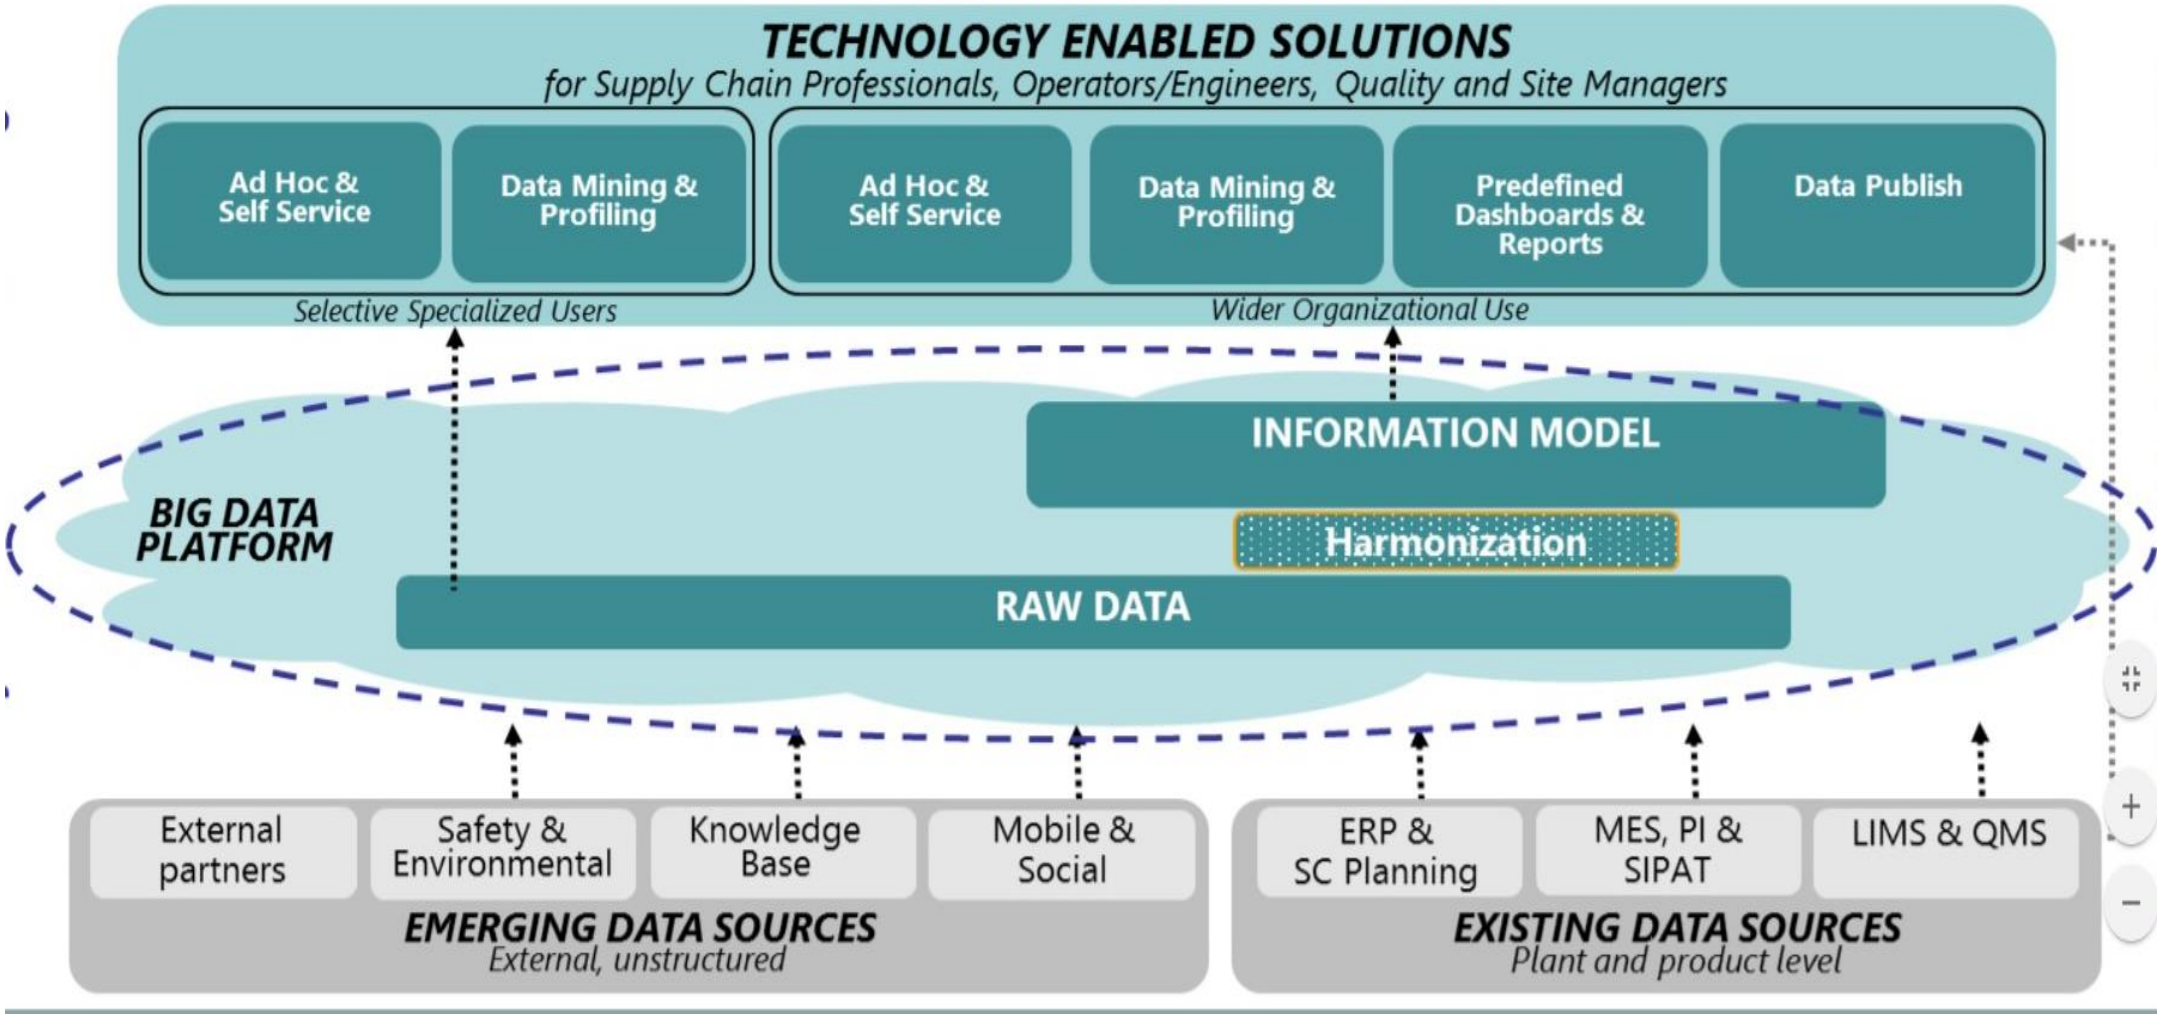

Some ML Applications

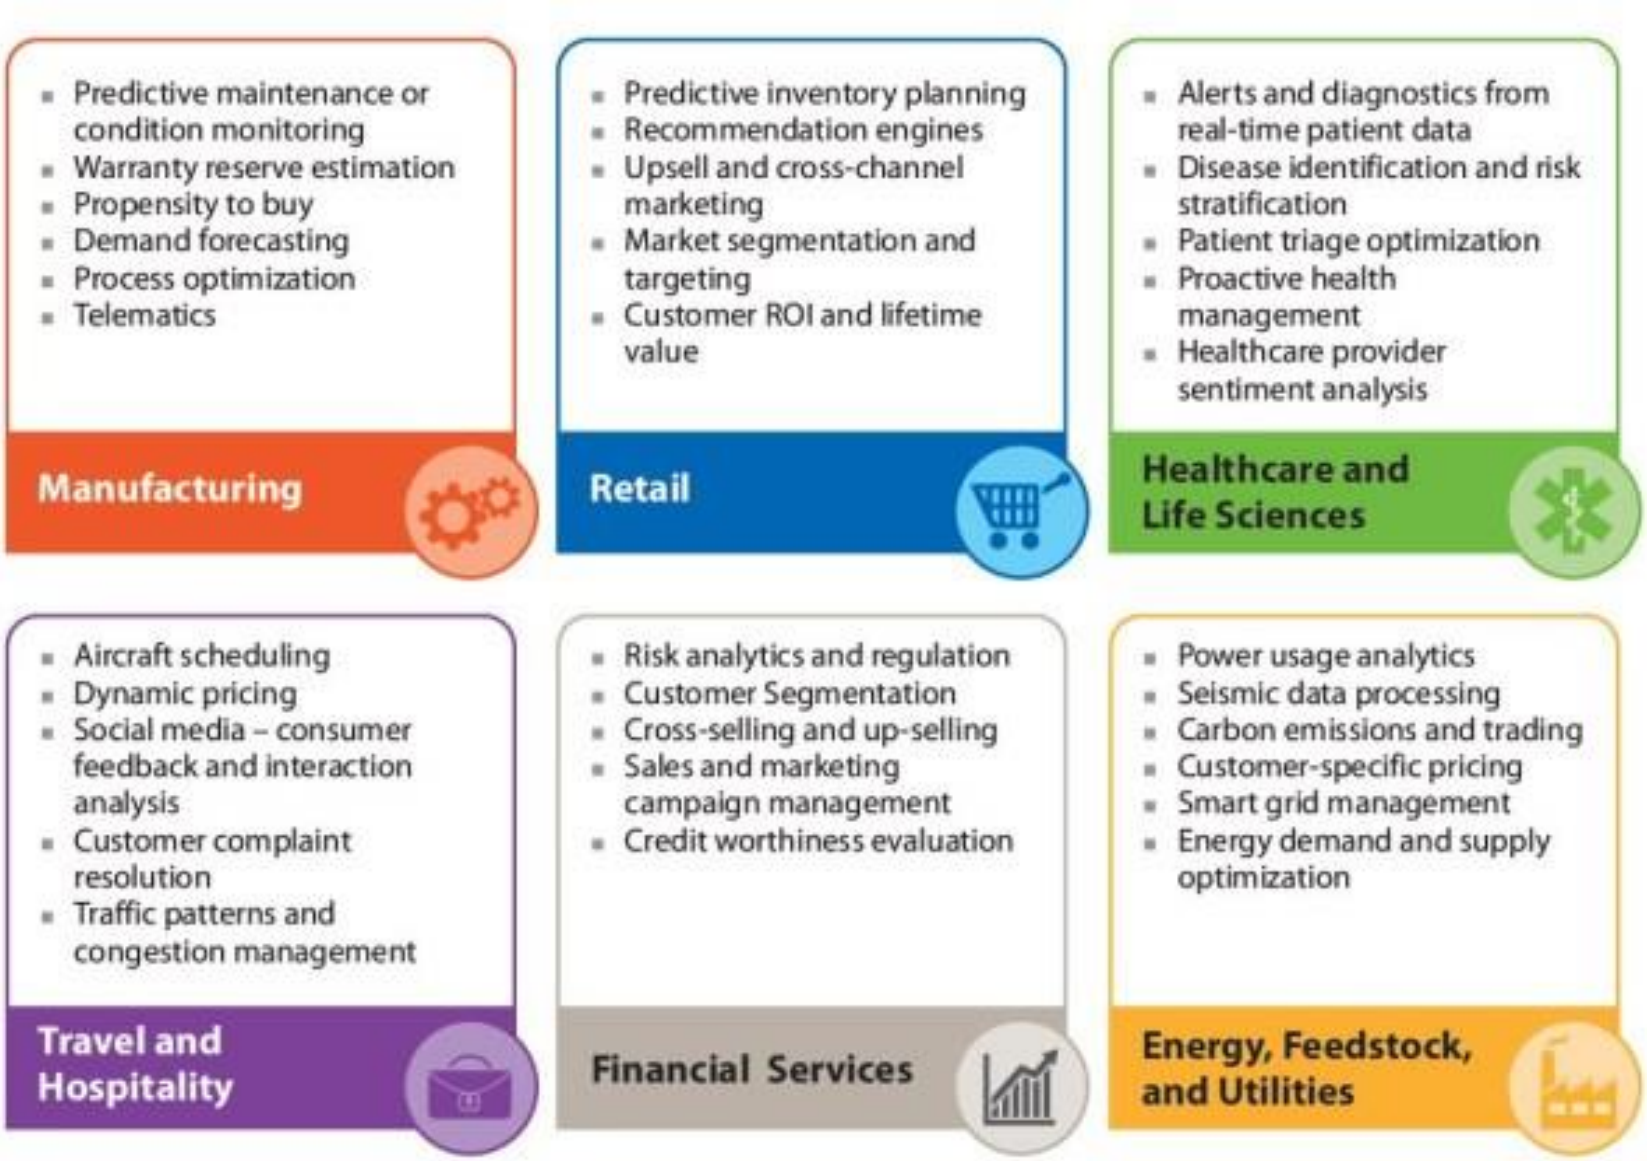

ML in Healthcare

Old World

- Diagnostic care after symptoms arise
- Cancer detected at late stages
- Patients admitted to ICU suffer high mortality rate after sudden deterioration
- Drugs developed without individual considerations

New World

- Real-time preventive care before illness becomes serious
- Patient lives saved due to early ICU admission
- Cancer detected early when still treatable
- Personalized medicine developed for individual needs

ML in Banking

Old World

- Rule-based Anti-Money Laundering (AML) has high false positive rate
- Credit risk scoring relies heavily on past credit history
- Fraud detection mechanisms adapt poorly to ever-changing fraud patterns

New World

- Pattern-based AML improves accuracy and reduces false positives
- Virtual help desk improves customer experience
- Fraudulent behavior is detected faster and with better accuracy
- Product recommendations are personalized and real-time

ML in Insurance

Old World

- Product recommendation from advisor experience
- Risk group classified using static data, e.g. age and medical history
- Labor-intensive claim processes take months to finish
- Customer churn analysis after the fact
- Many loopholes in insurance fraud detection

New World

- Product recommendation from deep analytics of customer data
- Risk group classified using broad spectrum of data, e.g. social media, click streams and web analytics
- Claims processed faster using automated image classification
- Customer churns predicted and intercepted in real time
- Insurance frauds more accurately detected from customer 360° analysis
- Increased revenue

ML in Marketing

Old World

- Product recommendation from experience
- Customer churn analysis after the fact
- Decide on Marketing campaigns based on experience
- Marketing resources allocated based on needs of the past

New World

- Product recommendation from deep analytics of purchase behavior
- Customer churns predicted and intercepted in real time
- Increased revenue
- Marketing resources allocated based on needs of the future

ML in Telecom

Old World

- Reactive Maintenance
- Network optimization with human intervention
- Centralized intelligence
- Security attack repair
- Backlogged customer tickets

New World

- Predictive Maintenance
- Self-optimizing network
- Optimal network quality
- Intelligence at the edge
- Security attack prediction
- Improved customer experience

The Analytics Process Model

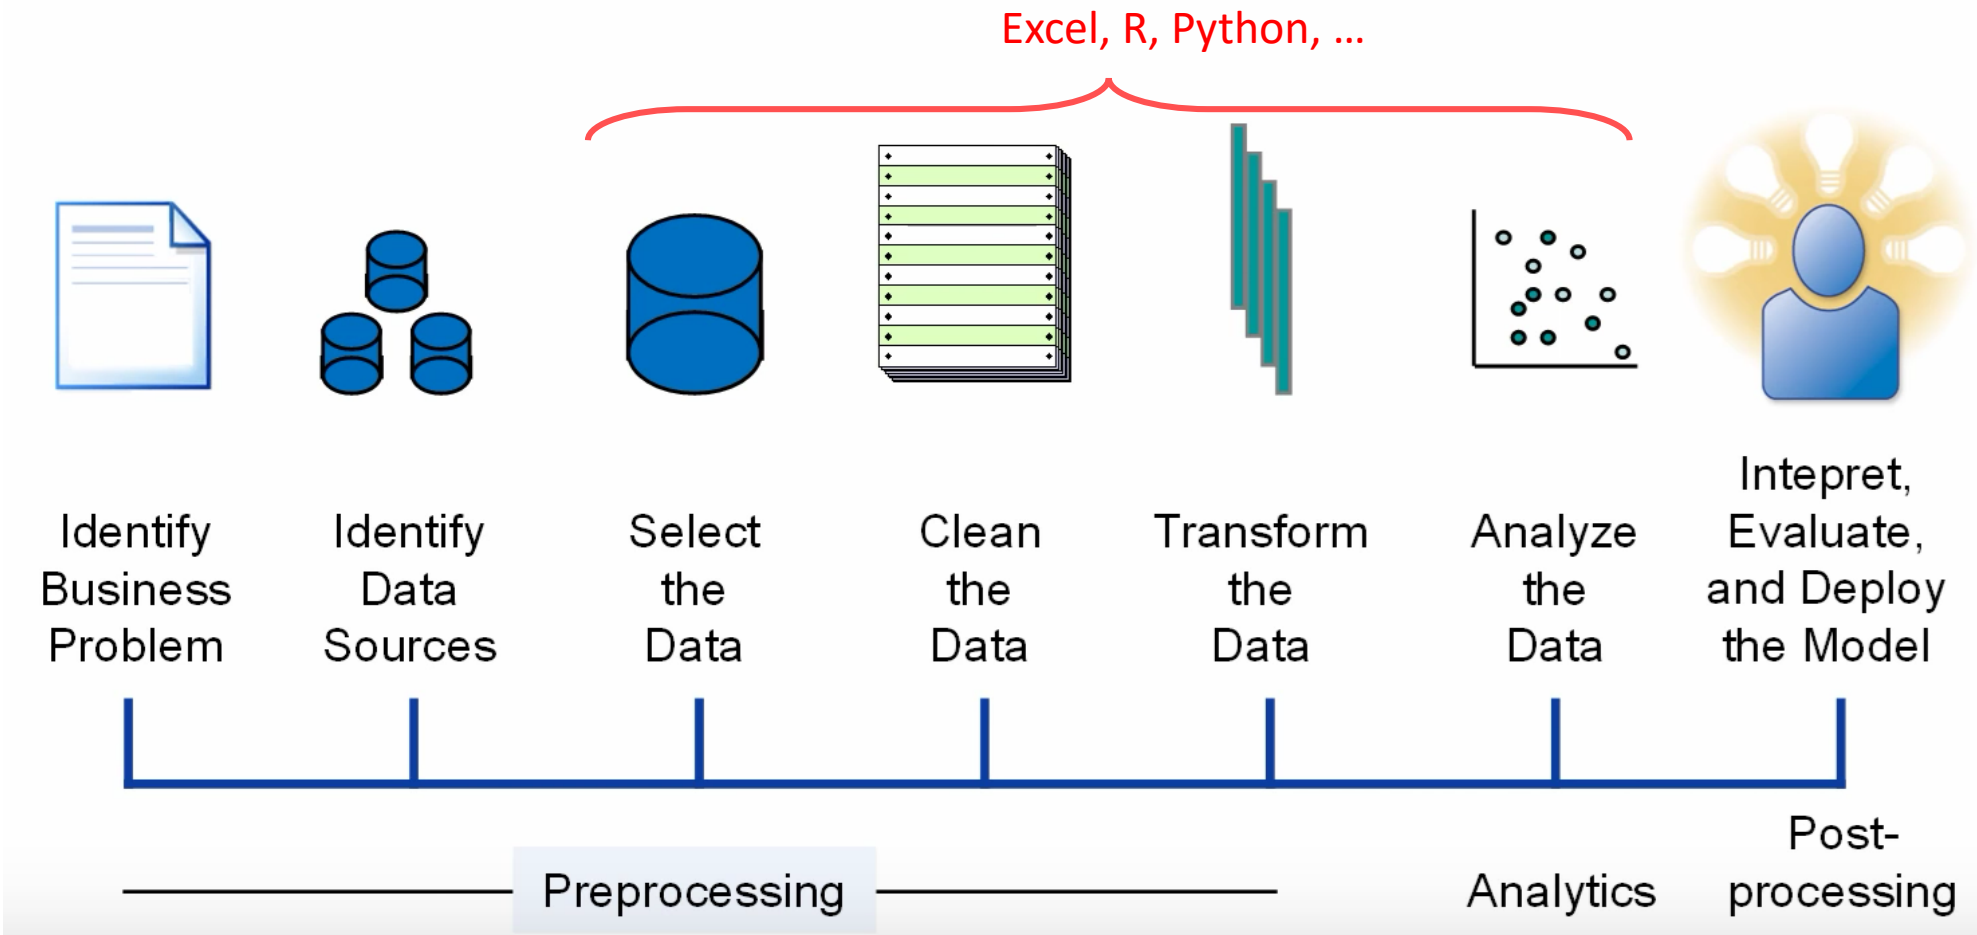

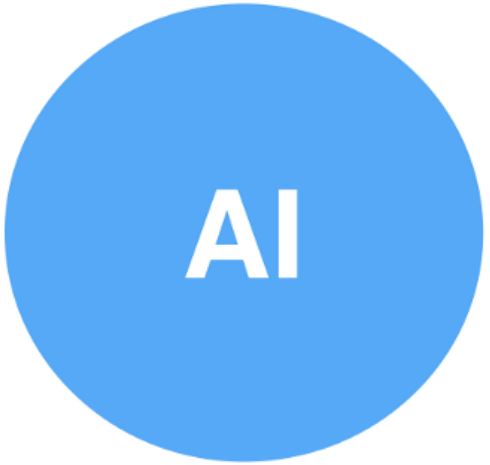

- sounds sexy
- gets us money from VCs
- what we all hope is the future

- the only real “AI”
- traditionally an academic discipline
- not concerned with real-world software

- applies machine learning to create actual products
- deals with real-world complexity

What is Data Science?

- Tools & Methods to produce business value
  - Ingest
  - Understand/Visualize
  - Wrangle/Prepare
  - Model
  - Communicate/Publish
- Support/automate business processes via:
  - Revenue Increase
  - Cost Reduction
  - Risk Management
- Deliver:
  - Dashboards/reports
  - Predictive models/APIs
  - New products

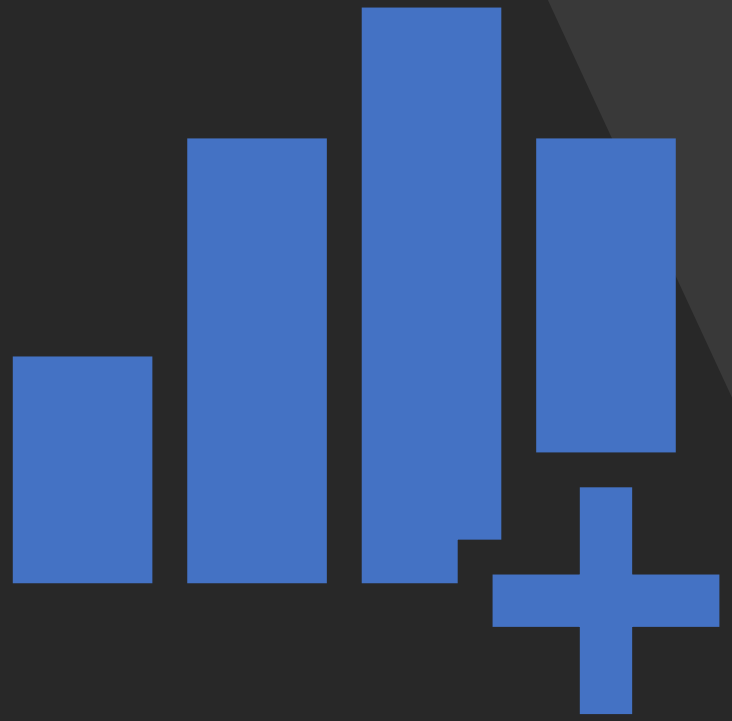

+people
+visualization
+communication

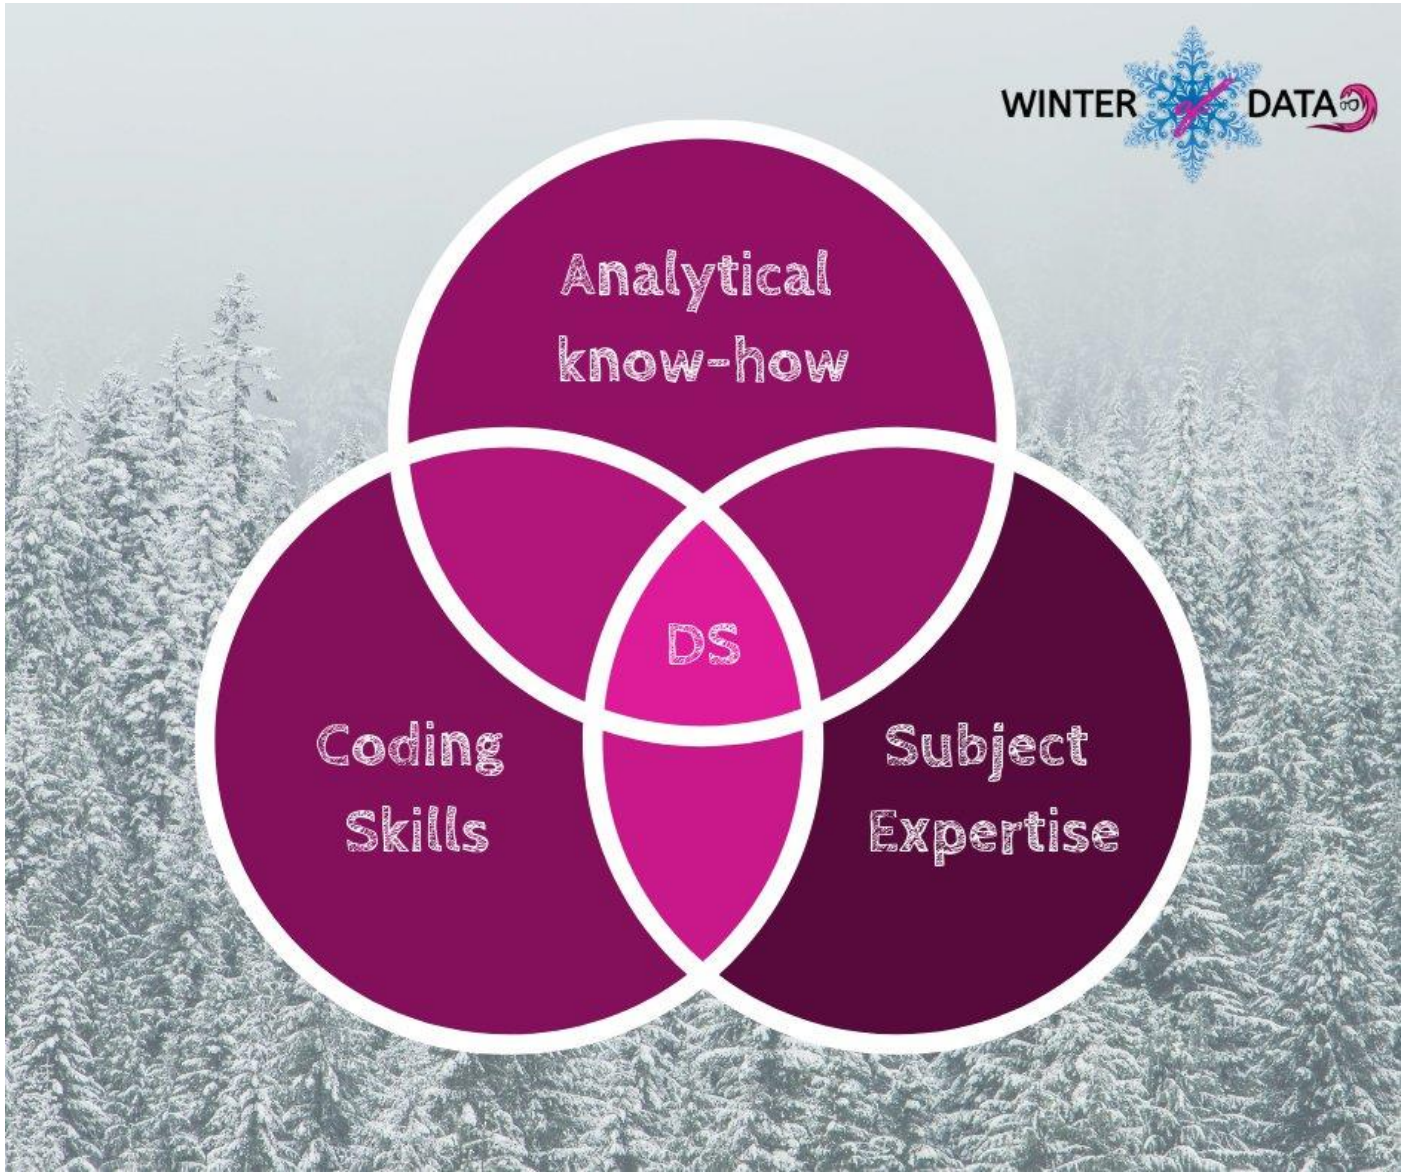


Author: Swami Chandrasekaran
1. Fundamentals
Matrices & Linear Algebra Fundamentals
Hash Functions, Binary Tree, O(n)
Relational Algebra, DB Basics
Inner, Outer, Cross, Theta Join
CAP Theorem
Tabular Data
Data Frames & Series
Entropy
Sharding
OLAP
Multidimensional Data Model
ETL
Reporting Vs BI Vs Analytics
JSON & XML
NoSQL
Regex
Vendor Landscape
Env Setup
5%
2. Statistics
Prob Den Fn (PDF)
ANOVA
Skewness
Continuos Distributions (Normal, Poisson, Gaussian)
Cumul Dist Fn (CDF)
Random Variables
Bayes Theorem
Probability Theory
Percentiles & Outliers
Histograms
Exploratory Data Analysis
Central Limit Theorem
Monte Carlo Method
Hypothesis Testing
p-Value
Chi² Test
Estimation
Confid Int (CI)
MLE
Kernel Density Estimate
Regression
Covariance
Correlation
15%
Factor Analysis
Install Pkgs
Descriptive Statistics (mean, median, range, SD, Var)
Pick a Dataset (UCI Repo)
3. Programming
Data Frames
Reading CSV Data
Reading Raw Data
Subsetting Data
Manipulate Data Frames
Functions
Lists
Factors
Arrays
Matrices
Vectors
Variables
Expressions
R Basics
R Setup R Studio
Rapid Miner
IBM SPSS
Python Basics
Working w/ Excel
4. Machine Learning
Regression
Perceptron
Linear Regression
Ranking
Logistic Regression
Classification
K-Nearest Neighbor
Naive Bayes Classifiers
Boosting
Decision Trees
Classification Rate
Trees & Classification
Bias & Variance
Overfitting
Lift
Prediction
Classifier
Training & Test Data
Concepts, Inputs & Attributes
Unsupervised Learning
Supervised Learning
Categorical Var
Numerical Var
What is ML?
30%
Euclidean Distance
Least² Fit
Causation
Pearson Coeff
Correlation
Clustering
Hierarchical Clustering
K-means Clustering
Neural Networks
Sentiment Analysis
Collaborative Filtering
Tagging
Vocabulary Mapping
Classify Text
50%
5. Text Mining / NLP
Corpus
Text Analysis
UIMA
Term Document Matrix
Term Document Matrix
Term Frequency & Weight
Support Vector Machines
Association Rules
Market Based Analysis
Feature Extraction
Using Mahout
Using Weka
Using NLTK
8. Data Ingestion
Data Survey
Using ETL
How much Data?
Google OpenRefine
80%
Sampling
Stratified Sampling
Principal Component Analysis
Transformation & Enrichment
Data Fusion
Data Integration
Data Sources & Acquisition
Data Discovery
Summary of Data Formats
9. Data Munging
Denoising
Feature Extraction
Binning Sparse Values
Unbiased Estimators
Handling Missing Values
Data Scrubbing
Normalization
Dimensionality & Numerosity Reduction
6. Visualization
Data Exploration in R (Hist, Boxplot etc)
Uni, Bi & Multivariate Viz
ggplot2
Histogram & Pie (Uni)
Tree & Tree Map
Scatter Plot (Bi)
Line Charts (Bi)
Spatial Charts
Survey Plot
Timeline
Decision Tree
40%
Tableau
IBM ManyEyes
InfoVis
D3.js
7. Big Data
Job & Task Tracker
MR Programming
Sqoop: Loading Data in HDFS
Flume, Scribe : For Unstruct Data
SQL with Pig
DWH with Hive
Scribe, Chukwa For Weblog
Using Mahout
Name & Data Nodes
Setup Hadoop (IBM / Cloudera / HortonWorks)
Data Replication Principles
HDFS
Hadoop Components
Map Reduce Fundamentals
Zookeeper Avro
Storm: Hadoop Realtime
Rhadoop, RHIPE
rmr
Cassandra
MongoDB, Neo4j
60%
10. Toolbox
MS Excel w/ Analysis ToolPak
Java, Python
R, R-Studio, Rattle
Weka, Knime, RapidMiner
Hadoop Dist of Choice
Spark, Storm
Flume, Scibe, Chukwa
Nutch, Talend, Scraperwiki
Webscraper, Flume, Sqoop
tm, RWeka, NLTK
RHIPE
D3.js, ggplot2, Shiny
IBM Languageware
Cassandra, MongoDB
100%

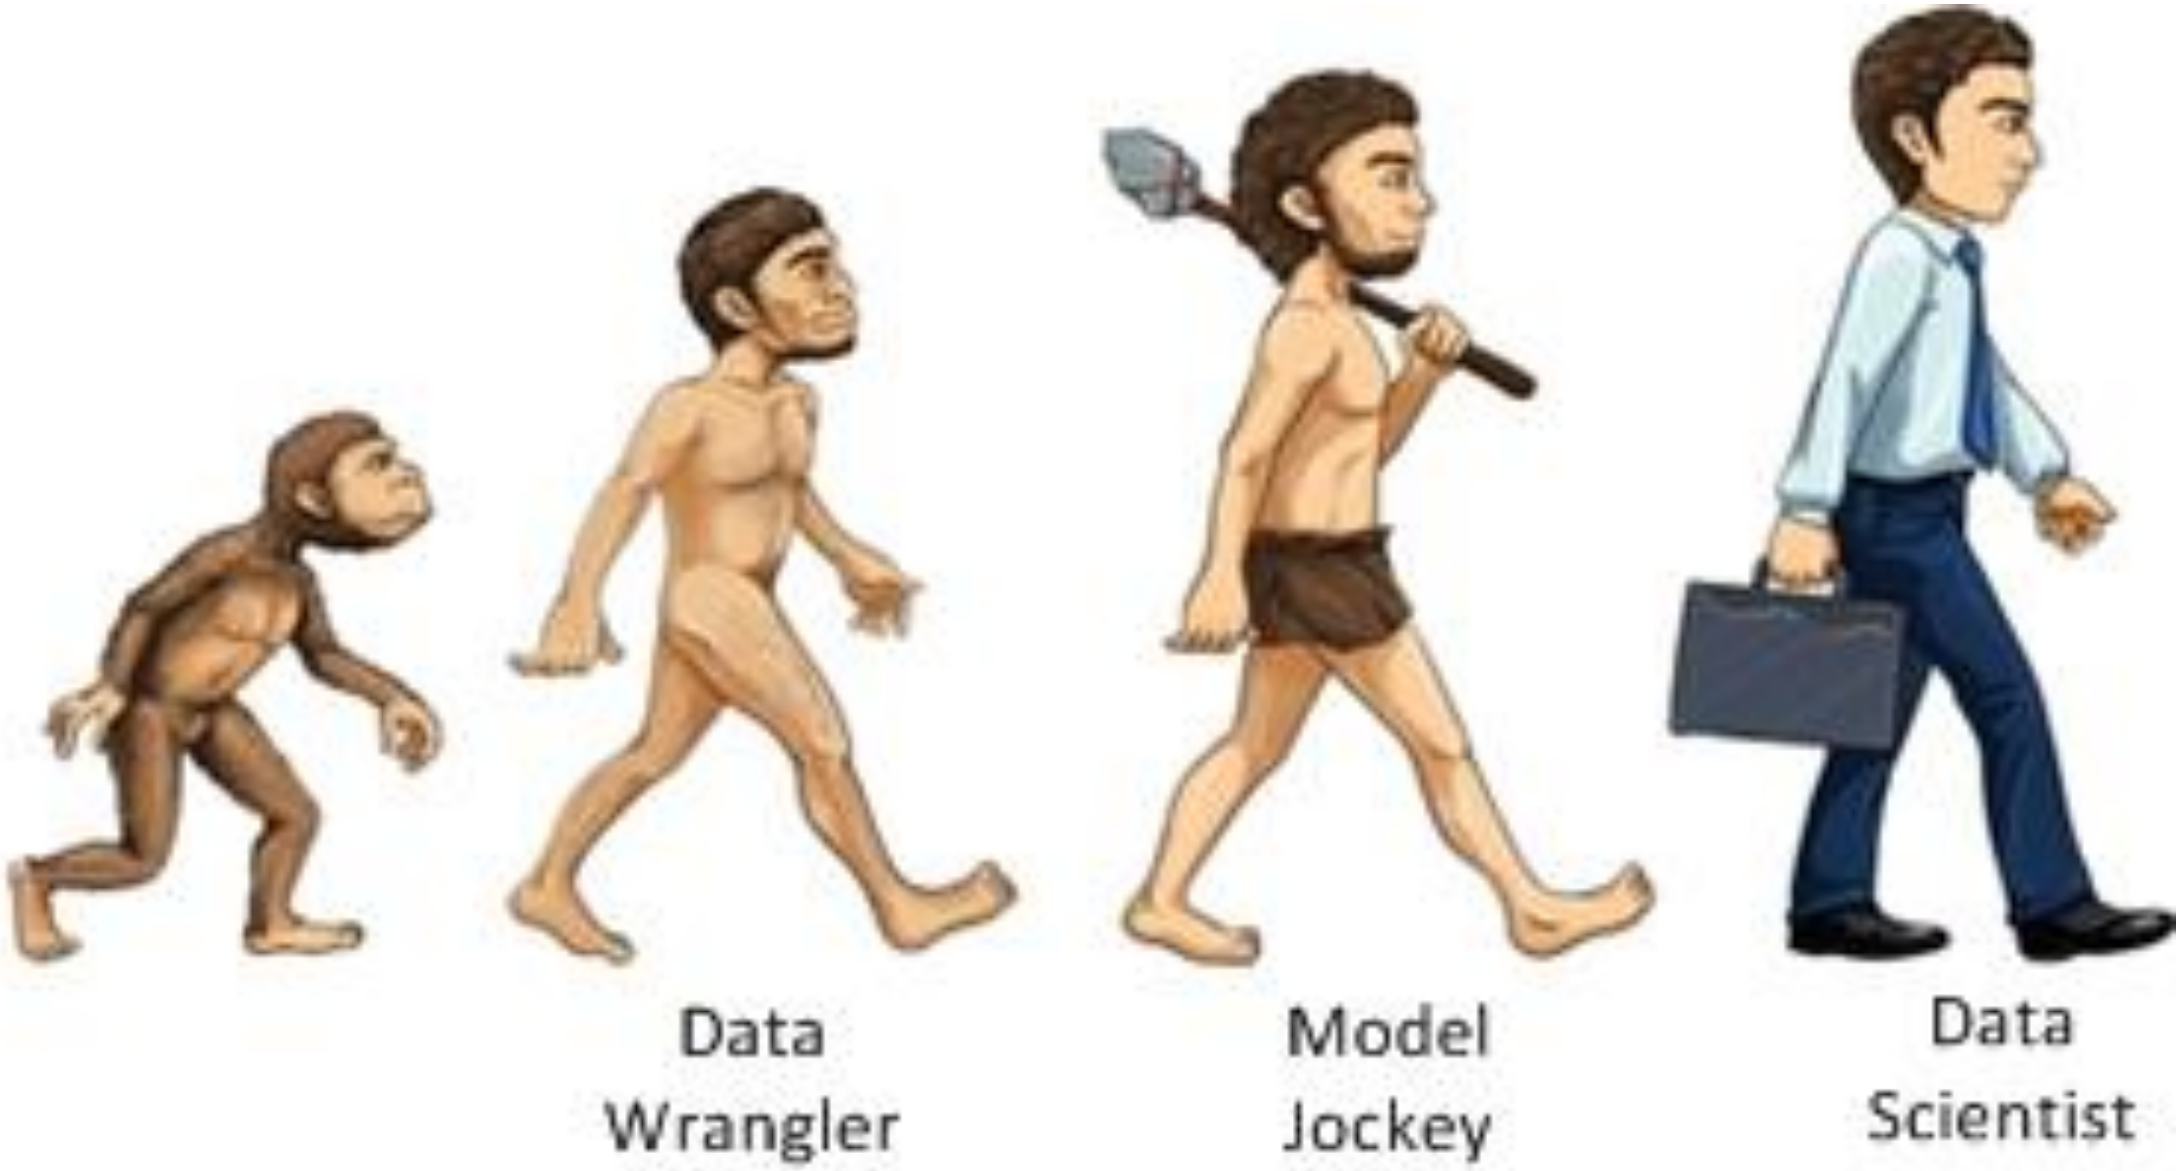
Data
Wrangler
Model
Jockey
Data
Scientist

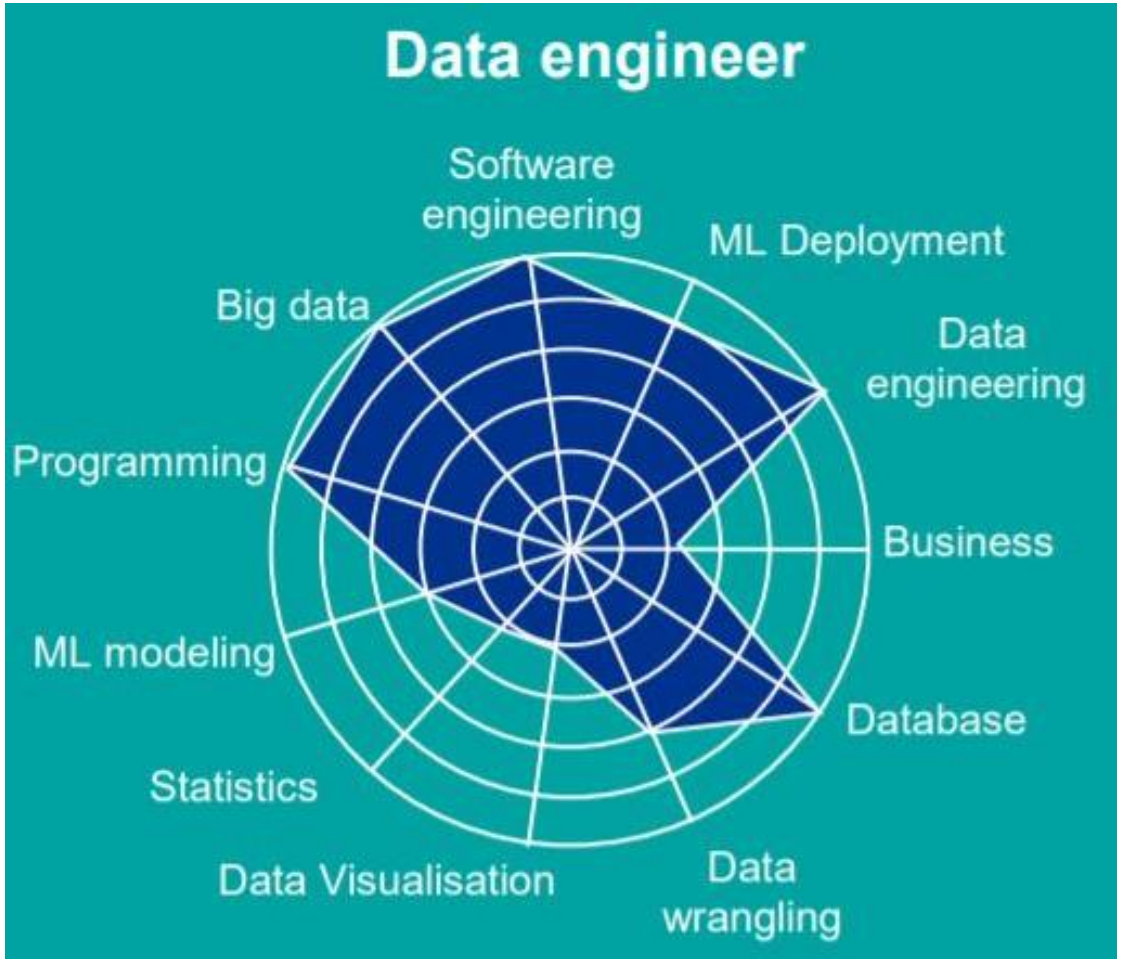
Data engineer
Software engineering
ML Deployment
Data engineering
Business
Database
Data wrangling
Data Visualisation
Statistics
ML modeling
Programming
Big data

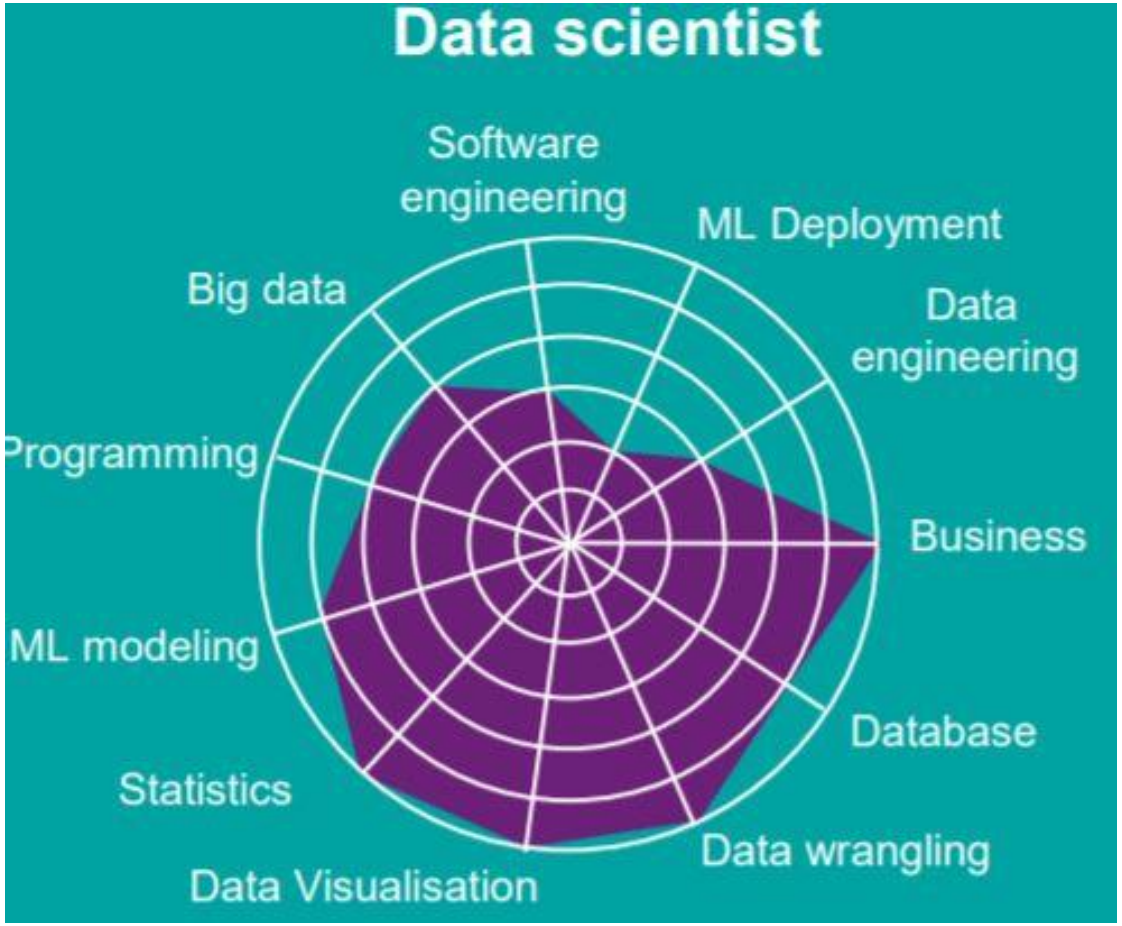
Data scientist
Software engineering
ML Deployment
Data engineering
Business
Database
Data wrangling
Data Visualisation
Statistics
ML modeling
Programming
Big data

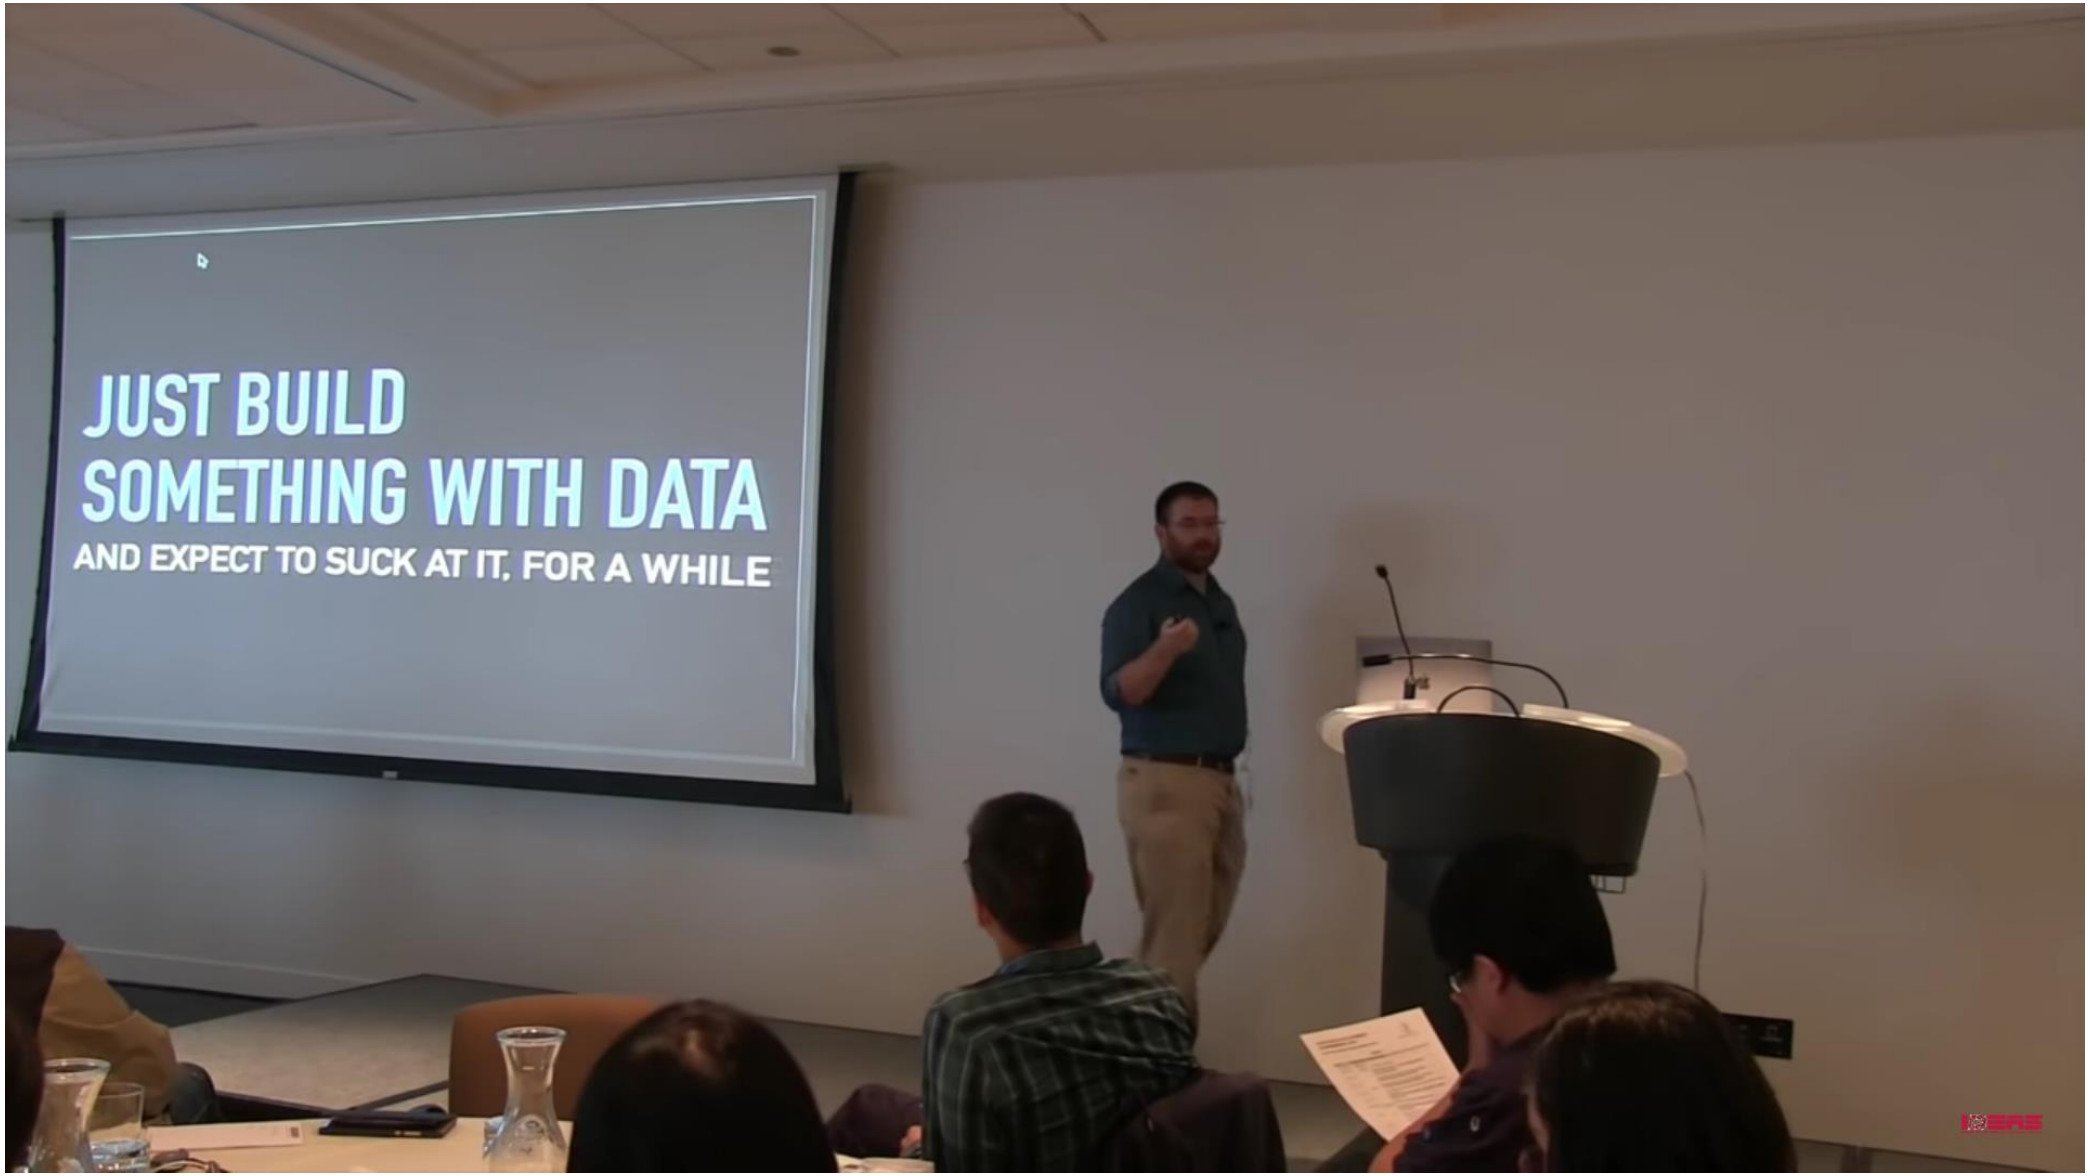
JUST BUILD
SOMETHING WITH DATA
AND EXPECT TO SUCK AT IT, FOR A WHILE

Challenges that Data Professionals have Faced in the Past Year

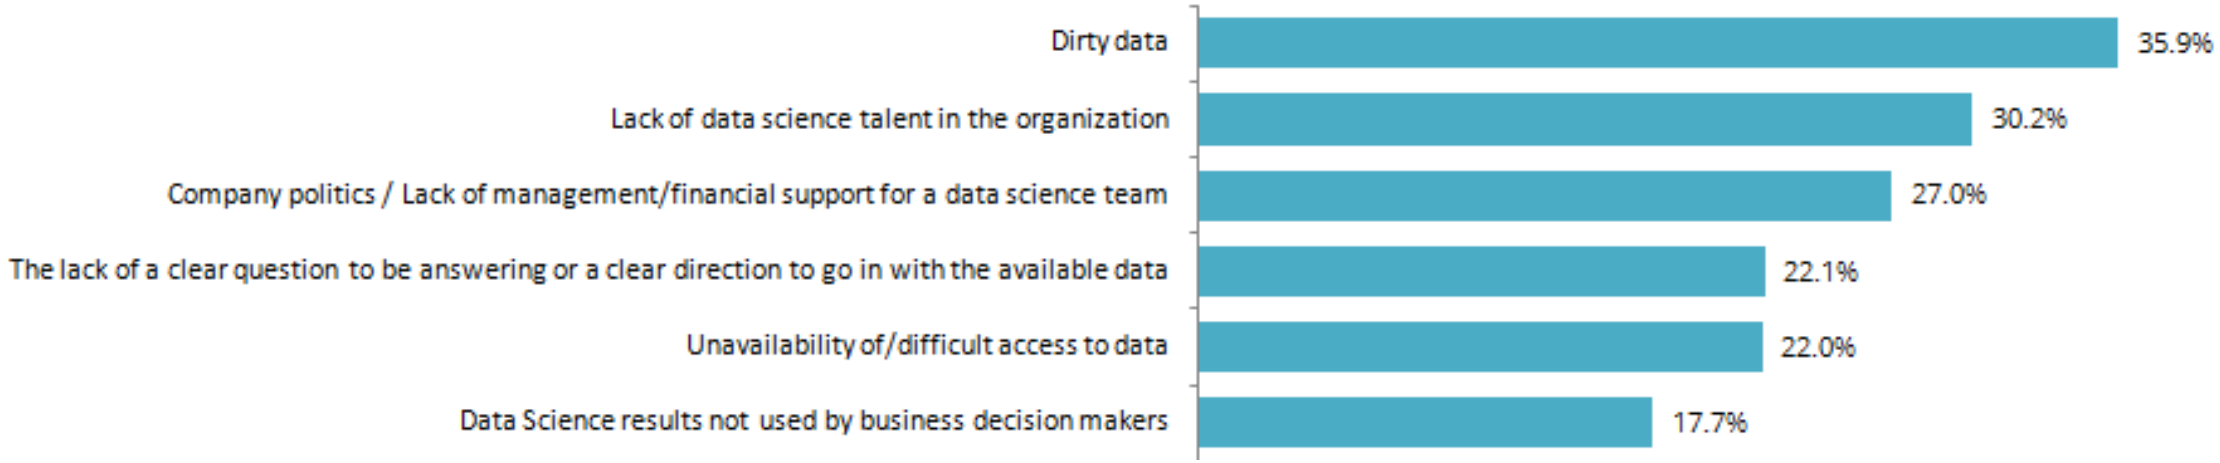

Baron Schwartz 

@xaprb

When you’re fundraising, it’s AI
When you’re hiring, it’s ML
When you’re implementing, it’s linear regression
When you’re debugging, it’s printf()

12:52 AM - Nov 15, 2017

 73  4,550  10,284

R Ecosystem

HISTORY AND EVOLUTION OF R

R has developed from the S language

S Version 1

S Version 2

S Version 3

S Version 4

Developed 30 years ago for research

applied to the high-tech industry

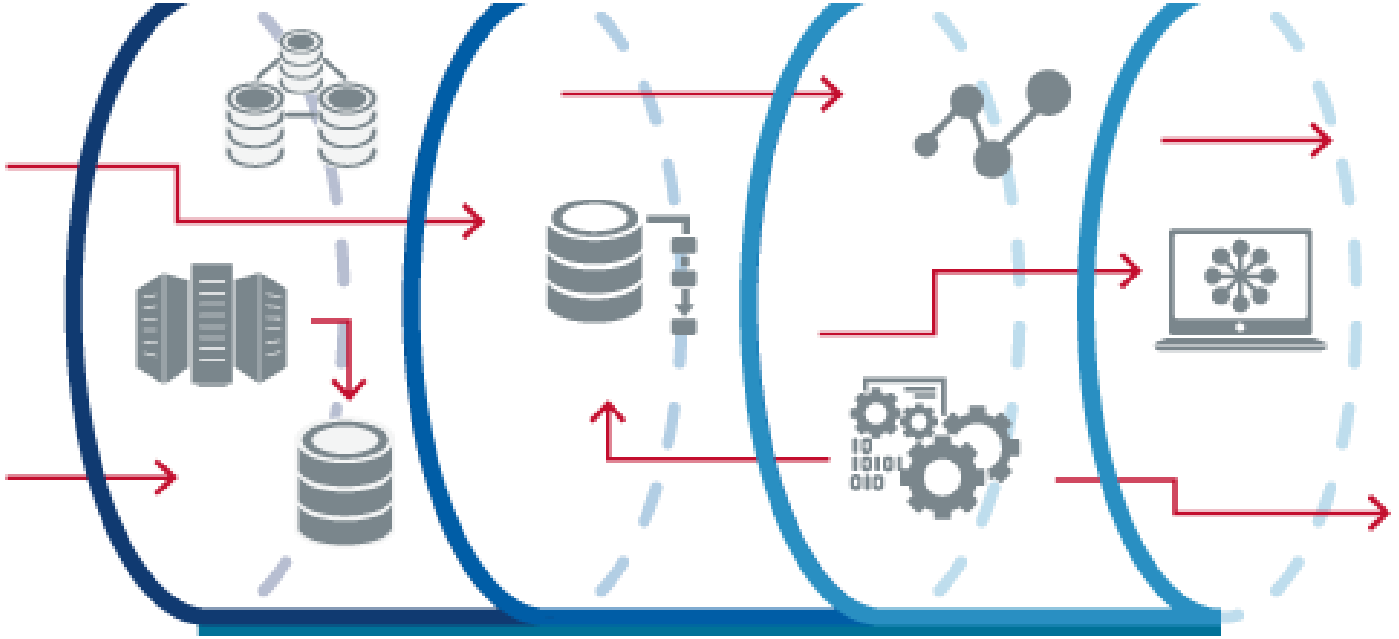

R ecosystem = Analytics Meat Grinder

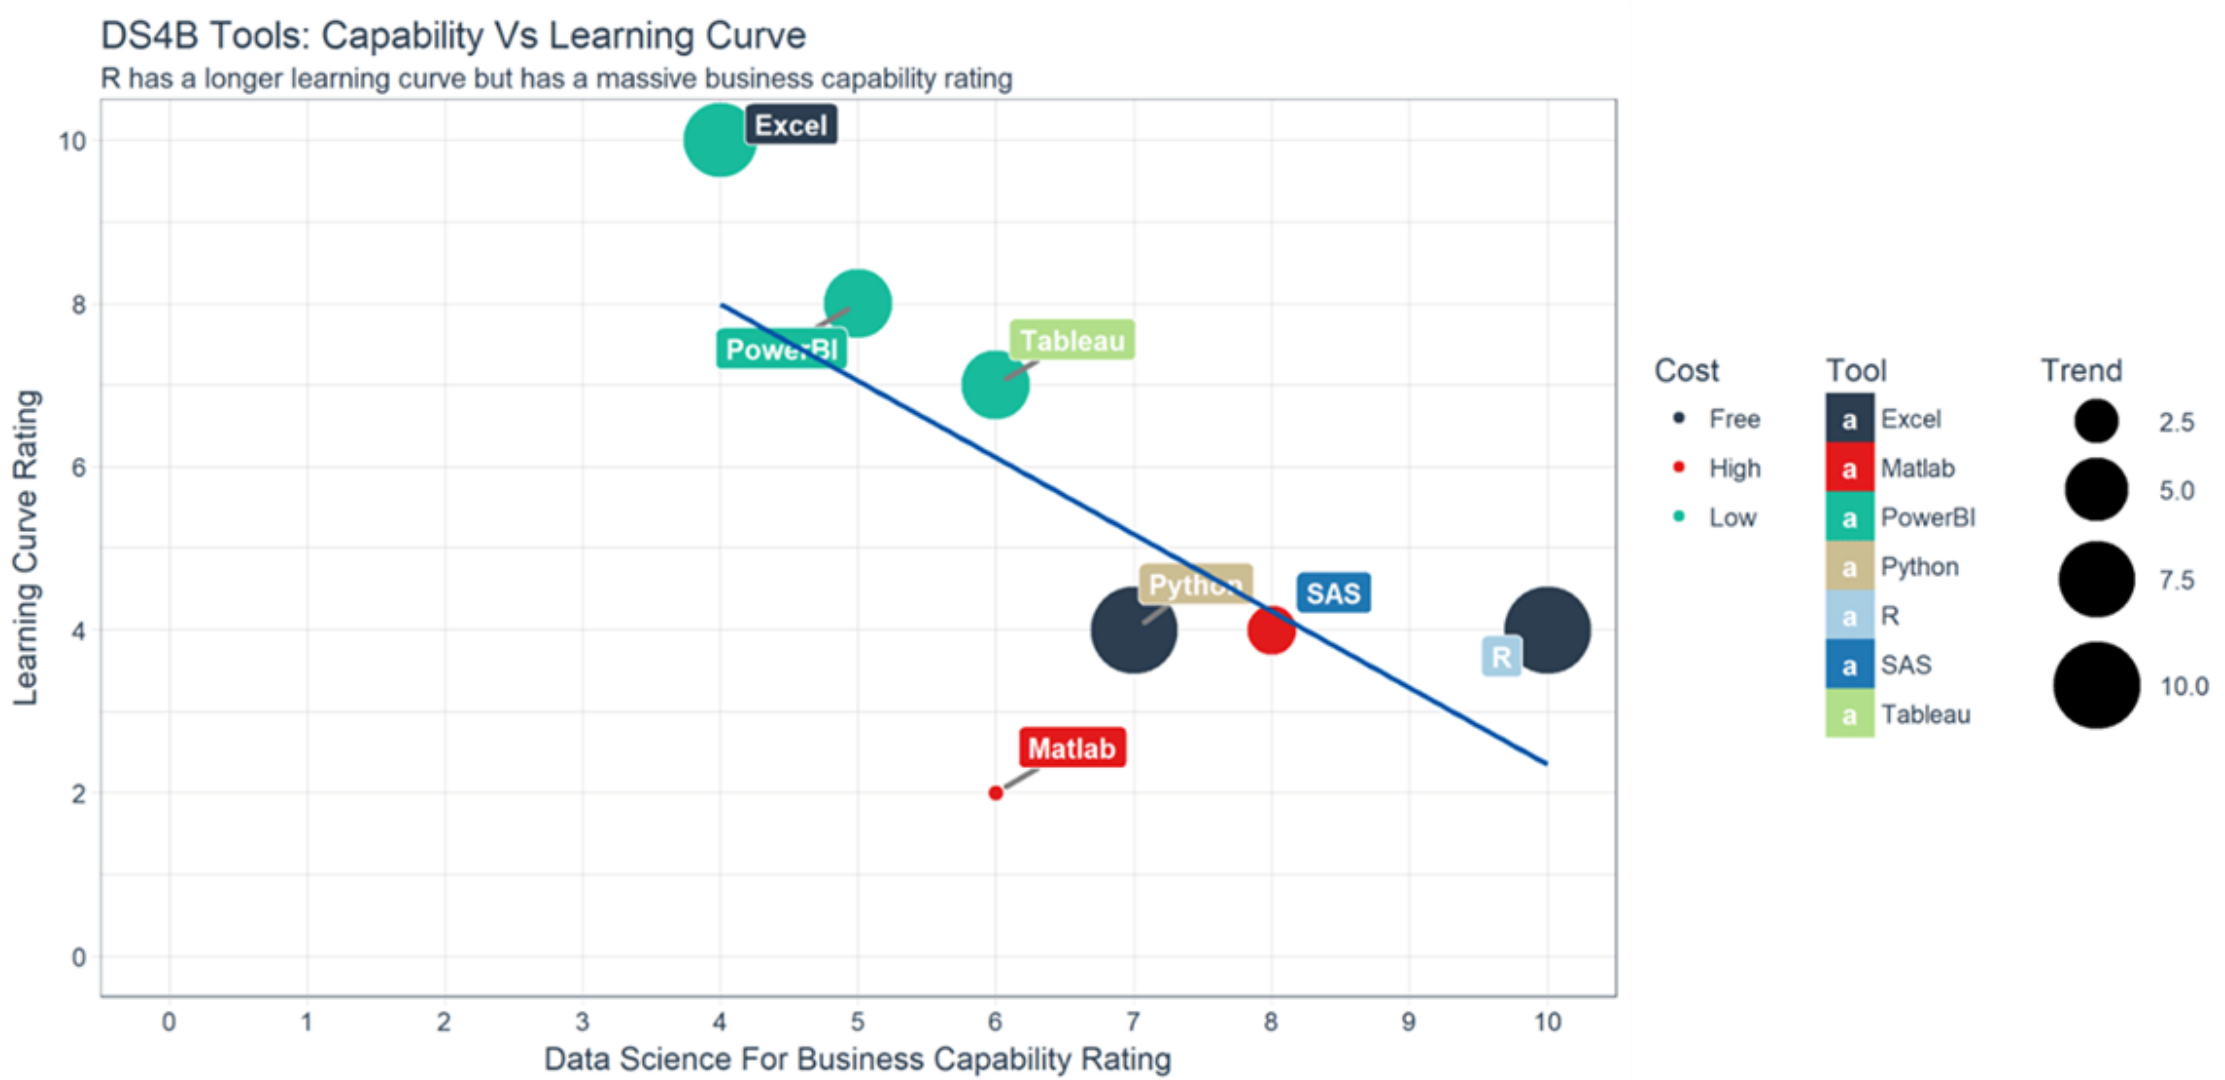
DS4B Tools: Capability Vs Learning Curve
R has a longer learning curve but has a massive business capability rating
Learning Curve Rating
Data Science For Business Capability Rating
Excel
PowerBI
Tableau
Python
SAS
R
Matlab
Cost
Free
High
Low
Tool
Excel
Matlab
PowerBI
Python
R
SAS
Tableau
Trend
2.5
5.0
7.5
10.0

R is the best platform for data munging or visualization

Szilard Pafka
Data Science and R Guru
Sta Monica, California

Companies that use R for Analytics

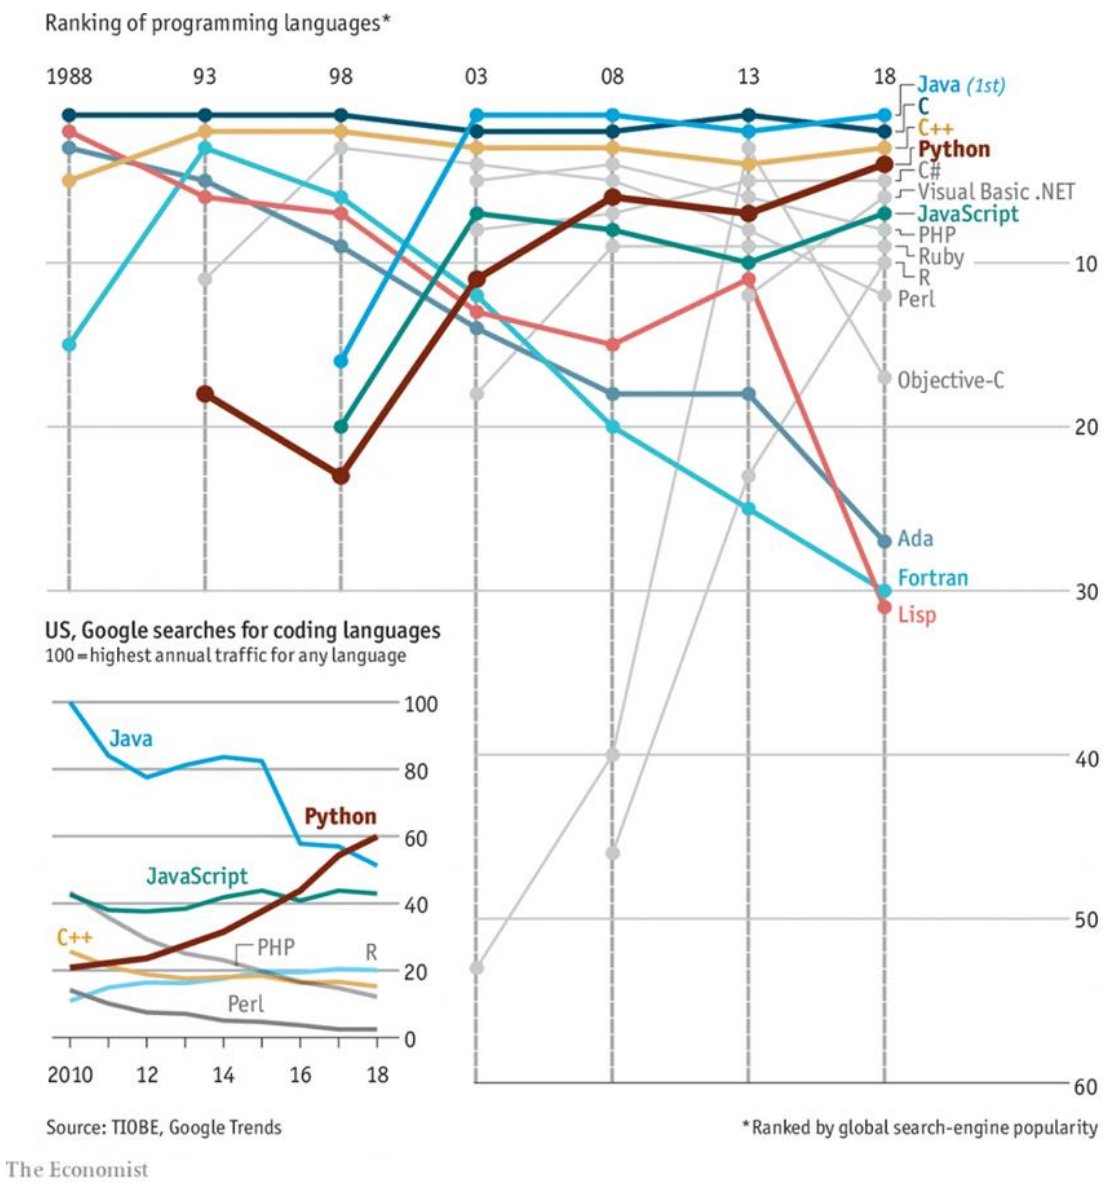
Ranking of programming languages*
1988
93
98
03
08
13
18
Java (1st)
C
C++
Python
C#
Visual Basic .NET
JavaScript
PHP
Ruby
R
Perl
Objective-C
Ada
Fortran
Lisp
10
20
30
40
50
60
US, Google searches for coding languages
100 = highest annual traffic for any language
Java
Python
JavaScript
C++
PHP
R
Perl
100
80
60
40
20
0
2010
12
14
16
18
Source: TIOBE, Google Trends
*Ranked by global search-engine popularity
The Economist

R is a “glue”

```{r}
idOper <- allTablesRowsPct %>%
  filter(col=="ID_OPER") %>%
  mutate(is702=map_int(sprintf("select count(*) cnt from %s where %s=%d",tab,col,702),
                       ~sqlQuery(dbhandle,.x)$cnt),
         is749=map_int(sprintf("select count(*) cnt from %s where %s=%d",tab,col,749),
                       ~sqlQuery(dbhandle,.x)$cnt))
```

UNIX

C

SQL

dplyr

lisp (via purrr)

Cambrian Explosion

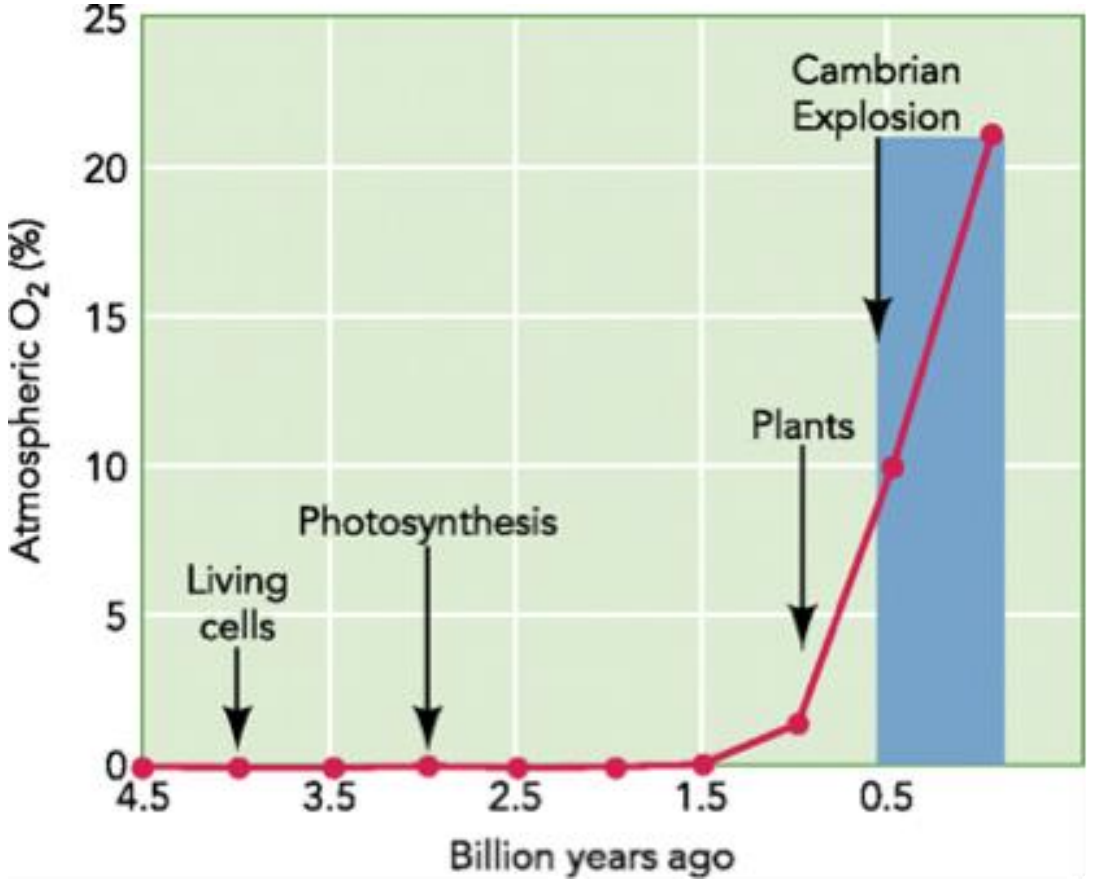

Package Explosion

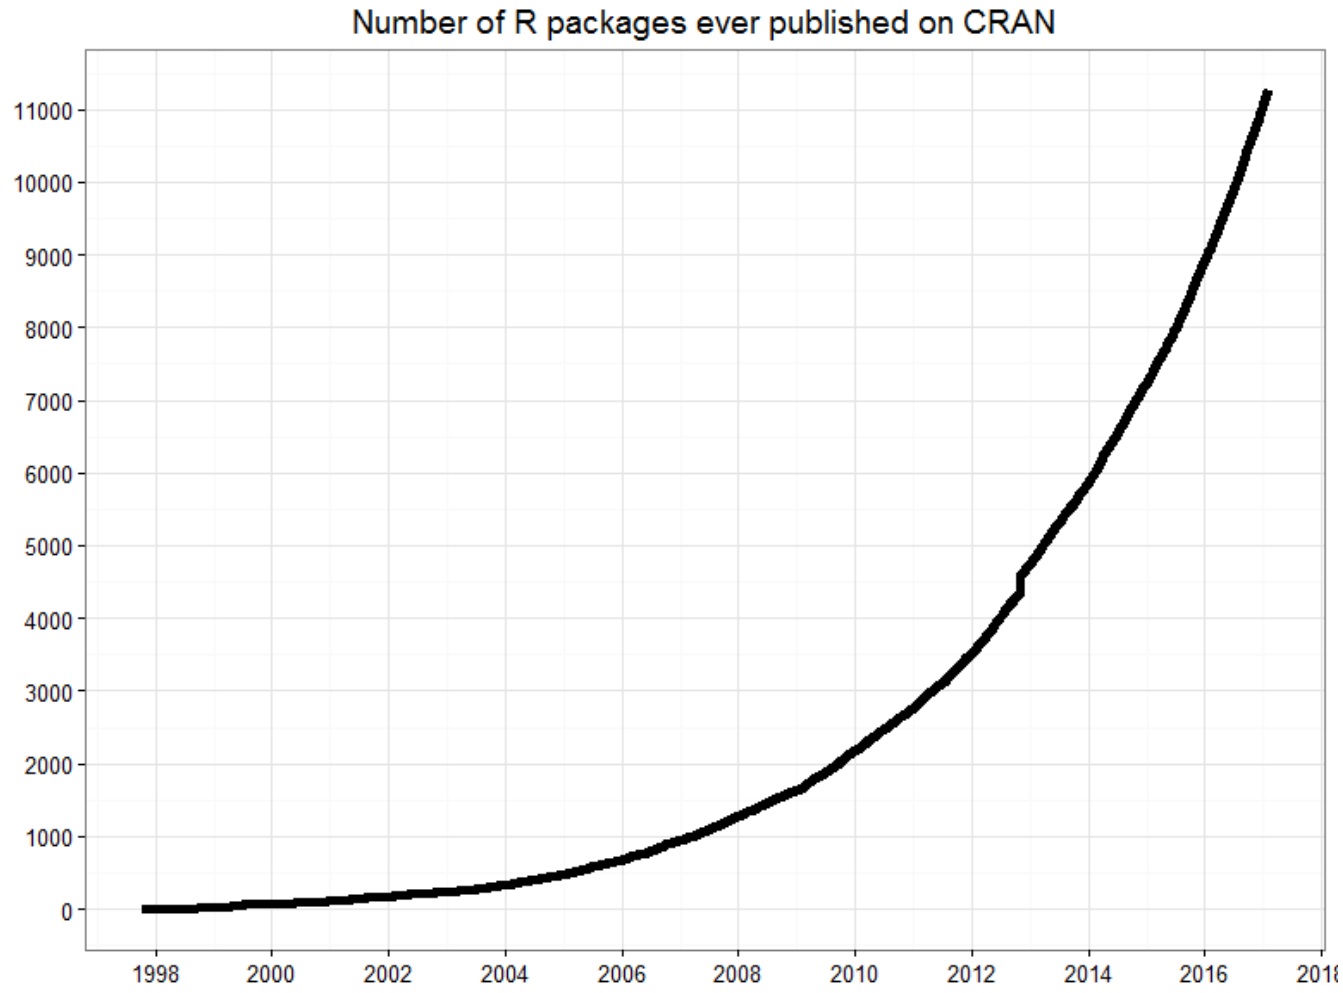

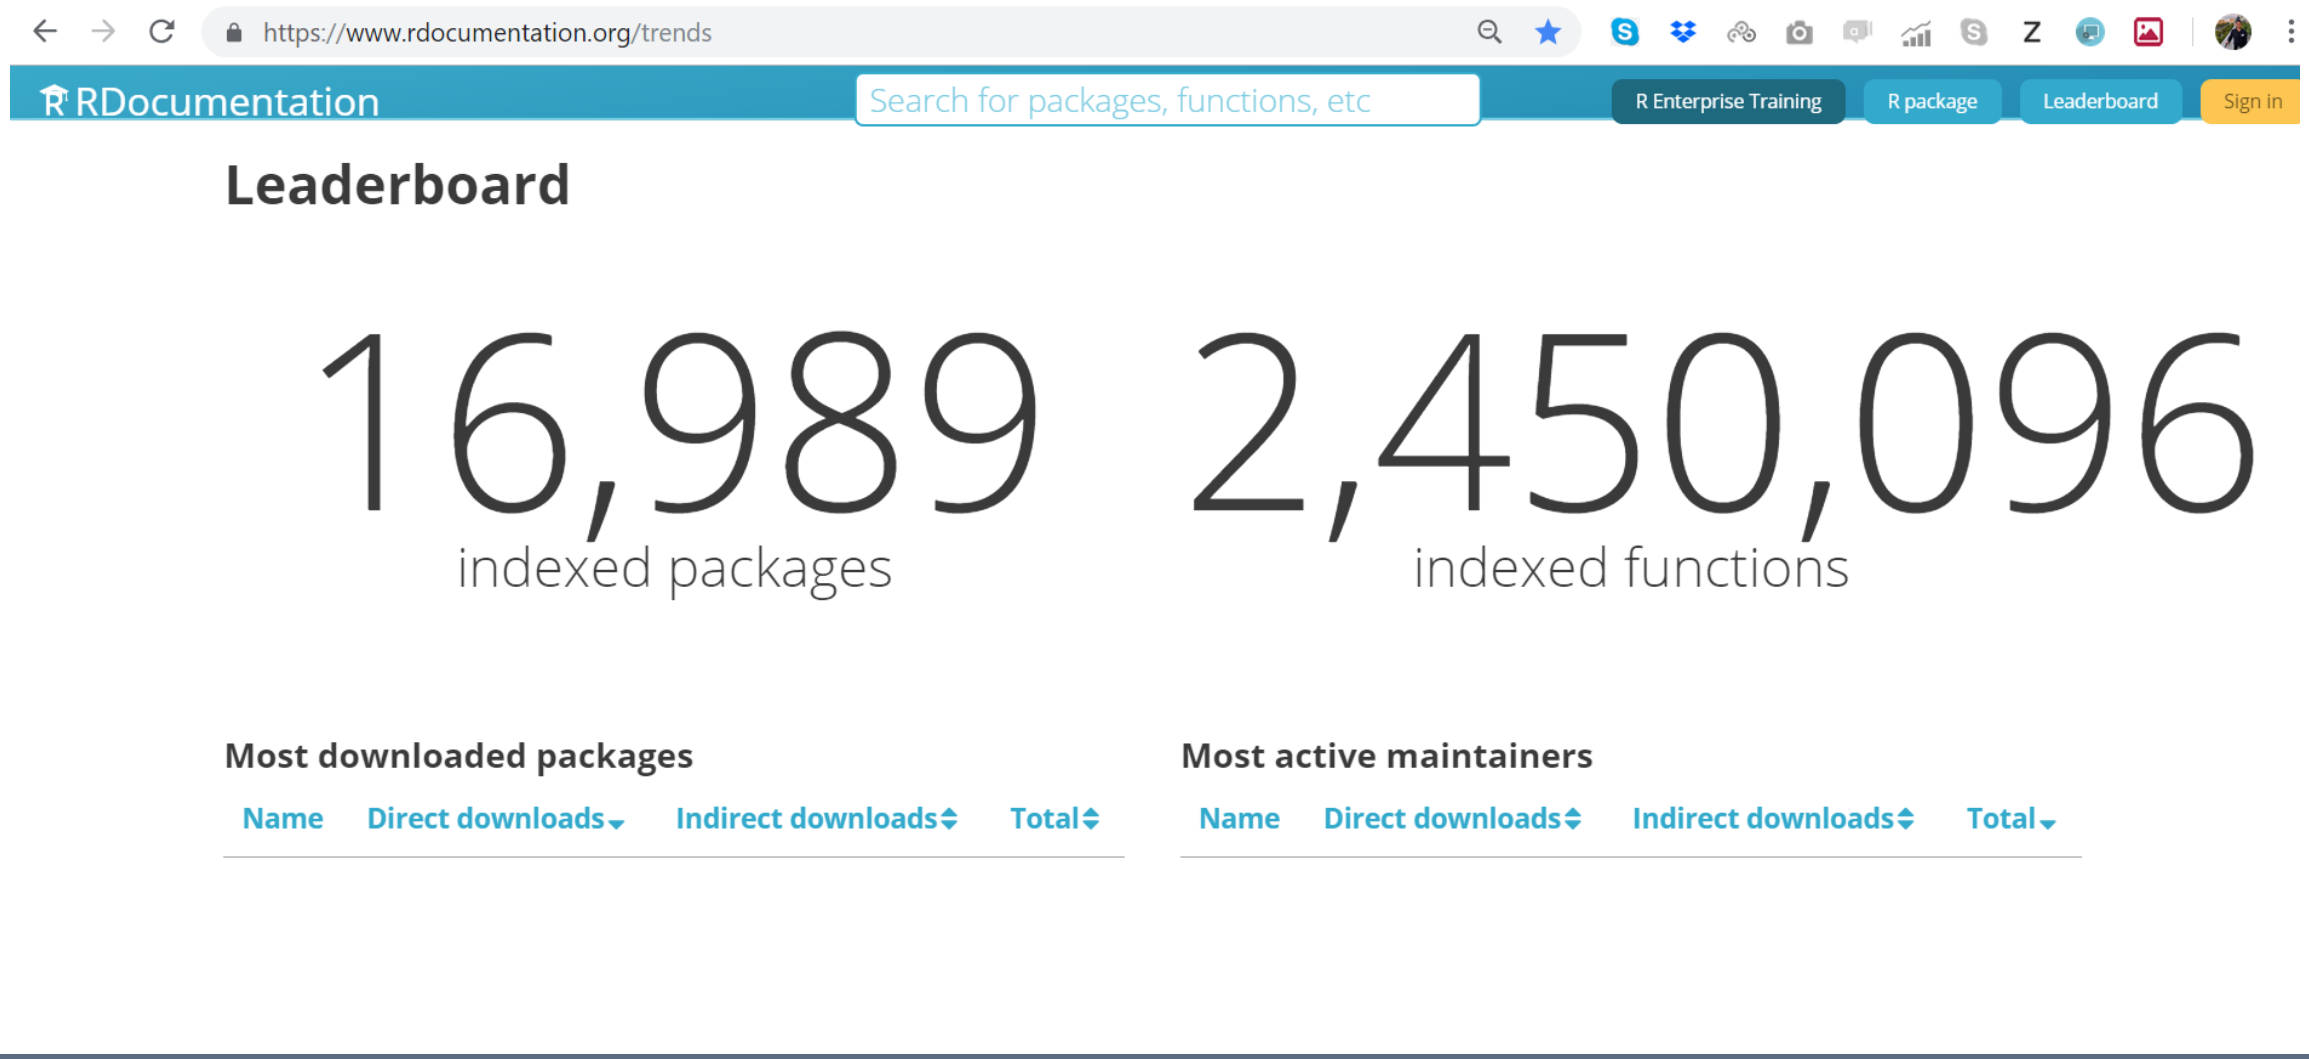
https://www.rdocumentation.org/trends
RDocumentation
Search for packages, functions, etc
R Enterprise Training
R package
Leaderboard
Sign in
Leaderboard
16,989
indexed packages
2,450,096
indexed functions
Most downloaded packages
Name
Direct downloads
Indirect downloads
Total
Most active maintainers
Name
Direct downloads
Indirect downloads
Total

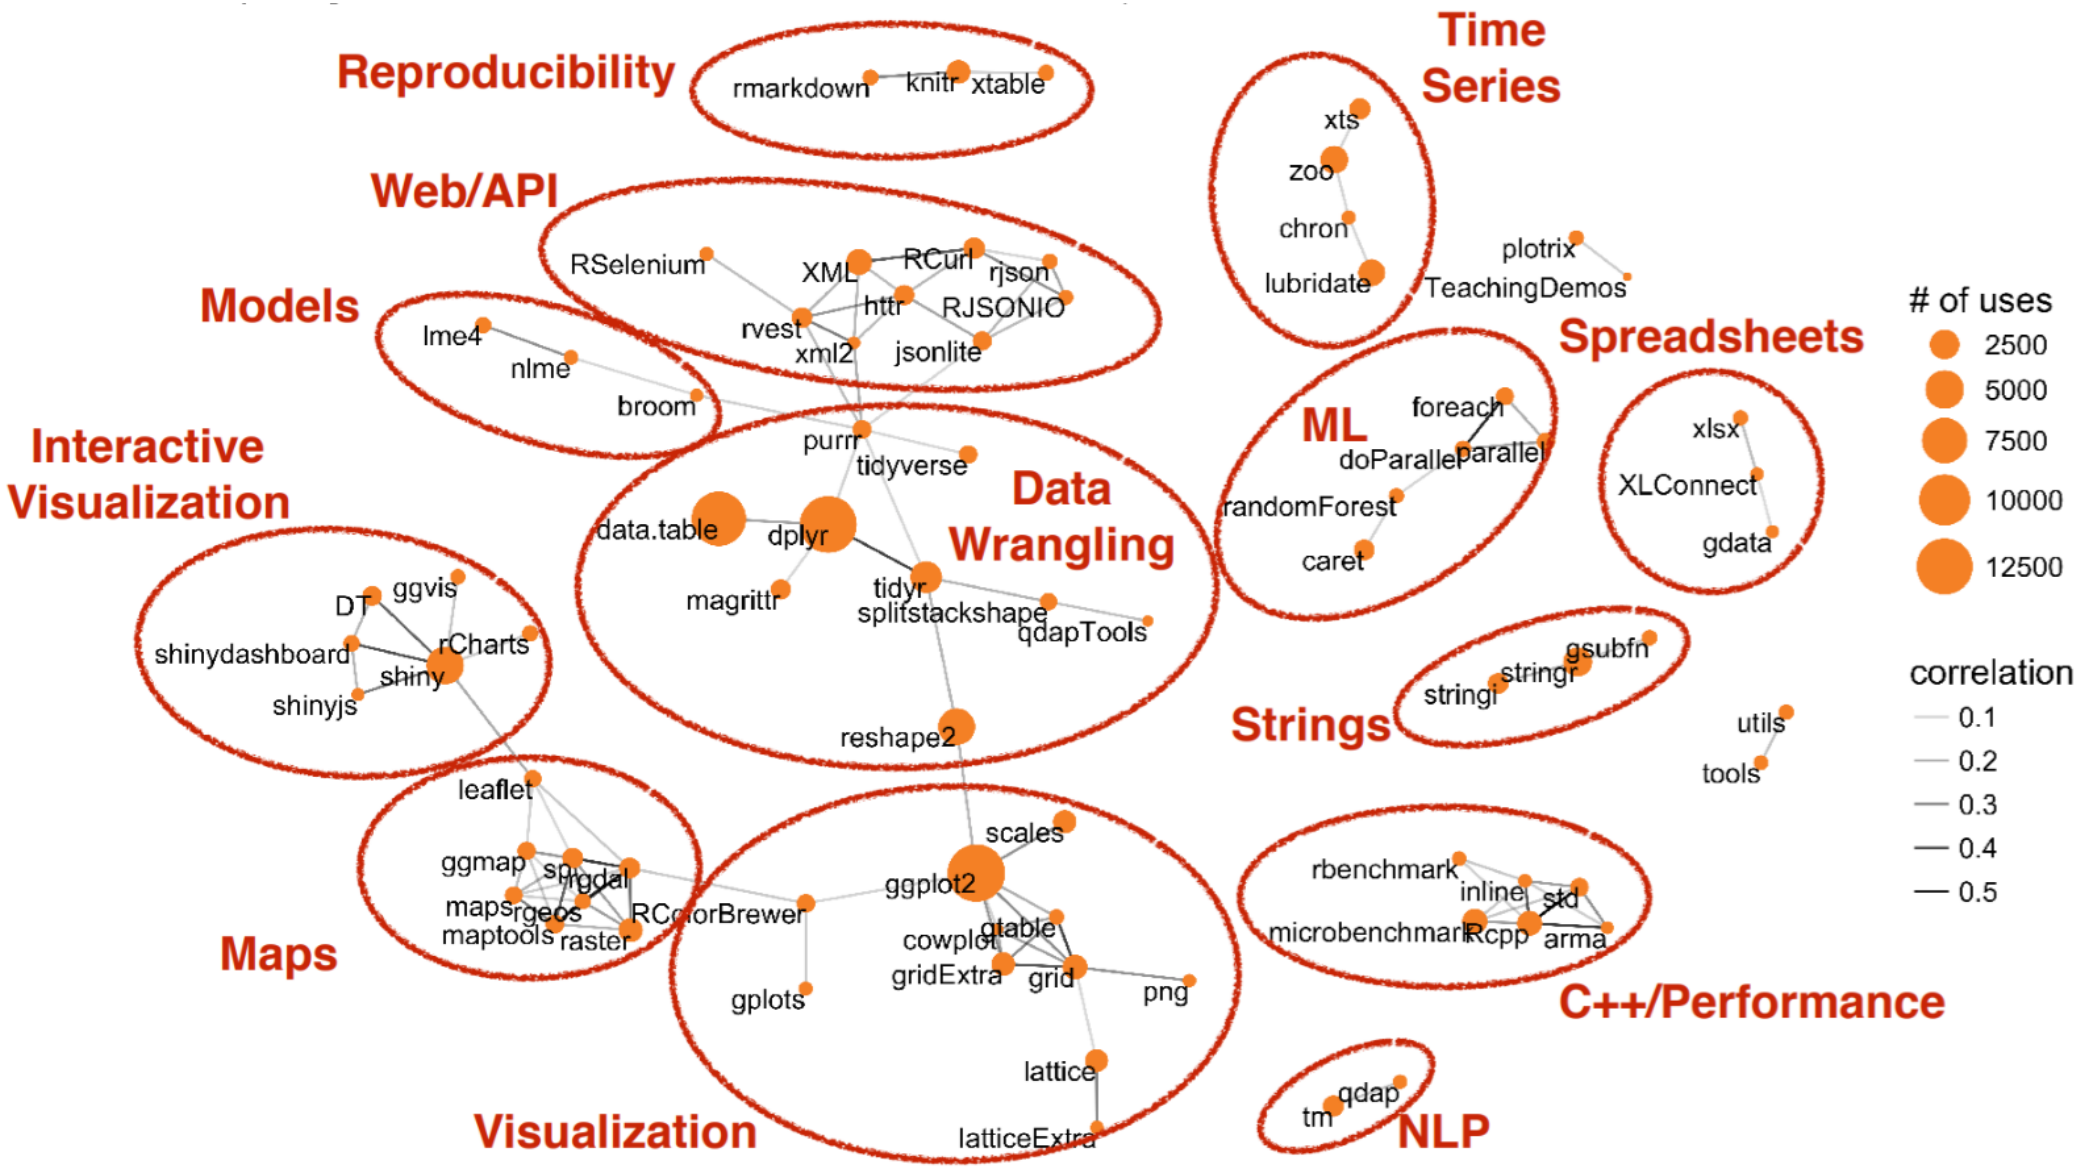
Reproducibility
rmarkdown
knitr
xtable
Web/API
RSelenium
XML
RCurl
rjson
httr
RJSONIO
rvest
xml2
jsonlite
Models
lme4
nlme
broom
Interactive Visualization
DT
ggvis
rCharts
shinydashboard
shiny
shinyjs
Data Wrangling
purrr
tidyverse
data.table
dplyr
tidyr
magrittr
splitstackshape
qdapTools
reshape2
Maps
leaflet
ggmap
sp
rgdal
maps
rgeos
maptools
raster
RColorBrewer
Visualization
scales
ggplot2
cowplot
gtable
gridExtra
grid
png
gplots
lattice
latticeExtra
Time Series
xts
zoo
chron
lubridate
plotrix
TeachingDemos
ML
foreach
doParallel
parallel
randomForest
caret
Spreadsheets
xlsx
XLConnect
gdata
Strings
gsubfn
stringr
stringi
utils
tools
C++/Performance
rbenchmark
inline
std
microbenchmark
Rcpp
arma
NLP
qdap
tm
# of uses
2500
5000
7500
10000
12500
correlation
0.1
0.2
0.3
0.4
0.5

CRAN

BioConductor

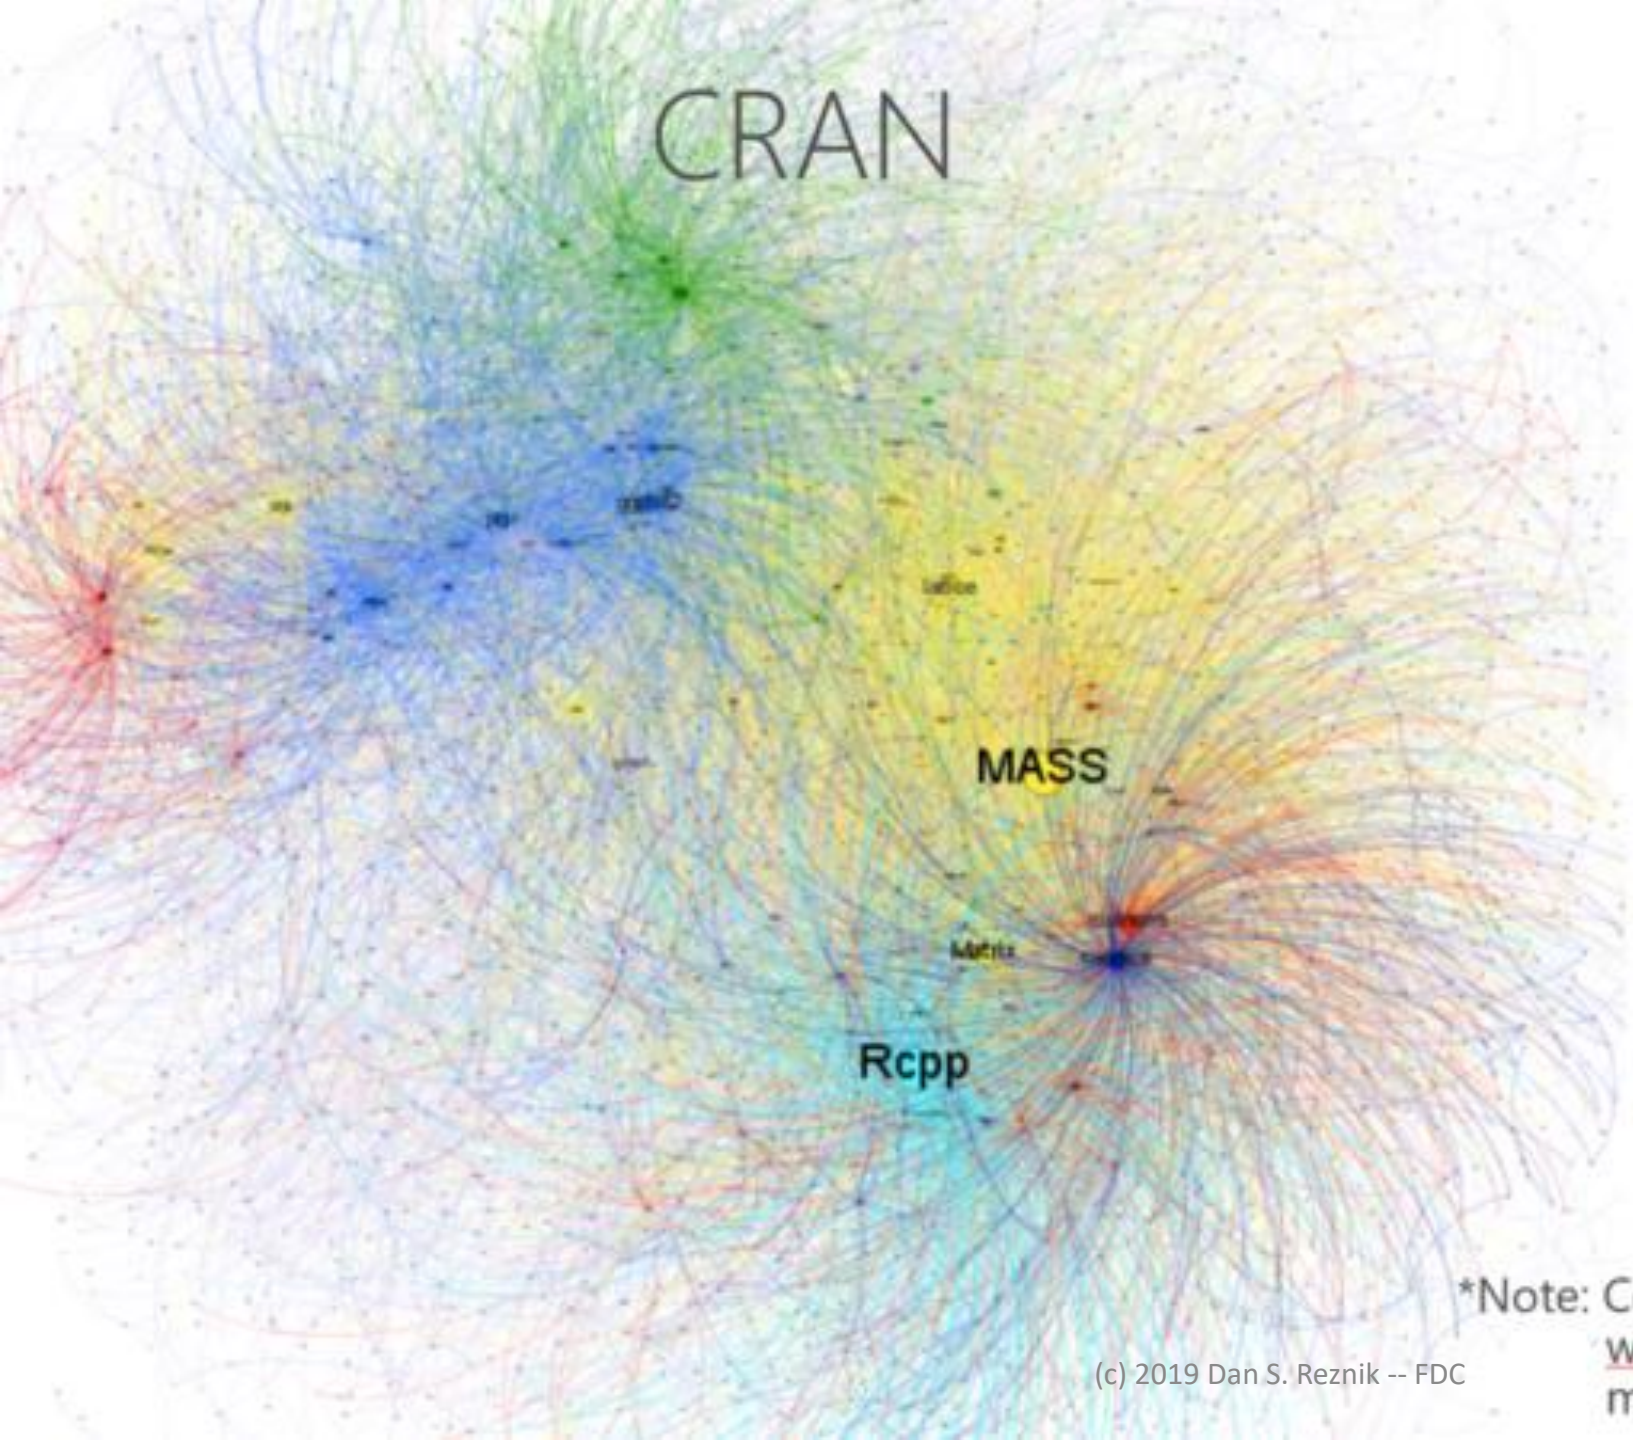

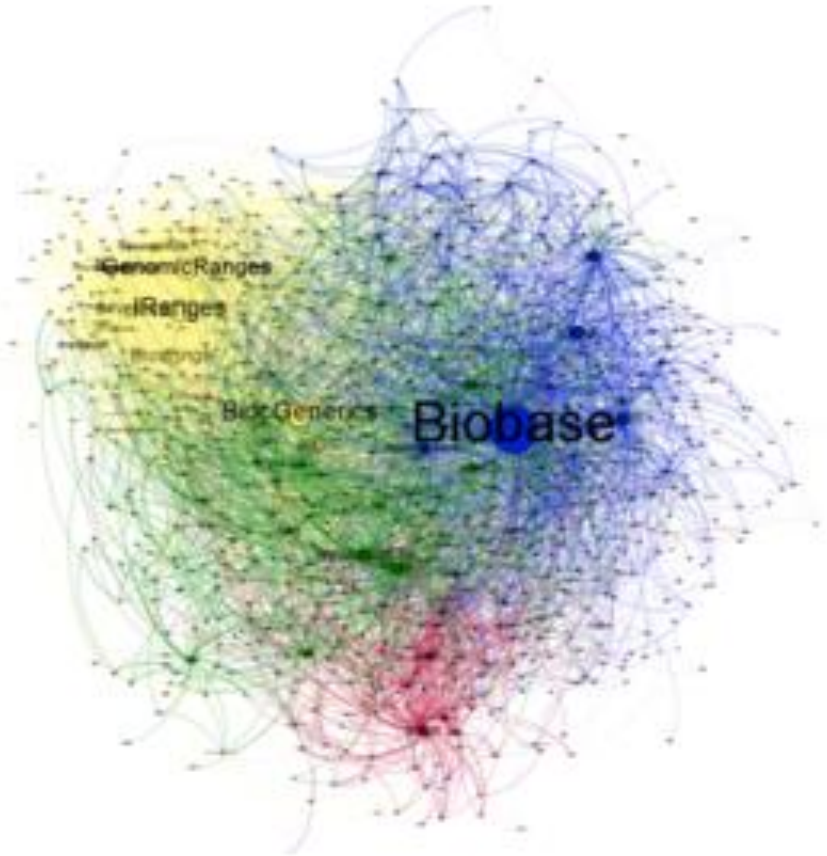

*Note: Colour indicates communities found by the walktrap algorithm, but has no common meaning in the two networks

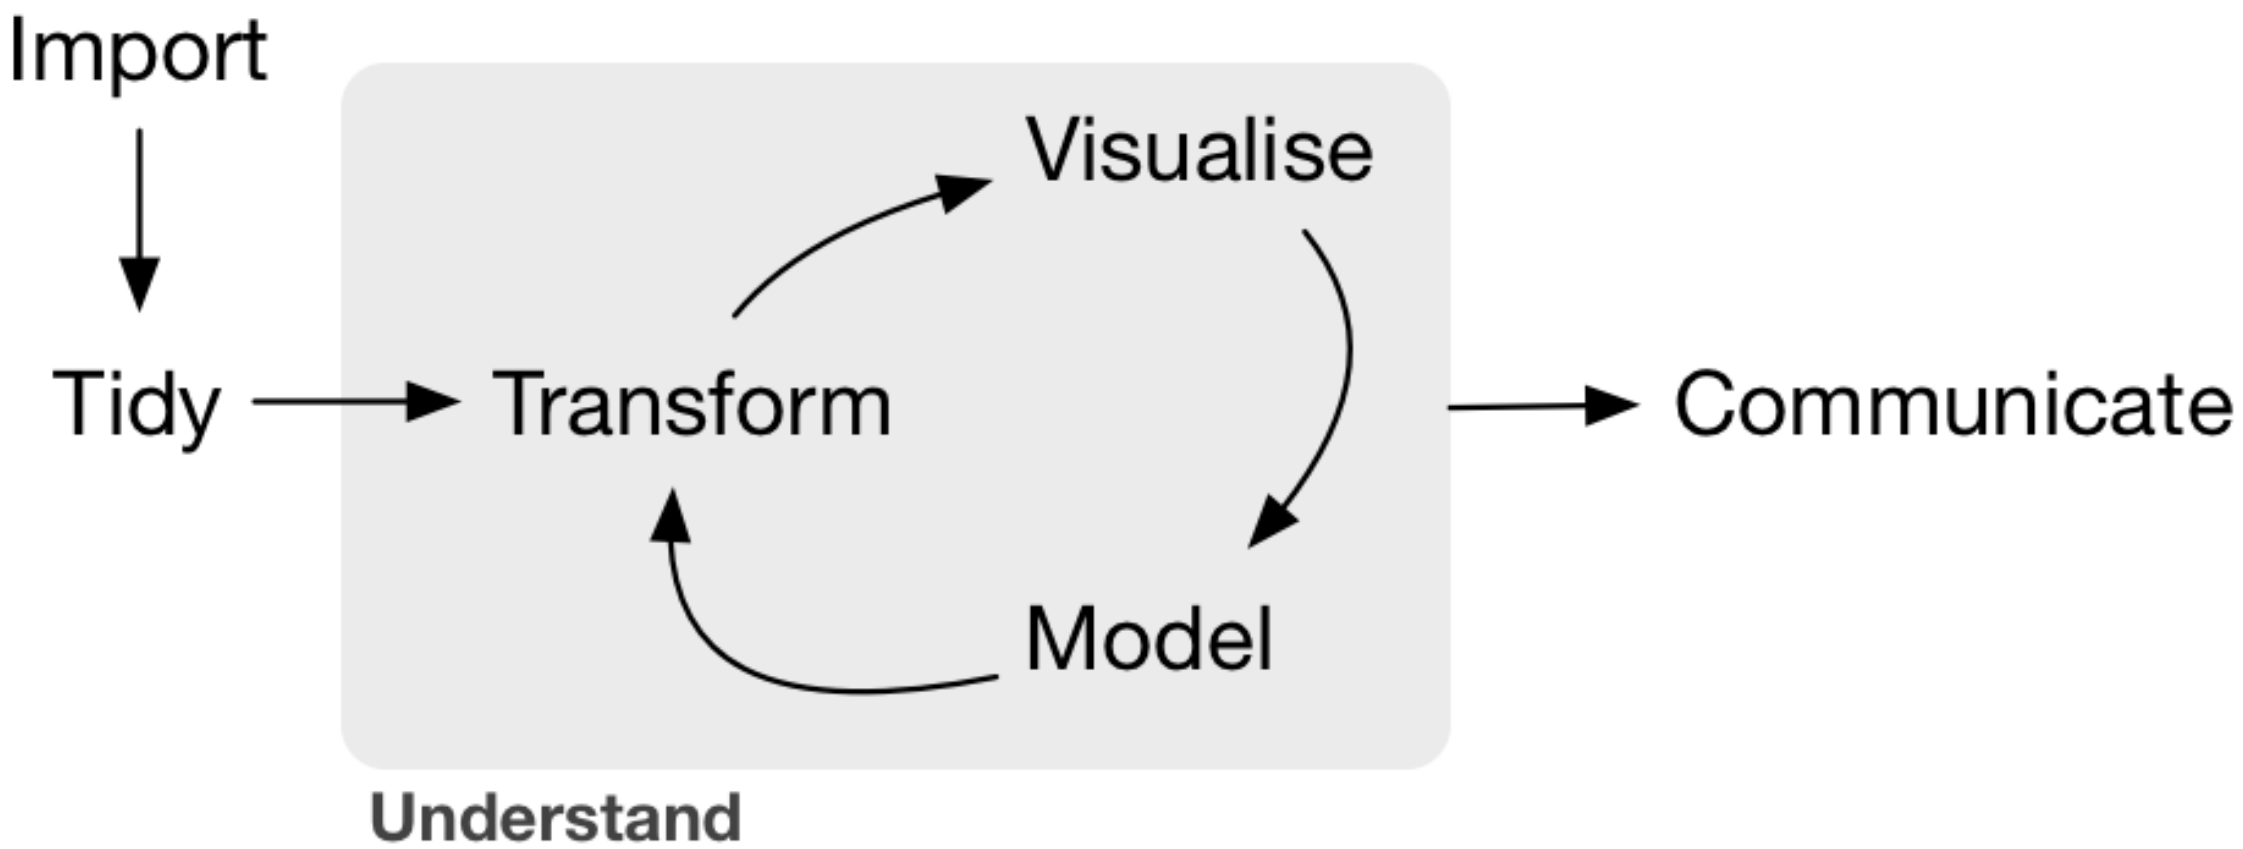
Import
Tidy
Transform
Visualise
Model
Communicate
Understand

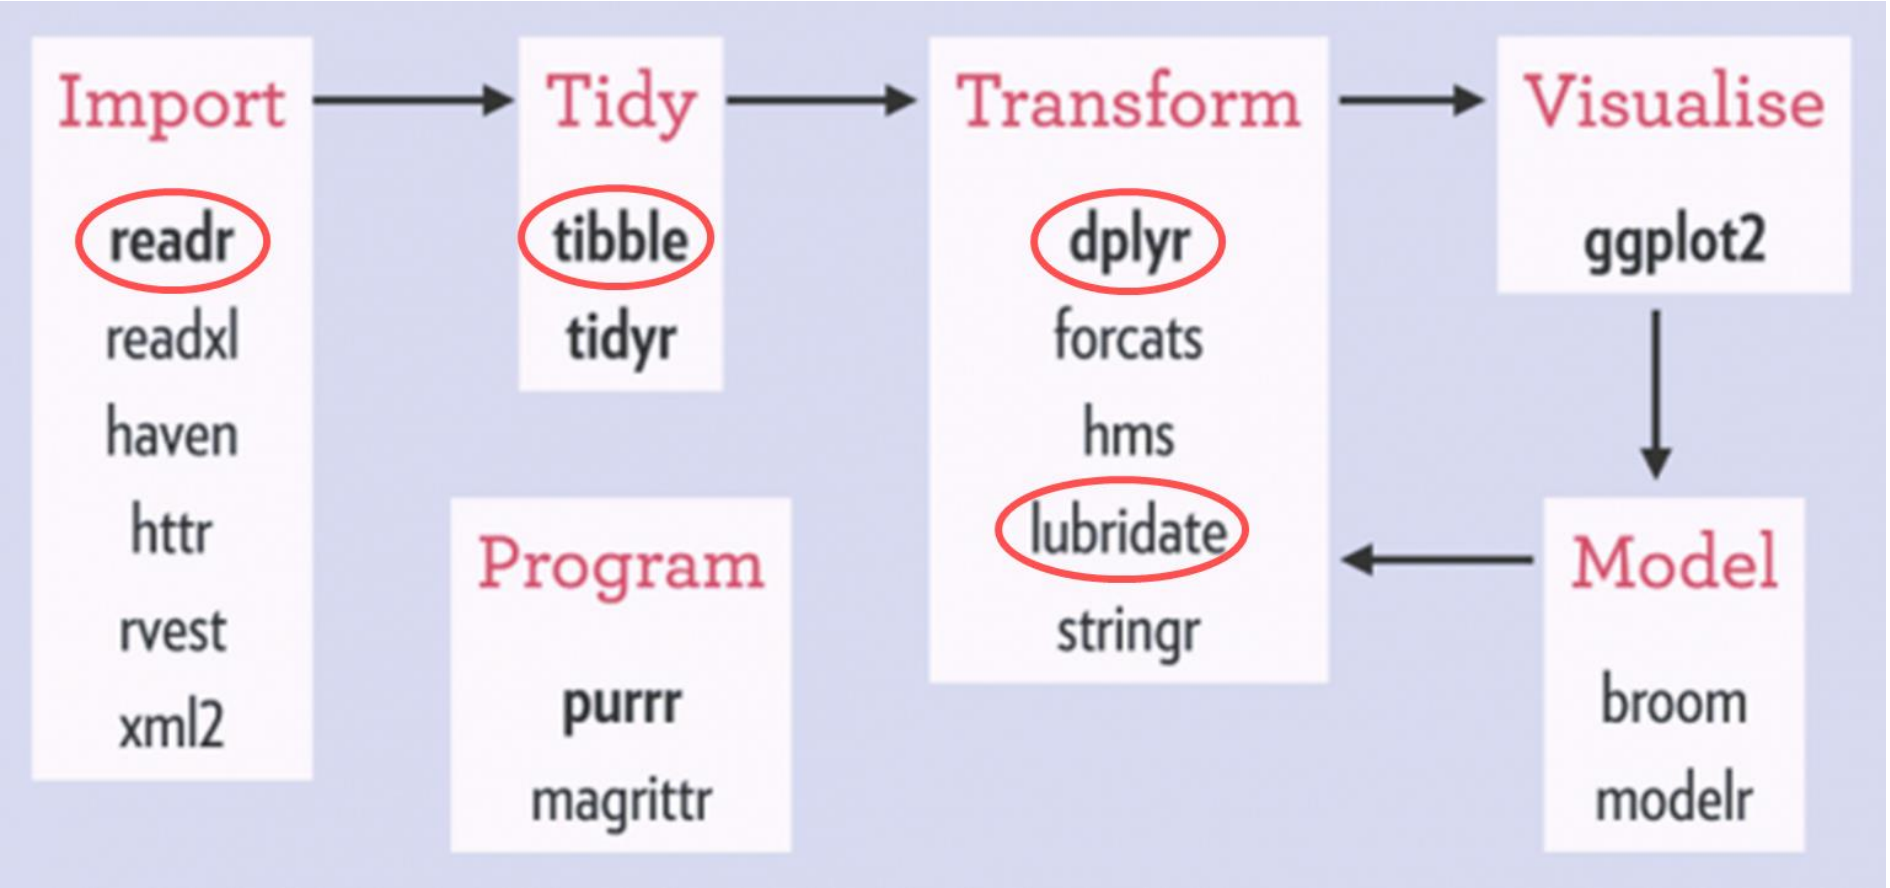

The tidyverse: packages for the data science workflow

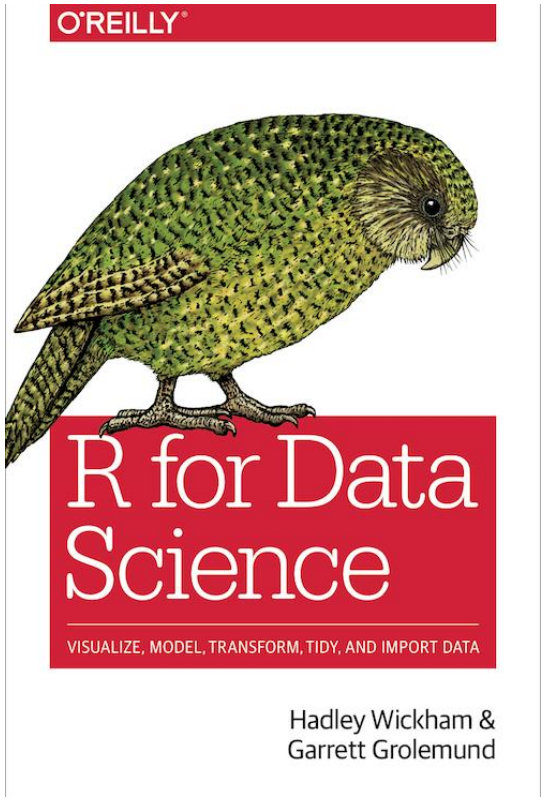

https://r4ds.had.co.nz/

RStudio

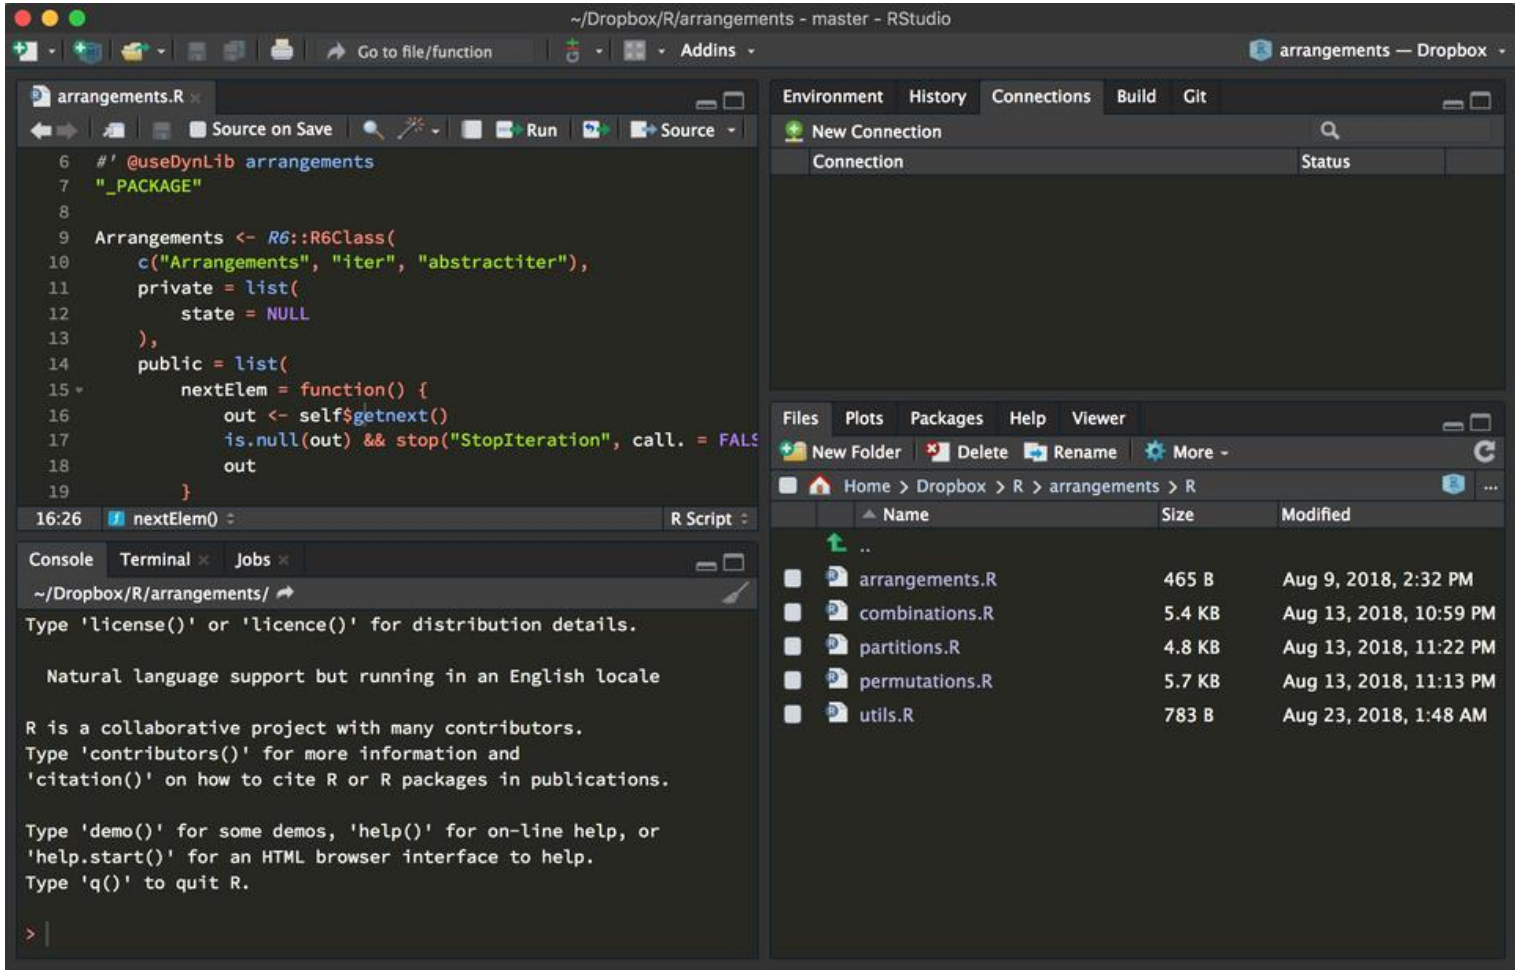
~/Dropbox/R/arrangements - master - RStudio
Go to file/function
Addins
arrangements — Dropbox
arrangements.R
Source on Save
Run
Source
#' @useDynLib arrangements
"_PACKAGE"
Arrangements <- R6::R6Class(
c("Arrangements", "iter", "abstractiter"),
private = list(
state = NULL
),
public = list(
nextElem = function() {
out <- self$getnext()
is.null(out) && stop("StopIteration", call. = FALS
out
}
16:26
nextElem()
R Script
Console
Terminal
Jobs
~/Dropbox/R/arrangements/
Type 'license()' or 'licence()' for distribution details.
Natural language support but running in an English locale
R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.
Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.
>
Environment
History
Connections
Build
Git
New Connection
Connection
Status
Files
Plots
Packages
Help
Viewer
New Folder
Delete
Rename
More
Home > Dropbox > R > arrangements > R
Name
Size
Modified
..
arrangements.R 465 B Aug 9, 2018, 2:32 PM
combinations.R 5.4 KB Aug 13, 2018, 10:59 PM
partitions.R 4.8 KB Aug 13, 2018, 11:22 PM
permutations.R 5.7 KB Aug 13, 2018, 11:13 PM
utils.R 783 B Aug 23, 2018, 1:48 AM

GitHub

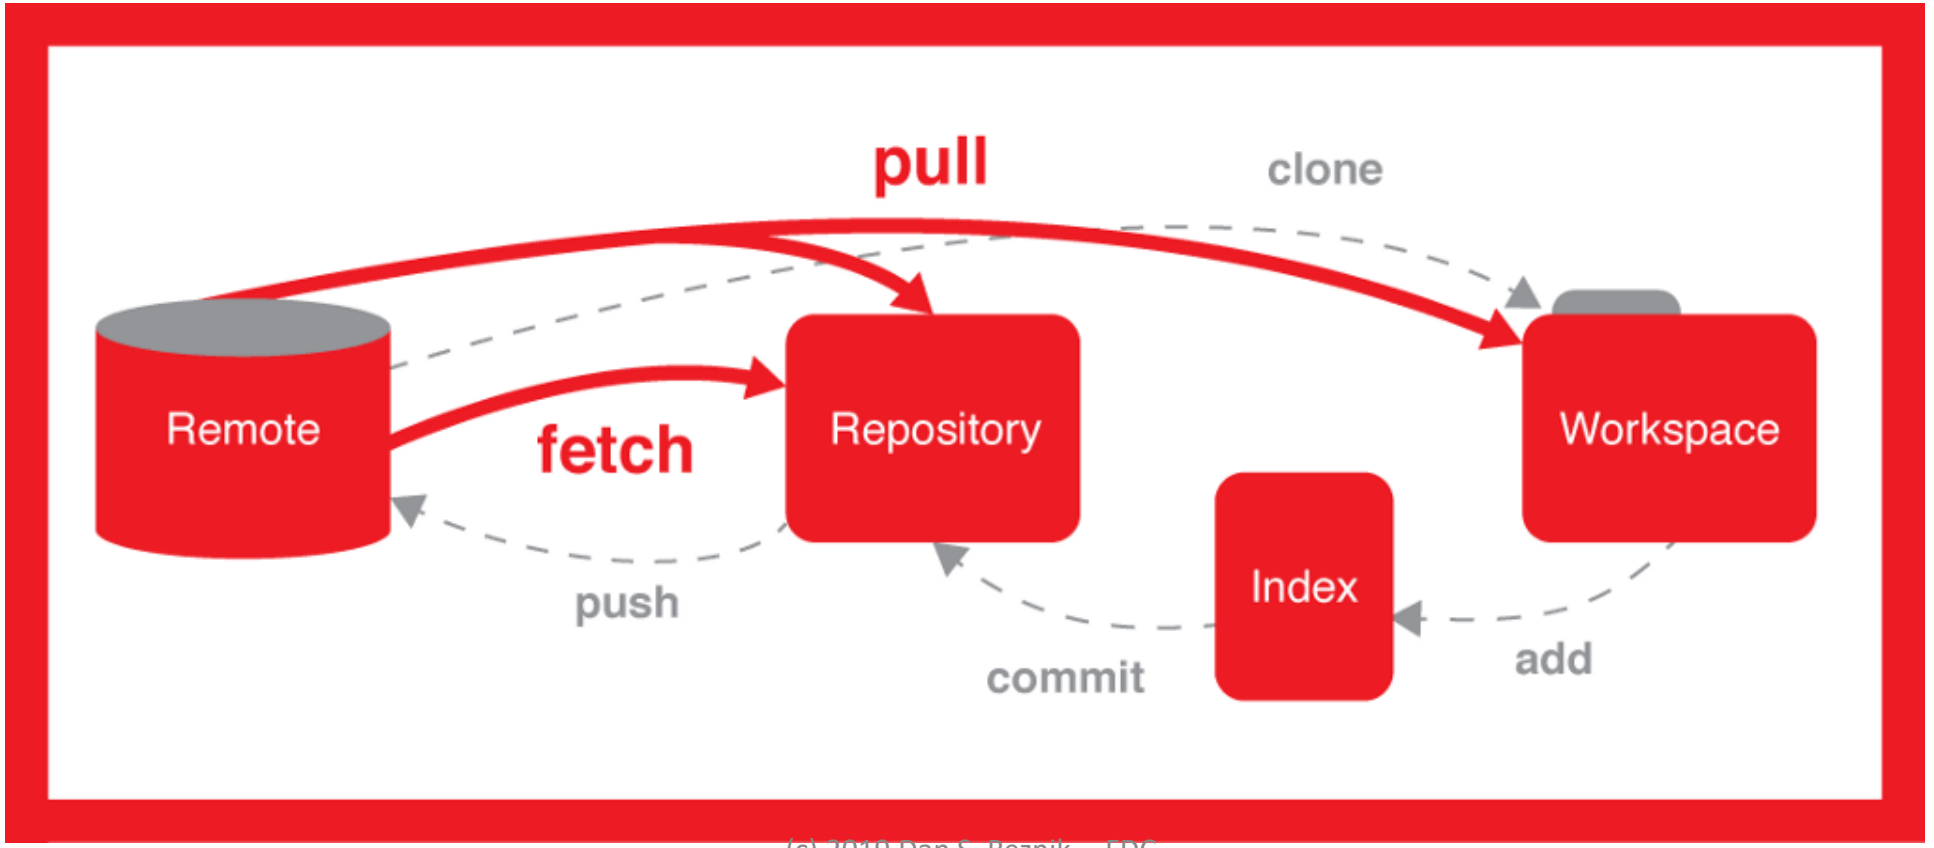
pull
clone
Remote
fetch
Repository
Workspace
Index
push
commit
add

Install R

R-3.6.1 for Windows (32/64 bit)

Download R 3.6.1 for Windows (81 megabytes, 32/64 bit)

Installation and other instructions
New features in this version

If you want to double-check that the package you have downloaded matches the package distributed by CR
to the fingerprint on the master server. You will need a version of md5sum for windows: both graphical and

Frequently asked questions

- Does R run under my version of Windows?
- How do I update packages in my previous version of R?
- Should I run 32-bit or 64-bit R?

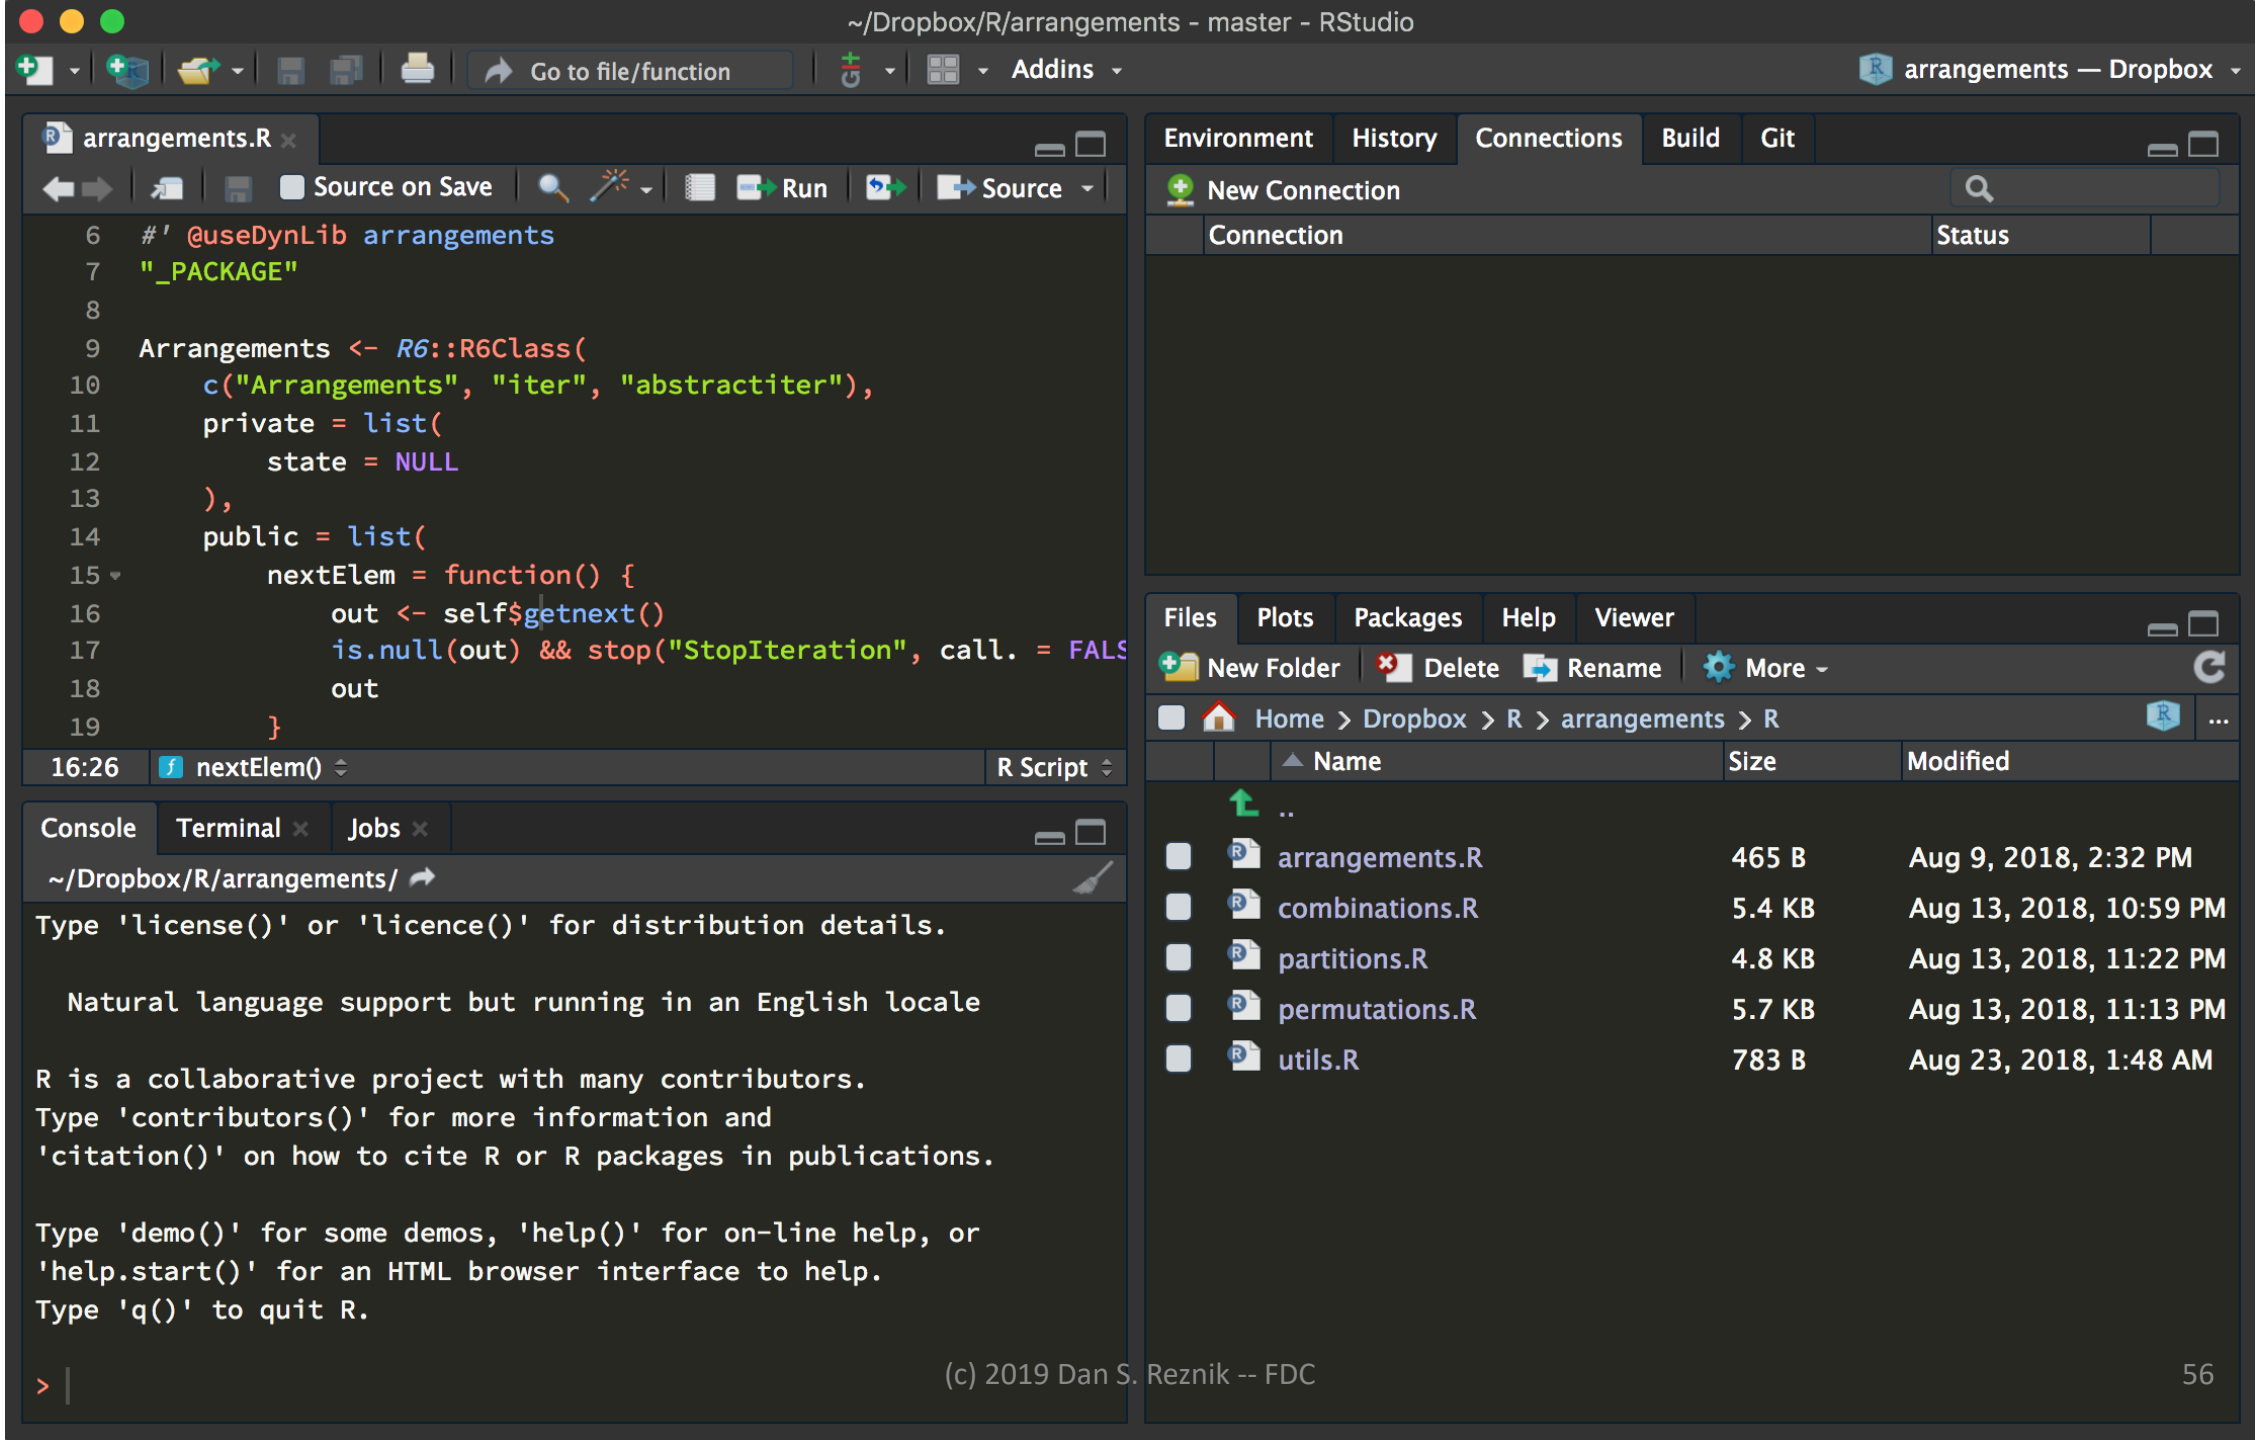
~/Dropbox/R/arrangements - master - RStudio
Go to file/function
Addins
arrangements — Dropbox
arrangements.R
Source on Save
Run
Source
#' @useDynLib arrangements
"_PACKAGE"
Arrangements <- R6::R6Class(
c("Arrangements", "iter", "abstractiter"),
private = list(
state = NULL
),
public = list(
nextElem = function() {
out <- self$getnext()
is.null(out) && stop("StopIteration", call. = FALS
out
}
16:26
nextElem()
R Script
Console
Terminal
Jobs
~/Dropbox/R/arrangements/
Type 'license()' or 'licence()' for distribution details.
Natural language support but running in an English locale
R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.
Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.
>
Environment
History
Connections
Build
Git
New Connection
Connection
Status
Files
Plots
Packages
Help
Viewer
New Folder
Delete
Rename
More
Home > Dropbox > R > arrangements > R
Name
Size
Modified
..
arrangements.R 465 B Aug 9, 2018, 2:32 PM
combinations.R 5.4 KB Aug 13, 2018, 10:59 PM
partitions.R 4.8 KB Aug 13, 2018, 11:22 PM
permutations.R 5.7 KB Aug 13, 2018, 11:13 PM
utils.R 783 B Aug 23, 2018, 1:48 AM

R DS Pipeline

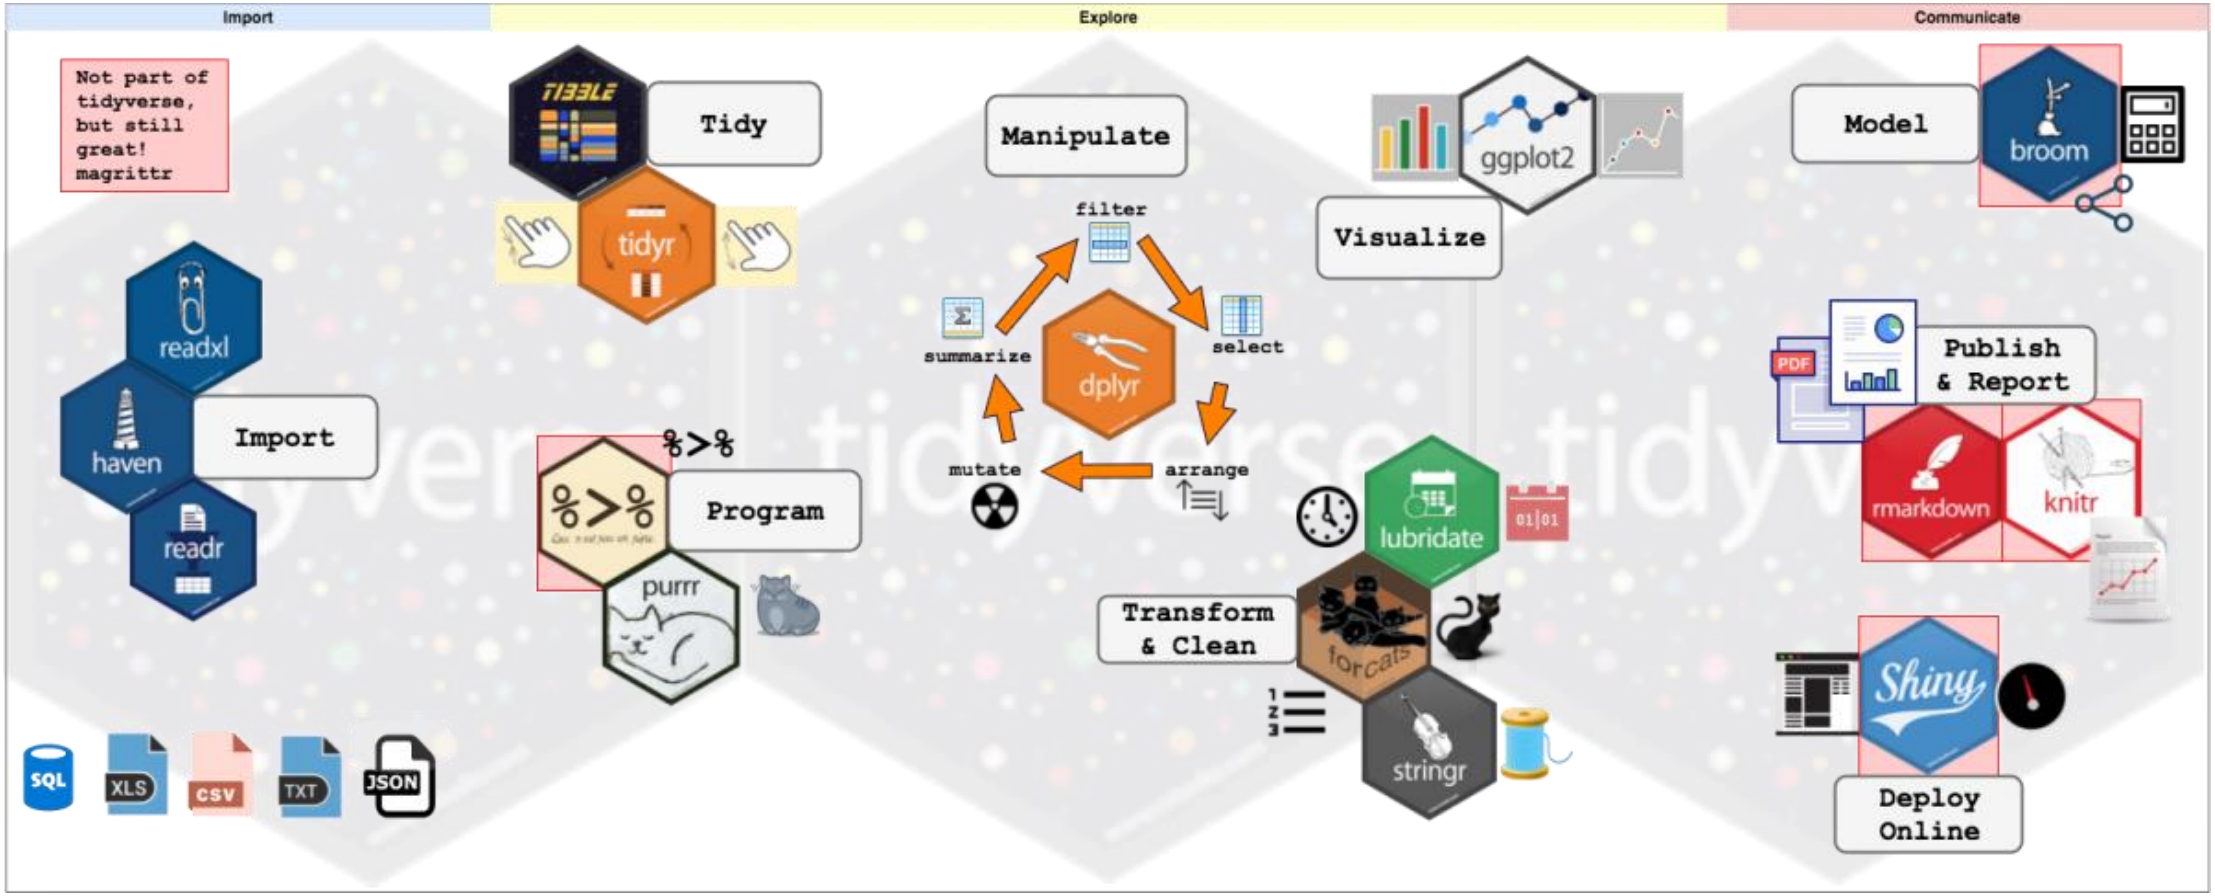

Takeaways

- Business Impact
  - Data Science
  - Machine Learning
- Tools for Working the Data
  - Cloud
  - GitHub
  - R Ecossystem
- Business Problem First, Technology Later

Thank you!

dan@dat-sci.com

www.dat-sci.com

Web: Data Science

- [book] Strategic Principles of DS https://lucidmanager.org/strategic-data-science/
- Towards Data Science: https://towardsdatascience.com/
- Automation Tools for DS: https://www.kdnuggets.com/2018/12/automation-data-science.html
- Hadley on how to become a DS: https://gist.github.com/hadley/820f09ded347c62c2864?fbclid=IwAR22uF47RIgAmC0ZcErb4MuVRVdHBQZKEn9yxn8S5tOBppxjE_Onfqjt5iQ
- Intro to GitHub: https://medium.com/@abhishekj/an-intro-to-git-and-github-1a0e2c7e3a2f
- Command-Line Data Science: https://adamdrake.com/command-line-tools-can-be-235x-faster-than-your-hadoop-cluster.html
- WEF: DS in the new economy: http://www3.weforum.org/docs/WEF_Data_Science_In_the_New_Economy.pdf
- Data Scientist vs Engineer vs Analyst: https://news.efinancialcareers.com/uk-en/3001517/data-science-careers-finance
- R for Data Science: https://r4ds.had.co.nz/

Web: ML/AI

- Data Science vs ML: https://www.kdnuggets.com/2018/12/learning-machine-learning-data-science.html
- Davenport: Shift Toward AI: https://www.forbes.com/sites/tomdavenport/2019/04/05/ai-is-destroying-traditional-business-thinking
- Amazing Innovation w/ Self-Driving Cars: https://www.cbinsights.com/research/startups-drive-auto-industry-disruption
- All ML Algos in one page: https://www.r-bloggers.com/101-machine-learning-algorithms-for-data-science-with-cheat-sheets/
- State of ML 2019: https://www.analyticsinsight.net/the-state-of-machine-learning-in-2019/
- ML Process: https://www.kdnuggets.com/2018/12/essence-machine-learning.html
- ML with R: https://bradleyboehmke.github.io/HOML/
- The AI “Clown Show”: https://blog.piekniewski.info/2019/05/30/ai-circus-mid-2019-update/
- How to recognize AI “snakeoil”: https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf

Web: R

- Is it Worth Learning R: https://www.quora.com/Is-it-worth-learning-R-programming/answer/Aaron-Kumar-35
- Ditch the Spreadsheet: https://lucidmanager.org/spreadsheets-for-data-science
- Data Exploration Package: https://boxuancui.github.io/DataExplorer/
- Using GitHub with R: https://happygitwithr.com/
- Feature Engineering with the Tidyverse: https://bookdown.org/max/FES/
- Programming with the Tidyverse: https://speakerdeck.com/lionelhenry/programming-in-the-tidyverse
- R tutorials: https://teachingr.com/
- EDA w/ R: https://www.r-bloggers.com/part-2-simple-eda-in-r-with-inspectdf/
- From Excel to R: https://rfortherestofus.com/2019/06/a-guide-to-r-for-excel-users/

Web: Viz

- Data Viz with R: https://socviz.co/
- Geocomputation with R: https://geocompr.robinlovelace.net/
- Gallery of gpplot2: https://www.r-graph-gallery.com/ggplot2-package.html
- How to pick the right plot: https://www.data-to-viz.com/
- Shiny Contest Winners: https://blog.rstudio.com/2019/04/05/first-shiny-contest-winners
- Shiny+Plotly: https://plotly-r.com/plotly_book.pdf
- Ggplot and Gganimate: https://djnavarro.github.io/satrdayjoburg/slides/

Videos: Data Science

- General:
  - Just Build Something w/ Data: https://youtu.be/MOdlp1d0PNA
- Viz:
  - Tour of Data Viz: https://youtu.be/9F1pZ4X8xPE
  - Chess Players over Time: https://youtu.be/z2DHpW79w0Y

Videos: R

- Gentle Intro to Stats w/ R/Tidyverse: https://resources.rstudio.com/webinars/a-gentle-introduction-to-tidy-statistics-in-r
- R then and now: https://youtu.be/c075aRNmkUk
- Wrangling w/ R: https://youtu.be/DwWH1mTerOc
- R vs Excel: https://youtu.be/S8wO4ppE9L8
- Deploying R APIs w/ Docker: https://youtu.be/cO3QvioMESo
- Six Reasons R is Ideal for DS: https://youtu.be/D4LPLS9yCmc