



Packaging Data Analytical Work Reproducibly Using R (and Friends)

Ben Marwick^a, Carl Boettiger^b, and Lincoln Mullen^c

^aUniversity of Washington, Seattle, WA; ^bUniversity of Wollongong, Wollongong, New South Wales; ^cUniversity of California, Berkeley, CA; ^dGeorge Mason University, Fairfax, VA

ABSTRACT

Computers are a central tool in the research process, enabling complex and large-scale data analysis. As computer-based research has increased in complexity, so have the challenges of ensuring that this research is reproducible. To address this challenge, we review the concept of the research compendium as a solution for providing a standard and easily recognizable way for organizing the digital materials of a research project to enable other researchers to inspect, reproduce, and extend the research. We investigate how the structure and tooling of software packages of the R programming language are being used to produce research compendia in a variety of disciplines. We also describe how software engineering tools and services are being used by researchers to streamline working with research compendia. Using real-world examples, we show how researchers can improve the reproducibility of their work using research compendia based on R packages and related tools.

ARTICLE HISTORY

Received May 2017
Revised August 2017

KEYWORDS

Computational science; Data science; Open source software; Reproducible research

1. Introduction

The purpose of this article is to show how the R package can be a suitable template for organising files into a research compendium to enhance the reproducibility of research. We first establish a working definition of reproducibility—a term with diverse meanings across different domains—and then describe the “what,” “why,” and “how” of using R packages as research compendia. We present real-world examples of research compendia of low, moderate, and high complexity to show how flexible the R package format is for this purpose. We conclude by reviewing some useful tools and templates that we have found to make it easy to create and share research compendia.

Long considered an axiom of science, the reproducibility of scientific research has recently come under scrutiny after some highly publicized failures to reproduce biomedical studies (Begley and Ellis 2012) and psychological experiments (Open Science Collaboration et al. 2015). This has resulted in extensive discussion of how to define reproducibility, and how to enhance it across diverse research contexts (Stodden et al. 2016; Goodman, Fanelli, and Ioannidis 2016; Marwick 2017). Here, we use the definition proposed by Claerbout (1992) and the National Science Foundation Subcommittee on Replicability in Science (Bollen et al. 2015): reproducing research is the calculation of quantitative scientific results by independent researchers using the original data and methods. Stodden (2014) further divided reproducibility into three dimensions: empirical (e.g., details of reagents, cell lines, sample identities, instrument settings), statistical (e.g., details of statistical tests, model parameters, threshold values, etc.) and computational (e.g., details about code, software, hardware, and implementation). Our focus here is on the computational dimension, which is important because it transcends many of the domain-specific

technical issues of empirical and statistical reproducibility. Indeed, the reproducibility of computational results has become an important consideration in fields such as digital history, digital humanities, and the social sciences (Graham, Milligan, and Weingart 2015, pp. 149,157; Allison 2016; Goldstone 2017).

Virtually all researchers use computers as a central tool in their workflow. However, our formal education rarely includes any training in how to organise our computer files to make it easy to reproduce results and share our entire analysis pipeline with others. Without clear instructions, many researchers struggle to avoid chaos in their file structures, and so are understandably reluctant to expose their workflow for others to see. This may be one of the reasons that so many requests for details about method, including requests for data and code, are turned down or go unanswered (Collberg and Proebsting 2016).

Scholarship will be strengthened if we are more open with research materials, yet we lack conventions and technical infrastructure for such openness. We see evidence of this need in the proliferation of commentary and how-to publications across different areas of science (Wilson et al. 2014; Stodden and Miguez 2014; Stanisic and Legrand 2014; LeVeque, Mitchell, and Stodden 2012; Donoho 2010; Sandve et al. 2013; Rokem, Marwick, and Staneva 2017; Munafo et al. 2017). There are also some forward-looking journal editors attempting to shift the norms of their research community by encouraging or requiring authors to make available their data and code files when submitting their manuscript for publication. For example, *Biostatistics* invites authors to submit materials for a reproducibility review after the article has been accepted (Peng 2009). Authors submitting to *Political Analysis* and the *Quarterly Journal of Political Science* are required to provide their code for review, with editors reproducing the reported analyses prior to

publication (Nosek et al. 2015). The *Journal of Cultural Analytics* requires authors to deposit code and dataset (Piper et al. 2017). A cross-discipline system of signaling open data and materials accompanying journal articles with badges is being used to incentivize openness in a dozen journals (Kidwell et al. 2016).

Our contribution to this drive to improve the openness and reproducibility of research is to show how R packages can be used as a research compendium for organising and sharing files. Although the R packaging system is traditionally a method for sharing statistical methods, we claim that R packages are suitable for use as research compendia that can help improve computational reproducibility. Our focus is on the R programming language because the R package structure is uniquely suitable to being easily adapted to solve problems of organising files and sharing them with other researchers. We describe how the conventional structure of R packages can be adapted for use as a research compendium, and illustrate this use with real-world examples.

We recognize that not all researchers are working in R, and that even those who are may not adopt our recommendations wholesale. Nevertheless, the principles of research compendia that we describe are generally applicable to projects that involve computation using any language, and many of our recommendations about specific software can be adopted individually.

2. What is a Research Compendium? Three Generic Principles

The goal of a research compendium is to provide a standard and easily recognizable way for organizing the digital materials of a project to enable others to inspect, reproduce, and extend the research. There are three generic principles that define research compendia, independent of particular software tools, and disciplinary contexts.

1. A research compendium should organize its files according to the prevailing conventions of the scholarly community, whether that be an academic discipline or a lab group. Following these conventions will help other people recognize the structure of the project, and also support tool building which takes advantage of the shared structure.
2. A research compendium should maintain a clear separation of data, method, and output, while unambiguously expressing the relationship between those three. In practice, this means data files must be separate from code files. This is important to let others easily identify how the original researcher operated on the data to generate the results. Keeping data and method separate treats the data as “read-only,” so that the original data are untouched and all modifications are transparently documented in the code. The output files should be considered as disposable, with a mindset that one can always easily regenerate the output using the code and data. The relationship between which code operates on which data in which order to produce which outputs must be specified as well. In his advice to industry data scientists, Ben Baumer’s article in this collection similarly highlights the importance of keeping data separate from the presentation of data, or research outputs.

3. A research compendium should specify the computational environment that was used for the original analysis. At its most basic, this could be a plain text file that includes a short list of the names and version numbers of the software and other critical tools used for the analysis. In more complex approaches, described below, the computational environment can be automatically preserved or reproduced as well.

Some of the earliest examples of these principles in action can be found in the work of Claerbout (1992) and his colleagues in the early 1990s. Claerbout introduced these practices in his Stanford geophysics research group (Buckheit and Donoho 1995). Other early examples of compendia appeared in econometrics (Koenker 1996), (Vinod 2001), epidemiology (Peng, Dominici, and Zeger 2006), and signal processing (Vandewalle, Kovacevic, and Vetterli 2009). Two influential papers are Gentleman’s (2005) article “Reproducible Research: A Bioinformatics Case Study” and Temple Lang’s article “Statistical Analyses and Reproducible Research” (Gentleman and Temple Lang 2004, 2007). The impact of these two papers is due to their clear description of an easy-to-use set of readily available tools based on the R programming language.

One of the most compelling tools described by Gentleman and Temple Lang is Sweave (Leisch 2002). This is an environment that facilitates writing “dynamic documents.” A dynamic document is a single document that includes both the narrative text of an article, and the R code that generates the plots and tables it includes. This document can be executed to run the R code and format the text to produce a PDF document that includes the output of the R code in the source document. By enabling authors to write text and code in a single document, Sweave enables Knuth’s (1992) literate programming methods for R users.

Since those papers were published, there has been a substantial increase in the use of R as a research tool in many fields (Tippmann 2014). At the same time, there have been substantial improvements in the ease of making dynamic documents in R (especially knitr and rmarkdown; Xie 2015; Baumer and Udwin 2015; Allaire et al. 2016) and in making R packages (especially devtools and roxygen2; Wickham and Chang 2016; Gandrud 2013). This means that making a research compendium based on an R package is now a practical solution to the challenges of organising and communicating research results.

That said, there are related efforts to support making research compendia that do not use the R package structure. For example, Gandrud (2013) advocates a reproducible research project structure that is consistent with the principles above, but does not use any elements of an R package. Likewise, ProjectTemplate (White 2014), recommended by Hilary Parker in her paper in this collection, is an R package that includes functions to generate a pre-determined directory structure to organize files, and auto-load packages, data and functions. Other similar packages for generating empty project directory structures include make Project (Silverman 2015), prodigenr (Johnston 2016), and workflowr (Blischak et al. 2017). A notable effort to promote reproducible research at the undergraduate level is Project TIER (Teaching Integrity in Empirical Research) for economics majors at Haverford College. Project TIER provides a detailed file-structure protocol for students working on their

senior thesis, aiming to teach students to document their data management and analysis to enable an independent researcher to reproduce the student's data processing and analysis (Ball and Medeiros 2012). A third influential approach to organising research compendia can be found in lessons developed by the Reproducible-Research-Curriculum community (Curriculum 2016), the Data Carpentry (Teal et al. 2015) and Software Carpentry (Wilson 2013) organisations, Jenny Bryan's (n.d.) materials for her STAT545 course, and Karl Broman's tutorials (2016). These lessons are less prescriptive than Gandrud (2013) and ProjectTemplate, focusing more on high-level guidance similar to the generic principles we described above.

3. Why Create a Research Compendium?

Using research compendia has benefits both for you as a researcher, and for the community that you work in. Among the most important of these benefits is that a research compendium is a convenient way to publicly share data and code (McKiernan et al. 2016). Papers with publicly available datasets may receive a higher number of citations than similar studies without available data (Piwowar, Day, and Fridsma 2007; Piwowar and Vision 2013; Henneken and Accomazzi 2011; Dorch 2012; Sears 2011; Gleditsch and Strand 2003). In addition to increased citations for data sharing, Pienta et al. (2010) found that data sharing is associated with higher publication productivity. They examined 7,040 NSF and NIH awards and concluded that the typical research grant award produces a median of five publications, but when data are archived a research grant award leads to a median of ten publications. Other benefits are documented elsewhere (Markowetz 2015; Donoho 2010).

In our own experience as researchers, we enjoy benefits of increased efficiency through simplified file management and streamlined analytical workflows. These help us to work more transparently and efficiently. While reproducibility is no guarantee of correctness, making our results reproducible lets us inspect our own work for errors. Because we follow the same pattern in multiple projects, the startup and reentry costs for each project are significantly reduced. The result is that we save time, and minimize the risk of errors that might result from results which cannot be reproduced. A compendium makes it easier to communicate our work with others, including collaborators and (not least of all) our future selves. As members of research communities, by including compendia with our published papers we benefit from increased impact and visibility of our research among our peers.

In the specific case of using an R package as a research compendium, you can benefit from the quality control mechanisms required to successfully build an R package. By organizing files into an R package, you follow conventions that save you time thinking about the best way to organize your project. Writing a function saves us from introducing errors that often result from repeated copying and pasting of code. This is because you can type one line of code (the name of your function) that calls hundreds of lines of code (your function). Writing functions for packages makes it easy to document the use of the code (particularly if you use roxygen2). This documentation can help you be productive more quickly when returning to work on a project

after stepping away from it. When you iteratively develop a package, you can easily run comprehensive checks on your code with `R CMD check` at the terminal or `devtools::check()` at the R console. Regularly checking your package like this can help identify problems before they become very frustrating to solve.

4. How to Make a Research Compendium?

At its simplest, a research compendium might consist of a single file of R code with in-line comments documenting the workflow. A slightly more complex approach might be a R markdown file with text and code in the same source document, and accompanying data files. In many cases, these simple approaches will be sufficient, and more elaborate organization would add unnecessary complexity and points of failure. However, as a project grows you might start with a single file, then add subdirectories which have consistent and readily meaningful names such as "analysis," "data," and so on. There are few strict rules here. The key principle is to organize the compendium so that another person can know what to expect from the plain meaning of the file and directory names. Using widely held conventions, such as the R package structure, will help other people to understand how your files relate to each other without having to ask you. Naming objects is difficult to do well, so it is worth to put some effort into a logical and systematic file naming convention if you have a complex project with many files and directories (e.g., a multi-experiment study where each experiment has numerous data and code files).

The R package structure can help with providing a logical organisation of files, by providing a set of standard locations for certain types of files. An ideal package-based file organization for a more complex project would look like this:

- A `README.md` file that describes the overall project and where to get started. It can be helpful to include graphical summary of the interlocking pieces of the project.
- Script files with reusable functions go in the `R/` directory. If these functions are documented using Roxygen, then the documentation will be automatically generated in a `man/` directory.
- Raw data files are kept in the `data/` directory. If your data are very large, or streaming, an alternative is to include a small-sample dataset so that people can try out the techniques without having to run very expensive computations.
- Analysis scripts and reports files go in the `analysis/` directory. In many cases it can be useful to give the analysis scripts ascending names, for example `001-load.R`, `002-clean.R` etc. This kind of file-naming helps with organisation, but it does not capture the full tree of dependencies in the way a `Makefile` or an R Markdown file does. To manage more complex workflows, the `analysis/` directory could include either an R markdown file, a `Makefile` or a `Makefile.R` file. These files are important because they control the order of the code execution. In more complex projects, careful use of caching or a `Makefile` can save time by only running code that has not changed since it was last run.

- A DESCRIPTION file in the project root provides formally structured, machine- and human-readable information about the authors, the project license, the software dependencies, and other metadata of the compendium. When a DESCRIPTION file is included along with the other items above, then the compendium is also a formal, installable R package. When your compendium is an R package, you can take advantage of many time-saving tools for package development, testing, and sharing (e.g., the devtools package that we noted above). R's built-in citation() function can use that metadata, along with references to any publications that result from the project, to provide users with the necessary information to cite your work.

We have developed the rrtools package to bootstrap the creation of research compendia packages that conform to the minimum ideal requirements listed above (Marwick et al. 2017c).

5. Examples of Real-World Research Compendia Using R Packages

In this section, we describe a selection of real-world package-based research compendia of varying complexity. Our hope is that these examples will give practical guidance about how you can organize your projects into package-based compendia.

5.1. Small Compendia

A simple example might look like Figure 1.

Duffy's (2015) parasite study provides an excellent real-world example of this simple research compendium format. Her compendium is archived on zenodo (<http://dx.doi.org/10.5281/zenodo.17804>), with the development repository on GitHub (<https://github.com/duffymeg/BroodParasiteDescription>). This format fulfills the three generic principles described above: it organizes the files according to a convention, separates data and code, and specifies the computational environment. The DESCRIPTION file shows that R version 3.2.0 or higher is required for this project, and the packages that the project used are also listed there.

Although Duffy's compendium is an excellent template for many research projects, it might not qualify as a compendium as defined by Gentleman and Temple Lang because it does not contain a dynamic document. An R script file, however, can be converted to an R Markdown document

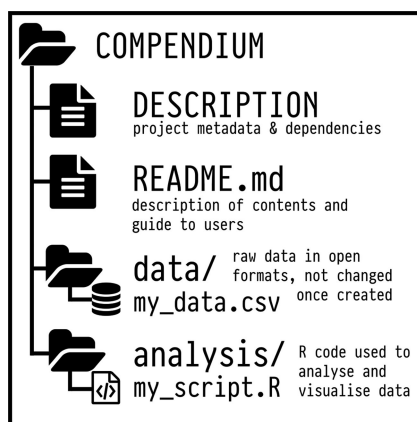


Figure 1. An example of a small research compendium.

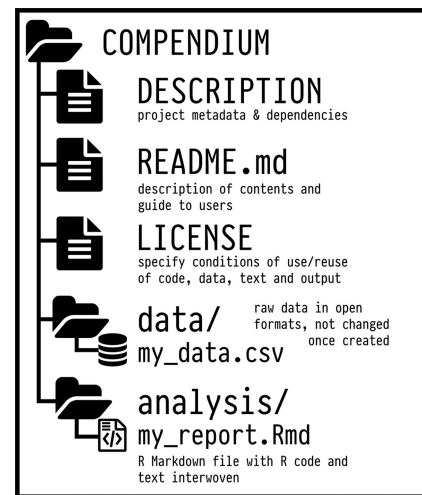


Figure 2. An example of a small research compendium that also includes a license file and a dynamic document.

which includes both text and code. The diagram below shows an intermediate-level compendium. Compared to the simple compendium described above, the `my_scripts.R` file has become `my_report.Rmd`. Another change is the addition of a LICENSE file that specifies how the contents of the compendium are allowed to be reused (more on this below in the section How to Share a Research Compendium?).

The example in Figure 2 is based on the real-world compendium of the zooarcheological study by Conrad et al. (2016). Their compendium is archived at the University of New Mexico (<http://repository.unm.edu/handle/1928/26730>) and the development version is on GitHub (https://github.com/cylerc/AP_SC). The R Markdown file generates all the data visualisations and statistical contents of the publication, with some basic commentary on the methods. However, it does not contain the complete article as submitted to the publisher.

5.2. Medium Compendia

While the two examples above qualify as minimal R packages, our next example (Figure 3) shows a more complete use of the R package structure. This includes the R/ and man/ directories. The R/ directory contains custom functions that are used repeatedly throughout the project. The man/ directory contains the manual (i.e., documentation) for the use of the functions. The NAMESPACE file is a feature of R packages that exports the functions in a package, so that they can be used in the current project and by other packages.

A real-world compendium demonstrating this more complete use of the R package structure is Hollister et al. (2016). Their archived compendium is online at zenodo (<http://dx.doi.org/10.5281/zenodo.40271>) and the development repository is at GitHub (<https://github.com/USEPA/LakeTrophicModelling>). It differs slightly from the diagram above because Hollister et al. have their main article in the vignettes/ directory. The vignette is a traditional component of an R package for providing detailed examples of how package functions can be used. By including the article as a vignette, it can be automatically generated from the R Markdown document when the package is installed. Their vignette document includes the full

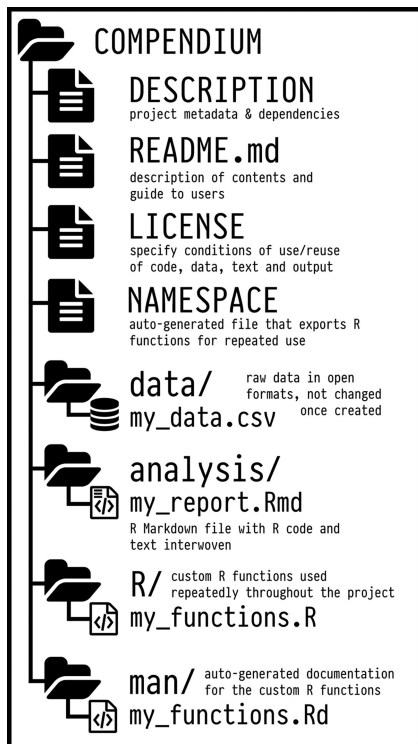


Figure 3. An example of a more complete use of the R package structure for a research compendium.

text of the article, as well as the figures and tables generated by R, along with the usual scholarly apparatus such as cross-references, citations and a reference list. A similar example of this use of an R package as a research compendium can be seen in Negre et al. (2016).

5.3. Large Compendia

The diagram in Figure 4 outlines one possible structure of a compendium at the more complex end of the spectrum. This example includes continuous integration; `.travis.yml` is a configuration file for the Travis continuous integration service. The `Dockerfile` is a file that instructs the Docker software how to create a virtual environment for running the R code in an isolated and portable context. The `Makefile` is a file with instructions for executing R code in the compendium in a way that avoids unnecessary repetition. These three files contain organizing metadata that are intended to be both machine- and human-readable, but are not written in the R language. Unit tests, contained in the `tests/` directory, are R code scripts to test that specific functions in the compendium produce their expected outputs, given known inputs. The addition of these tools, initially developed for software engineering, solves the problems of specifying the computational environment and the relationships between data, code, and output.

A real-world example of this more complex type of research compendium can be seen in Boettiger et al. (2015). Their archived compendium is online at zenodo (<http://dx.doi.org/10.5281/zenodo.12669>) and the development repository is at GitHub (<https://github.com/cboettig/nonparametric-bayes>). Boettiger et al. have a top-level `manuscripts/` directory that holds the files that generate the journal article and supplementary documents, as well as

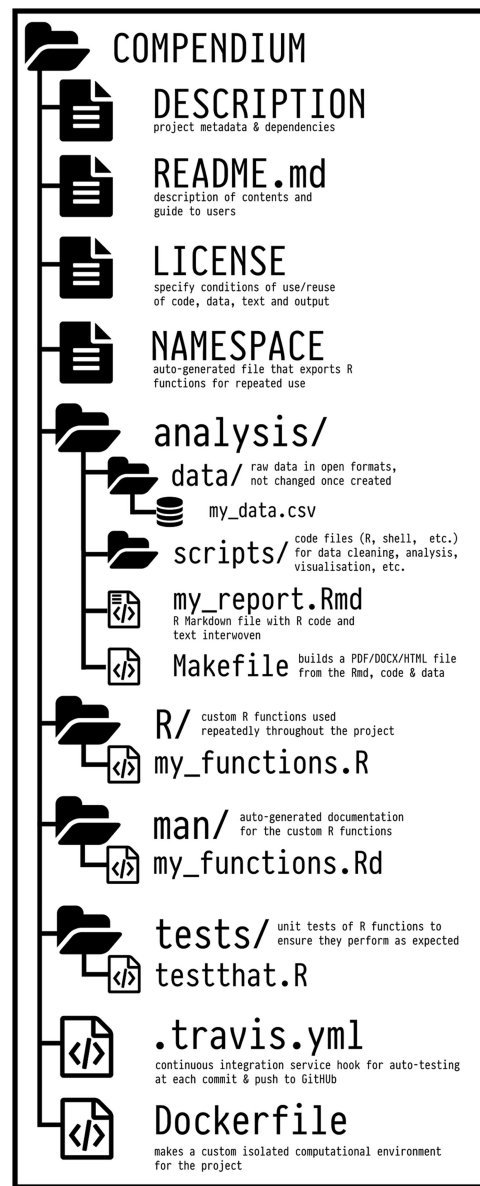


Figure 4. An example of a more elaborate research compendium based on the R package structure.

a `Makefile` and a `Dockerfile`. Other notable items at the top level of this compendium are the `.drone.yml`, `.zenodo.json` and `.travis.yml` files. The drone file contains configuration details for the Drone continuous integration service that operates in Docker containers. In this case, each time a commit is made to the repository, the Drone web service automatically renders the manuscript files (including running the R code in the manuscript) to PDF in a Docker container defined by the `Dockerfile` in `manuscripts/`. This provides an automatic check to see if a PDF can be generated from the manuscript R markdown files after the latest commit. The `.travis.yml` file performs a similar purpose, but is focused on whether or not the R package can successfully build in a generic Linux environment. The `.zenodo.json` file provides machine-readable metadata about the compendium. This is useful for automatically archiving the compendium at zenodo. Additional examples of similarly complex compendium can be found in Clarkson et al. (2015), Kamvar, et al. (2017), Marwick

et al. (2016), Marwick et al. (2017a), Marwick et al. (2017b) and Eglen (2016). These examples indicate that this use of the R package is a sustainable and efficient approach to packaging research reproducibly.

5.4. Makefiles

The `Makefile` in the example above and in Boettiger's compendia uses the `make` language (rather than R) to concisely specify the relationship between data, the output, and the code that generates it. A `Makefile` defines outputs (called "targets") in terms of inputs (called "dependencies") and the code necessary to produce them (called "recipes"). The program GNU Make then creates the outputs from that specification, see Broman (2013) for more details. In addition to GNU Make, R packages such as `remake` (FitzJohn 2016) and `drake` (Landau 2017) allow you to write Make-like instructions entirely in R.

5.5. Isolating the Computational Environment

The `Dockerfile` in these more complex examples helps to specify an isolated computational environment that includes all the dependencies necessary for your analysis to run (Boettiger 2015; Boettiger and Eddelbuettel 2017). Using a lightweight virtualization system such as Docker minimizes problems with changes in R packages that might break your code, since Docker will preserve the exact versions of R and the packages that you used for your analysis. Docker makes it easy to share your project because when you share your Docker container with collaborators, you do not need to worry about differences between operating systems. This is important if your analysis requires software outside of the R programming environment.

While Docker will virtualize R, its packages, and the entire operating system, there are other solutions that are focused exclusively on the R environment. The `packrat` package, for example, downloads the source code of the exact versions of the packages you have used in your project, and stores them with your project (Ushey et al. n.d.). Another approach is provided by `checkpoint` (Microsoft Corporation 2016) which installs specific previous versions of R packages from the Microsoft R Archived Network. The `devtools` package (Wickham and Chang 2016) similarly enables installation of previous versions of packages from CRAN (the Comprehensive R Archive Network, the default R package repository). These packages can protect you from updates to packages that might stop your code from working, and help you and your collaborators work with the same package versions.

5.6. Continuous Integration

With continuous integration, each Git commit you make to your repository on GitHub (or similar service) triggers a script that builds your R package, and reports to you if the build succeeded or not. This is convenient because it saves you from having to manually build your package after each update to your code. The complex research compendia examples cited above make use of Travis CI, Drone.io, and CircleCI services. These provide free remote continuous integration services for public GitHub

projects. Another advantage of these services is that they provide badges for your repository webpage that signal the current state of your compendium, based on the last build attempt (e.g., "build passing" or "build failing").

6. How to Share a Research Compendium?

To share a research compendium you need to think about four separate but related issues: (1) licensing, (2) version control, (3) persistence, and (4) metadata.

A license file indicates to others how your work may be reused. You should of course seek specific intellectual property advice relevant to your work from your organization, as well as the policies of any publisher distributing material from the compendium (such as an associated academic article). While research compendia typically use open source software licenses (see <http://opensource.org>), authors should be aware that such licenses are designed for software and may not always be well suited for other elements of a compendium. In many jurisdictions, data are considered facts rather than a creative work, so they are not protected by copyright law. Stodden (2009) recommended using the Creative Commons Public Domain declaration (CC-0) for data. A compendium may also include documentation such as a scientific article and figures that are more appropriately licensed under a Creative Commons Attribution (CC-BY) license. Note that some data repositories (such as DataDryad.org or figshare.com, discussed below) will require a CC-BY or CC-0 license for all submitted content.

A version control system such as Git is the best way to preserve the history of changes to the research compendium. Version control facilitates both private collaboration among colleagues on the project and the distribution and maintenance of the compendium in the future (Ram 2013; Jones 2013; Loeliger and McCullough 2012). Bryan's article in this collection is a detailed and accessible introduction to using Git in research contexts.

Once you are ready to share your compendium, the best way is to archive a specific commit at a repository that issues persistent URLs, such as a Digital Object Identifier (DOI), for file archives (e.g., osf.io, figshare.com or zenodo.org). DOIs are designed to be far more persistent than other URLs, which often break as web pages change over time (Klein et al. 2014). There are many discipline-specific DOI-issuing repositories, listed at re3data.org. You should survey the literature of your research community to see what repositories are preferred.

DOIs provided by data repositories have another advantage over other web links: when a data repository provides an author with a DOI, it must also provide basic metadata information to a central registry which oversees the creation of DOIs: either CrossRef (primarily journal articles) or DataCite (data and other products). Once your compendium is hosted at a repository, you include the DOI to the compendium in your reports and publications. This means you have a publicly available snapshot of the code that matches results in the article. A DOI also simplifies citation of the compendium, so you can cite it in your article (and others can cite it in their work).

Code development can continue after the article is published, but with a DOI that links to a specific commit, you and other users of your code can be confident that they have the version

that matches the article. Putting the compendium on Dropbox, Google Drive, or similar services is another way to make it easily available. However, these services do not offer the same persistence as a DOI-issuing repository and are suitable only for short-term, temporary storage during development.

While CRAN is one of the biggest and best-known systems for archiving and distributing R packages, we do not recommend it for research compendium packages. The main reason that CRAN is not suitable is that it is very strict about the directory structure and contents for the R packages that it accepts. For example, the top-level `analysis/` directory in several of the examples above would not be allowed in a package hosted on CRAN. There are ways around this (e.g., use the `inst/` or `vignettes/` directory, see Hollister's compendium), but this makes the package structure less intuitive. This has the disadvantage of making the compendium harder to navigate and so less accessible to others.

A second reason why CRAN is not suitable for research compendium packages is that CRAN has a 5 MB size limit for package data and documentation. Hollister's package is 6 MB after installation (from GitHub), and many research compendia exceed this size because of moderate to large raw data files, image files, cached computation results, etc., so these could not be hosted on CRAN. The DOI-issuing repositories that we describe above are more suitable than CRAN for research compendium packages because they have more generous file size limits, and no constraints on directory organization.

7. Useful Tools and Templates for Making Research Compendia

Probably the most useful set of tools for making research compendia as R packages is the `devtools` package (Wickham and Chang 2016), combined with the RStudio integrated development environment. Both of these offer useful shortcuts for meeting the standards of an R package. Our `rrtools` package extends functions in `devtools` and provides instructions, templates, and functions for making a basic compendium-package suitable for doing reproducible research with R (Marwick et al. 2017). Wickham's (2015) book *R Packages* is a useful in-depth guide to the format of R packages. The manual *Writing R Extensions* is included with every R installation and describes the process of creating R packages.

More specifically, several people have developed templates for using R packages as research compendia. These templates are mostly for their personal use, and are works-in-progress, but are freely available for others to adapt and learn from:

- Jeff Hollister's manuscriptPackage (<https://github.com/jhollist/manuscriptPackage>)
- Carl Boettiger's template (<https://github.com/cboettig/template>)
- Francisco Rodriguez-Sanchez's template (<https://github.com/Pakillo/template>)
- Lincoln Mullen's example Makefiles in R-package like repositories

For writing research papers in R Markdown, a useful package is `bookdown` (Xie 2016), which enables figure and table captions and cross-referencing. There are many R Markdown templates for journal articles from various publishers in the `rticles` (Allaire

et al. 2016) package. These packages and templates, together with RStudio add-ins such as the `wordcountaddin` (Marwick et al. 2017d) for character, word and sentence counts, and `citr` (Aust 2016) for inserting formatted citations, provide a rich scholarly writing environment for working with R packages in RStudio. See also `rOpenSci's` Reproducibility Guide (rOpenSci et al. 2017) for additional community-contributed templates and resources.

8. Conclusion

In summary, we have defined the general principles of research compendia, and described a number of different ways that researchers have organized their work to meet those criteria. For a researcher working primarily in R, a compendium based on the R package standards, along with additional tools drawn from the software engineering toolkit, is a natural, efficient, and reliable way of organizing one's work. Researchers working in different languages can draw ideas from the general principles and example projects presented here. We encourage researchers to strive to organize their code in reproducible ways and treat their code and data as much a product of their scholarship worthy of distribution as the articles themselves.

Acknowledgments

This essay originated in discussions at the 2015 `rOpenSci` unconference (<https://github.com/ropensci/unconf/issues/11> and <https://github.com/ropensci/unconf/issues/31>). Thanks to those who were part of these discussions and helped to broaden our understanding of compendia and R packages. BM was also supported by an Australian Research Council Future Fellowship (FT140100101).

References

- Allaire, J. et al. (2016), "rmarkdown: Dynamic Documents for R", R package version 0.9.6, available at <https://CRAN.R-project.org/package=rmarkdown> [81,86]
- Allaire, J., R Foundation, H. Wickham et al. (2017), "rticles, R package," version 0.4.1, available at <https://CRAN.R-project.org/package=rticles> [xxxx]
- Allison, S. (2016), "Other People's Data: Humanities Edition," *Journal of Cultural Analytics*. Available at <http://culturalanalytics.org/2016/12/other-peoples-data-humanities-edition/> [80]
- Aust, F. (2016), "citr: 'RStudio' Add-in to Insert Markdown Citations," R package, version 0.2.0, available at <https://CRAN.R-project.org/package=citr> [86]
- Ball, R., and Medeiros, N. (2012), "Teaching Integrity in Empirical Research: A Protocol for Documenting Data Management and Analysis," *The Journal of Economic Education*, 43, 182–189. [82]
- Baumer, B., and Udwin, D. (2015), "R Markdown," *Wiley Interdisciplinary Reviews: Computational Statistics*, 7, 167–177. [81]
- Begley, C. G., and Ellis, L. M. (2012), "Drug Development: Raise Standards for Preclinical Cancer Research," *Nature*, 483, 531–533. [80]
- Blischak, J., P. Carbonetto, M. Stephens (2017), "workflow: A Workflow Template for Creating a Research Website," R package, Version: 0.11.0, available at <https://github.com/jdblischak/workflow> [81]
- Boettiger, C. (2015), "An Introduction to Docker for Reproducible Research," *ACM SIGOPS Operating Systems Review*, 49, 71–79. [85]
- Boettiger, C. and D. Eddelbuettel (2017), "An Introduction to Rocker: Docker Containers for R," *The R Journal*, 9, 527–536. [85]

- Boettiger, C., Mangel, M., and Munch, S. (2015), "Avoiding Tipping Points in Fisheries Management Through Gaussian Process Dynamic Programming," *Proceedings of the Royal Society of London, Series B*, 282, 20141631. [84]
- Bollen, K., Cacioppo, J., Kaplan, R., Krosnick, J., and Olds, J. (2015), "Social, Behavioral, and Economic Sciences Perspectives on Robust and Reliable Science: Report of the Subcommittee on Replicability in Science, Advisory Committee to the National Science Foundation Directorate for Social, Behavioral, and Economic Sciences," Retrieved from the *National Science Foundation*, available at www.nsf.gov/sbe/AC_Materials/SBE_Robust_and_Reliable_Research_Report.pdf. [80]
- Broman, K. (2013), "minimal Make: A Minimal Tutorial on Make," available at http://kbroman.org/minimal_make. [85]
- (2016), "Organize Your Data and Code," available at <http://kbroman.org/steps2rr/pages/organize.html> [82]
- Bryan, J. (2016), Stat545: Data Wrangling, Exploration, and Analysis With R, available at <http://stat545.com/>. [82]
- Buckheit, J. B., and Donoho, D. L. (1995), "Wavelab and Reproducible Research," *Wavelets and Statistics. Lecture Notes in Statistics*, 103, 55–81. [81]
- Claerbout, J. F., and Karrenfach, M. (1992), "Electronic Documents Give Reproducible Research a New Meaning," in *Society of Exploration Geophysics Technical Program Expanded Abstracts*, pp. 601–604, available at <http://sepwww.stanford.edu/doku.php?id=sep:research:reproducible:seg92> [80,81]
- Clarkson, C., Smith, M., Marwick, B., Fullagar, R., Wallis, L. A., Faulkner, P., Manne, T., Hayes, E., Roberts, R. G., and Jacobs, Z. (2015), "The Archaeology, Chronology and Stratigraphy of Madjedbebe (Malakunanja ii): A Site in Northern Australia With Early Occupation," *Journal of Human Evolution*, 83, 46–64. [84]
- Collberg, C., and Proebsting, T. A. (2016), "Repeatability in Computer Systems Research," *Communications of the ACM*, 59, 62–69. [80]
- Conrad, C., Higham, C., Eda, M., and Marwick, B. (2016), "Palaeoecology and Forager Subsistence Strategies During the Pleistocene Holocene Transition: A Reinvestigation of the Zooarchaeological Assemblage From Spirit Cave, Mae Hong Son Province, Thailand," *Asian Perspectives*, 55, 2–27. [83]
- Curriculum, R. S. (2016), "The Organization Lesson for the Reproducible Science Curriculum," available at <https://github.com/Reproducible-Science-Curriculum/rr-organization1> [82]
- Donoho, D. L. (2010), "An Invitation to Reproducible Computational Research," *Biostatistics*, 11, 385–388. Available at <http://biostatistics.oxfordjournals.org/content/11/3/385.short> [80,82]
- Dorch, S. (2012), "On the Citation Advantage of Linking to Data: Astrophysics," available at <https://halshs.archives-ouvertes.fr/hprints-00714715/> [82]
- Duffy, M. A., James, T. Y., and Longworth, A. (2015), "Ecology, Virulence, and Phylogeny of Blastulidium Paedophthorum, a Widespread Brood Parasite of Daphnia," *Applied and Environmental Microbiology*. Available at <http://aem.asm.org/content/early/2015/06/02/AEM.01369-15.abstract> [83]
- Eglen, S. J. (2016), "Bivariate Spatial Point Patterns in the Retina: A Reproducible Review," *Journal de la Société Française de Statistique*, 157, 33–48. [85]
- FitzJohn, R. (2016), *Remake: Make-Like Build Management*, R package version 0.2.0. Available at <https://github.com/richfitz/remake> [85]
- Gandrud, C. (2013), *Reproducible Research With R and RStudio*, Boca Raton, FL: CRC Press. [81]
- Gentleman, R., and Temple Lang, D. (2004), "Statistical Analyses and Reproducible Research," *Bioconductor Project Working Papers*, p. 2. [81]
- (2012), "Statistical Analyses and Reproducible Research," *Journal of Computational and Graphical Statistics*, 16, 1–23. [81]
- Gentleman, R. (2005), "Reproducible Research: A Bioinformatics Case Study," *Statistical Applications in Genetics and Molecular Biology*, 4, 1034. [81]
- Gleditsch, N. P., and Strand, H. (2003), "Posting Your Data: Will You be Scooped or Will You be Famous?" *International Studies Perspectives*, 4, 72–107. [82]
- Goldstone, A. (2017), "From Reproducible to Productive," *Journal of Cultural Analytics*. Available at <http://culturalanalytics.org/2017/02/from-reproducible-to-productive/> [80]
- Goodman, S. N., Fanelli, D., and Ioannidis, J. P. A. (2016), "What Does Research Reproducibility Mean?" *Science Translational Medicine*, 8, 341ps12–341ps12. [80]
- Graham, S., Milligan, I., and Weingart, S. (2015), *Exploring Big Historical Data: The Historian's Macroscope*, Hackensack, NJ: Imperial College Press. [80]
- Henneken, E. A., and Accomazzi, A. (2011), "Linking to Data: Effect on Citation Rates in Astronomy," *CoRR*, abs/1111.3618. Available at <http://arxiv.org/abs/1111.3618> [82]
- Hollister, J. W., Milstead, W. B., and Kreakie, B. J. (2016), "Modeling Lake Trophic State: A Random Forest Approach," *Ecosphere*, 7, e01321. [83]
- Johnston, L. (2016), "prodigenr: Research Project Directory Generator," R package, version 0.3.0 Available at <https://cran.r-project.org/package=prodigenr> [81]
- Jones, Z. M. (2013), "Git/GitHub, Transparency, and Legitimacy in Quantitative Research," *The Political Methodologist*, 21, 6–7. Available at <http://zmjones.com/static/papers/git.pdf> [85]
- Kamvar, Z. N., Amaradasa, B. S., Jhala, R., McCoy, S., Steadman, J. R., Everhart, S. E., (2017), "Population Structure and Phenotypic Variation of *Sclerotinia sclerotiorum* from Dry Bean (*Phaseolus vulgaris*) in the United States," *PeerJ*, 5, e4152. [84]
- Kidwell, M. C., Lazarevi, L. B., Baranski, E., Hardwicke, T. E., Piechowski, S., Falkenberg, L.-S., Kennett, C., Slowik, A., Sonnleitner, C., Hess-Holden, C., Errington, T. M., Fiedler, S., and Nosek, B. A. (2016), "Badges to Acknowledge Open Practices: A Simple, Low-Cost, Effective Method for Increasing Transparency," *PLoS Biol*, 14, 1–15. [81]
- Klein, M., Van de Sompel, H., Sanderson, R., Shankar, H., Balakireva, L., Zhou, K., and Tobin, R. (2014), "Scholarly Context Not Found: One in Five Articles Suffers From Reference Rot," *PLOS ONE*, 9, 1–39. [85]
- Knuth, D. E. (1992), "Literate Programming," *CSLI Lecture Notes*, Stanford, CA: Center for the Study of Language and Information (CSLI), 1992, 1. [81]
- Koenker, R. (1996), "Reproducible Econometric Research, Department of Econometrics, University Of Illinois, Urbana-Champaign," Technical Report, IL, Tech. Rep. Available at <http://www.econ.uiuc.edu/~roger/research/repro/> [81]
- Landau, W. M. (2017), "drake: Data Frames in R for Make," R package, version 5.0.0, available at <https://cran.r-project.org/package=drake> [85]
- Leisch, F. (2002), "Sweave: Dynamic Generation of Statistical Reports Using Literate Data Analysis," in *Compstat*, New York: Springer, pp. 575–580. [81]
- LeVeque, R. J., Mitchell, I. M., and Stodden, V. (2012), "Reproducible Research for Scientific Computing: Tools and Strategies for Changing the Culture," *Computing in Science & Engineering*, 14, 13–17. Available at <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6171147> [80]
- Loeliger, J., and McCullough, M. (2012), *Version Control with Git: Powerful Tools and Techniques for Collaborative Software Development*, Sebastopol, CA: O'Reilly Media, Inc. [85]
- Markowitz, F. (2015), "Five Selfish Reasons to Work Reproducibly," *Genome Biology*, 16, 274. [82]
- Marwick, B. (2017), "Computational Reproducibility in Archaeological Research: Basic Principles and a Case Study of Their Implementation," *Journal of Archaeological Method and Theory*, 24, 424. [80]
- Marwick, B., Clarkson, C., O'Connor, S., and Collins, S. (2016), "Early Modern Human Lithic Technology From Jerimalai, East Timor," *Journal of Human Evolution*, 101, 45–64. [85]
- Marwick, B., Hayes, E., Clarkson, C., and Fullagar, R. (2017a), "Movement of Lithics by Trampling: An Experiment in the Madjedbebe Sediments, Northern Australia," *Journal of Archaeological Science*, 79, 73–85. [85]
- Marwick, B., Van Vlack, H., Conrad, C., Shoocongdej, L., Thongcharoenchaikit, C., and Kwak, S. (2017b), "Adaptations to Sea Level Change and Transitions to Agriculture at Khao Toh Chong Rockshelter, Peninsular Thailand," *Journal of Archaeological Science*, 77, 94–108. [85]
- Marwick, B., Knitter, D., Kennedy, P., Muller-Scheessel, N., Hinz, M., Schmid, C., Braun, R., and Francuzik, W. (2017c), "rrtools: Tools for

- Writing Reproducible Research in R”, R package version 0.1.0. Available at <https://github.com/benmarwick/rrtools> [83,86]
- Marwick, B., JooYoung, S., and Bengtsson, H. (2017d), “wordcountaddin: Word Counts and Readability Statistics in R Markdown Documents,” R package, version 0.2.0, available at <https://github.com/benmarwick/wordcountaddin> [86]
- McKiernan, E. C., Bourne, P. E., Brown, C. T., Buck, S., Kenall, A., Lin, J., McDougall, D., Nosek, B. A., Ram, K., Soderberg, C. K., Spies, J. R., Thaney, K., Updegrave, A., Woo, K. H., and Yarkoni, T. (2016), “How Open Science Helps Researchers Succeed,” *eLife*, 5, e16800. Available at <https://dx.doi.org/10.7554/eLife.16800> [82]
- Microsoft Corporation (2016), “checkpoint: Install Packages From Snapshots on the Checkpoint Server for Reproducibility”, R package version 0.3.16. Available at <https://CRAN.R-project.org/package=checkpoint> [85]
- Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., Simonsohn, U., Wagenmakers, E.-J., Ware, J. J., and Ioannidis, J. P. A. (2017), “A Manifesto for Reproducible Science,” *Nature Human Behaviour*, 1, 0021. [80]
- Negre, J., Munoz, F., and Lancelotti, C. (2016), “Geostatistical Modelling of Chemical Residues on Archaeological Floors in the Presence of Barriers,” *Journal of Archaeological Science*, 70, 91–101. [84]
- Nosek, B. A., Alter, G., Banks, G. C., Borsboom, D., Bowman, S. D., Breckler, S. J., Buck, S., Chambers, C. D., Chin, G., Christensen, G., Contestabile, M., Dafoe, A., Eich, E., Freese, J., Glennerster, R., Goroff, D., Green, D. P., Hesse, B., Humphreys, M., Ishiyama, J., Karlan, D., Kraut, A., Lupia, A., Mabry, P., Madon, T., Malhotra, N., Mayo-Wilson, E., McNutt, M., Miguel, E., Levy Paluck, E., Simonsohn, U., Soderberg, C., Spellman, B. A., Turitto, J., VandenBos, G., Vazire, S., Wagenmakers, E. J., Wilson, R., and Yarkoni, T. (2015), “Promoting an Open Research Culture: Author Guidelines for Journals Could Help to Promote Transparency, Openness, and Reproducibility,” *Science*, 348, 1422–1445. [81]
- Open Science Collaboration (2015), “Estimating the Reproducibility of Psychological Science,” *Science*, 349, aac4716. [80]
- Peng, R. D. (2009), “Reproducible Research and Biostatistics,” *Biostatistics*, 10, 405–408. [80]
- Peng, R. D., Dominici, F., and Zeger, S. L. (2006), “Reproducible Epidemiologic Research,” *American Journal of Epidemiology*, 163, 783–789. [81]
- Pienta, A. M., Alter, G. C., and Lyle, J. A. (2010), “The Enduring Value of Social Science Research: The Use and Reuse of Primary Research Data,” Available at <https://deepblue.lib.umich.edu/handle/2027.42/78307> [82]
- Piper, A. et al. (2017), “Data Sharing Policy,” *Journal of Cultural Analytics*. Available at <http://culturalanalytics.org/about/about-ca/> [81]
- Piwowar, H. A., Day, R. S., and Fridsma, D. B. (2007), “Sharing Detailed Research Data is Associated With Increased Citation Rate,” *PLoS ONE*, 2, e308. Available at <http://dx.plos.org/10.1371/journal.pone.0000308> [82]
- Piwowar, H. A., and Vision, T. J. (2013), “Data Reuse and the Open Data Citation Advantage,” *PeerJ*, 1, e175. [82]
- Ram, K. (2013), “Git can Facilitate Greater Reproducibility and Increased Transparency in Science,” *Source Code for Biology and Medicine*, 8, 7. [85]
- Rokem, A., Marwick, B., and Staneva, V. (2017), “Assessing Reproducibility,” in J. Kitzes, D. Turek, and F. Deniz, eds, *The Practice of Reproducible Research: Case Studies and Lessons from the Data-Intensive Sciences*, CA: University of California Press. [80]
- rOpenSci et al. (2017), “Reproducibility Guide,” available at <http://ropensci.github.io/reproducibility-guide/> [86]
- Sandve, G. K., Nekrutenko, A., Taylor, J., and Hovig, E. (2013), “Ten Simple Rules for Reproducible Computational Research,” *PLoS Comput Biol*, 9, e1003285. Available at <http://dx.doi.org/10.1371/journal.pcbi.1003285> [80]
- Sears, J. (2011), “Data Sharing Effect on Article Citation Rate in Paleooceanography,” in *AGU Fall Meeting Abstracts* (Vol. 1), available at <http://adsabs.harvard.edu/abs/2011AGUFMIN53B1628S>, p. 1628. [82]
- Silverman, N. (2015), “makeProject: Creates an Empty Package Framework for the LCFD Format,” R package, version 1.0, available at <https://CRAN.R-project.org/package=makeProject> [81]
- Stanisic, L., and Legrand, A. (2014), “Effective Reproducible Research With Org-Mode and Git,” in *Euro-Par 2014: Parallel Processing Workshops*, Springer, pp. 475–486. Available at http://link.springer.com/chapter/10.1007/978-3-319-14325-5_41. [80]
- Stodden, V. (2009), “The Legal Framework for Reproducible Scientific Research: Licensing and Copyright,” *Computing in Science & Engineering*, 11, 35–40. [85]
- (2014), “What Scientific Idea is Ready for Retirement? Reproducibility,” *Edge*. Available at <https://www.edge.org/response-detail/25340> [80]
- Stodden, V., McNutt, M., Bailey, D. H., Deelman, E., Gil, Y., Hanson, B., Heroux, M., Ioannidis, J. P., and Taufer, M. (2016), “Enhancing Reproducibility for Computational Methods,” *Science*, 354, 1240–1241. [80]
- Stodden, V., and Miguez, S. (2014), “Best Practices for Computational Science: Software Infrastructure and Environments for Reproducible and Extensible Research,” *Journal of Open Research Software*, 2. Available at <http://openresearchsoftware.metajnl.com/articles/10.5334/jors.ay/> [80]
- Teal, T. K., Cranston, K. A., Lapp, H., White, E., Wilson, G., Ram, K., and Pawlik, A. (2015), “Data Carpentry: Workshops to Increase Data Literacy for Researchers,” *International Journal of Digital Curation*, 10, 135–143. [82]
- Tippmann, S. (2014), “Programming Tools: Adventures With R,” *Nature*, 517, 109–110. [81]
- Ushey, K., McPherson, J., Cheng, J., Atkins, A., and Allaire, J. (n.d.), “packrat: A Dependency Management System for Projects and Their R Package Dependencies”, R package version 0.4.7-12. available at <https://github.com/rstudio/packrat/> [85]
- Vandewalle, P., Kovacevic, J., and Vetterli, M. (2009), “Reproducible Research in Signal Processing,” *IEEE Signal Processing Magazine*, 26, 37–47. [81]
- Vinod, H. (2001), “Care and Feeding of Reproducible Econometrics,” *Journal of Econometrics*, 100, 87–88. [81]
- White, J. M. (2014), “ProjectTemplate: Automates the Creation of New Statistical Analysis Projects”, R package version 0.6. Available at <https://CRAN.R-project.org/package=ProjectTemplate> [81]
- Wickham, H. (2015), *R Packages* (1st ed.), Sebastopol, CA: O’Reilly Media, Inc. [86]
- Wickham, H., and Chang, W. (2016), “devtools: Tools to Make Developing R Packages Easier”, R package version 1.12.0. Available at <https://CRAN.R-project.org/package=devtools> [81,85,86]
- Wilson, G. (2013), “Software Carpentry: Lessons Learned,” *CoRR*, abs/1307.5448. available at <http://arxiv.org/abs/1307.5448> [82]
- Wilson, G., Aruliah, D. A., Brown, C. T., Chue Hong, N. P., Davis, M., Guy, R. T., Haddock, S. H. D., Huff, K. D., Mitchell, I. M., Plumbley, M. D., Waugh, B., White, E. P., and Wilson, P. (2014), “Best Practices for Scientific Computing,” *PLoS Biology*, 12, e1001745. Available at <http://dx.plos.org/10.1371/journal.pbio.1001745> [80]
- Xie, Y. (2015), *Dynamic Documents With R and knitr* (2nd ed.), Boca Raton, FL: Chapman and Hall. ISBN 978-1498716963. Available at <http://yihui.name/knitr/> [81]
- (2016), *bookdown: Authoring Books and Technical Documents with R Markdown*, Boca Raton, FL: Chapman and Hall/CRC. [86]