

Language-agnostic data analysis workflows and reproducible research

Andrew John Lowe

28 April 2017

Abstract

This is the abstract for a document that briefly describes methods for interacting with different programming languages within the same data analysis workflow.

0.1 Overview

- This talk: language-agnostic (or polyglot) analysis workflows
- I'll show how it is possible to work in multiple languages and switch between them without leaving the workflow you started
- Additionally, I'll demonstrate how an entire workflow can be encapsulated in a markdown file that is rendered to a publishable paper with cross-references and a bibliography (and with raw a \LaTeX file produced as a by-product) in a simple process, making the whole analysis workflow reproducible

0.2 Which tool/language is best?

- TMVA, scikit-learn, h2o, caret, mlr, WEKA, Shogun, ...
- C++, Python, R, Java, MATLAB/Octave, Julia, ...
- Many languages, environments and tools
- People may have strong opinions
- A false polychotomy?
- Be a polyglot!

0.3 Approaches

- Save-and-load
- Language-specific bindings
- Notebook
- knitr

1 Save-and-load approach

1.1 Flat files

- Language agnostic
- Formats like text, CSV, and JSON are well-supported
- Breaks workflow
- Data types may not be preserved (e.g., datetime, NULL)
- New binary format Feather solves this

1.2 Feather

Feather: A fast on-disk format for data frames for R and Python, powered by Apache Arrow, developed by Wes McKinney and Hadley Wickham

```
# In R:
library(feather)
path <- "my_data.feather"
write_feather(df, path)
```

```
# In Python:
import feather
path = 'my_data.feather'
df = feather.read_dataframe(path)
```

Other languages, such as Julia or Scala (for Spark users), can read and write Feather files without knowledge of details of Python or R

1.3 RootTreeToR

- RootTreeToR [1] allows users to import ROOT [2] data directly into R
- Written by Adam Lyon (Fermilab), presented at useR! 2007
- cdcv.sfn.gov/redmine/projects/roottreator
- Requires ROOT to be installed, but no need to run ROOT

```
# Open and load ROOT tree:
rt <- openRootChain("TreeName", "FileName")
N <- nEntries(rt) # number of rows of data
# Names of branches:
branches <- RootTreeToR::getNames(rt)
# Read in a subset of branches (varsList), M rows:
df <- toR(rt, varsList, nEntries=M)
# Use writeDFToRoot to write a data.frame to ROOT
```

1.4 root_numpy

root_numpy is a Python extension module that provides an efficient interface between ROOT and NumPy

```
import ROOT
from root_numpy import root2array, root2rec, tree2rec
from root_numpy.testdata import get_filepath
filename = get_filepath('test.root')
# Convert a TTree in a ROOT file into a NumPy
# structured array:
arr = root2array(filename, 'tree')
# Convert a TTree in a ROOT file into a NumPy
# record array:
rec = root2rec(filename, 'tree')
# Get the TTree from the ROOT file:
rfile = ROOT.TFile(filename)
intree = rfile.Get('tree')
# Convert the TTree into a NumPy record array:
rec = tree2rec(intree)
```

2 Language-specific approaches

2.1 rPython (using Python from R)

```
library(rPython)
python.call( "len", 1:3 )
a <- 1:4
b <- 5:8
python.exec( "def concat(a,b): return a+b" )
python.call( "concat", a, b)
python.assign( "a", "hola hola" )
python.method.call( "a", "split", " " )
```

2.2 rpy2 (using R from Python)

```
from rpy2.robjects.lib.dplyr import (filter,
                                     mutate,
                                     group_by,
                                     summarize)

dataf = (DataFrame(mtcars) >>
         filter('gear>3') >>
         mutate(powertoweight='hp*36/wt') >>
         group_by('gear') >>
         summarize(mean_ptw='mean(powertoweight)'))

dataf
```

2.3 Others

- Calling Python from Julia: PyCall
- Calling R from Julia: RCall
- Calling Python from Excel: xlwings
- Calling Python from Ruby: Rupy
- ...

3 Notebook approach

3.1 Jupyter Notebook

- Each notebook has one main language
- Cells can contain other languages through “magic”
- For example: ggplot2 in Jupyter Notebook
- ROOTbooks: Notebooks running ROOT Jupyter kernel
- Assume that most people here have already seen these

3.2 Beaker Notebook

- Each notebook can contain any language
- Many languages are supported

- Auto translation of data (copied)
- <http://beakernotebook.com/>



Figure 1: “A universal translator for data scientists”

4 Reproducible research

4.1 Typical research pipeline

1. Measured data $\xrightarrow{\text{processing code}}$ analysis data
2. Analysis data $\xrightarrow{\text{analysis code}}$ computational results
3. Computational results $\xrightarrow{\text{presentation code}}$ plots, tables, numbers
4. Plots, tables, numbers + text $\xrightarrow{\text{manual editing}}$ **ARTICLE**
5. (Maybe) put the data and code on the web somewhere

4.2 Challenges with this approach

- Onus is on researchers to make their data and code available
- Authors may need to undertake considerable effort to put data and code on the web
- Readers must download data and code individually and piece together which data go with which code sections, etc.
- Typically, authors just put stuff on the web
- Readers just download the data and try to figure it out, piece together the software and run it
- Authors/readers must manually interact with websites
- There is no single document to integrate data analysis with textual representations; *i.e.* data, code, and text are not linked
- *Your experiment may impose restrictions on what kind of data can be shared publicly. Nevertheless, making an analysis reproducible benefits both your collaboration colleagues and your future self!*

4.3 Literate Statistical Programming

- Original idea comes from Donald Knuth
- An article is a stream of **text** and **code**
- Analysis code is divided into test and code “chunks”
- Presentation code formats results (tables, figures, *etc.*)
- Article text explains what is going on

Table 1: Effectiveness of Insect Sprays. The mean counts of insects in agricultural experimental units treated with different insecticides.

Count	**Spray**
A	14.500000
B	15.333333
C	2.083333
D	4.916667
E	3.500000
F	16.666667

- Literate programs are *weaved* to produce human-readable documents and *tangled* to produce machine-readable documents

4.4 knitr

- **knitr** [3] [4] [5] is a system for literate (statistical) programming
- Uses R as the programming language, but others are allowed
- Benefits: text and code are all in one place, results automatically updated to reflect external changes, *workflow supports reproducibility from the outset*
 - Almost no extra effort required by the researcher after the article has been written
- Like Notebooks, but output is publication quality, and you get the \LaTeX also as a freebie

4.5 R Markdown

This is an R Markdown [6] presentation. Markdown is a simple formatting syntax for authoring HTML, PDF, and MS Word documents. For more details on using R Markdown see <http://rmarkdown.rstudio.com>

It looks like \LaTeX , or to be more precise, it's a **Beamer** presentation, but this slide was created like:

```
## R Markdown
```

```
This is an R Markdown presentation. Markdown is a simple
formatting syntax for authoring HTML, PDF, and MS Word
documents. For more details on using R Markdown see
<http://rmarkdown.rstudio.com>
```

```
It looks like \LaTeX, or to be more precise, it's a
**Beamer** presentation, but this slide was created like:
```

4.6 R Markdown (*continued*)

The Markdown syntax has some enhancements. For example, you can include \LaTeX equations, like this:

$$i\hbar\gamma^\mu\partial_\mu\psi - mc\psi = 0 \tag{1}$$

We can also add tables and figures, just as we would do in \LaTeX :

4.7 Example code chunk

A code chunk using R is defined like this:

```
```{r}
print('hello world!')
```
```

A code chunk using some other execution engine is defined like this:


```
```{<execution-engine>}
print('hello world!')
```
```

For example:

```
```{python}
x = 'hello, python world!'
print(x.split(' '))
```
```

4.8 Example R code chunk

```
require(gapminder, quietly = TRUE)
require(ggplot2, quietly = TRUE)
ggplot(
  data = gapminder, aes(x = lifeExp, y = gdpPercap)) +
  geom_point(
    aes(color = continent, shape = continent)) +
  scale_y_log10()
```



4.9 Code execution engines

- Although R is the default code execution engine and a first-class citizen in the knitr system, many other code execution engines are available; this is the current list: *awk, bash, coffee, gawk, groovy, haskell, lein, mysql, node, octave, perl, psql, python, Rscript, ruby, sas, scala, sed, sh, stata, zsh, highlight, Rcpp, tikz, dot, c, fortran, fortran95, asy, cat, asis, stan, block, block2, js, css, sql, go*
- I've already shown examples of R and Python code chunks earlier in this talk
- Except for R, all chunks are executed in separate sessions, so the variables cannot be directly shared. If we want to make use of objects created in previous chunks, we usually have to write them to files (as side effects). For the bash engine, we can use `Sys.setenv()` to export variables from R to bash.
- Code chunks can be read from external files, in the event that you don't want to put everything in a monster-length markdown document
 - Use **bookdown** for books, theses and long reports

4.10 FORTRAN

Define a subroutine:

```
C Fortran test
  subroutine fexp(n, x)
    double precision x
C output
    integer n, i
C input value
    do 10 i=1,n
      x=dexp(dcos(dsin(dble(float(i))))))
10 continue
  return
end
```

4.11 FORTRAN (*continued*)

Call it from R:

```
# Call the function from R:
res = .Fortran("fexp", n=100000L, x=0)
str(res)
## List of 2
## $ n: int 100000
## $ x: num 2.72
```

4.12 Haskell

```
[x | x <- [1..10], odd x]

## [1,3,5,7,9]
```

4.13 C

C code is compiled automatically:

```
void square(double *x) {
    *x = *x * *x;
}
```

```
## gcc -std=gnu99 -I/usr/share/R/include -DNDEBUG      -fpic  -g -O2 -fstack-protector --param=ssp-buff
## gcc -std=gnu99 -shared -L/usr/lib/R/lib -Wl,-Bsymbolic-functions -Wl,-z,relro -o c59986f963dec.so c5
```

Test the square() function:

```
.C('square', 9)
```

```
## [[1]]
## [1] 81
```

4.14 Rcpp

The C++ code is compiled through the **Rcpp** package; here we show how we can write a function in C++ that can be called from R:

```
#include <Rcpp.h>
using namespace Rcpp;

// [[Rcpp::export]]
NumericVector timesTwo(NumericVector x) {
    return x * 2;
}
```

```
# In R:
print(timesTwo(3.1415926))
```

```
## [1] 6.283185
```

4.15 cat

A special engine `cat` can be used to save the content of a code chunk to a file using the `cat()` function, defined like this:

```
```${cat engine.opts=list(file='source.cxx')}
// Some code here...
```
```

4.16 BASH

```
echo hello bash!
echo 'a b c' | sed 's/ /\|/g'
# We can also compile the source.cxx file that we
# created in the previous code chunk; maybe compile
# against ROOT libraries, etc.
# Then run the application
```

```
## hello bash!
## a|b|c
```

We can run anything that we can run from the command line, including stuff we built in previous code chunks

4.17 Data exchange

Since the Python engine executes code in an external process, exchanging data between R chunks and Python chunks is done via the file system. If you are exchanging data frames, you can use the **Feather** package for very high performance transfer of even large data frames between Python and R:

```
import pandas
import feather
# Read flights data and select flights to O'Hare
flights = pandas.read_csv("flights.csv")
flights = flights[flights['dest'] == "ORD"]
# Select carrier and delay columns and drop rows with
# missing values
flights = flights[['carrier', 'dep_delay', 'arr_delay']]
flights = flights.dropna()
print flights.head(10)
# Write to feather file for reading from R
feather.write_dataframe(flights, "flights.feather")
```

4.18 Data exchange (*continued*)

```
library(feather)
library(ggplot2)

# Read from feather and plot
flights <- read_feather("flights.feather")
ggplot(flights, aes(carrier, arr_delay)) +
  geom_point() +
  geom_jitter()
```

4.19 Knitron

- The **knitron** package, available on GitHub but not yet on CRAN, is intended to allow users the ability to use IPython/Jupyter and **matplotlib** in R Markdown code chunks and render them with **knitr**
- According to the author's website, "Knitron works by lazily starting a global IPython kernel the first time a code chunk gets rendered by knitr and this kernel is reused for all consecutive chunks. This way all the computation done in any previous chunk is available in the current chunk, providing R-like behaviour for Python".

4.20 Knitron example

```
from mpl_toolkits.mplot3d import Axes3D
import matplotlib.pyplot as plt
import numpy as np
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
for c, z in zip(['r', 'g', 'b', 'y'], [30, 20, 10, 0]):
    xs = np.arange(20)
    ys = np.random.rand(20)
    cs = [c] * len(xs)
```

```

cs[0] = 'c'
ax.bar(xs, ys, zs=z, zdir='y', color=cs, alpha=0.8)
ax.set_xlabel('X')
ax.set_ylabel('Y')
ax.set_zlabel('Z')
plt.savefig("myplot.png")

```

GLib-GIO-Message: Using the 'memory' GSettings backend. Your settings will not be saved or shared w

4.21 Knitron example (*continued*)

Then plot the figure using a standard knitr mechanism, for example:

```

cat("![My matplotlib plot.](myplot.png)")

```




Figure 2: My matplotlib plot.

4.22 ROOT

- In the R tutorial that I gave at the IML Workshop last month, I spoke the possibility of using code execution engines other than R, but noted then that there is no way to use ROOT as a code execution engine

- Well, now there is!
- I have written an interface to allow ROOT to be used as a code execution engine
 - Simple proof-of-concept at the moment; nothing too sophisticated at this stage

4.23 ROOT code execution engine example

```

// Builds a graph with errors, displays it and saves it
// as image. First, include some header files (within,
// CINT these will be ignored).

#include "TCanvas.h"
#include "TROOT.h"
#include "TGraphErrors.h"
#include "TF1.h"
#include "TLegend.h"
#include "TArrow.h"
#include "TLatex.h"

void macro1(){
    // The values and the errors on the Y axis
    const int n_points=10;
    double x_vals[n_points]=
        {1,2,3,4,5,6,7,8,9,10};
    double y_vals[n_points]=
        {6,12,14,20,22,24,35,45,44,53};
    double y_errs[n_points]=
        {5,5,4.7,4.5,4.2,5.1,2.9,4.1,4.8,5.43};

    // Instance of the graph
    TGraphErrors graph(n_points,x_vals,y_vals,NULL,y_errs);
    graph.SetTitle("Measurement XYZ;lenght [cm];Arb.Units");

    // Make the plot looks better
    graph.SetMarkerStyle(kOpenCircle);
    graph.SetMarkerColor(kBlue);
    graph.SetLineColor(kBlue);

    // The canvas on which we'll draw the graph
    TCanvas* mycanvas = new TCanvas();

    // Draw the graph !
    graph.DrawClone("APE");

    // Define a linear function
    TF1 f("Linear law","[0]+x*[1]",.5,10.5);
    // Let's make the function line nicer
    f.SetLineColor(kRed); f.SetLineStyle(2);
    // Fit it to the graph and draw it
    graph.Fit(&f);
    f.DrawClone("Same");

    // Build and Draw a legend

```

```

TLegend leg(.1,.7,.3,.9,"Lab. Lesson 1");
leg.SetFillColor(0);
graph.SetFillColor(0);
leg.AddEntry(&graph,"Exp. Points");
leg.AddEntry(&f,"Th. Law");
leg.DrawClone("Same");

// Draw an arrow on the canvas
TArrow arrow(8,8,6.2,23,0.02,"|>");
arrow.SetLineWidth(2);
arrow.DrawClone();

// Add some text to the plot
TLatex text(8.2,7.5,"#splitline{Maximum}{Deviation}");
text.DrawClone();

mycanvas->Print("graph_with_law.pdf");
}

```

4.24

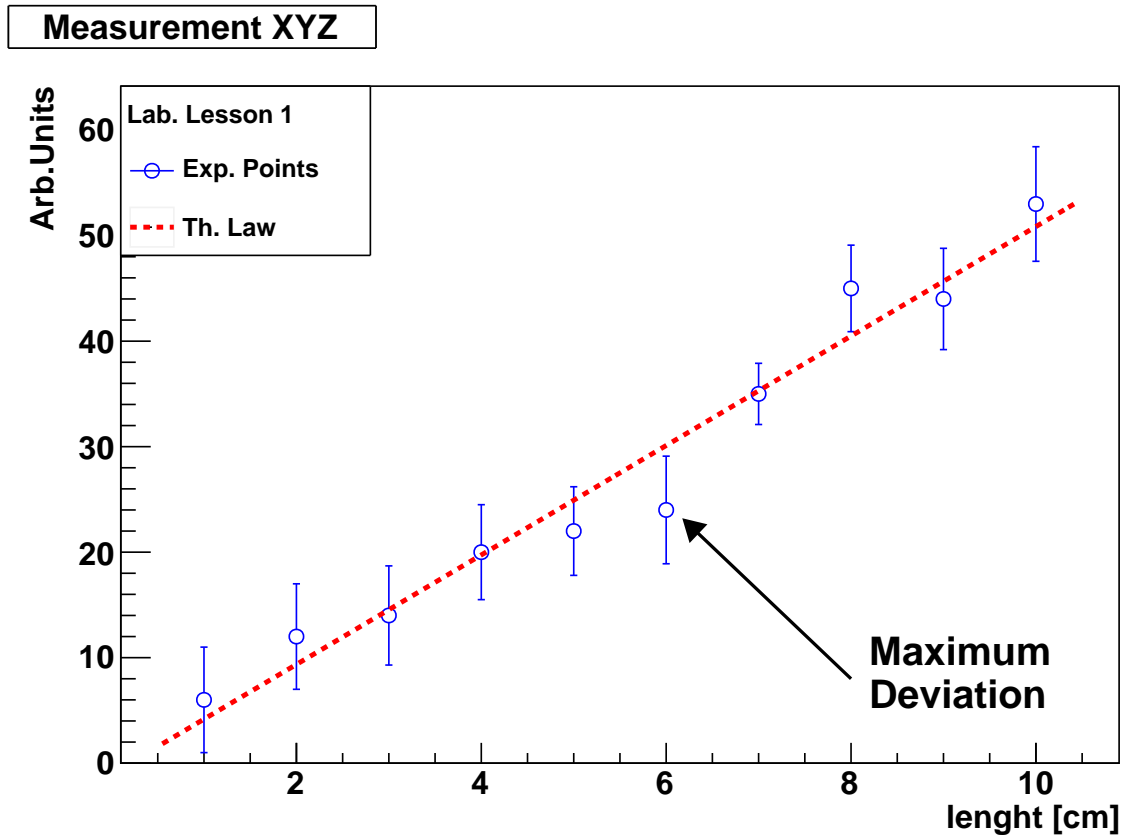


Figure 3: This is the graph

4.25 Messages to take away

- It is already possible to write a reproducible analysis in your favourite programming language and have your paper rendered fit for publication
- You can mix and match programming languages
- There are ways to exchange data between code chunks that are using different languages
 - ROOT, Python, R, ...
- You can embed ROOT code in a reproducible analysis

4.26 Acknowledgements

This work is supported by NKFIH OTKA grant K120660.

4.27 References

Back to introduction: 0.1.

References

- [1] Adam Lyon. *RootTreeToR*. R package version 4.1. URL: <https://github.com/lyonsquark/RootTreeToR>.
- [2] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth.*, A389:81–86, 1997. See also <http://root.cern.ch>. doi:10.1016/S0168-9002(97)00048-X.
- [3] Yihui Xie. *knitr: A General-Purpose Package for Dynamic Report Generation in R*, 2016. R package version 1.15.1. URL: <http://yihui.name/knitr/>.
- [4] Yihui Xie. *Dynamic Documents with R and knitr*. Chapman and Hall/CRC, Boca Raton, Florida, 2nd edition, 2015. ISBN 978-1498716963. URL: <http://yihui.name/knitr/>.
- [5] Yihui Xie. knitr: A comprehensive tool for reproducible research in R. In Victoria Stodden, Friedrich Leisch, and Roger D. Peng, editors, *Implementing Reproducible Computational Research*. Chapman and Hall/CRC, 2014. ISBN 978-1466561595. URL: <http://www.crcpress.com/product/isbn/9781466561595>.
- [6] JJ Allaire, Joe Cheng, Yihui Xie, Jonathan McPherson, Winston Chang, Jeff Allen, Hadley Wickham, Aron Atkins, Rob Hyndman, and Ruben Arslan. *rmarkdown: Dynamic Documents for R*, 2017. R package version 1.4. URL: <https://CRAN.R-project.org/package=rmarkdown>.