


EDUCATION

Ten simple rules for training scientists to make better software

Kit Gallagher^{1*}, Richard Creswell², Ben Lambert³, Martin Robinson², Chon Lok Lei⁴, Gary R. Mirams⁵, David J. Gavaghan^{1*}

1 Doctoral Training Centre, University of Oxford, Oxford, United Kingdom, **2** Department of Computer Science, University of Oxford, Oxford, United Kingdom, **3** Department of Statistics, University of Oxford, Oxford, United Kingdom, **4** Department of Biomedical Sciences, Faculty of Health Sciences, University of Macau, Macau, China, **5** Centre for Mathematical Medicine & Biology, School of Mathematical Sciences, University of Nottingham, Nottingham, United Kingdom

* gallagher@maths.ox.ac.uk (KG); david.gavaghan@dtc.ox.ac.uk (DJG)

OPEN ACCESS

Citation: Gallagher K, Creswell R, Lambert B, Robinson M, Lok Lei C, Mirams GR, et al. (2024) Ten simple rules for training scientists to make better software. *PLoS Comput Biol* 20(9): e1012410. <https://doi.org/10.1371/journal.pcbi.1012410>

Editor: Scott Markel, Dassault Systemes BIOVIA, UNITED STATES OF AMERICA

Published: September 12, 2024

Copyright: © 2024 Gallagher et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: KG, MR, DG acknowledge funding from the EPSRC CDT in Sustainable Approaches to Biomedical Science: Responsible and Reproducible Research - SABS:R3 (EP/S024093/1). GRM acknowledges support from the Wellcome Trust via a Wellcome Trust Senior Research Fellowship (grant no. 212203/Z/18/Z). CLL acknowledges funding support from the Science and Technology Development Fund, Macao SAR (FDCT) (reference no. 0048/2022/A) and the University of Macau (MYRG2022-00023-FHS). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Introduction

Computational methods and associated software implementations are central to every field of scientific investigation. Modern biological research relies heavily on the development of software tools to process and organize increasingly large data sets, simulate complex mechanistic models, provide tools for the analysis and management of data, and visualize and organize outputs [1–3]. Such software varies widely in its scope, complexity, and potential for re-use, from single-use analysis scripts that accompany journal publications to domain-specific packages (such as molecular dynamics simulators [4,5]), common numerical methods (such as finite element methods [6] or optimization algorithms [7]), and finally fundamental scientific software (such as the numerical methods package “numpy” [8]).

For valid usage in research, it is essential that this software is both openly available and accurately implements its intended functionality. Accessibility of code has improved significantly in recent years [9], and it is increasingly accepted that research papers should be accompanied by accurate code scripts, which are subject to peer review alongside the other methods of the research. However, this has simultaneously highlighted the role of computational science in the so-called “reproducibility crisis” [10], where multiple cross-disciplinary meta-analyses have indicated that less than half of published code may be run without errors [11–14], and as little as 5% can replicate the primary results of the associated paper [15].

This causes a multitude of negative effects on scientific research including a lack of transparency and open access [16], poor development and deployment practices [17], and a lack of executable reproducibility—where code cannot even be run [18]. This also undermines the productivity of the research software base, as any researchers wishing to use the same computational framework are then forced to re-implement this in their own software.

Beyond basic reproducibility, higher-quality software possesses additional qualities such as extensibility, reliability, and reusability. These characteristics arise from carefully designed, well-documented, and appropriately maintained code, and they enable research software to more thoroughly and efficiently support scientific progress—for example, by allowing a software package developed by one research group to be picked up by another which goes on to add additional features to address further scientific questions. The term sustainable (not to be confused with environmentally friendly software) has been adopted to refer to software that is reliable, reproducible, and reusable [19].

Given the importance of high-quality software to effective research in computational biology, there has been significant literature on ensuring reproducibility [20,21] and good

Competing interests: The authors have declared that no competing interests exist.

development practices [22–24] in computational research. Indeed, several other Ten Simple Rules articles have already provided excellent descriptions of best software development practices to aspire towards, and we point the reader to these guides on documentation [25], usability [24], robustness [26], and version control [27].

However, less attention has been devoted to specific teaching strategies which are effective at nurturing in researchers the complex skillset required to produce high-quality software that underpins both academic and industrial biomedical research. Biologists and computational researchers, even if aware of the importance of high-quality software to their research, are typically left to fend for themselves in developing the necessary skills to produce reusable software. Although training resources are available (for example, courses offered by the Software Sustainability Institute to UK-based researchers), many doctoral training programs overlook extensive formal education in effective software engineering.

Two recent articles in the Ten Simple Rules collection [28,29] have discussed the teaching of foundational computer science and coding techniques to biology students. We advance this discussion by describing the specific steps for effectively teaching the necessary skills a scientist needs to develop sustainable software packages that are fit for (re-)use in academic research or more widely. We advocate that future researchers receive extended training in software engineering, moving beyond few-day training sessions and forming a substantial and integrated portion of their scientific education. Although our advice is likely to be applicable to all students and researchers hoping to improve their software development skills, our guidelines are directed towards an audience of students who have some programming literacy but little formal training in software engineering, typical of early doctoral students. These practices are also applicable outside of doctoral training environments, and we believe they should form a key part of postgraduate training schemes more generally in the life sciences.

The following rules have been fine-tuned through generations of doctoral students at the EPSRC CDT in Sustainable Approaches to Biomedical Science: Responsible and Reproducible Research (SABS:R3) CDT at the University of Oxford. The SABS:R3 program trains doctoral students in cutting-edge, collaborative systems approaches to biomedical research, with a strong focus on computational methods. Perhaps uniquely, it provides both comprehensive training in advanced software development and software engineering to all of its students, and introduces them to state-of-the-art techniques applicable to industrially derived research within the biomedical sciences.

Students initially take a three-week classroom training program in the principles of software engineering. This course is designed to introduce important themes in software development and convey their importance; we discuss key aspects of this in the first 4 rules. National and international consortia can provide resources and expertise to support institutions in running these programs—the SABS:R3 classroom section was supplemented by training from the Software Sustainability Institute.

To make this training immediately relevant, all students subsequently undertake an industry-supported group software development project over their first year, allowing them to learn and apply their software skills in a realistic setting. These projects require students to develop high-quality software in support of some scientific or industrial project or investigation, involving a variety of fields including epidemiology, pharmacokinetics, and medical imaging. Two such projects have additionally led to scientific publications [30,31]. We have fine-tuned the design and implementation of these projects over a series of student cohorts and share some of the lessons we have learned in Rules 5 to 8, before finally discussing the output from these projects, and how they may be sustained after the students finish their training.

Rule 1: Emphasize the value of good software development

It has widely been accepted within industrial software development communities (and is increasingly accepted within academic environments) that investing additional time to uphold common software engineering practices improves overall productivity across a wide range of project scales and fields [32–34].

However, this translation of this acceptance to academia has been limited by insufficient prior training, as mentioned in the introduction, as well as a lack of its perceived value and importance. Within the results-focused setting of academia, good software development has often been considered a secondary priority—“nice if you have time, but not essential.” This may be attributed to the perceived lack of academic recognition for software outputs [35,36], diminishing the motivation and financial support for upholding good practices. This is exacerbated in the time-pressured environment of a PhD, where the limited time frame encourages a focus on obtaining immediate results over longer-term investments, such as a high-quality and reusable code base.

Much research software is initially developed for highly specialized applications (for example, a pipeline of idiosyncratic procedures to process a particular data set). In this context, reusability and extensibility may not seem essential at the outset of the software’s lifespan, but their utility may become apparent later on (e.g., several years later, if a new PhD student is tasked with processing a new data set with similar characteristics, are they able to draw upon the previous pipeline, or must they start over from scratch?). This increased productivity quickly pays back the initial time investment required and allows the easy integration of small pipelines into large and varied code bases, increasing their potential application (Fig 1). While this paper focuses on the development of large-scale software packages intended for immediate reuse, many of the skills we aim to inculcate in researchers will equally well serve them in the development of higher-quality analysis scripts.

By emphasizing the advantages of good software development with students at the start of their research careers, it is possible to dispel any preconceived assumptions that careful coding practices will inherently slow down the pace of research. Our experience, both as PhD supervisors and PhD students, is that while students are often enthusiastic to learn software development as a new skill, the benefits to scientific research of a formal grounding in sound software engineering practices may not be immediately apparent at the start of the PhD. The advantage of, for example, taking a test-driven approach is typically realized later in the PhD, as codes developed early in the PhD can be extended and built upon with confidence. Ultimately, good software practices such as automated testing and prudent refactoring can prevent bugs, flawed results, and retractions [37].

As we discuss further in Rule 8 below, these benefits can be reinforced through structured peer mentoring. Over time, this results in a body of robust and reusable open-source software underpinning the long-term research within the group, and these benefits become more immediately obvious and self-sustaining.

Rule 2: Support students to develop good coding practices

There are a variety of schools of thought for proper coding practices, and giving students a degree of freedom over the standards they adopt can encourage accountability and teach the role of such practices. However, it is essential to provide students with direction for what good code is. This will help student groups to come to a consensus that aligns with widely used programming practices, while still leaving enough room for interpretation. It is often helpful to introduce students to general standards used by existing projects/domains, before allowing them to adjust it based on the specific requirements of their work. This is a key concept in the




Fig 1. While sustainable approaches to software development can require more effort than a quick-and-dirty coding style initially, they quickly become time-saving over medium-to-long time scales.

<https://doi.org/10.1371/journal.pcbi.1012410.g001>

FAIR (Findable, Accessible, Interoperable and Reusable) principles for research software [38], where accordance with domain-specific standards supports code reuse by ensuring their code can interoperate with other codebases in the future.

When introducing coding standards, and encouraging students to think critically about how their code should be written, it can be helpful to direct them to think about points such as the following:

- Can you choose meaningful and understandable variable and function names, so that your code is somewhat self-documenting?
- What are the naming conventions in your programming language? (For example, in Python, the PEP 8 Style Guide prescribes the use of snake case for function and variable names, and camel case for class names.)
- How will you lay out your code? While aspects such as indentation are part of the syntax in some languages (like Python), this (along with the use of whitespace and allowed line length) is a stylistic choice in others.
- How will you document your code and communicate changes (such as refactoring) to others working on the same codebase?

The impact of such changes on the readability of the code (without changing its function) can be seen in Fig 2. Similarly, it can be helpful to warn students of so-called “code smells,” which suggest failures to adhere to good software practices and are associated with increased code maintenance [39,40]. Students should be asked to avoid situations such as the following (unless they have a good reason to break the guidelines):

- Functions should not contain more than 20 lines.
- Repeated code should be combined into a function for re-use.
- Unused features (created to anticipate future functionality that never gets implemented) should be removed.

```

A. def v(q, t, D):
    N=0
    while N<D & int(len(q)>0)==1:
        p=q[len(q)-1];q=q[0:len(q)-1]
        v2(p,t)
        N= N +1

```

```

B. def vacc_all(queue, time, doses):
    # Set a counter for the number of vaccinations given
    num = 0

    # Vaccinate until vaccines given is greater than doses
    # available, or the queue is empty
    while num < doses and len(queue) > 0:
        p = queue[len(queue)-1] # Get the last person in the queue
        queue = queue[0:len(queue)-1] # Remove the last person from the queue
        vaccOne(p, time) # Vaccinate one person and record the time
        num += 1 # Increment the counter

```

```

C. class Vaccination():
    def __call__(self, simulation_time: float):
        """Iterate through the priority queue removing people and
        vaccinating them.
        """
        number_of_vaccinations = 0

        while (number_of_vaccinations < self.daily_dose_capacity
               and not self._population.vaccine_queue.empty()):
            person = self._population.vaccine_queue.get()
            person.vaccinate(simulation_time)
            number_of_vaccinations += 1

```

Fig 2. Codes snippets A–C are all valid Python and implement identical functionality. However, they vary in readability: (A) is difficult to understand for anyone except the original author; (B) is much improved, but perhaps suffers in readability due to verbose commenting; (C), however, should be directly interpretable on account of meaningful variable names and sensible object-oriented architecture.

<https://doi.org/10.1371/journal.pcbi.1012410.g002>

- Standard library functions/data types (i.e., min/max functions or arrays) should not be reimplemented from scratch.
- Long parameter lists (i.e., with more than four parameters) may be collapsed into a single object.
- Functions/classes/methods should not have multiple responsibilities. If they do, their scope should be simplified by defining additional functions/classes/methods.

While these approaches may require some refactoring of the student’s code, this typically results in significant time savings over the lifetime of the code, reducing the amount of time spent debugging—the balance of such approaches is discussed by Balaban and colleagues [41]. These good practices may be developed interactively through exercises identifying antipattern in deliberately poorly written code and then rectifying the issues with it. This can also help

students appreciate the difficulties associated with working on poorly written code and motivate the importance of good coding practices before they begin working on larger-scale projects. Additionally, students should take the time to explore the features of static analysis tools and Integrated Development Environments to support their efforts to follow good coding practices (see [Rule 9](#)).

Rule 3: Advocate the role and importance of open-source and collaborative development

We encourage open-source development using sites such as GitHub ([Fig 3](#)) and find that most students are enthusiastic about sharing their work in this way. The benefits of open-source development for early career researchers [[42,43](#)], software quality [[44,45](#)], and research in general [[46,47](#)] are well documented.

This also prompts students to think about how their work may be shared, and we specifically teach the students how to package and distribute their code, including managing dependencies and publishing to package management systems such as pip or CRAN. The students also explore how they can facilitate other researchers' use of their code, such as by providing online documentation, installation instructions/executable files, and example workflows. Finally, we cover code maintenance, so that students are familiar with the complete life-cycle of a software product, and can make informed decisions on maintaining or withdrawing support for different operating systems or versions of external dependencies.

We recognize that open-source development may not always be completely possible, for example, if industrially partnered research projects involve proprietary data or methods that must be hidden, but in these situations believe it is helpful to train students to share their code to the extent possible. Collaborative work is still possible in this context, through the careful use of private repositories, or excluding sensitive data from public versions of the repository. Students are also taught about the importance of software licensing to allow the reuse of their software, and the different approaches that licenses take to redistribution and proprietization. Clear licensing and accessible software (including metadata) are both key features of the FAIR principles for research software [[38](#)].

Having a public-facing repository also means that software may attract those in the wider community to act as co-developers, which can also encourage students to take more ownership and care of the quality of the repository. As other users can see, use or even contribute to their

```
+ import math

def calculate_area(radius):
-   area = 3.14 * radius**2
+   if radius < 0:
+       raise ValueError()
+   area = math.pi * radius**2
    return area
```

Fig 3. Version control tools such as GitHub make it easy for developers to keep track of changes in a software codebase.

<https://doi.org/10.1371/journal.pcbi.1012410.g003>

code, software development principles discussed in Rule 2 such as ensuring code quality and good documentation have immediate benefits. The potential for wider involvement also prompts students to review how they might support third parties wishing to contribute, seek support, or report issues in the software. As part of this, students are guided to promote recommended community standards by providing public contribution guidelines, issue templates, discussion boards, and security policies for their software [48]. Previous teaching projects [49] within the SABS:R3 program have attracted external collaborators across the globe, as a result of effective publishing through GitHub and clear open-source contribution guidelines.

Rule 4: Use agile development practices

Among different development paradigms, so-called “agile” development has many features that are beneficial to scientific research. Project goals are typically not fully defined at the conception of a project and are often informed by preliminary and intermediate results, so the ability to easily refactor the codebase is essential. Agile development methodologies account for this through concise planning horizons (often termed “sprints” or “iterations”), working towards coding targets that can be feasibly accomplished in a reasonable time frame (typically 1 month). This approach facilitates the rapid creation of working prototypes and mitigates the phenomenon of “analysis paralysis,” a pitfall that arises when attempting to anticipate too many potential future requirements.

While there are several frameworks for agile development such as “Scrum” or “eXtreme Programming” [50], we emphasize the basic principles in our courses rather than any particular approach. The key principles are continuous integration (that is, always having a codebase that passes automated tests and is safe to deploy, even if lacking features), test-driven development (wherein tests are written before the source code itself, requiring students to consider the best interface for new code), short-range planning, and frequent code refactoring (reorganization and tidying). These characteristics provide a good framework to adapt to changing demands and evolving research questions; importantly, they produce software that is well tested—whose outputs do not change without developers realizing it [51–53].

Rule 5: Create a real-world project that is ambitious, relevant, and exciting

In general, we learn skill-based knowledge best by doing, rather than through taught lectures or textbooks [54,55]. For this reason, we developed annual group projects for students to consolidate the software development skills introduced in taught courses. These problems should not be toy projects; they should be challenging, real-world problems that have not been tackled before. We have found that this novelty motivates students to engage more directly with the project [56].

These projects should also have wider relevance outside of pure pedagogy, be it to academic research or industrial applications (Fig 4). This typically necessitates a larger scale problem than typical teaching/textbook problems, requiring a team of students to tackle it together. This resembles effective software engineering in industry, where good software is developed in large, multi-disciplinary groups. This will be in contrast with the primarily individual experience that students typically gain during undergraduate research projects, where individual scripts are written to answer a specific research question rather than creating a tool for exploring a wider field of research. These larger projects help to contextualize the role and uses of research software, and develop students’ experience and confidence in working with a large codebase. A few examples of projects we have previously run with student cohorts are linked




Fig 4. To keep students of all interests and backgrounds excited, we recommend software projects that are oriented towards appealing scientific subjects (e.g., biology, epidemiology, chemistry), rather than projects that dwell merely on more foundational computer science or algorithmic tasks.

<https://doi.org/10.1371/journal.pcbi.1012410.g004>

below, to give a perspective of both the wide range of research areas covered and the scope of possible outcomes of these projects:

- [Epidemiological agent-based modeling \(EpiABM\) software, based on CovidSim model developed at Imperial College London \[49\].](#)
- [An educational game based on the process of designing new drugs, developed in collaboration with Roche \[57\].](#)
- [Extensible quality control tool for clinical images, developed in collaboration with GE Healthcare \[58\].](#)

Ensuring that these projects have tangible outputs is also key to student engagement. The key output for these projects is a complete software package, which should be open-source if possible (see [Rule 3](#)) to enable the wider research community to engage with this work. Additionally, because the projects are drawn from real-world research questions, successful software outputs will involve insightful, novel results shedding light on some scientific problem. This gives students the motivation to apply rigorous scientific thinking to any modeling assumptions inherent in their software and to examine their software's behavior in some rigorous way (for example, by comparing its predictions to existing benchmark data or established comparator methods).

Furthermore, in SABS:R3, projects were developed in collaboration with industrial partners to tackle problems in their respective fields, and many of these industrial partners have

subsequently used the open-source research outputs in-house. The industrial partners have also supported the projects, funding access to compute resources or providing access to databases to enable these projects to have real-world impact. The impact may also be tracked more broadly, through a combination of direct software engagement (i.e., downloads or citations), or public event engagement (such as the attendance at a software workshop). Development of web-hosted user-friendly interfaces allows greater quantification of the impact through telemetry tracking the number of users, use frequency, and typical use cases.

Providing publishing opportunities, either within dedicated software journals or through scientific journals in cases where the software has been used for notable scientific applications, can also provide a concrete end goal for students. Thus far, 3 cohorts within the SABS:R3 program have published (or prepared for publication) descriptions of their software and associated scientific findings in reputable journals [30,31]; further scientific and software publications are anticipated in future years of the program.

Rule 6: Set up long-term projects for multiple cohorts of students

We have typically run multi-cohort projects, with a new cohort of students working to extend an existing codebase developed by a previous cohort in the preceding years. The benefits of this are twofold. Firstly, it allows more ambitious projects to be tackled which might not be attainable in a year (Fig 5). Secondly, developing someone else's codebase is an essential software skill and encourages students to write their code with future developers in mind.

While most teaching in academic settings focuses on developing a new codebase, this is rarely the reality of software development in academia or industry, and it is common for new starters to struggle to adapt to working on legacy code instead of developing new projects [59]. Providing experience of this in a teaching setting (and often with the luxury of access to the previous developers—see Rule 8) gives students the skills required to implement changes in long-standing research software codebases.

Projects which merely involve the maintenance of an existing codebase, without adding any new features, are in our experience less likely to spur significant student interest, so we recommend that legacy software projects always involve adding some new features.

Multi-cohort projects can also adapt to changes in the project motivation—for example, a project developing broad epidemiological inference software at the start of the COVID-19 pandemic [30], was redirected to develop agent-based modeling software [31] following the high-profile use of this approach in UK governmental response to the pandemic [60]. This software package was recently extended by a subsequent year group to interface with spatial data repositories and integrate accurate population information into epidemiological models [61]. This approach ensured the relevancy of the software being developed, and that it would continue to be supported and maintained after the original developers graduated from the teaching course.



Fig 5. While all year-long projects should produce a complete and usable product, these can then be further extended by subsequent cohorts. Such extensions should not be trivial “bolt-ons,” but rather fundamentally extend the functionality of the software output, which may require students to modify/refactor existing elements of the codebase.

<https://doi.org/10.1371/journal.pcbi.1012410.g005>

To develop skills in software maintenance, we also encourage students (even those not involved in peer supervision (see [Rule 8](#))) to revisit any maintenance-related concerns that arise after they complete the training program.

In addition to this focus on developing preexisting codebases, we recognize that the process of embarking on a “greenfield” software project—where students must create a new software library from scratch—still provides valuable lessons in how to set up the fundamental architecture and collaborative infrastructure of a software project. We have therefore found it valuable to set aside the first few days of the training for a greenfield “mini-project.” This approach also allows instructors to gain a detailed insight into student ability levels, strengths, and weaknesses, which can guide the goals and organization of the primary project.

Rule 7: Encourage collaboration within mixed-ability groups

Software development is rarely a solo endeavor, and this should be reflected in the teaching environment. Group work encourages students to learn from each other’s coding styles, as well as “soft skills” such as communication and teamwork within the specific context of software development.

Students are not a homogeneous group; however, they will enter any program with differing levels of experience. While this can be seen as a disadvantage, collaboration within groups allows students to share their various individual skill sets, becoming their own teachers as they share their particular specialties [62]. For example, students with more experience can gravitate towards leadership or design roles, while the less experienced students can learn from more experienced, for example, through pair programming [63,64]. Ideally, groups should be large enough to contain students with a variety of skill levels ([Fig 6](#)) and enable some degree of specialization between students, but small and cohesive enough for all students to maintain familiarity with all aspects of the project. We have found that groups of 4 to 5 students work very well.

These groups also give students a strong motivation to develop project management skills via planning discussions and delegation of responsibilities among themselves. Although formal training in project management techniques may be available at some institutions or from other sources, such training is typically classroom-based and often optional. This means that students will benefit from consolidating these skills in a practical context, to refresh concepts they may have previously encountered in taught modules. Group software work also requires the use of version control software, and developing skills such as branch management, bug tracking, and code review that students are taught at the start of the course, but may not have applied during solo projects.

With larger-scale projects (see [Rule 5](#)), it is inevitable (and indeed desirable) that projects will be modular, and readily broken into separate tasks. Despite this, it is important to ensure that the different aspects of the project still interact where possible so that learners are not working in isolation on their sections. This interaction often happens explicitly (for example, when developing a computational pipeline where elements must interface effectively), but ensuring individual students are not isolated working on separate aspects of the codebase is crucial to ensuring that the students learn from each other’s experience. Peer-to-peer learning and assessment may be strengthened by promoting a shared responsibility in code review, developing teamwork and communication skills within a computational environment as well as exposing students to new approaches and design patterns in their peers’ code. These approaches also help avoid premature specialization, where students only learn about their own features and miss out on skills from other areas; similarly, all learners should be comfortable documenting and testing their code, rather than leaving this to a single individual.

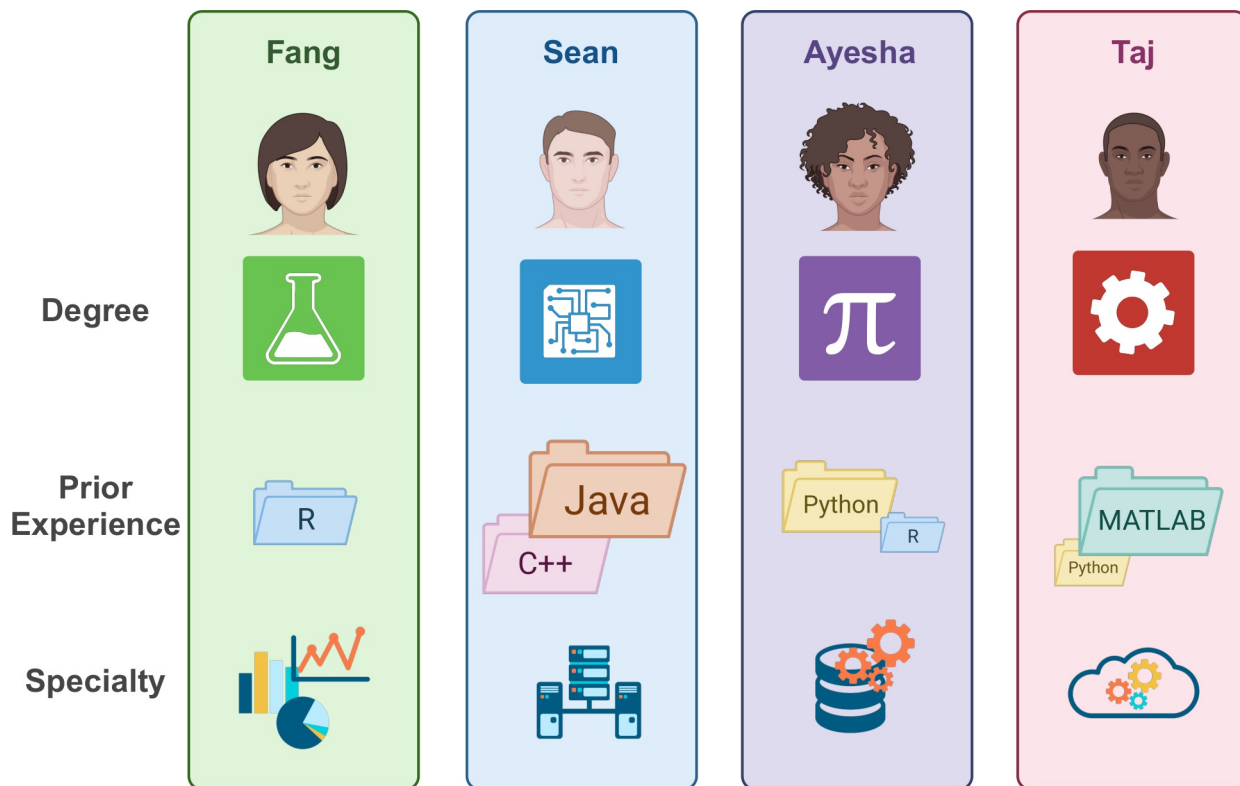


Fig 6. Cohort heterogeneity can take many forms, such as previous academic degrees/professional experience, prior fluency in different programming languages (which may or may not include the teaching language of the course), and familiarity with wider software development skills. Course leaders can take advantage of this range when designing and supervising projects so that all students can both utilize their own strengths and learn skills from others. Created with [BioRender.com](https://www.biorender.com).

<https://doi.org/10.1371/journal.pcbi.1012410.g006>

We have found it most effective for students to have the chance to work together in person, enabling collaborative activities such as regular stand-up meetings and pair programming. When students are working remotely, appropriate collaboration tools should be employed and project leaders should ensure that all students remain actively involved in the project (e.g., via holding regular group and one-on-one meetings, or monitoring contributions to identify any students who might be withdrawing from the project), for the benefits of collaborative activities such as pair programming to still be realized [65,66]. For example, we have found it effective when students engage in pair programming remotely using the screen-sharing facility offered by most video conferencing platforms, possibly in conjunction with the live-coding feature present in some modern IDEs such as VSCode.

Rule 8: Organize supervision from former students and peers

Practical courses at a graduate level are often limited in scope by the availability of senior academics to offer supervision. Instead, more senior students on the course can act in supervisory roles for these projects—able to provide hands-on supervision, and a more informal resource (away from senior academics) to answer questions about unfamiliar aspects of software development (such as unit testing). Drawing students from previous cohorts that worked on the same project is particularly advantageous (and is a proven technique in open source development: used, for example, in Google’s Summer of Code [<https://google.github.io/gsocguides/mentor/>]), as this ensures that student supervisors have a strong familiarity with the

underlying codebase. This relationship is also beneficial for the supervising student—as well as learning through their teaching they also have first-hand experience of the struggles new developers may have picking up their codebase and how they could mitigate this in future projects.

This supervision also mimics professional structures common in industry, where junior developers will typically submit work to senior developers for code review. Another good practice for learners (and supervisors) who may later wish to transition to careers in the tech sector is the formal use of version control features and code review. Code review can also help to update other members of the team on features within the code and formalize the protocol for introducing significant changes into the main codebase, including those that affect wider functionality.

Encouraging peer code review, where learners begin to review each other's pull requests, is also highly beneficial to ensure continual feedback on code outputs. This form of comprehensive peer supervision can be supported by continuous integration metrics such as test coverage and code quality, reducing some of the supervision requirements in student projects by automatically highlighting issues in the students' code that they must resolve themselves [67]. Peer-to-peer feedback can also be used to derive assessment metrics of student performance [68].

Rule 9: Focus on the process, not the final outcome

Although projects should be ambitious and exciting (Rule 5), the final outcomes of these projects should never come at the expense of the learning process. Course leaders should endeavor to foster a creative environment where students are focused on achieving the desired functionality in the best way possible (rather than the most functionality possible), enabling students to produce higher-quality outputs and develop stronger software development skills to take forward into their future projects.

Students will gain more from having the time and freedom to write higher-quality code at a slower pace, rather than rushing to complete targets. For this reason, we do not recommend formal assessment (i.e., grades) of student contributions to software training projects; instead, informal and constructive feedback on students' rate of progress, strengths, and areas for improvement should be provided by instructors and mentors. Meanwhile, students should also be continuously receiving detailed, practical feedback on their code via code review from their peers or course demonstrators.

This focus on personal development also provides students with a unique opportunity to learn features and patterns of the programming language with which they are unfamiliar, in a way that is rarely achieved elsewhere in their studies when working to deadlines. However, students should avoid using this freedom to delve into obscure, esoteric, or overly academic coding styles, and they should be encouraged to keep the practical goals of the project in mind as they explore more advanced concepts in their programming language (see [Rule 2](#)).

If the initial project aims turn out to be too ambitious, projects should be rescoped to focus on some smaller goals that the students can more confidently achieve without ever feeling the need to forego good software engineering practices. Alternatively, students can be divided into different streams working on different aspects of the project, or initially be assigned an easier toy project to develop various software development skills before launching into more ambitious year-long projects where they learn to balance best practices with time and resource limitations.

Another key benefit of this teaching program is the opportunity to explore and develop familiarity with common software development tools, as proficiency in these will boost their productivity throughout both their PhD and subsequent career choices in computational fields. As introduced in Rule 7, one key aspect of this is familiarity with version control tools

such as Git, and the importance of this skill has been addressed in “Ten Simple Rules for Taking Advantage of Git and GitHub” [27]. Students should be taught and encouraged to utilize features of GitHub widely, for example, to document bugs and planned features through issues or peer-review each other’s code through pull requests. An additional advantage of students engaging with GitHub features is that they will build up a public history of code commits, issues, pull request reviews, and other software-related tasks; such track records are extremely valuable evidence of software-related competency during the hiring process for both industry and academic jobs involving software development.

Similarly, familiarity with the functionality of Integrated Development Environments (IDEs), including built-in debugging and refactoring tools, is important for efficient software development and is associated with students’ improved learning and productivity [69]. Learners may also find it helpful to use static analysis tools, such as PVS-Studio (for C, C++, and Java) or flake8 (for Python), which provide feedback on code style to develop awareness of clean code practices.

Furthermore, outcomes may not simply be limited to the academic output or even the function of the code. The importance of software development practices can be emphasized, including software development goals as well as functionality—for example, a group may target publishing complete documentation online, or 100% unit test coverage. When a software project is used as part of a scientific publication, the quality of the underlying software design and engineering (for example, in user-friendliness, extensibility, and reproducibility) should be highlighted in the publication.

Rule 10: Build a community to “future-proof” your project

Our experience over the last 2 decades highlights the potential benefits of providing advanced training in software development in the context of a “live” applied research project. Our first experience with this approach was in 2005, when we used the initial development of the Chaste (Cancer, heart, and soft tissue environment) physiological modeling software [70–73] to teach test-driven, agile software development practices to a cohort of 10 PhD students and several postdoctoral researchers. Over the intervening period, Chaste has become one of the leading such platforms and is in continual use by research groups across the world (Fig 7). Over 50 early career researchers have contributed to its development, and several of the original Chaste developers now use it within their own research groups in academia, industry, and at regulatory authorities.

Although this success is unlikely to be replicated in every such training-driven and hence student-led development project, our similar experience in the development of PINTS (Probabilistic inference for noisy time series) [74] suggests that this approach has the potential to deliver sustainable and reusable software platforms to the research community with surprising frequency. The key here, we believe, is the generation of a community of researcher-developers who take collective ownership of the code and, through joint publications and frequent code releases, have a joint interest in its continued maintenance and use. This community stemmed from an early focus on developing the user base of this software from the original developers to a wider, multi-institutional community. For further details on this development, we refer readers to excellent Ten Simple Rules articles that have previously addressed community-building [75], and promotion/sponsorship [76].

Communities of software users, developers, and maintainers can take many forms and may range in scale from local collaborations within a university up to the more global network we observed with Chaste (Fig 7). Although such communities will inevitably tend to develop alongside existing collaborative relationships, it is also possible to start building connections



Fig 7. We downloaded the first 50 citing papers of Chaste [72] according to Google Scholar (ordered by relevance), and, for each paper, we plot the location of the city of the first corresponding author's institution as an orange dot, with orange lines indicating the great circle path from that location to Oxford, England where Chaste was originally developed. Public domain continents data from Natural Earth.

<https://doi.org/10.1371/journal.pcbi.1012410.g007>

with developers in other regions of the world using online, collaborative development tools such as GitHub (see [Rule 3](#)).

Building such communities can future-proof a code base, and perhaps suggest a new and more sustainable model for the development of the computational tools and software that increasingly underpin research.

Conclusions

Instilling good software development practices in early career researchers is invaluable in ensuring the impact of their computational outputs, but software development does not always receive as much attention in student scientific computational training as it deserves. Our experience running doctoral training programs focused on software development for 15 years has shown that this is best achieved through active learning on ambitious group projects targeting real-world problems, mentored by older students on the program. Feedback from students has supported these views, with many students being enthusiastic about learning and applying sustainable software development practices using the approaches we have discussed here and continuing to apply the practices in their work after the conclusion of the course. This teaching pathway has demonstrated a track record of success in open-source software development, with a range of widely used software packages developed by students during their PhDs, which is available at <https://github.com/SABS-R3/software-outputs>. For example, students who

developed an automated pipeline to build antibody databases during this teaching program have subsequently used this domain-specific experience to publish highly cited antibody structure prediction tools [77], antibody language models [78], and antibody paratope predictors [79]. In each case, the incorporation of techniques such as automated workflows, contribution templates, and module packaging that they learned during the teaching course has been instrumental in the widespread sharing and reuse of their software outputs.

We hope that other graduate training schemes may consider adopting these strategies in their own institutions, to promote the value of sustainable software development, and the research benefits it can bring.

Acknowledgments

We acknowledge the work of Steve Crouch and James Graham from the Software Sustainability Institute, for their helpful instruction on the topic of software engineering in the graduate training courses. We also thank David Augustin for providing helpful comments on the notes which led to this paper.

References

1. Ghosh S, Matsuoka Y, Asai Y, Hsin KY, Kitano H. Software for systems biology: from tools to integrated platforms. *Nat Rev Genet.* 2011; 12(12):821–832. <https://doi.org/10.1038/nrg3096> PMID: 22048662
2. Markowitz F. All biology is computational biology. *PLoS Biol.* 2017; 15(3):e2002050. <https://doi.org/10.1371/journal.pbio.2002050> PMID: 28278152
3. Baxter SM, Day SW, Fetrow JS, Reisinger SJ. Scientific Software Development Is Not an Oxymoron. *PLoS Comput Biol.* 2006; 2(9):e87. <https://doi.org/10.1371/journal.pcbi.0020087> PMID: 16965174
4. Šulc P, Romano F, Ouldridge TE, Rovigatti L, Doye JPK, Louis AA. Sequence-dependent thermodynamics of a coarse-grained DNA model. *J Chem Phys.* 2012; 137(13):135101. <https://doi.org/10.1063/1.4754132> PMID: 23039613
5. Thompson AP, Aktulga HM, Berger R, Bolintineanu DS, Brown WM, Crozier PS, et al. LAMMPS—a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput Phys Commun.* 2022; 271:108171. <https://doi.org/10.1016/j.cpc.2021.108171>
6. Guyer JE, Wheeler D, Warren JA. FiPy: Partial Differential Equations with Python. *Comput Sci Eng.* 2009; 11(3):6–15. <https://doi.org/10.1109/MCSE.2009.52>
7. Gad AF. PyGAD: an intuitive genetic algorithm Python library. *Multimed Tools Appl.* 2023. <https://doi.org/10.1007/s11042-023-17167-y>
8. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, et al. Array programming with NumPy. *Nature.* 2020; 585(7825):357–362. <https://doi.org/10.1038/s41586-020-2649-2> PMID: 32939066
9. Cadwallader L, Mac Gabhann F, Papin J, Pitzer VE. Advancing code sharing in the computational biology community. *PLoS Comput Biol.* 2022; 18(6):e1010193. <https://doi.org/10.1371/journal.pcbi.1010193> PMID: 35653366
10. Baker M. 1500 scientists lift the lid on reproducibility. *Nature.* 2016; 533(7604):452–454. <https://doi.org/10.1038/533452a> PMID: 27225100
11. Trisovic A, Lau MK, Pasquier T, Crosas M. A large-scale study on research code quality and execution. *Sci Data.* 2022; 9(1). <https://doi.org/10.1038/s41597-022-01143-6> PMID: 35190569
12. Konkol M, Kray C, Pfeiffer M. Computational reproducibility in geoscientific papers: Insights from a series of studies with geoscientists and a reproduction study. *Int J Geogr Inf Sci.* 2018; 33(2):408–429. <https://doi.org/10.1080/13658816.2018.1508687>
13. Chang AC, Li P. Is Economics Research Replicable? Sixty Published Papers From Thirteen Journals Say “Often Not”. *Crit Financ Rev.* 2022; 11(1):185–206. <https://doi.org/10.1561/104.00000053>
14. Stodden V, Seiler J, Ma Z. An empirical analysis of journal policy effectiveness for computational reproducibility. *Proc Natl Acad Sci U S A.* 2018; 115(11):2584–2589. <https://doi.org/10.1073/pnas.1708290115> PMID: 29531050
15. Pimentel JF, Murta L, Braganholo V, Freire J. A Large-Scale Study About Quality and Reproducibility of Jupyter Notebooks. In: 2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR). IEEE; 2019. p. 507–517. Available from: <http://dx.doi.org/10.1109/MSR.2019.00077>.

16. Howison J, Bullard J. Software in the scientific literature: Problems with seeing, finding, and using software mentioned in the biology literature. *J Assoc Inf Sci Technol*. 2015; 67(9):2137–2155. <https://doi.org/10.1002/asi.23538>
17. Katz DS, Choi SCT, Lapp H, Maheshwari K, Löffler F, Turk M, et al. Summary of the First Workshop on Sustainable Software for Science: Practice and Experiences (WSSSPE1). *J Open Res Softw*. 2014; 2(1):e6. <https://doi.org/10.5334/jors.an>
18. Strijkers R, Cushing R, Vasyunin D, de Laat C, Belloum ASZ, Meijer R. Toward Executable Scientific Publications. *Procedia Comput Sci*. 2011; 4:707–715. <https://doi.org/10.1016/j.procs.2011.04.074>
19. Crouch S, Hong NC, Hettrick S, Jackson M, Pawlik A, Sufi S, et al. The Software Sustainability Institute: Changing Research Software Attitudes and Practices. *Comput Sci Eng*. 2013; 15(6):74–80. <https://doi.org/10.1109/mcse.2013.133>
20. Sandve GK, Nekrutenko A, Taylor J, Hovig E. Ten Simple Rules for Reproducible Computational Research. *PLoS Comput Biol*. 2013; 9(10):e1003285. <https://doi.org/10.1371/journal.pcbi.1003285> PMID: 24204232
21. Peng RD. Reproducible Research in Computational Science. *Science*. 2011; 334(6060):1226–1227. <https://doi.org/10.1126/science.1213847> PMID: 22144613
22. Osborne JM, Bernabeu MO, Bruna M, Calderhead B, Cooper J, Dalchau N, et al. Ten Simple Rules for Effective Computational Research. *PLoS Comput Biol*. 2014; 10(3):e1003506. <https://doi.org/10.1371/journal.pcbi.1003506> PMID: 24675742
23. Brack P, Crowther P, Soiland-Reyes S, Owen S, Lowe D, Williams AR, et al. Ten simple rules for making a software tool workflow-ready. *PLoS Comput Biol*. 2022; 18(3):e1009823. <https://doi.org/10.1371/journal.pcbi.1009823> PMID: 35324885
24. List M, Ebert P, Albrecht F. Ten Simple Rules for Developing Usable Software in Computational Biology. *PLoS Comput Biol*. 2017; 13(1):e1005265. <https://doi.org/10.1371/journal.pcbi.1005265> PMID: 28056032
25. Lee BD. Ten simple rules for documenting scientific software. *PLoS Comput Biol*. 2018; 14(12):e1006561. <https://doi.org/10.1371/journal.pcbi.1006561> PMID: 30571677
26. Taschuk M, Wilson G. Ten simple rules for making research software more robust. *PLoS Comput Biol*. 2017; 13(4):e1005412. <https://doi.org/10.1371/journal.pcbi.1005412> PMID: 28407023
27. Perez-Riverol Y, Gatto L, Wang R, Sachsenberg T, Uszkoreit J, da Veiga LF, et al. Ten Simple Rules for Taking Advantage of Git and GitHub. *PLoS Comput Biol*. 2016; 12(7):e1004947. <https://doi.org/10.1371/journal.pcbi.1004947> PMID: 27415786
28. Carey MA, Papin JA. Ten simple rules for biologists learning to program. *PLoS Comput Biol*. 2018; 14(1):e1005871. <https://doi.org/10.1371/journal.pcbi.1005871> PMID: 29300745
29. Reyes RJ, Hosmane N, Ihorn S, Johnson M, Kulkarni A, Nelson J, et al. Ten simple rules for designing and running a computing minor for bio/chem students. *PLoS Comput Biol*. 2022; 18(7):e1010202. <https://doi.org/10.1371/journal.pcbi.1010202> PMID: 35834439
30. Creswell R, Augustin D, Bouros I, Farm HJ, Miao S, Ahern A, et al. Heterogeneity in the onwards transmission risk between local and imported cases affects practical estimates of the time-dependent reproduction number. *Philos Trans A Math Phys Eng Sci*. 2022; 380(2233). <https://doi.org/10.1098/rsta.2021.0308> PMID: 35965464
31. Gallagher K, Bouros I, Fan N, Hayman E, Heirene L, Lamirande P, et al. Epidemiological Agent-Based Modelling Software (Epiabm). *J Open Res Softw*. 2024; 12. <https://doi.org/10.5334/jors.449>
32. Blackburn JD, Scudder GD, Wassenhove LNV. Improving speed and productivity of software development: a global survey of software developers. *IEEE Trans Softw Eng*. 1996; 22(12):875–885. <https://doi.org/10.1109/32.553636>
33. Canedo ED, Santos GA. Factors Affecting Software Development Productivity: An Empirical Study. In: Proceedings of the XXXIII Brazilian Symposium on Software Engineering. SBES '19. New York, NY, USA: Association for Computing Machinery; 2019. p. 307–316. Available from: <https://doi.org/10.1145/3350768.3352491>.
34. Chue Hong N, Hettrick S, Pringle K, Ainsworth R, Aragon S, Crouch S, et al. Software Sustainability Institute Midterm Review; 2023. Available from: <https://zenodo.org/record/8205595>.
35. Way GP, Greene CS, Carninci P, Carvalho BS, de Hoon M, Finley SD, et al. A field guide to cultivating computational biology. *PLoS Biol*. 2021; 19(10):e3001419. <https://doi.org/10.1371/journal.pbio.3001419> PMID: 34618807
36. Hafer L, Kirkpatrick AE. Assessing open source software as a scholarly contribution. *Commun ACM*. 2009; 52(12):126–129. <https://doi.org/10.1145/1610252.1610285>
37. Miller G. A Scientist's Nightmare: Software Problem Leads to Five Retractions. *Science*. 2006; 314(5807):1856–1857. <https://doi.org/10.1126/science.314.5807.1856> PMID: 17185570

38. Barker M, Chue Hong NP, Katz DS, Lamprecht AL, Martinez-Ortiz C, Psomopoulos F, et al. Introducing the FAIR Principles for research software. *Sci Data*. 2022; 9(1). <https://doi.org/10.1038/s41597-022-01710-x> PMID: 36241754
39. Yamashita A, Moonen L. Do code smells reflect important maintainability aspects? In: 2012 28th IEEE International Conference on Software Maintenance (ICSM). IEEE; 2012. p. 306–315. <https://doi.org/10.1109/icsm.2012.6405287>
40. Lacerda G, Petrillo F, Pimenta M, Guéhéneuc YG. Code smells and refactoring: A tertiary systematic review of challenges and observations. *J Syst Softw*. 2020; 167:110610. <https://doi.org/10.1016/j.jss.2020.110610>
41. Balaban G, Grytten I, Rand KD, Scheffer L, Sandve GK. Ten simple rules for quick and dirty scientific programming. *PLoS Comput Biol*. 2021; 17(3):e1008549. <https://doi.org/10.1371/journal.pcbi.1008549> PMID: 33705383
42. Allen C, Mehler DMA. Open science challenges, benefits and tips in early career and beyond. *PLoS Biol*. 2019; 17(5):e3000246. <https://doi.org/10.1371/journal.pbio.3000246> PMID: 31042704
43. McKiernan EC, Bourne PE, Brown CT, Buck S, Kenall A, Lin J, et al. Point of View: How open science helps researchers succeed. *Elife*. 2016; 5:e16800. <https://doi.org/10.7554/eLife.16800> PMID: 27387362
44. Fitzgerald. The Transformation of Open Source Software. *MIS Q*. 2006; 30(3):587. <https://doi.org/10.2307/25148740>
45. Paulson JW, Succi G, Eberlein A. An empirical study of open-source and closed-source software products. *IEEE Trans Softw Eng*. 2004; 30(4):246–256. <https://doi.org/10.1109/tse.2004.1274044>
46. von Krogh G, Spaeth S. The open source software phenomenon: Characteristics that promote research. *J Strateg Inf Syst*. 2007; 16(3):236–253. <https://doi.org/10.1016/j.jsis.2007.06.001>
47. Pearce JM. The case for open source appropriate technology. *Environ Dev Sustain*. 2012; 14(3):425–431. <https://doi.org/10.1007/s10668-012-9337-9>
48. Nakakoji K, Yamamoto Y, Nishinaka Y, Kishida K, Ye Y. Evolution patterns of open-source software systems and communities. In: Proceedings of the International Workshop on Principles of Software Evolution. IWPSE02. ACM; 2002. p. 76–85. <https://doi.org/10.1145/512035.512055>
49. Gallagher K, Hayman E, Fan N, Heirene L, Lamirande P, Bouros I, et al. SABS-R3-Epidemiology/epiabm: Zenodo Release; 2022. Available from: <https://zenodo.org/record/7327444>.
50. Beck K, Andres C. Extreme Programming Explained: Embrace Change. XP Series. Pearson Education; 2004. Available from: <https://books.google.co.uk/books?id=-DNcBAAQBAJ>.
51. Ahmed A, Ahmad S, Ehsan N, Mirza E, Sarwar SZ. Agile software development: Impact on productivity and quality. In: 2010 IEEE International Conference on Management of Innovation & Technology. IEEE; 2010. p. 287–291. <https://doi.org/10.1109/icmit.2010.5492703>
52. Sletholt MT, Hannay J, Pfahl D, Benestad HC, Langtangen HP. A Literature Review of Agile Practices and Their Effects in Scientific Software Development. In: Proceedings of the 4th International Workshop on Software Engineering for Computational Science and Engineering. SECSE '11. New York, NY, USA: Association for Computing Machinery; 2011. p. 1–9. <https://doi.org/10.1145/1985782.1985784>
53. Madeyski L, Szala Ł. The Impact of Test-Driven Development on Software Development Productivity—An Empirical Study. In: Lecture Notes in Computer Science. Berlin Heidelberg: Springer; 2007. p. 200–211. https://doi.org/10.1007/978-3-540-75381-0_18
54. Freeman S, Eddy SL, McDonough M, Smith MK, Okoroafor N, Jordt H, et al. Active learning increases student performance in science, engineering, and mathematics. *Proc Natl Acad Sci U S A*. 2014; 111(23):8410–8415. <https://doi.org/10.1073/pnas.1319030111> PMID: 24821756
55. Prince M. Does Active Learning Work? A Review of the Research. *J Eng Educ*. 2004; 93(3):223–231. <https://doi.org/10.1002/j.2168-9830.2004.tb00809.x>
56. de los Rios I, Cazorla A, Diaz-Puente JM, Yagüe JL. Project-based learning in engineering higher education: two decades of teaching competences in real environments. *Procedia Soc Behav Sci*. 2010; 2(2):1368–1378. <https://doi.org/10.1016/j.sbspro.2010.03.202>
57. Quast N, Shepard K, Raja A, Adlard D, Durant G, Zhao S, et al. Drug Discovery Game App; 2023. Available from: <https://github.com/SABS-Group-2-2021-22/drug-discovery-game-app>.
58. Wei Y, Turnbull O, Oladokun E, Wee H, Robinson M, Krol J, et al. Extensible-Clinical-Imaging-QC-Tool; 2023. Available from: <https://github.com/SABS-Group-2-2021-22/drug-discovery-game-app>.
59. Craig M, Conrad P, Lynch D, Lee N, Anthony L. Listening to Early Career Software Developers. *J Comput Sci Coll*. 2018; 33(4):138–149.

60. Ferguson N, Laydon D, Nedjati Gilani G, Imai N, Ainslie K, Baguelin M, et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand; 2020. Available from: <http://spiral.imperial.ac.uk/handle/10044/1/77482>.
61. Herriott L, Capel HL, Ellmen I, Schofield N, Zhu J, Lambert B, et al. EpiGeoPop: A Tool for Developing Spatially Accurate Country-level Epidemiological Models; 2023. Available from: <https://arxiv.org/abs/2310.13468>.
62. Porter L, Bouvier D, Cutts Q, Grissom S, Lee C, McCartney R, et al. A Multi-Institutional Study of Peer Instruction in Introductory Computing. In: Proceedings of the 47th ACM Technical Symposium on Computing Science Education. SIGCSE '16. New York, NY, USA: Association for Computing Machinery; 2016. p. 358–363. <https://doi.org/10.1145/2839509.2844642>
63. Hanks B, Fitzgerald S, McCauley R, Murphy L, Zander C. Pair programming in education: a literature review. *Comput Sci Educ*. 2011; 21(2):135–173. <https://doi.org/10.1080/08993408.2011.579808>
64. Salleh N, Mendes E, Grundy J. Empirical Studies of Pair Programming for CS/SE Teaching in Higher Education: A Systematic Literature Review. *IEEE Trans Softw Eng*. 2011; 37(4):509–525. <https://doi.org/10.1109/tse.2010.59>
65. Cockburn A, Williams L. The costs and benefits of pair programming. *Extreme programming examined*. 2000; 8:223–247.
66. Schümmer T, Lukosch SG. Understanding tools and practices for distributed pair programming. *J Univ Comput Sci*. 2009; 15(16):2009.
67. Sus JG, Billingsley W. Using continuous integration of code and content to teach software engineering with limited resources. In: 2012 34th International Conference on Software Engineering (ICSE). IEEE; 2012. p. 1175–1184. <https://doi.org/10.1109/icse.2012.6227025>
68. Clark N, Davies P, Skeers R. Self and peer assessment in software engineering projects. In: Proceedings of the 7th Australasian conference on Computing education-Volume 42. 2005. p. 91–100.
69. Dyke G. Which Aspects of Novice Programmers' Usage of an IDE Predict Learning Outcomes. In: Proceedings of the 42nd ACM Technical Symposium on Computer Science Education. SIGCSE '11. New York, NY, USA: Association for Computing Machinery; 2011. p. 505–510. <https://doi.org/10.1145/1953163.1953309>
70. Pitt-Francis J, Pathmanathan P, Bernabeu MO, Bordas R, Cooper J, Fletcher AG, et al. Chaste: A test-driven approach to software development for biological modelling. *Comput Phys Commun*. 2009; 180(12):2452–2471. <https://doi.org/10.1016/j.cpc.2009.07.019>
71. Osborne JM, Walter A, Kershaw S, Mirams G, Fletcher A, Pathmanathan P, et al. A hybrid approach to multi-scale modelling of cancer. *Philos Trans A Math Phys Eng Sci*. 1930; 2010(368):5013–5028.
72. Mirams GR, Arthurs CJ, Bernabeu MO, Bordas R, Cooper J, Corrias A, et al. Chaste: An Open Source C++ Library for Computational Physiology and Biology. *PLoS Comput Biol*. 2013; 9(3):e1002970. <https://doi.org/10.1371/journal.pcbi.1002970> PMID: 23516352
73. Cooper FR, Baker RE, Bernabeu MO, Bordas R, Bowler L, Bueno-Orovio A, et al. Chaste: cancer, heart and soft tissue environment. *J Open Source Softw*. 2020; 5(47). <https://doi.org/10.21105/joss.01848> PMID: 37192932
74. Clerx M, Robinson M, Lambert B, Lei CL, Ghosh S, Mirams GR, et al. Probabilistic Inference on Noisy Time Series (PINTS). *J Open Res Softw*. 2019. <https://doi.org/10.5334/jors.276>
75. Sholler D, Steinmacher I, Ford D, Averick M, Hoyer M, Wilson G. Ten simple rules for helping newcomers become contributors to open projects. *PLoS Comput Biol*. 2019; 15(9):e1007296. <https://doi.org/10.1371/journal.pcbi.1007296> PMID: 31513567
76. Prlić A, Procter JB. Ten Simple Rules for the Open Development of Scientific Software. *PLoS Comput Biol*. 2012; 8(12):e1002802. <https://doi.org/10.1371/journal.pcbi.1002802> PMID: 23236269
77. Abanades B, Georges G, Bujotzek A, Deane CM. ABlooper: fast accurate antibody CDR loop structure prediction with accuracy estimation. *Bioinformatics*. 2022; 38(7):1877–1880. <https://doi.org/10.1093/bioinformatics/btac016> PMID: 35099535
78. Olsen TH, Moal IH, Deane CM. AbLang: an antibody language model for completing antibody sequences. *Bioinform Adv*. 2022; 2(1). <https://doi.org/10.1093/bioadv/vbac046> PMID: 36699403
79. Chinery L, Wahome N, Moal I, Deane CM. Paragraph—antibody paratope prediction using graph neural networks with minimal feature vectors. *Bioinformatics*. 2022; 39(1). <https://doi.org/10.1093/bioinformatics/btac732> PMID: 36370083