A Reproducible MEEG Data Analysis Workflow with conda, Snakemake, and
R Markdown

Evgenii Kalenkovich, Egor Levchenko
HSE University

Abstract

This tutorial is devoted to computational reproducibility, which is an ability to recreate the reported results using the original data and code. Previous studies show that this is impossible for a large percentage of studies with published data and code. We find this situation to be a serious problem for science in general and for the cognitive neuroscience in particular. In this tutorial, we focused on three sources of irreproducibility: differences in software environment, utilization of out-of-date derivative files, and human errors during manual copying of figures, tables, and numbers to the manuscript. We describe three tools that solve these issues: conda, Snakemake, and R Markdown, respectively. Together, they form an effective toolkit that can help researchers achieve reproducibility of their analyses. We demonstrate an application of this toolkit by reimplementing a published data analysis pipeline applied to an open MEEG dataset. Main strengths and weaknesses of our and other approaches are discussed.

A Reproducible MEEG Data Analysis Workflow with conda, Snakemake, and R Markdown

Introduction

In this tutorial, we describe an approach to performing of MEEG analyses that are computationally reproducible, i.e. identical results can be obtained by running the same code on the same data. We describe different aspects of trying to achieve this goal: (a) creating and maintaining reproducible computational environments with software versions as close as possible to those used originally, (b) using workflow management systems that help ensure that all the derivative files are continuously up-to-date and do so efficiently, (c) using literate programming so that all numbers and plots in the manuscript are derived directly from the data and not copy-pasted manually. Combining all three of these aspects ensures computational reproducibility of your results and makes it possible for your future self, for your colleagues and for other researchers to reproduce your results once the data and code are published.

Here, we define *reproducibility* (more precisely, *computational reproducibility*) as an ability to obtain exactly the same results that were reported in the publication, using the same data and code that were used by the original authors. This is in contrast to *replicability*: obtaining similar experimental results (in terms of hypothesis test results and effect sizes) by using the same experimental design and analysis but applying it to new data. The necessity of replicability is taken for granted by some (Open Science Collaboration, 2015; Pashler & Harris, 2012) and considered more skeptically by others (Devezer, Navarro, Vandekerckhove, & Ozge Buzbas, 2021). However, we find it hard to imagine anyone arguing that computational reproducibility is not essential for good scientific practice. The whole replicability question is moot if it is not possible to reproduce the original results in the first place.

The issue of reproducibility might seem trivial at first: publish your data and your code - and voilà - anyone can reproduce your results. Unfortunately, the reality is far from this naive understanding. Multiple studies (Hardwicke et al., 2018; Obels, Lakens, Coles, Gottfried, & Green, 2020; Stockemer, Koehler, & Lentz, 2018) showed that, in a

large percentage of cases it was impossible to reproduce the reported results, even though the data and the code were available. We are not aware of any empirical studies testing the computational reproducibility of MEEG, so we have to rely on anecdotal evidence and common sense here. The attempts of current study's authors to reproduce several MEEG results have not been successful even once: it usually took considerable time to even make the code run at all - only for it to output results that were vastly different from the ones reported. Also, we believe that everyone has either first-hand or second-hand experience with "I haven't changed anything but the code does not work anymore or produces different results". In this tutorial, we present our approach by applying it to the analysis presented in Jas et al. (2018). In that paper, the authors show how a comprehensive analysis of an MEG dataset reported earlier (Wakeman & Henson, 2015) can be recreated in a reproducible manner. Despite their explicit effort to make the analysis reproducible and the fact that the code accompanying the paper is well-documented and clearly organized, we were not able to run their code without updating it and rerunning the whole analysis (which takes about a week to complete) several times. Finally, MEEG analysis is usually quite involved and relies on solutions from multiple toolboxes and custom code, so we have no reason to believe that the situation should be any better than in other fields.

We find the following sources of irreproducibility to be the main culprits: changes in the software environment, out-of-date derivative files, and human error during copying the analysis outputs to the manuscript.

More often than not, someone else's, or your own older code will not run at all and throw errors instead. The reason is that exactly the same code will not run in a computationally different environment. As the software is updated, its API changes and, for example, the same function calls can require a different number of parameters, optional parameters can become mandatory, etc. You can update the code to reflect these changes but it is likely that the output of the function will change as well. With the software updates also comes the dependence on new pieces of software or on other versions of the previously existing dependencies. So, updating the code is often

insufficient. Often, despite the fact that the versions of the main pieces of software were reported in the manuscript, it is still impossible to reproduce the environment.

Sometimes the description is insufficient because seemingly minor but important pieces of software were not reported. Other times, separate pieces might have incompatible dependencies. Finally, reproducing the environment could be financially prohibitive in the case when you have a different version of commercial software (MatLab, SPSS, etc.).

The first part of the solution is to rely on open-source, free scientific software as much as possible, so that anyone can have access not only to the software itself but also to the specific version that you used. Once you do this, you can use an environment management system to manage all your software dependencies, including dealing with their complicated inter-dependencies and keeping track of all the versions. In this tutorial, we will use a tool named *conda* for that.

In practice, reproducibility comes in two distinct flavors: the *final-product* reproducibility of an already finished analysis and the *ongoing* reproducibility of an analysis as it evolves from a script that gathers the experimental data all the way to the publishable manuscript. The final-product reproducibility can, in theory, be checked by running a single master script. With MEEG, it can take a while, but if you are lucky enough and the results do not change, you will only need to run it once at the time, as you are not changing the analysis anymore, anyway. Realistically, however, you will not be running your analysis just once when the analysis pipeline is completed, but rather run it incrementally, at different stages of the analysis development. Here, however, the master script approach breaks: it is impractical to wait for the whole analysis to be rerun when you need to continue working on the other steps. There are ad-hoc solutions to this problem involving commenting out pieces of code, adding code to compare file modification times, adding custom configuration files, etc. But these solutions are extremely error-prone. Instead, we recommend delegating dealing with these issues to a workflow management system. Such tools help you run only the specific bits of the analysis that were affected by any changes you might have made. In this tutorial, we will use a tool named *Snakemake* for that. In addition, it will automatically make your

analysis execution more efficient through automatic parallelization.

The final important source of irreproducibility is the manual copying of numbers and figures from the analysis software to the text editing software. We recommend using literate programming approaches, in which the text is interspersed with bits of code that produce all the numbers and the plots that go into the final rendered document. This eliminates the possibility of making a manual copying error, by eliminating the copying step altogether: all data-derived outputs in the manuscript are directly connected to the underlying data. In this tutorial, we will use a tool named *R Markdown* for that.

In sum, we will show how to make your MEEG analysis reproducible, using *conda* to make the software environment reproducible, *Snakemake* to make and keep the analysis outputs up to date, and *R Markdown* to keep the manuscript connected to the underlying data. The usage of *conda* and *R Markdown* is the same, whether you are working with MEEG or other types of data, so we will be brief and refer you to other sources of information where necessary. With *Snakemake*, we will go into much more detail because this tool was developed and is mainly used in the bioinformatics context, which does not translate easily to MEEG. After describing the tools, we will walk you through applying them to the analysis reported in Jas et al. (2018) (which, itself, is a reanalysis of the Wakeman and Henson (2015) dataset).

Tools

In this section, we describe the tools that help us tackle the three causes of irreproducibility: (a) *conda* will help avoid differences in the software environment, (b) *Snakemake* will help keep all the derivative files up-to-date and do so efficiently, and (c) *R Markdown* will help to remove the disconnect between the numbers and figures in the manuscript and the underlying data. Out of these three, only *Snakemake* requires consideration specific to MEEG analysis. Therefore, for *Snakemake*, we provide a comparably detailed guide while we describe *conda* and *R Markdown* rather briefly. At the end, we mention several other helpful tools: *git*, *BIDS*, *renv*, and *Docker*.

Software Environment Management

In order for your code to produce the same results upon later execution, it is necessary that it is run using exactly the same versions of as many pieces of software as possible. A practical and user-friendly solution to this problem is to use the software package management system, *conda*, and the so-called *conda environments* it allows one to manage. Conda is probably most widely known as that part of the Anaconda data science distribution that lets you install Python packages. However, conda is not just a Python package manager, rather it is a software package management system or, citing the conda documentation: "Package, dependency and environment management for any language - Python, R, Ruby, Lua, Scala, Java, JavaScript, C/ C++, FORTRAN, and more"¹. As long as you stick to free (not always necessarily open-source) software, conda will cover most of your reproducible software environment needs.

Technically, the conda environment is a directory that contains a number of *conda packages* that you install into this environment. You can *activate* this environment in a terminal, after which the code that you run from that terminal, will use the software from that directory. The power of conda environments comes from the fact that a *conda package* might be a Python package, but it can also be Python itself, R, a number of other languages, a command-line utility, etc. Moreover, all of these will be installed in an isolated environment, in which your code is run against the specific versions of these packages, while not affecting other code run on your system.

Conda can be installed as part of the Anaconda distribution or through its small, bootstrap version *miniconda*. You can check that conda is available on the terminal by running `conda --version` in it. Consult the conda documentation in case of any issues. A new environment called, for example, `reproducible_MEEG`, can be created and activated with

```
conda create -n reproducible_MEEG
conda activate reproducible_MEEG
```

¹ <https://docs.conda.io/en/latest/>

After the environment is activated, (`reproducible_MEEG`) is added at the beginning of the command line. You can install Python of a specific version (3.7 here) with

```
conda install python=3.7
```

To see all the packages installed in the environment, you can run `conda list`. To save a list with all the packages to a file, you can use

```
conda env export > environment-full.yml
```

This command creates a YAML²-formatted file called `environment-full.yml` containing the list. Unfortunately, we cannot directly use this file to reproduce the environment later: due to differences in operating systems and hardware architecture, a small proportion of the packages will differ from one computer to another. Some packages, e.g. `vs2015_runtime`, are necessary on Windows but are not needed and actually do not exist for Linux/macOS. So, if you create `environment-full.yml` on Windows, you will not be able to recreate the environment on other systems. As a workaround to this problem, we will use a manually populated file that we will call `environment.yml` (the environment definition files can be named anything but using `environment.yml` will signal to other users what this file contains). In this file, we will keep track of the packages that we specifically install, but not their dependencies. This approach leads to less reproducible environments in the sense that some packages might differ in their versions. In another sense, however, they are more reproducible because it will be more likely that an environment can be recreated on a different computer, albeit with some differences. In practice, these differences are usually inconsequential (because the main packages *are* tracked). For cases where they are not, we recommend keeping the conda-populated `environment-full.yml` as well. This allows one to pin down possible execution differences to specific packages. Following is an example of a minimal `environment.yml` file:

² "YAML is a human friendly data serialization language for all programming languages" (from <https://yaml.org/>)

dependencies:

- python=3.7

When conda installs a package, it downloads it from the Anaconda package repository. This repository is divided into subsections called *channels*. Unless configured otherwise, `conda install` looks for packages in the channel called *defaults*. A large number of scientific software packages is located in the *conda-forge* channel; the workflow management system, *Snakemake*, that we will later use, is in the *bioconda* channel. To install a package from a specific channel, use

`conda install -c <channel_name> <package_name>=<version>`. For example, *mne 0.23.0*, from the *conda-forge* channel, can be installed with

```
conda install -c conda-forge mne=0.23.0
```

As you install more packages into an environment, installing a new package might lead to the change in the version of an earlier installed package, rendering `environment.yml` invalid. To avoid this situation and catch the package dependency conflicts early, we recommend to avoid using `conda install` altogether and instead update the environment based on the `environment.yml` file. To add a package from a specific channel, use `<channel_name>::<package_name>=<version>`. For example, adding *mne 0.23.0* from the *conda-forge* channel and *snakemake-minimal 6.2.1* from the channel *bioconda* (we use the minimal version of Snakemake here because - unlike the full version - it is available for Windows as well) will look as follows:

dependencies:

- python=3.7
- bioconda::snakemake-minimal=6.2.1
- conda-forge::mne=0.23.0

To update the environment after introducing changes to `environment.yml`, use

```
conda env update -n reproducible_MEEG -f environment.yml --prune
```

Flag `-n` signals the environment name, `-f` - the name of the environment definition file. The last flag, `--prune`, instructs conda to remove packages that are no longer needed. It does not change anything in this case but it is useful when you remove a package or change some packages' versions. It is a good practice to use this flag with any update and let conda figure out whether it was necessary.

Tracking the packages you add in `environment.yml` is the bare minimum requirement for the environment reproducibility and it is a good idea to track more. There is no simple rule to decide which packages to track. Some rules of thumb include tracking packages that you directly use in your code and tracking widely used packages, such as *pandas*, *scipy*, *matplotlib*, *numpy*, etc. Another reason to add additional packages to `environment.yml` is the hidden incompatibilities of some packages with newer versions of some of their dependencies. For example, *snakemake-minimal 6.2.1* is incompatible with new versions of its dependency, *smart_open*. This can be remedied by downgrading it to version 3 by adding `conda-forge::smart_open=3` to `environment.yml` and updating the environment.

Some Python packages are not available directly through conda but only through the Python Package Index (*PyPI*) and installed using the *pip* package manager. In this case, add *pip* itself to `environment.yml` and any PyPI packages under a separate *pip* node:

dependencies:

```
...
- pip=21.0.1
- pip:
  - pybids=0.13.2
```

Finally, if you use R, install it through conda as well: multiple versions are available in the conda-forge channel. It is possible to install and track R packages directly through conda but this approach might be limiting. Instead, we recommend using a dedicated R package manager, *renv*.

Summing up using conda environments:

1. Use a separate environment for each project.
2. Add new packages to `environment.yml` without specifying the version unless you need a specific version.
3. Create or update the environment file with

```
conda env update -n <environment-name> -f environment.yml --prune
```
4. After each change, update the full environment specification with

```
conda env export > environment-full.yml
```

and pin the newly installed package versions in `environment.yml`.
5. Add additional packages to `environment.yml` to increase reproducibility, consulting `environment-full.yml` for their version.

Snakemake

Let us start with an oversimplified illustrative example. Suppose we have electroencephalography (EEG) data from 20 subjects, recorded during an oddball task and all we want to do is compare P300 between the standard and the deviant conditions with a paired t-test. For each subject, we have a single recording file, `data/subXX_raw.fif`, where XX is the subject code going from 01 to 20. We will run our analysis in three steps, assuming that we already have the code that does each one of them:

1. Clean and epoch each `data/subXX_raw.fif` file and save the results to `epoched/subXX_data.fif` (done by function `clean_and_epoch(path_to_raw, path_to_epoched)`).
2. Using each `epoched/subXX_data.fif`, calculate P300 in both conditions and save to `P300/subXX_P300.csv` (done by function `calculate_P300(path_to_epoched, path_to_P300)`).

3. Using all of the `P300/subXX_P300.csv` files, create a pdf report, `manuscript.pdf`, with the results of the test and a plot of P300 values (done by function `make_report(path_to_P300_list, path_to_manuscript)`).

This analysis pipeline is schematically drawn in Figure 1 below.

Figure 1. Each node in the graph represents a file and arrows are Python function calls. First bunch of files (first line of blue nodes) are raw files for subjects from 1 to 20, the second bunch of files (light blue) is a result of applying the *clean_and_epoch* function, the third bunch is obtained after calculating P300 properties (orange), and the final file is a manuscript (light gray).

Often, such an analysis would be managed through a *master script* that might look like this:

```

for subject in subjects:
    clean_and_epoch(path_to_raw=..., path_to_epoched=...)
for subject in subjects:
    calculate_P300(path_to_epoched=..., path_to_P300=...)

```

```
make_report(path_to_P300_list=[... for subject in subjects],  
            path_to_manuscript='manuscript.pdf')
```

(In the code above, the ellipses stand for the string-manipulation code, for assembling the appropriate path strings for a given subject code.) In our oversimplified example, the master script is very short and simple, but in a real-life scenario it quickly gets large and unwieldy. It still, however, does the job, if you want to rerun the whole analysis assuming there are no errors and no missing files. There are a couple of problems with this approach though.

First of all, it is inefficient: a lot of the function calls could be done in parallel. This is true, both for individual calls of, for example, `clean_and_epoch`, but also for calls of `calculate_P300`, once most of the `clean_and_epoch` calls are finished and there are idle processors.

Second, a lot of times there *are* errors *and* missing files. Considering that a lot of MEEG processing steps take considerable time, it is impractical to rerun the whole analysis when an error is fixed or a missing file is provided, which only affects a part of the pipeline. The same is true for when a piece of code is updated. So, in practice, the researcher will either not use the master script at all, after a fix or a change, because it would take too long to run, or they would comment out some of the lines and tweak the script in other ways, so that it only runs what is necessary. Both options are very error-prone and are very likely to lead to derivative files that are out of date and non-reproducible outputs.

Third, the pipeline usually quickly becomes too large to hold in one's head and requires visualization to describe and analyse. With the master script approach, this usually requires hand-made drawing of all the interwoven dependencies in the pipeline. These are hard to keep up-to-date with the master script and check for consistency.

Of course, all the issues described above could be solved by adding pipeline managing code of varying difficulty to the master script: parallelization code is relatively straight-forward, managing file state and visualizing the inter-dependencies of individual steps would be more involved. That would be an unnecessary waste of the

researcher's time and effort spent on coding. Or, they could be solved by the researcher being extra careful. Which would be very error-prone while also being an unnecessary waste of the researcher's time and effort spent on *executing* the code. A more reliable and simpler solution is to outsource these efforts to a *workflow management system*. In this tutorial, we will show an example of such a system called *Snakemake* (for those familiar with *make*, Snakemake is *make* with the added flexibility brought by allowing Python code in the Makefiles).

Snakemake is most prevalently used in bioinformatics (Mölder et al., 2021). It shares some features with MEEG research: large files, complicated analysis pipelines, combination of multiple tools in the same analysis, parallelizable independent steps. This makes Snakemake a good fit for our purposes. With Snakemake, you define the analysis pipeline in terms of *rules* for each of the steps that transform a number (one or more) of existing files into a number (one or more) of new ones. Snakemake then takes care of running the pipeline with parallelization, without rerunning unnecessary steps, with monitoring outdated files and it can visualize the pipeline in the form of a directed acyclic graph (DAG).

The pipeline description is written to a file usually called simply "Snakefile" (no extension). This file contains a set of rules to create each of the derived files in the pipeline. Most often a rule consists of four parts: its name, the *output* - the file(s) that this rule helps to create, the *input* - the file(s) that are required to create the output, and the instructions on how to generate the output from the input. The instructions can come in different forms (more on that later), for this example we will assume that it is a piece of Python code. A basic rule has the following form (note the colons and the indents):

```
rule <rule_name>:
    input:
        <input_file>
    output:
        <output_file>
```

```
run:

    <python_code>
```

For example, a rule to create `epoched/sub01_epo.fif` might look like this:

```
rule clean_and_epoch_subject_01:
    input:
        'data/sub01_raw.fif'
    output:
        'epoched/sub01_epo.fif'
    run:
        clean_and_epoch(input[0], output[0])
```

Note that the Python code in the `run` section does not use explicit paths to the input/output files but instead refers to them via Python lists `input` and `output`. We will soon see that this not only helps us avoid repeating ourselves but also makes rules flexible. The reason we have to add `[0]` after `input` and `output` is that generally there can be multiple input and output file paths and thus `input` and `output` refer to the lists of the paths even when they contain a single element.

To create `epoched/sub01_epo.fif`, we can ask Snakemake to make it for us with

```
snakemake epoched/sub01_epo.fif --cores 1
```

In this case, the file to be created is referred to as the *target file*. Here, we assumed that the code is run in the same directory where the **Snakefile** is located and that it is actually called **Snakefile**. In this case, there is no need to specify the location explicitly. In other cases, use `-f <path_to_Snakefile>` command-line argument. The argument `--cores 1` tells Snakemake to use a single processing core. Throughout this section, we will assume there is only one core; adjust this parameter in case you want to use more cores.

Another way to create `epoched/sub01_epo.fif` is to tell Snakemake to run the rule `clean_and_epoch_subject_01`:

```
snakemake clean_and_epoch_subject_01 --cores 1
```

In this case, the rule to be run is referred to as the *target rule*. Finally, you could simply run

```
Snakemake --cores 1
```

Without additional specifications, Snakemake will use the first rule it finds in the Snakefile as the target rule.

Note that the rule `clean_and_epoch_subject_01` could theoretically be used on any subject if the subject number could be allowed to vary. Snakemake makes this possible via the use of *wildcards*, which are named variable parts of the input/output paths. Using a wildcard, we can define a single rule that is able to create any of the epoched files (note that the wildcard name `subject_number` is put inside curly brackets):

```
rule clean_and_epoch:
    input:
        'data/sub{subject_number}_raw.fif'
    output:
        'epoched/sub{subject_number}_epo.fif'
    run:
        clean_and_epoch(input[0], output[0])
```

We can then ask Snakemake to create any of the epoched files. For example, after the following command:

```
snakemake epoched/sub06_epo.fif
```

Snakemake would

1. Look for rules whose output matches `epoched/sub06_epo.fif` for some value of the wildcard

2. Determine that `clean_and_epoch` is the only such rule.
3. Determine that the value of the wildcard `subject_number` is 06.
4. Plug this wildcard into the input resulting being `data/sub06_raw.fif`.
5. Check that the input exists.
6. Run the code in the `run` section of the rule.

Note that step 2 would fail if we deleted `epoched/sub01_epo.fif` and asked Snakemake to recreate it because there would be two rules that could create this file. See Snakemake documentation for solutions to this problem should it arise in your own pipeline. Here, we will simply virtually delete `clean_and_epoch_subject_01` to avoid this collision.

Unlike with the rule `clean_and_epoch_subject_01`, it does not make sense to ask Snakemake to just run the rule `clean_and_epoch` because while the rule can create any of the `*_epo.fif` files, it does not specify any specific ones that it would output. This applies to any rule with wildcards in the output: Snakemake needs to know which files to create so these rules cannot be used as target rules. Instead, it is customary to have a dummy rule `all` as the first rule in the file and have any output files in the current state of the pipeline as its inputs:

```
rule all:
    input:
        [f'epoched/sub{subject_number}_epo.fif'
         for subject_number in subject_numbers]
```

Where `subject_numbers` is the list of character strings: `'01', '02', ..., '20'`. Plugging a number of possible wildcard values is a common action when working with Snakemake, so it has a special function that simplifies the process - `expand`. Instead of the list comprehension above, one could simply write

```
rule all:
```

```
input:

    expand('epoched/sub{subject_number}_epo.fif',
          subject_number=subject_numbers)
```

In case of multiple wildcards, you can add them as another keyword argument to `expand` - in that case Snakemake will use all the possible combinations of the elements of the supplied lists. Sometimes, what you want instead is to plug the wildcards in a row-wise fashion: plug the first element from each of the wildcard lists, then the second, etc. This can be achieved by supplying `expand` with the second positional argument `zip` (this can be any function that creates combinations of elements from multiple lists, `zip` is the one to achieve row-wise selection). Finally, sometimes you want to plug in some wildcards but not the others. In this case, supply only those you want to plug as keyword arguments and add `allow_missing=True` at the end.

If you run command `snakemake` in the console, it will:

1. Find the target rule - the first rule in the Snakefile which is the rule `all`.
2. Determine which files are missing from its input.
3. Determine that the rule `clean_and_epoch` is capable of producing each of these files.
4. For each of the necessary files, prepare a separate independent invocation of `clean_and_epoch` which is referred to as a *job*.
5. Run all the `clean_and_epoch`-based jobs in an arbitrary order.
6. Run an `all`-based job (as this is a dummy rule does not really do anything, neither will this job).

Before running Snakemake, it is advisable to check how many jobs will be run from each rule, and check whether the numbers match your expectations without running the jobs themselves, in what is referred to as *dry run*. To do so, add `--dry-run` flag to the Snakemake call, or `-n` for a shorter version.

The dummy `all`-based job can only be run once all the `clean_and_epoch`-based jobs are done - when its input files are all created. The `clean_and_epoch`-based jobs, however, can be run independently and thus can be parallelized. Snakemake does this automatically if you tell it to use more processing cores with the `--cores <n>` argument.

The rest of the pipeline can be defined with the following rules:

```
rule all:
    input:
        'manuscript.pdf'

rule clean_and_epoch:
    input:
        'data/sub{subject_number}_raw.fif'
    output:
        'epoched/sub{subject_number}_epo.fif'
    run:
        clean_and_epoch(input[0], output[0])

rule calculate_P300:
    input:
        'epoched/sub{subject_number}_epo.fif'
    output:
        'P300/sub{subject_number}_P300.csv'
    run:
        calculate_P300(input[0], output[0])

rule make_report:
    input:
        expand('P300/sub{subject_number}_P300.csv',
              subject_number=subject_numbers)
    output:
        'manuscript.pdf'
```

```
run:

    make_report(input, output[0])
```

There are several things to note about this version of Snakefile. First, the rule `all` includes only the `manuscript.pdf` file but not the epoched files or the P300 csv files. The reason for this is that the latter files become redundant once we have the whole pipeline defined. When `manuscript.pdf` is the target file, Snakemake will have to run `make_report` to create it. To do that, it will need all the P300 csv files, to create each of which it will need to run a `calculate_P300` job, which will require the corresponding epoched files. To create each of those, it will need to run a `clean_and_epoch` job, which will require the corresponding raw files. Thus, the only file we need to explicitly ask Snakemake for is `manuscript.pdf`. Second, note that `make_report` does not have any wildcards in the output specification and thus can be run directly as the target rule without the need to run the dummy `all` rule. It is advisable to nevertheless keep rule `all`. Finally, in the `run` section of the rule `make_report` we use `input` instead of `input[0]` because we need to supply the whole list of input files to the `make_report` function.

There are a few ways that Snakefile can be cleaned a bit in order to avoid repetition and increase clarity. We will create variables containing the wildcard templates for the file paths instead of repeating them in every rule. We will name the input and outputs so that we can refer to them in the `run` section with `input.<name>` instead of `input[0]`. We will refer to the outputs of previous rules instead of defining the paths directly. For this, we can use a variable `rules` that allows us to refer to the field of any rule defined above. For example, output of the rule `clean_and_epoch` can be referred to as `rules.clean_and_epoch.output` in any of the rules defined below that rule. Here is what a cleaner Snakefile might look like:

```
raw_template = 'data/sub{subject_number}_raw.fif'
epoched_template = 'epoched/sub{subject_number}_epo.fif'
p300_template = 'P300/sub{subject_number}_P300.csv'
manuscript_filename = 'manuscript.pdf'
```

```
rule all:
    input:
        pdf = manuscript_filename

rule clean_and_epoch:
    input:
        raw = raw_template
    output:
        epoch = epoch_template
    run:
        clean_and_epoch(input.raw, output.epoch)

rule calculate_P300:
    input:
        epoch = rules.clean_and_epoch.output.epoch
    output:
        p300 = p300_template
    run:
        calculate_P300(input[0], output[0])

rule make_report:
    input:
        p300_list = expand(rules.calculate_P300.output.p300,
                           subject_number=subject_numbers)
    output:
        pdf = manuscript_filename
    run:
        make_report(input.p300_list, output.pdf)
```

With a pipeline this small, it is not hard to hold it whole in the working memory. As pipelines get larger, however, it is easier to visualize all the dependencies in the pipeline using a graph similar to those depicted in Figure 1. Snakemake provides ways to create such images both at the rule level (as in Figure 2) and at the job level (as in

Figure 3). To do so, first install the `graphviz` utility (also available as a conda package in the `conda-forge` channel). To create a rule-level graph, use

```
snakemake --rulegraph | dot -Tpng > rulegraph.png
```

And for the job-level graph, use

```
snakemake --dag | dot -Tpng > jobgraph.png
```

Figure 2. Rules graph of P300 example.

Both the `--rulegraph` and `--dag` flags tell Snakemake to produce text-based graphs in the graph description language DOT. The `dot` utility from Graphviz then translates it into visual form. The `-Tpng` flag tells `dot` to save the output in a png graphics format. Consult the Graphviz documentation for other output formats. The `rulegraph.png` and `jobgraph.png` names are arbitrary, you can use anything else.

Let us assume that the way we calculate the P300 has to be changed (for example, it was calculated using the wrong set of sensors). After making the appropriate changes to the `calculate_P300` function, all we need to do is run Snakemake again telling it that the `calculate_P300` rule has to be rerun using the `--forcerun` flag (or, shorter `-R`) flag:

```
snakemake --forcerun calculate_P300
```

Figure 3. Graph of jobs of P300 example.

This will result in the recalculation of all the `*_P300.csv` files and recreation of `manuscript.pdf`. Importantly, the rule `clean_and_epoch` will not be run at all because our change to the code of the `calculate_P300` function does not influence this earlier step in the pipeline. So, how does Snakemake understand that the rule `make_report` has to be rerun but the rule `clean_and_epoch` - does not? In this specific case, we told Snakemake explicitly to rerun the `calculate_P300` rule, so it can assume that its outputs - the `*_P300.csv` files - will be updated. Thus, the input of the `make_report` rule will have changed which will *trigger* its rerun. This is similar to the thought process underlying a manual execution of a pipeline like this: the researcher figures out what lines in the master script must be rerun because of the change in the code, then comments out the lines that came before and reruns the master script. In this particular case, the manual solution is simple enough. However, as the processing pipeline becomes more involved and especially when it becomes non-linear, the manual solution becomes increasingly complicated and, more importantly, highly error-prone. Snakemake takes care of tracking the inter-dependence of individual steps and figures out which of them have to be rerun automatically.

Upon changing the code of the function `calculate_P300`, you might want to test the corresponding rule without triggering the *downstream* rules - rules that depend on the outputs of `calculate_P300` either directly or through a sequence of rules. To do so,

one might create a temporary rule whose input includes all the *_P300.csv files.

Another way is to tell Snakemake to start running the whole pipeline but stop after all the jobs *spawned* by the `calculate_P300` rule are finished using the `--until` flag:

```
snakemake --forcerun calculate_P300 --until calculate_P300
```

Once you have checked that the code and the corresponding rule work as expected, you can run Snakemake without any additional flags and Snakemake will run all the rules that have to be rerun but not any other rules. In this case, Snakemake is not told explicitly that the rule `calculate_P300` was updated so it will rely on the more general principle: any rule that the target rule depends on will be rerun if the modification date of any of its inputs is later than the modification date of some of its outputs. In our case, Snakemake will see that the target rule `all` depends on the rule `make_report` whose inputs (the updated *_P300.csv files) were modified after its output (`manuscript.pdf`) was. Therefore, Snakemake will rerun the rules `make_report` and `all`. The rule `make_report` depends on the rule `calculate_P300` which, in turn, depends on the rule `clean_and_epoch`. However, both of these rules have outputs that were updated after their inputs, so Snakemake will not rerun them.

The same logic of comparing the modification times applies when only some of the files of the same type have been updated. Let us assume, for example, that one of the files with raw data was mislabelled: `data/sub02_raw.fif` is actually a copy of `data/sub01_raw.fif` and so you substitute the correct file for it. Again, you can account for this in a master script: you will have to restrict all the subject-level codes to subject "02", leave the group-level code as is and rerun the script. This, however, is even more error-prone than the case above. Snakemake will take care of it for you automatically: you just need to set the modification date of the new correct `data/sub02_raw.fif` to the current time so that it is newer than any of the derivative files and rerun Snakemake without any additional flags. It will then spawn a single job from each of the subject-level rules (`clean_and_epoch` and `calculate_P300`) and rerun the group-level rules (`make_report` and `all`).

Using Snakemake also helps when there is an execution error in one of the jobs. Suppose that the mislabelled file `data/sub02_raw.fif` did not just have the wrong data but was also the wrong kind of file altogether. Most probably, in this case, the function `clean_and_epoch` will throw an error and halt. By default, when a single job spawned by a rule does not finish successfully, Snakemake will halt the overall execution and delete all the outputs of that rule and any of the downstream rules that might have finished. However, sometimes, including the current example, this is undesirable: the subject-level rules `clean_and_epoch` and `calculate_P300` can be run on all the other subjects. To tell Snakemake to run as many jobs as it can when it stumbles upon an error, add a flag `--keep-going`.

Up to this point, we have only been using Python for the actual processing. When you do so, you can put any Python code under the `"run"` node of a rule. However, as long as your code can be run from the command line, you can use Snakemake to run it. In this case, the command should be put as a string in the `shell` node. For example, if the `clean_and_epoch` function was implemented in Matlab and saved to a `clean_and_epoch.m` file, we could use it as follows:

```
rule clean_and_epoch:
    ...
    shell:
        "matlab -r 'clean_and_epoch({input.raw}, {output.epochs})'"
```

Note that because we used the `shell` node, the command is quoted, because for Snakemake it is a Python string, and the inputs and outputs are referred to using the curly brackets. There is also an execution node `script` which allows one to use a Python, R, Julia, or Rust script (Matlab is not supported) or an R Markdown document (it will be rendered according to the first output option in the yaml front matter). In all these scripts, there will be a special object `snakemake` in the environment that will allow the script to access the input and output file paths of the corresponding rule. See Snakemake documentation for more details.

In conclusion, Snakemake allows one to define the processing pipeline independently from the implementation of individual steps and execute that pipeline with automatic parallelization all while ensuring that all the derivative files are continuously up-to-date. Using Snakemake improves the replicability of your project directly - by allowing you to keep the pipeline updated efficiently with only the necessary computational steps rerun, and indirectly - by offloading the task of tracking the file dependencies from the researcher to their computer and thus minimizing the chance of error.

Literate Programming

Multiple reproducibility problems stem from simple mistakes made during copying analysis outputs (numbers, figures) from the analysis software to the text editing software. Tracking what has or has not been updated in the manuscript manually is extremely error-prone. A solution to this problem is using literate programming: interleaving the text with pieces of code that produce all the data-driven outputs. These pieces are executed during the process called *rendering* resulting in the final product that contains all the numbers and figures directly derived from your data.

There are multiple implementations for literate programming, the one we recommend is called *R Markdown*. It is run using the R software environment but is capable of executing code in Python, SQL, and multiple other languages. The code pieces can either come between text paragraphs in which case they are called *chunks* or directly inside the paragraphs where they are called *inline code*. A code chunk is put between lines ````\r```` and `````, a piece of inline code - between ``r` and ``` (`r` tells R Markdown that you want to execute this code using R). At the very beginning of an R Markdown document, there is a section called *YAML front matter* with `---` delineating its beginning and end. The front matter contains the document's meta information, e.g. the title, and rendering options, e.g. the output format. After the front matter, comes the body, which consists of text formatted using the Markdown language and possibly inline code and the code chunks.

Below is a minimal example of an R Markdown document. We will assume that the data is in a file called `data.csv` and that there is an R file `functions.r` that contains three functions: `load_data`, `report_t_test`, and `plot_data` with the latter two requiring only the output of the first one to function. The functions are loaded in the line `source("functions.r")`.

```
---
title: Reproducible MEEG
output: pdf_document
---
```{r}

source("functions.r")

data <- load_data("data.csv")
```

The test showed the following results: `r report_t_test(data)`.

The raw data are depicted below:

```{r}

plot_data(data)
```
```

To render the document, you will need R and pandoc installed (use conda packages *r-base* and *pandoc*). After installing R, install the R package *rmarkdown*: run R in the console by printing `r`, then run `install.packages('rmarkdown')` (select any CRAN (R packages repository) mirror if prompted). To render the document, save it to a file, then run inside R:

```
rmarkdown::render('<filename.rmd>', envir = new.env())
```

The latter argument ensures that the code in the document is run in a new R environment thus letting you avoid potential problems when your R session changes certain options, variables, functions, etc.

In case there is an error during rendering saying that a LaTeX distribution is missing on your machine, you will need to install a *LaTeX* (a typesetting system) distribution as well. The simplest way to obtain one is to install the R package `tinytex` (run `install.packages('tinytex')`) and then run `tinytex::install_tinytex()`. To quit R after the rendering is done, run `quit(save='no')`.

There are multiple packages that extend the capabilities of R Markdown. One of them is the package called `papaja` that allows you to render R Markdown to a pdf or a MS Word document formatted according to the APA style guide. The result can then be submitted to the journal as is. Install `papaja` with

```
install.packages('remotes')
remotes::install_github('crsh/papaja')
```

On Windows, you might have to install two additional tools: git and RTools. Read the error messages (if any) for the instructions. Once the installation is complete, run

```
rmarkdown::draft("reproducible-MEEG_papaja.rmd", "apa6",
                 package = "papaja", create_dir = FALSE, edit = FALSE)
```

This will create an R Markdown document called `reproducible-MEEG_papaja.rmd` that can be rendered into an APA-compliant document. Go through this document and update the fields in the YAML front matter to reflect your own information. Once you are done, render it with

```
rmarkdown::render('reproducible-MEEG_papaja.rmd')
```

Both editing and rendering R Markdown documents is facilitated by using an integrated development environment (IDE) for R called RStudio. This is not necessary, however: you can always edit the `.rmd` file in any plain-text editor and render it from the terminal. In fact, we recommend using the console even if you use RStudio to avoid RStudio settings influencing the rendering. Or, better yet, wrap the console code in a Snakemake rule so that the number and the figures are not only derived directly from

the analysis outputs, but those outputs themselves are kept up-to-date with your analysis code and raw data. For more information about R Markdown, read "R Markdown: The Definitive Guide" by Xie, Allaire, and Golemund (2019).

Other tools

Here, we will briefly mention tools that are helpful when you aim for analysis reproducibility but do not directly contribute to it (git, BIDS) and tools that are simply out of scope of the current tutorial (Docker, renv).

Git. Git is a system for version controlling your code and any other text files. It can control non-text-based files as well but is less efficient and useful for that. Remember that reproducibility boils down to applying the same code to the same data and getting the same results. The problem is that it is very hard to check whether the code is the same or not and what has changed if anything has. The situation is even worse in a collaborative setting. Git efficiently solves these problems: it shows you what has changed, when any change happened, what the state of your code was when you ran the analysis last time, etc. And it does all this while allowing for relatively easy collaboration. We highly recommend using git for all your code and code-related (e.g., R Markdown files) needs. While we find git to be a helpful tool for your reproducible analyses, others find it an essential one (Peikert & Brandmaier, 2021). Many people find git's learning curve to be steep which we believe to be in part due to the insistence on the part of many tutorials on using git from a terminal. We recommend starting with a graphical user interface (*GUI*), such as GitHub Desktop, GitKraken, SourceTree, GitX, and learn the command-line usage in parallel. This way, you can already use git while still climbing the learning curve. We recommend to start with a very basic guide by Roger Dudler³ and then continue with a more comprehensive one by Sam Livingston-Gray⁴.

³ <https://rogerdudler.github.io/git-guide/>

⁴ <http://think-like-a-git.net/>

BIDS. BIDS (Brain Imaging Data Structure) is a standard for organizing both raw and derived neuroimaging data. It does not force you to use any specific data storing formats but instead provides a clear unified way to arrange your data within a specific directory structure and write down the necessary metadata so that your data can be understood and subsequently used by anyone. Using BIDS does not necessarily help with reproducing your own analysis but it certainly facilitates checking the reproducibility and reusing the data by other researchers. Other benefits of formatting your dataset according to the BIDS standard include the ease of reading and handling your data with many tools now being aware of the BIDS data structure, the ease of uploading the dataset to a public repository such as openneuro.org, and the hidden benefit of being forced to store the metadata that you do not use now but you or someone else might require later.

For the toy example above, the BIDS-formatted dataset would look like this:

```
BIDS_dataset/
|-- dataset_description.json
|-- sub-01/
|   |-- meg/
|   |   |-- sub-01_task-oddball_meg.fif
|-- sub-02/
|   ...
|-- sub-20/
```

Notice that the only new file is the `dataset_description.json`. Otherwise, this is exactly the same dataset except for the filenames and the folder structure.

Renv. Renv is an R package that manages R dependencies of your projects. Like conda environments, renv environments are separate folders that contain all the installed R packages for your project without affecting other projects. Similarly to the `environment-full.yml` file, renv creates a "renv.lock" file that tracks all the installed R packages together with their versions. Unlike conda, however, this file is cross-platform and there is no need for a separate shortened `environment.yml` file. Even though a

large number of R packages can be installed through conda, we recommend using renv for R package management because it is more powerful and more convenient to use. It can't however manage versions of R itself so for that we still need conda.

Docker. Docker - like conda and renv - allows one to create and manage isolated software environments - *containers* to use the docker term. Conceptually, docker containers are close to conda and renv environments: they allow you to use a separate set of software pieces without affecting other processes on your machine and it allows other people to recreate this set on their own machine. The difference is that docker containers allow for an even deeper level of isolation and for installation of almost arbitrary software (tracking software versions requires the corresponding installation files to be available for download in the future). In a sense, docker simulates a computer within your computer making it similar to virtual machines. There are certain technical differences between docker containers and virtual machines that make docker containers much easier to use and thus making it practical to run your analysis through them. These two properties together: deeper recreation of the software environment and an ability to use them for everyday work make them a perfect solution for the software environment reproducibility. The only reason that we did not include using docker into this tutorial is that many MEEG researchers - especially in the developing countries (based on admittedly anecdotal evidence from Russia) - still use Windows 7 and Windows Server 2012 in which the use of docker, while possible through virtual machines, is impractical for everyday use. This means that they will neither be able to maintain the ongoing reproducibility of their projects nor check or utilize the final-product reproducibility of the docker-based analyses of other researchers. This fact, of course, does not preclude one from using docker as an additional final-product reproducibility layer once the analysis is done. This way, there is an option to reproduce your environment both with the docker that contains conda and with conda alone. We recommend the following guide for setting up conda within a docker container⁵.

⁵ <https://towardsdatascience.com/conda-pip-and-docker-ftw-d64fe638dc45>

Application to an Open MEEG Dataset

In this part of the tutorial, we go through the steps of applying our approach to the analysis reported in Jas et al. (2018) which is a reanalysis of the MEEG part of the "multi-subject, multi-modal human neuroimaging dataset" (Wakeman & Henson, 2015) publicly available at openfmri.org (legacy version) and openneuro.org (current version). The article by Jas et al. (2018) had an aim closely similar to ours: to show a way to run "A reproducible MEG/EEG Group Study" (from the article title). However, they focused on different aspects of the task than we do. They showcased the MNE software package for MEEG analysis, provided a tutorial that goes beyond the usual single-subject data, applied and compared different strategies at various stages of the analysis, and made all the scripts necessary to reproduce their analysis publicly available.

In addition, they shared their software environment using a conda environment YAML file (what we name `environment-full.yml` in Section Software Environment Management), ran all the individual scripts (such as `05-run_ica.py` for fitting ICA models for each participant's data) using a dependency-tracking utility GNU Make (which was the inspiration for Snakemake), and generated dynamically-populated reports using the reStructuredText format and the documentation building software tool Sphinx. Thus, their code, in theory, deals with all three areas that we focused on in this tutorial: software environment reproducibility, workflow management, and dynamically populated reporting. However, none of the employed tools are taught in the article. Which is where our tutorial comes in: it extends the one by Jas et al. (2018) by teaching the researcher the tools to organize their MEEG analysis in a continuously reproducible way. In other words, Jas et al. (2018) showed *what* one can do with an MEEG group dataset whereas we will show you *how* to do so reproducibly.

We would like to highlight several differences in which tools we used and how we used them. First, Jas et al. (2018) used a conda environment to manage their software environment and shared it in full in the form of the `environment.yml` that contained all the packages in the environment. This approach, while sensible on the surface level,

led to us not being able to reproduce the environment due to one package that was installed from a non-standard channel that does not exist anymore and a few packages incompatible with the operating systems we used (Windows 7 and Windows Server 2012). With some guesswork we could have tried to reproduce the environment closely but, on the advice of a colleague of the original authors, we used the current environment file used for the mne-python package when trying to reproduce the analysis as-is. That is why our advice is to use a shorter environment specification in `environment.yml` for environment recreation and the full specification in `environment-full.yml` for potential tracking of environment differences. Second, while Jas et al. (2018) used the dependency-tracking utility GNU Make, they did not use it to track the file dependencies in their analysis and thus avoid unnecessary reruns of the analysis steps. Our tool of choice for this is Snakemake which was inspired by GNU Make but took it a step forward by adding automatic parallelization, more flexible wildcard usage and generally allowing for more flexibility by allowing the user to intersperse the rule definitions with Python code. Third, Jas et al. (2018) used Sphinx for creating reports whose content stemmed directly from the data. This is a versatile tool that is capable of creating comprehensive interactive reports as can be readily seen on the site accompanying their paper. However, here we chose R Markdown instead because its output can be directly submitted to the journal and its learning curve is considerably flatter.

We would like to emphasize that our alternative approach is not meant as a critique of the one taken by Jas et al. (2018) but instead as a proposal for an improvement with stronger reproducibility and higher efficiency. Their analysis was well-documented making it relatively easy for us to make occasional changes to the code so that the analysis could be run as a whole. However, we do believe that with a little time investment in learning the tools we describe here, one can make both their life and the life of other people trying to reproduce an analysis easier. This is especially true for learning to use a workflow management system such as Snakemake. On the two-core Linux machine we used to run the Jas et al. (2018) analysis in its original form (with

occasional code changes), the whole run took around a week. So, every time there was an error, we lost a considerable amount of time: either on rerunning the whole part of the analysis or carefully selecting which parts did not have to be rerun all while not running some other parts that did not depend on the step that produced an error. This is precisely where Snakemake steps in, see Section Snakemake for a more detailed discussion of the matter.

Finally, there are some differences between the analysis of Jas et al. (2018) and the one reported here. First, our conda environment is different as described above. Second, instead of the original Wakeman and Henson (2015) dataset, we used an updated version that is organized according to the BIDS standard (Gorgolewski et al., 2016) and contains only 16 out of 19 participants. These, however, were the same 16 participants analyzed in both the original paper and in the Jas et al. (2018) reanalysis. Third, using BIDS-formatted data additionally meant that we did not have to extract the stimulus events from the data as this information is now already present in the dataset. Note that in the original dataset the uploaded "raw" data had been cleaned with MaxFilter (Taulu & Kajola, 2005) while in the updated dataset, the raw data is truly raw and the MaxFiltered data are under the `derivatives/meg_derivatives` folder. Fourth, for the purposes of the source modeling of the MEEG data, Jas et al. (2018) performed reconstruction of the brain structure using the FreeSurfer software distribution (Fischl, 2012). Our results of the same procedure were slightly different, possibly due to the difference in the FreeSurfer version which was not specified in Jas et al. (2018) article. To avoid this problem, we provide a public repository with the results of the reconstruction and suggest that such data should always be published together with the raw data. Fifth, we had a technical problem with the sliding estimator analysis and skipped this step entirely.

Setting up

The first step is setting up the software environment and the folder structure. Create a folder to store the code for the project. In the beginning, we will install and

track the minimal set of software: python, snakemake (requires channels conda-forge and bioconda), and mne (requires channel conda-forge). Here is what our `environment.yml` will look like:

```
name: reproduction

channels:

  - defaults

  - conda-forge

  - bioconda

dependencies:

  - python

  - bioconda::snakemake-minimal

  # Otherwise, snakemake throws an error (see https://git.io/JRckv)

  - conda-forge::smart_open=3

  - conda-forge::mne
```

There are two non-obvious lines here. First, we install the *snakemake-minimal* package here because it is available for Windows as well as Linux/MacOS but still provides all the necessary core functionality. Second, we had to pin the version of the conda package `smart_open` because its newer versions are incompatible with the *snakemake-minimal* version available from conda-forge at the time of writing (6.2.1). Such tweaks are unfortunately an inevitability but they are usually easily made as necessary with the use of `environment-full.yml`.

For your own analysis, create the environment based on the file above, then create its full description in `environment-full.yml` and then pin the package versions in `environment.yml`. To recreate *our* environment, start by pinning the versions in `environment.yml` as follows: Python 3.9.4, *snakemake-minimal* 6.2.1, and *mne* 0.2.2. Next, create a folder where you will store the raw and the derivative data and set an environment variable to its path. This way, the code can refer to this folder without relying on its location relative to the code location. This is especially helpful when

multiple researchers use a shared folder for the data but not the code. We will use the environment variable `reproduction_data`.

Next, you have two options for getting the data. First, you can download the whole dataset from openneuro.org to the data folder. Second, you can rely on a rule that downloads the data for you. The downside of the first option is the increased storage requirement, of the second - connection interruptions that might require multiple passes. Note that the two approaches are not mutually exclusive: if you download the dataset manually, the downloading rule will simply not be triggered. If you choose the manual download route, store the data in the `bids` subfolder.

Analysis pipeline

As an overview of the analysis, here are the major steps in the analysis in Jas et al. (2018) that we recreate:

1. get cleaned stimulus-locked epochs,
2. calculate evoked response,
3. do the forward models for source reconstruction,
4. do source reconstruction using two alternative techniques,
5. visualize evoked responses and the reconstructed sources.

Before we start implementing the pipeline proper, we will first create an empty Snakefile and start filling it with a few variables that we will certainly use whatever the rules will be. These are the `bids_dir` and `subject_codes` variables that contain the folder that contains the raw data (subfolder `bids` of the folder stored in the environment variable `reproduction_data`) and the participant codes ("01", "02", etc.) respectively.

There is no defined way to start defining the pipeline. If this is an analysis you have already done on a different dataset earlier, then you might be able to define the whole pipeline before any one of the rules is actually implemented. Start with the rule `all` with a single input: `manuscript.pdf` and then back-trace it all the way to the raw

files. To create `manuscript.pdf`, you will need a rule `make_manuscript` that will take as its input: (1) an R Markdown document, (2) the tabulated data for the descriptive and inferential statistics and any figures created within the document, (3) figures that are created outside of the document. For each of the data and figure inputs, there needs to be another rule with its own inputs that will require more rules to be created, and so on and so on. Until you hit the end which is the raw data. Once all the rules and the corresponding file dependencies are set up, you can start implementing the `run` (or `shell`, or `script`) parts of the rules now going forward from the raw files. To avoid errors, just tell Snakemake to stop at the last rule having been implemented using the `--until` flag.

Cleaning and Epoching. We believe that the situation described above, where you are able to anticipate all the things that you will need to do with the data before even touching them, is ideal. It is also utopian. More likely, you will know only parts of it, for example, the way you will clean the data. In this case, decide in what form you will need the cleaned data in and start with that. For the Jas et al. (2018) analysis, this is a single file for each of the participants that contains the stimulus-locked data. Loosely following the BIDS standard, we will put each such file in the following folder:

```
subj_preproc_dir = (bids_dir / 'derivatives'
                    / 'preprocessing'
                    / 'sub-{subject_number}'
                    / 'ses-meg' / 'meg')
```

We refer the reader to the BIDS documentation for the reasons why this path is as rather complicated as it is. It's only property important here is that this folder is subject-specific, hence we used the `subject_number` wildcard so that the rules we create are applicable to any of the participants. The path to the files themselves will be

```
epochs_cleaned_template = (subj_preproc_dir /
                             'sub-{subject_number}_ses-meg_task-facerecognition_epoCleaned.fif')
```

This path does not follow the BIDS standard because of the non-standard `epoCleaned` but it is close enough for our purposes. Again, the only thing important here is that the filename is subject-specific as well and so we used the same `subject_number` wildcard as above. When Snakemake plugs in the wildcard values, the ones with the same name will get the same value. As you can see, the BIDS standard leads to rather long and complicated paths. This is not a problem in practice but to save space, we will use ellipses for large parts of the paths from here on out. See the accompanying code for the full versions. Now we are ready to start defining rules `all` and `clean_epochs`:

```
rule all:
    input:
        clean_epochs = expand(epochs_cleaned_template,
                               subject_number=subject_numbers)

rule clean_epochs:
    input:
        ??? = ???

    output:
        clean_epochs = epochs_cleaned_template

    run:
        pass
```

Now, we need to trace back our analysis from the `epochs_cleaned_template` back to the raw data. We will not go through each rule and will only focus on the ones whose input-output structure is slightly more complicated than what we described in Section Snakemake. Let us start with finishing the rule `clean_epochs`. As you can see, the cleaning is done in two ways: the linear filtering is applied at an earlier stage and at the `clean_epochs` stage, the already filtered and epoched data (input 1) is additionally filtered by removing ICA components (the results of fitting ICA to the data is input 2) that we earlier selected as artifactual (the list of artifactual components is input 3). For

the actual cleaning, we will use a function `clean_epochs(ica_path, artifact_components_path, epochs_path, epochs_cleaned_path)` that takes all the input and output paths as arguments, does the cleaning and save the results to a file at the output path. That the rule and the function share a name will not cause errors, whether you use the same names to highlight their connection or use different names to avoid confusion is the matter of taste. For the purposes of this tutorial, the body of this function (and all the other functions for this matter) is irrelevant. Here is what the next iteration (only the new or changed lines of code) of Snakefile might look like:

```

epoched_template = ... / 'sub-{subject_number}..._epo.fif')
ica_template = ... / '{subject_number}..._ica.fif')
artifact_components_template = (... /
    'sub-{subject_number}..._artifactComponents.npz')
rule all:
    input:
        clean_epochs = expand(epochs_cleaned_template,
                               subject_number=subject_numbers)
def clean_epochs(ica_path, artifact_components_path, epochs_path,
                 epochs_cleaned_path):
    ...
rule clean_epochs:
    input:
        epochs = epoched_template,
        ica = ica_template,
        artifact_components = artifact_components_template
    output:
        clean_epochs = epochs_cleaned_template
    run:
        clean_epochs(
            ica_path=input.ica,
```

```

artifact_components_path=input.artifact_components,
epochs_path=input.epochs,
epochs_cleaned_path=output.clean_epochs)

```

Where `epoched_template`, `ica_template`, and `artifact_components_template` are the variables containing the generalized subject-level paths to the epoched data, fitted ICA model, and select artifactual components respectively. We will not list any more templates and function signatures to save space, all of them can be found in the accompanying code. To implement the pipeline further, we will need rules for creating all of the input files of the `clean_epochs` rule, for creating their input files, etc. You can see the graph of these rules in Figure 4.

Figure 4. The rule graph of the pipeline until `clean_epochs` rule.

In our implementation, they all end up depending on the root rule `download_from_openneuro`. This is not necessary, simply skip such a rule if your data

are not published yet or you simply want to download them manually. Here, we added this rule mostly to show that rules do not have to have inputs! This rule creates file based on their path relative to `bids_dir` only:

```
rule download_from_openneuro:
    output:
        file_path = bids_dir / '{openneuro_filepath}'
    wildcard_constraints:
        openneuro_filepath = openneuro_filepath_regex
    run:
        ...
```

Another reason for adding this rule is to highlight a common if not very frequent problem: multiple rules are seemingly capable of creating the same files. Here, `download_from_openneuro` seems to be able to download any file in `bids_dir` which is unfortunate because we will put all our derivative files under the same folder. In this case, Snakemake will not be able to determine which rule to use to create some of the files and will throw an error informing you of this ambiguity. One of the possible solutions to this situation is to add so called wildcard constraints - regular expressions that limit what values can a wildcard take. This way, Snakemake will only use `download_from_openneuro` for the files that satisfy these constraints - the ones that we specify to be the source data files for our analysis. Consult Snakemake documentation for more ways to deal with ambiguous rules.

The rules `apply_linear_filter` and `extract_bad_channels` (the latter one extracts the information about the channels marked as bad by the dataset authors) are straightforward with one input and one output file. A more interesting one is rule `ica`. Each participant in the dataset has six MEG recordings - *runs* in the BIDS vernacular. The data from all of them is put together back-to-back before fitting the ICA model. Additionally, before the ICA model is fitted, the data need to be highpassed with 1 Hz being most often used as the lower frequency limit. However, for the rest of the analysis, the data will not be highpassed at all resulting in the files with the filtered

data being fully described by three wildcards:

```
filtered_template = (... /
    'sub-{subject_number}_..._run-{run_id}_filteredHighPass{
        l_freq}.fif')
```

The `run_id` can take values from "01" to "06" and `l_freq` can either be "1" or "None". For the input of the `ica` rule, we will need all the participant's runs highpassed at 1 Hz. Importantly, we need to keep the `subject_number` wildcard so that the rule can be used on any participant's data. For that, we will supply only the `run_id` and `l_freq` to the `expand` function (see Section Snakemake) and tell it to ignore the fact that we skipped `subject_number`:

```
rule ica:
    input:
        runs = expand(filtered_template,
                      run_id=run_ids, l_freq="1",
                      allow_missing=True)

    output:
        ica = ica_template

    run:
        calculate_ica(input.runs, output.ica)
```

Where `run_ids` is the list of possible run codes: "01", "02", ..., "06". A similar approach is used in the rule `concatenate_runs` that concatenates non-highpass-filtered data for any participant. It also showcases two other common patterns. First, a rule can have multiple outputs. Second, we can mark outputs as temporary so that Snakemake deletes them when all the rules that depend on them have finished. The concatenated runs are rather redundant: they are quick to create, contain no additional information and take a lot of unnecessary space. If they were only needed for one other rule, we would simply create them in memory as we did when we fitted ICA. But in this case, there are two rules that depend on the outputs of `concatenate_runs`:

`make_artifact_epochs` and `make_epochs`. Therefore, it is convenient to have these outputs saved temporarily and then deleted once they are no longer needed.

```
rule concatenate_runs:
    input:
        filtered = expand(filtered_template,
                           run_id=run_ids, l_freq=EPOCHS_L_FREQ,
                           allow_missing=True),
        bads = expand(bad_channels_template,
                      run_id=run_ids, allow_missing=True),
        events = expand(events_template,
                         run_id=run_ids, allow_missing=True)
    output:
        raw = temp(concatenated_raw_template),
        events = temp(concatenated_events_template)
    run:
        concatenate_runs(...)
```

The other rules for cleaning and epoching the data:

`select_artifact_components`, `make_artifact_epochs`, and `make_epochs` do not require any approaches that have not been already described so we refer the reader to the accompanying code for their implementation.

Evoked Responses. The evoked responses are first calculated at the subject level by rule `make_evoked` and then on the sample level by the rule `group_average_evoked`s. Neither of these require any special treatment.

Forward Modeling. To create participant-specific forward models for the source reconstruction, Jas et al. (2018) first reconstructed the brain structures from the MRI images using the FreeSurfer software distribution. FreeSurfer is not available for Windows, the reconstruction is rather slow, it is often run on a computer that is different from the one where the analysis is run and it is done only once, and only a very small proportion of the output files is used in the forward modeling. For these

reasons, we decided to not include the reconstruction part in the pipeline and share the results in a public repository. We believe that this is a reasonable strategy not only in the case of this tutorial but in the case of any analysis involving the use of FreeSurfer - treat the outputs almost as raw files and share them together with the collected data.

The forward modeling itself is done with the rules `estimate_transformation_matrix` and `run_forward`, the latter of which does not require any non-standard handling. The first rule requires as its input a single participant's run so we plug in `run_id "01"` into the template and set `allow_missing` to `True` so that `expand` ignores the missing `subject_number` wildcard. We already saw this pattern with the rule `ica`. Here, the difference is that we only need a single path while `expand` always returns a list. This is solved by simply taking the first element of that list with `[0]`:

```
rule estimate_transformation_matrix:
    input:
        run01 = expand(run_template, run_id=|01|,
                      allow_missing=True)[0],
```

Inverse Modeling. Inverse modeling is implemented in the rules `make_inverse` and `group_average_source` with the dSPM algorithm and `lcmv_beamformer` and `group_average_lcmv` with the LCMV beamformer algorithm.

Report

To simplify the R Markdown file, we opted to create all the figures externally using Python and the rules `plot_erp`, `plot_dspm`, and `plot_lcmv`, none of which required any special handling. Plotting, however, required selecting certain properties: which sensor to use for plotting the evoked response, what colors to use for each condition, etc. We save all these properties to JSON-formatted text files and then read them in R Markdown so that the captions always reflect the values used during the creation of the figures. Consult `report.Rmd` in the accompanying code for details.

The rendering of the R Markdown is done by the rule `make_report`:

```
rule make_report:
    input:
        rmd = 'report.Rmd',
        erp = Path(rules.plot_erp.output.png).as_posix(),
        erp_properties = Path(
            rules.plot_erp.output.properties).as_posix(),
        dspm = Path(rules.plot_dspm.output.png).as_posix(),
        dspm_properties = Path(
            rules.plot_dspm.output.properties).as_posix(),
        lcmv = Path(rules.plot_lcmv.output.png).as_posix(),
        lcmv_properties = Path(
            rules.plot_lcmv.output.properties).as_posix(),
    output:
        'report.html'
    shell:
        ('Rscript -e "rmarkdown::render(\'{input.rmd}\', '
        'output_file = \'{output}\', '
        'params = list('
        'erp = \'{input.erp}\', '
        'erp_properties = \'{input.erp_properties}\', '
        'dspm = \'{input.dspm}\', '
        'dspm_properties = \'{input.dspm_properties}\', '
        'lcmv = \'{input.lcmv}\', '
        'lcmv_properties = \'{input.lcmv_properties}\''
        '))"')
```

There are a couple of things in this rule that are worth discussing. Snakemake is capable of rendering R Markdown files via `script: 'report.Rmd'` which is of course considerably more concise than our version. There are several reasons why we chose to

go with the `shell` option instead. First, we have experienced technical problems with the `script` option and thus chose to avoid it. Second, the `script` option is limiting for R Markdown: it renders only to the first output option in the YAML front matter of the document, it precludes ones from using R startup files (`.Rprofile`, `.Renv`, etc.). Finally, we wanted to use this rule to show a few peculiarities of using the `shell` node. Some shells, most notably `cmd` on Windows, do not support multi-line commands. Therefore, the most reproducible option is to have the full command on one line. One way to achieve this while maintaining a semblance of readability is to break the lines into multiple Python strings in parentheses without commas between them. Next, `Rscript` utility with the `-e` flag, runs a piece of R code within the quotation marks. As we already used single quotes to delineate the Python strings, we used double quotes to tell `Rscript` where the R code starts and ends. Snakemake plugs special symbols such as `{input.rmd}` without any quotes around them so we need to add them ourselves so that R knows that these are strings. We can't use double quotes anymore so that we do not signal the end of R code to `Rscript`. We thus used the single quotes again but escaped them by prepending a backslash so that we do not signal the end of Python code to Snakemake. Navigating the different quotation marks is cumbersome but the errors usually give enough information to fix the code and highlighting features in integrated development environments (IDEs) make the task relatively easy. Next, using `shell` meant that we would not have a special variable `snakemake` containing the information about the paths in the R environment. Thus, we added empty parameters under the `params` node in the YAML front matter and then supplied the parameters via the `params` argument of the `rmarkdown::render` function. On Windows, the paths will not be valid R strings because they will contain backslashes which should be escaped in R string literals. To avoid this problem, we modified the inserted paths using `Path(<path>).as_posix()` which forces the paths to use forward slashes on any platform. Note that you will need to first import class `Path` from the `pathlib` package. All the above considerations make it effortful to get the rule to execute without errors. We suggest copying our rule verbatim and then modifying it to your needs.

Conclusion

We hope that these instructions in conjunction with the accompanying code provide enough information to start applying conda, Snakemake, and R Markdown to make your analysis pipelines more reproducible. See their descriptions in section 2 for more learning resources. Additionally, we suggest you ask questions on *StackOverflow.com* in regards to any of these tools.

Discussion

Multiple scientific fields, cognitive neuroscience included, are faced with the crises of replicability: results of prior experiments often cannot be replicated after being repeated by another or even by the same laboratory. Setting aside the issue of whether replicability is a goal worth pursuing in and of itself (Devezer et al., 2021), the purported crises are a strong indication of substantive issues in the scientific process. One of such issues is reproducibility - an ability to get the same results from the same data. It is unclear what replicability entails when the results being replicated cannot be reproduced in the first place. This makes reproducibility a prerequisite of replicability.

In this tutorial we focus on three important threats to reproducibility of data analyses and three tools that help handle them. Conda helps to keep the software environment reproducible so that packages and utilities versions are kept constant. Snakemake helps to manage a data analysis pipeline so that no outdated files are mistakenly used. Finally, R Markdown helps to directly connect the data analysis outputs to the final manuscript so that no mistakes are made during manual copying. We find the combination of conda, Snakemake, and R Markdown to be an effective set of tools that solves the outlined reproducibility issues, but we do not insist on our specific choice. Other tools can solve these problems and we encourage you to use any of the alternative tools that manage your software environment, the analysis pipeline and allow you to use literate programming.

We appreciate that learning these tools requires considerable additional effort on the researcher's part. We thus would like to first point out that these tools are

independent and can be adopted sequentially: introducing even one of them to your workflow will already increase the data analysis reproducibility. Second, it is our experience that their use saves time and reduces cognitive load in the long run. For example, the pipeline management simplifies updating or extending the analysis, literate programming saves time by obviating the need for manual copying.

We focused on analyses implemented in Python (and partially R), but our approach can be combined with Matlab, SPSS, etc. as long as all the steps are scripted. We do, however, suggest avoiding using proprietary software. The problem is that it is likely to be unavailable to other researchers trying to reproduce an analysis. Even when they do have it installed, it might be of a different version, or might have different components installed. For example, Taulu and Kajola (2005) report that the same EEGLAB-based analysis produced different results on the same computer depending on whether a Matlab component called Signal Processing Toolbox was installed. We also suggest avoiding GUI-based analysis tools, unless they provide a way to export the analysis steps to a configuration file which can be later used to rerun the analysis. It is, of course, possible to document a GUI-based analysis in enough detail for it to be reproducible but it is a highly effortful and error-prone process. In general, switching from using GUI-based tools to running scripts substantially increases the analysis reproducibility. Our tutorial mainly addresses the researchers who already use script-based analyses.

This is not to say that using scripts automatically leads to reproducible analyses. This is exactly the reason why the approach described here or similar ones are required. And even applying those is not a guarantee of seamless reproducibility. Our approach, however, provides a way to get you reasonably close to this ideal. We thus believe that it is well worth the effort of learning and implementing it.

References

- Devezer, B., Navarro, D. J., Vandekerckhove, J., & Ozge Buzbas, E. (2021, March). The case for formal methodology in scientific reform. *Royal Society Open Science*, 8(3), rsos.200805, 200805. doi: 10.1098/rsos.200805
- Fischl, B. (2012, August). FreeSurfer. *NeuroImage*, 62(2), 774–781. doi: 10.1016/j.neuroimage.2012.01.021
- Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., ... Poldrack, R. A. (2016, December). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, 3(1), 160044. doi: 10.1038/sdata.2016.44
- Hardwicke, T. E., Mathur, M. B., MacDonald, K., Nilsson, G., Banks, G. C., Kidwell, M. C., ... Frank, M. C. (2018, August). Data availability, reusability, and analytic reproducibility: Evaluating the impact of a mandatory open data policy at the journal *Cognition*. *Royal Society Open Science*, 5(8), 180448. doi: 10.1098/rsos.180448
- Jas, M., Larson, E., Engemann, D. A., Leppäkangas, J., Taulu, S., Hämäläinen, M., & Gramfort, A. (2018, August). A Reproducible MEG/EEG Group Study With the MNE Software: Recommendations, Quality Assessments, and Good Practices. *Frontiers in Neuroscience*, 12, 530. doi: 10.3389/fnins.2018.00530
- Mölder, F., Jablonski, K. P., Letcher, B., Hall, M. B., Tomkins-Tinch, C. H., Sochat, V., ... Köster, J. (2021, April). Sustainable data analysis with Snakemake. *F1000Research*, 10, 33. doi: 10.12688/f1000research.29032.2
- Obels, P., Lakens, D., Coles, N. A., Gottfried, J., & Green, S. A. (2020, June). Analysis of Open Data and Computational Reproducibility in Registered Reports in Psychology. *Advances in Methods and Practices in Psychological Science*, 3(2), 229–237. doi: 10.1177/2515245920918872
- Open Science Collaboration. (2015, August). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716–aac4716. doi: 10.1126/science.aac4716

- Pashler, H., & Harris, C. R. (2012, November). Is the Replicability Crisis Overblown? Three Arguments Examined. *Perspectives on Psychological Science*, 7(6), 531–536. doi: 10.1177/1745691612463401
- Peikert, A., & Brandmaier, A. M. (2021, May). A Reproducible Data Analysis Workflow. *Quantitative and Computational Methods in Behavioral Sciences*, 1, e3763. doi: 10.5964/qcmb.3763
- Stockemer, D., Koehler, S., & Lentz, T. (2018, October). Data Access, Transparency, and Replication: New Insights from the Political Behavior Literature. *PS: Political Science & Politics*, 51(4), 799–803. doi: 10.1017/S1049096518000926
- Taulu, S., & Kajola, M. (2005, June). Presentation of electromagnetic multichannel data: The signal space separation method. *Journal of Applied Physics*, 97(12), 124905. doi: 10.1063/1.1935742
- Wakeman, D. G., & Henson, R. N. (2015, December). A multi-subject, multi-modal human neuroimaging dataset. *Scientific Data*, 2(1), 150001. doi: 10.1038/sdata.2015.1
- Xie, Y., Allaire, J. J., & Grolemund, G. (2019). *R Markdown: The definitive guide*. Boca Raton: CRC Press, Taylor and Francis Group.