

A Manifesto on Resource Re-Use in Interactive Information Retrieval

Maria Gäde
Humboldt-Universität zu Berlin
Berlin, Germany
maria.gaede@ibi.hu-berlin.de

Marijn Koolen
Royal Netherlands Academy of Arts
and Sciences
Amsterdam, Netherlands
marijn.koolen@di.huc.knaw.nl

Mark Hall
The Open University
Milton Keynes, United Kingdom
mark.hall@open.ac.uk

Toine Bogers
Aalborg University Copenhagen
Copenhagen, Denmark
toine@hum.aau.dk

Vivien Petras
Humboldt-Universität zu Berlin
Berlin, Germany
vivien.petras@ibi.hu-berlin.de

ABSTRACT

This perspective paper on resource re-use intends to draw the attention of the interactive information retrieval (IIR) community to the challenges of research documentation and archiving for future use. Resources are understood as encompassing research designs, research data and research infrastructures. It proposes eight principles for improving the re-use of resources in the IIR community and presents concrete steps on how to achieve them. A five-level system for data archiving and documentation envisions increasingly open and stable documentation and access infrastructures.

KEYWORDS

Research data, research design, research software, re-use, data documentation, data sharing practices

ACM Reference Format:

Maria Gäde, Marijn Koolen, Mark Hall, Toine Bogers, and Vivien Petras. 2021. A Manifesto on Resource Re-Use in Interactive Information Retrieval. In *Proceedings of the 2021 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '21), March 14–19, 2021, Canberra, ACT, Australia*. ACM, New York, NY, USA, ?? pages. <https://doi.org/10.1145/3406522.3446056>

1 INTRODUCTION

What makes a manifesto a manifesto? While consulting research manifestos across different disciplines [? ? ? ? ? ? ? ? ? ?], we found that a manifesto (1) identifies and defines a problem or area that has not gathered much traction in the research community, (2) motivates and argues why the problem needs to be tackled, (3) develops principles or recommendations for the community to address the challenge, and (4) ends with concrete calls to action.

This manifesto on resource re-use wants to draw the attention of the interactive information retrieval (IIR) community to the challenges of research documentation and archiving for future use.

It advocates the uptake of sharing and re-use in the IIR research community in order to support validation, standardization and ultimately comparability or even reproducibility of IIR research. While IIR prides itself on the heterogeneity of research approaches and the inclusiveness of the community, it is a sign of a maturing research area that methodologies and research designs increasingly adhere to certain standards and protocols. Some standardization of empirical research can already be observed in the IIR community, such as the evaluation of an experiment and its appropriate documentation.

While our vision of an ideal future does not include total conformity of research questions or methodologies, this manifesto does stress the importance of well-documented and—in the best of all cases—re-usable and shareable resources in IIR research in order to progress the discipline. Only if we successfully trace, document and reference our research designs, data and—optimally—infrastructure, will we be able to build stably on top of the empirical and experiment-based research foundation that has already been established for IIR and move the state of the art forward without the risk of repeating research efforts needlessly.

This manifesto is the result of a development process that started several years ago, when several IIR evaluation campaigns were organized at CLEF, including the Interactive Task of the Cultural Heritage at CLEF lab (CHiCi) [? ? ?] and the interactive Social Book Search track (iSBS) [? ? ?], where organizers and participants grappled with preserving research designs and research data across the different years the track was running. The iSBS track's goal was to aggregate a large pool of experimental interaction data from a book search portal in order to support data analysis with multilingual and multicultural user groups. Data gathering was performed by different research groups using the same platform that was provided by the organizers several years in a row. The challenges that were encountered in just three years of attempting to maintain continuity in the research design and making the research data available, reinforced the idea that research data management, documentation and archiving are important issues that the IIR community has yet to tackle effectively [?]. These ideas were elaborated in two workshops at CHIIR, the *Barriers to Interactive Information Retrieval Resource Re-use* workshop series (BIIRRR) [? ?], where the discussion focused on the challenges of re-using IIR resources—the aspects or components of an IIR experiment that should be documented and re-used. The BIIRRR workshop also tested the publication of

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHIIR '21, March 14–19, 2021, Canberra, ACT, Australia

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8055-3/21/03.

<https://doi.org/10.1145/3406522.3446056>




Figure 1: The three resource types that make up research: research design, research infrastructure, and research data

so-called *experience papers*, which allowed authors to discuss and document the research designs of previously published studies in a more detailed way—with room for describing failed experiments—which were not possible in the original outcomes-focused research publications.

This perspective paper—the Manifesto on Resource Re-use in IIR—is the culmination of these community discussions and thought processes. We first define the concepts re-use and resource (Section ??) before we motivate why re-use and documentation are important steps in moving IIR research forward (Section ??). We then propose eight principles for improving the re-use of resources in the IIR community and present concrete steps on how to achieve them (Section ??). Finally, we outline a plan on how to achieve the principles (Section ??).

Our manifesto is intended to draw the IIR community’s attention to aspects in our research projects that are sometimes overlooked in favor of ever more creative and innovative design work. We hope to spark a discussion within the wider community on the principles which, at the very least, should increase awareness of research documentation aspects in our community and, ideally, have a lasting and positive effect on own research practices.

2 DEFINING RESOURCE AND RE-USE

As we will expand upon further in Section ??, we believe that clearly defining the terminology used in a research study is essential to its successful future documentation and re-use. We therefore start our manifesto by defining the terminology used in the rest of the paper.

2.1 Resources

Research is a complex process that is supported by three main resource types that should be documented and re-used. As shown in Figure ??, we view these resources as interlocking puzzle pieces representing the different elements of the research process: (1) research design, (2) research infrastructure, and (3) research data.

Research design, the first type of resource, is defined here as the methods and techniques used to collect and analyse the empirical data [?]. This includes the design of the research environment (e.g., location of study, study participants, tasks), any data collection protocols or other instruments used (including the questions and scales) as well as the data analysis methods and measures [?].

We do not consider the physical research environment to be part of the research design, but part of the research infrastructure instead. In many research fields and in IIR in particular, a technical infrastructure is also required and designed to handle the experimental procedure and tools and support the overall execution of

the research design. In IIR, the *research infrastructure* usually provides access to an IR system (including software, interfaces and collections) as well as the application of the data collection techniques such as user management, pre- and post-test questionnaires, assignment and randomization of user tasks, interaction logging, and other components [?].

Finally, the type of resource that has arguably received the most attention in recent years is *research data*. An increasing focus on open access and open science has been responsible for promoting research data publication and the further development of research data repositories [?]. Research data is typically defined quite broadly, for instance by [?], p. 3], who defines it as “the output of any systematic investigation involving a process of observation, experiment or the testing of a hypothesis, which when assembled in context and interpreted expertly will produce new knowledge”, a heterogeneity that is echoed by [?]. We define research data as any data that has been collected, observed, generated or created during or as a results of the research process.

For the purposes of the IIR community and this manifesto, we view research as a holistic process and consider the term ‘resource’ to cover any component belonging to one of these three resource types: research design, research infrastructure and research data.

2.2 Re-use

Several recent articles discuss and disagree on what exactly constitutes re-use. [?] define *re-use* as use of data by someone other than the creator of the data. When the creator uses the data beyond the original purpose for which it was created, it is still embedded in the creator’s context, which they define as *use*. In the context of reproducibility there are different types of re-use, e.g., “reanalyzing published data, repeating the study, reprocessing the ‘raw’ data, or replicating the findings under different conditions” [?]. Broader re-use includes activities such as “returning to one’s own data for later comparisons, acquiring datasets from public or private sources to compare to newly collected data, surveying available datasets as background research for a new project, or conducting reanalyses of one or more datasets to address new research questions” [?]. [?] therefore see re-use as a process, not as a single action.

In defining different levels of re-use associated with different curation intensities, the [?, p. 20] propose a continuum of local to global re-use, which ranges from *research data collections* that are focused on the specific research project for which they were made with minimal curation, to *resource collections* that serve a community and with more curation, all the way to *reference collections* that serve a wider community and follow robust and comprehensive standards. The characteristics of the resource collection determine its potential for re-use by others, with re-use in a wider community requiring increasing levels of curation and documentation.

There is a set of related concepts regarding re-usable research, which [?] considers distinguishable by three variables: research question, research data and research method. In [?]'s conceptual analysis, *re-use of data* happens when it is used for a different research question using a different method. Using the same data with the same method for a different research question would be *reinterpretation*. Using the same data with the same question and method would be *replication*, while using a different dataset with

the same question and method would be *reproduction*. Re-use is here defined not by the data producers or re-users, but by the context within which the use takes places.

Van de Sandt et al. [?] compared 20 definitions of re-use found in the literature and provided an analysis of the discourse around use and re-use to identify and compare their characteristics. They find that there is no clear way to distinguish between use and re-use, as there are no attributes or characteristics that clearly delineate the two. This is particularly true in disciplines where research data is not necessarily generated for a single project, but for community use, such as social science panels on societal questions. Van de Sandt et al. define (re-)use as “the use of any research resource regardless of when it is used, the purpose, the characteristics of the data and its user”.

For the purposes of the IIR community and this manifesto, we define re-use in its broadest sense as use of research data, research designs or infrastructure for more than an individual purpose.

3 THE IMPORTANCE OF RE-USE

Research cultures that promote re-use and support it through sustainable data sharing habits, policies and infrastructures allow its researchers to spend more of their time working on knowledge discovery and innovation instead of repetition of development work. It also allows for the replication and reproduction of earlier studies, thereby vetting those results and solidifying the joint foundation of scientific knowledge that future studies can build on [?]. Young researchers that aim to enter a research field and interdisciplinary projects building upon previous research will particularly benefit from data sharing. The formulation of the FAIR standards in 2016 with the explicit goal of ensuring the sustainable use of research data exemplifies this mindset [?]. FAIR means that research data should be (1) findable, (2) accessible, (3) interoperable and (4) re-usable—with these principles applying equally to researchers (data producers) and repositories (data providers). Our manifesto is closely aligned with the FAIR principles [?], but focuses on the specific issues encountered in IIR, while at the same time expanding the scope beyond research data to encompass research designs and infrastructure.

While research data access, re-use, and reproducibility have been promoted for years, there is still much we do not know about the intentions, incentives, practices and barriers to re-use [?]. One of the major barriers to optimal re-use of research data is often the lack of contextualization and implicit knowledge transfer [?]. Sharing research data means more than just sharing the original input and output data to an experiment. For a complete understanding and reliable interpretation of the data, the context of data production also needs to be documented. This includes information about the environment in which the experiments were designed and took place as well as implicit knowledge that was gained and drawn upon before, during, and after the study.

For contextual information, the research designs of data producing studies also need to be documented, archived, and shared to enable their re-use. While complete and transparent documentation of research designs is important for all types of studies, it is perhaps even more important in the context of qualitative research, which is part of the IIR method portfolio. Here, re-using the research design

is often the only possibility to reproduce or validate previous studies, since it is often more problematic to make qualitative research data available for re-use due to ethical or legal issues, in particular privacy protection—such as interview data for example [?].

The development of the research infrastructure represents a major investment of time and effort, much of which is duplicated or could be reduced through adaptation and re-use. The difficulty is that—except for very generic components, such as the underlying information retrieval system—re-use generally requires rewriting significant parts of the code base [?], resulting in a tendency to rewrite from scratch. Even for more generic aspects of the IIR research infrastructure, such as software to manage the participants flow through an experiment, the attempts at developing a re-usable solution [??] have not achieved any uptake. This mirrors the state of research software re-use in general [??].

Some disciplines, such as the social sciences or psychology, have already progressed far in documenting and providing access to research data and research designs [?]. Resource re-use is no less important for the IIR field: IIR studies are typically complex due to the interaction between users and the search system(s) and the combination of user- and system-centered evaluation [?]. This necessitates sophisticated research designs with rich combinations of data collection and analysis methods and a research infrastructure with many interdependent IR system components. In the past few decades, there have been several multi-year IIR campaigns that have attempted to re-use the experimental setup from year to year within the campaign, such as the TREC Interactive Track (1997–2002) [?], the INEX Interactive Track (2004–2010) [???], the Cultural Heritage in CLEF (CHC) Interactive Task (2013) [??], and the interactive Social Book Search (iSBS) task (2014–2016) [??].

While these campaigns have provided a great deal of insight into the challenges and opportunities for longer-term re-use of IIR resources [?], there has been relatively little re-use of resources related to research design or research infrastructure from campaign to campaign. Despite an ongoing interest in standardising IIR experimentation and evaluation within evaluation campaigns, this shows that establishing and maintaining a collaborative process and platform to support the re-use of IIR resources remains an unsolved issue. This is further compounded by a need for flexibility in being able to tailor the research design to the specific search scenario of each study or campaign.

Nevertheless, there have been efforts to collect, archive, and make available certain types of IIR study resources. For instance, the Repository of Assigned Search Tasks (RepAST)¹ collects, categorises, and provides access to a large collection of search tasks mined from publications of IIR studies [?]. Apart from the RepAST initiative, however, no continuity or satisfying re-use practices appears to have been established throughout the years [?]. Our goal with this manifesto is to provide both guiding principles as well as concrete steps to stimulate and strengthen the re-use of IIR resources.

¹<https://ils.unc.edu/searchtasks/search.php>, last visited February 25, 2021.




Figure 2: Principles to Support Re-Use in IIR

4 PRINCIPLES TO SUPPORT RE-USE IN IIR

We now present eight principles for improving resource re-use in the IIR community, which are visualized in Figure ?? . At the core lie the principles on *Knowledge* (#1) and *Terminology* (#2) that underpin everything else. The remaining principles are split into those aimed at the individual researcher—*Plan* (#3), *Document* (#4), and *Publish* (#5)—and those aimed at the IIR community as a whole—*Require* (#6) and *Reward* (#7). Around all of these sits the *Strive for more!* principle (#8) that encourages all of us, both individually and as a community, to move re-use forward. For each principle, we describe who we argue is primarily responsible for it as well as which of the three resource types the principle addresses; Table ?? contains an overview of this information. All principles apply to all resource types (see Section ??), but may be more relevant for some than for others (see Section ??).

4.1 Principle 1: Make tacit knowledge explicit

All implicit knowledge, choices, and study aspects should be noted, codified, and made explicit.

Transforming existing protocols and research methods, with all the distinct elements that they consist of, into a study requires a large number of choices, that are often based upon implicit knowledge and experience. These implicit aspects of the study are often hard to codify, but they need to be documented and made explicit to enable any re-use. This extends into the study itself, where checklists and notes should be used to capture exceptions and unusual aspects or events [?].

Inspired by the Manifesto for Research Services [?], this principle helps not only with reproducibility, but also reduces the risk of biases and other issues that later call into question the study’s results. Additionally, by describing the tacit knowledge of a study and comparing it with others, it becomes possible to gradually develop and merge that tacit knowledge into the research design, formalisms and procedures, which can then be more easily codified [?] and become common ground.

Concrete actions.

- Prepare a checklist of the practical matters to take into account before, during, and after the experiment.
- Document your exact study goals, research questions, premises, and (if applicable) hypotheses.
- Take notes, screenshots (and photos or videos where appropriate) during your experiments.
- Take note of actions and decisions that are not part of standardized research designs and not on the checklist.

- Take note of aspects that are difficult to reproduce and describe them.

4.2 Principle 2: Define your terminology

All concepts and terms should be precisely defined, ideally referring to existing, published definitions.

Even though many concepts and terms are common in the community, there are often multiple definitions with subtle or not so subtle differences. The specific definitions used in a study should be made explicit, including references to publications where they exist. Where existing definitions are not used, a clear explanation should be provided for why new definitions are required. The terminology should also be used precisely and consistently throughout the study.

This is particularly important as IIR is a highly interdisciplinary field, with researchers coming from various academic backgrounds, who might not all be familiar with the same concepts and terminology, or who would assume different definitions based on their disciplinary background and training. Reproducibility, for instance, is defined by [?] as the same researchers using the same experimental setup, whereas the ACM definition [?] refers to this as repeatability. Instead, ACM defines reproducibility as a different research group using a different setup to confirm the reported outcomes and findings. The potential misunderstandings that could arise from such subtle differences could be avoided by providing clear and specific definitions.

Concrete actions.

- Identify which concepts and terminology you will use while creating your research design.
- Look for alternative definitions or related concepts.
- Discuss differences in definitions and uses of concepts and terms that you find across publications.
- Discuss to what extent they apply to your work.
- Argue for the definitions you choose or why you introduce your own.
- Cite papers that define concepts that you use including those with alternative definitions that you do not use.

4.3 Principle 3: Plan the data life cycle

The full data life cycle, covering collection, processing, analysis, publication, and archiving, should be planned before starting any data collection.

Data collection is often time- and resource-consuming. Therefore, IIR research should start with an initial Data Management Plan (DMP) that clearly describes what kind of research data will be collected or generated during the study, how it will be stored, managed, described and analysed, and how it will be shared and archived at the end of the study. This ensures that the collection and usage of data is transparent to the participants and research community.

In the fields of medical and health studies, presenting such a DMP in advance is generally a basic requirement [? ? ?] and these plans are usually made publicly available. Increasingly, funding

| Principle | Primary responsibility | | Resource types addressed | | |
|----------------|------------------------|-----------|--------------------------|------|----------------|
| | Researcher | Community | Designs | Data | Infrastructure |
| #1 Knowledge | ✓ | | ✓ | | |
| #2 Terminology | ✓ | | ✓ | | |
| #3 Plan | ✓ | | | ✓ | |
| #4 Document | ✓ | | ✓ | ✓ | ✓ |
| #5 Publish | ✓ | | ✓ | ✓ | ✓ |
| #6 Require | | ✓ | ✓ | ✓ | ✓ |
| #7 Reward | | ✓ | ✓ | ✓ | ✓ |
| #8 Strive | ✓ | ✓ | ✓ | ✓ | ✓ |

Table 1: Overview of the primary responsibility for each of the eight principles (left) as well as the different resource types addressed by them (right).

organizations require a DMP to be submitted together with a project proposal in most data-producing research areas [?]. There are several tools, which will support the creation of a DMP², which will also provide guidelines on how to document and archive data appropriately.

Concrete actions.

- Start with a data management plan and document what data will be collected and for which purpose.
- Make sure that your data collection practice follows current data protection and privacy regulations such as the General Data Protection Regulation (GDPR) in the European Union.
- Publish your data management plan as part of your research project and/or publication of results.

4.4 Principle 4: Document your research resources

All aspects of the research resources should be documented in detail.

In principle, everything required to reproduce an experiment, including the raw research data, analyses and systems used should be documented. As an initial minimum requirement, the research design must be documented to enable a basic level of resource reuse. This includes the experimental setup (e.g., location of study, study participants, tasks), any data collection protocols or other instruments used (including the questions and scales) as well as the data analysis and measures [?].

Documenting the research design must include implicit knowledge and terminology definitions as covered by principles #1 and #2. Additionally, contextual information, such as the study goals, research questions, or associated publications should also be documented. Where available, standardized and machine-readable documentation formats should be used, e.g., the User Study Exchange Format [?].

In a more advanced re-use context, the research data and research infrastructure (including software) should also be documented in similar ways.

²See for instance <https://dmponline.decc.ac.uk/>, last visited February 25, 2021.

Concrete actions.

- Review existing practices and guidelines (in the social sciences it is common to archive codebooks for research designs, for example).
- Define your data model or use a community standard for research resource documentation.

4.5 Principle 5: Archive & publish your research resources

All documentation (see principle #4) should be archived in open access, long-term storage to be available for publication reviews and should then be published in a referenceable location.

Publications are generally limited in the amount of space available to authors and details on the research designs, data and infrastructure are frequently the first victim of this limit [?]. There are several platforms where additional research data and documents can be formally published, e.g., the Open Science Framework³, Zenodo⁴, or figshare⁵. Many funding organisations already require that publications must be made available via public open access repositories—some of them can also be used to archive the research designs and other research resources.

A dedicated repository for research designs in IIR does not (yet) exist, but for some resources related to the research designs, specific repositories exist, for example RePaST for tasks [?]. Tasks used in studies should therefore also be provided to those repositories to provide more access points for the community. This may involve some duplication of work, but if machine-readable formats are used, this should reduce the effort required.

Concrete actions.

- Identify an appropriate open access archive repository.
- Separately archive those aspects of the resources that can be freely shared and those that cannot for legal or other reasons.

³<https://osf.io>

⁴<https://zenodo.org>

⁵<https://figshare.com/>

- Include links to the archived resources in the publication submission.
- After publication of research outcomes, publish those resource aspects that can be freely shared and update the publication to refer to those.
- Publish parts of the resources to specialised repositories, where such exist.

4.6 Principle 6: Require basic sharing practices

Publication venues should require that resources be archived and available to reviewers.

The foundations for this principle are often already in place with publishers [?]. What is needed is that journals and conferences in the IIR field adopt them. This principle also goes hand-in-hand with principle #7, which is focused on rewarding good resource sharing behaviour. Together they should produce better, more cooperative outcomes [?]. Optionally, publication venues could also require that the archived resources are published after the main publication.

To ensure uptake, minimally the research design documentation must become a required part of the publication process. For complete transparency, also the research data and infrastructures should be made available first to reviewers and then—if at all possible—to the public. If a data archival section is already a requirement for a publication, the research design documentation can be integrated there, but if not, then there should be a required section in the paper that includes a referenceable link to the documentation.

Concrete actions.

- Require that at least documented research designs for IIR studies are available to reviewers at CHIIR.
- Petition journals relevant to the IIR field to add research design and data archiving into their publication workflow.

4.7 Principle 7: Reward good sharing practices

Researchers who follow good sharing practices should be visibly rewarded.

Complementing the basic requirements in principle #6, the IIR community should also develop a range of rewards to highlight good reusability practice. Changing the incentives for responsible sharing practices is a problem that requires a coordinated effort by all stakeholders to alter existing reward structures—authors, reviewers, chairs/editors, publishers, funders, institutions, and societies.

One particular reward mechanism that shows promise are *badges* to acknowledge reusability practices. Promoting or even requiring such badges would signal that the IIR community values these practices. An example of the successful use of such a badging system is the *Psychological Science* journal, where the adoption of open data badges has had a positive effect, increasing data sharing by a factor of more than ten [?]. Another such initiative is by the Association for Computing Machinery (ACM), which proposes a set of badges to denote functional, reusable, and available research artifacts [?].

Another approach for rewarding good practices would be to provide incentives towards providing more detail when reporting methodology, study procedures, i.e., the research design, and results. This is achieved by either allowing authors to add extra pages to a publication or provide additional publication venues or categories, dedicated to the reporting of research designs. This would remove the current incentive structure towards providing simplified, ‘clean’ narratives, which can result in the incomplete reporting of research designs or results.

Concrete actions.

- Allow for adding pages to a publication that are focused on reporting research designs.
- Institute a badging system with different levels at CHIIR, ISIC, and topically related journals.
- Institute an award for best-documented research paper at conferences like CHIIR.
- Add a publication category focused on reporting research designs, such as the experience papers at the 2019 BIIRRR workshop.

4.8 Principle 8: Strive for more!

Unlike the previous seven principles, which cover the basics required to embed re-use practices into the IIR community, this principle covers where the IIR community could go next. As such, the principle does not contain any concrete steps, only potential actions.

Potential actions.

- If you have shared your research design and data management plan, the next step is to share your research data and infrastructure. If you have shared your data and research infrastructure, the next step is to add documentation. Documentation can have different levels of detail and explicitness and different levels of structure and adherence to conventions. In the next section we offer concrete suggestions for increasing levels of documentation and sharing.
- Beyond sharing the designs, infrastructure and data of your research, you could pre-register your IIR studies, including hypotheses that are to be tested. There is a call for more transparency, including pre-registration of studies and hypotheses, across an increasing number of fields [? ? ?]. Although a significant part of our research is exploratory, where we do not have a very specific hypothesis to test, there are many studies where we have very concrete and specific questions and hypotheses. Pre-registration of the design, setup and hypotheses of such a study helps others to understand the original motivation behind the design and setup of the study, so that its components can more easily be re-used and compared against.
- Another step beyond sharing resources is to try to reproduce an IIR study. This requires a deep understanding of the theory as well as the research design of the original study. With proper re-use of IIR research resources, this type of reproducibility study could be used for graduate work. Organizing dedicated reproducibility tracks at conferences for Master’s or PhD students and their supervisors could provide great

training opportunities as well as a deeper understanding of which original studies truly contribute to our understanding of information seeking behavior.

5 THE WAY FORWARD

5.1 Five Levels of IIR Resource Re-Use

Inspired by the Five Star Linked Data system [?], which describes a hierarchy of steps to increase the level of openness and linking in data, we envision a five-level system for data archiving and documentation to increase the re-use of IIR resources. While level five is the ultimate objective for ensuring open documentation and re-use, every level contributes to a more standardized, documented and re-usable resource space for IIR. The levels build upon each other: a level can only be reached if the levels that come before it have already been accomplished. The last three levels are advanced levels, which depend on open and long-term accessible repositories and data models for structured documentation. While there are examples for research design and research data repositories available, there is no dedicated repository for IIR resources.

- *Level 1: Documentation of the resources in the research publication.* The description of the research design, data and infrastructure in the publication describing the study is already a community standard. We have included this as a first level, because very often, the description of the methodology and data lacks sufficient detail or is necessarily abbreviated because of typical length restrictions in conference and journal papers. Length restrictions should not preclude authors from at least precisely describing their research designs.
- *Level 2: Extended documentation of the resources on a project website.* To overcome information gaps and therefore challenges in reproducibility and re-use because of brevity restrictions, the study should be documented extensively on a publicly reachable website, which the authors create and maintain. Even though there is a risk that such websites may disappear, we argue that this level represents an important step. First, even if the documentation does completely or partially disappear after some time, this is still preferable over not having had any documentation at all. Second, ‘free’ and long-term options such as GitHub Pages are widely available, ensuring at least medium-term availability.
- *Level 3: Structured documentation of research design in a repository.* Documenting the research design in an openly accessible repository overcomes the challenges of maintenance and continuity within a research group, when researchers are leaving. Documenting research designs in a structured way—using a community-accepted data model, which includes the necessary information to reproduce an experiment—does not only provide a guideline on how to document research designs, but also makes them comparable across studies.
- *Level 4: Structured and documented archival of research design + research data.* While documenting a research design in an open repository is relatively straight-forward, because it shouldn’t incur any rights or permission issues, this next level requires not only the documentation of the research design, but also the documentation and the archiving of the generated research data in a repository. Preparing research

data for publication on a repository does not only include the documentation and preparation of the data, but also involves rights clearance, which sometimes is more difficult than the documentation (see principle 3).

- *Level 5: Structured and documented open archival of research design + research data + research infrastructure.* Ultimate re-use opportunities are provided, if not only the research design and research data are made available, but also the research system, so that other groups can experiment with it. Since research software is rarely stable and well-documented [?], this is may be an even bigger challenge than documenting and archiving research designs and research data, as evident from the lack of re-use of IIR research software [? ? ?] and research software in general [? ? ?].

Our five levels of archiving and documentation align well with the continuum of local to global (re-)use proposed by the US National Science Board [?] referenced also in Section 2.2. The more types of resources are shared and the more they are documented in a structured and standardized manner, the more we move IIR resources from local use to global re-use.

5.2 The Manifesto and the Future

Despite broad agreement over the necessity of being able to reproduce experimental research, the current academic infrastructure does not reward reproducibility studies the same way it rewards original work. This results in the overwhelming majority of research effort being directed at new research before the community has been able to properly verify the veracity and generalizability of recent research. If original work is required for publication, one might argue that making research resources re-usable may not be a worthwhile objective for the community to pursue.

While we believe that documenting and re-using research resources will also support original research and improve good scientific practice, another important application area for resource documentation and sharing is education and community growth. Students, young researchers and others entering the community will equally benefit from having resources available to attempt to reproduce IIR studies for their introduction to the field.

The aim of this paper is to enable a community discussion on the role of sharing and re-use of resources in IIR research, what we should strive for and how to get there. This manifesto should be discussed by the entire community and modified where necessary to ensure broad consensus. Moreover, as our eighth principle indicates, our shared understanding of how to best facilitate resource re-use develops over time, so our principles should be adapted according to new insights. Therefore, the manifesto should be treated as a living document, with this paper being merely the first version. One idea is to maintain a community repository with up-to-date best practices and resources, including this and future versions of the manifesto.