

---

# Full Computational Reproducibility in Biological Science: Methods, Software and a Case Study in Protein Biology

Les Hatton<sup>1\*</sup> Gregory Warr<sup>2\*\*</sup>

**1 Faculty of Science, Engineering and Computing, Kingston University, London, UK**

**2 Medical University of South Carolina, Charleston, South Carolina, USA and Division of Molecular and Cellular Biosciences, National Science Foundation, Arlington, VA 22230.**

\* [lesh@oakcomp.co.uk](mailto:lesh@oakcomp.co.uk) \*\* [gwarr@nsf.gov](mailto:gwarr@nsf.gov)

## Abstract

Independent computational reproducibility of scientific results is rapidly becoming of pivotal importance in scientific progress as computation itself plays a more and more central role in so many branches of science. Historically, reproducibility has followed the familiar Popperian [38] model whereby theory cannot be verified by scientific testing, it can only be falsified. Ultimately, this implies that if an experiment cannot be reproduced independently to some satisfactory level of precision, its value is essentially unquantifiable; put brutally, it is impossible to determine its scientific value. The burgeoning presence of software in most scientific work adds a new and particularly opaque layer of complexity [29]. In spite of much recent interest in many scientific areas, emphasis remains more on procedures, strictures and discussion [12, 14, 16, 29, 30, 37, 41], reflecting the inexperience of most scientific journals when it comes to software, rather than the details of how computational reproducibility is actually achieved, for which there appear to be relatively few guiding examples [6, 10, 17]. After considering basic principles, here we show how full computational reproducibility can be achieved in practice at every stage using a case study of a multi-gigabyte protein study on the open SwissProt protein database, from data download all the way to individual figure by figure reproduction as an exemplar for general scientific computation.

## Principles

This interdisciplinary paper falls squarely between two juggernauts: the sciences which consume computation as a necessary tool and the study of computation itself. We hope it successfully bridges both to provide a research scientist with enough information on exactly how to achieve reproducibility in practice and provides a computer scientist with some context in how important the simplest of tools can become.

We will first try to present typical viewpoints. To a scientific user of software, software can be an annoyance - it is frustrating, time-consuming and essential but it is “only” programming and therefore not particularly difficult compared to the real science. As such it is generally under-resourced in scientific work and is often left up to hard-pressed research assistants to do in their spare time in spite of being generally untrained in computational tools and techniques.

On the other hand, to a computer scientist, reproducibility is merely an interesting deployment of existing tools and techniques - but not particularly novel. *In contrast actually guaranteeing software is free of error is not only novel but fundamentally intractable in the sense that nobody knows how to do it, or even if it can be done.* As such, verifying even a modest sized program as correct in function is far harder than doing the science in the first place. This remains an unpleasant truth in spite of the extraordinary progress of Computer Science in the last fifty years or so, and CS researchers still find the elimination of defects from programs and the quantification of the results of a scientific computation challenging to say the least. Quoted residual error rates for software, that is error rates *after* it has been released, remain stubbornly high and may be anywhere from around 0.1 per kSLOC (thousand lines of code) for the very best [34] to 100 times worse, [3,19]. The low end of these figures represent impressively reliable systems but the rapid growth in size of software exemplified by a few of the examples in Table 1 suggests that there are still likely to be a significant number of residual and generally unquantifiable defects, even when packages are open source and as widely inspected as these. The code written by researchers is likely to be rather worse than this for no other reason than that insufficient time and resources is available to get them this good. Even public submissions to prestigious algorithm collections such as CALGO, where it might be expected that a higher degree of preparation would be present, are often found wanting, with various repeating defects and missing tests, [27].

Package	Language	kSLOC
gimp-2.6.11	C	902.7
perl-5.14.0	C	270.0
mysql-5.5.13	C	773.5
R-2.13.0	C	433.9
gcc-4.6.2	C	3761.8
gnat-3.15p-src	Ada	575.7
kdebase-runtime-4.6.3	C++	170.4
kdelibs-4.6.3	C++	1291.9
lapack-3.3.1	Fortran	953.9
linux-2.6.39	C	13263.6
openjdk-7-ea-src-b143-20_may_2011	Java	4289.4

**Table 1.** Typical sizes of packages in wide use in scientific research along with the revision numbers [23].

It is probably worth noting that if a researcher ever managed to produce an error-free program of any significant size, a) they would not know it because they would be unable to prove it, b) they would not live long enough to be able to test it completely and c) they would be unable to repeat the process systematically.

Indeed it is a common misconception to think of a program as *tested*. Unfortunately, it is difficult even to define this concept in CS

- Tested under what conditions ?
- What were the inputs ?
- What were the outputs ?
- How many of the possible paths in the program were exercised by the tests ?

With regard to the last of these, Adams [1] in a celebrated study, demonstrated on IBM's then world-wide mainframe distribution that about a third of all the faults found

---

took *longer than 5,000 execution years* to manifest themselves, greatly undermining any possible role for dynamic testing as a means of verifying code. This means that a fault occurring this rarely would be seen approximately once in 50 years in just one of 100 machines operating independently. For a recent review, see Hatton [22].

In parallel with the growth of ancillary packages such as those in Table 1, researchers write more and more software in their own projects, mostly written in programming languages which have been in use for decades, [26, 40]. These latter references also give some important insights into modern methods of eliminating as many defects as possible in scientific environments. However, such methods cost a lot of money and time, and hard-pressed researchers rarely have the budget, technical interest and resources, and probably not the know-how to perform such packaging and verification. It would be tantamount to asking a particle physicist to do their own brain surgery.

## Code sharing

Unfortunately, premier scientific journals continue to conflate simple *code sharing* with full computational reproducibility even in influential editorials [12]. Even on the question of code sharing the language is equivocal, and almost apologetic - “*Nature journals have decided that, given the diversity of practices in the disciplines we cover, we cannot insist on sharing computer code in all cases*”. This policy makes clear that even (or perhaps especially) elite journals are quite flexible when it comes to accepting degraded standards of scientific practice in certain fields whose results they wish to publish.

We understand the logistical problems but code sharing alone is a relatively small part of full computational reproducibility and it is simply not correct to think otherwise. To see the difference between simple code sharing and providing the complete means to build and run code, we have only to look at the world of software escrow [5] whereby commercial software source code is held by a third party as part of a commercial agreement. The intent of an escrow agreement is to give an end-user access to the source code of a proprietary system in very special circumstances controlled by the escrow agent, for example if a software supplier becomes insolvent. Such escrow agreements must not only include the source code but also the complete means to build the system in a form which the end-user can use and to satisfy a hopefully rigorous set of test procedures to give confidence that it is functioning appropriately. This is a non-trivial step and it is a sobering thought that the computational reproducibility study which we exemplify here used a Linux machine, the GNU compiler, perl, R and additional other packages so is built on top of at least 20 million lines of code. Printed out at 100 lines of code per page, it would produce a stack of paper around 20m high.

## Virtual machines

This is worth mentioning in that putting up results using a virtual machine in the cloud somewhere is a possible way of distributing a scientific result. Unfortunately, it is effectively useless as a mechanism for establishing computational reproducibility without appropriate packaging. If the machine is simply a binary engine which will produce the results when a third party presses the button, then it is not *reproducibility* but simply *replicability* about which we will have more to say below, and tells you nothing about the nature of the computations which led to the result. If the machine is an image of the researchers machine without any of the appropriate packaging described below, it is simply code sharing and remains unquantifiable. In short, a virtual machine is only a convenient vehicle for delivering computational reproducibility; it still must follow the processes described in what follows.

---

## Nomenclature

There are also problems with nomenclature. The terms code or computercode [12] are simply too ambiguous to be useful. The *binary code* (machine-readable) is not useful because it is unquantifiable. The *source code* (human-readable) is a minimum requirement but quite apart from the difficulties described above, there are other significant drawbacks [29]. We define *full computational reproducibility* as a set of open deliverables which provide the *complete* means to reproduce all of the published results. As such the reproducibility deliverables should not only include all source code, but should also include all ancillary source code necessary to build and run the analysis programs, ancillary source code to download the datasets (unless they are included as part of the package), ancillary source code scripts to run statistical or other analyses and *finally a regression suite produced by the original research team against which any locally reproduced copies will be tested automatically*. Anything less than this surely fails to meet the definition of rigorous science; the reader is without the means to verify, even at the simplest level, the validity of the conclusions.

Full computational reproducibility is not difficult in principle, particularly with the modern availability of open source tools and data, but it is very technical and time-consuming work. Although much can be done to streamline the process, it is clear that many practicing scientists still either question or are unaware of its importance. We emphasize that full computational reproducibility is a vital extension of the scientific method as espoused by Karl Popper [38]. Researchers have long noted the importance of Popper's work in various contexts [44] but in a scientific age increasingly reliant (and in many fields completely dependent) on computation, full computational reproducibility has, we would argue, become a *sine qua non*. Without the implementation of full computational reproducibility, it will be impossible to determine whether, and to what degree, any computationally-dependent results have been corrupted. Unfortunately the reader has no way to determine which results in the published literature are correct and which are not. The true costs of non-reproducibility are now beginning to surface [16] and appear to be high indeed; thus any initiatives to improve reproducibility [13,14] are laudable. Unfortunately, statements such as "... just because it is not reproducible does not necessarily make it wrong." [13] are remarkably rash and unhelpful from a premier journal.

That there is serious problem is highlighted in a recent news item in Nature [2], where 52% of the respondents thought there was a 'significant crisis', however of the factors cited for this crisis by respondents, software problems rated only a sub-category under methods. When asked which factors could boost reproducibility, software transparency was not even mentioned. This represents a woefully large chasm between the providers of software methods, the computer scientists, and the consumers of software methods, the research scientists. As the news item stated, it is perfectly true that statistical techniques for example are frequently inadequately or even incorrectly applied, however, this is a failure of peer review. We have had decades of careful statistical science to guide such methods. In contrast, peer review of software is in its infancy to say the least. Software methods are indeed in danger of becoming the "elephant in the room" in discussions of scientific reproducibility.

Before discussing the general principles of computational reproducibility, it is worth emphasizing the following point:

*Full computational reproducibility is the very MINIMUM requirement for a computationally dependent result to qualify as rigorous science. It does not guarantee that results are correct but it does guarantee that not only can they be repeated to an acceptable level of significance, but that every step of the process is open for independent checking. This is exactly analogous to*

---

*conventional experimental reproducibility, the cornerstone of scientific progress for many years.*

Of course, even if the complete means for computational reproducibility is available, there is no guarantee that anyone will bother to examine the study [11]. However, that should not be taken as an excuse and if it is made as simple as possible, the chances of independent reproduction of a computationally dependent result are larger.

## Computation in science

We will discuss the general role of computation in an archetypal scientific experiment using Figure 1 where the contributing elements in enabling full computational reproducibility are shown from the left. We discuss each of these from a general point of view here and will later specifically attach openly available pieces of software to them with a discussion of how they work before showing finally, the results of them working.

**Check computational environment** The computer environment consists of the set of programs which the researcher may use to produce their results and the hardware on which they run. These might be stand-alone packages such as *gnuplot* (plotting data) or *R* (statistical inference), interpreters such as *perl* or *python* which run the researcher's own or other perl or python programs, or compilers such as *gcc*, the GNU compiler for numerous scientific languages including C, C++, Fortran and Ada. An important aspect is to make sure that the computer does arithmetic correctly to the internationally standardised IEEE 754 [4].

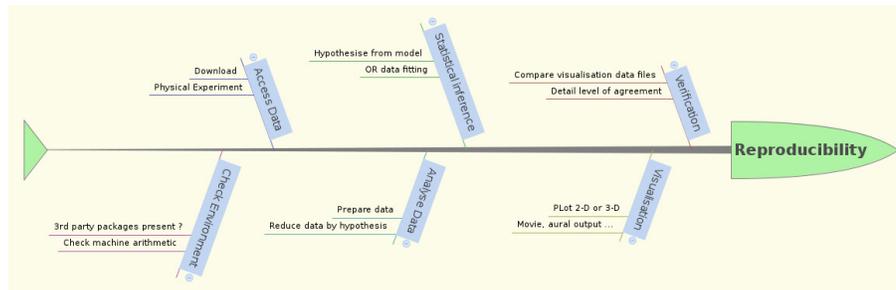
**Access Data** All scientific experiments are underpinned by data that can often be downloaded as is the case with the protein case study we describe shortly to exemplify these principles. It may equally be data acquired by the researcher using lab equipment, although to add a further layer of complexity, a researcher might be astonished to find how much software is actually embedded in the lab equipment itself, which is not without its own problems [42]. This proprietary software is generally unquantifiable and has to be taken on trust at present although we will expand on this later.

**Analyse data** Here the raw data is processed in various ways to isolate subsets of importance. This is particularly so in modern science where datasets may be immense. The case study we use to exemplify reproducibility utilises a raw dataset of around 3Gb, (taken as  $3 \times 10^9$  bytes).

**Statistical Inference** In all scientific experiments, a measure of statistical significance is essential to match theory with observation. This process might use *R*, *SPSS* or less commonly, hand-crafted software for example.

**Visualisation** As a glance through almost any recently published journal will reveal, articles increasingly contain (especially in the Supplementary Materials) extremely large datasets and complex illustrations. These may have been produced in numerous ways and may include complex 3-D representations, animations, sound recordings and so on.

**Regression** This is truly fundamental and provides a simple way for researchers trying to reproduce the results of a scientific paper to measure objectively how close they are. Simple inspection of graphical outputs, however exotic, is insufficient [24].



**Figure 1.** The necessary elements to achieve full computational reproducibility in a typical scientific experiment.

## Some comments on reproducibility

The object is to integrate as closely as possible the stringent reproducibility recommendations of Ince et. al [29] with the above model to allow the full reproduction of all the computational results of a paper for independent corroboration. In this sense we will go much further than detailing guidelines for good practice [41] by providing an actual example, although this naturally embodies a number of these guidelines. The model and examples will we hope be adaptable by others for their own experiments.

We are not the first to attempt this. Other researchers have produced similar initiatives, [6,10]. Notably, Jon Claerbout has been grappling with this thorny problem for 30 years with his group at Stanford “an article about computational results is advertising, not scholarship. The actual scholarship is the full software environment, code and data, that produced the result.” [6], but the practice is as yet by no means widespread in the face of the now formidable levels of computation in some fields. More recently Dave Donoho published an invitation in the journal *Biostatistics* to researchers [9] so there is some momentum growing but relatively slowly given the general reluctance of some journals to grasp this particular nettle.

We will distinguish the following types of principal computational object which might be used by a researcher in their work, (there are others but these give useful definitions for our discussion).

**Specific programs and configuration** These include everything written by the researchers in either compiled languages such as Fortran, C, C++ and scripted languages such as *perl* or *python*. Any configuration files necessary to run these hand-crafted programs will also have been written by the researchers. In addition, the environmental files such as the *make* scripts to allow them to be made locally would also be included.

**Third-party programs and configuration** These include programs readily available to the researcher but not written by them. These would include applications such as *perl*, the statistics program *R*, the plotting program *gnuplot* or whatever.

**Operating System and configuration** This is the operating system on top of which the specific programs and the packaged programs would be run. It might be virtualised, proprietary such as Windows or OS X or open source such as Linux or BSD.

The essence of computational reproducibility is to enable a reasonably knowledgeable researcher in the same field to set up a computer without too much effort using a reproducibility package supplied by the original researchers in such a way that

---

the original computational results can be reproduced to some reasonable level of accuracy, which should be stated, just as if they were repeating a physical experiment. This should consist of all source code actually written by the original researchers and as much as practicable of the rest as is discussed below. It should NOT simply be copies of the original binary executables. The output of these is not amenable to independent verification unless the source code is available - neither binaries nor an algorithmic description are sufficient, [29]. In the case of open source, it would be sufficient to quote the version number of packaged programs. for example, the GNU compiler *gcc*, or the *perl* environment in use, because the original source code of most such releases is readily accessible.

It may be that source code of third-party programs is not available. In that case, the ability to reproduce the results is degraded and must be replaced by trust in such programs. We would argue that in view of the preceding discussion, it is naive to trust any computer program as being 'correct', and in an ideal world, envisage the source code of all programs to be continually revisited to reduce their errors and increase confidence in their output incrementally inline with the scientific method. This employs a process known as source code inspection, the efficacy of which has been understood for almost 40 years, [15, 28, 35]. Clearly, the use of open source has revolutionised the possibility of reproducing computational results and this opportunity should be grasped with both hands. Without this, and to repeat our earlier perhaps rather bleak conclusion, *if the software used in a scientific experiment is not independently verifiable, then the results fall outside the definition of rigorous science and their value is unknowable.*

The experiment described here was carried out entirely with open source and is therefore theoretically reproducible at the highest possible level since everything can be inspected if necessary, and reproduced to the limits of single-precision floating point arithmetic, since the arithmetic on this machine is IEEE P754 conformant as noted shortly.

This protocol is not intended to be prescriptive in any way, it simply details what was done to allow the results to be fully reproducible, thus we hope that it approaches the "Gold standard" described by Peng [37] and gives a useful guide to other researchers as the principles are the same even if their software environment, data flow and analysis are somewhat different.

## Software environment

Whilst it is possible to run this software in a Windows environment, this entails the installation of a number of open source packages [7, 8] and their execution on a system which is proprietary and therefore closed. This in no way implies that Windows is an inappropriate environment. It simply makes the point that Windows of itself is not independently assessable for incorrectness. Linux on the other hand is open source and **is** at least in theory, independently verifiable allowing us to claim that every piece of software used in the production of the scientific results described here is capable of independent open scrutiny. This does not imply that either is more reliable than the other, simply that Linux is more independently accessible.

We don't for a moment imagine that any reader of this is going to download the 14 or so million lines of code in the current Linux kernel accompanied by several million assorted lines of code in the packaged programs which were used such as *perl*, the *bash* shell and the redoubtable GNU *gcc* compiler and check through them before embarking on their project, but the point is that they could. In any case, these pieces of software are independently scrutinised every day by many application experts. Equally, it would not be expected that the researchers at CERN would take apart every single component of the LHC before running an experiment. The scientific method does not demand perfection. It simply demands that we collectively might continue incrementally

---

verifying layer upon layer until something is beyond reasonable doubt and that nothing should ever impede this process. Software has greatly complicated this accessibility by reason of its rapidly growing size and complexity but open source at least guarantees that it can if necessary be done. This seems to be on the right track in the long term but the necessity for researchers to release in a useable form, the source code they themselves wrote, remains paramount.

## Third-party packages

All modern systems allow the version number of their components to be identified. For example, Linux has simple although rather variable ways of detailing system versions as revealed by the following, detailing the principle third-party packages which were used, (commands appearing in bold) and their output.

```
more /etc/*-release PRETTY_NAME="openSUSE 12.3 (Dartmouth) (i586)"
```

```
uname -a Linux linux-qa7l.site 3.7.10-1.45-desktop1 SMP PREEMPT Tue Dec 16  
20:27:58 UTC 2014 (4c885a1) i686 i686 i386 GNU/Linux
```

```
perl -v This is perl 5, version 16, subversion 2 (v5.16.2) built for  
i586-linux-thread-multi
```

```
gcc -v gcc version 4.7.2 20130108 [gcc-4.7-branch revision 195012] (SUSE Linux)
```

```
gnuplot --version gnuplot 4.6 patchlevel 1
```

```
R --version R version 3.1.2 (2014-10-31) – "Pumpkin Helmet". (We also used the  
third-party R packages tgp, akima and plot3D which may be installed within the  
R environment using the command install.packages("tgp"), and so on.)
```

This information is sufficient to take a standard PC and install it to the same specification as was used in our study. This is not an onerous activity and should be mandatory as the essence of reproducibility is to provide the means to reproduce results. Whether or not a researcher chooses to go to such lengths is up to them, however for a really pivotal result, there is no reason not to do it and relevant budget and resources should be part of the overall research study, [29].

## Arithmetic verifiability

Although it should not really be necessary in this day and age, it is always comforting to verify that the chipset / compiler combination actually does arithmetic reasonably correctly. Many scientists will take this for granted but sufficient implementation errors have appeared in computational units over the years, from the notorious FDIV Pentium bug onwards, ([http://en.wikipedia.org/wiki/Pentium\\_FDIV\\_bug](http://en.wikipedia.org/wiki/Pentium_FDIV_bug)), quite apart from implementation deficiencies in languages such as Java as to make it worthwhile [31, 32]. In addition, compilers get upgraded from time to time, perhaps not always correctly and it is always worth checking as it only takes a moment. There is an international standard for floating point arithmetic (IEEE 754) and numerous freely available test programs [4].

As all of the above named programs, perl, R, gnuplot and so on in common with many open source packages used by scientists are written in the programming language C, its intrinsic model of arithmetic computation is the bedrock of much open source computation. An open source method for verifying arithmetic performance in both standard and embedded systems is *paranoia*, which only requires an ISO compliant C compiler and for which *gcc* serves admirably [21, 33].

---

Running it on the environment used in the accompanying paper [25] gave in summary:-

```
Embedded System Paranoia SUMMARY
SINGLE PRECISION 32 bits
Closest separation = 5.9604645e-08

Number of FAILURES encountered      = 0
Number of SERIOUS DEFECTs discovered = 0
Number of DEFECTs discovered        = 0
Number of FLAWs discovered          = 0
```

```
PASSED : No failures, defects nor flaws have been discovered.
Rounding appears to conform to IEEE standard P754.
PASSED : The arithmetic diagnosed seems Excellent.
```

```
Rating ...

=====> Excellent
        Very good
        Good
        Acceptable
        Unacceptable
        Broken
```

```
...
```

This speaks for itself.

Note that this is not quite as trivial as it might appear even though very little C was actually written for the case study we describe. However, the main application packages, notably *perl* and *R* used here are themselves written in C and use the arithmetic model of C. This verification step is therefore more important than it might first appear as it is the arithmetic bedrock on which a very large amount of practical computation sits.

The position of C as a computational *lingua franca* used to program the vast majority of open source projects, (probably greater than 75% of all code [23]), means that such verification is very likely relevant in many projects.

## Data formats

In analyses of large data, it proves convenient to keep all intermediate files in some kind of human-readable form. Conversion costs between binary and human readable form can generally be ignored and proved completely negligible on this project and human readable means that:

- Testing is much simpler
- Many modern data analysis tools such as *perl* and *grep* work directly on human readable formats.
- Human readable formats compress very well using *zip*, *gunzip* or an equivalent if storage is at a premium.

With regard to data formats, the need for human readability goes far beyond just the raw data. Any proprietary file format should be considered anathema because they do not have one of the most important requirements for reproducibility - longevity. To

---

illustrate, one of us (LH) has examples of C code written to ISO C 9899:1990 which is over 35 years old. The source code files (in ASCII) can today not only be read by any of a number of editors, closed source or open source, but it will compile and run with current versions of the GNU C compiler and give the same answers as it did over 35 years ago. In stark contrast, the same author has examples of technical papers written in an old but proprietary version of Microsoft Word for Apple less than 15 years ago which cannot be read with any current version of Word on Apple or PC or indeed anything else. The only thing that can be retrieved is the text - everything else, figures, equations and formatting is lost forever.

Journals which insist on Word submissions should take note. In contrast, Latex files are human readable and the Latex system itself is open source. Not only that, but as explained later, they can be used as the basis of a system which can not only reproduce the scientific results, but reproduce the final published format of a journal article, all from the same file, including equations, tables, references and figures.

## Specific programs and source code control

Researchers have many implementation languages from which to choose. Which is the best of the hundreds available ? This question has been teasing computer scientists for decades and there is no simple answer. Choice breaks down roughly to:-

**Compiled languages** These include the juggernauts of the programming world, *C*, *C++*, *Fortran*, *Ada* and so on. Writing in these languages requires specialist knowledge of the myriads of ways of getting it wrong, [18,20]. They can be very highly portable and thanks to the GNU C compiler, ubiquitous. Fortran in its most recent incarnations F90, F95 is particularly good at dealing with vector and matrix quantities and has a fairly sane arithmetic environment.

Although there are libraries available for these languages, the situation is nowhere near as good as it is for the next category.

**Interpreted languages** These include languages like perl and semi-interpreted languages like Java. Perl, like, C++ and so on, is old but it has borne the test of time well. There are myriads of freely available libraries, most of them with inbuilt test suites, and it is probably unmatched as a pattern recognition language. When the data is in human readable form, perl is quite literally invaluable. Java is harder work, derived as it is from C and C++. The interpreted languages as a whole have attractive properties when compared with the compiled languages, [39]

**Hybrid systems** It is not difficult to produce systems which are heterogeneous in the sense of having different languages for different parts of the system. These can communicate loosely through shared files or be more tightly integrated.

Although language choice invites often passionate responses, the authors' view is that it doesn't really make much difference. It's usually better to pick something you are already good at and its helpful to think of the act of choosing programming languages as a fashion statement rather than anything much to do with engineering.

Source code control is of absolutely vital importance so that which revision of software actually in use (and there will typically be many even on a simple research project) can be readily identified. Developers will generally find that little and often (known as incremental development) is much more controllable than big bang (loads of new features, switch on, step back and see what happens). For local systems, revision control using the Linux utility *rcs* is very robust. For bigger and more distributed teams, *git* or *subversion* is highly recommended.

---

## Packaging for reproducibility

Packaging for reproducibility is a painstaking and laborious process. Even when the workflow in a study is carefully documented, it may take a very considerable amount of time to reproduce a result [17].

It is particularly difficult to do when the research is continuing because computationally based research into text analysis as is the case with protein databases is about tool-building. To investigate a potential data pattern, a *perl* tool can be assembled in a couple of hours without performance fears and used to explore the data. There is therefore an overwhelming tendency to rush ahead leaving streams of small and often discarded tools behind. Even computer scientists do this. In stark contrast, to facilitate independent verification, reproducibility is a careful reconstruction of the simplest and most comprehensible path to reproducing the claims, which are typically figures or tables of data analysis.

In our case, the original working directory for our project contained 5,679 lines of *perl*, shell and *R* scripts, C programs and makefiles. The packaging process distilled this down to 3,812 lines measured the same way. This process of rationalisation is essential because research inevitably builds up clutter whilst ideas are pursued, especially software clutter because it is so easy to write, and although the whole directory could be simply bundled up and released as reproducible, in practice it would not be because it is relatively unstructured.

We re-iterate that *the essential purpose of reproducibility after all, is not to tick boxes so that it is merely possible, but instead to encourage people to actually do it, improving the quality of all computational science as projects build one upon another.* Attempting to make sense of 5,679 lines of relatively unstructured research code is not for the faint-hearted and does not represent a reproducibility deliverable. Research code can be much larger than this and results must be as easy to reproduce as the originating research team can make them.

On another point, it was not immediately obvious but efficiency turns out to be very important to this process because assembling the packaging means a seemingly endless repetition of clean installs and incremental testing to make sure that the figures and tables are faithfully reproduced and, not insignificantly, still agree with the claims in the original paper ! Unless these can be run quickly and efficiently, it is too easy to cut corners. In our case, it took a full three days of running scripts repeatedly, streamlining and checking and re-running before the packaging was complete. As the entire suite now runs in less than a minute on a reasonable desktop, there were literally hundreds of re-runs as the packaging was shaped into its downloadable form.

Naturally, this effort takes a lot of time and we would urge, again following [29], that in any funded research, budget be allocated for this or the value of the science will be greatly undermined. Even for a practising computer scientist, the overall exercise might reveal surprises.

## Full Computational Reproducibility: a case study

We will now put flesh on these concepts by demonstrating how it can be achieved in practice in a computationally intensive research study of the SwissProt database, [25]. The resulting reproducibility package is freely downloadable from <http://leshatton.org/><sup>1</sup>. In an ideal world, “Ideally, an external researcher should be able to download a reproducible research work and reproduce all the results with a simple mouse click” [43]. In theory, we could have done this but it seems more appropriate to a researcher to be able to home in on a particular Figure or Table as well as being able to generate all the

---

<sup>1</sup>see [http://www.leshatton.org/pone\\_0125663\\_reproducibility.html](http://www.leshatton.org/pone_0125663_reproducibility.html)

---

Figures and Tables. Doing this from the command line might be considered “nerdy” [9], but it is no more nerdy than addressing an envelope and it has the benefit of complete transparency. Introducing some kind of graphical user interface with its associated libraries on top simply adds another large and potentially error prone layer of software, (around 1.5 million lines of C++ in the case of KDE, Table 1).

## Deliverability package download and environmental work

First the reproducibility package itself is downloaded as a zip file onto a suitable Linux platform and unzipped in a convenient directory. The package is organised into the following phases with a preceding digit to indicate the phase at which it should be run, starting with 0 (0- indicates environmental checking, 1- data download, 2- data preparation, 3- figure and table reproduction, 4- verification). They should be run in this increasing numerical order and we will specifically reference the components here, comparing them directly against Figure 1 from the left.

### Check environment

**Command-line:** `./0-environment.sh` This command does two things.

- 3rd party package checking. Using a series of statements like the following:-

```
which gnuplot          2> /dev/null
if test $? -ne 0
then
    echo "gnuplot MISSING"
    status=1
fi
```

This simply uses the Linux command *which* to see if the package *gnuplot* is present and report it if not. This is repeated for all the 3rd party packages actually used in the study.

- Check machine arithmetic. After verifying that all the necessary packages are present, the script then makes the built-in copy of *paranoia* and runs it. The report on the machine arithmetic is produced in the file *esp.out*.

If all the necessary components are present and the machine passes *paranoia* acceptably, the experiment is ready to be reproduced.

### Access data

This phase loads the third-party datasets. Such datasets might be generated locally but in this case, these were downloaded from the standard SwissProt repository at <http://uniprot.org/> and the Selene data repository for Post-Translationally Modified (PTM) amino acids at <http://selene.princeton.edu/PTMCuration>.

**Command-line:** `./1-downloadUniProt.sh` This carries out the following.

- First it downloads the complete current distribution using the following command.

```
wget ftp://ftp.uniprot.org/pub/databases/uniprot/
current_release/knowledgebase/complete/uniprot_sprot.dat.gz
```

- 
- It then decompresses it using *gunzip* leaving a raw data set *uniprot\_sprot.dat* of some 2.7Gb. This dataset follows best practice by being human readable in a format which is relatively easy to parse using simple *perl* programs.
  - Finally, it unpacks it into kingdoms and species using a locally produced *perl* program, **uniprotSplit.pl**. This program is less than 100 lines of perl and splits the raw data set using the built-in ID and OC headers. In addition instructions are given on how to download the much smaller Selene PTM dataset for merging with the SwissProt dataset.

## Analyse data

This is a relatively complicated part of most scientific experiments. The data has to be carefully organised, reduced if possible for efficiency and then analysed. In Figure 1, we describe this in two stages.

- Prepare data. At 2.7Gb, SwissProt is a fairly large raw dataset given that it would be processed thousands of times in our researches, so some simple experiments were carried out to pack the data into a form which would take advantage of pipelining in modern chip sets. The first point to note is that instead of feeding individual species files in one by one in *.fasta* format, (which contains much extraneous information for our needs), they would be slimmed down by parsing out extraneous information, reducing the parts that were needed to a common *.csv* (comma separated variable) format with one record per protein. This is illustrated below.

```
for file in *.fasta
do
    cat $file | processing ...
done
```

is **much** less efficient than

```
cat $combined_file | processing ... sieving as necessary
```

Here *combined\_file* is produced just once and then repeatedly analysed. It is worth spending a little time at the start of a research project doing this because it really pays dividends later.

The *.csv* format is a particularly easy text format to work with and there are many support tools available. In addition, fields were added within this record to identify the kingdom, species and protein name so individual groupings could be trivially extracted using simple filters like *grep*. During this stage, the independent Selene information was also merged into the *.csv* records for each protein. The data could have been loaded into a MySQL database or similar but *csv* format is perfectly suitable and very much more efficient on this scale.

A second stage of data preparation concatenated all the files into one. From now on, every analysis stage started by piping this single file into a chain of other filters as above. This reduced a single processing stage from around 30 minutes to only seconds on our systems. Repeating the same processing thousands of times while exploring the dataset is then entirely feasible. We refer to this as *data rummaging* as being similar to data mining but when you don't really know what you are looking for as is often the case in the early days of a research study.

- Reduce data by hypothesis. By this we mean extracting those parts of the dataset which are identified by any underlying theory as most relevant. When datasets are large and ours was quite large, this is an important step. In our study [25], our theory predicts that in proteins, the total length of the  $i^{th}$  protein in amino acids,  $t_i$  and its unique number of amino acids including any post-translationally modified amino acids  $a_i$ , formed a pair of canonical variables related by

$$t_i \sim a_i^{-\beta} \tag{1}$$

To confirm this prediction, the data was reduced on these variables using the locally produced *perl* program **txtToTok.pl**. There is some other infrastructure of shell scripts involved in this as shown in the reproducibility package but the bottom line is to produce a reduced canonical dataset from which Figures and Tables can be produced which explore how well this relationship is satisfied by the data.

This is all carried out by the following script which also has an argument N.

**Command-line:** `./2-prepareData.sh N` The value of N considers the three different perturbations to the generated tokens to probe the robustness of our results are described in our paper [25] relating to the counting of Methionine start codons in the proteins, the excision of signal peptides and a Monte-Carlo estimate of unique amino acid alphabet counting. These can each be obtained by modifying the value of N as follows:

- 0** The data is processed without adulteration, in other words as it appears in the SwissProt release.
- 1** The data is processed by normalising all Methionine start codons
- 2** The data is normalised by excising all signal peptides.
- 3** The data is normalised by applying Monte-Carlo methods to the unique alphabet counting of the otherwise unadulterated data.

The net effect of this data preparation phase and canonical reduction is to reduce the original 2.7Gb. dataset to a far more manageable one of some 76Mb., a reduction factor of almost 36.

## Statistical inference and visualisation

We separated these in Figure 1 but here it is convenient to merge them into one stage as much of this was done with the excellent open source *R* program, <http://r-project.org>.

Ideally, the final paper and the results should both be reproduced from the same file. This is quite simple to achieve in Latex by embedding some recognisable unique pattern in the .tex source file which annotates the reproducibility for the reader but also can be pre-processed to generate the set of executables which actually then generate all the figures and tables.

One way of doing this is to use a Latex macro, for example:

```
\newcommand{\reproduce}[2]{\it The reproducibility script is \bf #1 #2}}
```

In a table this might be referenced in the table or figure caption, or perhaps in the text after the first occurrence in those journals which require figure captions to be separated from the main text, as

---

```
... \reproduce{./3-table3.sh}{  
... \reproduce{./3-fig1.sh}{0}
```

(Here the fig generating script has up to two arguments.) The manuscript.pdf file can then be generated by the standard Latex command

```
pdflatex manuscript.tex
```

and the embedded macros will expand within the published manuscript itself for the above example as “*The reproducibility script is ./3-table3.sh*” and “*The reproducibility script is ./3-fig1.sh 0*” respectively, providing a cross-reference between each claim and the full means to reproduce it.

Similarly, all the necessary scripts to be run can be generated from the same Latex manuscript source file by the rather more eye-watering

```
cat InformationInProteomes.tex \  
| perl -ne \  
| '/reproduce{([\^}]+)}{([0-9]*)}/ && printf "%s %s\n", $1, $2;' \  
| sort
```

This extracts the necessary scripts in the right order from the .tex file. In a scientific paper, the ability to repeat the figures and tables one by one like this seems completely fundamental; every diagram and table is accompanied by the complete means to reproduce it locally. The methods for doing this can be inspected internally throughout the flow, including the statistical analyses, but the primary output is the diagrams in encapsulated postscript form, (.eps). Assuming the unadulterated data only is being processed, (option 0 on the figure command line generators),

- Command-line: **./3-fig1.sh 0**
- Command-line: **./3-fig2.sh 0**
- Command-line: **./3-fig3.sh 0**
- Command-line: **./3-fig4.sh 0**
- Command-line: **./3-fig5.sh 0**
- Command-line: **./3-fig6.sh 0**
- Command-line: **./3-fig7.sh 0**
- Command-line: **./3-table3.sh**
- Command-line: **./3-table5.sh**

Each script and its supporting scripts (which are all included in the deliverability package), then produces the individual components of each figure, sub-figure and table which involved computation. For example Figure 1 comprised Figures 1A, 1B and 1C. Each is produced individually. so three .eps files result.

Figures 2, 3 and 4 each include 3-D plots produced using *R*. The *R* graphics scripts to generate these are also included, so the output of the above scripts on a compliant machine will be exactly as reproduced in the original article in PLOS ONE [25].

Finally the 3-table3,5.sh scripts collect the data and apply statistical analysis using the *R* scripts provided.

We note in passing that the above process is significantly more difficult to achieve on proprietary systems such as Windows although perfectly possible [36], albeit without the complete transparency available in fully open source systems. In particular we re-iterate the dangers of *any* proprietary data formats.

---

## Verification

No reproducibility method would be complete without the means to regress the locally produced results with those produced in the published paper, (i.e. the benchmark versions). Regression is a testing method whereby a freshly generated set of test results is compared with some Gold standard of test results, for example the published and packaged results, and differences quantified and investigated. To that end, all the generated data (.dat) files used to create the figures and tables can be verified against the benchmark versions included in a sub-directory of the reproducibility package. The appropriate command in this case is

- Command-line: `./4-verification.sh`

This script merely automates the versatile Linux *diff* program as follows

```
diff -yw $gold_standard.dat $locally_produced.dat > differences.dat
```

Any differences are then logged in the differences.dat file and can be inspected. In cases where exact agreement is not expected, for example comparing results run on a 32-bit machine with those on a 64-bit machine, the differences can be quantified by piping the output through a simple program (hypothetical here as we compared integer results in our study and did not need this) as below:

```
diff -yw $gold_standard.dat $locally_produced.dat \
| quantify_differences.pl > differences.dat
```

By this means, all locally generated data files can be quickly checked for agreement with the benchmark version. We note in passing that it is easily possible to verify any .eps graphics files in the same way as .eps files are human readable although their format is not exactly transparent to the casual viewer, but it is harder to quantify any differences and it is probably sufficient to ensure that the data used to draw the final figures agrees with the benchmark version.

It goes without saying that this is not verifying correctness of research; it only verifies the reproducibility of the research, [37].

## Conclusion

Generating fully reproducible research whereby both a manuscript and all the figures and tables can be produced automatically from a single manuscript file is entirely feasible as described here. It is painstaking but not particularly difficult work but developments in modern open source at least make it possible. Even though it takes careful preparation, it is greatly comforting to take a single manuscript file as published and from it generate everything, diagram by diagram, statistical analysis by statistical analysis, regress them against a master copy and get the answers as published.

It may be that scientists in other disciplines will need technical and financial resources to achieve this and this should be allowed in the funding, but the real question if we are to take reproducibility and the scientific method seriously amidst the unprecedented growth of computation in science, is can we afford not to do this ?

## Acknowledgements

- We would like to thank all the open source contributors to Linux, *perl* and *R* who in the last 20 years, along with the foresight of the biological community in

---

providing open access to the biological data itself in a compliant form, have allowed computationally intensive biological research to be taken out of the lab and straight onto any reasonable desktop where it can be independently reproduced using similar principles to those espoused here. These are opportunities of which scientists as a whole should make the most.

- **Competing Interests:** The authors declare that they have no competing financial interests.
- **Correspondence:** Correspondence and requests for materials should be addressed to Les Hatton (email: lesh@oakcomp.co.uk).
- This material is based in part on work supported by the National Science Foundation. Any opinion, finding, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

## References

1. E. Adams. Optimising preventive service of software products. *IBM Journal of Research and Development*, 1(28):2–14, 1984.
2. M. Baker. 1,500 scientists lift the lid on reproducibility. *Nature*, 533(7604), May 2016.
3. V.R. Basili and B.T. Perricone. Software errors and complexity: an empirical investigation. *Comm. ACM*, 27(1):42–52, January 1984.  
<http://www.lsmode.de/~bernhard/cvs/text/dipl/papers/p42-basili.pdf>.
4. N.H.F. Beebe. IEEE 754 floating-point test software, 2007.  
<http://www.math.utah.edu/~beebe/software/ieee/>, accessed 12-Jul-2015.
5. J. Binder, E. Weigl, S. Strodl, A. Draws, and A. Rauber. A Framework for Automated Verification in Software Escrow. In *Proceeding IPRES 2013 10th International Conference on Preservation of Digital Objects*, volume 1, pages 95–103. Biblioteca Nacional de Portugal, 2013.
6. J.F. Claerbout and M. Karrenbach. Electronic documents give reproducibility a new meaning. In *Proc. 62nd Ann. Int. Meeting*, pages 601–604. Soc. of Exploration Geophysics, 1992.
7. Cygwin. The cygwin project, 2015. Latest release <https://www.cygwin.com/>.
8. Delorie. The dgjpp project, 2015. Latest release <http://www.delorie.com/dgjpp/>.
9. D.L. Donoho. An invitation to reproducible computational research. *Biostatistics*, 11(3):385–388, 2010.
10. D.L. Donoho, A. Maleki, I.U. Rahman, M. Shahram, and V. Stodden. Reproducible research in computational harmonic analysis. *Computing in Science and Engineering*, 8(18), 2009.
11. S.M. Easterbrook. Open code for open science ? *Nature Geoscience*, 7:779–781, 2014. doi:10.1038/ngeo2283.
12. Editorial. Code share. *Nature*, 514(536), 2014. doi:10.1038/514536a.

- 
13. Editorial. Journals unite for reproducibility. *Science*, 346(6210):679, 2014. doi:10.1038/514536a.
  14. Editorial. Journals unite for reproducibility. *Nature*, 515(7), 2015. doi:10.1038/515007a.
  15. M. Fagan. Design and code inspections to reduce errors in program development. *IBM Systems Journal*, 2:182–211, 1976.
  16. L.P. Freedman, I.M. Cockburn, and T.S. Simcoe. The Economics of Reproducibility in Preclinical Research. *PLoS Biol*, 2015. doi:10.1371/journal.pbio.1002165.
  17. D. Garijo, S. Kinnings, L. Xie, L. Xie, Y. Ahang, P.E. Bourne, and Y. Gill. Quantifying reproducibility in computational biology: The case of the tuberculosis drugome. *PLOS ONE*, 2013. DOI: 10.1371/journal.pone.0080278.
  18. L. Hatton. *Safer C: Developing software in high-integrity and safety-critical systems*. McGraw-Hill, 1995. ISBN 0-07-707640-0.
  19. L. Hatton. The T experiments: Errors in scientific software. *IEEE Computational Science and Engineering*, 4(2):27–38, April 1997.
  20. L. Hatton. Safer language subsets: an overview and a case history, MISRA C. *Information and Software Technology*, 46:465–472, 2004.
  21. L. Hatton. Embedded system paranoia: a tool for testing embedded system arithmetic. *Information and Software Technology*, 47(8):555–563, 2005.
  22. L. Hatton. Scientific computation and the scientific method: a tentative road map for convergence. In *IFIP / SIAM / NIST Working Conference on Uncertainty Quantification in Scientific Computing*, pages 123–38, August 2011. Preprint available at [http://www.leshatton.org/Hatton\\_IFIP2011\\_19Oct2011.html](http://www.leshatton.org/Hatton_IFIP2011_19Oct2011.html).
  23. L. Hatton. Conservation of Information: Software’s Hidden Clockwork. *IEEE Transactions on Software Engineering*, 40(5):450–460, May 2014.
  24. L. Hatton and A. Roberts. How accurate is scientific software ? *IEEE Transactions on Software Engineering*, 20(10), 1994.
  25. L. Hatton and G. Warr. Protein Structure and Evolution: Are They Constrained Globally by a Principle Derived from Information Theory ? *PLOS ONE*, 2015. doi:10.1371/journal.pone.0125663.
  26. G.J. Holzmann. Landing a spacecraft on Mars. *IEEE Software*, 30(2):83–86, 2013. doi:10.1109/MS.2013.32.
  27. T.R. Hopkins. A comment on the presentation and testing of calgo codes and a remark on algorithm 639: To integrate some infinite oscillating tails. *ACM Trans. Math. Softw*, 28(3):285–300, 2002. Preprint available at <https://kar.kent.ac.uk/13818/>.
  28. W. Humphrey. *A discipline of software engineering*. Addison-Wesley, 1995. ISBN 0-201-54610-8.
  29. D.C. Ince, L. Hatton, and J. Graham-Cumming. The case for open program code. *Nature*, 482:485–488, February 2012. doi:10.1038/nature10836.

- 
30. L.N. Joppa, G. McInerny, R. Harper, L. Salido, K. Takeda, O'Hara. K., D. Gavaghan, and S. Emmott. Troubling trends in scientific software use. *Science*, 340:814–815, 2013.
  31. W. Kahan. Desperately needed remedies for the Undebuggability of Large Floating-Point Computations in Science and Engineering. In *IFIP / SIAM / NIST Working Conference on Uncertainty Quantification in Scientific Computing*, 2011.
  32. W. Kahan and J.D. Darcy. How java's floating point hurts everyone everywhere. *Originally presented at ACM 1998 Workshop on Java for High-Performance Network Computing*, July 2004.
  33. W.M. Kahan, B.A. Wichmann, D.M. Gay, T. Sumner, and R. Karpinski. paranoia: a tool to test floating-point arithmetic, 1986. <http://www.netlib.org/paranoia>, accessed 12-Jul-2015.
  34. T.W. Keller. Achieving error-free man-rated software. *Second International Software Testing, Analysis and Review Conference*, 1993. Monterey, USA.
  35. A. Güneş Koru, Khaled El Emam, Dongsong Zhang, Hongfang Liu, and Divya Mathew. Theory of relative defect proneness. *Empirical Softw. Engg.*, 13(5):473–498, 2008.
  36. J. Mesirov. Research reproducibility through genepattern, 2009. <http://themindwobbles.wordpress.com/2009/06/27/research-reproducibility-through-genepattern-ismb-dam-sig-2009/>, accessed 04-Aug-2015.
  37. R.D. Peng. Reproducible Research in Computational Science. *Science*, 334(6060):1226–1227, 2011. doi:10.1126/science.1213847.
  38. K.R. Popper. *The logic of scientific discovery*. Hutchinson, London, 1959.
  39. L. Prechelt. Comparing java vs. c/c++, efficiency differences to inter-personal differences. *CACM*, pages 109–112, October 1999.
  40. D. Rousseau. The software behind the Higgs boson discovery. *IEEE Software*, 29(5):11–15, 2012. doi:10.1371/journal.pone.0125663.
  41. G.K. Sandve, A. Nekrutenko, J. Taylor, and E. Hovig. Ten Simple Rules for Reproducible Computational Research. *PLoS Comput Biol*, 9(10), 2013. doi:10.1371/journal.pcbi.1003285.
  42. M. van Genuchten. The impact of software growth on the electronics industry. *IEEE Computer*, 40(1):106–108, 2007. doi:10.1109/MC.2007.36.
  43. P. Vandenwalle, J. Kovacevic, and M. Vetterli. Reproducible research in signal processing—what, why, and how. *IEEE Signal Processing Magazine*, 26:37–47, May 2009.
  44. A.M. Ziolkowski. Further Thoughts on Popperian Geophysics—the Example of Deconvolution". *Geophysical Prospecting*, 30:p.155–165, 1982.