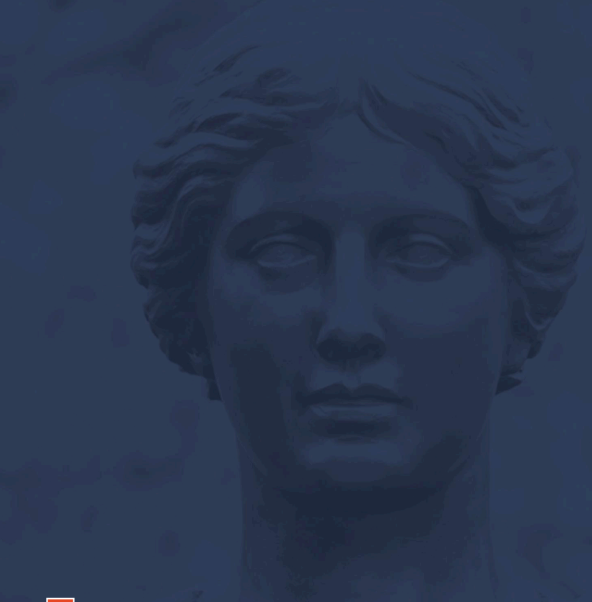

R for Machine Learning

An Brief Overview with Examples

Dirk Eddelbuettel

6th COST Conference on AI in Industry and Finance

9 September 2021, Zurich University of Applied Sciences (ZHAW), CH

https://dirk.eddelbuettel.com/papers/zhaw_cost_sep2021.pdf

ILLINOIS

Outline

Today's Talk

- R as "interface"
- Rcpp as glue
- Machine Learning *Interfacing* Examples

Quick Bio: My day job

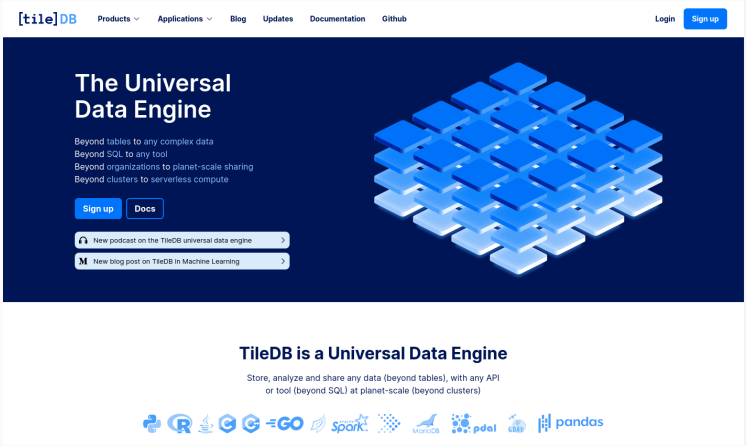

Serverless, Planet-scale

Any data, any size

Local, or multi-cloud

Via an API into many languages and applications

Talk to us

Yes, we're hiring

Quick Bio: My other jobs

Academic

- (Adjunct) Clinical Professor, University of Illinois
  - teaching STAT 447, a Data Science Programming Methods class

Open Source

- Debian developer
  - since 1995, currently maintaining about 170+ packages
- R contributor / package author
  - since 2002, author / maintainer of 60+ CRAN packages, R Foundation Board Member
- Rocker Project co-founder
  - Docker for R, including official 'r-base' image

Why R: View from Academia

Computer-Age Statistical Inference

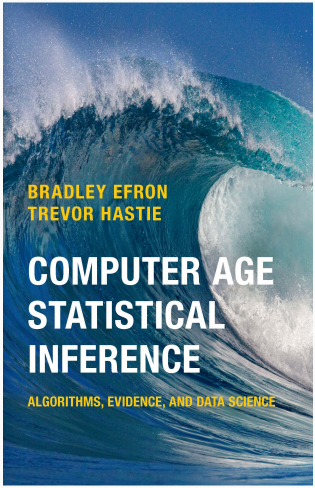

Almost all topics in twenty-first-century statistics are now computer-dependent [...]

Here and in all our examples we are employing the language R, itself one of the key developments in computer-based statistical methodology.

Efron and Hastie, 2016
pages xv and 6 (footnote 3)

Another View: James et al, ISLR (2nd ed)

[...] we have devoted a section within each chapter to R computer labs. In each lab, we walk the reader through a realistic application of the methods considered in that chapter. [...] We have used R because it is freely available and is powerful enough to implement all of the methods discussed in the book. It also has optional packages that can be downloaded to implement literally thousands of additional methods. Most importantly, R is the language of choice for academic statisticians, and new approaches often become available in R years before they are implemented in commercial packages.

Another View: James et al, ISLR (2nd ed)

[...] we have devoted a section within each chapter to R computer labs. In each lab, we walk the reader through a realistic application of the methods considered in that chapter. [...] We have used R because it is freely available and is powerful enough to implement all of the methods discussed in the book. It also has optional packages that can be downloaded to implement literally thousands of additional methods. Most importantly, R is the language of choice for academic statisticians, and new approaches often become available in R years before they are implemented in commercial packages.

Computational Statistics in Practice

- Statistics is now computational (Efron & Hastie, 2016)
- Within (computational) statistics, reigning tool is R (James *et al*, ISLR, 2021)
- Given R, Rcpp key for two angles:
  - *Performance* always matters, ease of use a sweetspot
  - *"Extending R"* (Chambers, 2016)

Why R: View from Practitioners

Why R? Pat Burn's View

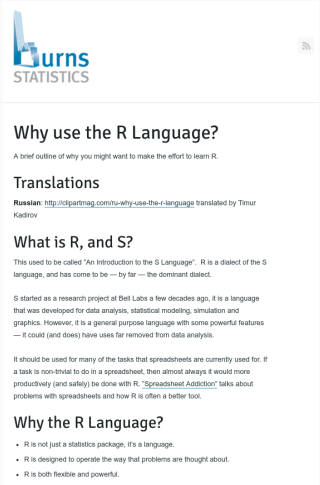

Burns
STATISTICS

Why use the R Language?

A brief outline of why you might want to make the effort to learn R.

Translations

Russian: http://clipartmag.com/ru-why-use-the-r-language translated by Timur Kadirov

What is R, and S?

This used to be called "An Introduction to the S Language". R is a dialect of the S language, and has come to be — by far — the dominant dialect.

S started as a research project at Bell Labs a few decades ago, it is a language that was developed for data analysis, statistical modelling, simulation and graphics. However, it is a general purpose language with some powerful features — it could (and does) have uses far removed from data analysis.

It should be used for many of the tasks that spreadsheets are currently used for. If a task is non-trivial to do in a spreadsheet, then almost always it would more productively (and safely) be done with R. "Spreadsheet Addiction" talks about problems with spreadsheets and how R is often a better tool.

Why the R Language?

- R is not just a statistics package, it's a language.
- R is designed to operate the way that problems are thought about.
- R is both flexible and powerful.

Why the R Language?

Screen shot on the left part of short essay at Burns-Stat

His site has more truly excellent (and free) writings.

The (much longer) R Inferno (free pdf, also paperback) is highly recommended.

Why R? Pat Burn's View

Why the R Language?

- R is not just a statistics package, it's a language.
- R is designed to operate the way that problems are thought about.
- R is both flexible and powerful.

Source: https://www.burns-stat.com/documents/tutorials/why-use-the-r-language/

Why R? Pat Burn's View

Why R for data analysis?

R is not the only language that can be used for data analysis.

Why R rather than another? Here is a list:

- interactive language
- data structures
- graphics
- missing values
- functions as first class objects
- packages
- community

Source: https://www.burns-stat.com/documents/tutorials/why-use-the-r-language/

Why R: Programming with Data

R as an Extensible Environment

- As R users we know that R can
  - ingest data in many formats from many sources
  - aggregate, slice, dice, summarize, ...
  - visualize in many forms, ...
  - model in just about any way
  - report in many useful and scriptable forms
- It has become central for programming with data
- Sometimes we want to extend it further than R code goes

R as central point

R as central point

From any one of

- csv
- txt
- xlsx
- xml, json, ...
- web scraping, ...
- hdf5, netcdf, ...
- sas, stata, spss, ...
- various SQL + NOSQL DBs
- various binary protocols

via

into any one of

- txt
- html
- latex and pdf
- html and js
- word
- shiny
- most graphics formats
- other dashboards
- web frontends

Why R: Historical Perspective

R as 'The Interface'

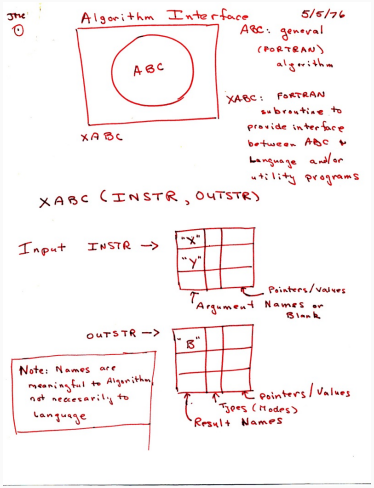

A design sketch called 'The Interface'

AT&T Research lab meeting notes

Describes an outer 'user interface' layer to core Fortran algorithms

Key idea of abstracting away inner details giving higher-level more accessible view for user / analyst

Lead to "The Interface"

Which became S which lead to R

Source: John Chambers, personal communication

Why R? : *Programming with Data* from 1977 to 2016

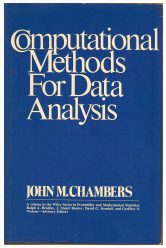

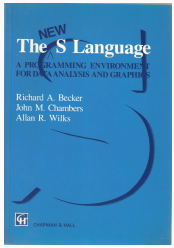

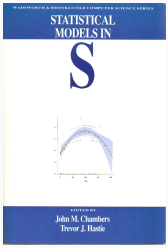

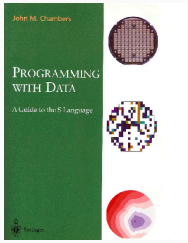

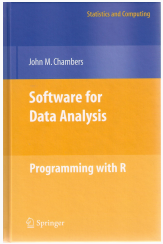

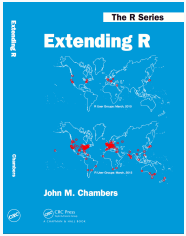

Thanks to John Chambers for high-resolution cover images. The publication years are, respectively, 1977, 1988, 1992, 1998, 2008 and 2016.

Chambers (2008)

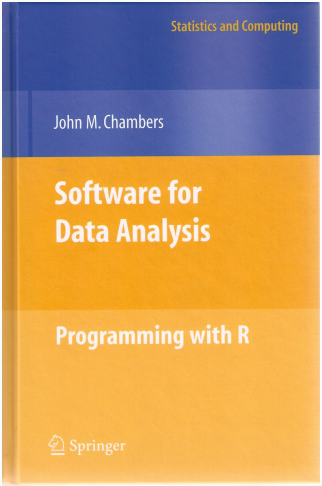

Software For Data Analysis

Chapters 10 and 11 devoted to *Interfaces I: C and Fortran* and *Interfaces II: Other Systems*.

Chambers (2016)

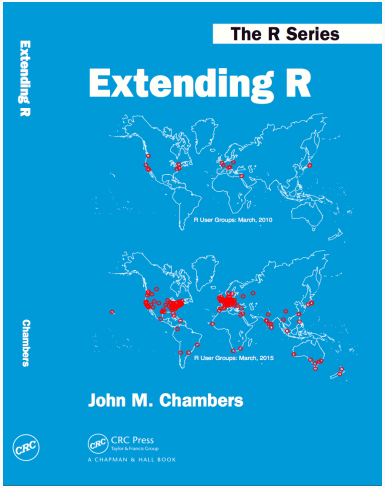

Extending R

Object: Everything that exists in R is an object

Function: Everything happens in R is a function call

Interface: Interfaces to other software are part of R

Chambers (2016)

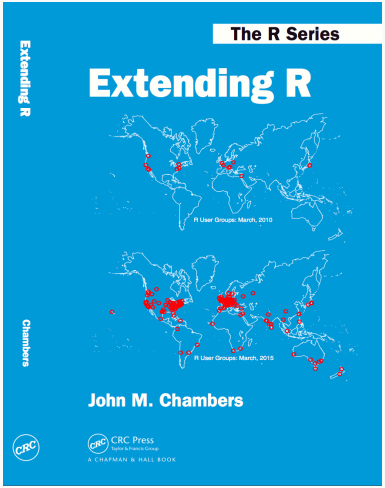

Extending R, Chapter 4

The fundamental lesson about programming in the large is that requires a correspondingly broad and flexible response. In particular, no single language or software system os likely to be ideal for all aspects. Interfacing multiple systems is the essence. Part IV explores the design of of interfaces from R.

So ...

From a system called ‘interface’ to a language where interfaces are a natural part

Why Rcpp?

A good fit, it turns out

- A good part of R is written in C (besides R and Fortran code)
- The principle interface to external code is a function `.Call()`
- It takes one or more of the high-level data structures R uses
- ... and returns one. Formally:

```
SEXP .Call(SEXP a, SEXP b, ...)
```

A good fit, it turns out (cont.)

- An SEXP (or S-Expression Pointer) is used for *everything*
- (An older C trick approximating object-oriented programming)
- We can ignore the details but retain that
  - everything in R is a SEXP, and the SEXP is self-describing
  - can matrix, vector, list, function, ... for a total of 27 types
- The key thing for Rcpp is that via C++ features we can map
  - each of these (few) SEXP types to a specific C++ class representing that type
  - and the conversion is automated back and forth

Other good reasons

- It is *fast* – compiled C++ is hard to beat in other languages
  - (That said, you can *of course* write bad and slow code....)
- It is *very general* and widely used
  - *many libraries*
  - many tools
- It is fairly universal:
  - just about anything will have C interface so C++ can play
  - just about any platform / OS will have it

Key Features

- (Fairly) Easy to learn as it really does not have to be that complicated
  there are numerous examples, tutorials, books, ...
- Easy to use as it avoids build and operating system build complexities
  simply by relying the robust build infrastrucure R itself has
- Expressive as it allows for *vectorised* C++ using *Rcpp Sugar*
- Seamless access to all R objects: vector, matrix, list, S3/S4/RefClass, Function, ...
- Fast as Rcpp excels at tasks where R struggles: loops, function calls, ...
- Extensions facilitates access to external libraries directly or via eg *Rcpp modules*

Rcpp Speed Illustration

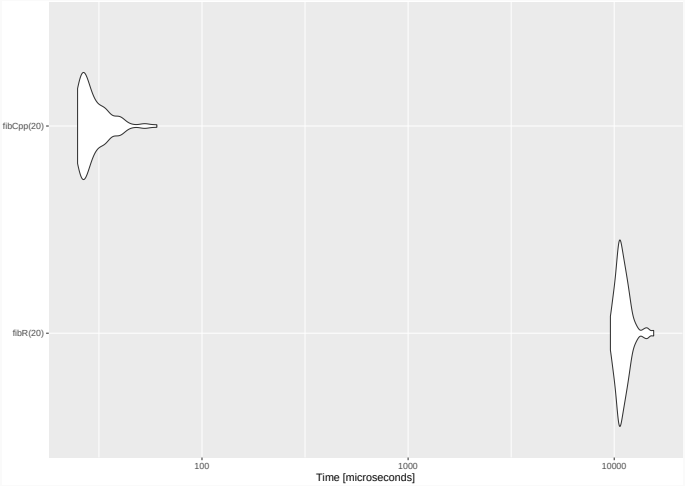

Benchmark on Fibonacci(20) between C++ and R – note the log scale!

Growth

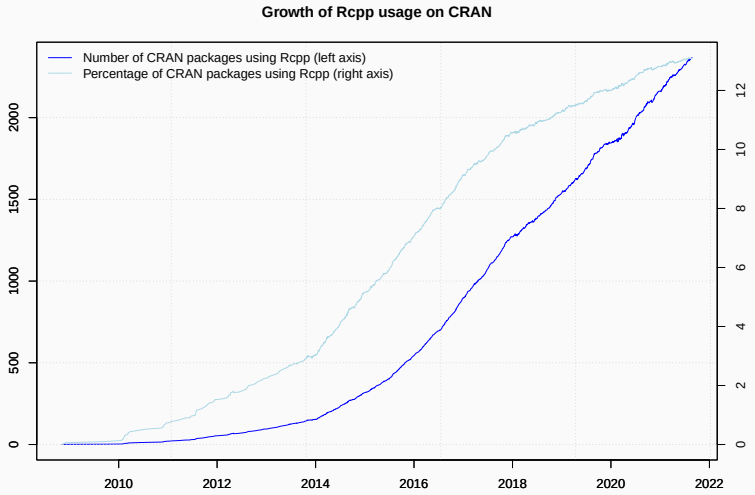

Source: Our calculations. Data current as of August 29, 2021.

Users on Core Repositories

Rcpp is currently used by

- 2367 CRAN packages
- 231 BioConductor packages
- an unknown (but “large”) number of GitHub projects
- over 50 million downloads (from the RStudio operated CRAN mirrors alone)

Pagerank

```
suppressMessages(library(utils))
library(pagerank)   # cf github.com/andrie/pagerank

cran <- "http://cloud.r-project.org"
pr <- compute_pagerank(cran)
round(100*pr[1:5], 3)
```

```
##     Rcpp  ggplot2    dplyr     MASS magrittr
##    2.790    1.483    1.216    1.176    0.781
```

Pagerank

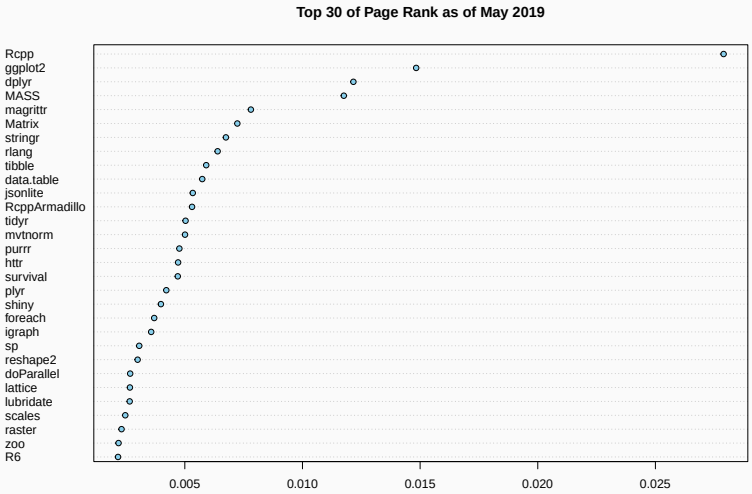
Top 30 of Page Rank as of May 2019
Rcpp
ggplot2
dplyr
MASS
magrittr
Matrix
stringr
rlang
tibble
data.table
jsonlite
RcppArmadillo
tidyr
mvtnorm
purrr
httr
survival
plyr
shiny
foreach
igraph
sp
reshape2
doParallel
lattice
lubridate
scales
raster
zoo
R6
0.005
0.010
0.015
0.020
0.025

Percentage of Compiled Packages

```
db <- tools::CRAN_package_db()
nTot <- nrow(db)
## all direct Rcpp reverse depends, ie packages using Rcpp
nRcpp <- length(tools::dependsOnPkgs("Rcpp", recursive=FALSE,installed=db))
nCompiled <- table(db[, "NeedsCompilation"])["yes"]
propTot <- nRcpp / nTot * 100
propComp <- nRcpp / nCompiled * 100
data.frame(tot=nTot, totRcpp = nRcpp, totCompiled = nCompiled,
           RcppPctTot = propTot, RcppPctOfCompiled = propComp)
```

```
##     tot totRcpp totCompiled RcppPctTot RcppPctOfCompiled
## 1 18112    2367        4338    13.0687           54.5643
```

Machine Learning via R(cpp): Three Examples

Vowpal Wabbit

Vowpal Wabbit

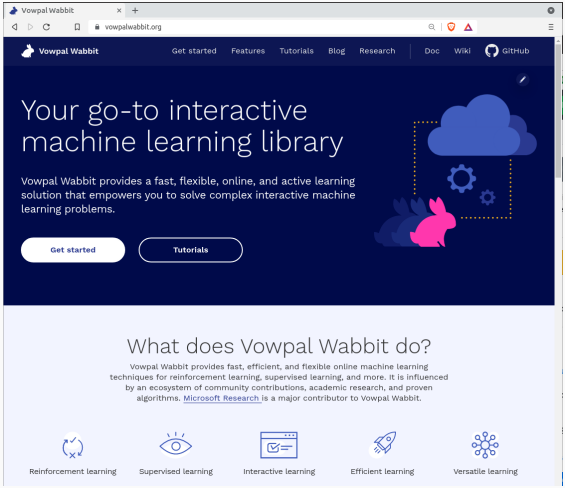

Well-known, established project

Initially at Yahoo! Research, now at Microsoft Research (same NY lab)

Lead by John Langford

Fast online learner, popular as a classifier

Now much more included distributed setups

Many papers, tutorials, ...

C++ core

Vowpal Wabbit in R

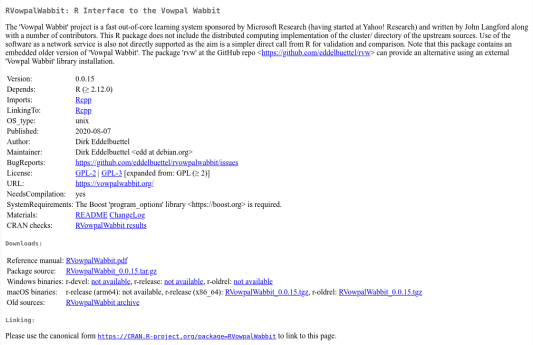

RVowpalWabbit: R Interface to the Vowpal Wabbit

The 'Vowpal Wabbit' project is a fast out-of-core learning system sponsored by Microsoft Research (having started at Yahoo! Research) and written by John Langford along with a number of contributors. This R package does not include the distributed computing implementation of the cluster/ directory of the upstream sources. Use of the software as a network service is also not directly supported as the aim is a simpler direct call from R for validation and comparison. Note that this package contains an embedded older version of 'Vowpal Wabbit'. The package 'rvw' at the GitHub repo <https://github.com/eddelbuettel/rvw> can provide an alternative using an external 'Vowpal Wabbit' library installation.

Version:	0.0.15
Depends:	R (≥ 2.12.0)
Imports:	Rcpp
LinkingTo:	Rcpp
OS_type:	unix
Published:	2020-08-07
Author:	Dirk Eddelbuettel
Maintainer:	Dirk Eddelbuettel <edd at debian.org>
BugReports:	https://github.com/eddelbuettel/rvowpalwabbit/issues
License:	GPL-2 | GPL-3 [expanded from: GPL (≥ 2)]
URL:	https://vowpalwabbit.org/
NeedsCompilation:	yes
SystemRequirements:	The Boost 'program_options' library <https://boost.org> is required.
Materials:	README ChangeLog
CRAN checks:	RVowpalWabbit results

Downloads:

Reference manual:	RVowpalWabbit.pdf
Package source:	RVowpalWabbit_0.0.15.tar.gz
Windows binaries:	r-devel: not available, r-release: not available, r-oldrel: not available
macOS binaries:	r-release (arm64): not available, r-release (x86_64): RVowpalWabbit_0.0.15.tgz, r-oldrel: RVowpalWabbit_0.0.15.tgz
Old sources:	RVowpalWabbit archive

Linking:

Please use the canonical form https://CRAN.R-project.org/package=RVowpalWabbit to link to this page.

Several Phases

Initial package **RVowpalWabbit** on CRAN since Sep 2011 (!!), took a copy of the library (plus some editing)

At the time ‘VW’ library had an awkward / unusual build

Taking a copy is quick and simple … and can leads to stale code

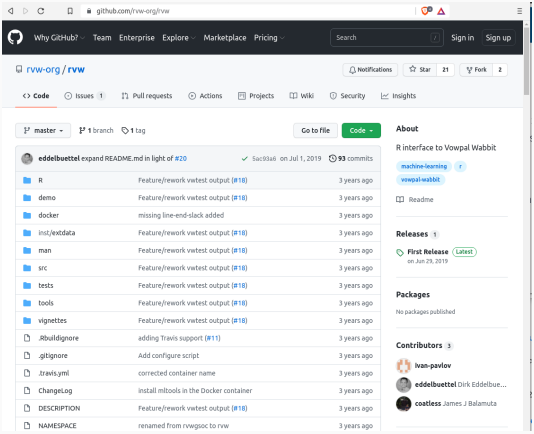

Second Version

Google Summer of Code project by Ivan Pavlov (supervised by DE + JJB) building on earlier work by Selim Raboudi

Used external ‘VW’ library (which now build better)

Taking an external library is more elegant … but harder on the user who needs to install it – so for example *not* at CRAN

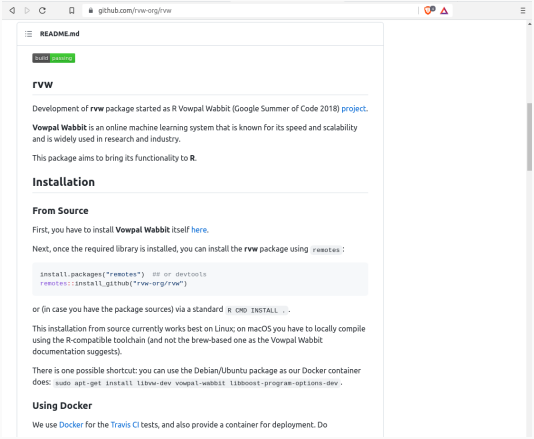

Second Version

It is however in a pretty good state

Contributions would be welcome, a project worth revisiting

As a fallback and alternative build approach, we could internalise building the library in order to get the package onto CRAN

Corels

Corels

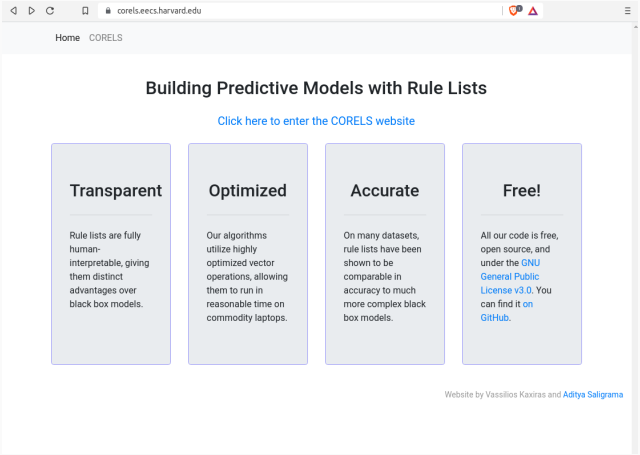

Very nice (smaller) project out of Harvard/Duke/UBC

Well documented with a couple of papers

Lead by Margo Seltzer and Cynthia Rudin

I wrote a simple wrapper package, and *documented my steps* in a what became a small paper

C++ core

Corels in R

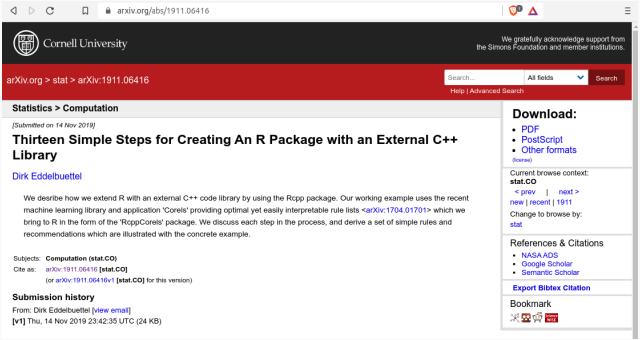

Little how-to tutorial as a side effect of wrapping Corels into RcppCorels

Paper now also a vignette in the Rcpp package

RcppCorels intgegrated with the rest of Corels at `https://github.com/corels`

Corels in R

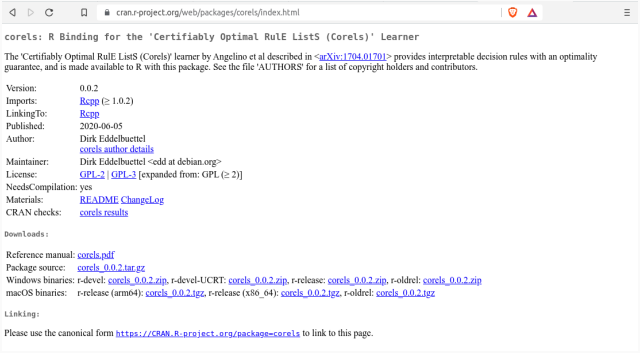
cran.r-project.org/web/packages/corels/index.html

corels: R Binding for the 'Certifiably Optimal RulE ListS (Corels)' Learner

The 'Certifiably Optimal RulE ListS (Corels)' learner by Angelino et al described in <arXiv:1704.01701> provides interpretable decision rules with an optimality guarantee, and is made available to R with this package. See the file 'AUTHORS' for a list of copyright holders and contributors.

Version:	0.0.2
Imports:	Rcpp (≥ 1.0.2)
LinkingTo:	Rcpp
Published:	2020-06-05
Author:	Dirk Eddelbuettel corels author details
Maintainer:	Dirk Eddelbuettel <edd at debian.org>
License:	GPL-2 | GPL-3 [expanded from: GPL (≥ 2)]
NeedsCompilation:	yes
Materials:	README ChangeLog
CRAN checks:	corels results

Downloads:

Reference manual:	corels.pdf
Package source:	corels_0.0.2.tar.gz
Windows binaries:	r-devel: corels_0.0.2.zip, r-devel-UCRT: corels_0.0.2.zip, r-release: corels_0.0.2.zip, r-oldrel: corels_0.0.2.zip
macOS binaries:	r-release (arm64): corels_0.0.2.tgz, r-release (x86_64): corels_0.0.2.tgz, r-oldrel: corels_0.0.2.tgz

Linking:

Please use the canonical form https://CRAN.R-project.org/package=corels to link to this page.

Now also on CRAN as

`https://cran.r-project.org/package=corels`

(with the package name normalised to **`corels`**)

Some work remaining to integrate better with external data sets

MLPACK

MLPACK

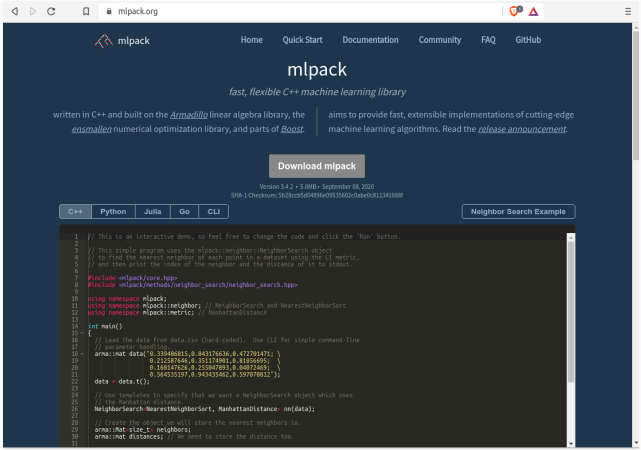

- Mature and established project
- Lead by Ryan Curtin with many contributors
- Very complete collection of algorithm
- Multiple language bindings

MLPACK

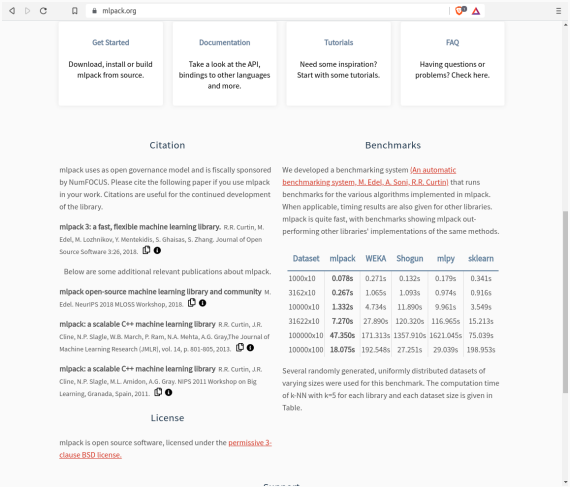

Many models implemented

(Run-time) Performance usually better than competitors in Java or Python

C++ core

MLPACK IN R


cloud.r-project.org/web/packages/mlpack/index.html

mlpack: 'Rcpp' Integration for the 'mlpack' Library

A fast, flexible machine learning library, written in C++, that aims to provide fast, extensible implementations of cutting-edge machine learning algorithms. See also Curtin et al. (2018) <doi:10.21105/joss.00726>.

Version: 3.4.2.1
Depends: R (≥ 4.0.0)
Imports: Rcpp (≥ 0.12.12)
LinkingTo: Rcpp, RcppArmadillo (≥ 0.8.400.0), BH (≥ 1.58), RcppEnsmallen (≥ 0.2.10.0)
Suggests: testthat (≥ 2.1.0)
Published: 2020-12-18
Author: Yashwant Singh Parihar [aut, cth, cph], Ryan Curtin [aut, cth, cph, cre], Dirk Eddelbuettel [aut, cth, cph], James Balamuta [aut, cth, cph], Bill March [cth, cph], Dongryeol Lee [cth, cph], Nishant Mehta [cth, cph], Parikshit Ram [cth, cph], James Cline [cth, cph], Sterling Peet [cth, cph], Matthew Amidon [cth, cph], Neil Slagle [cth, cph], Ajinkya Kale [cth, cph], Vlad Grantcharov [cth, cph], Noah Kauffman [cth, cph], Rajendran Mohan [cth, cph], Trironk Kiatkungwanglai [cth, cph], Patrick Mason [cth, cph], Marcus Edel [cth, cph], Mudit Raj Gupta [cth, cph], Sumedh Ghaisas [cth, cph], Michael Fox [cth, cph], Ryan Birmingham [cth, cph], Siddharth Agrawal [cth, cph], Sahob Motiani [cth, cph], Yash Vadalia [cth, cph], Abhishek Laddha [cth, cph], Vahab Akbarzadeh [cth, cph], Andrew Wells [cth, cph], Zhihao Lou [cth, cph], Udit Saxena [cth, cph], Stephen Tu [cth, cph], Jaskaran Singh [cth, cph], Hritik Jain [cth, cph], Vladimir Glazachev [cth, cph], QiaoAn Chen [cth, cph], Janzen Brewer [cth, cph], Trang Dinh [cth, cph], Tham Ngoc Wei [cth, cph], Grzegorz Krajewski [cth, cph], Joseph Mariadasan [cth, cph], Pavel Zhigulin [cth, cph], Andy Fang [cth, cph], Barak Pearlmutter [cth, cph], Ivan Horn [cth, cph], Dhawal Arora [cth, cph], Alexander Lenoff [cth, cph], Palash Ahuja [cth, cph], Yannis Mentekidis [cth, cph], Ranjan Mondal [cth, cph], Mikhail Lozhnikov [cth, cph], Marcus Pividori [cth, cph], Kevin Kim [cth, cph], Nilay Jain [cth, cph], Peter Lehner [cth, cph], Anurag Kaundal [cth, cph], Ivan Georgiev [cth, cph], Shikhar Bhardwaj [cth, cph], Yashu Seth [cth, cph], Mike Izbicki [cth, cph], Sudhanshu Ranjan [cth, cph], Piyush Jaiswal [cth, cph], Dinesh Raj [cth, cph], Vivek Pal [cth, cph], Prasanna Patil [cth, cph], Lakshya Agrawal [cth, cph], Praveen Ch [cth, cph], Kirill Mishchenko [cth, cph], Abhinav Moudgil [cth, cph], Thyrix Yang [cth, cph], Sagar B Hathwar [cth, cph], Nishanth Hegde [cth, cph], Parminder Singh [cth, cph], Codeck [cth, cph], Franciszek Stokowacki [cth, cph], Samskeksya Chand [cth, cph], N Rajiv Vaidyanathan [cth, cph], Karthik Nighania [cth, cph], Eugene Freyman [cth, cph], Manish Kumar [cth, cph], Haritha Sreedharan Nair [cth, cph], Sourabh Varshney [cth, cph], Prajyal Dev [cth, cph], Nikhil Goel [cth, cph], Shikhar Jaiswal [cth, cph], B Karthesh Reddy [cth, cph], Atharva Khandait [cth, cph], Wenhao Huang [cth, cph], Roberto Hueso [cth, cph], Prabhat Sharma [cth, cph], Tan Jun An [cth, cph], Mohsin Jain [cth, cph], Manthan-K-Sheth [cth, cph], Namrata Mukhija [cth, cph], Conrad Conrad Sanderson [cth, cph], Thanasis Mattas [cth, cph], Shashank Shekhar [cth, cph], Yasmine Dumouchel [cth, cph], Girman Girman Lancioni [cth, cph], Arush Alsghari [cth, cph], Ayush Ayush Chamoli [cth, cph], Tommi Laivamaa [cth, cph], Kim SangYeon [cth, cph], Nitesz Shah [cth, cph], Toshal Agrawal [cth, cph], Dan Dan Timson [cth, cph], Miguel Cantero [cth, cph], Bishwa Karki [cth, cph], Mehul Kumar Nirala [cth, cph], Heet Sankesara [cth, cph], Jeffin Sam [cth, cph], Vikas S Shetty [cth, cph], Khairi Siddiqui [cth, cph], Tejaswi Tomar [cth, cph], Jai Agarwal [cth, cph], Ziyang Jiang [cth, cph], Rohit Kartik [cth, cph], Aditya Viki [cth, cph], Kartik Dutt [cth, cph], Sriram S K [cth, cph], Manoranjan Kumar Bharti (<Nakul Bharti>) [cth, cph], Saransh Tandon [cth, cph], Gaurav Singh [cth, cph], Lakshya Ojha [cth, cph], Bisakh Mondal [cth, cph], Benson Muite [cth, cph], Sarthak Bhardwaj [cth, cph], Aakash Kaushik [cth, cph], Anush Kini [cth, cph], Nippun Sharma [cth, cph], Rishabh Garg [cth, cph], Sudhakar Brar [cth, cph], Alex Nguyen [cth, cph]
Maintainer: Ryan Curtin <ryan at ratml.org>
BugReports: https://github.com/mlpack/mlpack/issues
License: BSD_3_clause + file LICENSE
URL: https://www.mlpack.org/doc/mlpack-3.4.2/r_documentation.html, https://github.com/mlpack/mlpack
NeedsCompilation: yes
SystemRequirements: A C++11 compiler. Versions 4.8.*, 4.9.* or later of GCC will be fine.
In views: MachineLearning
CRAN checks: mlpack results

Downloads:

Reference manual: mlpack.pdf
Package source: mlpack_3.4.2.1.tar.gz
Windows binaries: r-devel: mlpack_3.4.2.1.zip, r-devel-UCRT: mlpack_3.4.2.1.zip, r-release: mlpack_3.4.2.1.zip, r-oldrel: mlpack_3.4.2.1.zip
macOS binaries: r-release (arm64): mlpack_3.4.2.1.tgz, r-release (x86_64): mlpack_3.4.2.1.tgz, r-oldrel: mlpack_3.4.2.1.tgz
Old sources: mlpack archive

Reverse dependencies:

Reverse imports: text2map
Reverse suggests: pensim.hurt

Linking:

Please use the canonical form https://CRAN.R-project.org/package=mlpack to link to this page.

Several earlier attempts at creating a package stopped for lack of a (external) library (say at CRAN)

Excellent Google Summer of Code work by Yashwant Singh Parihar in 2020; R is now a first-class supported language

Also on CRAN as `https://cran.r-project.org/package=mlpack`

Summary

Key Points

- R excels as *data-centric language* invented by statisticians
- Interfaces to other systems are key part of its design
- Many promising libraries either are in C++, or have C++ interface
- The Rcpp smoothes the interaction with another (compiled) language
- We looked at three concrete examples:
  - Vowpal Wabbit
  - Corels
  - MLPACK

Thank you!

slides	https://dirk.eddelbuettel.com/presentations/
web	https://dirk.eddelbuettel.com/
mail	dirk@eddelbuettel.com
github	@eddelbuettel
twitter	@eddelbuettel