

Time: the Cost of Reproducibility

Frédéric Clavert 

C²DH, University of Luxembourg

Elisabeth Guerard 

C²DH, University of Luxembourg

Mirjam Pfeiffer 

C²DH, University of Luxembourg

Andreas Fickers 

C²DH, University of Luxembourg

Introduction

5

In a previous [editorial](#), we have tried to reflect on “what it means to sustain a digital project such as the *Journal of Digital History* within our current digital environment, which is unstable by nature”, while referring to the notion of *updatism*. One of the ways to face this unstable digital environment is to ensure, the best we can, the reproducibility of all articles that we publish.

6

Since the beginning of our journal, we have invested a lot in guaranteeing the reproducibility of scholarly research outputs, enabling readers to use the original data to regenerate the same results. To ensure the reproducibility - specifically of the articles’ code - the JDH team has defined a thorough technical review, enabling all published articles to be run on the reader’s computer or to be run by services like MyBinder or Google Colab with a single click.

7

Though we are happy of the work achieved, the continuous development of our complex publishing platform while running the journal remains challenging. After four years of back-

end and front-end development and almost three years of editorial life, we know that ensuring reproducibility comes at a cost: time. While we do our utmost to keep the timespan between submission of an article and its sustainable online publication as short as possible, the complex workflow is often more time-consuming than expected. As editors and developers, we aim at making the itinerary process as smooth and efficient as possible, but we are aware of the fact that some of our authors would have loved to see their innovative work being published more rapidly.

8

As transparency has been a key epistemic virtue driving the establishment of the Journal of Digital History, this editorial aims at reflecting on the various challenges in the “lifespan” of an JDH article. In addressing the multiple steps of the workflow and explaining the many technical as well as editorial procedures involved, we aim at making the publication workflow more transparent and comprehensible for future authors and readers.

9

The lifespan of a *Journal of Digital History* article

10

Our publishing workflow can be described in several steps:

11

- **abstract submission** – We ask all potential authors to submit an abstract proposal. This is the occasion to meet with a member of the editorial team, to ensure that the article fits the *Journal of Digital History* specificities, and to emphasize the multi-layered structure of JDH which comes with certain affordances. As a first point of contact between authors and the editorial team, this step helps to manage expectations on both sides and to build trust.
- **writing** – Once the abstract is accepted, authors write their article. This step takes as long as the authors wish, unless they are publishing in a special issue which has a specific deadline. Yet even in the case of a special issue there is a certain temporal flexibility as we do not publish all articles at the same time. We are conscious that writing an article for the JDH is challenging. This journey starts with installing the necessary writing environment, made of, among other pieces of software, Jupyter notebooks, Zotero, GitHub, etc. We acknowledge that this process has been challenging for some of our authors in the past. Nevertheless, since 2020, we have tried to be proactive in organizing “writing workshops”, and organizing seminars and presentations at conferences, making future collaboration with potential authors easier. Today, authors can install this writing environment through Docker, enabling them to set up a JupyterLab installation based on R or Python depending on their needs. In order to facilitate this process, we have produced step-by-step instructions and video tutorials and provide best practices in terms of folder organization and GitHub templates. Last but not least, we offer one-to-one support sessions if needed. During this whole process of setup, the editorial team remains available in case of difficulties or problems.

- **article submission** – The submission of the article takes the form of an email, sent to the managing editor, including the URL to a GitHub repository. All articles of the *Journal of Digital History* are stored on GitHub for now.
- **technical review** – Once the article is submitted, it undergoes a so-called “technical review”, mostly realized by Elisabeth Guérard, our main journal developer. In managing this complex task, she is sometimes supported by a student-assistant or the managing editor. The technical review is complex and can sometimes be lengthy. As you will see below, it took around 280 days in some (extreme) cases (for instance for the first article we received with code written in R). At the time, we obviously lacked the know-how to technically assess a notebook in R, and R’s sensitivity to its library versions was another matter of complication. Yet experience matters, and we are now far better and quicker in dealing with articles using R. But as the code-environment is continuously changing, adapting to this environment asks for a constant learning and updating of skills on our side.
- **peer review** – Ensuring high quality articles requires a thorough peer review process. As all academic journals today, we face the situation that finding qualified peers to critically assess and review articles is a hard and very time consuming exercise. As many of the articles published in JDH ask for both a topical *and* at least basic technical expertise, the number of potential candidates to fulfil this time-consuming job is even more reduced. While we discussed this issue both with the international members of the editorial board and with our publisher DeGruyter-Brill, we didn’t find the “golden solution” how to incentivise peers and - even less so - how to solve the pathological lack of time and self-enslavement of academic staff in the neo-liberal university system.
- **final preparation** – This step includes a final technical review, a design review, the copy-editing of the article, and the preparation of its advertisement on social media (today mainly Facebook and BlueSky). The design review is becoming an ever more challenging step, as it concerns the readability of tables and dataframes (including their labelling), as well as the quality-check of figures and data-visualisations. Finally, all authors need to perform a final check of their article before we publish and advertise it.

12

Let’s now have a look at the duration of those different steps, from the article submission one (Figure 1 and 2).

13

14

We had to reconstruct the lifespan of each of the 16 articles (we did not include editorials, nor the articles of this Varia issue) that were published until now by the *Journal of Digital History*. One of our problems is that we are using three different systems to deal with articles, which can sometimes be confusing: emails, manuscript central (ScholarOne) and our own backend. As a consequence, the [data that we are using here](#) might contain mistakes. Furthermore, as we used another system for the first few months (Slack) to discuss with authors, some information is also there, though we are not using it anymore for more than one year now. This data has some limitations: the peer review duration encompasses different

cases: articles that were accepted after one or two rounds of evaluation, as well as those that proved more difficult to evaluate. Furthermore, we are here using the data concerning only published articles. In other words, articles that will be published soon are more prone to have longer technical and peer reviews durations.

```
import pandas as pd
import plotly.graph_objects as go
from datetime import datetime

# Define a date parser function
def date_parser(date_string):
    return pd.to_datetime(date_string, format='%d/%m/%Y', errors='coerce')

# Read the CSV file with explicit date parsing
df = pd.read_csv('JDH_review_chronology.csv',
                 parse_dates=['Abstract submission', 'Article submission',
                              'Technical review begins', 'Peer review start', 'Peer review end',
                              'Publication Date'],
                 date_parser=date_parser)

# Calculate durations
df['Technical Review duration'] = (df['Peer review start'] - df['Article submission']).dt.days
df['Peer Review Duration'] = (df['Peer review end'] - df['Peer review start']).dt.days
df['Review to Publication'] = (df['Publication Date'] - df['Peer review end']).dt.days

# Calculate total duration
df['Total Duration'] = (df['Publication Date'] - df['Article submission']).dt.days

# Sort by Issue and then by Total Duration within each Issue
df = df.sort_values(['Issue', 'Total Duration'], ascending=[True, True])

# Shorten titles, create clickable links, and include issue information
df['Short Title'] = df['Title'].apply(lambda x: x[:40] + '...' if len(x) > 43 else x)
df['Clickable Title'] = df.apply(lambda row: f'<b>{row["Issue"]}</b> <a href="{row["URL"]}" target="_blank">{row["Short Title"]}</a>', axis=1)
```

Article Submission and Review Process Duration (Sorted by Issue)

16

17

As you can see in Figure 1, the delays to publish an article in JHD can be quite long. The final part ('Review to publication') contains all the Article final preparation (last technical review, design review, copy-editing, social media advertisement preparation). There are quite some exceptions that would need explanation: the imagineRio article, for instance, has been published almost two years after its submission, due to an editorial decision: submitted before

the publication of the Digital Tools call for papers, whereas its peer review had already started, we decided, in agreement with the authors and the invited editors of the special issue, that it would best fit in this issue.

Sometimes, an article's peer review can also be delayed due to unforeseen technical issues. *Dialects of Discord* has had a very long technical review as it was the first article in R that we published and we lacked some know-how to technically review it faster. Articles submitted in R today are just as normal as articles in Python for us today, though.

18

table 1

```
import pandas as pd
import matplotlib.pyplot as plt

# Define a date parser function
def date_parser(date_string):
    return pd.to_datetime(date_string, format='%d/%m/%Y', errors='coerce')

# Read the CSV file with explicit date parsing
df = pd.read_csv('JDH_review_chronology.csv',
                 parse_dates=['Abstract submission', 'Article submission',
                              'Technical review begins', 'Peer review start', 'Peer review end',
                              'Publication Date'],
                 date_parser=date_parser)

# Calculate durations
df['Technical Review'] = (df['Peer review start'] - df['Article
submission']).dt.days
df['Peer Review Duration'] = (df['Peer review end'] - df['Peer review
start']).dt.days
df['Review to Publication'] = (df['Publication Date'] - df['Peer review
end']).dt.days

# Print rows where 'Review to Publication' is negative
print("Rows with negative 'Review to Publication' duration:")
print(df[df['Review to Publication'] < 0][['Title', 'Issue', 'Peer review
end', 'Publication Date', 'Review to Publication']])

# Group by issue and calculate mean durations
issue_means = df.groupby('Issue')[['Technical Review', 'Peer Review
Duration', 'Review to Publication']].mean()

print("\nMean durations by issue:")
print(issue_means)
```

Mean Duration of Review Stages by Issue

20

21

Nevertheless, as you can see with both figure 1 & 2, the timespan between submission and publication has increased between our first issue and the two following ones. There are several explanations to this. The first is that we had to deal with less articles than we do now. At least 12 articles are actually in peer review, *i.e.* two thirds of the number of articles we have already published.

For more general statistics about the Journal of Digital History, [you can consult this page](#).

22

Identifying bottlenecks

23

So, for more than six months, we have tried to identify bottlenecks in our workflow:

24

- The peer review platform we were using, ScholarOne (Manuscript Central), lacked flexibility. As our partner De Gruyter did acquire Ubiquity Press, an open access publisher of peer-reviewed academic journals, books and data based on OJS, we decided to switch to OJS for managing our peer review. Yet this switch takes time, as articles which started to be reviewed in the ScholarOne system have to be handled in this system until the review process is over.
- From the start, our managing editor, Frédéric Clavert, was dealing with the peer review process on its own. As this is a very time-consuming activity which asks for a lot of commitment and responsiveness, we looked for support and are very happy that Hannah Smyth (UCL) and Lorella Viola (VU Amsterdam) accepted to join the editorial team. We are confident that their support will have a positive impact on speeding up the peer-review process.
- In a similar vein, our backend developer, Élisabeth Guérard, was alone in performing the complex technical reviews. We therefore decided to hire a student assistant to ease her workload and we plan to hire other student assistants in the future. While the switch to OJS will make the managing editor's workflow a bit lighter, we are also currently in the process of hiring a second developer for the journal. For sure, this reinforcement of the team in terms of staff will make the production process more effective and quicker.
- Last but not least, we switched from a double-blind to a single-blind peer review process. The double-blind peer review process led to the existence of competing versions of the articles, one stored on the authors' GitHub repository, and the other one on the JDH's GitHub account. This added huge complexity to the anonymisation process - especially when dealing with different versions of the article (due to revisions required by the reviewers). Concerning the peer review process, we are thinking about introducing some sort of open peer review in the next few years, as well as reflecting on a framework to evaluate data visualisation.

Though those bottlenecks are specific to the JDH, we experienced one major difficulty which is common to many other journals nowadays: the recruitment of reviewers! Sometimes, we had to reach out up to 14 different colleagues in order to find one person accepting to do the review. Of course, we do not blame our colleagues here, we all have to refuse evaluating an article at one point in our career. But it is a symptom of a structural problem in academia - which is being discussed more critically in recent times. In a 2022 “Inside Higher Education” (a Times Higher Education spin-off) article, the “peer review crises” in the humanities and social sciences has been discussed at length. The [statement](#) by Gale Sinatra, professor of psychology and the Stephen H. Crocker Professor of Education at the University of Southern California, reflects pretty much our own experiences: “‘I’ve had a good run. I’ve done three journals, [and] I’ve enjoyed all three of these experiences. But I’ve peaked out because it’s just become too difficult,’ Sinatra said of trying to find reviewers, chasing down late reviews and, worst of all, apologizing to the scholars who understandably want to know if and when their delayed articles will be published.” Together with our editorial board and our publisher, DeGruyter-Brill, we discussed the possibility of paying reviewers. Yet most of the colleagues think, that offering a small amount of money will not be enough of an incentive to motivate people doing the job. It’s time that matters, not money! Yet we finally agreed with our publisher, that all reviewers will receive a book of our “Digital History & Hermeneutics”-series for free as a small and rather symbolic recognition for their important and crucial work. In addition, the editorial board decided that it would be a good idea to publish the names of colleagues who have been doing reviews on our journal website (if they accept to be published). This is just another small initiative on our side to symbolically value and recognize the important work that reviewers are doing to ensure academic quality in our field.

Sticking to reproducibility, respecting authors

When we started the *Journal of Digital History*, we knew we would face multiple challenges. We may have underestimated some, while being surprised of having solved others rather rapidly. Our development method, that one could qualify as ‘agile’, may also have proven insufficient, though it provided the required flexibility for a project that were in fact two: developing a publication platform and creating an academic journal.

There were also some elements that we did not really foresee as they later happened. If you read our About Page, there isn’t anything about reproducibility of what are executable papers. Reproducibility was implicit, but is today a major point of the JDH, and will stay a major point. Nevertheless, what is at stake when we talk reproducibility is the temporal costs to us and to our authors. In an epoch of ultra-fast circulation of information, historians are pressed to work faster. And some of our authors have no other choice in publishing their results as soon as possible - especially in the field of digital and computational history, where codes or algorithms can become obsolete in a few months.

29

Yet we need to collectively think and discuss about the costs of reproducibility, and how to make the different temporalities at stake when publishing at JDH visible and transparent. As some authors were asking us to publish the date of submission, acceptance and of publication together with their article, we did discuss this option with our Editorial Board where quite some variety of opinions were addressed. If everybody agreed that being more transparent on our review process would be a strong advantage, some of our colleagues emphasised that it could be interpreted as an indirect omission of the journal to the “publish or perish” pressure that we all face today.

30

It is up to us to find the right balance between the rigour that is required to publish an academic journal, our willingness to publish reproducible and executable data papers, and the need of our authors to publish *timely*. To us, a central element in finding that balance, beyond identifying bottlenecks and fixing problems on the run, is the need to communicate even better and more with our authors, and to produce even clearer and intuitive guidelines. This is what we promise to put our energy on in the coming months!

31

Reproducibility and our readers

32

Reproducibility is first and foremost something to be ensured by authors and editors. But its aim is also to make the reading of articles impactful for our readers’ research. In other words, if reproducibility is something that is ensured behind the scene (*i. e.* not visible on the website), during discussions between authors, editors, and technical editors, we wanted to translate it into an enhanced reading experience too, which implied the redesign of our reading interface.

33

For the first time, this new design - version 3 (beta) - can now be explored in Maria Eriksson’s, Tomas Skotare’s and Pelle Snickars’ article, [Tracking and tracing audiovisual reuse: Introducing the Video Reuse Detector](#). Again, this article presented serious challenges to the JDH team: while its peer-review took only a few weeks, the technical review took us months. The v3-design offers a more intuitive reading of the multi-layered articles by offering a “reading flow” between the narrative, hermeneutic, and code/data layer. We are keeping – this is our DNA – the principle of the layers, but their intertwinement is made more visible. In addition, we have introduced two new features in the code/data layer: readers are now able to run and edit the code directly within the article (without opening a separate window that displays the article as a code notebook), and for those who are less code literate, we have integrated an AI-generated explanation of the code. With this important new feature of the journal, we hope to make the data-driven scholarship showcased in JDH more accessible to the general field of history and - ideally - contribute to improve our readers’ code literacy in a simple but effective way.

34

This new design and its new functionalities are still in beta. One current limitation is, that running the code works only if the reader runs it from the first cell. While being a rather logical move - if a cell of code requires a variable that is created above, the cell where this variable is created should be run before - this limitation ends up messing the reader's experience. Furthermore, the functioning of the 'explain code' button is based on a call to [groq](#)'s API (not to be confused with xAI's grok, that we refuse to use for ethical reasons). For the time being, we have only restricted access to this API (one call per minute), making the reader's experience less optimal. But we are confident that this will change in the near future.

35

As mentioned, this is only the beginning of our v3-design, a design that will put even more emphasis on our authors' efforts to make the *Journal of Digital History*'s articles reproducible, and to make this reproducibility an asset for both authors and readers. As much as history is an unfinished business, the development and design of an online journal in the rapidly changing landscape of digital history and humanities remains a challenge in continuous engineering!

36

The *Journal of Digital History* team is happy to be part of [Revue 3.0](#), a project led by [Marcello Vitali-Rosati](#) (Université de Montréal) that "aims to develop the expertise needed to envision, design, and create protocols, infrastructures, and tools for academic journal publishing".

37

A few months ago, our colleague Lars Wieneke left the C²DH to start a new phase of his career outside academia. As head of the centre's digital infrastructure, Lars was an enthusiastic member of the team which built the *Journal of Digital History*. We wish to thank him for his fantastic contribution to this project.