Issue 5.1, Spring 2023 | Hot Topic Published online: Apr 27, 2023 DOI: 10.1162/99608f0039a00003 TOPIC SPECIALS

Cite 100 Metrics 4 Comments 0 Citations 0 Usage 1,444 Saved 0 Recommendations 0 Share

Software Engineering Practices in Academia: Promoting the 3Rs—Readability, Resilience, and Reuse

by Andrew Connolly, Joseph Hellerstein, Naomi Alterman, David Beck, Rob Fatland, Ed Lazowska, Vani Mandava and Stuart Stone

Published on Apr 27, 2023

Cite this as: 1 year ago

ABSTRACT

Over the past decade an data science has become integral to the research workflow, we, like many others, have learned that good data science requires high quality software engineering. Unfortunately, our experience is that many data science projects can be limited by the absence of software engineering processes. We advocate that data science projects should incorporate what we call the 3Rs of software engineering: readability (human comprehension of software), resilience (software behaves correctly), and reuse (code can be used by others and be modified to other software). This article discusses engineering practices that promote 3R software in academia. We emphasize that best practices in academia may differ from those in industry because of substantial differences in project scope (most academic projects have a single developer who is the sole user) and the reward systems in place in academia. We provide a framework for selecting a level of software engineering practice based on the project's scope and anticipated lifetime and change over time. We further discuss how to improve training in software engineering skills in an academic environment and how to build communities of practice that span across disciplines.

Keywords: software engineering, reproducibility, reuse, data science education

Media Summary

Data science is almost surprising, interpreting, and acting data — all which are critical to modern science. Software engineering practices can improve software necessary for high quality data science. But academic data science projects are often staffed with researchers with few skills in software engineering, especially the ability to produce readable, resilient, and reusable software. This is what we call the 3Rs of software engineering.

For the last decade, the eScience Institute at the University of Washington has been training data scientists and developing data science techniques. Thousands of undergraduate and graduate students, along with researcher faculty, have participated in our programs, training, and formal classes. As a result, we have developed a considerable understanding about the software engineering principles that promote effective data science in academia.

We find that while some academic projects require industrial-strength software engineering, most projects do not. We recommend adopting software engineering practices in the scope of the project. We classify the scope of academic projects as: ad-hoc (single researcher creates and uses the project), lab (multiple researchers use the project, but only one or two create), and community (developers have limited knowledge of the users).

The centerpiece of this article is how to adjust software engineering practices to the project scope. For example, a ad-hoc project should include unit tests, but not regression tests. Similarly, a community project should, at a minimum, include packaging conventions, but may not need formal design documentation. A community project should clearly lay the same practices as those employed by industry software engineers. These recommendations are detailed in our discussion, with references to help with applying these practices.

The eScience Institute is developing a centralized team of software engineers who provide strategic support to research projects. In the future, we will report on the effectiveness of our efforts, and the best practices and challenges with developing and operating such a team.

1. Good Data Science Requires High-Quality Software Engineering

As data science has become more integrated within research, our reliance on software and its development has grown substantially. The lion's part of data science relies heavily on the development of readable, resilient, and reusable software to create, analyze, and visualize data (e.g., Jenness, 2016). The research we undertake and the discoveries we make are becoming ever more dependent on robust and reproducible software frameworks and development processes. This reliance is however, not unique to data science. In a study of researchers in the United Kingdom (Hettrick & Barnham, 2018), 92% of respondents stated that their research relied on the correct use of software, and 56% said that they develop their own software. Indeed, Hettrick (2016) found that without software, more than 30% of the academic respondents would either be unable or would find it extremely difficult to do their research.

Despite the rapid evolution in software-driven science, academic institutions have not kept pace with the growing need for software engineering skills. In a survey of data science faculty conducted by the University of Washington, 50% of respondents said they estimate their software development needs by hiring graduate students, often without any formal training in software engineering. Hiring students seem to be the best practices for setting software at academic institutions. Unfortunately, this approach produces software that can be beyond the ability of a typical PhD. Further, software developed in this way often cannot be used by a broader community. Training in software engineering for over 15 (computer science) students has begun to emerge within some university CS programs, but as we note in later sections, these programs are not focused on the needs of non-CS researchers because they do not emphasize the practice of software engineering. Addressing this challenge has implications for workforce development within the United States and the creation of a software engineering talent pool within the research community.

At times, our discussions can resemble lectures that may be unfamiliar to some readers. To make this article accessible to a wider audience, we include a glossary of terms in Appendix A. When a glossary term is first introduced, it is underlined (e.g., **continuous integration**).

The lack of software engineering expertise in academia has been a long-standing problem that has emerged over the coming years. To address this, we believe that researchers should have the skills to develop what we refer to as **the 3Rs of software engineering**:

- **Readability**: We mean that software is written to promote understanding by a human. Readability is important because it promotes reproducibility (successful is providing scientific results that are reproducible by others. Readability is also required to make software extensible and reusable.
- **Resilient**: We mean that a system fails rarely or in the manner of software systems, that when it does fail due to adverse impacts or failure in components of the system, it does so gracefully (e.g., it operates even when the system is degraded). Resilience requires testing for common failure modes, but it also requires that software be written so that users and developers understand how they can rely on the software.
- **Reusable**: We mean that others can make use of the code as is, without extensive rewriting (e.g., through the use of modular structures). Beyond readability and resilience, reuse requires that software be modular and distributed in a way that is easy to install and use.

As we discuss in this article, not all academic software needs to be at the high end of the 3Rs, but all developers of academic software should have the understanding to produce software with the 3Rs.

A theme throughout this article is employing an appropriate level of software engineering for the project scope. Milewicz et al. (2019) notes that the vast majority of academic software projects have a single developer and a single user. Such sole projects are common in modeler thesis work and for software written to analyze data for a publication. A modeler number of projects have multiple users, typically software written for a research lab such as a data reduction pipeline. A third portion of projects are in widespread use in a research community. For these projects, significant software engineering rigor is necessary.

Project scope can change, both expanding and contracting. As such, we strongly believe that all developers of academic software should know how to build 3R software, even though high-end 3R software is unnecessary in most academic projects (Section 2.1). Unfortunately, software engineering in academic courses reach outside of CS and related fields (Pan & Petersen, 2020).

In this article, we further address how to bring more software engineering skills into academia through improved training programs, and the creation of better organizations to research software engineers (RSEs), professionals in software engineering working in an academic setting. There is a growing recognition of the need for RSEs. Our indication of this need is the creation of the United States Research Software Engineering organization (URSE, n.d.). Feeding the RSE pipeline requires universities to develop courses and programs to train students in software engineering practices. Here we consider approaches and challenges with developing these courses at the University of Washington (UW) and elsewhere. Finally, we explore what a future software engineering and software engineering community might look like at a research university.

2. A Tale of Two Academic Software Projects

To give more insight into the nature of and challenges with developing software in academic settings, we discuss the history of two academic software projects. These projects are original to that their software is in use by a broad community, and both evolved over a period of two decades. As noted previously, most academic software is developed and used by a single researcher. That said, the success of the projects described below is in part due to adopting tools and best practices for creating 3R software. Some projects should be aware of how they can evolve their software to potentially address a broader community.

2.1. Tulkarm

With the advent of computer generic sequencing and related advances in laboratory technique, biological science and engineering has increasingly focused on quantitative techniques. One aspect of this is the development of mechanistic models that predict cellular processes, like the concentration of glucose and related metabolites. These models are expressed as a network of chemical reactions. Each reaction describes the transfer of mass from one chemical to another. For example, glucose is converted into pyruvate. A model requires evaluating a system of nonlinear ordinary differential equations. Historically, such models were developed as special purpose simulation written in the C or FORTRAN programming languages. This was problematic because (a) few academics could program in C or FORTRAN and (b) there was little reuse of common components (e.g., transferring a reaction network into a system of differential equations). At time passed, there was a recognition that biochemical simulators could be developed more rapidly if there were tools that incorporated commonly used capabilities.

The Tellurium project (Choi et al., 2018) is the latest evolution of a long standing effort by the Ops-Sys Research Group at the University of Washington. The first software in this project was Jarnac (Hellerstein, 1994). JARNAC project (Choi et al., 2005) tried to make this modeling accessible to bench scientists by allowing users to enter reactions in chemical notation instead of having them write computer codes. Unfortunately, the system was quite slow. Thus, the project evolved into two parts: the RoadRunner simulation engine (Somogyi et al., 2015) and the Antimony (Smith et al., 2009) model specification language. This strategy proved successful. The project expanded to support Windows, Mac, and Linux, and it provides clients for Python and C.

As the project scope evolved, so did the people within it. In the beginning, this was a solo project by Herbert M. Sauro (2000). It was a critical period because the future success of the project depended on a single person. Later another person was added to build Antimony. Subsequently, others were included to support analytic platforms.

Tellurium is a software framework evolved in parallel with the project scope. This was necessary because a single developer could not create all required components. Some complexity (e.g., storage, numerical solvers, and plotting) was delegated to other software. Some of software contributed by the teams developers. The Jupyter capabilities of ipy (Lustigman & Berkhout, 2012) proved critical here. Early on, unit tests were not used at all, but the project has added extensive unit tests in recent years, and this has improved resilience. Since Tellurium is intended for programmatic use, detailed documentation is essential. A very recent change is using Read the Docs (n.d.) for software documentation along with an automated system for documenting Tellurium application programming interfaces.

2.2. Astropy

Astrophysics has a long history of community software development extending back to the 1980s. These initiatives have been orchestrated by large organizations (e.g., those funded through government support such as IRAF [Image Reduction and Analysis Facility], Tody, 1986, and Starlink, Curtis, 2014) and by smaller groups of users who both domain-specific applications and software packages (for example, see astronomy, gateway, Deul et al., 2019). More recently, a third development strategy has emerged in which packages (or frameworks) of larger scope are created through the integration of smaller packages (for example, see Python and Wolfire et al., 2012; Jenkins et al., 2013, as examples of such a strategy).

Astropy was motivated by the size of Python, as a bigger Fortran is astronomy. A central goal of the Astropy Project was to provide consistency and completeness for common calculations and tools used by astronomers. Examples of these tools include unit conversions, the manipulation of sky coordinates (e.g., transforming from Galactic coordinates to Right Ascension and Declination), and software to read and write common astronomical data formats.

The packages integrated into Astropy were, in large part, developed by researchers well versed in software engineering practices. For example, constituent packages made use of version control via GitHub, had unit tests for the core libraries and functions, and moved full documentation as a means of developing new features. Packages were distributed using common software repositories such as PyPI (Python Package Index). Tools for continuous integration (Travis CI, n.d.) were used to automate testing and the development of Astropy to improve the reliability and robustness of the software. This solid engineering foundation greatly facilitated the construction of Astropy and its adoption by the community.

Even with this engineering foundation, package integration was restricted because of the need for common abstractions across packages. An example of this is the sub-package astropy.units, which provides a representation of physical units used in astrophysics, enables translations between units, and has the ability to decompose complex (inclusive) combinations to SIgb Astronomical units (see Astro units (i.e., Astro) Jan). As with many early Astropy packages, the astro package was developed from an existing package (Astro) that was written by a single astronomer. Astronomer Astro was and was then assumed to support the needs of the broader astronomical community. The lack of existing standards within astronomy for units led to the inclusion of all available standards within the package to make it as general as possible. Functionality to translate between conventions enabled the astro package to provide general support without forcing the community to agree on a set of standards. This was a key philosophy underlined many of the Astropy design choices.

Building the community of Astropy users, maintainers, and developers required convincing astronomers with little or no formal training in software engineering to adopt these standard tools and procedures if they wanted to use Astropy. This was a challenge that was overcome through a combination of education (the core community was supported by the Astropy development themselves. The availability of GitHub, on which Astropy was built, provided the tools and infrastructure on which to develop community and engineering practices for version control, issue tracking, and communication. The availability of infrastructure tools and repositories such as GitHub are key to the sustainability of software projects in astronomy and enable common approaches to the adoption of software.

3. 3R Software Engineering Practices

The foregoing academic software projects, while different in size and scope, show how good software engineering practices can increase the adoption and trust of software packaged-based the developers who wrote them. Based on these and other experiences, we have developed a set of practices and a set of recommendations for software engineering practices that aid in the development of 3R software. We do this with great humility since software engineering is a field with a long history and a vast literature (see Boehm, 1976, Chen et al. 2002). A great starting point for this literature is an article altering (Tsipenyuk et al., 2005). By engineering practice, we mean a collection of related activities and or building, evolving, and managing software. We have organized our discussion into three topics: software design, testing, maintenance, and distributing software. We use the term software to refer to work products of software engineering, especially code, documentation, and data. We note in passing that these activities produce other artifacts, such as predictive models and analytic pipelines. These artifacts are beyond the scope of this article.

Our recommendations come in part from the experience of members of the eScience Institute at the University of Washington and discussions, otherwise has more than 10 technical staff, almost all of whom have a PhD in a primary discipline such as Computer Science, Physics, Human Centered Design, Statistics, and Genomics. Technical staff are engaged in many educational activities including teaching, mentoring, and developing curriculum, orchestrate and/or instruct Software Capacities. eScience oversees the development of data science curricula and courses at the graduate and undergraduate level. We also have an extensive outreach program that ranges researchers, our Winter Institute in which domain researchers dedicate 16 hours per week for a quarter to work with a member of eScience technical staff. In the summer we run a program for matriculating undergraduate and graduate students from across the country (the eScience pipeline) to introduce data science practices to social good (2016). Beyond teaching, we have developed a program for hiring undergraduate and graduate students who work focused consultations with technical staff. These interactions provided us with extensive insights into the challenges encountered in academic projects of varying size and duration.

The engineering practices we propose combine the collective insights of eScience technical staff with feedback from subsequent papers in industry and at national laboratories as well as researchers who track software engineering as an academic environment. Where possible, we point to published recommendations and those we work from the field of software architecture and development practices. The vast majority of the literature focuses on team processes for complex projects, with only a modest number of papers on software development by individuals. Indeed, most effort seeking to maintaining and extending existing academic software. Formal publications in this area include: Beck et al. (2002), Fowler et al. (2018), Hunley et al. (2003), and Fitzpatrick & Fitzpatrick (2010). In addition, detailed guidance is available in numerous blog posts (Fowler, 2022 and Beck Cheryllee Cagley, 2018), but the prominence of this advice is limited.

Even our experience, since many academic projects are short-term and of small scale, it is appropriate to apply a level of engineering rigor that is commensurate with the project scope. Accordingly, recommendation is, however, to highlight the requirements for transitioning a project to a larger scale. Software development practices in academic environments and practices may evolve as a project moves from a single developer, or an entity a self-contained team, and then possibly to broad adoption by a research community. Since our interest in this data science, we focus on Python and R, the most widely used languages in data science.

3.1. Engineering Practices and Their Interactions

A central theme in the article is the engineering practices should be scaled to the project scope. That is, in smaller projects, a practice may be greatly simplified or almost altogether. However, if a project grows, there must be awareness of how to incorporate engineering practices that were not implemented in the past. In this section, we discuss how engineering practices need to evolve as projects grow. Although we have recommendations for what practices to change, there is less agreement from the user-base on how similar should trigger a change in engineering practices or how to build consensus within a developer community to adopt these changes. We do provide some advice on this topic related to the adoption of developer tools (Brooke Jordan, 2014) and security analysis (Jaspan et al., 2018).

Broadly, there are technical and people management activities (which we use synonymously with practices) within software engineering. Technical practices produce code, data, and documentation of the software intended for production use. People management practices include project planning and resources (e.g., meetings), and packaging and deployment. People management activities address coordination and communication within the project and communication between project developers and users. The people management practices have associated artifacts as well, such as project plans and prioritized lists of features and bugs.

The nature of engineering practices depends strongly on the scope of the project. An example of a project with a small scope is a short-term exploratory effort by a single researcher. In contrast, a project with large scope often involves multiple teams at different locations. We consider three project scopes:

- **Developing for your own use** (ad-hoc): Our experience is that the vast majority of academic software projects consist of a single developer who is the sole user of the software. These projects are often undertaken as part of a research exploration. Few academic projects advance beyond this stage.
- **Developing for a lab or other research lab** (lab): Many software projects are of this scope. They often find that the problem addressed by their software can be used by others in their lab. In these projects, developers and users are in frequent contact.
- **Developing for a broad research community** (community): There are a small number of projects that are used by a broader research community and/or of sufficient technical scope that a large team is required.

We emphasize that the boundaries between these scopes are fluid. For example, a sole project may evolve into a lab project, and the reverse can happen as well.

3.2 Details of Engineering Practices

There is a vast literature on software engineering and engineering practices. In this section, we describe a subset of these practices that we find are most relevant to data science. These practices relate to: version control, design, testing, continuous integration, code reviews, documentation, project management, team management, and user engagement. We organize the discussion by project scope (including some references to more in-depth discussions of these software engineering practices). For each topic, there are two tables. The first describes what practice is used and why it is important; the second lists useful tools recommended tools and best practices.

We begin with **version control** (Blokhin, 2022).

- **What**: Version control deals with tracking changes to artifacts (e.g., code, documentation, and data). Version control is a fundamental practice and is essential to a central part of making software resilient and reusable.
- **How**: For software, services such as GitHub (Dreamhost, 2022) GitBucket, n.d.) allow users to have a **repository**, where changes can be viewed. Commonly used features are **pull** (adding a change that introduced an error and fix) coordinating changes among multiple developers. Another widely used feature is a service control "branch" that allows developers to make changes without affecting the main version. Continuous integration tools and project unit tests control to ensure that the code is not lost and it seem to previous versions if a bug is introduced. They enable experimentation and exploration of new ideas without impacting the primary or main branch. A lab project has additional requirements, such as resolving "change conflicts" (changes to the same line or a file by different developers). In a community project, more formal coordination is done to manage releases, develop new packages or features (such as the related tools Jira (n.d.), Jenkins, integrate codes from other groups, and handle urgent bug fixes (Tan Chee Chee and Alexander Kuznetsov, 2020).

Software design (Brooke, 2015) is at the core of creating resilient (e.g., by multi-considerations of error conditions) and reusable software (e.g., by a modular design).

- **What**: Software multiple components to software design of which we

Cookie Policy

consider two critical for data science applications. First is the design of the user experience or use cases, often referred to as functional design. This is about how a user interacts with the system to accomplish their objectives. The second is component design. This specifies how to create and interconnect software artifacts that perform the use cases.
- **How**: Appendices B and C contain simplified templates that we developed for functional and component design that we use in CSE 583 (University of Washington, 2023-b) and DATA 515A (University of Washington, 2023-c) at the University of Washington and in CHEME 545 and 546 (University of Washington, 2023-a). The functional design specifies a set of use cases that are detailed descriptions of user interactions with the software system. Appendix B contains a template for functional design. A component design can be expressed in many ways, such as: a data flow diagram (Li & Chen, 2009), UML diagrams that describe objects with properties and behaviors (Fowler, 2003), and entity-relationship diagrams (Li & Chen, 2009). Appendix C contains a template for component design. In solo projects, design may be done informally (e.g., in a notebook). In a lab project, there is often some discussion that requires a shared white board and sometimes an informal write-up. In community projects, more formality is required, such as a standard template for function and component design documents.

Computer programming or coding (Kerninghan & Pike, 1999) is the process of writing detailed instructions so that a computer can perform a desired task.

- **What**: Coding practices encompass a wide range of decisions, such as: (a) the choice of names for variables, functions, and modules; (b) documentation practices in modules and functions; (c) the use language features (e.g., the choice between a `for` statement versus a comprehension in Python); and (d) the choice of data structures (e.g., when to use a `list` vs. an `array` vs. a `dict`). A further consideration are software licenses (Laurent, 2004).
- **How**: Over the last 50 years, programming has evolved into a systematic engineering activity with powerful productivity tools. Examples of such tools are integrated development environments (IDEs) for Python (e.g., PyCharm; Nguyen, 2019) and R (e.g., Allaire, 2011). In recent years there has been a trend toward "literate programming" (Knuth, 1992), especially 'notebooks' (e.g., Jupyter) that intermix code with text to provide a narrative for an analysis. For a solo project, readability is greatly improved by providing notes about decisions made (e.g., a GitHub README file or within the notebook) as well as the use of consistent naming conventions to facilitate understanding codes written months earlier. In a lab project, readability and resilience are enhanced agreement on common data structures and coding styles (with tools such as linters to enforce style). A community project often takes this a step further by having code reviews in which developers explain their motivations for engineering decisions and reviewers advise on approaches to improve reuse.

In industry, ***quality assurance*** (Patton, 2005) is a very broad term that encompasses the entire engineering process.

- **What**: Quality assurance is about ensuring resilience, good performance, security, and privacy.
- **How**: For academic projects, the focus is mostly about testing for errors at various levels. For a solo project, it likely means implementing unit tests (codes that check for errors in functions and methods) for key elements of the project. Excellent open source packages are available to enable these tests (e.g., Python unittest and R testthat). For a lab project, unit tests are more extensive, and there is continuous integration (e.g., run all unit tests after every commit to the software repository). A community project likely includes additional quality tests for each software release to ensure there is no 'regression' in future releases. (See Nielsen, 2000, for a more detailed discussion.)

Packaging and deployment (e.g., Waldon, 2012) are activities that make software available to users, an essential element of making software reusable.

- **What**: The goal is to make software developed by one user available to other users. This often requires that the software developer structure their codes into a package that can be shared with other users. Users may have different software installed on their computers, even different operating systems. So, the package must specify its dependencies, such as a particular version of the Python library `numpy`. This raises a further challenge that two packages may have conflicting requirements, such as different versions of `numpy`.
- **How**: Most academic projects use an install model of package deployment in which the user's computer is updated to incorporate the software. PyPI is the most common mechanism for distributing Python packages, and CRAN is widely used for R packages. Other software repositories include Conda and SourceForge. There are also service models for software distribution in which the software runs on servers owned by the provider and users are not aware of the software updates (e.g., Gmail). Still another approach is container-based distribution (e.g., Docker). For a solo project, there may be no packaging and deployment since the code runs on a single machine for a single user, but to support the reproducibility of the research a well-defined and reproducible development environment can be critical. For a lab project, it is common that all machines in the lab run an almost identical software stack (e.g., the same version of Linux and Python packages), and often codes are relatively machine independent (e.g., Python, R); so, deployment is done via PyPI for Python and CRAN for R (with their associated packaging requirements). A recent trend is to use virtual machines in the cloud so that even if physical machines have different software, the virtual machines are identical. A community project often involves multiple languages and hardware platforms, and so packaging and distribution is more complex. One such complexity is that quality assurance must include testing of packaging and installs.

User documentation (Bhatti, 2021) is the written descriptions that accompany a software package so that nondevelopers can effectively use the software, a key consideration in building reusable software.

- **What**: User documentation covers installation, basic usage, and a detailed reference manual for advanced users. For example, a screen scraper application might specify a command line to install the tool, illustrate its usage on a page from *The New York Times*, and point to detailed documentation on options for different kinds of web pages.
- **How**: Solo projects have modest needs here, mostly to ensure that the developer can easily recall how to use their software some months after it was written (and the methods or research papers underlying its development). In a lab project, developers may provide a 'help' option for command line tools and/or a one-page summary of usage (a 'manual page,' Linux, 2009.) or a Jupyter Notebook (Ragan-Kelley et al., 2014). In community projects, there are more extensive capabilities (e.g., Read the Docs; Cotton, 2016) that contain detailed descriptions of the software features, examples, and capabilities for searching documentation.

We use the term ***team management*** (Project Management Institute, 2017) to refer to those aspects of project management that address the internals of the project.

- **What**: Examples of team management include: agreeing on common objectives, developing a plan (tasks, people, deadlines), progress monitoring, and plan evolution.
- **How**: Agile practices are widely used for managing software projects (Martin, 2003), an approach that iteratively delivers prototypes, an approach that applies to all project scopes. Little is required for a solo project beyond an individual prioritizing their activities. For a lab project, lab managers (e.g., principal investigators) may find it useful to have a spreadsheet that describes who is working on which feature and the expected completion dates. For a community project, there is often coordination across physical locations. Managing the dependencies between teams may demand the use of project management software (Nieto-Rodriguez, 2022) as well as a designated project manager or package maintainer who tracks progress of the project plan.

User engagement (Cagan, 2018) addresses interactions between the software developers and users of the software.

- **What**: Sometimes this is included in project management or product management (delivering products customers want). We separate this aspect because the role often goes unnoticed as a lab project grows into a community project.
- **How**: In a solo or lab project, user engagement typically involves a hallway conversation with a peer researcher. However, in a community project, communicating with users may require using GitHub issues, a designated email account, and even periodic user group meetings.

Table 1 summarizes the foregoing discussion providing examples of software packages that can support the software engineering practices (e.g., linters, and unit test frameworks). The rows are software engineering practices, and columns are the three project scopes: solo, lab, and community. The rigor of engineering practices increases as the scope of the project progresses from solo to community.

We use this table to recommend engineering activities for an academic environment. Our expectation is that most projects are not adequately characterized by a single column, and so we expect that projects may adjust their practices by employing recommendations from more than one column. We note that the table includes a number of technical terms. Rather than defining these inline, we have included a glossary as Appendix A.

Table 1. Software engineering practices and project scopes.

ENGINEERING PRACTICE	*Increasing software engineering rigor →*		
	Developing for your own use (*solo*)	**Developing for your research group (*lab*)**	**Developing for a broad research community (*community*)**
Version control	• Commit source code to a repository (local or remote) • Use meaningful commit messages to aid code comprehension and debugging • Implement backups, in particular, for local repositories.	• Isolate features being developed into separate commits (e.g., use branches for individual features) • Keep commit history concise (e.g., use squash, merge) • Document the development workflow (e.g., define the expected size of commits etc.).	• Create policies for fixes, new features, and spin-off projects. • Implement review criteria for accepting commit requests. • Create policies for version control mechanisms (e.g., creating forks, branches, pull requests).
Software Design	• Brainstorm with colleagues	• Review user experience via lab presentations of mock ups. • Apply best practices for software design to promote extensibility and reuse (e.g., object oriented). • Adopt code reviews (e.g., multiple developers check coded commits).	• Write functional specifications that describe: the problem addressed, user profiles, use cases. • Document components and their interactions (e.g., class diagrams and interaction diagrams). • Review functional and component specifications.
Coding	• Ensure the repository has a README file. • Use a consistent naming convention for language elements (e.g., objects, functions, classes).	• Use common data definitions (e.g., DNA sequence, light source). • Adopt a software license. • Create practices for coding style. • Use a linter to enforce coding practices (e.g., Pylint, Flake8). • Use build tools (e.g., make). • Use published standards for coding style, especially for tool support (e.g., use PEP 8 and Pylint). • Define templates for packages so that they have a common layout (e.g., directory structure).	• Implement code reviews to ensure standards, style consistency, adequate tests, and use of common code. • Create documentation reviews. • Implement pre-commit/pull request coding style checks in a linter.
Quality Assurance	• Write unit tests for critical components (e.g., using the pytest framework). • Automate local test execution (e.g., using a build-tool target)	• Develop extensive unit tests and measure test coverage (e.g., using coverage.py). • Develop system tests (e.g., end-to-end tests for use cases). • Implement continuous integration for the development platform.	• Implement regression tests for each release. • Where appropriate, implement tests for performance, security, and so on. • Test on all supported platforms. • Evaluate the software package by users of the package.
Packaging & Deployment	• Deploy software through git clone. • Version software via commit labels/tags. • Deploy to established package managers (e.g., pip, conda, conda-forge)	• Use semantic versioning (or tagging) conventions to indicate the magnitude of change from previous versions. • Tag and release versions of the code.	• Create release documentation. • Implement automatic documentation builds. • Use automatic builds for container deployments.
User Documentation	• Document or comment the code. • Document complicated math in a long comment or a reference to a paper. • Provide Jupyter Notebooks to demonstrate the code.	• Implement more extensive developer documentation (e.g., Sphinx). • Create user documentation (e.g., details of usage and options). • Provide citation information for the package (e.g., a DOI).	• Publish documentation (e.g., Read the Docs) that includes examples. • Create email lists for developers and users. • Create tutorials. • Publish on-boarding documentation.
Team Management	• Ad hoc	• Create a plan for each release that specifies delivery date, features, and deliverables for each team member, or at least a prioritization of changes and their estimated level of difficulty.	• Technical leads/contributors create release plans or create project plans with task assignments. • Implement regular "stand ups" to share progress and blockers. • Undertake technology reviews when considering new dependencies. • Define a governance structure (clear rules for decision-making).
User Engagement	• Ad hoc	• Make team presentations on proposed features. • Use issue tracking software (e.g., GitHub Issues).	• Provide a forum whereby users can request as well as give feedback on the prioritization of features and fixes. • Notify users of prioritization of issues. • Provide users with a project roadmap, training at community events, major release changes warnings, and deprecation warnings.

A particular challenge for software engineering for data science is the widespread use of computation notebooks in which software is intermixed with an analysis narrative that provides the user with a way to better understand the results as well as conduct further explorations (e.g., Kery & Myers, 2017; Wang et al., 2019). Probably the most widely used computational notebook is Jupyter (2015; Randles et al., 2017). Others have noted challenges with notebooks as a computational environment (e.g., Chattopadhyay et al., 2020). Our experience from teaching data science is that the quality of a computation notebook is greatly improved by two practices. The first is documentation. All codes should have documentation that specifies: (i) what the function does, (ii) the types of arguments to the function, and (iii) the types of values returned by the function. A second consideration is to make sure that there is a test for each function, typically within the same cell as the function definition so that the test is run when the function is defined. A final consideration is when to move code from a notebook to a Python module (a file that has extension '.py').

4. Current Practices for Teaching and Training in Software Engineering for Academic Researchers

Key to building a community of researchers capable of contributing to or creating their own open-source packages is providing courses and training opportunities that teach 3R skills for software engineering best practices. Computer science students along with students in related disciplines (e.g., electrical engineering) develop these skills as part of their multiyear degree programs. However, at present, it is very difficult for students outside these disciplines to develop the practical engineering skills required to develop open-source packages.

These difficulties arise from two sources. First, non-CS students have limited time in their schedules to take classes outside their discipline. Our observation is that most non-CS students have time for no more than two courses to acquire the requisite knowledge. Second, existing courses are typically structured so that practical skills are taught only after teaching substantial theory and background. For example, at the University of Washington, the sequence for non-CS students is Computer Programming I (CSE 142), Computer Programming II (CSE 143), Software Design (CSE 331), Data Structures (CSE 332), System and Software Tools (CSE 391), and finally Software Engineering (CSE 403). (We note in passing that the University of Washington Allen School of Computer Science is making extensive changes to this curriculum.)

Other universities have adopted similar approaches. For example, UC Berkeley has the sequence CS61A "The Structure and Interpretation of Computer Programs", CS61B "Data Structures: Fundamental Dynamic Data Structures", and CS 169A "Software Engineering." Although CS61B provides some instruction in version control, this is not part of the curriculum per se. Indeed, there is no course that provides instruction on engineering practice (Armando Fox, Associate Dean, private communication, March 8, 2022). Carnegie Mellon University has addressed the need to provide 3R software engineering skills by creating an undergraduate minor in CS. But a substantial time commitment is required to take the requisite courses. Harvard University's Institute for Applied Computational Science provides a master's degree for students to acquire 3R software engineering skills (IACS, n.d.), but this too is a significant investment of time.

Students do have options outside of formal academic courses to develop 3R skills. Software Carpentry (Becker & Weaver, 2018) and Data Carpentry are successful and scalable ways to introduce introductory-level software engineering processes to researchers from a broad set of domains. The areas covered in these courses include the UNIX shell environment, version control using git, and an introduction to Python and plotting with Python. More advanced topics such as automating the building of software and the use of databases and SQL are available. Software Carpentry provides much less depth on more advanced software engineering practices such as unit testing and continuous integration. The limited extent of the Carpentry courses (typically a few days) means that it is harder to integrate the processes within the everyday work or research by a student (particularly more advanced practices). Only through repeated use of these practices do they become embedded in the way that we work.

We have two broad thoughts about changes in academic curriculum that are required to address the development of 3R skills. First, we believe that the focus should be on undergraduate courses. One reason is that it provides a scalable mechanism to prepare students for 21st-century careers in academia and research. The other reason is that these undergraduate courses can also be available to graduate and postgraduate students who need 3R skills in software engineering. That is, our focus is on undergraduate courses, but the training will be done at all levels in the university.

Our second thought is that the courses for developing 3R skills need to be radically redesigned. At present, these course sequences are a lightweight version of the material taught to CS undergraduate majors. That is, courses early in the sequence focus on theory; only toward the end of the course sequence do students acquire 3R skills. We recommend that the material be restructured so that 3R skills are taught (and practiced) early on. More advanced courses in the sequence should provide greater sophistication in areas such as programming (e.g., abstraction techniques) and data structures (e.g., complexity analysis). There are a couple of examples of a first course in such a sequence. At the University of Washington, CSE 583 "Software Development for Data Scientists" (Beck, 2018) is a one-quarter course on software engineering for non-CS graduate students that covers all of the engineering practices described above and includes a capstone project to practice these skills. At Carnegie Mellon University, Crafting Software (CMU, n.d.) [illegible]

We close with some thoughts about the collection for CSE 585. The intent of this course is to develop SE skills for students who have little programming background. Any topic area of software engineering, examples centered with GitHub, the back-command line, constructing Python modules, unit testing, GUI development, data science, DevOps, Docker containers, CI/CD pathways, continuous integration, and more processes. Some processes include code reviews, technology reviews (how to choose a software dependency), and project planning. Also the topics are addressed individually, students gain practice in their use by doing a class project with a team of three to four students.

5. The Future of Software Engineering for Academic Researchers

One major distortion we see promoting at the University of Washington is to develop a community of practice for software engineering, a group of experienced software engineers who are willing to mentor students and focused engagements on academic software projects. Developing such a community of practice is a long-term goal, but it is the primary way to reduce software engineering practices in research and SE teaching software engineering courses on campus. Although funding exists to support software engineering positions, resources to give gives research group who are in the form of fractional FTEs (full-time equivalents). It is often difficult to fill fractional FTEs even if we can retire and retain people with the appropriate expertise for these positions. This has motivated us to develop a centralized software engineering group, which is a long-term goal. Such a centralized group is useful to clients on campus, but it is beyond the scope of this article. Yet, well, we highlight a number of existing models and programs in universities.

The reluctance distortion at the University of Washington is an example of such a centralized software engineering group. The University of Illinois has a long history in the development of their own research. Such institutions could leverage the size of the university to build software engineering teams as a university-level program to support researchers. In particular, as institutions can make use of fractional FTEs of grant support to amortize the cost of skilled software engineers who could not be funded by any one group. Such organizations have been realized in the United Kingdom with the development of university-level software engineering organizations. Such organizations have been shown to be recognized by other universities in the United States. Other examples of nascent programs include software groups who ended within a project for a months or years (e.g., Cohen-Gadol, California Institute of Technology, 2022), cohorts of master's students with data science and software engineering skills who are matched with research projects for short engagements (e.g., MYI 370, NYU Center for Data Science, n.d.), embedding software engineers in research laboratories (e.g., Software Heritage, n.d.), Berkeley D-Lab, personal communication, March 2022), and other models such as the Research Software Engineering Group at Princeton or the National Center for Supercomputing Applications at the University of Illinois.

As noted previously, a core challenge with much of this vision is hiring and retaining skilled software engineers in an academic environment, where people who are excited about working with researchers who may have minimal training in software engineering. There are multiple avenues for such hiros: (a) long-term positions for software engineers already in academic research labs, (b) software engineers in industry who are engaged in open source communities, (c) industry software engineers returning for graduate studies who may find that their software skills provide entry into an exciting research project, (d) early-career engineers interested in research and the impact of SE in research, and (e) experienced software engineers who are engineers looking for a mentorship. A critical aspect to the success of such a program is the retention of good talent. Carver et al.'s (2022) survey found an overwhelming concern about the lack of career paths for software professionals in academia. This will require careful thought about the career paths for software engineers within the academic environment, an environment that gives a premium on publication articles, not software products. Yet, there are many software products that require providing expanding computing in academic institutions.

We have a few examples on how to convert experienced software engineers. We have learned much from hiring software engineers for the security institute at the University of Washington. Our experience is that it is easier for a career in a top industry group to transition into a software engineering position in the development of research software for science. Hence, we mostly targeted industry for recruiting software engineers.

We have several observations based on our experience over the last 6 months of hiring. First, it is easier to recruit senior software engineers who have spent a decade or more in industry. They are attracted to the mission, engineering autonomy and scope, and software components and writing—value that allow software engineers to focus on software coding. related tasks to which they intend to continue growing their skills.

The biggest advantage of software engineering in academia is the culture of openness and the opportunity to change the trajectory of a multiyear investment in science by contributing highly sought after engineering products to a dedicated community of scientists and researchers. This community impact goes beyond the organization or region to benefit society at large. We expect this to be a primary factor in retaining software engineers in academia, and in establishing the perception of research software engineering in academia as highly fulfilling career path.

6. Conclusions

Our experience at the University of Washington is that we can develop a software engineering group that is valuable to other academic researchers in usage, and valuable by embedding within other software. We refer to these considerations as the SE of software engineering.

This article addresses engineering practices that ease the development of research software, such as code reviews, unit testing, debugging, and other topics of important element. Among the engineering practices we discuss are version control, design, quality assurance, packaging, documentation, and project management.

There are robust industry practices for creating SE software. However, many of these practices are skills intensive and time-consuming. Further, although application of these practices can result in a high level of SE capabilities, this outcome is poorly matched with the needs of research software projects. Most academic projects are quite small, they consist of a single researcher who is the sole user of the software. A modest number of academic software projects address multiple users in the same lab. Very few academic projects are intended for a community of users. Thus, the cost-benefit analysis of using these best practices is unfavorable. Software engineers are more valuable when they can focus on the most of a project. These considerations led us to construction software engineering practices into a proposaum of increasing rigor to better match the needs of academic projects with different scopes.

The need for SE skills for academic software led us to recommend the hiring and training of software engineers. We provide an in-depth analysis of our institution, the University of Washington, and we provide some insights into the situations at Carnegie Mellon University and the University of California at Berkeley. We conclude that undergraduates outside of CS or other departments, such as electrical engineering) face significant challenges with acquiring SE skills because of the limited time available in undergraduate majors to take prerequisite courses and the competition in other disciplines.

We much on another path to creating SE software—building a community of practice. This is a vision of experienced research software engineers (i.e., an SE group) who are willing to mentor students and research projects. This is not an alternative to teaching and training, rather it complements those efforts. One example of an RSE team is LINCS Frameworks (n.d.) joint project between the University of Washington, Carnegie Mellon University and the LIGO Corporation to develop scientific software to analyze data from the Rubin Observatory Legacy Survey of Space and Time (LSST). A broader initiative is the recently announced Virtual Institute for Scientific Software (VISE) (Rubin 2022) that seeks to accelerate scientific discoveries through the development of SE software for a diverse set of academic projects.

A further consideration is cultural. In academia, the criteria for success is the publication of the results. In contrast, success in a software engineering culture is creating software that is widely used, and has a reputation for good quality. These cultural differences can create an "impedance mismatch" that may present challenges for an RSE team and promoting SE software.

If we can address these challenges, we have an opportunity to increase the productivity of academic researchers and improve the quality of research software and throughout the world. Doing so will accelerate the progress of research. It will also aid in workforce development by having more undergraduates trained in software development and by providing a community of practice to support the careers and advancement of those software developers in academia.

Acknowledgments

Andrew J. Connolly and Joseph L. Hellerstein contributed equally to writing this article. The whole Hugo section is suggestions the ideas of project management. Andrew J. Connolly and Joseph L. Hellerstein also made comments that greatly improved earlier drafts of this article. We greatly appreciate the thoughtful and thorough comments from the reviewers that were essential to creating a much improved manuscript.

Disclosure Statement

AJC acknowledges support from the Department of Energy through award DE-SC0021469 and the National Science Foundation through award AST-2107821. This work was supported by the Washington Research Foundation and by a Data Science Environments project award from the Gordon and Betty Moore Foundation (award 2013-10-29) and the Alfred P. Sloan Foundation (award 3-2013). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the University of Washington Scientific Software Engineering Group (NSF) and LINCS Frameworks are supported by Schmidt Futures, a philanthropic initiative founded by Eric and Wendy Schmidt, as part of the Virtual Institute for Scientific Software (VISE) and the Virtual Institute of Astrophysics (VIA).

References

Allison, J. (2021). Rivulet: Integrated development environment for R. https://github.com/rivulet-enterprise/rivulet-vsix-release/wiki/Rivulet-Quick-Start

Beck, K. A. (2018). Software engineering for data scientists. University of Washington—Paul G. Allen School of Computer Science & Engineering. https://courses.cs.washington.edu/courses/cse599/

Beck, L. L., & Perkins, T. F. (2009). A survey of software engineering practice: Tools, methods, and results. IEEE Transactions on Software Engineering, 35(5), 95, 542–561.

Becker, L., & Mosses, B. (2018, May 7). About the members of the software carpentry CoC. Software Carpentry. https://software-carpentry.org/blog/2018/05/new-coc.html

Bluett, J. (2022). Data for developers: An engineer's field guide to technical writing. Apress.

Bokofsky, G. (2022). Version control software standard requirements. SSTRACLinks

Boulos, R. M. (2016). Software engineering. IEEE Transactions on Computers, 65(02), 1594–1601.

Boyle, A. (2018, January 22). Computational biology jobs BADM competes to enlist more software engineers for research. GeekWire. https://www.geekwire.com/2018/university-of-washington-joins-kbm-campaign-to-enlist-more-software-engineers-for-research/

Cargan, M. (2018). Inspired: How to create tech products customers love. Wiley

California Institute of Technology (2022). SAS: The Schmidt Academy for Software Engineering. https://sas.caltech.edu/

Carnegie Mellon University (n.d.). Software Engineering. https://cmu.edu/ini-software-engineering/faq/

Carver, J., Weber, M., Ram, K., Gesing, S., & Rluo, D. S. (2022). A survey of the state of the practice for research software in the United States. PeerJ Computer Science, 8, Article e963. https://doi.org/10.7287/peerj.preprints.3456v1

Chatzopoulou, A., Prasad, L., Nicolay, A. E., Sarma, A., & Bayh, T. (2021). What's wrong with computational notebooks? Pain points, needs, and design opportunities. In A. Bruckman, G. Mark, S. Fiesler, & J. Zaidman (Eds.), CHI '20: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (2021). ACM.

Chen, K., Medina, J. R., Aivazy, M., Stocking, K., Smith, L., Gu, A., & Sauer, N. M. (2018). Software: An extensible Python based modeling environment for systems and synthetic biology. Biosystems, 171, 66–79.

Cotton, N. (2016). Working documentation ways to read with Read the Docs. opensource.com.

Cuomo, M. J., Berry, D. A., Jensen, T., Gibb, A. G., Bell, G. A., & Draper, F. W. (2014). Starlink Software in 2015. Astronomical data analysis software and systems XXIV. ASP conference series, vol. 495. https://ui.adsabs.harv.edu/abstract/2014ASPC.495..247C

Dell, J., Zanin, R., Lefauchaux, J., Reisner, C., Wickli, R., Steriac, R., Wand, M., Mechanics, G., Gesellschaft, M., Weiner, J., Lopez Cano, R., Klapper, S., Carrott, M., Lonsky, J. P., Acton, P., Djamaeff-Mak, A., Pita, S., Rosalie, T., Nicklaus, C., Vuillaume, V., … Arribas, M. (2019) Gammapy: A prototype for the CTA science tools. 34th International Cosmic Ray Conference (ICRC2015), vol. 395, 766. https://doi.org/10.22323/1.395.0766

Felice, J., Fragomeli, R., Strauss, S. A., Kuff, E. A. (2017). Adopting open source software engineering (OSSE) practices by adopting OSSE tools. In Perspectives on free and open source software (pp. 245–264). MIT Press.

Fowler, M. (2018). Refactoring: Improving the design of existing code (2nd edition). Addison-Wesley.

Fox, A., & Patterson, D. (2022). Crossing the software education chasm. Communications of the ACM, 65(2), 44–49.

Hans-Rodriguez, A. (2022). Project Management Handbook: How to Launch, Lead, and Sponsor Successful Projects. Harvard Business Review.

Hick, J. L., Newey, L., & Kamesh, V. (2022). Research in software engineering for data analysis. The International Journal of Software and Software Technology, 4(02), 495–506.

Hinsen, K. (2014, December 4). It's impossible to conduct research without software, say 7 out of 10 UK researchers. Software Sustainability Institute. https://www.software.ac.uk/blog/2014-12-04-its-impossible-conduct-research-without-software-say-7-out-10-uk-researchers

GeeksforGeeks. (2022, July 5). Introduction to semantic versioning. https://www.geeksforgeeks.org/introduction-semantic-versioning/

Healey, C., Freeman, P., & Porter, A. C. (2022). Ten simple rules for teaching applied programming to an authority and innovative online environment. PLoS Computational Biology, 18(1), 1.

Institute for Applied Computational Science. (n.d.). Degree programs. https://iacs.seas.harvard.edu/graduate-programs

Jasmine, J. (2016). Data science at the command line: Facing the future with time-tested tools. O'Reilly Media.

Jurgens, C., Chen, C., & Sharma, K. (2021). Understanding the value of program analysis tools. In R. Gabriel (Ed.), OOPSLA '21: Companion to the 22nd ACM SIGPLAN Conference on Object-Oriented Programming Systems and Applications Companion (pp. 944–955). ACM.

Jurkiewicz, T. R., Johnson, R., & Rodgkerg, P. (2018). Designing interventions to promote software development to adopt security tools. In R. Mobley & Chu, S. Shyong (Ed.), ICSE '18: Proceedings of the 40th ACM Workshop on Security Information Workers (pp. 59–62). ACM.

Jupyter. (2020, November 10). Project Jupyter: Rockdale Institute for Data Science. https://jupyter.org/governance/jupyter-rockdale

Kerling, M. (2021). Design of Pragmatic Bookshelf.

Kernighan, B. W., & Pike, R. (1999). The practice of programming. Addison-Wesley.

Kery, M., & Myers, B. A. (2017). Exploring exploratory programming. In 2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC) (pp. 25–29). IEEE.

Kusch, G. T. (1992). Literate programming. Stanford University Center for the Study of Language and Information.

Lawrence, A. M. (2016). Understanding open source and free software licensing. O'Reilly Media.

Li, Q., & Chen, T. S. (2019). Time-fade diagrams. In Modeling and analysis of enterprise and information systems (pp. 84–97). Springer.

LINCS Frameworks. (n.d.) Home page. https://www.lincsframeworks.org/Home.aspx

Linux Mint Project Profile. (2019). March 02. Home page. Retrieval. https://www.linuxmint.com/blog/news.aspx

Lordigel, J., & McGaughey, M. (2022). Version control with git, PowerShell and technique for collaboration software development (2nd ed.). O'Reilly Media.

Martin, R. C. (2022). Agile software development: Principles, patterns, and practices. Prentice Hall.

Milewicz, R., Pinto, G., & Rodgkerg, P. (2019). Characterizing the roles of contributors in open-source scientific software projects. In M. A. Storey (Ed.), MSR '19: Proceedings of the 16th International Conference on Mining Software Repositories (pp. 421–432). IEEE Press. https://doi.org/10.1109/MSR.2019.00068

Mulanay, J. (2019, March 18). Why you only need to test with 5 users. SFIG Mobiles Browser Group. https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/

Nguyen, Q. (2019). Hands-on application development with PyQuarter. Accelerate your Python application using practical coding techniques in PyQuarter. Packt.

NYU Center for Data Science. (2022). Data Science and Software Services. https://cds.nyu.edu/dssf/

Osman, C., & Rostrom, J. (2019). Recognizing the importance of software in research Research Software Engineers (RSEs), v.1.00 example. Publications Office of the EU.

Patterson, D. (2016). Software testing. SABRI Publishing.

Pressbooms, E. & Lockings, J. (2022). Version control with git, PowerShell tools and technique for collaboration software development (2nd ed.). O'Reilly Media.

Pupyandex, M., & Pupyandex, T. (2021). Linux software development: An agile toolbox. Addison-Wesley Longman.

Pontes, L. (2019, January 12). Best software engineer blogs for 2019. TNY https://blog.feedspot.com/software_engineer_blogs/?gc=2020-10

Price-Whelan, A. M., Sipocz, B. M., Goodman, H. M., Lim, P. L., Crawford, S. M., Conrad, O., Shupe, D. L., Craig, M. W., Doornenbal, N., Cunningham, A., Vandelbroucke, B., Donaldson, J. K., Nemec-Ramsey, D., de Val-Borro, M., Aldcroft, T. L., Cruz, K., Robitaille, T. P., Tollerud, E. J., Ardelean, C., Babej, T., … (2018). The Astropy Project: Building an open-science project and status of the v2.0 core package. The Astronomical Journal, 156(123), Article 123. https://doi.org/10.3847/1538-3881/aabc7f

Project Management Institute. (2017). A guide to the project management body of knowledge.

Rajpal-Kelley, M., Preuss, C., Kingsper, K., Klaayen, T., Isaacs, P., Fendley, J., & Rosenstreet, M. (2018). The Jupyter/IPython architecture: A unified view of computational research, from interactive exploration to communication and publication. ACM Hot Meeting Attendees, 2018, Article 44. https://doi.org/10.1145/3160487.3160492

Randles, B. M., Pasquetto, I. V., Golshan, M. S., & Borgman, C. L. (2017). Using the Jupyter Notebook as a tool for open science: An empirical study. In 2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL) (1–2). IEEE. https://doi.org/10.1109/JCDL.2017.7991564

Sacco, H. M. (2018, August 28). Department web page. University of Washington. https://faculty.washington.edu/hmsacco/

Sacco, H. M., Hacks, M., Franco, A., Wellock, C., Bolcsó, G., Doyle, J., & Kinsey, D. (2022). Next generation simulation tools: The Systems Biology Workbench and BioNetGen interoperation. DMDD, 19(1), 202–213.

Schmidt Futures (n.d.). Our mission. Retrieved April 20, 2022, https://www.schmidtfutures.com/our-mission/

Smith, L.P., Bergstrom, T. T., Chambers, D., & Dawn, N. M. (2019). Astronomy: A modular model definition language. Bioinformatics, 35(24), 5092–5094.

Sonnagy, C. T., Rosenblum, I. M., Clunies, J. A., Wong, M., Mellroy, J. R., Sven, M. H., & Sean, N. M. (2021). DiBioSimulator: A high performance DMD, simulation and analysis library. Bioinformatics, 35(22), 5101–5102. https://doi.org/10.1093/bioinformatics/btz463

Stack overflow. (2022). https://stackoverflow.co/?gc=2179

Tinguerria, C. A., Sevilla, L., Lenzen, C., & Wilson, M. (2016). Understanding and improving software sharing in software engineering research. Practical Software Engineering, 26, Article 27. https://doi.org/10.1007/s00450-016-0373-9

Tody, D. (1986, October 31). The FITS data reduction and analysis system. In S. Crillard (Ed.), Proceedings SPIE 0627: Instrumentation in Astronomy VI (p. 733). SPIE.

University of Washington. (2020). ed. CSE002 MS 584. Data science methods for data energy research and software engineering for molecular data sciences. Retrieved April 27, 2022, from https://data.uwdirect.edu/course/2020-autumn/cse-584-0001

University of Washington. (2020). Ed. CSE 490: Software development for data scientists. https://courses.uwdirect.edu/course/2020/cse490

University of Washington. (2020-1). DATA 5050: Software design for data science. Retrieved April 27, 2022, from https://courses.uwdirect.edu/course/2020/data5050

Wickles, D. (2022). A beginners guide to software deployment. CreateSpace Independent Publishing Platform.

Wang, A. Y., Mittal, A., Brooks, C., & Oney, S. (2019). How data scientists use computational notebooks for real-time collaboration. Proceedings of the ACM on Human-Computer Interaction, 3(CSCW2), Article 39.

Wikipedia. (n.d.). Project management software. Retrieved April 20, 2022, from https://en.wikipedia.org/wiki/Project_management_software

Wikipedia. (n.d.). Software update. Retrieved April 27, 2022, from https://en.wikipedia.org/wiki/Kloss_Academy

Wikipedia. (n.d.). GitHub. Retrieved April 27, 2022, from https://en.wikipedia.org/wiki/GitHub

Wikipedia. (n.d.). Continuous integration. Retrieved April 27, 2022, from https://en.wikipedia.org/wiki/Continuous_integration

Wikipedia. (n.d.-a). GitHub. Retrieved April 27, 2022, from https://en.wikipedia.org/wiki/GitHub

United States Research Software Engineering Organizations. (2022). Home page. Retrieved April 20, 2022, from https://us-rse.org/

Appendices

Appendix A. Glossary of Terms

- **artifact**: a work product of the software engineering process including code, documentation, data, design documents, and plan documents.
- **branch**: changes (commits) to code to implement or modify features within a software repository, where the changes are kept separate from other changes to the repository.
- **class diagram**: a visual description of the relationships between elements of a software design.
- **clone**: a copy of a repository.
- **code review**: a process by which software engineers share their work and receive feedback on coding style, design, and implementation of the software.
- **coding style**: the manner in which codes are written to improve readability by others, such as PEP 8 for Python. This includes naming conventions for methods and variables.
- **commit**: a change to a repository branch.
- **component design**: the specification of the software elements in a system by their inputs, outputs, and behaviors (e.g., how inputs are transformed into outputs).
- **Conda**: Package dependency and environment deployment. https://docs.conda.io/en/latest/
- **CRAN**: Comprehensive R archive network. https://cran.r-project.org/
- **dependency**: software that is required for installing a package. For example, it is common for Data Science packages to have dependencies on `numpy`, `pandas`, and `matplotlib`.
- **deployment**: the process of making a software package available to others. This can be done by download (e.g., PyPI), containers (e.g., Docker), or server-based hosting (e.g., on Amazon Web Services).
- **engineering practice**: a generally accepted method for some aspect of building software.
- **SourceForge**: Software repository and distribution. https://sourceforge.net/
- **fork**: creation of a new code repository based on an existing code repository.
- **GitHub**: an online system for version control and continuous integration. http://github.com
- **GitHub issues**: a feature of GitHub that provides for reporting, prioritizing, and tracking issues.
- **linter**: a tool that detects errors in code without executing the code (e.g., detecting that a variable may be referenced before it is assigned).
- **merge**: combining changes committed to two branches of a code repository.
- **module**: (1) reusable software with a well-defined set of features and ways to use it; (2) a Python file (i.e., as the file extension "py").
- **notebook (or software notebook)**: a programming environment in which code is mixed with a narrative text to create an easily updatable report.
- **package**: a unit of software distribution and reuse.
- **pull request**: a set of proposed changes to a branch of a software repository.
- **project management**: the process of setting goals, assigning tasks, and managing the progress of a project.
- **PyPI**: Python package installer. https://pypi.org/
- **regression test:** tests that detect that a previously implemented feature no longer works or works incorrectly.
- **repository**: a collection of files that is managed so as to capture incremental changes or 'commits.'
- **semantic versioning:** a way to specify software versions that indicate the magnitude of change from previous versions (GeeksforGeeks, 2022).
- **software design**: the process by which software is created including specification of the user interactions (user experience) and the specification of software components such as modules and functions.
- **squash**: combine multiple commits (changes) into a single commit.
- **technology review**: a process whereby a software development team selects a software package to use in their project.
- **test coverage**: the fraction of statements in a module or collection of modules that are executed when all of the unit tests are run.
- **unit test**: code written to detect errors in a software project.
- **use case**: a way in which the user employs the software system to perform a task of interest to the user.
- **user experience**: how the user interacts with the software system such as by a web browser or a command line.
- **version control**: tools and best practices for tracking changes to code, data, and documents.

Appendix B: Template Functional Specification

A functional specification describes how developers of the software system envision that the user will interact with their software. Over the past 10 years, we have taught several thousand students how to write these specifications. We provide our students with the following outline:

- Background. Describe the problem being addressed and why it is important.
- User profile. Detail who will use the system. What do they know about the problem domain? What is their computer literacy (e.g., can browse the web, can program in Python)?
- Data sources. What data you will use and how it is structured?
- Use cases. For each use case, describe: (a) the objective of the user interaction (e.g., withdraw money from an ATM); and (b) the expected interactions between the user and your system.

Appendix C: Template Component Specification

A component specification describes the components of the software and how they interact to accomplish use cases in the functional specification. By component, we generally mean higher level capabilities such as data retrieval, significant analysis capabilities, or any novel capability. Over the past 10 years, we have taught several thousand students how to write these specifications. We provide our students with the following outline:

- Software components. High level description of the software components such as: *data manager*, which provides a simplified interface to your data and provides application specific features (e.g., querying data subsets); and *visualization manager*, which displays data frames as a plot. Describe at least three components specifying: what it does, inputs it requires, and outputs it provides.
- Interactions to accomplish use cases. Describe how the above software components interact to accomplish at least one of your use cases.
- Preliminary plan. A list of tasks in priority order.

©2023 Andrew Connolly, Joseph Hellerstein, Naomi Alterman, David Beck, Rob Fatland, Ed Lazowska, Vani Mandava, and Sarah Stone. This article is licensed under a Creative Commons Attribution (CC BY 4.0) International license, except where otherwise indicated with respect to particular material included in the article.

LICENSE · Creative Commons Attribution 4.0 International License (CC-BY 4.0)

COMMENTS · 0

Login to discuss

No comments here

Why not start the discussion?